



# **DANZ Monitoring Fabric User Guide**

**Arista Networks**

[\*www.arista.com\*](http://www.arista.com)

*DANZ Monitoring Fabric User Guide*

*DOC-06660-02*

Headquarters	Support	Sales
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
<a href="http://www.arista.com/en/">www.arista.com/en/</a>	<a href="mailto:support@arista.com">support@arista.com</a>	<a href="mailto:sales@arista.com">sales@arista.com</a>

© Copyright 2023 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks are subject to Arista Network Terms of Use Policy, available at [www.arista.com/en/terms-of-use](http://www.arista.com/en/terms-of-use). Use of marks belonging to other parties is for informational purposes only.



# Contents

<b>Chapter 1: Introduction to DANZ Monitoring Fabric.....</b>	<b>1</b>
1.1 Overview.....	1
1.2 Out-of-band Monitoring with DANZ Monitoring Fabric.....	1
1.3 Using the DANZ Monitoring Fabric CLI.....	2
1.4 Using the DANZ Monitoring Fabric GUI.....	5
1.4.1 DMF Features Page.....	9
1.4.2 Dashboard Layout.....	22
1.4.3 Controller Health.....	23
1.4.4 Switch Health.....	29
1.4.5 Policy Health.....	36
1.4.6 Smart Node Health.....	38
1.4.7 Empty State.....	41
<b>Chapter 2: Managing DMF Switches and Interfaces.....</b>	<b>43</b>
2.1 Overriding the Default Configuration for a Switch.....	43
2.1.1 CLI Configuration.....	43
2.2 DMF Interfaces.....	43
2.2.1 Using the GUI to Configure a DMF Filter or Delivery Interface.....	44
2.2.2 Using the CLI to Configure a DANZ Filter or Delivery Interface.....	46
2.2.3 Using the CLI to Identify a Filter Interface using Destination MAC Rewrite.....	48
2.2.4 Using the GUI to Identify a Filter Interface using Destination MAC Rewrite.....	50
2.3 Using Interface Groups.....	54
2.3.1 Using the GUI to Configure Interface Groups.....	54
2.3.2 Using the CLI to Configure Interface Groups.....	55
<b>Chapter 3: Managing DMF Policies.....</b>	<b>56</b>
3.1 Overview.....	56
3.2 Policy Elements.....	56
3.3 Configuring a Policy.....	57
3.3.1 Configure a Policy Using the GUI.....	57
3.3.2 Configure a Policy Using the CLI.....	63
3.3.3 Define Out-of-band Match Rules Using the CLI.....	65
3.3.4 Stop, Start, and Schedule a Policy Using the CLI.....	65
3.3.5 Clear a Policy Using the CLI.....	66
3.3.6 View Policies Using the CLI.....	66
3.4 Define a Policy Using Interfaces GUI.....	67
3.5 Define a Policy Using Host Tracker GUI.....	68
3.6 Using the Packet Capture Action in a Policy.....	69
3.7 Define a Policy with a Packet Capture Action Using the CLI.....	70
<b>Chapter 4: Viewing Information about Monitoring Fabric and Production Networks.....</b>	<b>72</b>
4.1 Monitoring DMF Interfaces.....	72
4.1.1 Using the GUI to Monitor DMF Interfaces.....	72
4.1.2 Viewing Oversubscription Statistics.....	72
4.1.3 Using the CLI to Monitor Interface Configuration.....	73

4.2 Viewing Devices Connected to the Monitoring Fabric.....	75
4.2.1 Using the GUI to View Fabric Connected Devices.....	75
4.2.2 Using the CLI to View Switch Configuration.....	76
4.3 Viewing Information about a Connected Production Network.....	76
4.3.1 Using the CLI to View Connected Devices and LAGs.....	77

## **Chapter 5: Using the DMF Service Node Appliance.....78**

5.1 Overview.....	78
5.2 Changing the Service Node Default Configuration.....	78
5.3 Using SNMP to Monitor DPDK Service Node Interfaces.....	79
5.4 Configuring Managed Services.....	80
5.4.1 Using the GUI to Define a Managed Service.....	80
5.4.2 Using the CLI to Define a Managed Service.....	81
5.4.3 Monitoring Managed Services.....	83
5.5 Deduplication Action.....	84
5.6 Header Strip Action.....	85
5.6.1 Configuring the Post-service Match.....	87
5.7 IPFIX and Netflow Actions.....	87
5.7.1 IPFIX Template.....	88
5.7.2 IPFIX Keys.....	88
5.7.3 IPFIX Fields.....	89
5.7.4 Active and Inactive Timers.....	89
5.7.5 Example Flowset.....	89
5.7.6 Using the GUI to Define an IPFIX Template.....	90
5.7.7 Using the CLI to Define an IPFIX Template.....	92
5.7.8 Using the GUI to Define an IPFIX Service Action.....	93
5.7.9 Using the CLI to Define an IPFIX Service Action.....	93
5.8 Packet-masking Action.....	94
5.9 Arista Analytics Node Capability.....	95
5.9.1 Configuring the Arista Analytics Node Using the GUI.....	96
5.9.2 Configuring the Arista Analytics Node Using the CLI.....	97
5.10 Pattern-drop Action.....	99
5.11 Pattern-match Action.....	100
5.12 Slice Action.....	102
5.13 Packet Slicing on the 7280 Switch.....	103
5.13.1 Using the CLI to Configure Packet Slicing - 7280 Switch.....	104
5.13.2 Using the GUI to Configure Packet Slicing - 7820 Switch.....	105
5.13.3 Troubleshooting Packet Slicing.....	111
5.14 Session-slice Action.....	112
5.15 Timestamp Action.....	113
5.16 UDP-replication Action.....	114
5.17 Redundancy of Managed Services in Same DMF Policy.....	115
5.17.1 Using the GUI to Configure a Backup Managed Service.....	115
5.17.2 Using the CLI to Configure a Backup Managed Service.....	116
5.18 Application Identification.....	116
5.18.1 Using the CLI to Configure Application Identification.....	116
5.18.2 Using the CLI to Configure app-id.....	117
5.18.3 Using the CLI to Configure app-id-filter.....	117
5.18.4 Using the GUI to Configure app-id and app-id-filter.....	118
5.18.5 Dynamic Signature Updates (Beta Version).....	124
5.18.6 CLI Show Commands.....	126
5.18.7 Syslog Messages.....	126
5.18.8 Scale.....	126
5.18.9 Troubleshooting and Considerations.....	127
5.19 Redundancy of Managed Services Using Two DMF Policies.....	127

5.20 Cloud Services Filtering.....	128
5.20.1 Cloud Services Filtering Configuration.....	128
5.21 Multiple Services Per Service Node Interface.....	129
<b>Chapter 6: Using the DMF Recorder Node.....</b>	<b>131</b>
6.1 Overview.....	131
6.2 Configuration Summary.....	131
6.3 Indexing Configuration.....	132
6.4 Pre-buffer Configuration and Events.....	133
6.5 Using an Authentication Token.....	133
6.6 Using the GUI to Add a Recorder Device.....	133
6.6.1 Configuring a Node to Use Local Storage.....	136
6.6.2 Configuring a Node to Use External Storage.....	137
6.7 Configuring a Recorder Node Interface.....	139
6.8 Using the GUI to Assign a Recorder Interface to a Policy.....	140
6.9 Using the GUI to Define a Recorder Query.....	141
6.10 Viewing Query History.....	143
6.11 Using the CLI to Manage the DMF Recorder Node.....	144
6.11.1 Basic Configuration.....	144
6.11.2 Authentication Token Configuration.....	146
6.11.3 Configuring the Pre-buffer.....	146
6.11.4 Triggering a Recorder Node Event.....	146
6.11.5 Terminating a Recorder Node Event.....	147
6.11.6 Viewing Recorder Node Events.....	147
6.12 Using the CLI to Run Recorder Node Queries.....	147
6.12.1 Packet Replay.....	147
6.12.2 Packet Data Query.....	148
6.12.3 Packet Object Query.....	148
6.12.4 Size Query.....	148
6.12.5 Window Query.....	149
6.12.6 Stopping a Query.....	149
6.13 Using RBAC to Manage Access to the DMF Recorder Node.....	149
6.13.1 RBAC Configuration Using the CLI.....	150
6.13.2 RBAC Configuration Using the GUI.....	150
6.14 Using the CLI to View Information About a Recorder Node.....	152
6.14.1 Viewing the Recorder Node Interface.....	152
6.14.2 Viewing Recorder Node Operation.....	152
6.14.3 Viewing Errors and Warnings.....	153
6.15 Using the GUI to view Recorder Node Statistics.....	154
6.16 Changing the Recorder Node Default Configuration.....	160
6.17 Large PCAP Queries.....	161
6.18 Recorder Node Management Migration L3ZTN.....	162
6.19 Recorder Node CLI.....	163
6.20 Multiple Queries.....	163
6.21 Ability to Deduplicate Packets - Query from Recorder Node.....	165
6.21.1 Using the CLI to Deduplicate Packets.....	165
6.21.2 Using the GUI to Deduplicate Packets.....	166
6.21.3 Limitations.....	167
<b>Chapter 7: Link Aggregation.....</b>	<b>169</b>
7.1 Configuring Link Aggregation.....	169
7.1.1 Using the GUI to Configure Link Aggregation Groups.....	169
7.1.2 Using the CLI to Configure Link Aggregation Groups.....	171
7.1.3 Configuring Hashing Fields.....	172

7.1.4 L2 GRE Key Hashing.....	173
7.1.5 VxLAN Hashing.....	176
7.2 Pseudo Multi-Chassis Link Aggregation.....	179
7.3 MLAG Components.....	180
7.4 MLAG Limitations.....	180
7.5 Configuring an MLAG via the CLI.....	181
7.6 MLAG Link Discovery.....	183
7.7 Configure MLAG via GUI.....	183
7.8 Create MLAG Policy from GUI.....	185
7.9 Viewing Policy Statistics in the GUI.....	185
7.10 Viewing MLAG Links in the GUI.....	186
7.11 Using LAG Interfaces as Members in MLAG Interfaces.....	187
7.12 Overlapping Policies in LAGs.....	187
<b>Chapter 8: Tunneling Between Data Centers.....</b>	<b>190</b>
8.1 Understanding Tunneling.....	190
8.2 Encapsulation Type.....	191
8.3 Using Tunnels in Policies.....	193
8.4 Using the GUI to Configure a GRE Tunnel.....	193
8.5 Using the CLI to Configure a GRE Tunnel.....	194
8.5.1 Using the CLI to Rate Limit the Packets on a GRE Tunnel.....	195
8.5.2 Using the CLI to View GRE Tunnel Interfaces.....	195
8.6 Using the GUI to Configure a VxLAN Tunnel.....	196
8.7 Using the CLI to Configure a VxLAN Tunnel.....	197
8.7.1 Using the CLI to Rate Limit the Packets on a VxLAN Tunnel.....	198
8.7.2 Using the CLI to View VxLAN Tunnel Interfaces.....	198
8.8 Viewing or Modifying Existing Tunnels.....	199
8.9 Using a Tunnel with User-defined Offsets.....	199
<b>Chapter 9: Integrating vCenter with the DANZ Monitoring Fabric.....</b>	<b>201</b>
9.1 Overview.....	201
9.2 Using SPAN to Monitor VMs.....	202
9.3 Using ERSPAN to Monitor VMs.....	202
9.3.1 Configuration Summary for vCenter Integration.....	203
9.4 Defining a Tunnel Endpoint.....	203
9.5 Wildcard Tunnels for VMware vCenter Monitoring.....	205
9.5.1 Configuration.....	205
9.5.2 Using the CLI to create Wildcard Tunnels.....	205
9.5.3 Using the GUI to create Wildcard Tunnels.....	207
9.5.4 Limitations.....	207
9.6 Using the GUI to Integrate a vCenter Instance.....	207
9.6.1 Using a vCenter Instance as the Traffic Source in a DMF Policy.....	211
9.7 Using the CLI to Integrate a vCenter Instance.....	211
9.8 Using the GUI to View vCenter Configuration.....	214
9.9 Using the CLI to View vCenter Configuration.....	215
9.10 Minimum Permissions for Non-admin Users.....	216
<b>Chapter 10: Advanced Fabric Settings.....</b>	<b>222</b>
10.1 Configuring Advanced Fabric Settings.....	222
10.2 Managing VLAN Tags in the Monitoring Fabric.....	224
10.3 Auto VLAN Mode.....	225
10.3.1 Configuring Auto VLAN Mode using the CLI.....	226
10.3.2 Configuring Auto VLAN Mode using the GUI.....	228

10.4 Auto VLAN Range.....	231
10.5 Auto VLAN Strip.....	233
10.5.1 Auto VLAN Strip using the CLI.....	234
10.5.2 Auto VLAN Strip using the GUI.....	241
10.6 CRC Check.....	242
10.6.1 CRC Check using the CLI.....	243
10.6.2 CRC Check using the GUI.....	243
10.7 Custom Priority.....	244
10.7.1 Configuring Custom Priority using the GUI.....	245
10.7.2 Configuring Custom Priority using the CLI.....	246
10.8 Device Deployment Mode.....	246
10.8.1 Configuring Device Deployment Mode using the GUI.....	246
10.8.2 Configuring Device Deployment Mode using the CLI.....	248
10.9 Inport Mask.....	248
10.9.1 InPort Mask using the CLI.....	248
10.9.2 Inport Mask using the GUI.....	249
10.10 Match Mode.....	250
10.10.1 Setting the Match Mode Using the CLI.....	251
10.10.2 Setting the Match Mode Using the GUI.....	251
10.11 Retain User Policy VLAN.....	252
10.11.1 Retain User Policy VLAN using the CLI.....	252
10.11.2 Retain User Policy VLAN using the GUI.....	254
10.12 Tunneling.....	255
10.12.1 Configuring Tunneling using the GUI.....	255
10.12.2 Configuring Tunneling using the CLI.....	256
10.13 VLAN Preservation.....	257
10.13.1 Using the CLI to Configure VLAN Preservation.....	257
10.13.2 Using the GUI to Configure VLAN Preservation.....	259
10.13.3 Troubleshooting.....	262
10.14 Reuse of Policy VLANs.....	263
10.15 Rewriting the VLAN ID for a Filter Interface.....	263
10.16 Reusing Filter Interface VLAN IDs.....	264
10.17 Using Push-per-filter Mode.....	265
10.18 Tag-based Forwarding.....	265
10.19 Prefix Optimizations.....	265
10.20 Switch Dual Management Port.....	266
10.20.1 Overview.....	266
10.20.2 Configuring Dual Management Using the CLI.....	267
10.20.3 Configuring Dual Management Using the GUI.....	267
10.20.4 Management Interface Selection.....	270
10.20.5 Switch Fabric Management Redundancy Status.....	271
10.20.6 Additional Notes.....	271
10.21 Controller Lockdown.....	271

## **Chapter 11: Advanced Policy Configuration.....273**

11.1 Advanced Match Rules.....	273
11.1.1 Match Fields and Criteria.....	273
11.1.2 Match-except Rules.....	274
11.1.3 Matching with IPv6 Addresses.....	275
11.1.4 Port and VLAN Range Matches.....	275
11.1.5 User-defined Filters.....	276
11.2 Using the Filter and Delivery Role with MAC Loopback for a Two-stage Policy.....	278
11.2.1 Using the GUI To Configure a Filter and Delivery Interface with MAC Loopback.....	279
11.2.2 Using the CLI To Configure a Filter and Delivery Interface with MAC Loopback.....	281
11.3 Rate Limiting Traffic to Delivery Interfaces.....	282

11.3.1 Rate Limiting Using the GUI.....	282
11.3.2 Rate Limiting Using the CLI.....	283
11.4 Configuring Overlapping Policies.....	285
11.4.1 Configuring the Policy Overlap Limit Using the GUI.....	286
11.4.2 Configuring the Overlapping Policy Limit Using the CLI.....	287
11.4.3 Using the CLI to View Overlapping Policies.....	288
11.4.4 Configuring the Policy Overlap Limit Strict using the GUI.....	288
11.4.5 Configuring the Policy Overlap Limit Strict using the CLI.....	290
11.5 Viewing Information about Policies.....	290
11.5.1 Viewing Policy Flows.....	290
11.5.2 Viewing Packets Dropped by Policies.....	290
11.6 Using Rule Groups.....	290
11.6.1 Using the GUI to Configure Rule Groups.....	291
11.6.2 Using the CLI to Configure Interface Groups.....	293
11.7 PTP Timestamping.....	294
11.7.1 Platform Compatibility.....	294
11.7.2 Configuration.....	295
11.7.3 Configuring PTP Timestamping using the CLI.....	295
11.7.4 Configuring PTP Timestamping using the GUI.....	295
11.7.5 PTP Interface Configuration.....	298
11.7.6 Policy Configuration for Timestamping.....	299
11.7.7 Using the CLI Show Commands.....	299
11.7.8 Configuration Validation Messages.....	301
11.7.9 Troubleshooting.....	301
11.7.10 Limitations.....	302

## **Appendix A: Stenographer Reference for DMF Recorder Node.....303**

A.1 Stenographer Query Syntax.....	303
A.2 Example Stenographer Queries.....	306

## **Appendix B: DMF Recorder Node REST API..... 308**

B.1 Authentication.....	308
B.1.1 Basic HTTP Authentication.....	308
B.1.2 Authentication with an Authentication Token.....	308
B.2 DMF Recorder Node API Headers.....	310
B.3 DMF Recorder Node REST APIs.....	310
B.3.1 /ready.....	310
B.3.2 /query/window.....	310
B.3.3 /query/size.....	311
B.3.4 /query/application.....	311
B.3.5 /query/packet.....	311
B.3.6 /query/analysis[filter="<stenographer-query-string>"][type="<analysis-type>"].....	312
B.3.7 /query/replay/request[filter="<stenographer-query-string>"][real-time="<boolean>"].....	312
B.3.8 /query/replay/request[filter="<stenographer-query-string>"][speed-mbps="<int>"].....	313
B.3.9 /query/replay/done.....	313
B.3.10 /erase/packet[filter="<stenographer-query-string>"].....	313
B.3.11 /event/update[trigger="<boolean>"][name="<string>"][pre-buffer-minutes="<int>"].....	314
B.3.12 /abort/query.....	314
B.3.13 /abort-all/query.....	314
B.3.14 /queries.....	315
B.3.15 /status/query.....	315
B.3.16 /status/all.....	316

<b>Appendix C: Configuring Third-party Services.....</b>	<b>317</b>
C.1 Services in the DANZ Monitoring Fabric.....	317
C.1.1 Using the GUI to Configure a DMF Unmanaged Service.....	318
C.1.2 Using the CLI to Configure a DMF Unmanaged Service.....	321
C.2 Service Insertion and Chaining in a DMF Policy.....	323
 <b>Appendix D: References.....</b>	 <b>325</b>
D.1 Related Documents.....	325

---



## Introduction to DANZ Monitoring Fabric

---

This chapter introduces the DANZ Monitoring Fabric (DMF) and the user interfaces provided for out-of-band monitoring and configuration.

### 1.1 Overview

DANZ Monitoring Fabric (DMF) is a cloud-first Network Packet Broker (NPB) which provides a single pane-of-glass with an integrated visibility fabric. The DMF solution includes NPB functionality with the DMF Recorder Node and the Analytics Node for deeper monitoring and pervasive security of out-of-band workloads in hybrid cloud deployments.

DMF leverages an SDN-controlled fabric using high-performance, open networking (white box/brite box) switches and industry-standard x86 servers to deploy highly scalable and flexible network visibility and security solutions. Traditional, box-based, hardware-centric NPBs are architecturally limited when trying to meet the evolving security and visibility demands of Cloud Native data centers. DMF addresses the challenges of traditional NPB solutions by enabling a scale-out fabric for enterprise-wide security and monitoring, a single pane of glass for operational simplicity, and multi-tenancy for multiple IT teams, including NetOps, DevOps, and SecOps.

### 1.2 Out-of-band Monitoring with DANZ Monitoring Fabric

As data center networks move toward 40/100G designs, cloud computing, hyper scale data analytics, and 5G mobile services, traffic monitoring must also transition to next-generation designs. To manage the modern data center, a large portion of network traffic must be copied and aggregated from TAP or SPAN ports and forwarded to monitoring and analysis tools. These tools, used for managing network performance, application performance, security, and compliance leverage data recorders, intrusion detection systems, data leakage detectors, SLA measurement devices, and other traffic analyzers like Wireshark.

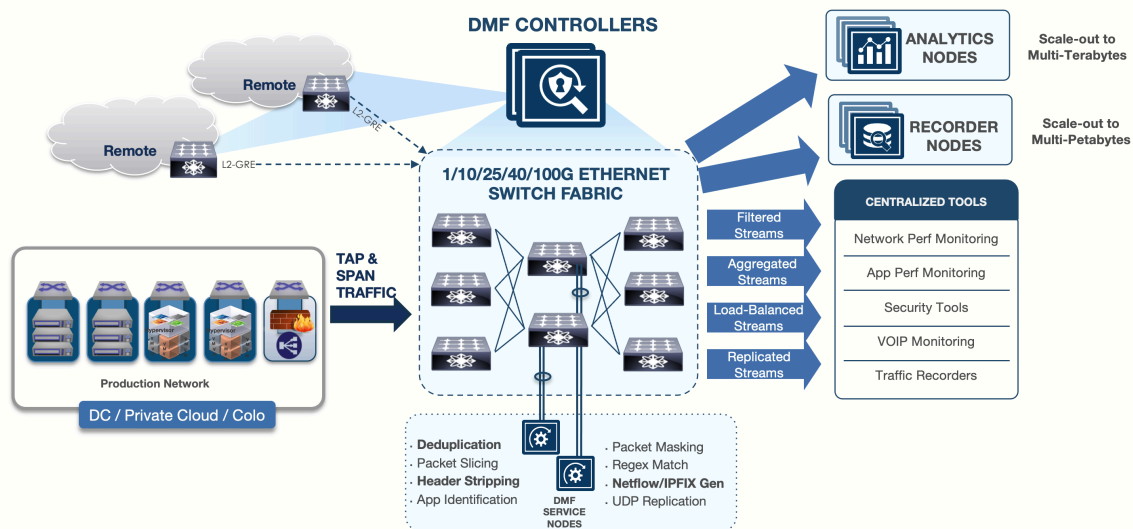
DMF delivers an open, production-grade, and scalable monitoring solution based on Software Defined Networking (SDN) technology using high-performance open-networking switches. The centralized DMF controller provides flexibility, simplifies policy management and monitoring fabric configuration, and supports cost-effective monitoring of data centers and remote sites or branches with up to several thousand TAP and SPAN ports.

DMF architecture is inspired by hyper scale networking designs and consists of the following components:

- HA pair of SDN-enabled DMF Controllers (VMs or hardware appliances), which enable simplified and centralized configuration, monitoring, and troubleshooting.
- Arista Networks SDN-enabled Switch Light OS, a production-grade, ONIE-deployable, lightweight OS, which runs on DMF Ethernet switches.
- Open Ethernet switches (white box/brite box) which use the same merchant silicon ASICs used by most incumbent switch vendors and have been widely deployed in production data center networks. These switches ship with Open Network Install Environment (ONIE) for automatic and vendor-agnostic installation of third-party network OS.
- DANZ Service Nodes (optional), a Data Plane Development Kit (DPDK)-powered, x86-based appliances that connect to the DMF, either singly or as part of a service node chain. The service node provides advanced packet functions, such as deduplication, packet slicing, header stripping, regex matching, packet masking, UDP replication, and IPFIX/NetFlow generation.

- DANZ Recorder Nodes (optional), x86-based appliances that connect to the DMF and are managed via the DMF Controller to provide petabyte packet recording, querying, and replay functions.
- Analytics Nodes (optional), x86-based appliances that integrate with the DMF to provide multi-terabit, security, and performance analytics with configurable, historical time-series dashboards.

**Figure 1-1: Out-of-Band Monitoring with DANZ Monitoring Fabric**



DMF lets a network operator easily deploy data center-wide monitoring with the following benefits:

- Organization-wide visibility: delivers traffic from any TAP to any tool at any time across one or multiple locations.
- Flexible, scale-out fabric deployment: supports a large number of 1G, 10G, 25G, 40G, and 100G ports (thousands per fabric).
- Multi-tenant tap and tool sharing: supports monitoring by multiple teams to enable Monitoring Fabric as a Service.
- Massive operational simplification: provides a single pane of glass for provisioning, management, monitoring and debugging through a centralized SDN controller. This feature eliminates the need for box-by-box configuration.
- Centralized programmability: a REST-based API architecture enables event-based, centralized policy management and automation for integrated end-to-end IT work flows. This feature leverages DMF Service Nodes, Analytics Nodes, and Recorder Nodes.
- Dramatic cost savings: multi-fold reduction of total cost of ownership due to open Ethernet switch economics, industry-standard x86 servers, optimized usage of tools, and SDN-enabled operations and automation.

## 1.3 Using the DANZ Monitoring Fabric CLI

Before connecting to the DMF controller, make sure the DMF application is running. Log in to the DMF controller using the local console or SSH to the address assigned to the DMF controller during installation.



**Note:** All configuration changes related to fabric switches must be made through the controller CLI, which provides configuration options in the `config-switch` submode for each switch. Do not log in to the switch to make changes directly using the switch CLI.

CLI commands are divided into modes and submodes, which restrict commands to the appropriate context. The main modes and their available commands are as follows:

- **login mode:** commands available immediately after logging in, with the broadest possible context.

- **enable mode:** commands that are available only after entering the enable command.
- **config mode:** commands that have a more significant effect on system configuration and that can only be entered after entering the **configure** command. The user can also access submodes from this mode.

Enter submodes from config mode to provision specific monitoring fabric objects. For example, the switch **switchname** command changes the CLI prompt to **(config-switch) #** and lets the user configure the switch identified by the switch name.

When the user logs in via SSH to the controller, the CLI appears in login mode, where the default prompt is the system name followed by a greater than sign (>), as shown below:

```
controller-1>
```

To change the CLI to enable mode, enter the **enable** command. The default prompt for enable mode is the system name followed by a pound sign (#), as shown below:

```
controller-1> enable
controller-1#
```

To change to config mode, enter the **configure** command. The default prompt for config mode is the system name followed by **(config) #**, as shown below:

```
controller-1> config
controller-1(config)#
```

To change to a submode, enter the command from config mode, followed by any object identifier required, as in the following example:

```
controller-1(config)# switch filter-switch-1
controller-1(config-switch)# interface ethernet54
controller-1(config-switch-if)#
```

To return to enable mode, type **end**, as shown below:

```
controller-1(config)# end
controller-1#
```

To view the path to the current CLI prompt, enter the **show this** command from any nested submode, as in the following example:

```
controller-1(config-switch-if)# show this
! switch
switch filter-switch-1
interface ethernet54
```

To view details about the configuration, enter the **show this details** command, as in the following example:

```
controller-1(config-switch-if)# show this details
! switch
switch filter-switch-1
!
interface ethernet54
no force-link-up
no optics-always-enabled
no shutdown
```

To view a list of available commands in the current or submode, enter the **help** command.

```
controller-1> help
For help on specific commands: help <command>
Commands:
%<n>                Move job to foreground
debug
echo                Print remaining arguments
enable              Enter enable mode
exit                Exit submode
help                Show help
history             Show commands recently executed
logout              Logout
no                  Prefix existing commands to delete item
ping                Send echo messages
ping6               Send echo messages
profile             Configure user profile
reauth              Reauthenticate
set                 Manage CLI sessions settings
show
support             Generate diagnostic data bundle for technical support
terminal            Manage CLI sessions settings
topic               Show documentation on topic
upload              Upload diagnostic data bundle for technical support
watch               Show output of other commands
whoami              Identify the current authenticated account
workflow            Show workflow documentation
controller-1>
```

To view detailed online help for the command, enter the **help** command followed by the command.

```
controller-1> help support
Support Command:      Generate diagnostic data bundle for technical support
Support Command Syntax: no support skip-switches skip-cluster skip-service-
nodes
                        skip-recorder-nodes sequential support [[skip-switch
es]
                        [skip-cluster] [skip-service-nodes]
                        [skip-recorder-nodes] [sequential]]
Next Keyword Descriptions:
sequential:           Use sequential (non-parallel) fallback collection
mode, which will be slower
                        but use fewer resources.
skip-cluster:         Skip cluster information from the collection.
skip-recorder-nodes:  Skip recorder nodes information from the collection.
skip-service-nodes:   Skip service nodes information from the collection.
skip-switches:        Skip switches information from the collection.
Support Command:      Generate diagnostic data bundle for technical support
Support Command Syntax: no support skip-switches skip-cluster skip-service-
nodes
                        skip-recorder-nodes sequential support [[skip-switch
es]
                        [skip-cluster] [skip-service-nodes] [skip-recorder-
nodes] [sequential]]
Next Keyword Descriptions:
sequential:           Use sequential (non-parallel) fallback collection
mode, which will be slower
                        but use fewer resources.
skip-cluster:         Skip cluster information from the collection.
skip-recorder-nodes:  Skip recorder nodes information from the collection.
skip-service-nodes:   Skip service nodes information from the collection.
skip-switches:        Skip switches information from the collection.
```

```
controller-1>
```

To display the options available for a command or keyword, enter the command or keyword followed by a question mark (?).

```
controller-1> support ?
<cr>
sequential                Use sequential (non-parallel) fallback collection
mode, which will be slower but use fewer resources.
skip-cluster              Skip cluster information from the collection.
skip-recorder-nodes       Skip recorder nodes information from the collection.
skip-service-nodes        Skip service nodes information from the collection.
skip-switches             Skip switches information from the collection.
controller-1>
```

To view the permitted values or keywords for any command, enter the command followed by a space, and press the <Tab> key. The command completion feature displays a concise list of permitted values, as in the following example:

```
controller-1> support <TAB>
<cr> sequential skip-cluster
skip-recorder-nodes skip-service-nodes skip-switches
controller-1>
```

For information about managing administrative access to the DMF controller, refer to the **DANZ Monitoring Fabric 8.4 Deployment Guide**.

## 1.4 Using the DANZ Monitoring Fabric GUI


The DANZ Monitoring Fabric GUI (DMF GUI) lets the user perform similar operations to the CLI using a graphic user interface instead of text commands and options. The DMF GUI can be used with recent versions of any of the following supported browsers:

- Mozilla Firefox
- Google Chrome
- Microsoft Edge
- Internet Explorer
- Apple Safari

To connect to the DMF GUI, use the DMF controller IP address. Use the virtual IP (VIP) assigned to the cluster if that was configured during deployment. Using the VIP ensures that the user connects to the current active controller, regardless of any failover that may have occurred.

The active controller must be used for all configuration operations and for reliable information when monitoring DMF. The standby controller is provided only for redundancy in case the active controller becomes unreachable. No configuration can be performed from the standby controller, and the information displayed may not be accurate. The figure below illustrates connecting to the DMF GUI using HTTPS (**port 443**) at the IP address: **192.168.17.233**

**Figure 1-2: Connecting to the DANZ Monitoring Fabric GUI**



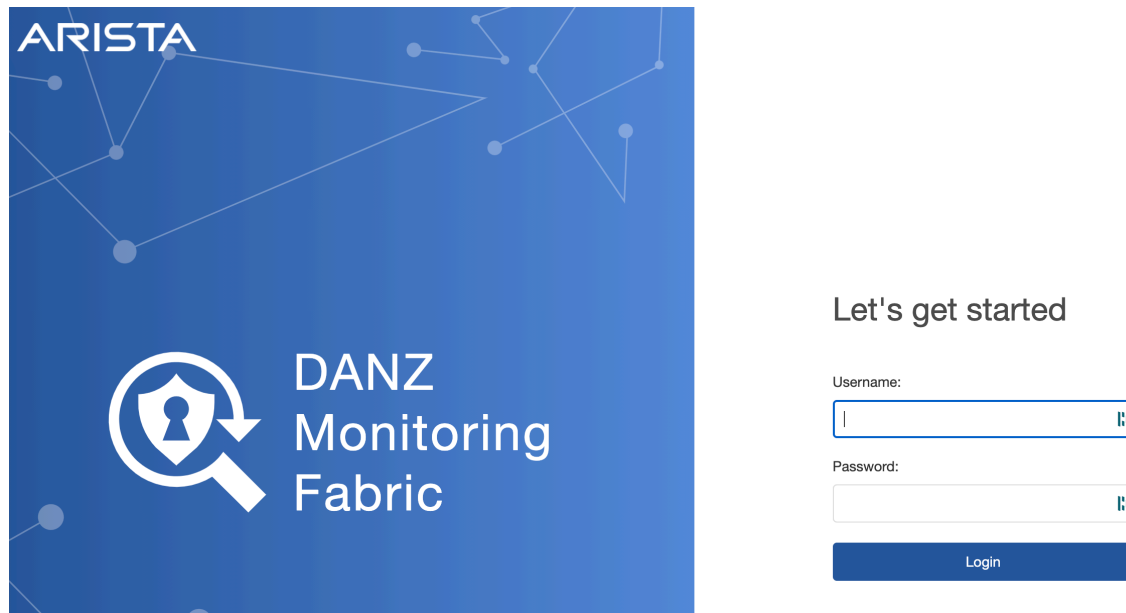
The first time a user connects to the controller, they may receive a prompt to make a security exception because the controller HTTPS server is using an unknown (self-signed) certificate authority.



**Note:** When using Internet Explorer, if the system time is different than the controller time, the login attempt may fail. To fix this, make sure the system being used to log in to the controller is synchronized with the controller.

After accepting the prompts, the system displays the login prompt, shown in figure below.

**Figure 1-3: DANZ Monitoring Fabric GUI Login Prompt**



Use the admin username and the password that was configured for the DMF controller during installation, or any user account and password configured with administrator privileges. A user in the read-only group will have access to options for monitoring fabric configuration and activity but will not be able to change the configuration.

**Figure 1-4: DANZ Monitoring Fabric GUI Main Menu**

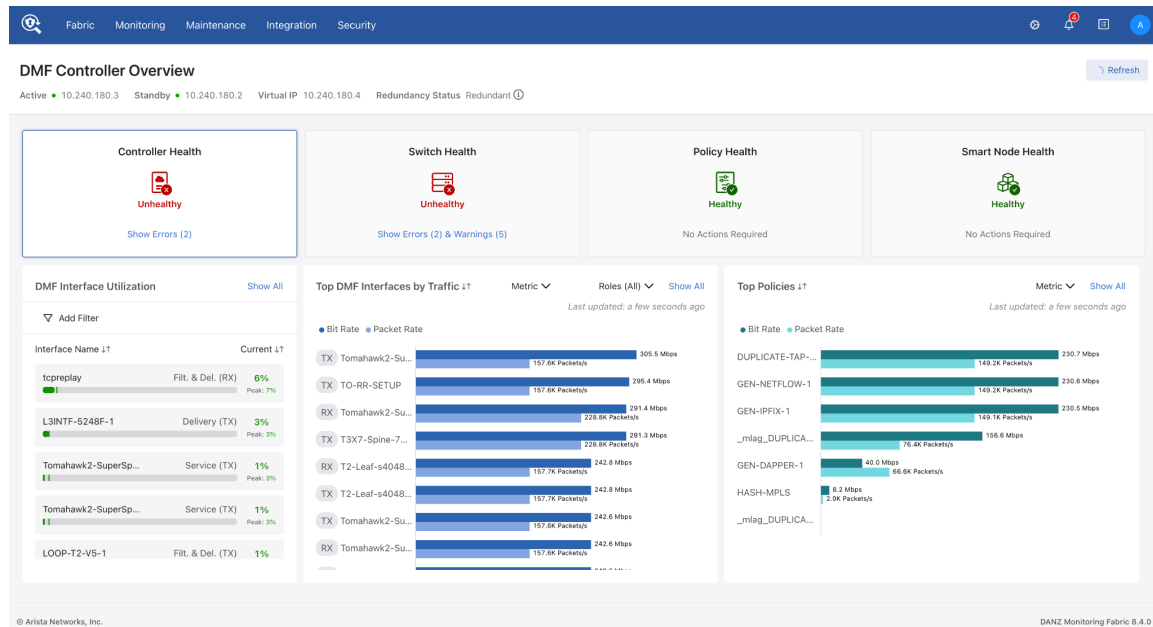


When a user logs in to the DMF GUI, they see a landing page. This page shows the **DMF Controller Overview**, dashboard and a menu bar at the top (shown above) with sub menus containing options for setting up DMF and for monitoring network activity. The menu bar includes the following sub menus:

- **Fabric:** manage DMF switches and interfaces.
- **Monitoring:** manage DMF policies, services, and interfaces.
- **Maintenance:** configure fabric wide settings (clock, SNMP, AAA, sFlow, Logging, Analytics Configuration).
- **Integration:** manage integration of vCenter instances to allow monitoring traffic using DMF.
- **Security:** manage administrative access.
- A profile page, that allows the user to display or change user preferences, change the password, or sign out.

The newly designed dashboard for DMF 8.4 displays information about the controller, including switches, interfaces, policies, and Smart Nodes.

**Figure 1-5: DMF Controller Overview**



The header displays the following basic information about the controller:

- Active IP address
- Standby IP address
- Virtual IP address
- Redundancy Status - The status contains an informational tool tip that can be hovered for more details.

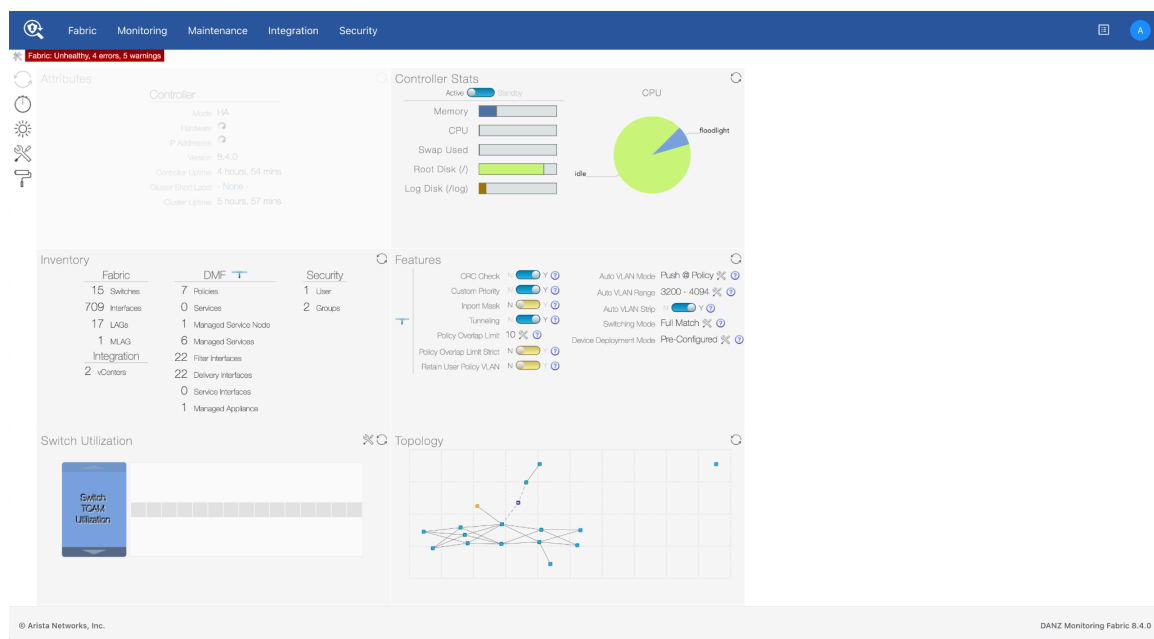
Four cards control the type of content displayed on the main section of the page. The cards are:

- Controller Health
- Switch Health
- Policy Health
- Smart Node Health



**Note:** This dashboard is on by default in the **Settings** page under the Navigation section. Toggling off displays the previous dashboard as illustrated below.

**Figure 1-6: Legacy Dashboard**



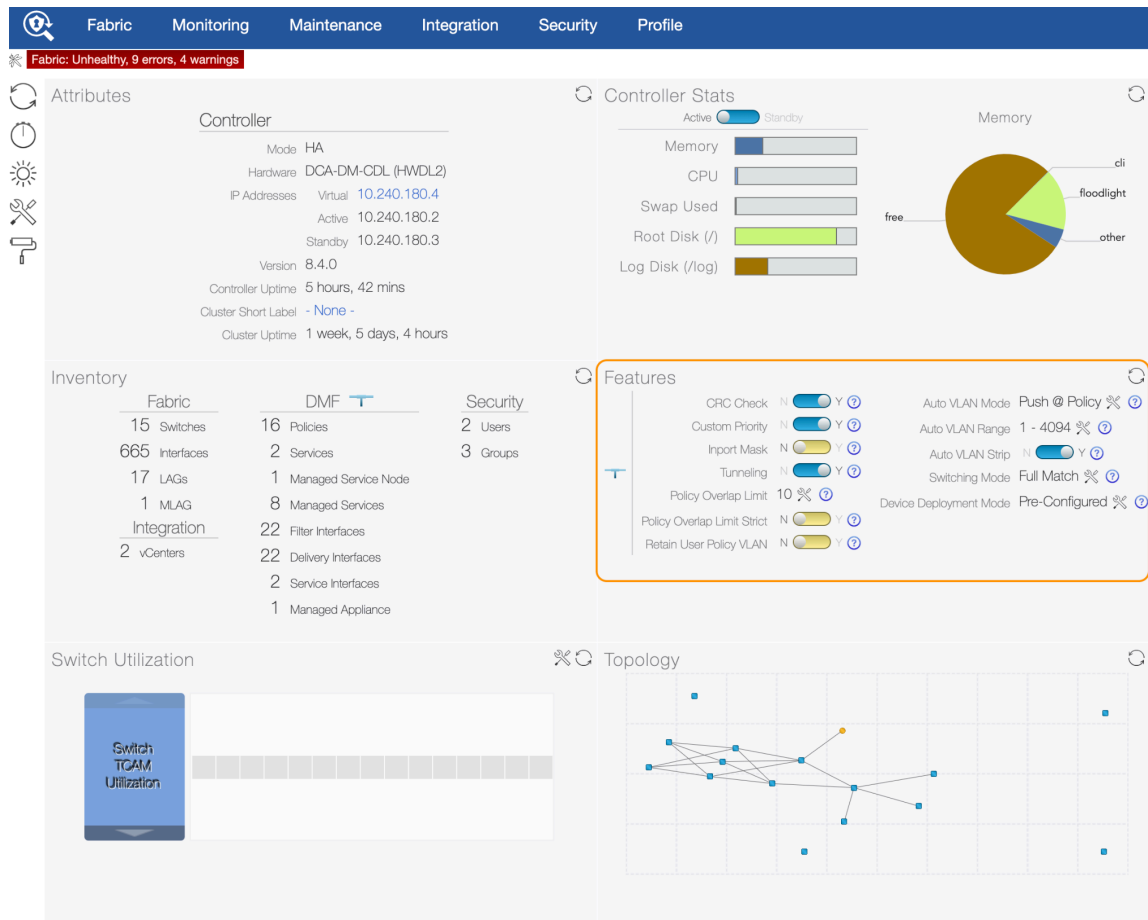


## 1.4.1 DMF Features Page

### Overview

Prior to the DMF 8.4 release, the fabric-wide settings, specifically the Features section as shown below were available on the home page after logging in.

**Figure 1-7: DMF Legacy Page (pre 8.4)**



In DMF 8.4, a newly designed Dashboard replaces the former home page. The **Features** section is now the new **DMF Features** page. To navigate to the DMF Features Page, click on the **gear icon** in the navigation bar.

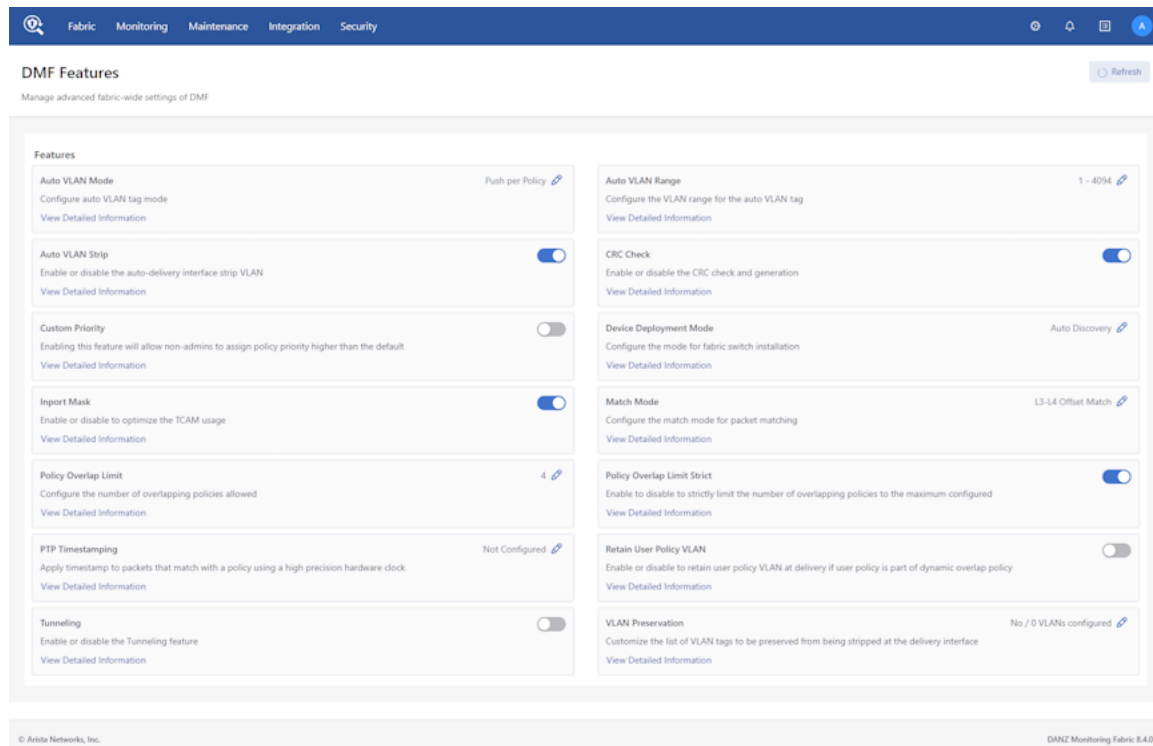
**Figure 1-8: Gear Icon**



## Page Layout

All fabric-wide configuration settings required in advanced use cases for deploying DMF policies appear in the new DMF Features Page.

**Figure 1-9: DMF Features Page**




The fabric-wide options used with DMF policies include the following:

**Table 1: Feature Set**

Auto VLAN Mode	Auto VLAN Range
Auto VLAN Strip	CRC Check
Custom Priority	Device Deployment Mode
Inport Mask	Match Mode
Policy Overlap Limit	Policy Overlap Limit Strict
PTP Timestamping	Retain User Policy VLAN
Tunneling	VLAN Preservation

Each card on the page corresponds to a feature set.

**Figure 1-10: Feature Set Card**



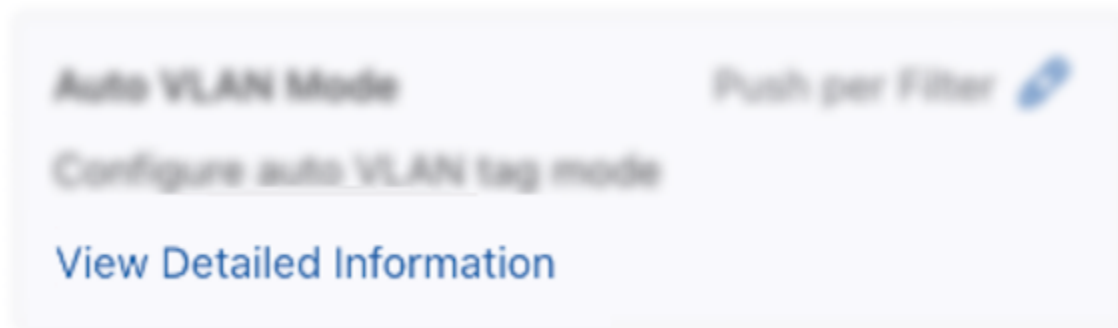
The UI displays the following:

- Feature Title
- A brief description
- View / Hide detailed information link
- Current Setting
- Edit Link - Use the **Edit** configuration button (**pencil icon**) to change the value.

#### View Detailed Information

Each configuration option has detailed information. For more details, click the **View Detailed Information** link on each card.

**Figure 1-11: View Detailed Information**



### Feature Settings

#### Auto VLAN Strip

1. A **toggle button** controls the configuration of this feature. Locate the corresponding card and click the **toggle** switch.

**Figure 1-12: Toggle Switch**




2. A confirm window pops up, displaying the corresponding prompt message. Click the **Enable** button to confirm the configuration changes or the **Cancel** button to cancel the configuration. Conversely, to disable the configuration, click **Disable**.

**Figure 1-13: Confirm / Enable**



3. Review any warning messages that appear in the confirmation window during the configuration process.

**Figure 1-14: Warning Message - Changing**




The following feature sets work in the same manner as the **Auto VLAN Strip** feature described above.

- **CRC Check**
- **Custom Priority**
- **Inport Mask**
- **Policy Overlap Limit Strict**
- **Retain User Policy VLAN**
- **Tunneling**

#### **Auto VLAN Mode**

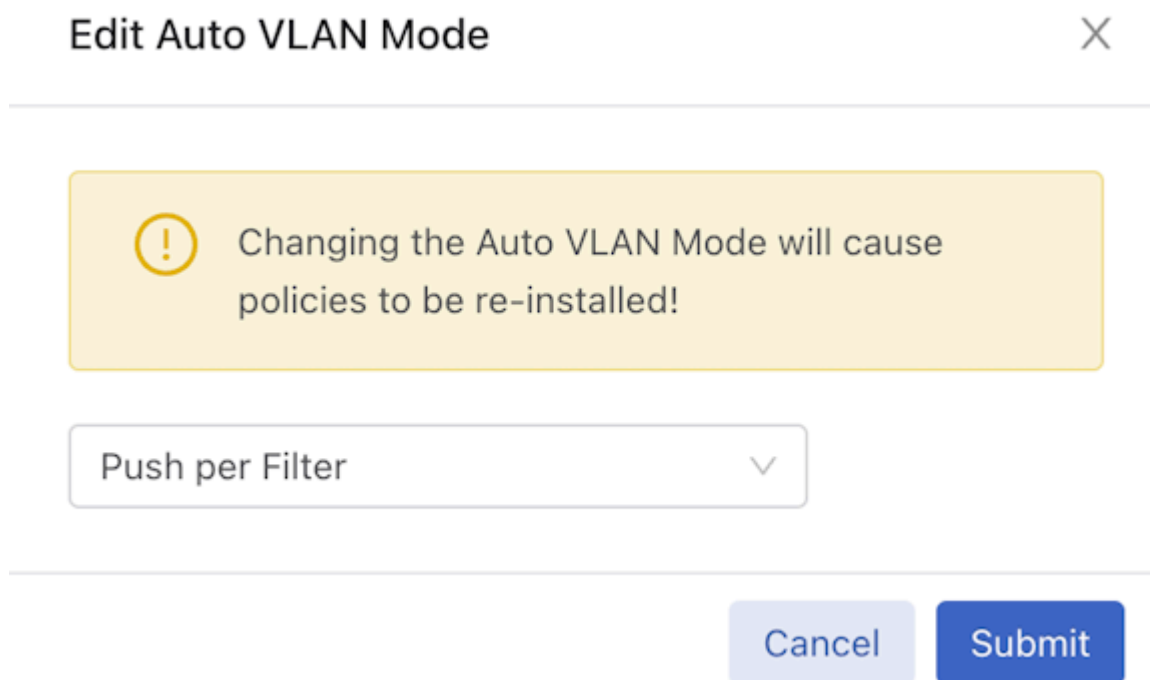
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 1-15: Auto VLAN Mode Config**



2. A confirmation edit dialogue window appears displaying the corresponding prompt message.

**Figure 1-16: Edit VLAN Mode**



The dialog box is titled "Edit Auto VLAN Mode" with a close button (X) in the top right corner. Below the title bar is a yellow warning box containing a yellow exclamation mark icon and the text: "Changing the Auto VLAN Mode will cause policies to be re-installed!". Below the warning box is a dropdown menu with the text "Push per Filter" and a downward arrow. At the bottom right of the dialog are two buttons: "Cancel" (light blue) and "Submit" (dark blue).

3. To configure different modes, click the drop-down arrow to open the menu.


**Figure 1-17: Drop-down Example**



A single dropdown menu with the text "Push per Filter" and a downward arrow.

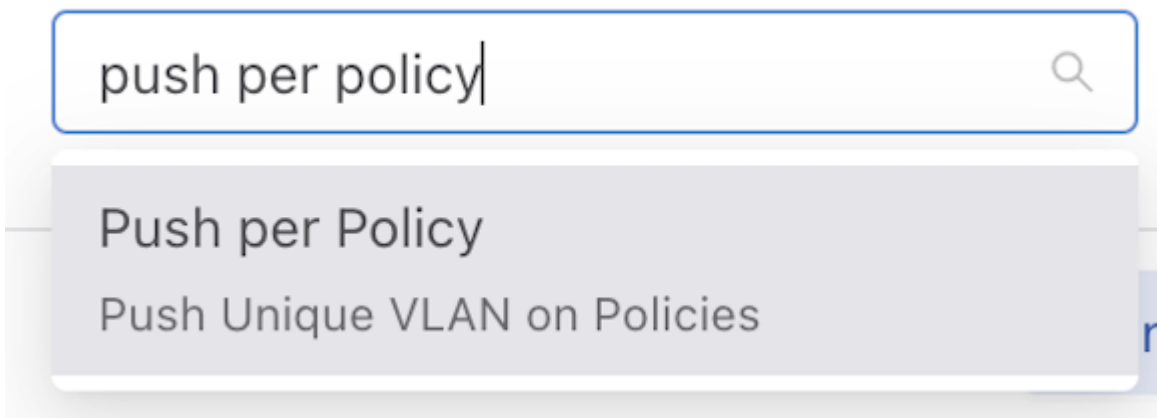
4. From the drop-down menu, select and click on the desired mode.

**Figure 1-18: Push Per Policy**



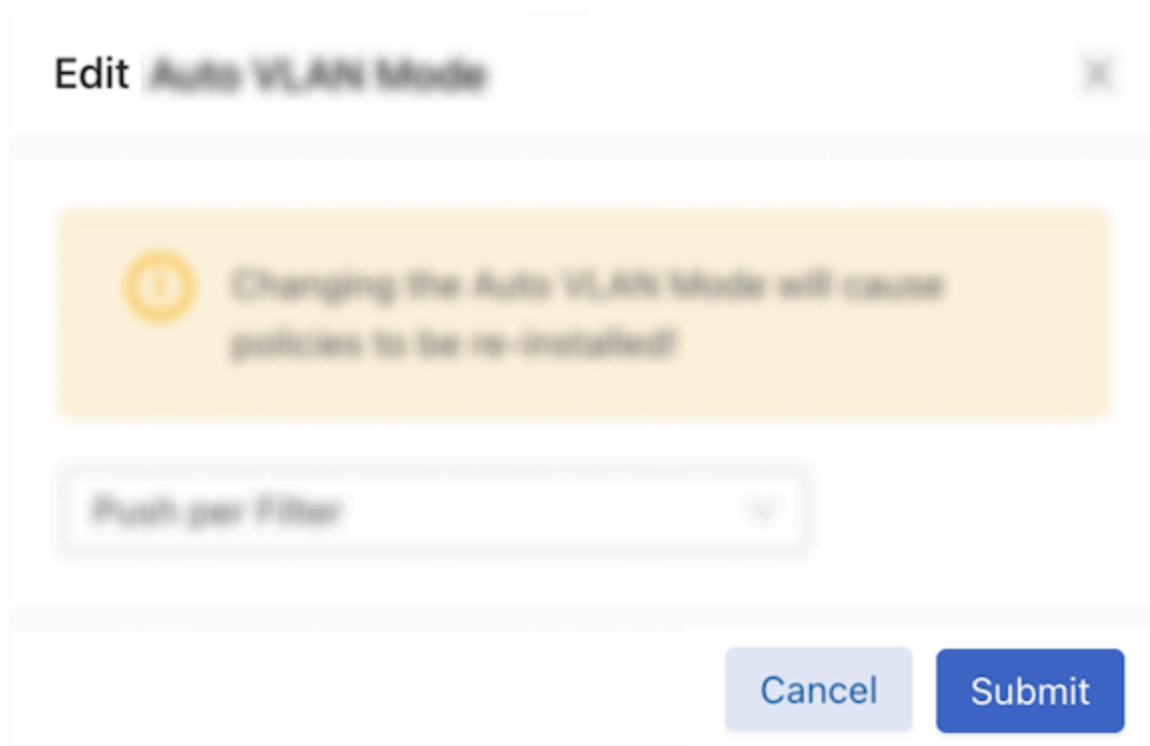
5. Alternatively, you can directly input the desired mode name in the input area.

**Figure 1-19: Push Per Policy**



- Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.

**Figure 1-20: Submit Button**



- The current configuration status displays next to the edit button after successfully setting the configuration.

**Figure 1-21: Current Configuration Status**



The following feature sets work in the same manner as the **Auto VLAN Mode** feature described above.

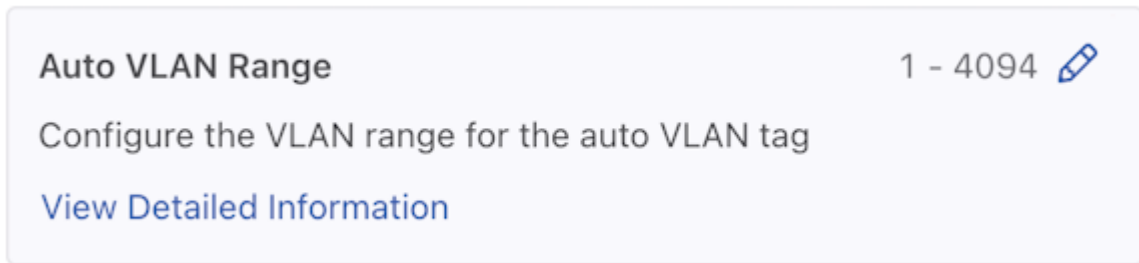
- **Device Deployment Mode**
- **Match Mode**

---

## Auto VLAN Range

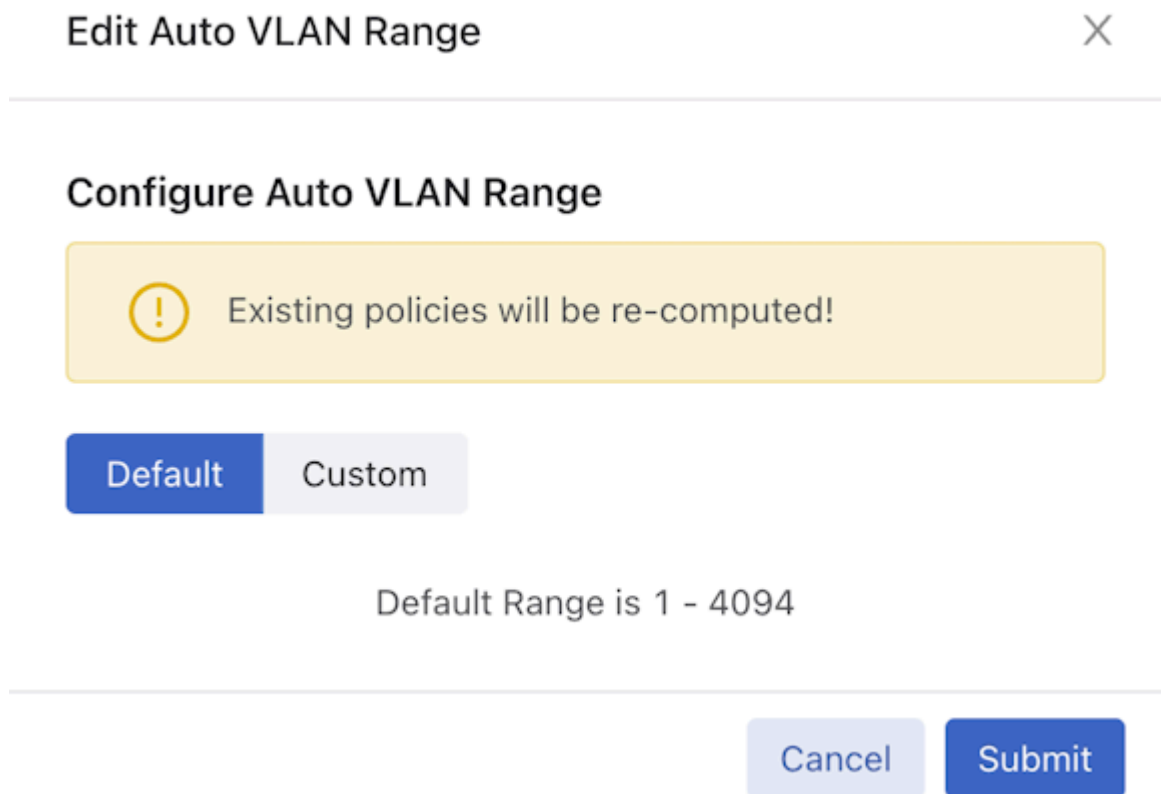
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 1-22: Edit Auto VLAN Range**



2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. The Auto VLAN Range defaults to 1 - 4094.

**Figure 1-23: Edit Auto VLAN Range**





- Click on the **Custom** button to configure the custom range.

**Figure 1-24: Custom Button**

The interface shows two buttons at the top: 'Default' and 'Custom'. The 'Custom' button is highlighted in blue. Below the buttons is a slider control. On the left of the slider is a text box containing the number '1'. On the right is a text box containing the number '4094'. A blue horizontal line connects the two text boxes, with a circular knob positioned at the far right end, indicating the current range.

- Adjust range value (minimum value: 1, maximum value: 4094). There are three ways to adjust the value of a range:
  - Directly enter the desired value in the input area, with the left side representing the minimum value of the range and the right side representing the maximum value.
  - Adjust the value by dragging the **slider** using a mouse. The left knob represents the minimum value of the range, while the right knob represents the maximum value.
  - Use the up and down arrow buttons in the input area to adjust the value accordingly. Pressing the up arrow increments the value by 1, while pressing the down arrow decrements it by 1.
- Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.
- After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 1-25: Configuration Change Success**

The interface shows a configuration card for 'Auto VLAN Range'. The title 'Auto VLAN Range' is on the left, and the value '200 - 500' is on the right next to a pencil icon. Below the title is the description 'Configure the VLAN range for the auto VLAN tag'. At the bottom is a link 'View Detailed Information'.

### Policy Overlap Limit

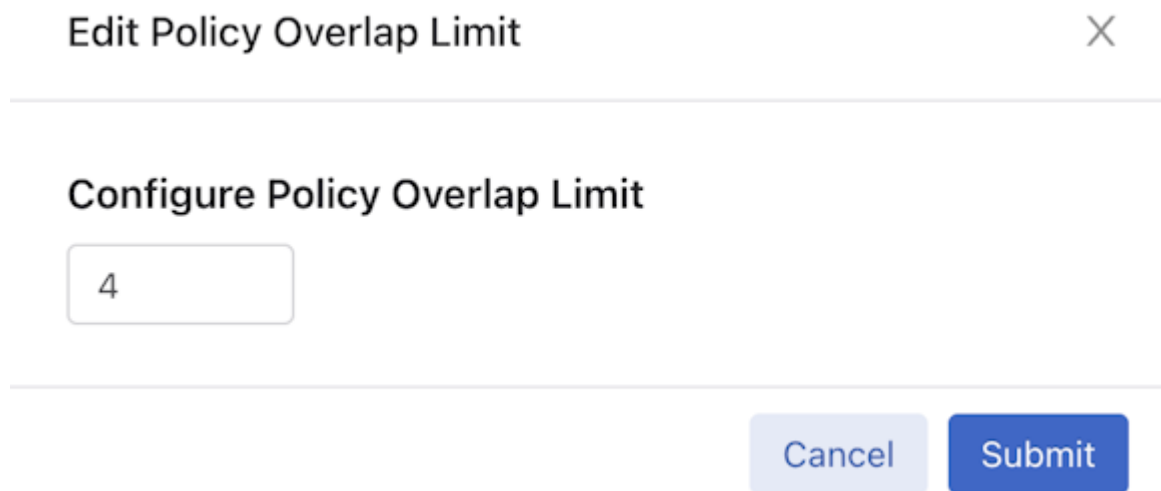
- Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 1-26: Policy Overlap Limit**

The interface shows a configuration card for 'Policy Overlap Limit'. The title 'Policy Overlap Limit' is on the left, and the value '5' is on the right next to a pencil icon. Below the title is the description 'Configure the number of overlapping policies allowed'. At the bottom is a link 'View Detailed Information'.

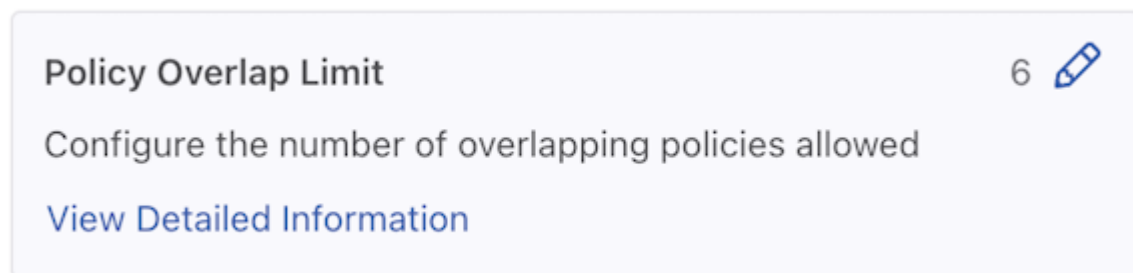
2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. By default, the Policy Overlap Limit is 4.

**Figure 1-27: Edit Policy Overlap Limit**



3. Adjust the Value (minimum value: 0, maximum value: 10). There are two ways to adjust the value:
  - Directly enter the desired value in the input area.
  - Use the up and down arrow buttons in the input area to adjust the value accordingly. Pressing the up arrow increments the value by 1, while pressing the down arrow decrements it by 1.
4. Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.
5. After successfully setting the configuration, the current configuration status displays next to the edit button.

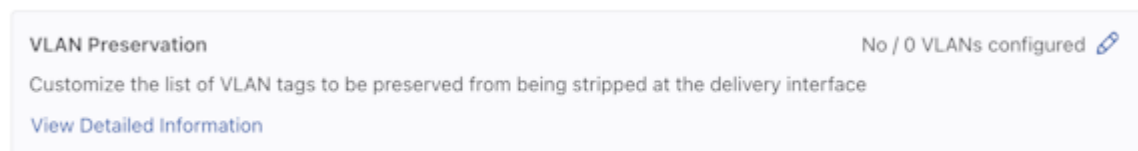
**Figure 1-28: Policy Overlap Limit Change Success**



### VLAN Preservation

1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 1-29: VLAN Preservation Feature Set**



2. A configuration edit dialogue window appears displaying the corresponding prompt message. The VLAN Preservation defaults to:

- Preserve User Configured VLANs: Off
  - Preserve VLAN: No VLAN Configured
3. To configure **Preserve User Configured VLANs**, **toggle** on the switch.

**Figure 1-30: Edit VLAN Preservation Configuration**

Edit VLAN Preservation Configuration

Preserve User Configured VLANs

VLAN Preserve

No VLAN Configured

+ Add VLAN

Cancel Submit

4. To configure **Preserve VLAN**, click the **Add VLAN** button to add a configuration area for preserving the VLAN value.

**Figure 1-31: Preserve VLAN - Add VLAN**

Preserve VLAN

+ Add VLAN

5. Click the drop-down button. There are two ways to configure the preserved VLAN value (minimum value: 1, maximum value: 4094) and a method to delete an entry.

**Figure 1-32: VLAN Single Example**

VLAN

Single

Range


- **Add Single:** Choose Single in the VLAN drop-down menu, and type in the value in the input area.

**Figure 1-33: Add Single VLAN**

- **Add Range:** Choose Range in the VLAN drop-down menu, and type in the input area's minimum and maximum values.

**Figure 1-34: Add VLAN Range**

- **Delete:** Since there must be a corresponding number in the value input area when submitting the configuration, when accidentally adding multiple redundant VLAN configuration areas, delete the

corresponding rows by clicking the red trash can icon .



**Note:** The feature supports combinations of any number of single values and any number of range values.

6. Click the **Submit** button confirm the configuration changes or the **Cancel** button to discard the changes.
7. After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 1-35: Preserve VLAN Configuration Change**

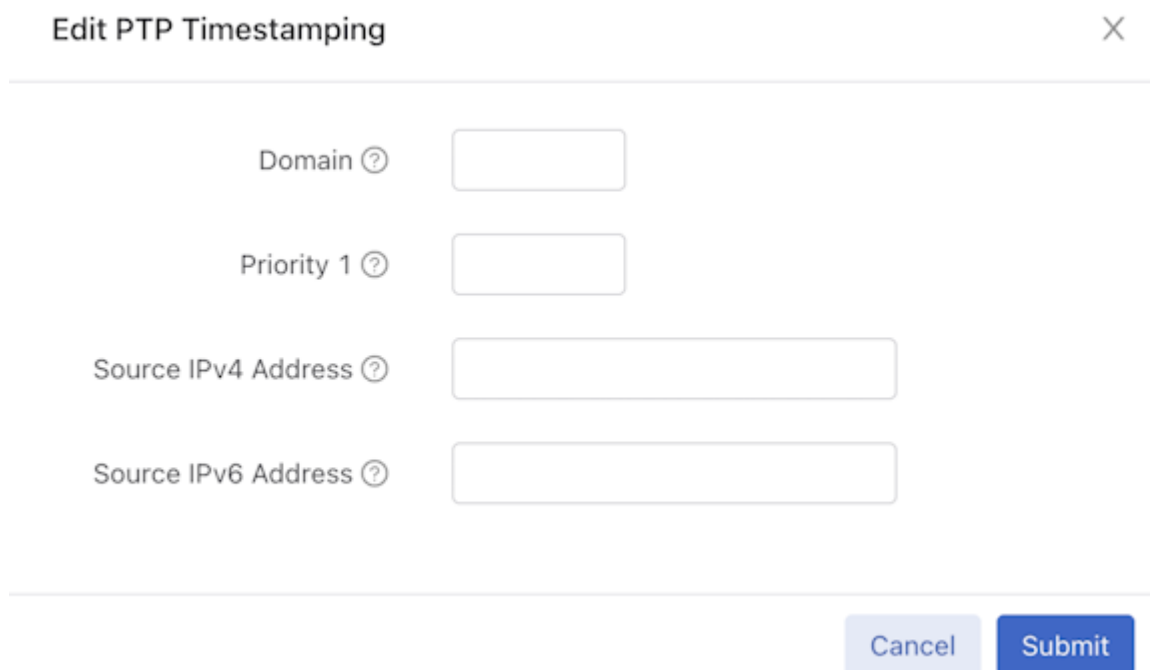
## PTP Timestamping

1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 1-36: PTP Timestamping**

2. A configuration edit dialogue window appears displaying the corresponding prompt message. By default, these features are not configured. Enter the desired configuration value in the corresponding input area. Hover over the **question mark** icon to obtain additional explanatory information.

**Figure 1-37: Edit PTP Timestamping**



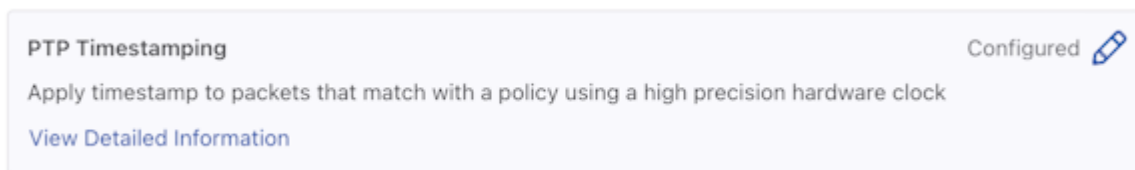
The dialog box titled "Edit PTP Timestamping" has a close button (X) in the top right corner. It contains four input fields, each with a question mark icon for help:

- Domain ?
- Priority 1 ?
- Source IPv4 Address ?
- Source IPv6 Address ?

At the bottom right, there are two buttons: "Cancel" and "Submit".

3. Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.
4. After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 1-38: PTP Timestamping Configuration Change**

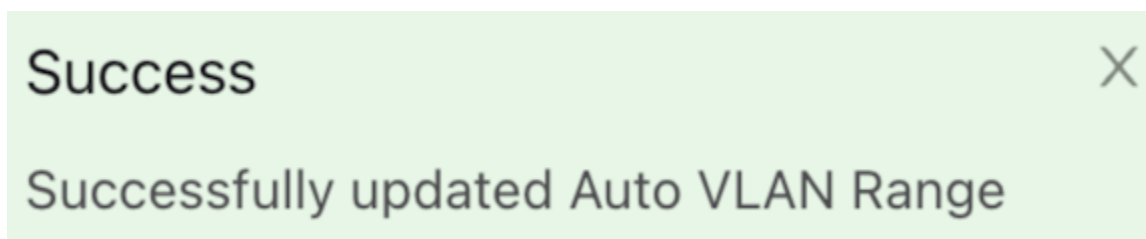


The notification box shows the configuration status for "PTP Timestamping". It includes the text "Apply timestamp to packets that match with a policy using a high precision hardware clock" and a link "View Detailed Information". The status is "Configured" with a pencil icon.

### Feature Setting Notification Message

Whenever successfully configuring a feature, a success notification message pops up with specific details.

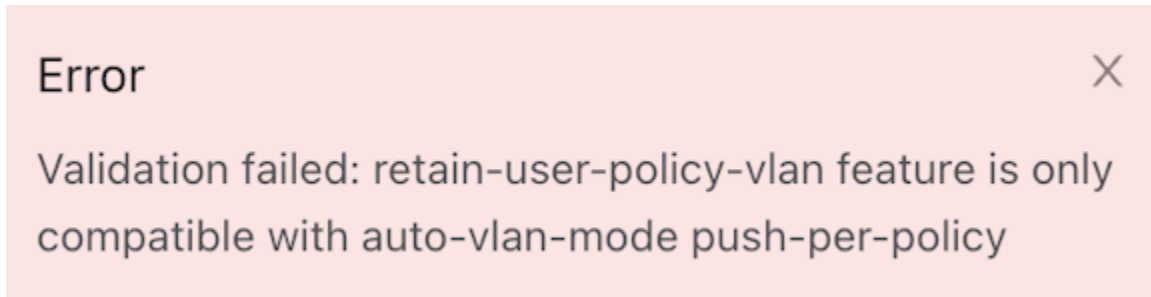
**Figure 1-39: Success Message**



The success notification message is displayed in a green box with a close button (X) in the top right corner. It contains the text "Success" and "Successfully updated Auto VLAN Range".

Whenever an error occurs during the configuration of a feature, an error notification message pops up along with specific details about the error.


**Figure 1-40: Failure Message**



## 1.4.2 Dashboard Layout

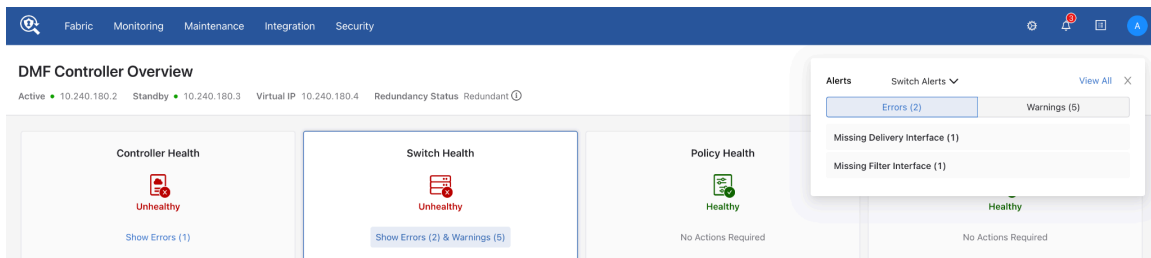
The dashboard data displays four tabs: Controller Health, Switch Health, Policy Health, and Smart Node Health. Each tab has health indicators for that category, and accessing the tab displays the relevant data below.

**Figure 1-41: DMF Controller Tabs**



If a category contains errors or warnings, clicking on the message in the tab opens up a details window and displays the number of errors or warnings filtered by tab category.

**Figure 1-42: Filtered by Category**



Review errors by clicking the bell icon on the right side of the Navigation bar, and it will list all fabric errors and warnings instead of filtering by an individual tab.

**Figure 1-43: Notification Bell**



### 1.4.3 Controller Health

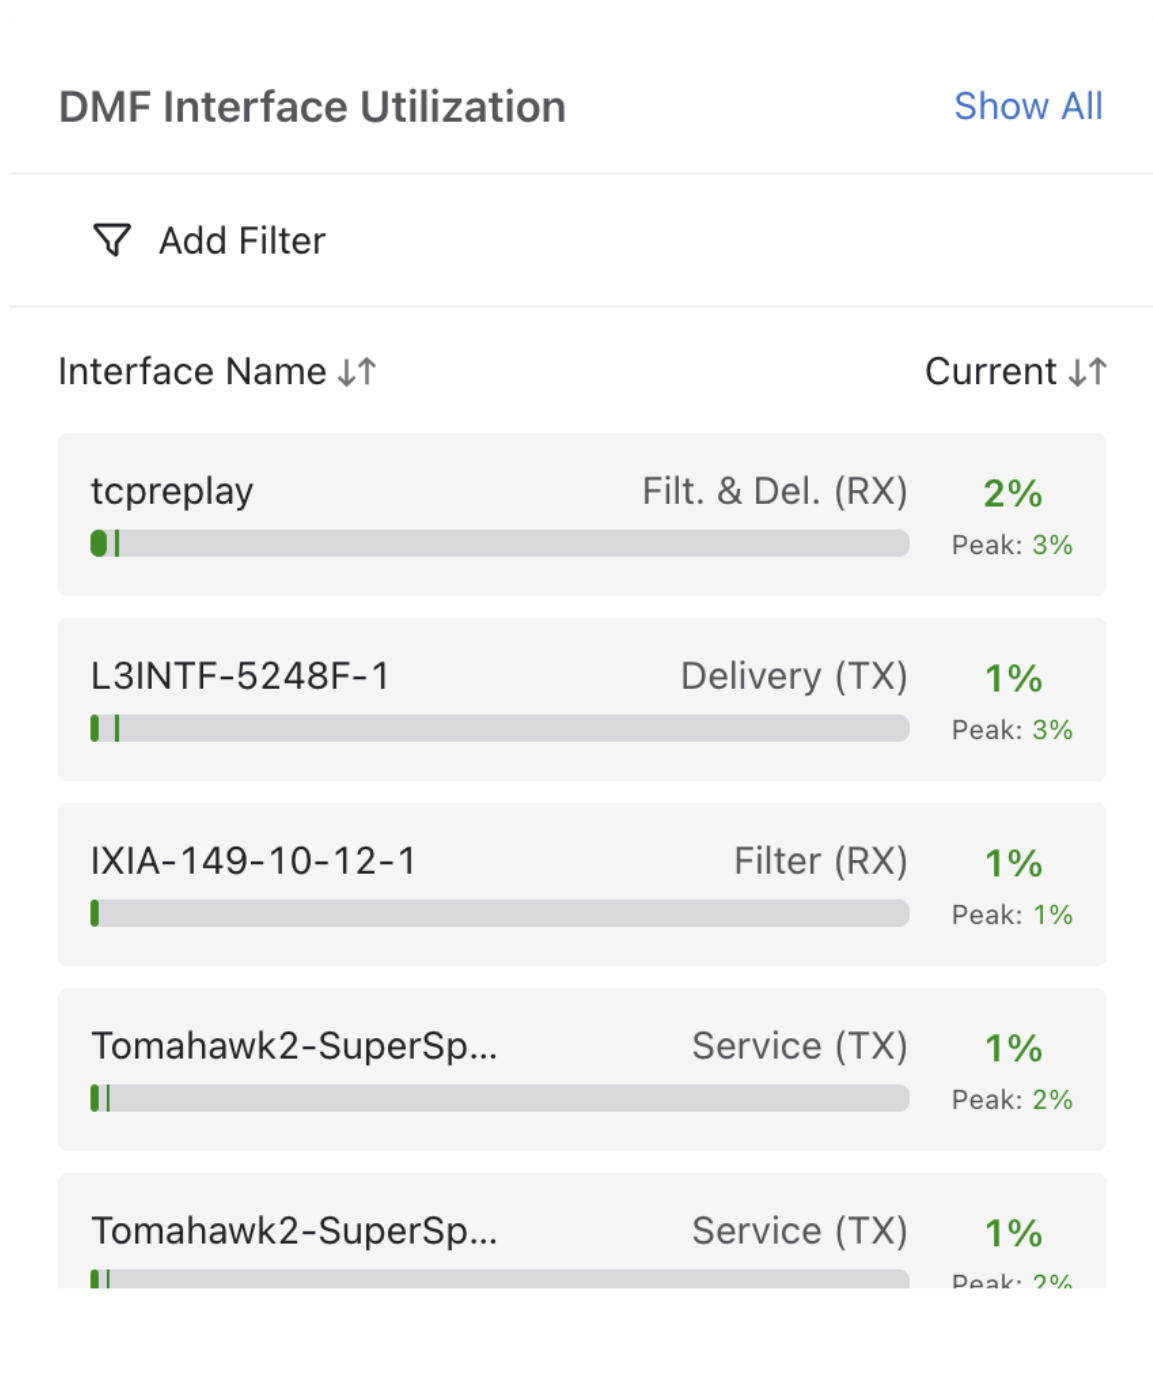
#### **DMF Interface Utilization**

This widget displays the utilization of each DMF interface as follows:

- DMF Interface Name
- Interface Role
- Traffic Direction
- Current Utilization (%)

- Peak Utilization (%)

**Figure 1-44: DMF Interface Utilization**



The bar indicates the current utilization and shows peak utilization with a vertical line. The color of the bar and percentage changes depending on the utilization:

- Red means the utilization percentage is greater than 95%.
- Yellow means the utilization percentage is greater than 70%.
- Green means the utilization percentage is less than 70%.

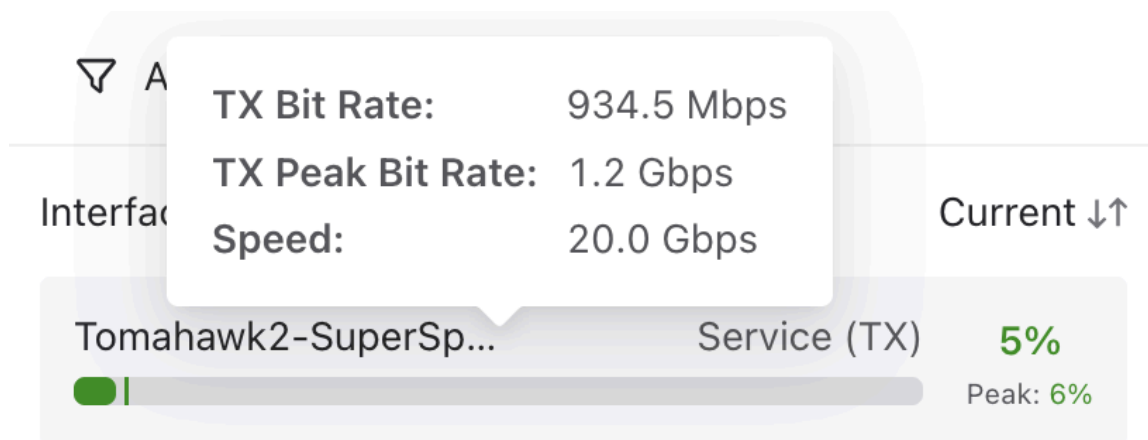
Filter interfaces display only RX traffic, while delivery interfaces display only TX traffic. Other roles with bidirectional data can have one item for each direction of traffic, RX or TX.



The **Show All** button directs you to the DMF Interfaces page.

On hover, the bar shows the interface's Bit Rate, Peak Bit Rate, and Speed in bits per second.

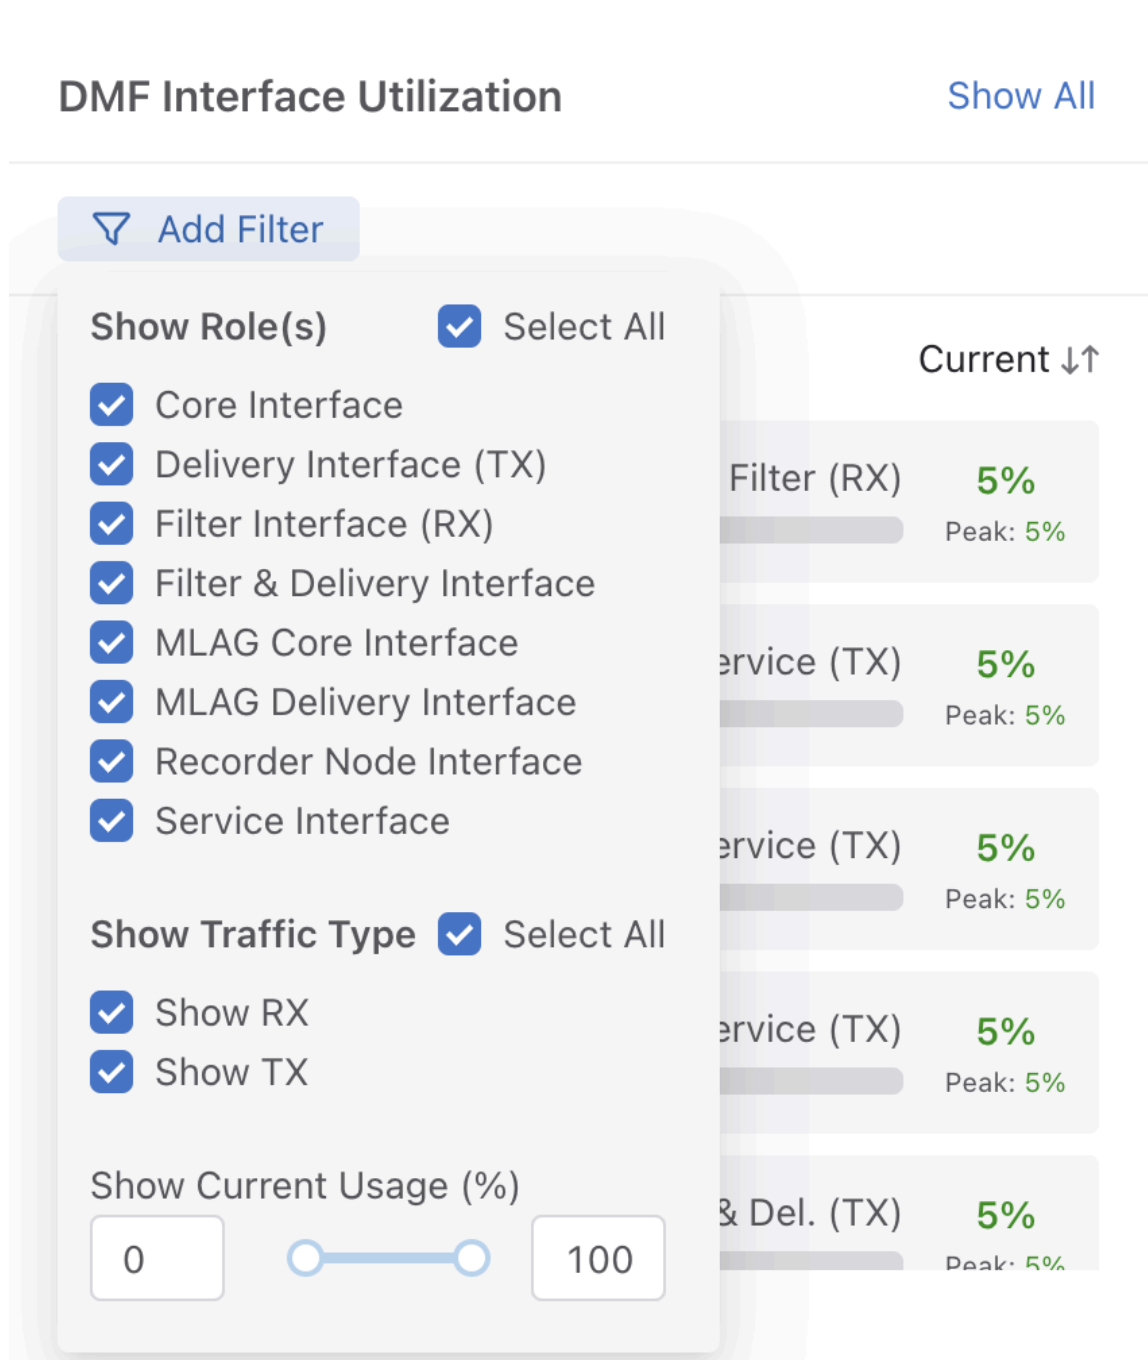
**Figure 1-45: DMF Interface Utilization Hover Details**



Sort interfaces by Interface Name or Current Utilization. The interfaces are sorted by current utilization (descending order) by default.

Display the interfaces by filtering using Role, Traffic Direction, and Current Utilization.


**Figure 1-46: Sort Roles**



Top DMF Interfaces by Traffic

This visualization displays each DMF interface's traffic (bit rate and packet rate).

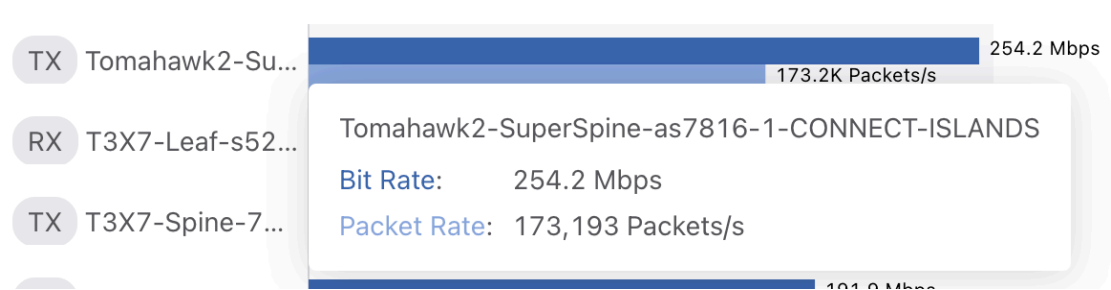
Figure 1-47: Top DMF Interfaces by Traffic



The widget shows each interface's traffic direction, DMF Interface name, bit rate, and packet rate. The **Show All** button directs you to the DMF Interfaces page. Sort interfaces by Bit Rate and filter by Metric and Role. By default, the data is sorted in descending order of bit rate.

On hover, the widget shows the DMF name, bit rate, and packet rate.


Figure 1-48: Top DMF Interfaces by Traffic Hover Details



Top Policies

The widget displays the top policies in DMF. For each policy, traffic is determined by totaling the traffic of each of its configured filter interfaces.

Figure 1-49: DMF Top Policies

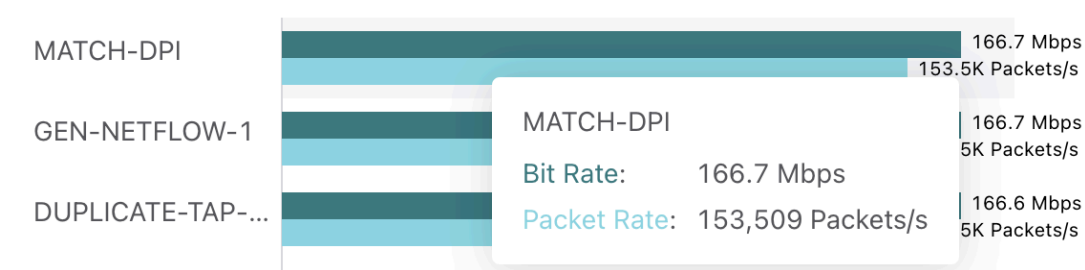


For each policy, the bar chart shows the following:

- Policy Name
- The sum of the bit rates of all filter interfaces associated with the policy.
- The sum of the packet rates of all filter interfaces associated with the policy.

On hover, the bar displays the policy name, bit rate, and packet rate.

Figure 1-50: DMF Policies Hover Details



Sort policies by Bit Rate and filter by Metric. By default, policies are sorted in descending order of bit rate. The **Show All** button directs you to the Policies page.

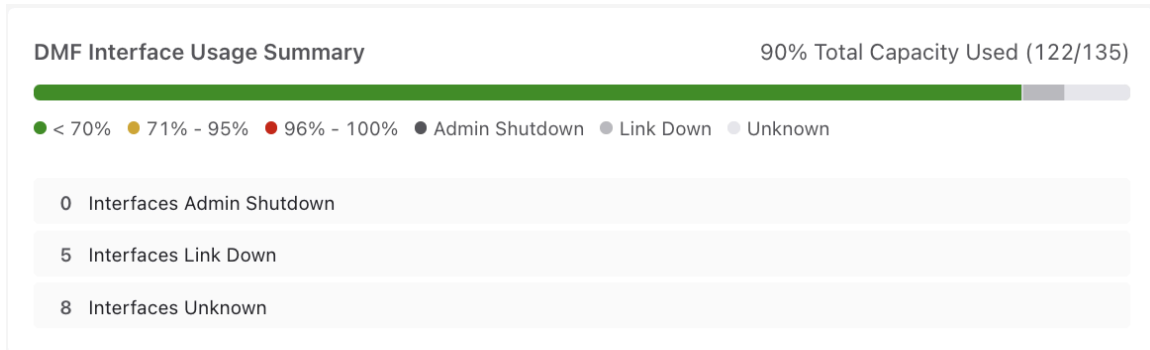
## 1.4.4 Switch Health

### Interface Usage Summary

This widget displays the usage statistics for all DMF interfaces. The interface utilization groups all active interfaces:

- Red means that the utilization percentage is greater than 95%.
- Yellow means that the utilization percentage is greater than 70%.
- Green means that the utilization percentage is less than 70%.

**Figure 1-51: DMF Interface Usage Summary**

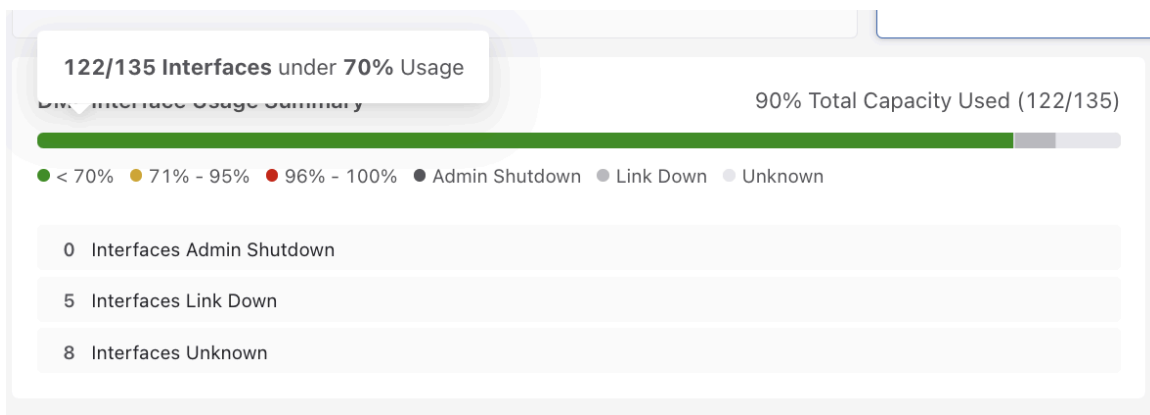


There are three other categories for DMF Interfaces with no traffic. These appear beneath the Usage Bar:

- Admin Shutdown
- Link Down
- Unknown - when Interface Speed is undefined or not known
- Total Capacity Used displays with Total Capacity defined as the number of Active DMF Interfaces divided by the Number of Total DMF Interfaces

On hover, the number of interfaces in each category displays in the respective usage bar.

**Figure 1-52: DMF Interface Usage Hover Details**



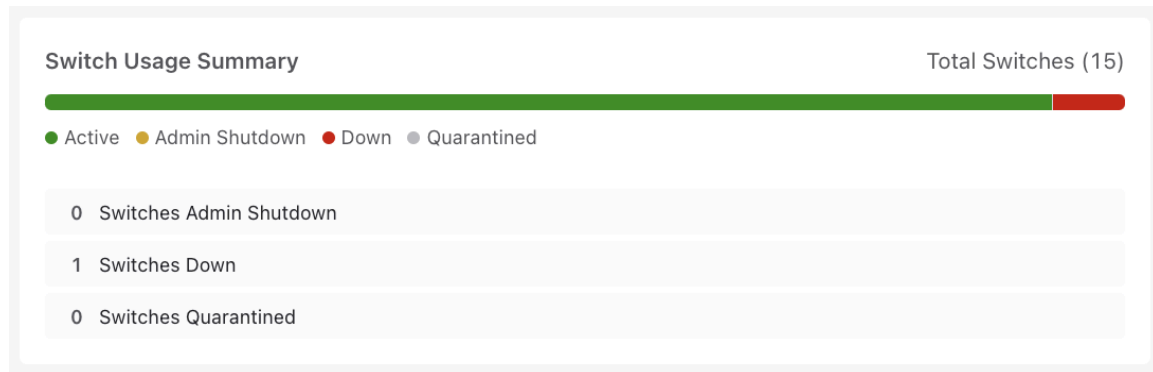
### Switch Usage Summary

This widget displays the usage statistics for each switch. All switches are grouped by:

- Active (Green)
- Admin Shutdown (Yellow)
- Down (Red)

- Quarantined (Grey)

**Figure 1-53: DMF Switch Usage Summary**



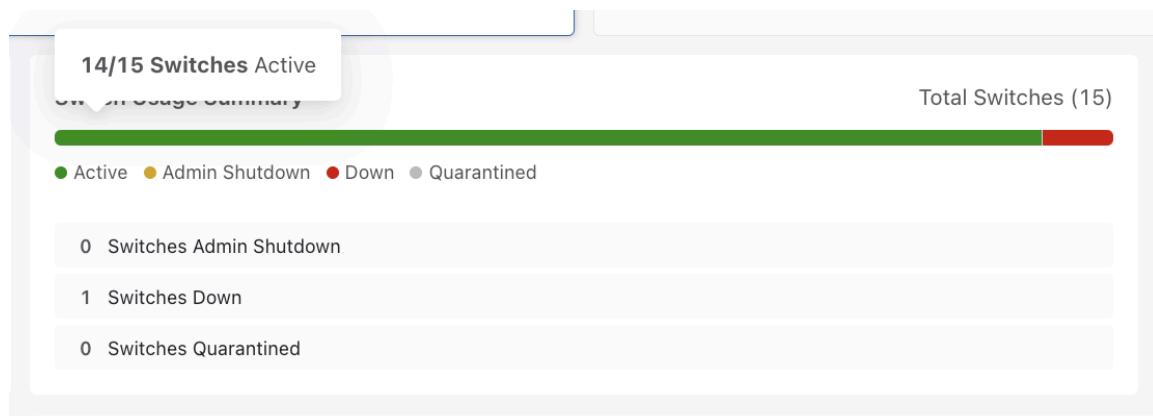
The total number of switches is displayed.

Three list items display the number of:

- Switches Admin Shutdown
- Switches Down
- Switches Quarantined

On hover, the number of switches in each category displays in the respective usage bar.

**Figure 1-54: DMF Switch Usage Hover Details**




## TCAM Usage Summary

This widget displays the usage statistics for the TCAM of each switch and groups all active TCAMs by usage:

- Red means that the utilization percentage is greater than 95%.
- Yellow means that the utilization percentage is greater than 70%.
- Green means that the utilization percentage is less than 70%.
- Grey means that the utilization is Unknown.

- A switch is grouped in the Unknown category when no TCAM usage statistics are available, generally from a switch being shut down or disconnected.

**Figure 1-55: DMF TCAM Usage Summary**



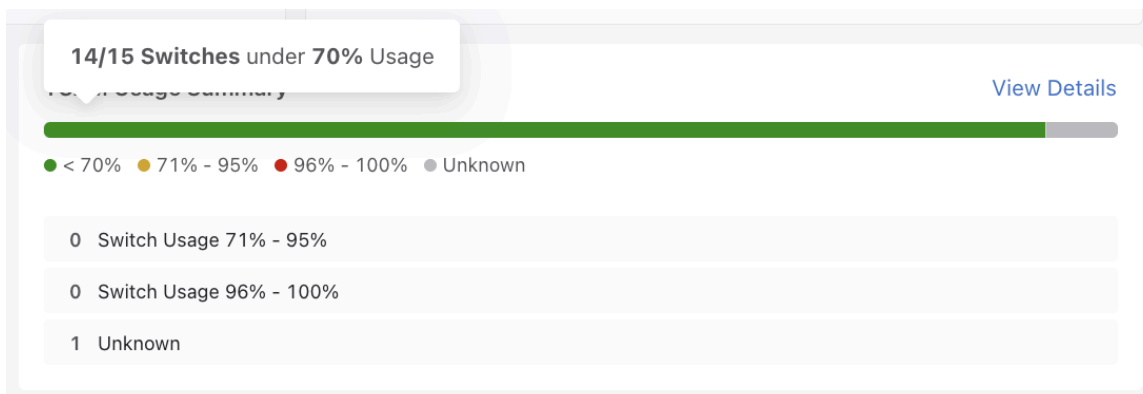
The **View Details** link directs you to the TCAM Utilization tab of the Switches page.

Below the Usage Bar, there are three list items displaying:

- Switch Usage 71% - 95%
- Switch Usage 96% - 100%
- Unknown

On hover, the number of switches in each category displays in the respective usage bar.

**Figure 1-56: DMF TCAM Utilization Hover Details**



## DMF Interface Utilization

DMF Interface Utilization is similar to the data displayed in the Controller Health tab. Please refer to its description for more information.

### Switch Utilization

This widget contains two tabs:

- Switch Usage
- TCAM Usage

#### Switch Usage

The Switch Usage tab of the Switch Utilization displays essential information for each switch, including the use of each switch interface and alerts for any warnings or errors.

**Figure 1-57: DMF Switch Usage Tab**



The widget displays the following data for each switch:

- Switch Name (contains a link to the Switches page for that specific switch).
- Switch Usage: Each section represents the number of interfaces with a specific role.
- Total Usage: Displays the Number of Interfaces with an assigned role divided by the Total Number of Interfaces on the switch.
- Alerts: This column displays any alerts related to interfaces.
  - The yellow badge indicates the number of warnings, while the red badge shows the number of errors.

The Switch Usage column contains the number of interfaces for each role:

- Filter
- Delivery
- Filter & Delivery
- Core
- Recorder Node
- Service
- PTP
- MLAG Core
- MLAG Delivery.

Three columns can sort the table:

- Sort the Switch Name column in alphabetical order.
- Sort the Total Usage column by percentage (%) usage (# used interfaces / # total interfaces).
- Sort the Alerts column by the total number of alerts (# warnings + # errors).

The default sort order for this table is the Alerts column in descending order and ensures the switches with the highest number of alerts are initially at the top.



On hover, the number of each interface displays.

**Figure 1-58: DMF Switch Usage Filter Interfaces Hover Details**

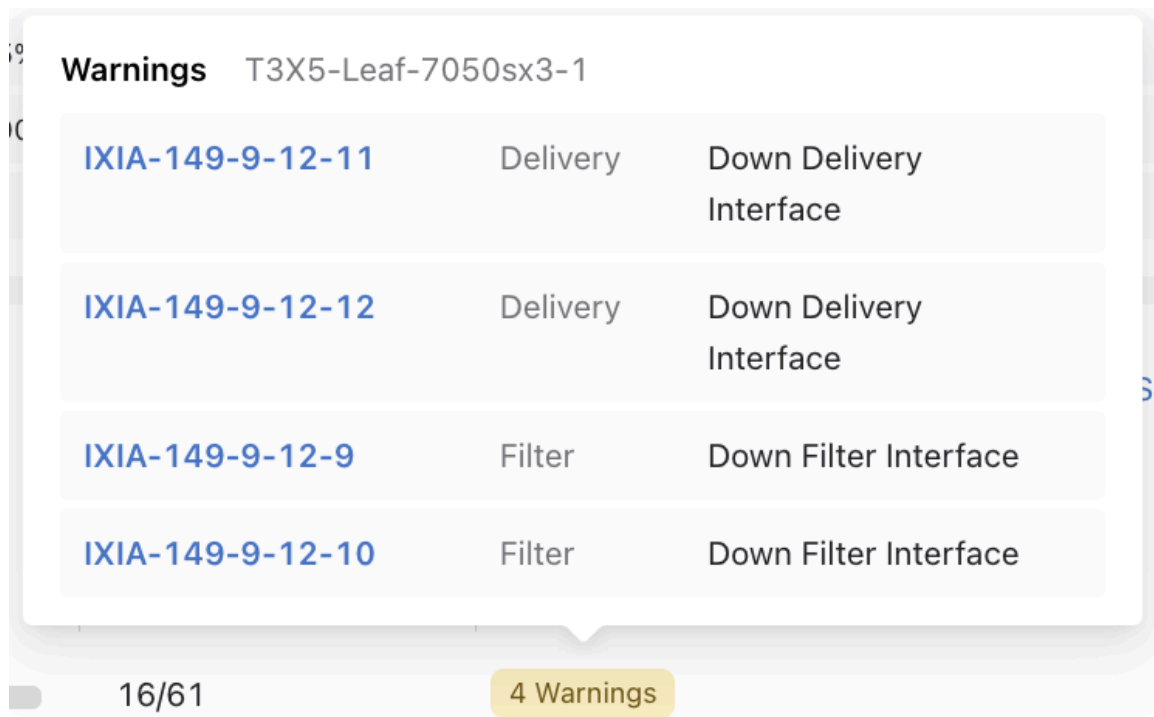


While hovering over the warnings or alerts badge, a table appears and displays Warnings for the yellow badge and Errors for the red badge, and it will also show the switch name.

Each row of the table contains the following:

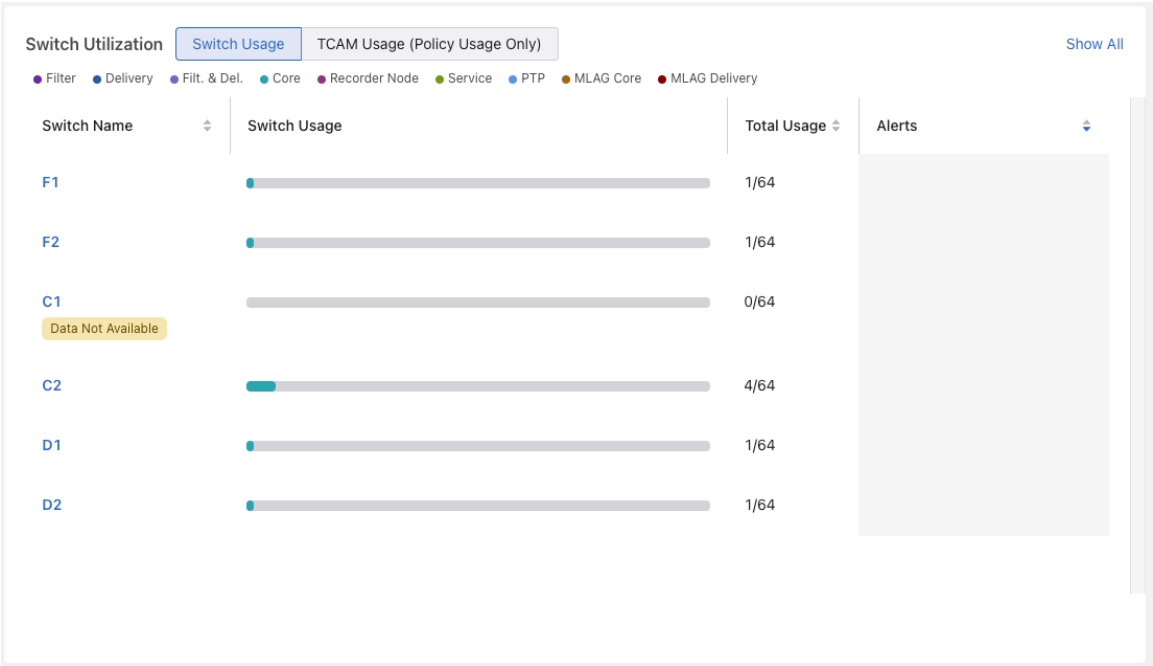
- Interface name (includes link to Interfaces/[INTERFACE-NAME] page)
- Interface role
- Alert type (e.g., Down Delivery Interface)

**Figure 1-59: DMF Interface Warnings**



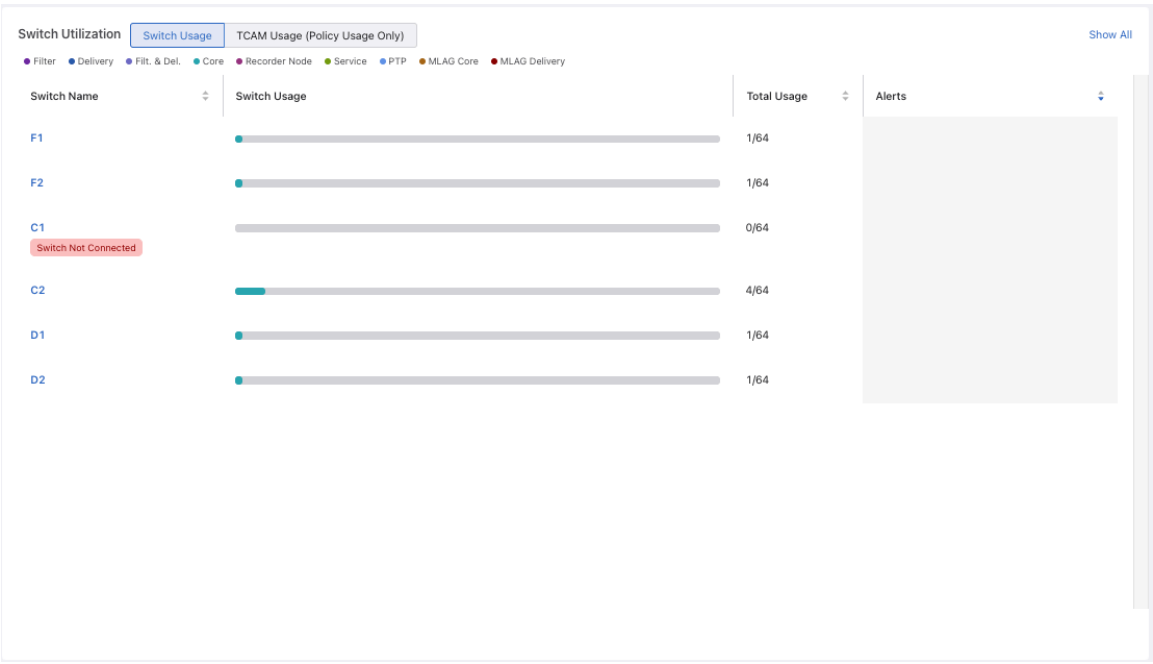
When data is unavailable for a switch (C1), there will be a yellow badge under the Switch Name that says **Data Not Available**. The Switch Usage column will have an empty usage bar, the Total Usage Column will show 0 (zero) for the number of currently used interfaces, and the Alerts column will be empty.

**Figure 1-60: DMF Switch Usage - Data Not Available**



When a switch is down, a red badge appears under the Switch Name that says **Switch Not Connected**. The other columns will be empty in the same way as the **Data Not Available** case.

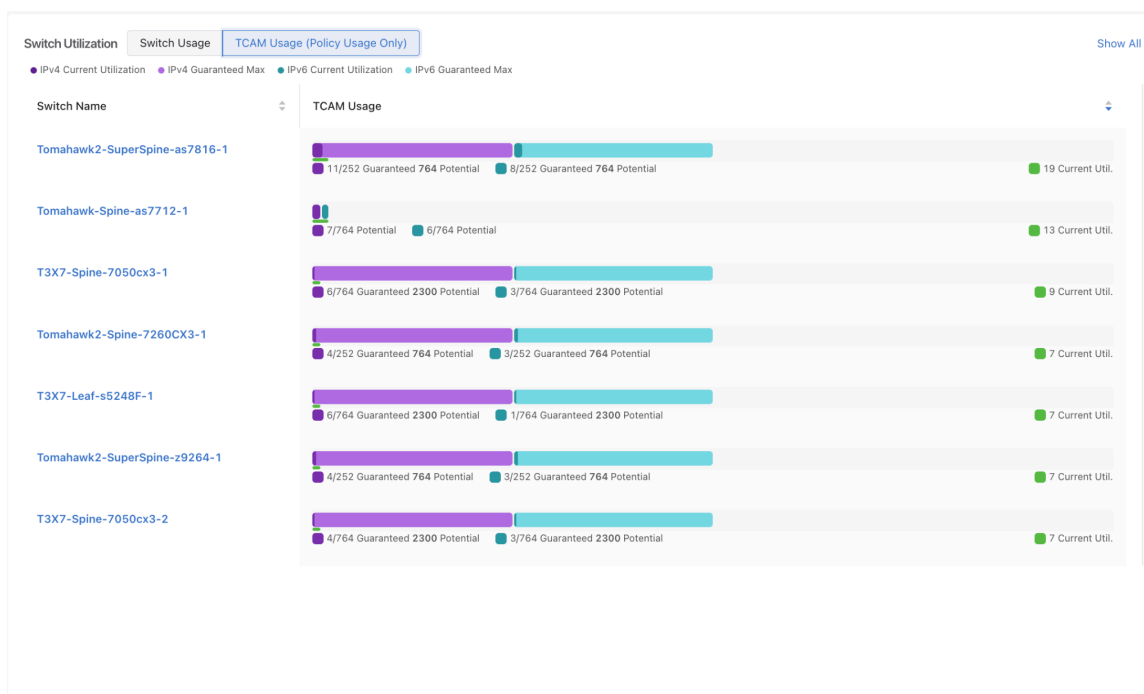
**Figure 1-61: DMF Switch Usage - Switch Not Connected**



**TCAM Usage**

The TCAM Usage widget displays the current utilization of the TCAMs for each active switch. A switch can have a TCAM for IPv4, IPv6, or both. Each TCAM has a guaranteed maximum usage and current utilization. This table compares the current utilization of each TCAM to its guaranteed maximum.

**Figure 1-62: DMF Usage (Policy Usage Only)**



This widget displays a TCAM Usage chart for each switch:

- The purple bar shows IPv4 Current Utilization and Guaranteed Maximum.
- The cyan bar shows IPv6 Current Utilization and Guaranteed Maximum.
- Each row will display Current Utilization (IPv4 + IPv6 Current Utilization)
- Sort by Switch Name and Current Utilization.
- Sort the Switch Name column alphabetically (descending and ascending).
- Sort the Current Utilization column in descending and ascending order (IPv4 + IPv6 Current Utilization).
- The default sort order for the table is the Current Utilization column in descending order, ensuring the switches with the highest current utilization display first.


### Top DMF Interfaces by Traffic

The visualization shows DMF interface traffic (bit rate and packet rate) color-coded by interface role. The roles displayed are:

- Core
- Delivery
- Filter
- Filter & Delivery
- MLAG Core
- MLAG Delivery
- Recorder Node

- Service

**Figure 1-63: DMF Top Interfaces by Traffic**




For each interface, the chart item shows:

- Interface role
- Traffic direction
- DMF interface name
- Bit rate
- Packet rate

The **Show All** button directs you to the DMF Interfaces page.

Sort the interfaces by bit rate, which, by default, are sorted in descending bit rate order. Filter interfaces by interface role using the drop-down.

**Figure 1-64: DMF Sort Interfaces by Bit Rate**



## 1.4.5 Policy Health


### Policies Usage by Traffic

This widget displays policy traffic. For each policy, the bar chart shows:

- Name of the policy
- Bit rate

- Packet rate

**Figure 1-65: DMF Policies Usage by Traffic**



On hover, similar information displays.

Sort policies by Bit Rate.

The **Show All** button directs you to the Policies page.

### Active Interfaces by Policy

The table displays DMF interfaces associated with policies. DMF interfaces that are not associated with a policy are not displayed.

**Figure 1-66: DMF Active Interfaces by Policy**

Active Interfaces by Policy Roles (All) Interfaces (All) Show All

Last updated: a few seconds ago

DMF Interface Name	Role	Policy Name(s)	Bit Rate	Packet Rate
TO-EOS-T4SWL-2	Delivery	MATCH-DPI	0.0 bps	0.0 Packets/s
vc-Tomahawk-Spine-as7712-1-filter-vcenter-c4f259a4	Filter	VC-ERSPAN	0.0 bps	0.0 Packets/s
vc-Tomahawk-Spine-as7712-1-filter-vcenter-c4f25c10	Filter	VC-ERSPAN	0.0 bps	0.0 Packets/s
vc-Tomahawk-Spine-as7712-1-filter-vcenter-c4f25c2f	Filter	VC-ERSPAN	0.0 bps	0.0 Packets/s
vc-Tomahawk-Spine-as7712-1-filter-vcenter-c6334690	Filter	VC-ERSPAN	0.0 bps	0.0 Packets/s
Tomahawk2-SuperSpine-as7816-1-MS-LAG-7-to-managed-service	Service	MATCH-DPI	RX: 40.0 bps TX: 216.0 bps	RX: 0.00 Packets/s TX: 0.00 Packets/s
Tomahawk2-SuperSpine-as7816-1-MS-LAG-5-to-managed-service	Service	GEN-DAPPER-1	RX: 40.0 bps TX: 216.0 bps	RX: 0.00 Packets/s TX: 0.00 Packets/s
Tomahawk2-SuperSpine-as7816-1-MS-LAG-3-to-managed-service	Service	GEN-IPFIX-1	RX: 40.0 bps TX: 216.0 bps	RX: 0.00 Packets/s TX: 0.00 Packets/s
Tomahawk2-SuperSpine-as7816-1-MS-LAG-1-to-managed-service	Service	GEN-NETFLOW-1	RX: 40.0 bps TX: 216.0 bps	RX: 0.00 Packets/s TX: 0.00 Packets/s
Tomahawk-Spine-as7712-1-ethernet21	Core	DUPLICATE-TAP-FLOWS	RX: 104.0 bps TX: 187.9 Mbps	RX: 0.00 Packets/s TX: 186.45K Packets/s

The table contains the following columns:

- DMF Interface Name: The DMF name of the switch interface.
- Role: The role of the interface.
- Policy Name(s): A list of the policies associated with the interface.
- Bit Rate: The bit rate of the interface.
- Packet Rate: The packet rate of the interface.

The **Show All** button directs you to the DMF Interfaces page.

Sort the table by each column; the items are sorted in descending bit rate order by default.

Two filters, **Roles** and **Interfaces**, allow data sorting by interface role and DMF interface name.

**Figure 1-67: DMF Active Interfaces by Policy - Roles**

Active Interfaces by Policy

DMF Interface Name	Role	Policy Name(s)	Bit Rate	Packets/s
TO-RR-SETUP	Delivery	DUPLICATE-TAP-FLOWS	669.1 Mbps	125.1 kbps
LOOP-TOM-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS	RX: 658.1 Mbps TX: 658.1 Mbps	125.1 kbps
LOOP-TOM2-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS	RX: 650.9 Mbps TX: 650.9 Mbps	RX: 129.63K Packets/s TX: 129.63K Packets/s

Roles (All) Interfaces Select All Show All

Show Role(s) Select All

Core Interface  
Delivery Interface (TX)  
Filter Interface (RX)  
Filter & Delivery Interface  
MLAG Core Interface  
MLAG Delivery Interface  
Recorder Node Interface  
Service Interface

**Figure 1-68: DMF Active Interfaces by Policy - Interfaces**

Active Interfaces by Policy

DMF Interface Name	Role	Policy Name(s)	Bit Rate	Packets/s
_MLAG-DELIVERY-SX3-S5248_T3X7-Leaf-s5248F-1_ethernet32	Mag Delivery	DUPLICATE-TAP-FLOWS _mlag_DUPLICATE-TAP-FLOWS_T3X7-Leaf-s5248F-1	357.8 Mbps	125.1 kbps
TO-RR-SETUP	Delivery	DUPLICATE-TAP-FLOWS	209.0 Mbps	125.1 kbps
V5-FEED	Filter	DUPLICATE-TAP-FLOWS	188.5 Mbps	125.1 kbps
LOOP-TOM2-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS	RX: 188.1 Mb	125.1 kbps
LOOP-T3X5-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS GEN-NETFLOW-1	RX: 181.8 Mb	125.1 kbps

Roles (All) Interfaces Select All Show All

Select Interface(s)

Search

V5-FEED  
TO-EOS-T4SWL-1  
LOOP-T2-V5-1  
LOOP-T3X7-V5-1  
L3INTF-5248F-1  
LOOP-TOM2-V5-1  
LOOP-TXX2-2  
LOOP-TOM-V5-1  
LOOP-T3X5-V5-1  
TO-RR-SETUP  
IXIA-149-10-12-1  
\_MLAG-DELIVERY-SX3-S5248\_T3X5-Leaf-7050sx3-1\_ethernet33  
\_MLAG-DELIVERY-SX3-S5248\_T3X7-Leaf-s5248F-1\_ethernet32  
Tromawiki?\_SumerSina-ac7816-1-MR-1 AG-F-in-managed-service

## 1.4.6 Smart Node Health

### Recorder Node

The Recorder Nodes table displays Recorder Node health and the following columns:

- Recorder Node Name
- IP Address
- MAC Address
- Recording
  - Indicates the recording configuration status, either Yes or No, for the Recorder Node.
- Storage Utilization
- Index and Packet disk storage utilization % (percentage) using the following colors:
  - Red means the utilization percentage is greater than 95%.
  - Yellow means the utilization percentage is greater than 70%.
  - Green means the utilization percentage is less than 70%.

**Figure 1-69: Recorder Nodes**

Recorder Nodes

Recorder Node Name	IP Address	MAC Address	Recording	Storage Util.	Backup Storage Util.	Virtual Disk Health	Recorder Node Fabric Interface	Zero Touch State	Alerts
PR-NewHW		78:ac:44:86:c2:08	Yes	Unknown	Unknown	Unknown			Errors (1)
RMA-CNrail	10.240.155.22	78:ac:44:4a:a5:28	Yes	Index: 12% Packet: 17%	Index: 0% Packet: 0%	Index: Good Packet: Good	RMA-CNrail-intf	Ok	

On hover, various details display depending on the column selected. These include:

- Free and Total Disk Usage
- Backup Storage Utilization
- Index and Packet backup disk storage utilization % (percentage) using the following colors.
  - Red means the utilization percentage is greater than 95%.
  - Yellow means the utilization percentage is greater than 70%.

- Green means the utilization percentage is less than 70%.
- Virtual Disk Health
- Status of Index and Packet virtual disks:
  - Green means the virtual disk's health is good.
  - Red means the value of the virtual disk's health is bad.
- Recorder Node Fabric Interface
  - Shows the DMF interface name and its status where the Recorder Node connects to the DMF Fabric.
- Switch, Interface, and status
- Zero Touch State
- Alerts
- Errors and warnings for the Recorder Node - Hovering over an error displays additional information about the errors and warnings.

The following are examples of the hover detailed information.

**Figure 1-70: Example - Index Disk Storage**

Recorder Nodes					
Recorder Node Name	IP Address	MAC Address	Recording	Index Disk Storage	Backup Storage Util.
DMF-RACK258-RN-1	10.240.180.120	78:ac:44:86:c6:9c	Yes	Free: 12.8 TB Total: 15.4 TB Index: 16% Packet: 18%	Index: 0% Packet: 0%

**Figure 1-71: Example - Index Backup Disk Storage**

Recorder Nodes					
Recorder Node Name	IP Address	MAC Address	Recording	Storage Util.	Index Backup Disk Storage
DMF-RACK258-RN-1	10.240.180.120	78:ac:44:86:c6:9c	Yes	Index: 23% Packet: 26%	Free: 0.0 B Total: 0.0 B Index: 0% Packet: 0%

**Figure 1-72: Example - Recorder Node Fabric Interface**

Recorder Nodes					
Recorder Node Name	IP Address	MAC Address	Recording	Storage Util.	Backup Storage Util.
DMF-RACK258-RN-1	10.240.180.120	78:ac:44:86:c6:9c	Yes	Index: 23% Packet: 26%	Index: 0% Packet: 0%

**Figure 1-73: Example - Errors**

Recorder Nodes					
Name	IP Address	MAC Address	Recording	Storage Util.	Backup Storage Util.
78:ac:44:86:c2:08			Yes	Unknown	Unknown

The **View All** link takes you to the **Recorder Node** page.

## Service Node

The Service Nodes table displays Service Node health and the following columns:

- Service Node Name

- IP Address
- Service Node Interface Load
- Zero Touch State

**Figure 1-74: DMF Service Nodes**

Service Nodes <span>Show All</span>			
Service Node Name	IP Address	Service Node Interface Load	Zero Touch State
DMF-SN-R740-1	10.240.180.124	Moderate (8) No (2)	Ok

Hovering over the Service Node Interface Load column displays:

- Interface Name
- Service Name
- Action

**Figure 1-75: DMF Service Nodes Hover Details**

Recorder Nodes			
Recorder Node Name	IP Address	MAC Address	
DMF-RACK258-RN-1	10.240.180.120	78:ac:44:86:c6:9c	
Service Nodes			
Service Node Name	IP Address		
DMF-SN-R740-1	10.240.180.124		

Interfaces with Moderate load (5)		
Interface Name	Service Name	Action
sni11	MS-IPFIX-ACTIVE-1	IPFIX
sni1	MS-NETFLOW-ACTIVE-1	NetFlow
sni3	MS-IPFIX-ACTIVE-1	IPFIX
sni5	MS-DAPPER-ACTIVE-1	TCP Analysis
eni0	MS-NETFLOW-ACTIVE-1	NetFlow

The **View All** link takes you to the **Service Node** page.

## Analytics Node

The Analytics Node table displays Analytics Node health and the following columns:

- IP Address: The IP address of the configured Analytics Node.
  - Clicking on the **IP Address** opens the Analytics Node UI.
- Redis Status
  - Displays the status in green if healthy, along with the last updated timestamp.
  - Displays the status in red if unhealthy, along with the latest updated timestamp.
- Replicated Redis Status
  - Displays the status in green if healthy, along with the latest updated timestamp.
  - Displays the status in red if unhealthy, along with the latest updated timestamp

**Figure 1-76: DMF Analytics Node**

Analytics Node <span>View Details</span>		
IP Address	Redis Status	Replicated Redis Status
<a href="#">10.240.180.5</a>	Healthy Last updated: 5 seconds ago	Healthy Last updated: 5 seconds ago



The **View Details** link takes you to the **Analytics Node** details page.

## Refreshing Data

Data automatically refreshes every minute, and interface topology data automatically refreshes every 10 seconds.

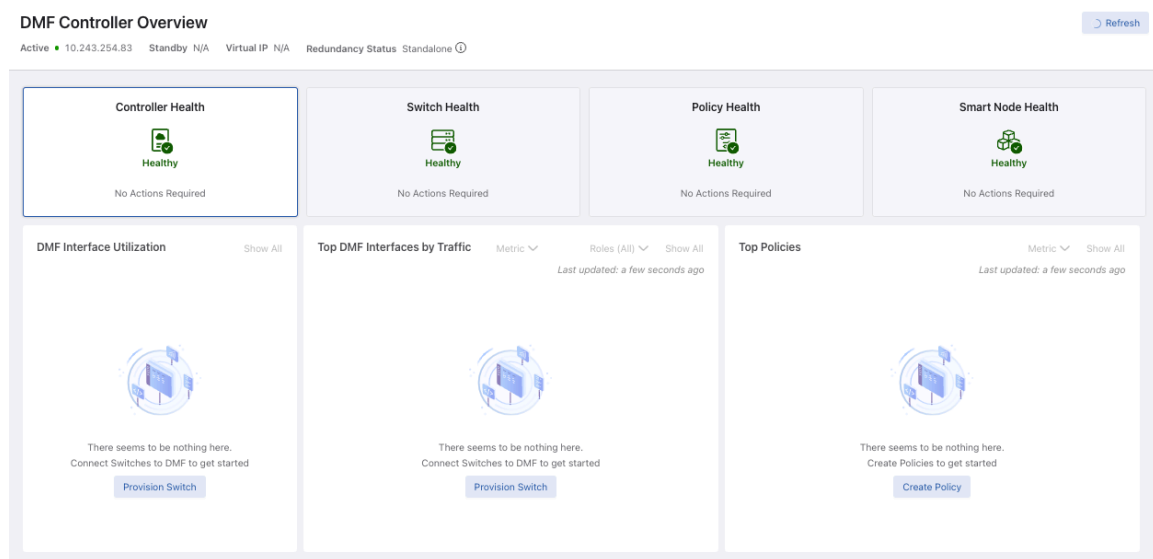
Manually refresh dashboard data using the **Refresh** button.

### 1.4.7 Empty State

When there are no provisioned switches, DMF Interface Utilization and Top DMF Interfaces by Traffic will display an Empty Component.

Each empty component contains a link to provision a switch. The system prompts the user to create a DMF interface if there are provisioned switches but no assigned DMF interfaces.

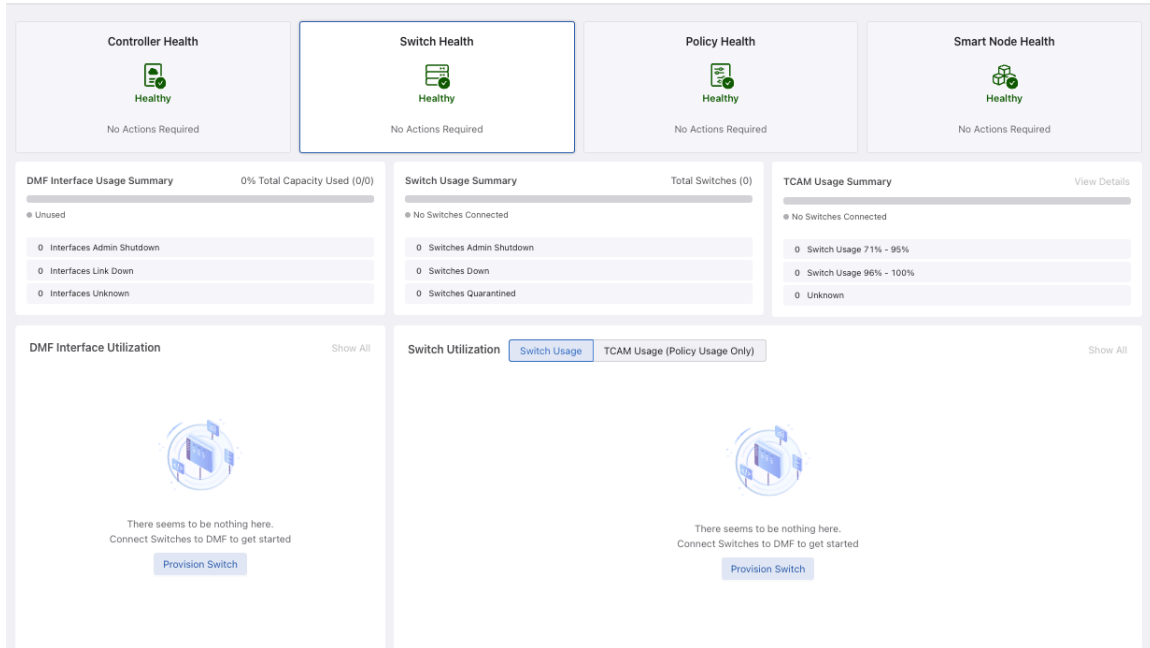
**Figure 1-77: DMF Controller Overview - Empty State**



Top Policies will display an Empty Component if no policies exist.

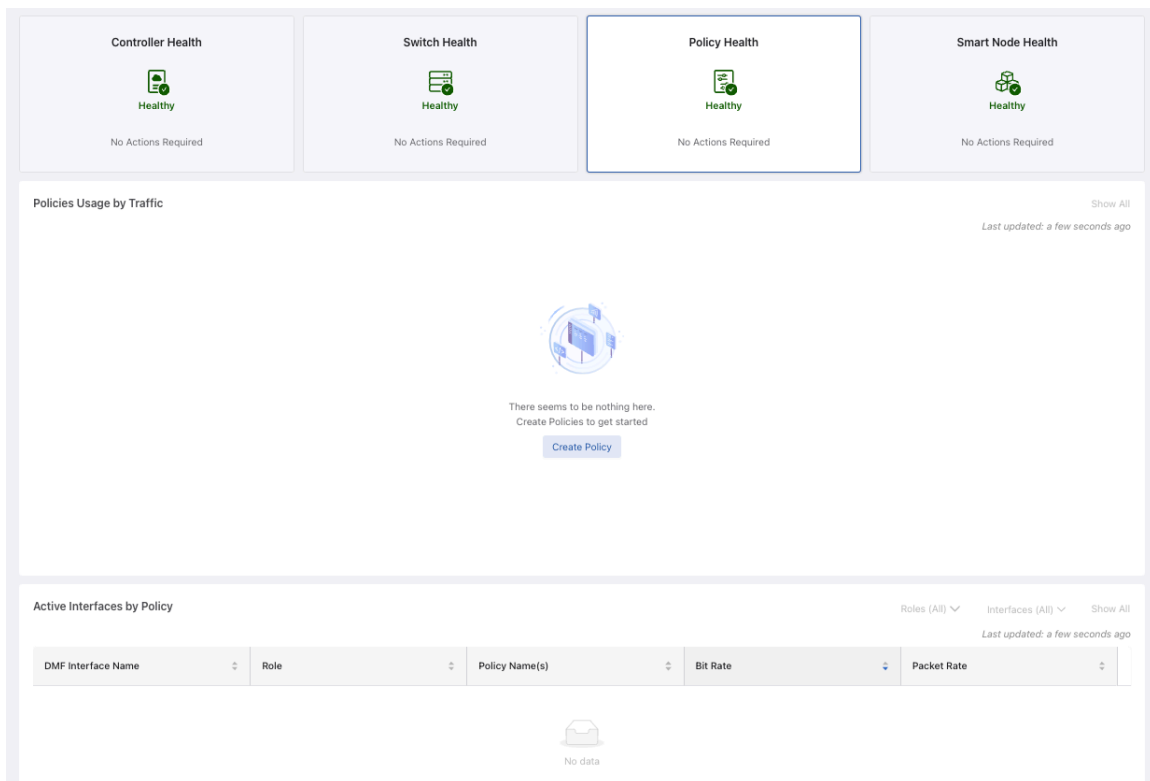
Use the **Create Policy** button to go to the **Create Policy** page.

**Figure 1-78: DMF Switch Health - Empty State**



The Usage Summary components will display Unused or No Switches Connected for the usage bar legend.

**Figure 1-79: DMF Policy Health - Empty State**



Policies Usage by Traffic displays the same Empty Component as Top Policies.

## Managing DMF Switches and Interfaces

---

This chapter describes the basic configuration required to deploy and manage DANZ Monitoring Fabric switches and interfaces.

### 2.1 Overriding the Default Configuration for a Switch

By default, each switch inherits its configuration from the controller. You can use pages two through eight of the **Configure Switch** dialog to override the following configuration options for a specific switch.

- Info
- Clock
- SNMP
- SNMP traps
- Logging
- TACACS
- sFlow
- LAG enhanced hash

#### 2.1.1 CLI Configuration

To use the CLI to manage switch configuration, enter the following commands to enter the `config-switch` submode.

```
controller-1(config)# switch <switch-name>
```

Replace **switch-name** with the alias you assigned to each switch during installation, as in the following example.

```
controller-1(config)# switch DMF-SWITCH-1  
controller-1(config-switch)#
```

From this submode, you can configure the specific switch and override the default configuration pushed from the DMF controller to the switch.

Detailed instructions on overriding default configuration on switches can be found in the **DANZ Monitoring Fabric 8.4 Deployment Guide**.

### 2.2 DMF Interfaces

To monitor traffic, you assign a role to each of the DANZ Monitoring Fabric (DMF) interfaces, which can be of the following four types:

- **Filter interfaces:** ports where traffic enters the DMF. Use filter interfaces to TAP or SPAN ports from production networks.
- **Delivery interfaces:** ports where traffic leaves the DMF. Use delivery interfaces to connect to troubleshooting and/or monitoring and compliance tools, including Network Performance Monitoring

(NPM), Application Performance Monitoring (APM), data recorders, security (DDoS, Advanced Threat Protection, Intrusion Detection, etc.), and SLA measurement tools.

- **Filter and delivery interfaces:** ports with both incoming and outgoing traffic. Use a filter and delivery interface when the port is placed in loopback mode to send outgoing traffic back into the switch for further processing. You can also use a filter and delivery interface when transmit and receive cables are connected to two separate devices to reduce cost.
- **Service interfaces:** interfaces connected to third-party services, or network packet brokers. This includes any interface that sends or receives traffic to or from an NPB.

In addition, interfaces connected to managed service nodes and DANZ recorder nodes can be referenced in the configuration directly without assigning a role explicitly. Also, Inter-Switch Links (ISLs) which interconnect DANZ monitoring switches are automatically detected and referred to as core interfaces.

## 2.2.1 Using the GUI to Configure a DMF Filter or Delivery Interface

To use the DANZ Monitoring Fabric GUI to configure a fabric interface as a filter or delivery interface, complete the following steps:

1. Select **Monitoring > Interfaces** from the main menu to display the DMF interfaces.

**Figure 2-1: DMF Interfaces**

Fabric

Monitoring

Maintenance

Integration

Security

Profile

Open

Close

DANZ Monitoring Fabric - Interfaces

DMF Interfaces

Filter table rows

Filter

	Name	Role	Filter	Delivery	Service	Switch ID	Switch Name	Interface	Description	Groups	Analysis	Route (LAN)	Data (LAN) or Core	Data Forwarding (LAN)	IP Address	Mac (IP)	Mac (LAN)	IP (IP)	Service	Status	Speed	Direction	Device Name		
1	LOOPBACK-INT-1	Filter/Delivery	✓	✓	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet10	---	---	---	sa	Default	---	---	---	---	---	sa	✓	Up	40 Gbps	tx	---	
2	LOOPBACK-INT-2	Filter/Delivery	✓	✓	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet11	---	---	---	sa	Default	---	---	---	---	---	sa	✓	Up	40 Gbps	tx	---	
3	POST-SERVICE-INT-1	Service	---	✓	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet12	---	---	---	sa	---	---	---	---	---	---	sa	unused	✓	Up	40 Gbps	---	---
4	PRE-SERVICE-INT-1	Service	---	✓	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet13	---	---	---	sa	---	---	---	---	---	---	sa	unused	✓	Up	40 Gbps	---	---
5	SW-INT-1	Filter	✓	---	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet14	---	---	---	0	---	---	172.16.48.1	---	---	---	sa	✓	Up	10 Gbps	tx	app-switch-01-vmware-1-ga.logswitch.com	
6	SW-INT-2	Filter	✓	---	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet15	---	---	---	0	---	---	---	---	---	---	sa	✓	Up	10 Gbps	tx	app-switch-01-vmware-1-ga.logswitch.com	
7	TOOL-INT-1	Delivery	---	✓	---	88-88-88-3c-37-13-48-28	88-88-88-3c-37-13-48-28	ethernet16	Connected to SW port 8/11	---	---	sa	Default	---	172.16.32.1	172.16.32.2	205.205.204.0	---	sa	✓	Up	1 Gbps	tx	---	

- Click the provision control (+) at the upper left of the **Interfaces** table to configure a new interface.

**Figure 2-2: Create Interface**

**Create Interface**

1. Interface ✓

2. Configure

The following settings may affect the availability of some configuration options.

Auto VLAN Mode push-per-policy

Strip Forwarding VLANs ✓ Enabled

Big Tap Switch

l2ztn-s4148T-1 (00:00:e4:f0:04:32:91:1a) ☐ Show Connected Switches Only

Optional: use to filter interface choices below

Interface \*

ethernet3

Choose from a list of interfaces not already configured for Big Tap

BACK NEXT RESET CANCEL SAVE

- Select the **Edit** option to change the configuration of an interface that has already been configured.
- Select the switch and interface from the selection lists and click **Next**.  
The system displays the second **Configuration** page.
- Assign a name, IP address, and subnet mask to the interface.
- Select a radio button to assign a role to the interface:
  - Filter
  - Delivery
  - Filter and Delivery
  - Service



**Note:**

- The options available are updated based on your selection.

For example, when you select Filter, the system displays the following dialog box.

**Figure 2-3: Create Interface: Filter**

Create Interface

1. Interface ✓

2. Configure ✓

The following settings may affect the availability of some configuration options.

Auto VLAN Mode push-per-policy

Strip Forwarding VLANs ✓ Enabled

☒ Filter

☐ Delivery

☐ Filter and Delivery

☐ Service

Interface Name \*

filter-1

Rewrite VLAN

Controller will automatically push VLAN.

Analytics

Disabled ☒ Enabled

IP Address

- IPv4 Address -

For receiving IP datagram traffic

BACK NEXT RESET CANCEL SAVE

- Analytics is enabled by default. To disable Analytics for the interface, move the slider to **Disabled**. For information about Analytics refer to the *Analytics Node User Guide*.
- For a filter or a filter and delivery interface, you can optionally enable the Rewrite VLAN option.
- When configuring a Filter interface, you can enable the Rewrite VLAN option by identifying the VLAN in the Rewrite VLAN field.

7. Complete the configuration for the specific interface role.

8. Click **Save** to save the configuration.

## 2.2.2 Using the CLI to Configure a DANZ Filter or Delivery Interface

To assign a filter or delivery role to an interface, complete the following steps:

1. From the config mode, enter the **switch** command, identifying the switch having the interface you want to configure.

```
controller-1(config)# switch DMF-FILTER-SWITCH-1
controller-1(config-switch)#
```



**Note:** You can identify the switch using the alias if that has been configured, as in this example. The CLI changes to the config-switch submode, which lets you configure the specified switch.

2. From the config-switch mode, enter the **interface** command, as in the following example:

```
controller-1(config-switch)# interface ethernet1
```

```
controller-1(config-switch-if)#
```



**Note:** To view a list of the available interfaces, enter the **show switch <switch-name> interface** command, or press the **Tab** key, and the command completion feature displays a concise list of permitted values. When you identify the interface, the CLI changes to the config-switch-if mode, which lets you configure the specified interface.

3. From config-switch-if submode, enter the **role** command to identify the role for the interface. The syntax for defining an interface role (delivery, filter, filter-and-delivery, or service) is as follows:

```
[no] role delivery interface-name <name> [strip-customer-vlan] [ip-address
<ip-address>]
[nexthop-ip <ip-address> <subnet> ]
[no] role filter interface-name <name> [ip-address <ip-address>] {[rewrite
vlan <vlan id (1-4094)>]} [no-analytics]
[no] role both-filter-and-delivery interface-name <name> {[rewrite vlan
<vlan id
(1-4094)>]} [noanalytics]
[no] role service interface-name <name>a
```

The interface-name command assigns an alias to the current interface, which typically would indicate the role assigned, as in the following example:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1
```



**Note:** An interface can have only one role and the configured interface name must be unique within the DANZ Monitoring Fabric.

The following examples show the configuration for filter, delivery, and service interfaces:

- **Filter Interfaces**

```
controller-1 (config)# switch DMF-FILTER-SWITCH-1
controller-1(config-switch)# interface ethernet1
controller-1(config-switch-if)# role filter interface-name TAP-PORT-1
controller-1(config-switch-if)# interface ethernet2
controller-1(config-switch-if)# role filter interface-name TAP-PORT-2
```

- **Delivery Interfaces**

```
controller-1(config-switch-if)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch-if)# interface ethernet1
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1
controller-1(config-switch-if)# interface ethernet2
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-2
```

- **Filter and Delivery Interfaces**

```
controller-1(config-switch-if)# switch DMF-CORE-SWITCH-1
controller-1(config-switch-if)# interface ethernet1
controller-1(config-switch-if)# role both-filter-and-delivery interface-
name loopback-
port-1
controller-1(config-switch-if)# interface ethernet2
controller-1(config-switch-if)# role both-filter-and-delivery interface-
name loopback-
port-2
```

- **Service Interfaces**

```
controller-1(config-switch-if)# switch DMF-CORE-SWITCH-1
controller-1(config-switch-if)# interface ethernet1
```

```
controller-1(config-switch-if)# role service interface-name PRE-SERVICE-PORT-1
controller-1(config-switch-if)# interface ethernet2
controller-1(config-switch-if)# role service interface-name POST-SERVICE-PORT-1
```



**Note:**

- a. An interface can have only one role and the configured interface name must be unique within the DANZ Monitoring Fabric.
- b. A delivery interface will show drops under a many-to-one scenario, i.e., multiple filter interfaces pointing to single delivery interface as per policy definition. These drops should be accounted for micro bursts at egress port. For example, consider a use case of three 10G ingress ports and one 25G egress port. Even if we send a total of 25Gbps of traffic by calculation from ingress to egress, each individual ingress port still operates at 10Gbps inside the BCM chip (i.e., a total of 30G on ingress; 5Gbps traffic is still running at 10Gbps speed on wire but with a bigger inter-frame gap). This means the ingress may potentially oversubscribe the egress due to the 30G to 25G traffic ratio. For example, if each ingress port receives one packet at the exact same time, it causes 30G-to-25G over-subscription or micro-bursting (5Gbps traffic still gets processed at the ingress port's native speed of 10Gbps). Because the egress can only process packets up to 25Gbps, one of the packets will not get dequeued in a timely manner and will be accumulated inside the egress TX queue. If this pattern keeps on repeating, the egress queue eventually drops packets due to the TX buffer becoming full. Therefore, this behavior is expected for the case of many-to-one forwarding. If you reconfigure and use only one 25G ingress port to one 25G egress port, there is no TX drop problem at all.

### 2.2.3 Using the CLI to Identify a Filter Interface using Destination MAC Rewrite

The Destination MAC (D.MAC) Rewrite feature provides an option to identify the Filter interface by overriding the destination MAC address of the packet received on the filter interface. Use this feature for auto-assigned and user-configured VLANs in push-per-filter and push-per-policy modes.



**Note:** The D.MAC Rewrite feature VLAN preservation applies to switches running SWL OS and does not apply to 7280R/7280R2 switches running EOS.

#### Global Configuration

Configure this function at the filter interface level and perform the following steps using the CLI.

1. Select a filter switch and enter the config mode.

```
(config)# switch filter1
```

2. Select an interface from the switch that will be acting as the filter-interface.

```
(config-switch)# interface ethernet5
```

3. Create a filter interface with a name and provide the MAC address to override.

```
(config-switch-if)# role filter interface-name f1 rewrite dst-mac
00:00:00:00:00:03
```

#### CLI Show Commands

The following **show** command displays the ingress flow for the filter switch.

In the Entry value column, the filter switch contains `dst MAC tlv: EthDst(00:00:00:00:00:03)`.

```
(config-policy)# show switch filter1 table ingress-flow-2
# Ingress-flow-2 Device name Entry key Entry
value
```



```

-|-----|-----|-----|-----|
-----|
1 0          filter1      Priority(6400), Port(5), EthType(34525) Name(p1),
  Data([0, 0, 0, 61]), PushVlanOnIngress(flags=[]), VlanVid(0x1), Port(1),
  EthDst(00:00:00:00:00:03)
2 1          filter1      Priority(6400), Port(5)                      Name(p1),
  Data([0, 0, 0, 62]), PushVlanOnIngress(flags=[]), VlanVid(0x1), Port(1),
  EthDst(00:00:00:00:00:03)
3 2          filter1      Priority(36000), EthType(35020)
  Name(__System_LLDP_Flow_), Data([0, 0, 0, 56]), Port(controller), QueueId(0)

```

The core and delivery switch in the Entry value column doesn't contain dst MAC tlv as shown in the following examples.

```

(config-policy)# show switch core1 table ingress-flow-2
# Ingress-flow-2 Device name Entry key
  Entry value
-|-----|-----|-----|-----|
-----|-----|
1 0          core1        Priority(6400), Port(1), EthType(34525),
  VlanVid(0x1) Name(p1), Data([0, 0, 0, 60]), Port(2)
2 1          core1        Priority(6400), Port(1), VlanVid(0x1)
  Name(p1), Data([0, 0, 0, 59]), Port(2)
3 2          core1        Priority(36000), EthType(35020)
  Name(__System_LLDP_Flow_), Data([0, 0, 0, 57]), Port(controller),
  QueueId(0)

```

```

(config-policy)# show switch delivery1 table ingress-flow-2
# Ingress-flow-2 Device name Entry key
  Entry value
-|-----|-----|-----|-----|
-----|-----|
1 0          delivery1    Priority(6400), Port(1), EthType(34525),
  VlanVid(0x1) Name(p1), Data([0, 0, 0, 64]), Port(6)
2 1          delivery1    Priority(6400), Port(1), VlanVid(0x1)
  Name(p1), Data([0, 0, 0, 63]), Port(6)
3 2          delivery1    Priority(36000), EthType(35020)
  Name(__System_LLDP_Flow_), Data([0, 0, 0, 58]), Port(controller),
  QueueId(0)

```

## Troubleshooting

To troubleshoot the scenario where the provided destination MAC address is attached incorrectly to the filter interface. The ingress-flow-2 table above will have a destination MAC rewrite tlv on the filter switch, but no such tlv appears on the core or delivery switch.

As an alternative, drop into the bash of the filter switch to check the flow and destination MAC rewrite.

Use the following commands for the ZTN CLI of the filter switch.

```

(config)# connect switch filter1
(ztn-config) debug admin
filter1> enable
filter1# debug bash

```

The following command prints the flow table of the filter switch.

```
root@filter1:~# ofad-ctl gt ING_FLOW2
```

**Figure 2-4: Filter Switch Flow Table**

```
***** Warning: this is a debug command - use caution! *****
***** Type "exit" or Ctrl-D to return to the Switch Light CLI *****

root@filter1:~# ofad-ctl gt ING_FLOW2

GENTABLE : ing_flow2

GENTABLE ID : 0x0017

[Table count: matched/lookup : 5743/5803

Entry count/limit : 3/4092

[IFP entry reserved count: 4
IPv4 guaranteed max: Not Supported, IPv4 potential max: 4092

IPv6 guaranteed max: Not Supported
IPv6 potential max: 4092

priority 36000 eth_type 0x88cc/0xffff cpu queue_id 0 5743p/1062455b eid 5
priority 6400 in_ports 5 eth_type 0x86dd/0xffff out_ports 1 push_vlan vlan_vid 1 eth_dst 00:00:00:00:00:03 0p/0b eid 8
priority 6400 in_ports 5 out_ports 1 push_vlan vlan_vid 1 eth_dst 00:00:00:00:00:03 0p/0b eid 9
  ipv4_noop      40p/2800b eid 3
  ipv4_drop      20p/1630b eid 4
  ipv6_drop      40p/2800b eid 7
```

The following command shows the policy flow from the filter switch to the delivery switch. The filter switch will have the assigned destination MAC in the match-field.

```
(config)# show policy-flow
# Policy Name Switch Pkts Bytes Pri T Match
Instructions
-|-----|-----|-----|-----|-----|-----|
-----|-----|-----|-----|-----|-----|
1 p1 core1 (00:00:52:54:00:15:94:88) 0 0 6400 1 eth-type
  ipv6,vlan-vid 1 apply: name=p1 output: max-length=65535, port=2
2 p1 core1 (00:00:52:54:00:15:94:88) 0 0 6400 1 vlan-vid 1
  apply: name=p1 output: max-length=65535, port=2
3 p1 delivery1 (00:00:52:54:00:00:11:d2) 0 0 6400 1 vlan-vid 1
  apply: name=p1 output: max-length=65535, port=6
4 p1 delivery1 (00:00:52:54:00:00:11:d2) 0 0 6400 1 eth-type
  ipv6,vlan-vid 1 apply: name=p1 output: max-length=65535, port=6
5 p1 filter1 (00:00:52:54:00:d5:2c:05) 0 0 6400 1
  apply: name=p1 push-vlan: ethertype=802.1Q (33024),set-field:
  match-field/type=vlan-vid, match-field/vlan-tag=1,output: max-length=65535,
  port=1,set-field: match-field/eth-address=00:00:00:00:00:03 (XEROX), match-
  field/type=eth-dst
6 p1 filter1 (00:00:52:54:00:d5:2c:05) 0 0 6400 1 eth-type
  ipv6 apply: name=p1 push-vlan: ethertype=802.1Q (33024),set-field:
  match-field/type=vlan-vid, match-field/vlan-tag=1,output: max-length=65535,
  port=1,set-field: match-field/eth-address=00:00:00:00:00:03 (XEROX), match-
  field/type=eth-dst
```

### Considerations

1. The destination MAC rewrite cannot be used on the filter interface where timestamping is enabled.
2. The destination MAC rewrite will not work when the filter interface is configured as a receive-only tunnel interface.

## 2.2.4 Using the GUI to Identify a Filter Interface using Destination MAC Rewrite

In the UI, configure the **Rewrite Dest. MAC Address** for a **Filter Interface** using one of the two workflows detailed below. The first workflow uses the **Monitoring > Interfaces** UI, while the second uses the **Fabric > Interfaces** UI. To use the second workflow, proceed to step 6, detailed below.

**Workflow One:** Using either the **Monitoring > Interfaces** (or the **Monitoring > Interfaces > Filter Interfaces**) page, proceed to the following workflow:

**1. Create Interface**

1. Click the table action **icon +** button to create a filter interface.
2. After selecting the switch interface in the interface tab, use the **Configure** tab to assign roles.
3. Select the **Filter** radio button for the interface and use the **Rewrite Dest.MAC Address** input to configure the MAC address to override.

**Figure 2-5: Create Interface**

The following settings may affect the availability of some configuration options.

Auto VLAN Mode push-per-filter  
Strip Forwarding VLANs ✓ Enabled

☒ Filter  
☐ Delivery  
☐ Filter and Delivery  
☐ Service  
☐ PTP

Interface Name \*

Rewrite VLAN Rewrite Dest. MAC Address

- MAC Address -

Analytics  
Disabled ☒ Enabled

sFlow  
Disabled ☒ Enabled  
sFlow sampling on this interface

IP Address  
- IPv4 Address -  
For receiving IP datagram traffic

BACK NEXT RESET CANCEL SAVE

4. Click **Save** to continue.
- 2. Edit Interface**

5. Select the row menu of the filter interface to configure or edit, and select **Edit**.

**Figure 2-6: DMF Interfaces - Edit**


The screenshot shows the 'DMF Interfaces' configuration page. At the top, there is a navigation bar with tabs: Fabric, Monitoring, Maintenance, Integration, and Security. Below the navigation bar, a status bar indicates 'Fabric: Healthy, 4 warnings'. The main heading is 'DMF Interfaces'. Below the heading, there is a search bar with the text 'is\_filter:true'. A table lists the DMF interfaces. The table has columns: Name, Role, Filter, Delivery, PTP, and MLAG. The table contains several rows, including 'IXIA-149-9\_12-9' and 'LOOP-T2-s4048'. A context menu is open over the 'IXIA-149-9\_12-9' row, showing options: Edit ..., Delete ..., Monitor Stats, Clear Stats, Clear Peak Stats, and Manage RBAC Group Membership.

		DMF Role				
	Name	Role	Filter	Delivery	PTP	MLAG
<input type="checkbox"/>	IXIA-149-9_12-9	Filter	✓	—	—	—
<input type="checkbox"/>	IXIA-149-9_12-10	Filter	✓	—	—	—
<input type="checkbox"/>			✓	—	—	—
<input type="checkbox"/>			✓	—	—	—
<input type="checkbox"/>		very	✓	✓	—	—
<input type="checkbox"/>		very	✓	✓	—	—
<input type="checkbox"/>		very	✓	✓	—	—
<input type="checkbox"/>		very	✓	✓	—	—
<input type="checkbox"/>	LOOP-T2-s4048	Filter/Delivery	✓	✓	—	—
<input type="checkbox"/>	LOOP-T3X5-7050sx3	Filter/Delivery	✓	✓	—	—

6. In the Configure step, use the Rewrite Dest. MAC Address input to configure the MAC address to override.
3. **Workflow Two:** When using the **Fabric > Interfaces** page, use the following workflow:

7. Select the row menu of the switch interface associated with the filter interface you want to configure and select **Configure**.

**Figure 2-7: Configure Interface**



8. In the DMF tab, select the **Rewrite Dest.MAC Address** field to enter the MAC address to be overridden.

**Figure 2-8: Edit Interface DMF Rewrite Dest. MAC Address**

The screenshot shows the 'Edit Interface' configuration window with the 'DMF' tab selected. The left sidebar shows '1. Port', '2. Traffic', and '3. DMF' with checkmarks. The main area has an orange header with the text 'The following settings may affect the availability of some configuration options.' Below this, there are radio buttons for 'None', 'Filter' (selected), 'Delivery', 'Filter and Delivery', 'Service', and 'PTP'. The 'Interface Name' field is 'IXIA-149-9\_12-9'. The 'Rewrite VLAN' section has the text 'Controller will automatically push VLAN.' and a dropdown for 'Rewrite Dest. MAC Address' set to '- MAC Address -'. There are toggle switches for 'Analytics' (Enabled) and 'sFlow' (Enabled). The 'IP Address' field is set to '- IPv4 Address -' with the note 'For receiving IP datagram traffic'. At the bottom are buttons for 'BACK', 'NEXT', 'RESET', 'CANCEL', and 'SAVE'.

9. Click **Save** to continue.

## 2.3 Using Interface Groups

You can create an interface group consisting of one or more filter or delivery interfaces. It is often easier to refer to an interface group when creating a policy than to explicitly identify every interface to which the policy applies.

You can use an address group in multiple policies, referring to the IP address group by name in match rules. If no subnet mask is provided in the address group, it is assumed to be an exact match. For example, for an IPv4 address group, no mask is interpreted as a mask of /32. For an IPv6 address-group, no mask is interpreted as /128.

You can identify only a single IP address group for a specific policy match rule. Address lists with both **src-ip** and **dst-ip** options cannot be used in the same match rule.

### 2.3.1 Using the GUI to Configure Interface Groups

To create an interface group from the **Monitoring > Interfaces** table, complete the following steps:

1. Select the **Monitoring > Interfaces** option.

**Figure 2-9: Creating Interface Groups from Monitoring > Interfaces**

2. On the **Interfaces** table, enable the check boxes for the interfaces to include in the group.
3. Click the **Menu** control at the top of the table and select **\*\*Group Selected Interfaces**.
4. Complete the dialog that appears to assign a descriptive name to the interface group.



**Note:** You can also define an interface group using the **Monitoring > Interface > Groups** option.

## 2.3.2 Using the CLI to Configure Interface Groups

The following example shows the configuration of two interface groups: a filter interface group **TAP-PORT-GRP** and a delivery interface group **TOOL-PORT-GRP**.

```
controller-1(config-switch)# filter-interface-group TAP-PORT-GRP
controller-1(config-filter-interface-group)# filter-interface TAP-PORT-1
controller-1(config-filter-interface-group)# filter-interface TAP-PORT-2
controller-1(config-switch)# delivery-interface-group TOOL-PORT-GRP
controller-1(config-delivery-interface-group)# delivery-interface TOOL-PORT-1
controller-1(config-delivery-interface-group)# delivery-interface TOOL-PORT-2
```

To view information about the interface groups in the DMF fabric, enter the **show filter-interface-group** command, as in the following examples:

### • Filter Interface Groups

```
controller-1(config-filter-interface-group)# show filter-interface-group
! show filter-interface-group TAP-PORT-GRP
# Name          Big Tap IF Name          Switch IF Name          Direction  Speed    State  VLAN  Tag
-----|-----|-----|-----|-----|-----|-----|-----|
1 TAP-PORT-GRP  TAP-PORT-1  DMF-CORE-SWITCH-1  ethernet17  rx        100Gbps  up     0
2 TAP-PORT-GRP  TAP-PORT-2  DMF-CORE-SWITCH-1  ethernet18  rx        100Gbps  up     0
controller1(config-filter-interface-group)#
```

### • Delivery Interface Groups

```
controller1(config-filter-interface-group)# show delivery-interface-group
! show delivery-interface-group DELIVERY-PORT-GRP
# Name          Big Tap IF Name          Switch IF Name          Direction  Speed    Rate  limit  State  Strip  Forwarding  Vlan
-----|-----|-----|-----|-----|-----|-----|-----|
1 TOOL-PORT-GRP  TOOL-PORT-1  DMF-DELIVERY-SWITCH-1  ethernet15  tx        10Gbps  up     True
2 TOOL-PORT-GRP  TOOL-PORT-2  DMF-DELIVERY-SWITCH-1  ethernet16  tx        10Gbps  up     True
controller-1(config-filter-interface-group)#
```

## Managing DMF Policies

---

This chapter describes how policies work in the DANZ Monitoring Fabric (DMF) and how to configure them.


### 3.1 Overview

A policy selects the traffic to be copied from a production network to one or more tools for analysis. To define a policy, identify the traffic source(s) (filter interfaces), the match rules to select the type of traffic, and destination tool(s) (delivery interfaces). The DMF controller automatically forwards the selected traffic based on the fabric topology. Define match rules to select interesting traffic for forwarding to the tools connected to the specified delivery interfaces. User can also send traffic to be processed by a managed service, such as time stamping, slicing, or deduplication, on a DMF service node. The output from the service node can then be forwarded for analysis to the appropriate tool.

Policies can be simple, but they can also be more complicated when optimizing hardware resources, such as switch TCAM space. Also, DMF provides different switching modes to optimize policies based on use cases and switch capabilities. It is recommended to plan the switching mode to use before configuring policies in a production deployment. For further information, refer to the chapter [Advanced Policy Configuration](#).

### 3.2 Policy Elements

Each policy includes the following configuration elements:

- **Filter interfaces:** these identify the ingress ports for the traffic that needs to be analyzed for this policy. Choose individual filter interfaces and/or one or more filter interface groups. You can also choose **Select All Filter Interfaces**, intended for small-scale deployments.
  - **Delivery interfaces:** these identify the egress ports for the traffic to be analyzed as part of this policy. Choose individual delivery interfaces and/or one or more delivery interface groups. Similar to filter interfaces, a **Select All Delivery Interfaces** option is available for small deployments.
  - **Action:** identifies the policy action to be applied to the inbound traffic. The following actions are available:
    - **Forward:** forwards matching traffic at filter ports to the delivery ports defined in a given policy. You must select at least one or more filter and delivery interfaces.
    - **Drop:** drops matched traffic at the Filter ports. A policy with a drop action is often used in combination with another lower-priority policy to forward all traffic except the dropped traffic to tools. Drop can also be used to measure the bandwidth of matching traffic without forwarding it to a tool. You must select at least one or more filter interfaces.
    - **Capture:** sends the selected traffic to a physical interface on the controller to be saved in a PCAP file. This option works only on a hardware controller appliance, and you must select at least one or more filter interfaces. A policy with a capture action can only run for a short period of time. For continuous packet capture, use the DANZ recorder node. Refer to the chapter [Using the DMF Recorder Node](#) for details.
-  **Note:** If an action is not selected, the policy will not be installed.
- **Match rules:** used to select traffic. The selected traffic is treated based on the action, with the most common action being **Forward**, i.e., forward matched traffic to delivery interfaces. If a match rule is not specified, the policy is not installed. Multiple match rules can be specified in one policy, and each rule is differentiated by its rule number.





**Note:** The rule numbers do not define an order in which the rules will be installed or processed. The numbering just allows the user to list them in order.

- **Managed services** (optional): identifies additional operations to perform, such as packet slicing, time stamping, packet deduplication, packet obfuscation etc., before sending the traffic to the selected delivery interfaces.
- **Status** (optional): enables or disables the policy by using the **active** or **inactive** sub-command from the config-policy submode. By default, a policy is active when initially configured.
- **Priority** (optional): unless a user specifies, all policies are installed with a priority of **100**. In case filter/ingress ports are being shared across policies, a policy with a higher priority will get access to matching traffic first. Traffic that is not matched by the policies with the higher priority then gets processed according to policies with lower priority. Overlapping policies are also not created when two policies have different priorities defined.
- **Push VLAN** (optional): when a user configures the Auto VLAN Mode push as push-per-policy (i.e., to **Push Unique VLAN on Policies**, every policy configured on DMF gets a unique VLAN ID. This VLAN ID is typically picked from the range **1-4094** and auto-increments by **1**. However, if you want a specific policy to have a specific VLAN ID, you should first define a smaller VLAN range using the command **auto-vlan-range** and then pick a VLAN outside of that range to attach to a specific policy. This attachment of a specific VLAN to a specific policy can be done in the CLI using the CLI command **push-vlan** or in the GUI by selecting **Push VLAN** from the **Advanced Options** drop-down and then specifying the VLAN ID.
- **Root switch** (optional): when a core switch (or core link) goes down, existing policies using that switch are rerouted using other core switches. When that switch comes back, we do not move any policy back. In some cases, this causes traffic overload. One way to overcome this problem is to specify a root switch in each policy. When the root switch goes down, policy is rerouted through other switches. When the root switch comes back, we reroute policy through the root switch again.

Policies can include multiple filter and delivery interfaces, and services are optional. Traffic that matches the rules in any policy associated with a filter interface is forwarded to all the delivery interfaces defined in the policy.

Except for a capture action policy, a policy runs indefinitely once activated. You can optionally schedule the policy by specifying a starting time and period for which the policy should run and specify the number of received packets in the tool, after which the policy is automatically deactivated.



**Note:** Create and configure all interfaces and service definitions before creating a policy that uses them. Use only existing interfaces and service definitions when creating a policy. If you create a policy with interfaces or service definitions that do not exist, the policy may enter an inconsistent state. If this happens, delete the policy, create the interfaces and service definitions, and then recreate the policy.

## 3.3 Configuring a Policy

### 3.3.1 Configure a Policy Using the GUI

Log in to the DMF GUI, then complete the following steps:

1. Click **Monitoring > Policies**. To add a policy, click the Provision control (+) at the upper-left corner of the table.

**Figure 3-1: Create Policy Page**

2. Optionally, you can enable or disable advanced options by selecting or deselecting them from the drop-down menu.

**Figure 3-2: Policy Configure: Enable/Disable Advanced Options**



3. Type a name for the policy, and optionally provide a text description.
4. Select the action to be taken on packets matching the policy: forward, drop, or capture.
5. Configure any of the following options, as required and click **Next**:
  - **Inactive/Active**: set inactive state using toggle switch to pre-provision the policy and enable it later.
  - **Start Policy**: specify when to begin the policy.
  - **Set Time**: set the schedule based on local time or UTC.
  - **Run Policy**: run the policy for the specified duration or as long as it is active state.

6. By default, the rule is to match on any traffic. To add a rule for matching packets in the current policy, click the **Create Rule** button. Optionally, if you have exported rules you can import them by clicking **Import Rule** and selecting the file.

**Figure 3-3: Policy Configure: Match Traffic**

- a. Add match criteria to the rule. The following options for matching traffic are added to the dialog.
- **Match Rule Shortcut:** select one or more rules.
  - **Match Rule Group:** select the rule group from the drop-down menu, if the rule group has been created in **Monitoring > Rule Groups**. Refer to the chapter [Using Rule Groups](#) for details.
  - **Custom Rule:** customize the rule.

**Figure 3-4: Policy Configure: Create Rule**

- b. Use the fields provided to select traffic based on the ethertype, IP protocol, IP DSCP, or IP fragments, etc.
- c. To match on the VLAN ID, select **Single**, **Range**, or **Untagged**. Enter the VLAN ID to match on a single VLAN or select a range to match on a range of VLAN IDS. Additionally, you can also select or deselect **Source**, **Destination** and **Offset Match** from the **Additional Configuration** drop-down menu.
-  **Note:** Matching on a VLAN range may use considerable hardware resources.
- d. Identify the source IP address, port number, or MAC address to use for selecting traffic.
-  **Note:** If the CIDR prefix is not provided, a full mask is assumed (for example, **255.255.255.255**).
- e. Identify the destination IP address, port number, or MAC address to use for selecting traffic.
- f. Select **Offset Match - Anchor** and enter offset, mask and value etc. You can add additional offset matches by clicking on **Add Offset Match Setting**.
- g. Click **Add Rule** on the **Create Rule** dialog.
7. Click **Next** on the **Match Traffic** tab of the **Create Policy** page.

- On the **Traffic Source** tab, define the source (feeds) for the traffic in the current policy.

**Figure 3-5: Policy Configure: Traffic Sources**

Create Policy

Advanced Options Policy Details Match Traffic **Traffic Sources** Services Destination Tools Cancel Back Next Create

Existing Configured Filter Interfaces Connected Devices List of vCenters

Select the Filter Interface(s) you want to configure ☐ Select All Filter Interfaces ☐ Hide Interface Groups Create

Filter Interface Name	Switch Name	Interface Name	Description	State
No Data				

To add filter interface(s) to this table, click the **Create** button and then click **Filter Interface** and complete the fields provided on the **Create Filter Interface** side drawer. Optionally, you can create the interface(s) for connected devices through the **Connected Devices** tab, entering the filter interface name(s) on the column **Configured DMF Interface Name**.

To use the filter interface(s) added, select the checkbox **Select All Filter Interfaces** to assign all filter interfaces, or select the checkbox for each desired filter interface.

**Figure 3-6: Policy Configure: Create Filter Interface(s)**

Create Filter Interface

\* Filter interface Name

\* Switch Name

\* Interface Name

Description

Input Description

☒ Enable Analytics ☒ Enable sFlow

IP Address

- (IPv4 Address -)

Cancel Create

- (Optional) Add filter interface group(s) for source traffic.

To add filter interface group(s) to the table, you should have at least one or more interfaces. Click the **Create** button, then click **Filter Interface Group**. Select the interface(s) by clicking on the + icon (or

clicking the - icon to deselect), then click the **Group Selected Interfaces** button. Complete the fields provided on the **Create Filter Interface Group** side drawer, then click the **Create** button.

**Figure 3-7: Policy Configure: Create Filter Interface(s) Group**

Filter Interface Name	Switch Name	Interface Name	Description	State
filter_1	filter1	ethernet1		Up
filter_2	filter2	ethernet1		Up
filter_3	filter3	ethernet1		Up
filter_4	core1	ethernet2		Up

10. (Optional) Add VMware vCenter(s) for source traffic.

To select vCenter instance(s) integrated with the DMF controller as the source of traffic for the current policy, click the **List Of vCenters** tab and select the vCenter(s) checkbox. Refer to the chapter [Integrating vCenter with the DANZ Monitoring Fabric](#) for details.

11. After creating or selecting the traffic source interface(s), click **Next**.
12. (Optional) On the **Services** tab, add managed or unmanaged services to apply the selected traffic to the current policy.

**Figure 3-8: Policy Configure: Services**

Sequence ID	Service Type	Service	Optional	Backup Service	Use as Delivery Service
1	Managed Service	Service Name	No	Backup Service Name	No
2	Un-Managed Service	Service Name	No	Backup Service Name	

- a. To add managed service(s), add one or more service(s) on the **Monitoring > Managed Services** page. Refer to the chapter [Configuring Third-party Services](#) for details.
- b. After adding the managed service(s), click **Add Row**, then select **Service Type Managed Service** from the drop-down menu.
- c. Select **Service Name** from the **Service** drop-down menu, which is already created.
- d. To add unmanaged service(s), add one or more service(s) on the **Monitoring > Services** page. Refer to the chapter [Services in the DANZ Monitoring Fabric](#) for details.
- e. After adding the unmanaged service(s), click **Add Row**, then select **Service Type Un-Managed Service** from the drop-down menu.
- f. Select **Service Name** from the **Service** drop-down menu, which is already created, then click the **Next** button.



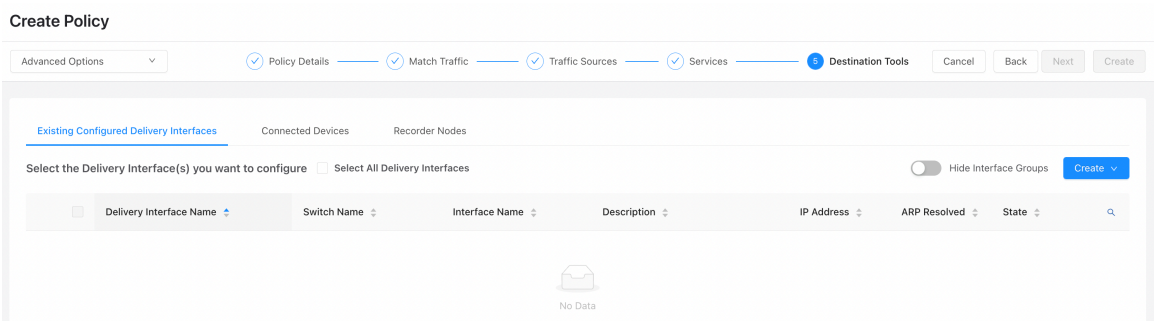
**Note:** When applying a NetFlow and UDP Replication managed service, move the **Managed Services** toggle switch to **Yes**.



**Note:** The **Use as Delivery Service** uses the configured L3 Delivery interface as a fabric exit point for the traffic egressing the Managed Service.

13. On the **Destination Tools** tab, define the delivery interface(s) for the traffic in the current policy.

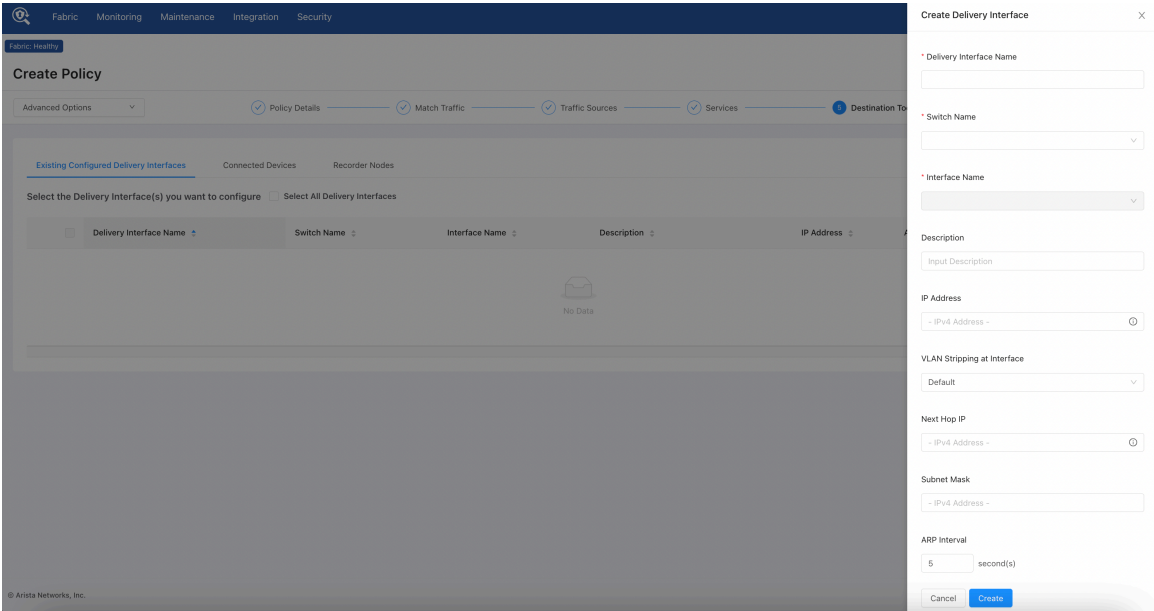
**Figure 3-9: Policy Configure: Destination Tools**



To add delivery interface(s) to this table, click the **Create** button, then click **Delivery Interface** and complete the fields provided on the **Create Delivery Interface** side drawer. Optionally, you can create the interface(s) for connected devices through the **Connected Devices** tab, entering the delivery interface name(s) on the column **Configured DMF Interface Name**.

To use the delivery interface(s) added, select the checkbox **Select All Delivery Interfaces** to assign all delivery interfaces, or select the checkbox for each desired delivery interface.

**Figure 3-10: Policy Configure: Create Delivery Interface(s)**



14. (Optional) Add delivery interface group(s) for destination traffic.

To add delivery interface group(s) to the table, you should have one or more interfaces. Click the **Create** button, then click **Delivery Interface Group**. Select the interface(s) by clicking on the **+** icon (or clicking



the - icon to deselect), then click the **Group Selected Interfaces** button. Complete the fields provided on the **Create Delivery Interface Group** side drawer, then click the **Create** button.

**Figure 3-11: Policy Configure: Create Delivery Interface(s) Group**

Delivery Interface Name	Switch Name	Interface Name	Description	IP Address	ARP Resolved	State
delivery_1	core1	ethernet8			No	Up
delivery_2	delivery2	ethernet1			No	Up
delivery_3	delivery3	ethernet1			No	Up
delivery_4	core1	ethernet1			No	Up

15. (Optional) Add Arista DMF recorder node(s) as delivery interface to forward the traffic to a recorder interface.

To select DMF recorder node(s) integrated with the DMF controller to receive the traffic, click the **Recorder Nodes** tab and select the **DMF Recorder Node(s)** checkbox. Refer to the chapter [Using the DMF Recorder Node](#) for details.

16. Click the **Create** button to create the policy.

### 3.3.2 Configure a Policy Using the CLI

Before configuring a policy, define the filter interfaces that will be used in the policy.

To configure a policy, login to DMF console or SSH to the IP address assigned and complete the following steps:

1. From config mode, enter the `policy` command to name the policy and enter config-policy submode, as in the following example:

```
controller-1(config)# policy POLICY1
controller-1(config-policy)#
```

This example creates the policy **POLICY1** and enters the config-policy submode.

2. Configure one or more match rules to identify the traffic to be aggregated from the filter interfaces assigned to the policy, as in the following example.

```
controller-1(config-policy)# 10 match full ether-type ip dst-ip 10.0.0.50
255.255.255.255
```

This match rule (**10**) selects IP traffic with a destination address of **10.0.0.50**.

3. Assign one or more filter interfaces, which are monitoring fabric edge ports connected to production network TAP or SPAN ports and defined using the `interface` command from `config-switch-if` submode.

```
controller-1(config-policy)# filter-interface TAP-PORT-1
```



**Note:** The filter interfaces used in the policy should be defined before configuring the policy.

To include all monitoring fabric interfaces assigned the filter role, use the **all** keyword, as in the following example:

```
controller-1(config-policy)# filter-interface all
```

4. Assign one or more delivery interfaces, which are monitoring fabric edge ports connected to destination tools and defined using the **interface** command from **config-switch-if** submode.

```
controller-1(config-policy)# delivery-interface TOOL-PORT-1
```

The delivery interfaces used in the policy should be defined before configuring the policy. To include all monitoring fabric interfaces assigned the delivery role, use the **all** keyword, as in the following example:

```
controller-1(config-policy)# delivery-interface all
```

5. Define the action to take on matching traffic, as in the following example:

```
controller-1(config-policy)# action forward
```

- The **forward** action activates the policy so matching traffic immediately starts being forwarded to the delivery ports identified in the policy. The other actions are **capture** and **drop**.
- A policy is active when the configuration of the policy is complete and a valid path exists through the network from a minimum of one of the filter ports to at least one of the delivery ports.
- If a service is inserted in the policy, the policy can only become active and begin forwarding when at least one delivery port is reachable from all the post-service ports defined within the service.

To verify the operational state of the policy enter the **show policy** command.

```
controller-1# show policy GENERATE-IPFIX-NETWORK-TAP-1
Policy Name : GENERATE-IPFIX-NETWORK-TAP-1
Config Status : active - forward
Runtime Status : installed
Detailed Status : installed - installed to forward
Priority : 100
Overlap Priority : 0
# of switches with filter interfaces : 1
# of switches with delivery interfaces : 1
# of switches with service interfaces : 0
# of filter interfaces : 1
# of delivery interfaces : 1
# of core interfaces : 0
# of services : 0
# of pre service interfaces : 0
# of post service interfaces : 0
Push VLAN : 3
Post Match Filter Traffic : -
Total Delivery Rate : -
Total Pre Service Rate : -
Total Post Service Rate : -
Overlapping Policies : none
Component Policies : none
~ Match Rules ~
# Rule
~|-----|
1 1 match any
~|-----|
# DMF      IF Switch      IF Name      State Dir Packets  Bytes      Pkt Rate Bit Rate Counter Reset Time
~|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 TAP-TRAFFIC-2 FILTER-SWITCH-1 ethernet16 up    rx  182876967 69995305364 0      -      2022-10-31 23:13:10.177000 PDT
~|-----|-----|-----|-----|-----|-----|-----|-----|-----|
# DMF      IF Switch      IF Name      State Dir Packets  Bytes      Pkt Rate Bit Rate Counter Reset Time
~|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 TAP-TRAFFIC-1 FILTER-SWITCH-1 ethernet15 up    tx  182876967 69995305364 0      -      2022-10-31 23:13:10.177000 PDT
~ Service Interface(s) ~
None.
~ Core Interface(s) ~
None.
~ Failed Path(s) ~
None.
controller-1#
```



**Note:** If two policies have the same filter and delivery interfaces, and also have the same priority with similar match conditions, then incorrect statistics could be displayed for one or both policies. To alleviate this issue, either increase the priority or change the match conditions in one of the policies.

Detailed status in **show policy** command shows detailed information about a policy status. If for any reason a policy fails, the detailed status shows why the policy failed. One cause of policy failure is the TCAM



getting full. When this happens, the detailed status shows a message like **Table ing\_flow2 is full** <switch\_DPID>.

- **ing\_flow1** is used for programming analytics tracking like DNS, DHCP, ICMP, TCP control packets, and ARP.
- **ing\_flow2** is the TCAM table used for programming data forwarding.
- To delete an existing policy, use the **no policy** command and identify the policy to delete, as in the following example:

```
controller-1(config-policy)# no policy policy-name-1
Warning: submode exited due to deleted object
```

- When a policy is deleted, all traffic rules associated with the policy are also deleted.

### 3.3.3 Define Out-of-band Match Rules Using the CLI

A policy can contain multiple match rules, and each rule is assigned a rule number. However, the rule number does not specify priority or the sequence in which the match rule is applied to traffic entering the filter ports included in a policy. Instead, if the traffic matches any of the match rules, all actions specified in the policy are applied to all matching traffic.

The following example adds two match rules to **dmf-policy-1**.

```
controller-1(config)# policy dmf-policy-1
controller-1(config-policy)# 10 match full ether-type ip dst-ip 10.0.0.50
255.255.255.255
controller-1(config-policy)# 20 match udp src-ip 10.0.1.1 255.255.255.0
controller-1(config-policy)# filter-interface filename2
controller-1(config-policy)# delivery-interface delname3
controller-1(config-policy)# action forward
```



**Note:** When changing an existing installed policy by adding or removing match rules, the change in policy flows is calculated and only the difference is sent to the switches in the path for that policy. The unmodified flows for that policy are not affected.

When more than one action applies to the same packet, DMF makes copies of the interesting packet. For details, refer to the chapter [Advanced Policy Configuration](#) for details.

### 3.3.4 Stop, Start, and Schedule a Policy Using the CLI

To enable or disable a policy, enter the **active** or **inactive** command from the **config-policy** submode.

To stop an action that is currently active, enter the **stop** command from the **config-policy** submode for the policy, as in the following example:

```
controller-1(config)# policy policy1
controller-1(config-policy)# stop
```

By default, if the policy action is **forward** or **drop**, the policy is active unless it is manually stopped or disabled.

To start a stopped or inactive policy immediately, enter the **start now** command from the **config-policy** submode for the policy, as in the following example:

```
controller-1(config)# policy policy1
controller-1(config-policy)# start now
```

For a policy with the **forward** action, the **start now** command causes the policy to run indefinitely. However, policies with the **capture** action run capture for **1** minute unless otherwise specified, after which

the policy becomes inactive. This prevents a capture from running indefinitely and utilizes the appliance storage capacity.

You can also use the **start** command with other options to schedule a stopped or inactive policy. The full syntax for this command is as follows:

```
start { now [ duration duration ] [ delivery-count delivery-packet-count ] | automatic | on-date-time start-time [ duration duration ] seconds-from-now start-time [ duration duration ] [ delivery-count delivery-packet-count ]
```

The following summarizes the usage of each keyword:

- **now**: starts the action immediately.
- **delivery-count**: runs until the specified number of packets are delivered to all delivery interfaces.
- **seconds**: starts the action after waiting the specified number of seconds. For example, **300+** starts the action in **5** minutes.
- **date-time**: starts the action on the specified date and time. Use the format %Y-%m-%dT%H:%M:%S.
- **duration**: if no duration is specified, a duration of **60** seconds is assigned by default. A value of **0** causes the action to run until it is manually stopped. If the **delivery-count** keyword is used with the **capture** action, the maximum duration is **900** seconds.

For example, to start a policy with the **forward** action immediately and run for **five** minutes, enter the following command:

```
controller-1(config-policy)# start now duration 300
```

The following example starts the action immediately and stops after matching **100** packets:

```
controller-1(config-policy)# start now delivery-count 100
```

The following example starts the action after waiting **300** seconds:

```
controller-1(config-policy)# start 300+
```

### 3.3.5 Clear a Policy Using the CLI

To remove a specific DMF policy, use the **no** keyword before the **policy** command, as in the following example:

```
controller-1(config)# no policy sample_policy
```

This command removes the policy **sample\_policy**.

To clear all policies at once, enter the following command:

```
controller-1(config)# clear-all-configured-policy
```

### 3.3.6 View Policies Using the CLI

To display the policies currently configured in the DMF fabric, enter the **show policy** command, as in the following example:

```
controller-1# show policy
```

#	Policy Name	Action	Runtime	Status	Type	Pri	Overlap	Pri	Push	VLAN	Filter	BW	Delivery	BW	Post	Match	Flt	Traf	Delivery	Traffic	Services
1	DUPLICATE-NETWORK-TAPS	drop	installed	to drop	Configured	100	0	1	-	10Gbps	917Mbps	-	-	-	-	-	-	-	-	-	-
2	GENERATE-IPFIX-NETWORK-TAP-1	forward	installed		Configured	100	0	3		10Gbps	10Gbps	-	-	-	-	-	-	-	-	-	-
3	GENERATE-NETFLOW-NETWORK-TAP-1	forward	installed		Configured	100	0	2		10Gbps	10Gbps	-	-	-	-	-	-	-	-	-	-
4	GENERATE-TCP-ANALYTICS-1	forward	installed		Configured	100	0	4		10Gbps	20Gbps	-	-	-	-	-	-	-	-	-	-
5	RECORD-TCP-TRAFFIC	forward	installed		Configured	100	0	5		10Gbps	10Gbps	-	-	-	-	-	-	-	-	-	DEDUPLICATE-BEFORE-RECORD

```
controller-1#
```

This output provides the following information about each policy.

- **#:** a numeric identifier assigned to the policy.
- **Policy Name:** name of the policy.
- **Action:** Forward, Capture, or Drop.
- **Runtime Status:** a policy is active only when the configuration of the policy is complete and a valid path exists through the network from a minimum of one of the filter ports to at least one of the delivery ports (and moves on through the service ports if that is specified). When a service is inserted in the policy, the policy can only become active/forwarding when a delivery port is reachable from all the post-service ports of the service.
- **Type:** configured or dynamic. Refer to the [Configuring Overlapping Policies](#) section for details about dynamic policies created automatically to support overlapping policies.
- **Priority:** determines which policy is applied first.
- **Overlap Priority:** the priority assigned to the dynamic policy applied when policies overlap.
- **Push VLAN:** a feature that rewrites the outer VLAN tag for a matching packet.
- **Filter BW:** bandwidth used.
- **Delivery BW:** bandwidth used.

The following is the full command syntax for the **show policy** command:

```
show policy [ name [filter-interfaces | delivery-interfaces | services | core | optimized-match | failed-paths | drops | match-rules | optimized-match ] ]
```

Event history can be used to determine the last time when policy flows were installed or removed. A value of **dynamic** for Type indicates that the policy was dynamically created for overlapping policies.

## 3.4 Define a Policy Using Interfaces GUI

Using the DMF Interfaces GUI, you can define a policy by selecting the interfaces from the **DMF Interfaces** table to which you want to apply the policy. The system generates a default policy that causes all traffic from the selected filter interfaces to be forwarded to the selected delivery interfaces. You can then modify the default policy to include rules for selecting the traffic and to insert one or more services between the filter and delivery interfaces.

To create a DMF policy, complete the following steps:

1. Select **Monitoring > Interfaces**. The system displays the table shown in the following figure.

**Figure 3-12: DMF Interfaces**

DMF Interfaces

Filter table rows

Name	Role	Filter	Delivery	Service	MLAG	Switch DPID	Switch Name	Interface	Description	Groups	Analytics	sFlow	Rewrite VLAN	Strip VLAN on Egress	Strip Forwarding VLAN	IP Address	Next IP
NETWORK-TAP-1	Filter	✓	—	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	NETWORK-TAP-1	—	—	✓	✓	0	—	NA	—	NA
TAP-8-15	Filter	✓	—	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet1	Connected to IXIA port 8/15	—	✓	✓	0	—	NA	—	NA
TAP-8-16	Filter	✓	—	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet2	Connected to IXIA port 8/16	—	✓	✓	0	—	NA	—	NA
TAP-TRAFFIC-1	Filter/Delivery	✓	✓	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet15	—	—	✓	✓	NA	Default	—	—	—
TAP-TRAFFIC-2	Filter/Delivery	✓	✓	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet16	—	—	✓	✓	NA	Default	—	—	—
TAP-TRAFFIC-3	Filter/Delivery	✓	✓	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet17	—	—	✓	✓	NA	Default	—	—	—
TAP-TRAFFIC-4	Filter/Delivery	✓	✓	—	—	00:00:e4:f8:04:32:91:1a	FILTER-SWITCH-1	ethernet18	—	—	✓	✓	NA	Default	—	—	—
to-analytics-direct	Delivery	—	✓	—	—	00:00:88:67:d4:a0:a2:00	DELIVERY-SWITCH-2	ethernet3	Connected to Analytics Node at 10.240.180.134	—	NA	NA	NA	Default	—	172.29.180.1	172.29.

2. Check the box of each interface that you want to define as the filter and delivery interfaces from the interfaces table.
3. Click the **Menu** control at the top of the table and select **Start Policy From Selected Interfaces**. The system displays the dialog shown in the following figure.

Edit Policy

1. Info

2. Rules

3. Feeds

4. Tools

5. Services

6. Packet Recorder

7. Summary

The following settings may affect the availability of some configuration options.

Switching Mode — Unknown

Policy

(required) Name policy-admin605384

Description —

(required) Action forward

(required) Priority 100

Rules

1 match any

Network-TAP-1

vCenters

- None -

Filter Interfaces

NETWORK-TAP-1

Delivery Interfaces

TOOL-1

**Edit Policy Summary:** this page summarizes policy that is generated automatically for the selected interfaces. This policy selects all traffic from any selected filter interface and forwards it to any selected delivery port.

- To modify the policy, click **Back**, or click the numbered dialog titles in the left column to go directly to a specific dialog.
- To apply the policy, click **Submit**.

### 3.5 Define a Policy Using Host Tracker GUI

Once the monitoring fabric is set up and connected to packet feeds from the production network, DMF starts to collect information about the production network. By default, DMF provides a view of all hosts in the production network visible from the filter interfaces.

For each host that is discovered, this table displays the MAC address, the IP address, the filter interface where the host was seen, and DHCP lease information. DMF can also track DNS names of hosts by capturing DNS packets and analyzing them.

The protocols used for discovery include DNS, DHCP, and ARP. To create a DMF policy through the GUI, complete the following steps:

1. Select **Monitoring > Host Tracker**. The system displays the page shown in the following figure.

**Figure 3-13: DMF Host Tracker**

DMF Host Tracker Do Not Track Track

Filter table rows FILTER 1 2 3 ... 50 <>

	Host Name	MAC Address	Vendor	IP Address	VLAN ID	Filter Interface	First Seen	Last Seen
<input type="checkbox"/>	▶ 10.2.18.25	ac:1f:6b:6d:25:82	Super Micro Computer, Inc.	10.2.18.25	0	TAP-TRAFFIC-3	12m	12m
<input type="checkbox"/>	▶ 10.2.18.26	ac:1f:6b:6d:24:42	Super Micro Computer, Inc.	10.2.18.26	0	TAP-TRAFFIC-1	2h	2h
<input type="checkbox"/>	▶ 10.2.18.27	ac:1f:6b:6d:25:a8	Super Micro Computer, Inc.	10.2.18.27	0	TAP-TRAFFIC-2	51m	51m
<input type="checkbox"/>	▶ 10.4.4.84	8c:ea:1b:94:0d:47	Edgecore Networks Corporation	10.4.4.84	0	TAP-TRAFFIC-1	19m	19m
<input type="checkbox"/>	▶ 10.4.100.35	70:72:cf:c6:f8:df	EdgeCore Networks	10.4.100.35	0	TAP-TRAFFIC-2	11m	11m
<input type="checkbox"/>	▶ 10.5.6.27	00:15:5d:06:1a:02	Microsoft Corporation	10.5.6.27	0	TAP-TRAFFIC-1	19m	19m
<input type="checkbox"/>	▶ 10.5.6.28	00:15:5d:06:1a:00	Microsoft Corporation	10.5.6.28	0	TAP-TRAFFIC-1	19m	19m
<input type="checkbox"/>	▶ 10.8.20.212	70:72:cf:dc:98:c8	EdgeCore Networks	10.8.20.212	0	TAP-TRAFFIC-3	53s	53s
<input type="checkbox"/>	▶ 10.8.20.214	70:72:cf:b6:00:36	EdgeCore Networks	10.8.20.214	0	TAP-TRAFFIC-1	10m	10m
<input type="checkbox"/>	▶ 10.8.22.37	d4:ae:52:d1:3d:ab	Dell Inc.	10.8.22.37	0	TAP-TRAFFIC-1	19m	19m

2. To create a DMF policy from this page, select the available host entries.
3. Check the boxes for the filter interfaces
4. Click the **Menu** control in the upper left corner above the table and select **Start Policy** from **Selected Hosts**. The system displays the **Edit Policy** dialog box with the information derived from the selected hosts.
5. To view the physical filter interfaces, click the expansion control.
 

**Note:** If you enable the host tracker option using the DNS protocol for tracking, DNS traffic cannot be forwarded by DMF policies to the delivery interfaces.
6. Use the other pages on the **Edit Policy** dialog to define the other components of the policy, such as filter and traffic destinations.
7. To modify the policy, click **Back**, or click the numbered dialog titles in the left column to go directly to a specific dialog.
8. To apply the policy, click **Submit**.

## 3.6 Using the Packet Capture Action in a Policy

You can capture packets into a PCAP file for later processing or analysis. The captured packets are stored on the DMF controller hardware appliance. This feature is intended for a quick look at a small amount of traffic. For continuous packet capture and storage, use the DMF Recorder Node, described in the chapter [Using the DMF Recorder Node](#).

**Note:** Storing PCAP files is supported only with the hardware appliance and it is not possible running the controller in a virtual machine. The DMF hardware appliance normally provides **200 GB** of storage capacity, but the hardware appliance is optionally available with **1 TB** of storage capacity.

To enable this feature, connect one of the DMF controller hardware interfaces to a fabric switch interface defined as a DMF delivery interface.

**Figure 3-14: DMF Controller Hardware Appliance**


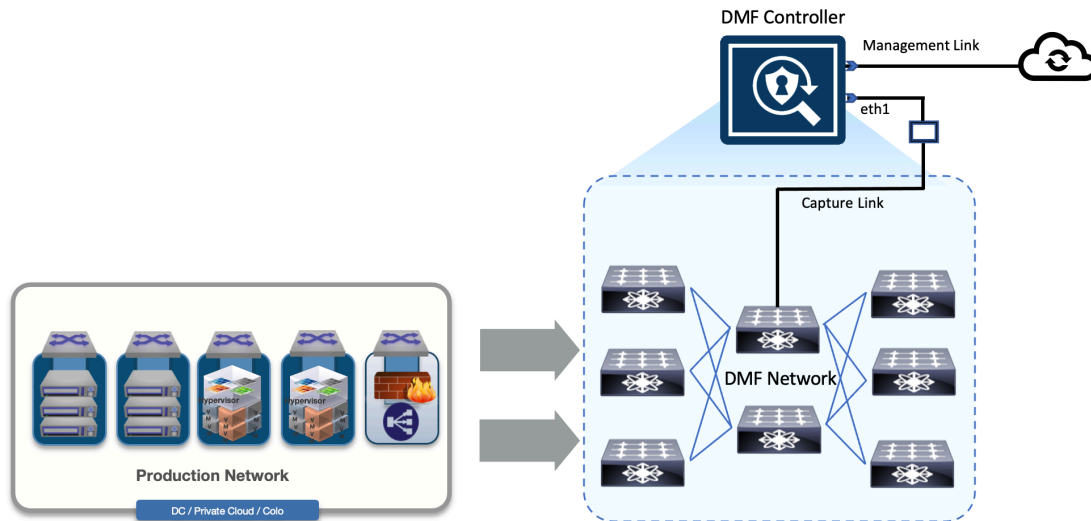


Table 2:

1 1G Management Port

2 10G Management Port

Figure 3-15: Capturing Packets on the DMF Appliance



To capture packets, define a policy with filter ports and match rules to select the interesting traffic. Specify the capture action in the policy, then schedule the policy for a duration or packet count. In the illustrated example, a service has been included in the policy to modify the packets before they are captured, but this optional.

By default, when the policy action is **capture**, the policy is not active until you schedule the policy. Packet captures are always saved on the master (active) controller. In case of HA failover, previous packet captures are left on the controller where they were initially saved.

By default, PCAP files are automatically removed after seven days. This default value can be changed via the following CLI command with command option.

```
controller-1(config)# packet-capture retention-days <tab-key>
<retention-days> Configure packet capture file retention period in days.
Default is 7 days
controller-1(config)#
```

### 3.7 Define a Policy with a Packet Capture Action Using the CLI

Use the **packet-capture-retention-days** command to change the number of days to retain PCAP files. To view the current setting, use the **show packet-capture retention-days <retention-days>** command.

To remove PCAP files immediately, use the **delete packet-capture files** command. You can delete the files associated with a specific policy, as in the following example:

```
controller-1(config-policy)# delete packet-capture files policy capture file
2022-02-24-07-31-25-34d9a85a.pcapng
```

The following command assigns the **capture** action to the current policy and schedules the packet capture to start immediately and run for **60** seconds.

```
controller-1(config-policy)# action capture
controller-1(config-policy)# start now duration 60
```

For a policy with the **forward** action, the **start now** command causes the policy to run indefinitely. However, policies with the **capture** action run capture for **1** minute unless otherwise specified, after which the policy becomes inactive. This prevents a capture from running indefinitely and utilizes the appliance storage capacity.

The following command starts the capture immediately and runs until it captures **1000** packets:

```
controller-1(config-policy)# start now delivery-count 1000
```

Once the packet capture is complete, the PCAP file can be downloaded via HTTP using the URL displayed when entering the **show packet-capture files** command, as shown in the following example.

```
controller-1(config-policy)# show packet-capture files
----- All Packet Capture Files -----
# Policy Name File Name File Size Last Modified URL
-----
1 capture 2022-11-01-03-03-19-cl06e6c.pcapng 258MB 2022-11-01 03:04:17.227000 PDT https://10.9.33.2/pcap/capture/2022-11-01-03-03-19-cl06e6c.pcapng
controller-1(config-policy)#
```

To view the storage used and remaining for PCAP files, enter the **show pcap-storage** command, as in the following example:

```
controller-1 > show packet-capture disk-capacity
Disk capacity : 196GB
controller-1> show packet-capture disk-usage
Disk usage : 258MB
controller-1>
```

To view the number of days PCAP files are retained before deletion, use the **show packet-capture retention-days** command as in the following example:

```
controller-1> show packet-capture retention-days
```

To view the history of packet captures, enter the following command:

```
controller-1(config-policy)# show policy capture history
# Time Event Detail PCAP File
-----
1 2022-11-01 03:03:19.382000 PDT installation complete capturing packets /pcap/capture/2022-11-01-03-03-19-cl06e6c.pcapng
2 2022-11-01 03:04:16.895000 PDT Configuration updated by admin. capturing packets inactive - outside configured runtime/duration,
scheduled to be started in 7sec if set active
3 2022-11-01 03:04:17.266000 PDT policy removed inactive - outside configured runtime/duration,
scheduled to be started in 6sec if set active
controller-1(config-policy)#
```





After enabling the **Peak Bit Rate** column, a column appears in the **Interfaces** table that indicates the relative bandwidth utilization of each interface. When less than **50%** of the bandwidth is used, the bar appears green; **50-75%** changes the bar to yellow, and over **75%** changes the bar color to red.

To display statistics for a specific interface, select **Monitor Stats** from the **Menu** control to the left of the row.

To reset the statistics counters, select **Clear Stats** from the **Menu** control.



**Note:** DMF 8.4 version controllers generate SNMP traps for link saturation and packet loss. Please refer to the DMF 8.4 Deployment Guide - SNMP Trap Generation for Packet Drops and Link Saturation chapter for more information.

### 4.1.3 Using the CLI to Monitor Interface Configuration

To display the currently configured interfaces, enter the **show interface-names** command, as shown in the following example.

```
Ctrl-2> show interface-names

~~~~~ Filter Interface(s) ~~~~~
# DMF IF      Switch      IF Name  Dir State Speed  VLAN Tag Analytics
  Ip address Connected Device
-|-----|-----|-----|---|----|-----|-----|---
1 Lab-traffic Arista-7050SX3-T3X5 ethernet7 rx  up   10Gbps 0      True

~ Delivery Interface(s) ~
None.

~ Service Interface(s) ~
None.

~~~~~ Recorder Fabric Interface(s) ~~~~~
# DMF IF      Switch      IF Name  Dir      State Speed
  Connected Device
-|-----|-----|-----|-----|-----|---
1 PR-NewHW-Intf Arista-7050SX3-T3X5 ethernet25 bidirectional up   25Gbps PR-NewHW ens1f0
2 RMA-CNrail-intf Arista-7050SX3-T3X5 ethernet35 bidirectional up   25Gbps RMA-CNrail ens1f0
```



**Note:** The name is used when configuring a policy.

To display a summary of the current DMF configuration, enter the **show fabric** command, as in the following example.

```
controller-1# show fabric

~~~~~ Aggregate Network State ~~~~~
Number of switches : 3
Inport masking : False
Start time : 2018-03-16 15:42:43.322000 PDT
Number of unmanaged services : 0
Filter efficiency : 0:1
Number of switches with service interfaces : 0
Total delivery traffic (bps) : 168bps
Number of managed service instances : 2
Number of service interfaces : 0
```

```

Match mode : l3-l4-offset-match
Number of delivery interfaces : 6
Max pre-service BW (bps) : 20Gbps
Auto VLAN mode : push-per-policy
Number of switches with delivery interfaces : 2
Number of managed devices : 1
Uptime : 5 hours, 4 minutes
Total ingress traffic (bps) : 160bps
Max filter BW (bps) : 221Gbps
Auto Delivery Interface Strip VLAN : True
Number of core interfaces : 12
Overlap : True
Number of switches with filter interfaces : 2
State : Enabled
Max delivery BW (bps) : 231Gbps
Total pre-service traffic (bps) : 200bps
Track hosts : True
Number of filter interfaces : 5
Number of active policies : 2
Number of policies : 5
~~~~~ Aggregate Interface Statistics ~~~~~
# Interface Type      Dir Packets Bytes  Pkt Rate Bit Rate
-|-----|---|-----|-----|-----|-----|
1 Filter Interface    rx  2444    455611 0         160bps
2 Delivery Interface  tx  4050    421227 0         168bps
-----example truncated-----
controller-1#

```



This page displays the devices discovered by the controller through LACP.

## 4.2.2 Using the CLI to View Switch Configuration

To verify switch interface configuration, enter the **show topology** command, as in the following example:

```
controller> show topology
~~~~~ Filter Interface(s) ~~~~~
# DMF IF      Switch IF  Name      state speed  Connected Device
-|-----|-----|-----|-----|-----|-----|
1 f1          filter-sw-1 s11-eth1 up    10 Gbps
2 f2          filter-sw-1 s11-eth2 up    10 Gbps
~~~~~ Delivery Interface(s) ~~~~~
# DMF IF      Switch IF  Name      state speed  Connected Device
-|-----|-----|-----|-----|-----|-----|
1 d1          filter-sw-2 s12-eth1 up    10 Gbps
2 d2          filter-sw-2 s12-eth2 up    10 Gbps
~~~~~ Service Interface(s) ~~~~~
# DMF IF      Switch IF  Name      Dir state speed  Connected Device
-|-----|-----|-----|-----|-----|-----|
1 post-serv-intf-1 core-sw-1 s9-eth2   up    10 Gbps
2 pre-serv-intf-1  core-sw-1 s9-eth1   up    10 Gbps
~~~~~ Core Interface(s) ~~~~~
# Src Switch  Src IF     Src Speed Dst Switch  Dst IF     Dst Speed
-|-----|-----|-----|-----|-----|-----|
1 core-sw-3    s13-eth2  10 Gbps  delivery-sw-2 s15-eth3  10 Gbps
2 core-sw-3    s13-eth1  10 Gbps  delivery-sw-1 s14-eth3  10 Gbps
3 filter-sw-1  s11-eth3  10 Gbps  core-sw-2     s10-eth1  10 Gbps
4 core-sw-2    s10-eth1  10 Gbps  filter-sw-1   s11-eth3  10 Gbps
5 delivery-sw-2 s15-eth3  10 Gbps  core-sw-3     s13-eth2  10 Gbps
6 core-sw-2    s10-eth2  10 Gbps  filter-sw-2   s12-eth3  10 Gbps
7 filter-sw-2  s12-eth3  10 Gbps  core-sw-2     s10-eth2  10 Gbps
8 delivery-sw-1 s14-eth3  10 Gbps  core-sw-3     s13-eth1  10 Gbps
~~~~~ Statistics ~~~~~
# DMF IF      Switch IF  Role      State  Packets  Bytes  Pkt Rate Bit Rate
-|-----|-----|-----|-----|-----|-----|
1 f1          filter-sw-1 s11-eth1 filter up    0      0      0      -
2 f2          filter-sw-1 s11-eth2 filter up    0      0      0      -
3 d1          filter-sw-2 s12-eth1 delivery up    8      600    0      -
4 d2          filter-sw-2 s12-eth2 delivery up    8      600    0      32 bps
6 -          core-sw-3    s13-eth1 core   up    3432   257400 0      32 bps
7 -          delivery-sw-2 s15-eth3 core   up    3431   257325 0      32 bps
8 -          delivery-sw-1 s14-eth3 core   up    3430   257250 0      32 bps
9 -          core-sw-2    s10-eth1 core   up    3429   257175 0      32 bps
10 -         filter-sw-1  s11-eth3 core   up    3431   257325 0      32 bps
11 -         core-sw-3    s13-eth2 core   up    3432   257400 0      32 bps
12 -         filter-sw-2  s12-eth3 core   up    3429   257175 0      32 bps
```

## 4.3 Viewing Information about a Connected Production Network

Once the monitoring fabric is set up and connected to packet feeds from the production network, DMF starts to gather information about the production network. By default, DMF provides a view of all hosts in the production network visible from the filter interfaces. You can get a view of this information on the GUI page, under **Monitoring > Host Tracker**. As shown below, you can see the host MAC address, IP address, when and on which filter interface traffic from the host was seen, and DHCP lease information. To display this information, enter the **show tracked-hosts** command, as shown in the following example.

```
# show tracked-hosts
# IP Address      MAC Address      Host name      Filter interfaces  VLANs  Last seen
-|-----|-----|-----|-----|-----|-----|
1 10.0.0.3        40:a6:d9:7c:9f:9f Apple         wireless-poe-1     0      1 hours
2 10.0.0.6        98:fe:94:1c:37:06 Apple         wireless-poe-1     0      42 min
3 10.0.0.6        dc:2b:61:81:64:45 Apple         wireless-poe-1     0      3 hours
4 10.0.0.7        20:c9:d0:48:f3:3d Apple         wireless-poe-1     0      2 hours
5 10.0.0.11       60:03:08:9b:4f:48 Apple         wireless-poe-1     0      13 min
6 10.0.1.3        14:10:9f:e4:e6:bf Apple         wireless-poe-1     0      51 min
-----output truncated-----
```

DMF also tracks the DNS names of hosts by capturing and analyzing packets using a number of different protocols. To manage host-name tracking, from config-analytics mode, use the **track** command, which has the following syntax:

**[no] track { arp | dns | dhcp | icmp }**

For example, the following command enables tracking using DNS:

```
controller-1(config)# analytics
```

```
controller-1(config-analytics)# track dns
```



**Note:** If you enable DNS for tracking, DNS traffic will not be included in DMF policies.

You can also exclude host tracking for a specific filter interface by using the **no-analytics** option with the **role** command.

```
controller-1(config)# switch DMF-FILTER-SWITCH-1
controller-1(config-switch)# interface ethernet20
controller-1(config-switch-if)# role filter interface-name TAP-PORT-01 no-
analytics
```

This command disables all host tracking on interface **TAP-PORT-01**.

### 4.3.1 Using the CLI to View Connected Devices and LAGs

Some information on devices in the production network, which is discovered using LLDP and CDP, can be seen using the **show connected-devices** command. The information helps in determining if filter interfaces are connected to the intended production device.

Use the **show connected-devices** command from login mode to display the devices connected to the DANZ Monitoring Fabric. This command displays information about devices connected to DMF switch interfaces. The information is extracted from link-level protocol packets such as LLDP, CDP, and UDLD. Expired link-level information is ignored.

```
controller-1# show connected-devices switch FILTER-SWITCH-1
# Switch
-----
1 FILTER-SWITCH-1 ethernet3 TOOL-1 False dmf-rack257-server-4 Ubuntu 18.04.6 LTS Linux 4.15.0-188-generic #199-Ubuntu SMP Wed d4:ae:52:c6:b7:1e eth1 fe80::15054:fff:fe8d:67a0 LLDP
2 FILTER-SWITCH-1 ethernet4 False dmf-rack257-server-4 Ubuntu 18.04.6 LTS Linux 4.15.0-188-generic #199-Ubuntu SMP Wed d4:ae:52:c6:b7:1e 90:e2:ba:28:ea:5a eth5 fe80::15054:fff:fe8d:67a0 LLDP
3 FILTER-SWITCH-1 ethernet5 False dmf-rack258-arista-2 Arista Networks EOS version 4.10.6 running on an Arista Network 00:1c:73:4c:62:83 Ethernet27 10.240.180.121 LLDP
4 FILTER-SWITCH-1 ethernet7 False dmf-rack258-arista-2 Arista Networks EOS version 4.10.6 running on an Arista Network 00:1c:73:4c:62:83 Ethernet28 10.240.180.121 LLDP
controller-1#
```

Users can see the most recent events related to particular connected devices via the CLI command **show connected-devices history device\_alias**.

```
controller-1# show connected-devices history DMF-RECORDER-NODE
-----
Device Name : DMF-RECORDER-NODE
Device Description : dmf-recorder-node, SN 33R8X02
Connection State : connected
Last Connected Switch : DELIVERY-SWITCH-1
Last Connected Interface : ethernet11
Chassis ID : 24:6e:96:b4:b5:10
Port ID : f8:f2:1e:00:0b:b0
Port Description : ens1f0
Management Address : 10.240.180.222
Protocol : LLDP

# Time of Event      Type of Event      Switch      IF Name      Device Name      Device Description      Connection State      Last Connected Switch      Last Connected Interface      Chassis ID      Port ID      Port Description      Management Address      Protocol
-----
1 2022-11-02 12:52:37.490000 PDT device-added DELIVERY-SWITCH-1 ethernet11 DMF-RECORDER-NODE dmf-recorder-node, SN 33R8X02 24:6e:96:b4:b5:10 f8:f2:1e:00:0b:b0 ens1f0 10.240.180.222 LLDP
2 2022-11-02 12:49:56.936000 PDT device-removed DELIVERY-SWITCH-1 ethernet11 DMF-RECORDER-NODE dmf-recorder-node, SN 33R8X02 24:6e:96:b4:b5:10 f8:f2:1e:00:0b:b0 ens1f0 10.240.180.222 LLDP
3 2022-11-02 12:45:56.865000 PDT device-added DELIVERY-SWITCH-1 ethernet11 DMF-RECORDER-NODE dmf-recorder-node, SN 33R8X02 24:6e:96:b4:b5:10 f8:f2:1e:00:0b:b0 ens1f0 10.240.180.222 LLDP
controller-1#

CLI command show connected-devices disconnected displays connected devices that are no longer connected to the fabric
controller-1# show connected-devices disconnected
-----
# Time Last Connected      Device Name      Device Description      Connection State      Last Connected Switch      Last Connected Interface      Chassis ID      Port ID      Port Description      Management Address      Protocol
-----
1 2022-11-05 05:42:44.605000 PDT VC7-ESX-4 VMware ESX1 Releasebuild-18426014 disconnected FILTER-SWITCH-1 NETWORK-TAP-1 c0:d6:82:ba:b5:31 v0x1@ TO-RELEASE-SETUP 10.240.180.143 LLDP
2 2022-11-02 12:52:24.736000 PDT VC7-ESX-5 VMware ESX1 Releasebuild-18426014 disconnected DELIVERY-SWITCH-1 ethernet11 f8:f2:1e:00:0b:b0 f8:f2:1e:00:0b:b0 10.240.180.143 LLDP
3 2022-11-02 12:46:59.884000 PDT VC7-ESX-6 VMware ESX1 Releasebuild-19450014 disconnected DELIVERY-SWITCH-2 ethernet3 f8:f2:1e:0f:08:01 f8:f2:1e:0f:08:01 10.240.180.143 LLDP
controller-1#
```

If a DMF switch interface is connected to a SPAN port, inaccurate information may be displayed because some vendor devices mirror link-level packets to the SPAN port.

To display details about the link aggregation groups connected to the DMF switch interfaces use the **show connected-lacp** command. The information is extracted from LACP protocol packets. Expired LACP information is ignored. The following is an example:

```
controller-1> show connected-lacp
# Switch      IF Name      DMF name      Device Mac      Port ID      Link Group      Key Partner Device Mac Partner Port ID Partner Link Group Key
-----
1 DMF-SWITCH-1 ethernet1      5c:16:c7:02:b8:01 23      10      00:00:00:00:00:00 0      0
2 DMF-SWITCH-1 ethernet14     5c:16:c7:02:b8:01 23      10      00:00:00:00:00:00 0      0
3 DMF-SWITCH-1 ethernet16     5c:16:c7:02:b6:01 42      3      00:00:00:00:00:00 0      0
4 DMF-SWITCH-1 ethernet2      5c:16:c7:02:b6:01 42      3      00:00:00:00:00:00 0      0
```

## Using the DMF Service Node Appliance

This chapter describes how to configure the managed services provided by the DMF Service Node Appliance.

### 5.1 Overview

The DMF Service Node comes with multiple interfaces that can be connected to traffic for processing and analysis. Each interface can be programmed independently to provide any of the supported managed-service actions.

To create a managed service, identify a switch interface connected to the service node, specify the service action, and configure the service action options.

Configure a DMF policy to use the managed service by name. This causes the controller to forward traffic selected by the policy to the service node. The processed traffic is returned to the monitoring fabric using the same interface and is forwarded to the tools (delivery interfaces) defined in the DMF policy.

If the volume of traffic selected by the policy is too much for a single service node interface, define a LAG on the switch connected to the service node, then use the LAG interface when defining the managed service. All the service node interfaces connected to the LAG are then configured to perform the same action. The traffic selected by the policy is automatically load balanced among the LAG member interfaces, and the return traffic is distributed in the same way.

### 5.2 Changing the Service Node Default Configuration

Configuration settings are automatically downloaded to the service node from the DMF controller to eliminate the need for box-by-box configuration. However, you can override the default configuration for a service node from the **config-service-node** submode for any service node.



**Note:** In the current release, these options are available only from the CLI, and are not included in the DMF GUI.

To change the CLI mode to **config-service-node**, enter the following command from config mode on the Active DMF controller:

```
controller-1(config)# service-node <service_node_alias>
controller-1(config-service-node)#
```

Replace **service\_node\_alias** with the alias you want to use for the service node. This alias is associated with the hardware MAC address of the service node, using the **mac** command. The hardware MAC address configuration is mandatory for the service node to interact with the DMF controller.

Use any of the following commands from **config-service-node** submode to override the default configuration for the associated service node:

- **admin password**: set the password to log in to the service node as an admin user.
- **banner**: set the service node pre-login banner message.
- **description**: set a brief description.
- **logging**: enable service node logging to the controller.
- **mac**: configure a MAC address for the service node.
- **ntp**: configure the service node to override default parameters.

- **snmp-server**: configure an SNMP trap host that will receive SNMP traps from the service node.

## 5.3 Using SNMP to Monitor DPDK Service Node Interfaces

Starting with **BMF Release 7.1.1**, you can directly fetch the counters and status of the service node interfaces handling traffic (DPDK interfaces). The following are the supported OIDs.

```
interfaces MIB: #.1.3.6.1.2.1.2#
ifMIBObjects MIB: #.1.3.6.1.2.1.31.1#
```



**Note:** SNI DPDK (traffic) interfaces are identified by a three-digit number between **101** and **116**.

In the following example, interface **sni5** (105) is handling data traffic. To fetch the packet count, use the following command:

```
snmpget -v2c -c public 10.106.6.5 .1.3.6.1.2.1.31.1.1.1.6.105
IF-MIB::ifHCInOctets.105 = Counter64: 10008
```

To fetch the counters for packets exiting the service node interface, enter the following command:

```
snmpget -v2c -c public 10.106.6.5 .1.3.6.1.2.1.31.1.1.1.10.105
IF-MIB::ifHCOutOctets.105 = Counter64: 42721
```

To fetch Link Up and Down status, enter the following command:

```
[root@TestTool anet]# snmpwalk -v2c -c onlon1 10.106.6.6 .1.3.6.1.2.1.
2.2.1.8.109
IF-MIB::ifOperStatus.109 = INTEGER: down(2)
[root@TestTool anet]# snmpwalk -v2c -c onlon1 10.106.6.6 .1.3.6.1.2.1.
2.2.1.8.105
IF-MIB::ifOperStatus.105 = INTEGER: up(1)
```

## 5.4 Configuring Managed Services

To view, edit, or create DMF managed services, select the **Monitoring > Managed Services** option.

**Figure 5-1: Managed Services**

The screenshot displays the 'DMF Managed Services' page in the Arista Monitoring Fabric interface. The left sidebar contains navigation options: Devices, Managed Services, GTP Correlation Profiles, APN Groups, IMEI Groups, IMSI Groups, MSISDN Groups, and IPFX Templates. The 'Managed Services' section is active, showing a table of services.

Name	Description	Action	Action Detail	Anchor/Offset	Insert Original Packet Length	Installed	Service Interface	Service Interface Status	Policy Count
MS-APP-FILTER-ACTIVE	---	Application ID	1 match amazon 2 match ebay 3 match facebook 4 match google 5 match microsoft 6 match netflix 7 match office365 8 match sharepoint 9 match skype 10 match skype	NA	NA	✓	superspine-z9264-1 / ethernet17/4	✓ Up	1
MS-APP-FILTER-BACKUP	---	Application ID	1 match amazon 2 match ebay 3 match facebook 4 match google 5 match microsoft 6 match netflix	NA	NA	✓	superspine-as7816-1 / ethernet13/4	✓ Up	0

This page displays the service node appliance devices connected to the DMF controller and the services that have been configured on the controller.

### 5.4.1 Using the GUI to Define a Managed Service

To create a new managed service, complete the following steps:

1. Click the **Provision** control (+) in the upper left corner of the **Managed Services** table. The system displays the **Create Managed Service** dialog, shown in the following figure.

**Figure 5-2: Create Managed Service: Info**

The 'Create Managed Service' dialog box is shown with the 'Info' tab selected. It contains the following fields and options:

- Name:** my-managed-service
- Description:** (empty text area)
- Switch:** (dropdown menu)
- Interface:** - Interface - (dropdown menu)
- Optional:** use to filter interface choices below
- Options:**
  - ☒ Show Managed Device Switches Only
  - ☐ Show Connected Switches Only

2. Assign a name to the managed service.
3. (Optional) Provide a text description of the managed service.
4. Select the switch and interface where the service will be provided.



The **Show Managed Device Switches Only** checkbox, which is enabled by default, limits the switch selection list to service node appliances. To limit the display to connected switches, enable the **Show Connected Switches Only** checkbox.

5. Select the action from the Action selection list, which provides the following options.
  - **Application ID**
  - **Deduplication:** Deduplicate selected traffic, including NATted traffic.
  - **GTP Correlation**
  - **Header Strip:** Remove bytes of packet starting from zero till selected Anchor and offset bytes
  - **Header Strip Cisco Fabric Path Header:** Remove Cisco Fabric Path encapsulation header
  - **Header Strip ERSPAN Header:** Remove Encapsulated Remote Switch Port Analyzer Encapsulation header
  - **Header Strip Genev1 Header:** Remove Generic Network Virtualization Encapsulation header
  - **Header Strip L3 MPLS Header:** Remove Layer 3 MPLS encapsulation header
  - **Header Strip LISP Header:** Remove Locator Separation Protocol Encapsulation header
  - **Header Strip VXLAN Header:** Remove Virtual Extensible LAN Encapsulation header
  - **IPFIX:** Generate IPFIX by selecting matching traffic and forward to specified collectors.
  - **Mask:** Mask sensitive information as specified by user in packets fields.
  - **NetFlow:** Generate a NetFlow by selecting matching traffic and forward to specified collectors.
  - **Pattern-Drop:** Drop matching traffic.
  - **Pattern Match:** Forward matching traffic.
  - **Session Slice:** Slice TCP sessions.
  - **Slice:** Slice the given number of bytes based on the specified starting point in the packet.
  - **TCP Analysis**
  - **Timestamp:** Identify the time that the packet is received by the service node.
  - **UDP Replication:** Copy UDP messages, such as syslog or NetFlow messages, to multiple IP destinations.
6. (Optional) Identify the start point for service actions.
 

You can identify the start point for the deduplication, mask, pattern-match, pattern-drop services, or slice services using one of the keywords listed below.

  - **packet-start:** add the number of bytes specified by the *integer* value to the first byte in the packet.
  - **l3-header-start:** add the number of bytes specified by the *integer* value to the first byte in the Layer 3 header.
  - **l4-header-start:** add the number of bytes specified by the *integer* value to the first byte in the layer-4 header.
  - **l4-payload-start:** add the number of bytes specified by the *integer* value to the first byte in the layer-4 user data.
  - **integer:** specify the number of bytes to offset for determining the start location for the service action, relative to the specified start keyword.
7. To assign a managed service to a policy, enable the checkbox on the **Managed Services** page of the **Create Policy** or **Edit Policy** dialog.
8. To create a backup service, select the backup service from the Backup Service selection list. The backup service is used when the primary service is not available.

## 5.4.2 Using the CLI to Define a Managed Service



**Note:** When connecting a LAG interface to the DMF service node appliance, member links should be of same speed and can span across multiple service nodes. The maximum number of supported member links per LAG interface is 32, but this number varies based on switch platform. Please refer to the hardware guide for the exact details of the supported configuration.

To configure a service to direct traffic to a DMF service node, complete the following steps:

1. Define an identifier for the managed service by entering the following command:

```
controller-1(config)# managed-service DEDUPLICATE-1  
controller-1(config-managed-srv)#
```

This enters the **config-managed-srv** submode, where you can configure a DMF managed service.

2. (Optional) Configure a description for the current managed service by entering the following command:

```
controller-1(config-managed-srv)# description "managed service for policy DEDUPLICATE-1"
```

The following are the commands available from this submode:

- **description**: provide a service description
  - **post-service-match**: select traffic after applying the header strip service
  - Action sequence number in range [**1 - 20000**]: identifier of service action
  - **service-interface**: associate an interface with the service
3. Use a number in the range [**1 - 20000**] to identify a service action for a managed service.

The following summarizes the available service actions. See the subsequent sections for details and examples for specific service actions.

- **dedup** {**anchor-offset** | **full-packet** | **routed-packet**}
- **header-strip** {**I4-header-start** | **I4-payload-start** | **packet-start** } [**offset**]
- **decap-cisco-fp** {**drop**}
- **decap-erspan** {**drop**}
- **decap-geneve** {**drop**}
- **decap-l3-mpls** {**drop**}
- **decap-lisp** {**drop**}
- **decap-vxlan** {**drop**}
- **mask** {**mask/pattern** } [{**packet-start** | **I3-header-start** | **I4-header-start** | **I4-payload-start**} **mask/offset**] [**mask/mask-start** **mask/mask-end**]
- **netflow** *Delivery\_interface Name*
- **ipfix** *Delivery\_interface Name*
- **udp-replicate** *Delivery\_interface Name*
- **tcp-analysis** *Delivery\_interface Name*



**Note:** The IPFIX, NetFlow, and udp-replicate service actions enable a separate submode in which you can define one or more specific configurations. One of these services must be the last service applied to the traffic selected by the policy

- **pattern-drop** *pattern* [{**I3-header-start** | **I4-header-start** | **packet-start** }]
- **pattern-match** *pattern* [{**I3-header-start** | **I4-header-start** | **packet-start** }]
- **slice** {**packet-start** | **I3-header-start** | **I4-header-start** | **I4-payload-start**} *integer*
- **timestamp**

For example, the following command enables packet deduplication on the routed packet:

```
controller-1(config-managed-srv)# 1 dedup routed-packet
```

4. Optionally, identify the start point for the mask, pattern-match, pattern-drop services, or slice services.
5. Identify the service interface for the managed service by entering the following command: controller-1(config-managed-srv)# **service-interface switch DMF-CORE-SWITCH-1 ethernet40**.

To increase the bandwidth available to the managed service, use a port channel instead of an interface. The following example enables lag-interface1 for the service interface:

```
controller-1(config-managed-srv) # service-interface switch DMF-CORE-SWITCH-1 lag1
```

6. Apply the managed service within a policy like any other service, as shown in the following examples for deduplication, NetFlow, pattern matching (forwarding), and packet slicing services.



**Note:** The same managed service can be used in multiple DMF policies. For example, if you configure a packet slicing managed-service, this can be used in multiple DMF policies.

### 5.4.3 Monitoring Managed Services

To identify managed services bound to a service node interface and the health status of the respective interface, use the following commands:

```
controller-1# show managed-service-device <SN-Name> interfaces
controller-1# show managed-service-device <SN-Name> stats
```

For example, the following command shows the managed services handled by the Service Node Interface (SNI):

```
controller-1#show managed-service-device <SN-Name> interfaces
----- Interfaces -----
# Managed Service Interface Switch      Switch Interface Service Name
-----
1 sni1      CORE-SWITCH-1 MS-LAG-1      MULTIPLE-SERVICES-ACTIVE-1
2 sni2      CORE-SWITCH-1 MS-LAG-2      MULTIPLE-SERVICES-BACKUP-1
3 sni3      CORE-SWITCH-1 MS-LAG-3      MULTIPLE-SERVICES-ACTIVE-2
4 sni4      CORE-SWITCH-1 MS-LAG-4      MULTIPLE-SERVICES-BACKUP-2
5 sni5      CORE-SWITCH-1 MS-LAG-5      MULTIPLE-SERVICES-ACTIVE-3
6 sni6      CORE-SWITCH-1 MS-LAG-6      MULTIPLE-SERVICES-BACKUP-3
7 sni7      CORE-SWITCH-1 MS-LAG-7
8 sni8      CORE-SWITCH-1 MS-LAG-8
9 sni9      CORE-SWITCH-1 MS-LAG-1      MULTIPLE-SERVICES-ACTIVE-1
10 sni10     CORE-SWITCH-1 MS-LAG-2      MULTIPLE-SERVICES-BACKUP-1
-----
controller-1# show managed-service-device <SN-Name> stats
----- Stats -----
# Device Int Name Service Name      Action Load      Rx packets Rx bytes      Rx Bit Rate Applied packets Applied bytes App Bit Rate Tx packets Tx bytes      Tx Bit Rate
-----
1 sni1      MULTIPLE-SERVICES-ACTIVE-1 netflow moderate 38830541102 22534951895775 331Mbps 33491332356 19392280709127 316Mbps 64547484 242557266936 4.77Mbps
2 sni1      MULTIPLE-SERVICES-ACTIVE-1 dedup moderate 40818176809 22762491495037 335Mbps 40790968622 22760851128054 335Mbps 38803332951 22533111539451 331Mbps
3 sni3      MULTIPLE-SERVICES-ACTIVE-2 dedup no 0 0 0 0 0 0 0 0 0
4 sni3      MULTIPLE-SERVICES-ACTIVE-2 ipfix no 0 0 0 0 0 0 0 17048 1619560 48bps
5 sni5      MULTIPLE-SERVICES-ACTIVE-3 dagger critical 14769229199 1070427935998 29.1Mbps 14769229193 1070427935998 29.1Mbps 1614349 13328879865 367Kbps
6 sni5      MULTIPLE-SERVICES-ACTIVE-3 slice critical 14769229199 8719723597030 183Mbps 14762965759 8719097432562 183Mbps 14769229199 1069801771998 29.1Mbps
7 sni9      MULTIPLE-SERVICES-ACTIVE-1 netflow low 38709416293 24203785663784 132Mbps 33247678310 21332298565174 116Mbps 64946609 242204089910 4.75Mbps
8 sni9      MULTIPLE-SERVICES-ACTIVE-1 dedup low 39352939982
-----
output truncated
```



**Note:** The `show managed-service-device <SN-Name> stats <Managed-service-name>` command filters the statistics of a specific managed service.

The Load column above can reflect no, low, moderate, high, and critical health indicators. These health indicators are represented by green, yellow, and red under **DANZ Monitoring Fabric > Managed Services > Devices > Service Stats**. They reflect the processor load on the service node interface at that instant of time, but do not show the bandwidth of the respective data port (SNI) handling traffic. A sample snapshot of the **Service Stats** output is shown below.

**Figure 5-3: Service Node Interface Load Indicator**

Service Stats

Show: 10 25 100 All (1 - 4 / 4)

Interface Name	Service Name	Table Name	Load	Rx Packet Count	Rx Byte Count	Rx Bit Rate	Tx Packet Count	Tx Byte Count	Tx Bit Rate	Applied Packets	Applied Bytes	Applied Bit Rate
sni1	mngsrv-640-2	mask	Critical	8,949,967,634	1,145,595,854,080	5 Gbit/s	8,949,967,634	1,145,595,854,080	5 Gbit/s	8,949,967,609	1,145,595,853,952	5 Gbit/s
sni2	mngsrv-640-1	slice	Moderate	24,486,431,717	1,714,050,220,190	7 Gbit/s	24,486,431,717	1,714,050,220,190	7 Gbit/s	24,486,431,717	1,714,050,220,190	7 Gbit/s
sni3	mngsrv-640-4	regex	Low	83,878,701	42,946,406,912	192 Mbit/s	0	0	0 bit/s	83,878,701	42,946,406,912	192 Mbit/s
sni4	mngsrv-640-3	strip	Moderate	4,294,441,685	2,198,754,142,720	10 Gbit/s	4,294,441,685	2,104,276,425,650	9 Gbit/s	4,294,441,685	2,198,754,142,720	10 Gbit/s

## 5.5 Deduplication Action

The DMF Service Node enhances the efficiency of network monitoring tools by eliminating duplicate packets. Duplicate packets can be introduced into the out-of-band monitoring data stream by receiving the same flow from multiple TAP or SPAN ports spread across the production network. Deduplication eliminates these duplicate packets and allows a more efficient use of passive monitoring tools.

The DMF Service Node provides three modes of deduplication for different types of duplicate packets.

- Full packet deduplication: deduplicates packets that are incoming and are exactly identical at the L2/L3/L4 layers.
- Routed packet deduplication: as packets traverse a IP network, the mac address changes from hop to hop. Routed packet deduplication allows the user to match on packets from the start of the L3 header onwards.
- Natted packet deduplication: to perform natted deduplication, the service node compares packets in the configured window that are identical from the start of the L4 payload onwards. To use natted packet deduplication, complete the following fields as required:
  - Anchor: Packet Start, L2 Header Start, L3 Header Start, or L3 Payload Start fields.
  - Offset: the number of bytes from the anchor where the deduplication check begins.

The time window within which the service node looks for duplicate packets is configurable. You can select between **2ms** (the default), **4ms**, **6ms**, and **8ms**.

### GUI Configuration

Figure 5-4: Create Managed Service>Action: Deduplication Action

The screenshot shows the 'Create Managed Service' configuration window. The 'Action' dropdown is set to 'Deduplication'. Under 'Packet Handling', the 'None' radio button is selected. The 'Anchor' dropdown is set to 'Packet Start'. The 'Offset' is set to '- Bytes -'. The 'Window Size' is set to '2 ms'. A diagram on the right shows the packet structure: Packet Start (blue), L3 Header (green), L4 Header (blue), and L4 Payload (orange).

### CLI Configuration

```
Controller-1(config)# show running-config managed-service MS-DEDUP-FULL-PACKET
! managed-service
managed-service MS-DEDUP-FULL-PACKET
description 'This is a service that does Full Packet Deduplication'
1 dedup full-packet window 8
service-interface switch CORE-SWITCH-1 ethernet13/1
```

```
Controller-1(config)#
```

```
Controller-1(config)# show running-config managed-service MS-DEDUP-ROUTED-PACKET
! managed-service
managed-service MS-DEDUP-ROUTED-PACKET
description 'This is a service that does Routed Packet Deduplication'
1 dedup routed-packet window 8
service-interface switch CORE-SWITCH-1 ethernet13/2
Controller-1(config)#
```

```
Controller-1(config)# show running-config managed-service MS-DEDUP-NATTED-PACKET
! managed-service
managed-service MS-DEDUP-NATTED-PACKET
description 'This is a service that does Natted Packet Deduplication'
1 dedup anchor-offset l4-payload-start 0 window 8
service-interface switch CORE-SWITCH-1 ethernet13/3
Controller-1(config)#
```



**Note:** The existing command is augmented to show deduplication percentage. The command syntax is **show managed-service-device <SN- name> stats <dedup-service-name>**

## 5.6 Header Strip Action

This action removes specific headers from the traffic selected by the associated DMF policy. Alternatively, you can define custom header stripping based on the starting position of the Layer-3 header, the Layer-4 header, the Layer-4 payload, or the first byte in the packet.

Use the following decap actions isolated from the header-strip configuration stanza:

- **decap-erspan:** remove the Encapsulated Remote Switch Port Analyzer (ERSPAN) header.
- **decap-cisco-fabric-path:** remove the Cisco FabricPath protocol header.
- **decap-l3-mpls:** remove the Layer-3 Multi-protocol Label Switching (MPLS) header.
- **decap-lisp:** remove the LISP header.
- **decap-vxlan [udp-port vxlan port]:** remove the Virtual Extensible LAN (VXLAN) header.
- **decap-geneve:** remove the Geneve header.



**Note:** For the Header Strip and Decap actions, you can apply post-service rules to select traffic after stripping the original headers.

To customize the header-strip action, use one of the following keywords to strip up to the specified location in each packet:

- **l3-header-start**
- **l4-header-start**
- **l4-payload-start**
- **packet-start**

You can also input a positive integer representing the offset from which the strip action begins. If the offset is omitted, the header stripping begins from the first byte in the packet.

## GUI Configuration

Figure 5-5: Create Managed Service: Header Strip Action

1. Info

2. Action ✓

3. Post-Service Match ✓

Action: Header Strip

Strip Type: Custom

Anchor: Packet Start

Offset: - Bytes -

L2 Header Options

If payload after header stripping is L3, an L2 header may be added to the packet according to the options below:

☒ Don't Add L2 Header

Packet is expected to have an L2 header after stripping.

☐ Use Original (Stripped) L2 Header

☒ Get EthType from Post-Strip L3 Header

☐ Use Original EthType from Stripped L2 Header

☐ Insert Inet EthType

☐ Add Custom EthType

☐ Customize L2 Header

EtherType: - Select By Name -

Dst. MAC: - MAC Address -

Src. MAC: - MAC Address -

After assigning the required actions to the header stripping service, click **Next** or click **Post-Service Match**. The system displays the **Post Service Match** page, which can be used with the header strip service action.

Figure 5-6: Create Managed Service: Post Service Match for Header Strip Action

1. Info

2. Action ✓

3. Post-Service Match ✓

Sequence: Any Traffic

No rules

## CLI Configuration

The header-strip service action lets you strip the header and replace it in one of the following ways:

- Add the original L2 **src-mac**, and **dst-mac**.
- Add the original L2 **src-mac**, **dst-mac**, and **ether-type**.
- Specify a completely custom **src-mac**, **dst-mac**, and **ether-type** to be added.

The following are examples of custom header stripping:

This example strips the header and replaces it with the original L2 **src-mac** and **dst-mac**.

```
! managed-service
managed-service MS-HEADER-STRIP-1
1 header-strip packet-start 20 add-original-l2-dstmac-srcmac
service-interface switch CORE-SWITCH-1 ethernet13/1
```

This example adds the original L2 **src-mac**, **dst-mac**, and **ether-type**.

```
! managed-service
managed-service MS-HEADER-STRIP-2
1 header-strip packet-start 20 add-original-l2-dstmac-srcmac-ethertype
service-interface switch CORE-SWITCH-1 ethernet13/2
```

This example specifies a completely custom **src-mac**, **dst-mac** and **ether-type** to be added.

```
! managed-service
managed-service MS-HEADER-STRIP-3
1 header-strip packet-start 20 add-custom-l2-header 00:11:01:02:03:04
  00:12:01:02:03:04
  0x800
service-interface switch CORE-SWITCH-1 ethernet13/3
```

### 5.6.1 Configuring the Post-service Match

The post-service match configuration option enables matching on inner packet fields after the DMF Service Node performs header stripping. This option is applied on the post-service interface after the service node completes the strip service action. Feature benefits include the following:

- The fabric can remain in L3/L4 mode. You do not need to change to offset match mode.
- Easier configuration.
- All match conditions are available for the inner packet.
- Policy requires only one managed service to perform the strip service action.

With this feature enabled, DMF knows exactly where to apply the post-service match. The following is an example of this configuration.

```
! managed-service
managed-service MS-HEADER-STRIP-4
service-interface switch CORE-SWITCH-1 interface ethernet1
1 decap-l3-mpls
!
post-service-match
1 match ip src-ip 1.1.1.1
2 match tcp dst-ip 2.2.2.0 255.255.255.0
! policy
policy POLICY-1
filter-interface TAP-1
delivery-interface TOOL-1
use-managed-service MS-HEADER-STRIP-4 sequence 1
```

## 5.7 IPFIX and Netflow Actions

IP Flow Information Export (IP FIX), also known as NetFlow v10, is an IETF standard defined in **RFC 7011**. The IPFIX generator (agent) gathers and transmits information about flows, which are sets of packets that contain all the keys specified by the IPFIX template. The generator observes the packets received in each flow and forwards the information to the IPFIX collector (server) in the form of a flowset.

Starting with the **BMF-7.1.0** release, NetFlow v9 (Cisco proprietary) and IPFIX/NetFlow v10 are both supported. Configuration of the IPFIX managed service is similar to configuration for earlier versions of NetFlow with the exception of the UDP port definition. NetFlow v5 collectors typically listen over **UDP port 2055**, while IPFIX collectors listen over **UDP port 4739**.

NetFlow records are typically exported using User Datagram Protocol (UDP) and collected using a flow collector. For a NetFlow service, the service node takes incoming traffic and generates NetFlow records. The

---

original packets are dropped by the service node and the generated flow records, containing meta data about each flow, are forwarded out the service node interface.

### 5.7.1 IPFIX Template

The IPXIF template consists of the key element IDs representing IP flow, field element IDs representing actions the exporter has to perform over IP flows matching key element IDs, the template ID number for uniqueness, collector information, and eviction timers

To define a template, configure keys of interest representing the IP flow, and fields which identify the values measured by the exporter, the exporter information, and the eviction timers . To define the template, select the **Monitoring > Managed Service > IPFIX Template** option from the DMF GUI, or enter the `ipfix-template <template-name>` command in config mode, replacing *template- name* with a unique identifier for the template instance.

### 5.7.2 IPFIX Keys

An IPFIX key lets you specify the characteristics of the traffic you want to monitor, such as source and destination MAC or IP address, VLAN ID, Layer-4 port number, and QoS marking. The generator includes flows in a flowset that have all the attributes specified by the keys in the template applied. The flowset is updated only for packets that have all the specified attributes. If a single key is missing, the packet is ignored. To see a listing of the keys supported in the current release of the DMF Service Node, select the **Monitoring > Managed Service > IPFIX Template** option from the DMF GUI, or type `help key` in *config-ipxif-template* submode. The following are the keys supported in the current release:

- `destination-ipv4-address`
- `destination-ipv6-address`
- `destination-mac-address`
- `destination-transport-port`
- `dot1q-priority`
- `dot1q-vlan-id`
- `ethernet-type`
- `icmp-type-code-ipv4`
- `icmp-type-code-ipv6`
- `ip-class-of-service`
- `ip-diff-serv-code-point`
- `ip-protocol-identifier`
- `ip-ttl`
- `ip-version`
- `policy-vlan-id`
- `records-per-dmf-interface`
- `source-ipv4-address`
- `source-ipv6-address`
- `source-mac-address`
- `source-transport-port`
- `vlan id`



**Note:** The `policy-vlan-id` and `records-per-dmf-interface` keys are Arista Proprietary Flow elements. The `policy-vlan-id` key helps to query per-policy flows information at Arista Analytics-node (Collector) in *push-per-policy* deployment mode. The `records-per-dmf-interface` key helps to identify filter interfaces tapping the traffic. The following limitations apply at the time of IPFIX template creation:

- Controller will not allow the key combination of `source-mac-address` and `records-per-dmf-interface` in *push-per-policy* mode.



- Controller will not allow the key combinations of **policy-vlan-id** and **records-per-dmf-interface** in *push-per-filter* mode.

### 5.7.3 IPFIX Fields

A field defines each value that is updated for the packets received by the generator that match the specified keys. For example, fields can be included in the template to record the number of packets, the largest and smallest packet size, or the start and end time of the flowset. To see a listing of the fields supported in the current release of the DMF Service Node, select the **Monitoring > Managed Service > IPFIX Template** option from the DMF GUI, or type `help` in **config-ipxif-template** submode. The following are the fields supported in the current release:

- flow-end-milliseconds*
- flow-end-reason*
- flow-end-seconds*
- flow-start-milliseconds*
- flow-start-seconds*
- maximum-ip-total-length*
- maximum-layer2-total-length*
- maximum-ttl*
- minimum-ip-total-length*
- minimum-layer2-total-length*
- minimum-ttl*
- octet-delta-count*
- packet-delta-count*
- tcp-control-bits*

### 5.7.4 Active and Inactive Timers

After the number of minutes specified by the active timer, the flowset is closed and forwarded to the IPFIX collector. The default active timer is one minute. During the number of seconds set by the inactive timer, if no packets are received that match the flow definition, the flowset is closed and forwarded without waiting for the active timer to expire. The default value for the inactive time is **15** seconds.

### 5.7.5 Example Flowset

The following is a Wireshark view of an IPFIX flowset.

**Figure 5-7: Example IPFIX Flowset in Wireshark**

```

▼ Set 2 [id=22222] (1 flows)
  FlowSet Id: (Data) (22222)
  FlowSet Length: 44
  [Template Frame: 5 (received after this frame)]
▼ Flow 1
  Dot1q Vlan Id: 1700
  SrcPort: 7000
  SrcAddr: 21.0.0.0
  DstPort: 8000
  DstAddr: 20.0.0.3
  Packets: 8514359
  MinTTL: 0
  MaxTTL: 0
  [Duration: 60.001000000 seconds (milliseconds)]
  StartTime: Oct 12, 2018 13:26:02.184000000 PDT
  EndTime: Oct 12, 2018 13:27:02.185000000 PDT
  ► Set 3 [id=22222] (1 flows)

```

The following is a running-config that shows the IPFIX template used to generate this flowset.

## Example IPFIX Template

```
! ipfix-template
ipfix-template Perf-temp
template-id 22222
key destination-ipv4-address
key destination-transport-port
key dot1q-vlan-id
key source-ipv4-address
key source-transport-port
field flow-end-milliseconds
field flow-end-reason
field flow-start-milliseconds
field maximum-ttl
field minimum-ttl
field packet-delta-count
```

### 5.7.6 Using the GUI to Define an IPFIX Template

To define an IPFIX template, complete the following steps:

1. Select the **Monitoring > Managed Services** option.
2. On the **DMF Managed Services** page, select **IPFIX Templates** from the left panel.




The system displays the **IPFIX Templates** section in the right panel.

**Figure 5-8: IPFIX Templates**

### DMF Managed Services

- Devices
- Managed Services
- GTP Correlation Profiles
  - APN Groups
  - IMEI Groups
  - IMSI Groups
  - MSISDN Groups
- IPFIX Templates

### IPFIX Templates



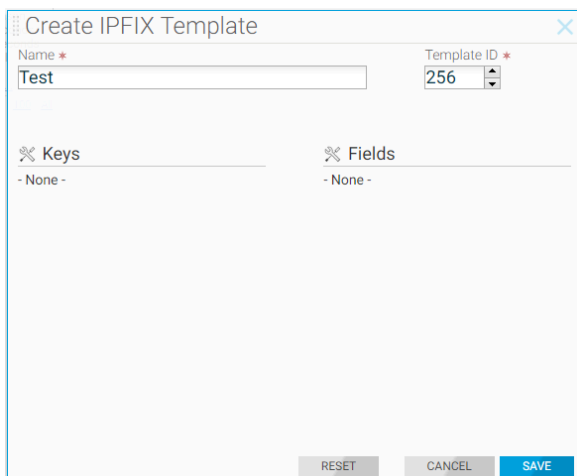
<input type="checkbox"/>	Name ▲	Template Id	Keys	Fields
<input type="checkbox"/>	IPFIX-IP	1974	destination-ipv4-address destination-ipv6-address ethernet-type source-ipv4-address source-ipv6-address	flow-end-milliseconds flow-end-reason flow-start-milliseconds minimum-ttl tcp-control-bits

Nov 26, 2020, 2:45:15am UTC

Show: [10](#) [25](#) [100](#) [All](#)

3. To create a new template, click the provision (+) icon in the upper left corner of the **IPFIX Templates** section.

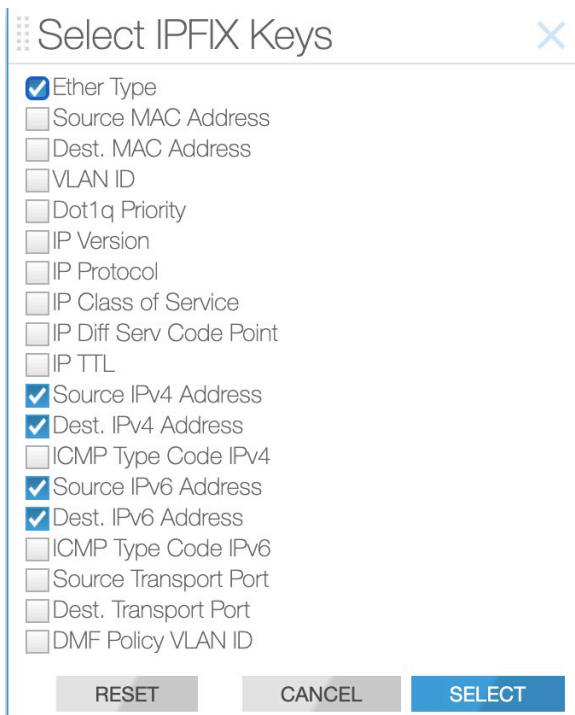
**Figure 5-9: Create IPFIX Template**



The 'Create IPFIX Template' dialog box has a title bar with a close button. It contains two input fields: 'Name' with the value 'Test' and 'Template ID' with the value '256'. Below these are two sections: 'Keys' and 'Fields', both showing '- None -'. At the bottom are three buttons: 'RESET', 'CANCEL', and 'SAVE'.

4. To add an IPFIX key to the template, click the **Settings** control in the **Keys** section. The system displays the following dialog.

**Figure 5-10: Select IPFIX Keys**



The 'Select IPFIX Keys' dialog box has a title bar with a close button. It contains a list of checkboxes for various IPFIX keys. The checked keys are: Ether Type, Source IPv4 Address, Dest. IPv4 Address, Source IPv6 Address, and Dest. IPv6 Address. The unchecked keys are: Source MAC Address, Dest. MAC Address, VLAN ID, Dot1q Priority, IP Version, IP Protocol, IP Class of Service, IP Diff Serv Code Point, IP TTL, ICMP Type Code IPv4, ICMP Type Code IPv6, Source Transport Port, Dest. Transport Port, and DMF Policy VLAN ID. At the bottom are three buttons: 'RESET', 'CANCEL', and 'SELECT'.

5. Enable each checkbox for the keys you want to add to the template and click **Select**.

- To add an IPFIX field to the template, click the **Settings** control in the **Fields** section. The system displays the following dialog:

**Figure 5-11: Select IPFIX Fields**

- Enable the checkbox for each field you want to add to the template and click **Select**.
- On the **Create IPFIX Template** page, click **Save**.

The new template is added to the **IPFIX Templates** table, with each key and field listed in the appropriate column. You can now use this customized template to apply when defining an IPFIX managed service.

### 5.7.7 Using the CLI to Define an IPFIX Template

- Create an IPFX template.

```
controller-1(config)# ipfix-template IPFIX-IP
controller-1(config-ipfix-template)#
```

This changes the CLI prompt to the **config-ipfix-template** submode.

- Define the keys to use for the current template, using the following command:

```
[ no ] key { ethernet-type | source-mac-address | destination-mac-address | dot1q-vlan-id | dot1q-
priority | ip-version | ip-protocol-identifier | ip-class-of-service | ip-diff-serv-code-point | ip-ttl |
sourceipv4-address | destination-ipv4-address | icmp-type-code-ipv4 | source-ipv6-address |
destination-ipv6-address | icmp-type-code-ipv6 | source-transport-port | destination-transport-port
}
```

The keys specify the attributes of the flows to be included in the flowset measurements.

- Define the fields to use for the current template, using the following command:

```
[ no ] field { packet-delta-count | octet-delta-count | minimum-ip-total-length | maximum-ip- total-
length | flow-start-seconds | flow-end-seconds | flow-end-reason | flow-start-milliseconds | flow-
end-milliseconds | minimum-layer2-total-length | maximum-layer2-total- length | minimum-ttl |
maximum-ttl }
```

The fields specify the measurements to be included in the flowset.

You can now use the template when defining the IPFIX action.

## 5.7.8 Using the GUI to Define an IPFIX Service Action

Select IPFIX from the **Action** selection list on the **Create Managed Service > Action** page.

**Figure 5-12: Selecting IPFIX Action in Create Managed Service**

The screenshot shows the 'Create Managed Service' window with the '2. Action' tab selected. The 'Action' dropdown is set to 'IPFIX'. The 'Delivery Interface' is 'DELIVERY-AS-TUNNEL'. The 'Collector IP' is '- IP Address -'. The 'Inactive Timeout' is '15 seconds'. The 'Active Timeout' is '1 minutes'. The 'UDP Port' is '4739'. The 'Template Timeout' is '30 seconds'. The 'MTU' is '1500'. The 'IPFIX Templates' section shows 'IPFIX-IP' selected. The bottom of the window shows a list of services and a 'NEXT' button.

Complete the following required configuration:

- Assign a delivery interface.
- Configure the collector IP address.
- Identify the IPFIX template.

The following configuration is optional:

- Inactive timeout:
- Active timeout: length of time between each IPFIX flowset for a specific flow.
- Source IP: source address to use for the IPFIX flowsets.
- UDP port: UDP port to use for sending IPFIX flowsets.
- MTU: MTU to use for sending IPFIX flowsets.

After completing the configuration, click **Next**, and then click **Save**.

## 5.7.9 Using the CLI to Define an IPFIX Service Action

Define a managed service and define the IPFIX action.

```
controller(config)# managed-service MS-IPFIX-SERVICE
controller(config-managed-srv)# 1 ipfix TO-DELIVERY-INTERFACE
controller(config-managed-srv-ipfix)# collector 10.106.1.60
```

---

```
controller(config-managed-srv-ipfix)# template IPFIX-TEMPLATE
```

The **active-timeout** and **inactive-timeout** commands are optional

To view the running-config for a managed service using the IPFIX action, enter the following command:

```
controller1# show running-config managed-service MS-IPFIX-ACTIVE  
! managed-service  
managed-service MS-IPFIX-ACTIVE  
service-interface switch CORE-SWITCH-1 ethernet13/1  
!  
1 ipfix TO-DELIVERY-INTERFACE  
collector 10.106.1.60  
template IPFIX-TEMPLATE
```

To view the IPFIX templates, enter the following command:

```
config# show running-config ipfix-template  
! ipfix-template  
ipfix-template IPFIX-IP  
template-id 1974  
key destination-ipv4-address  
key destination-ipv6-address  
key ethernet-type  
key source-ipv4-address  
key source-ipv6-address  
field flow-end-milliseconds  
field flow-end-reason  
field flow-start-milliseconds  
field minimum-ttl  
field tcp-control-bits  
-----output truncated-----
```

## 5.8 Packet-masking Action

The packet-masking action lets you hide specific characters in a packet, such as a password or credit card number, based on offsets from different anchors, and by matching characters using regular (regex) expressions.

The mask service action applies the specified mask to the matched packet region.

## GUI Configuration

**Figure 5-13: Create Managed Service: Packet Masking**

**Create Managed Service**

1. Info | **2. Action** | 3. Post-Service Match

Action: **Mask**

Match Pattern:

Match Characters:  Start  End

Anchor:

Offset:  - Bytes -

Packet Structure Diagram:

- Packet Start
- L3 Header
- L4 Header
- L4 Payload

Services Table:

Name	Description	Action	Action Details	Anchor/Offset	Insert Origin Packet Length	Installed
MS-PACKET-MASK	This service masks pattern matching an email address in payload with X	Mask	1 match amazon 2 match ebay 3 match facebook 4 match google 5 match microsoft 6 match netflix 7 match office365 8 match sharepoint 9 match skype 10 match skype			✓

Buttons: BACK, NEXT, RESET, CANCEL, SAVE

## CLI Configuration

```
Controller-1(config)# show running-config managed-service MS-PACKET-MASK
! managed-service
managed-service MS-PACKET-MASK
description "This service masks pattern matching an email address in payload
with X"
1 mask ([a-zA-Z0-9._-]+@[a-zA-Z0-9._-]+.[a-zA-Z0-9._-]+)
service-interface switch CORE-SWITCH-1 ethernet13/1
```

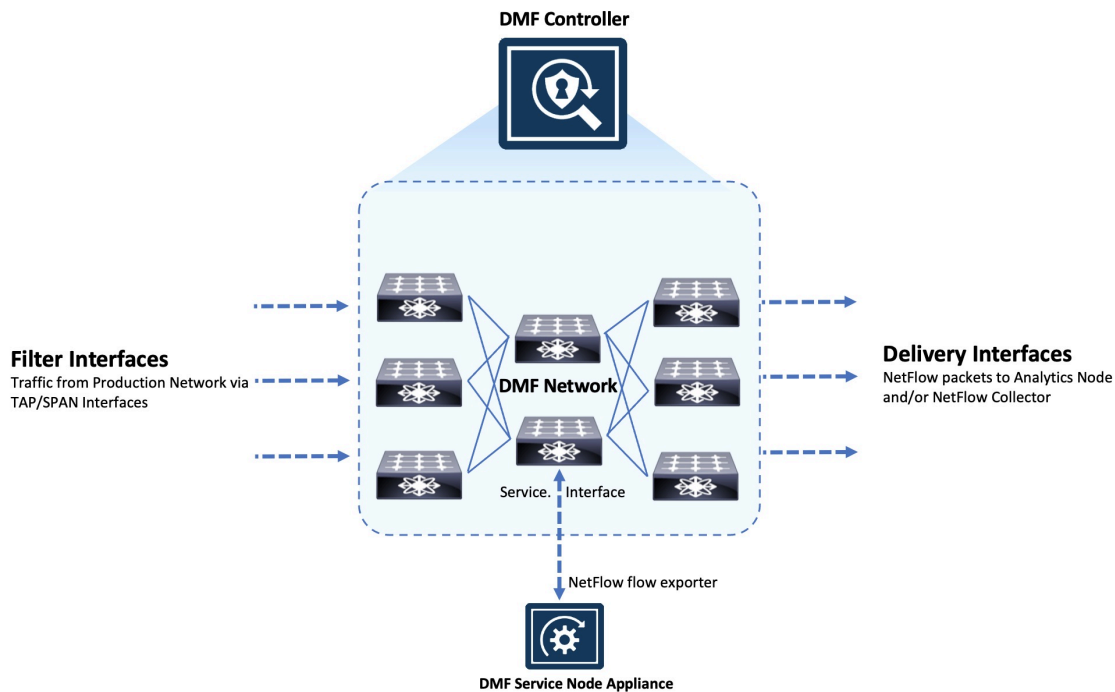
## 5.9 Arista Analytics Node Capability

Starting with **BMF-7.2.1** release Arista Analytics Node capabilities are enhanced to handle NetFlow V5/V9 and IPFIX Packets. All these flows data are represented with the Netflow index.



**Note:** Starting with **BMF Release 7.1.0**, NetFlow flow record generation is enhanced for selecting VxLAN traffic. For VxLAN traffic, flow processing is based on inner headers, with the VNI used as part of the key for flow lookup, because IP addresses can overlap between VNIs.

**Figure 5-14: NetFlow Managed Service**



NetFlow records are exported using User Datagram Protocol (UDP) to one or more specified NetFlow collectors. The NetFlow collector IP address and the destination UDP port must be configured on the DMF Service Node. The default UDP port is **2055**.



**Note:** No other service action, except the UDP replication service, can be applied after a NetFlow service action because part of the NetFlow action is to drop the packets.

### 5.9.1 Configuring the Arista Analytics Node Using the GUI

From the **Arista Analytics Node** dashboard, filter rules can be applied to display specific flow information.

The following are the options available on this page:

- **Delivery interface:** interface to use for delivering NetFlow records to collectors. Note that the next-hop address must be resolved for service to be active.
- **Collector IP:** identify the IP address of the NetFlow collector.
- **Inactive timeout:** use the `inactive-timeout` command to configure the interval of inactivity before NetFlow times out. The default is **15** seconds.
- **Source IP:** specify a source IP address to use as the source of the NetFlow packets.
- **Active timeout:** use active timeout to configure a period of time that a NetFlow can be generated continuously before it is automatically terminated. The default is one minute.
- **UDP port:** change the UDP port number used for the NetFlow packets. The default is **2055**.
- **Flows:** specify the maximum number of NetFlow packets allowed. The allowed range is **32768** to **1048576**. The default is **262144**.
- **Per-interface records:** identify the filter interface where the NetFlow packets were originally received. This information can be used to identify the hop-by-hop path from the filter interface to the NetFlow collector.



- MTU: change the Maximum Transmission Unit (MTU) used for NetFlow packets.

**Figure 5-15: Create Managed Service: NetFlow Action**

## 5.9.2 Configuring the Arista Analytics Node Using the CLI

Use the **show managed-services** command to display ARP resolution status.



**Note:** The DMF controller resolves ARP messages for each NetFlow collector IP address on the delivery interface that matches the defined subnet. The subnets defined on the delivery interfaces cannot overlap, and must be unique for each delivery interface.

When you enter the **1 netflow** command and identify the name of the configuration, the submode changes to the **config-managed-srv-netflow** mode, where you can view and configure a specific NetFlow configuration.

The DMF Service Node replicates NetFlow packets received without changing the source IP address. Packets that do not match the specified destination IP address as well as packets that are not IPv4 or UDP are passed through. To configure a NetFlow managed service, complete the following steps:

1. Configure the IP address on the delivery interface.

This IP address should be the next-hop IP address from the DANZ Monitoring Fabric towards the NetFlow collector.

```
CONTROLLER-1(config)# switch DMF-DELIVERY-SWITCH-1
CONTROLLER-1(config-switch)# interface ethernet1
CONTROLLER-1(config-switch-if)# role delivery interface-name NETFLOW-DELIV
ERY-PORT ip-address 172.43.75.1 nexthop-ip 172.43.75.2 255.255.255.252
```

2. Configure the rate-limit for the NetFlow delivery interface.

```
CONTROLLER-1(config)# switch DMF-DELIVERY-SWITCH-1
CONTROLLER-1(config-switch)# interface ethernet1
CONTROLLER-1(config-switch-if)# role delivery interface-name NETFLOW-DELIV
ERY-PORT ip-address 172.43.75.1 nexthop-ip 172.43.75.2 255.255.255.252
CONTROLLER-1(config-switch-if)# rate-limit 256000
```



**Note:** The rate limit must be configured when enabling Netflow. When upgrading from a version of DMF before release **6.3.1**, the Netflow configuration is not applied until a rate limit is applied to the delivery interface.

3. Configure the NetFlow managed service using the **1 netflow** command followed by an identifier for the specific NetFlow configuration.

```
CONTROLLER-1(config)# managed-service MS-NETFLOW-SERVICE CONTROLLER-1
(config-managed-srv)# 1 netflow NETFLOW-DELIVERY-PORT CONTROLLER-1
(config-managed-srv-netflow)#
```

The following commands are available in this submode:

- **active-timeout:** configure the maximum length of time the NetFlow is transmitted before it is ended (in minutes).
- **collector:** configure the collector IP address, and change the UDP port number or the MTU.
- **inactive-timeout:** configure the length of time that the NetFlow is inactive before it is ended (in seconds).
- **max-flows:** configure the maximum number of flows managed.

You can optionally limit the number of flows or change the inactivity timeout using the **max-flows** or **active-timeout**, or **inactive-timeout** commands.

4. Configure the IP address of the NetFlow collector using the following command:

```
collector <ip4-address>[udp-port<integer>][mtu <integer>][records-per-int
erface]
```

The IP address, in IPV4 dotted-decimal notation, is required. The MTU and UDP port are required only if you want to change these parameters from the defaults. You can enable the **records-per-interface** option to allow identification of the filter interfaces from which the Netflow originated. You can configure the Arista Analytics Node to display this information, as described in the **DMF User Guide**.

The following is an example of changing the Netflow UDPF port to **9991**.

```
collector 10.181.19.31 udp-port 9991
```



**Note:** The IP address must be in the same subnet as the configured next hop and must be unique. It cannot be the same as the controller, service node, or any monitoring fabric switch IP address.

5. Configure the DMF policy with the forward action and add the managed service to the policy.



**Note:** A DMF policy does not require any configuration related to a delivery interface for NetFlow policies because the DMF controller automatically assigns the delivery interface.

The example below shows the configuration required to implement two NetFlow service instances (MS-NETFLOW-1 and MS-NETFLOW-1).

```
! switch
switch DMF-DELIVERY-SWITCH-1
!
interface ethernet1
role delivery interface-name NETFLOW-DELIVERY-PORT-1 ip-address 10.3.1.1
nexthop-ip 10.3.1.2 255.255.255.0
interface ethernet2
```

```

role delivery interface-name NETFLOW-DELIVERY-PORT-2 ip-address 10.3.2.1
nexthop-ip 10.3.2.2 255.255.255.0
! managed-service
managed-service MS-NETFLOW-1
service-interface switch DMF-CORE-SWITCH-1 interface ethernet11/1
!
1 netflow NETFLOW-DELIVERY-PORT-1
collector-ip 10.106.1.60 udp-port 2055 mtu 1024
managed-service MS-NETFLOW-2
service-interface switch DMF-CORE-SWITCH-2 interface ethernet12/1
!
1 netflow NETFLOW-DELIVERY-PORT-1
collector-ip 10.106.2.60 udp-port 2055 mtu 1024
! policy
policy GENERATE-NETFLOW-1
action forward
filter-interface TAP-INTF-DC1-1
filter-interface TAP-INTF-DC1-2
use-managed-service MS-NETFLOW-1 sequence 1
1 match any
policy GENERATE-NETFLOW-2
action forward
filter-interface TAP-INTF-DC2-1
filter-interface TAP-INTF-DC2-2
use-managed-service MS-NETFLOW-2 sequence 1
1 match any

```

## 5.10 Pattern-drop Action

The pattern-drop service action drops matching traffic.

Pattern matching allows content-based filtering beyond Layer-2, Layer-3, or Layer-4 Headers. This functionality allows filtering on the following packet fields and values:

- URLs and user agents in the HTTP header
- patterns in BitTorrent packets
- encapsulation headers for specific parameters including GTP, VXLAN, and VN-Tag
- subscriber device IP (user-endpoint IP)

Pattern matching allows Session-aware Adaptive Packet Filtering (SAPF) and can identify HTTPS transactions on non-standard SSL ports. It can filter custom applications, and can separate control traffic from user data traffic.

Pattern matching is also useful for enforcing IT policies, such as identifying hosts using unsupported operating systems or dropping unsupported traffic. For example, the Windows OS version can be identified and filtered based on the user-agent field in the HTTP header. The user-agent field may appear at variable offsets, so a regular expression search is used to identify the specified value wherever it occurs in the packet.

## GUI Configuration

Figure 5-16: Create Managed Service: Pattern Drop Action

**Create Managed Service**

1. Info  
2. Action ✓  
3. Post-Service Match ✓

Action: **Pattern Drop**

Pattern:

Anchor:

Offset:

Packet Structure Diagram:

- Packet Start
- L3 Header
- L4 Header
- L4 Payload

Services Table:

Name	Description	Action	Action Details	Anchor/Offset	Insert Origin Packet Length	Installed	Supers
IP-FILTER-ACTIVE		Application ID	1 match email 2 match ebay 3 match facebook 4 match google 5 match microsoft 6 match netflix 7 match office365 8 match skype 9 match skype 10 match skype				
IP-FILTER-BACKUP		Application ID	1 match email 2 match ebay 3 match facebook				

Buttons: BACK, NEXT, RESET, CANCEL, SAVE

## CLI Configuration

```
Controller-1(config)# show running-config managed-service MS-PACKET-MASK
! managed-service
managed-service MS-PACKET-MASK
description "This service drops traffic that has an email address in its
payload"
1 pattern-drop ([a-zA-Z0-9._-]+@[a-zA-Z0-9._-]+\.[a-zA-Z0-9._-]+)
service-interface switch CORE-SWITCH-1 ethernet13/1
```

## 5.11 Pattern-match Action

The pattern-match service action matches and forwards matching traffic and is otherwise similar to the pattern-drop service action.

Pattern matching allows content-based filtering beyond Layer-2, Layer-3, or Layer-4 Headers. This functionality allows filtering on the following packet fields and values:

- URLs and user agents in the HTTP header
- patterns in BitTorrent packets
- encapsulation headers for specific parameters including GTP, VXLAN, and VN-Tag
- subscriber device IP (user-endpoint IP)

- Pattern matching allows Session Aware Adaptive Packet Filtering and can identify HTTPS transactions on non-standard SSL port. It can filter custom applications and can separate control traffic from user data traffic.

Pattern matching allows Session-aware Adaptive Packet Filtering (SAPF) and can identify HTTPS transactions on non-standard SSL ports. It can filter custom applications, and can separate control traffic from user data traffic.

Pattern matching is also useful for enforcing IT policies, such as identifying hosts using unsupported operating systems or dropping unsupported traffic. For example, the Windows OS version can be identified and filtered based on the user-agent field in the HTTP header. The user-agent field may appear at variable offsets, so a regular expression search is used to identify the specified value wherever it occurs in the packet.

## GUI Configuration

**Figure 5-17: Create Managed Service: Pattern Match Action**

The screenshot displays the 'Create Managed Service' configuration window. The 'Action' tab is active, showing the 'Pattern Match' configuration. The 'Pattern' field is empty. The 'Anchor' dropdown is set to 'Start'. The 'Offset' dropdown is set to '- Bytes -'. A diagram on the right illustrates the packet structure layers: Packet Start, L3 Header, L4 Header, and L4 Payload. The bottom of the window features a list of services and a set of action buttons: BACK, NEXT, RESET, CANCEL, and SAVE.

## CLI Configuration

Use the **pattern-match pattern** keyword for enabling the pattern-matching service action. Specify the pattern to match for packets to submit to the packet slicing operation.

The following example matches traffic with the string **Windows NT 5. (0-1)** anywhere in the packet and delivers the packets to the delivery interface **TOOL-PORT-TO-WIRESHARK-1**. This service is optional, and is applied to TCP traffic to destination **port 80**.

```
! managed-service
managed-service MS-PATTERN-MATCH
description 'regular expression filtering'
```

```

1 pattern-match 'Windows\\sNT\\s5\\. [0-1]'
service-interface switch CORE-SWITCH-1 ethernet13/1
! policy
policy PATTERN-MATCH
action forward
delivery-interface TOOL-PORT-TO-WIRESHARK-1
description 'match regular expression pattern'
filter-interface TAP-INTF-FROM-PRODUCTION
priority 100
use-managed-service MS-PATTERN-MATCH sequence 1 optional
1 match tcp dst-port 80

```

## 5.12 Slice Action

The slice service action slices the given number of packets based on the specified starting point in the packet. Packet slicing reduces packet size to increase processing and monitoring throughput. Passive monitoring tools process fewer bits while maintaining the vital, relevant portions of each packet. Packet slicing can significantly increase the capacity of forensic recording tools. Packet slicing can be applied by specifying the number of bytes to forward based on an offset from the following locations in the packet:

- Packet start
- L3 header start
- L4 header start
- L4 payload start

### GUI Configuration

**Figure 5-18: Create Managed Service: Slice Action**

The screenshot shows the 'Create Managed Service' window with the '2. Action' tab selected. The 'Action' dropdown is set to 'Slice'. Below it, the 'Insert Original Packet Length' section has a 'No' radio button selected. The 'Anchor' dropdown is set to 'Packet Start'. The 'Offset' is set to '- Bytes -'. The 'L2 Header Options' section shows 'Don't Add L2 Header' selected. The 'EthernetType' dropdown is set to '- Select By Name -'. The 'Dst. MAC' and 'Src. MAC' fields are set to '- MAC Address -'.

This page provides the option to insert an additional header containing the original header length.

## CLI Configuration

Use the **slice** keyword for enabling the packet slicing service-action.

You can insert an additional header containing the original header length, as shown in the following example:

```
! managed-service
managed-service my-service-name
1 slice l3-header-start 20 insert-original-packet-length
service-interface switch DMF-CORE-SWITCH-1 ethernet20/1
```

The following example truncates the packet from the first byte of the Layer-4 payload, preserving just the original Ethernet header. The service is optional and is applied to all TCP traffic from **port 80** with the destination IP address **10.2.19.119**

```
! managed-service
managed-service MS-SLICE-1
description 'slicing service'
1 slice l4-payload-start 1
service-interface switch DMF-CORE-SWITCH-1 ethernet40/1
! policy
policy slicing-policy
action forward
delivery-interface TOOL-PORT-TO-WIRESHARK-1
description 'remove payload'
filter-interface TAP-INTF-FROM-PRODUCTION
priority 100
use-managed-service MS-SLICE-1 sequence 1 optional
1 match tcp dst-ip 10.2.19.119 255.255.255.255 src-port 80
```

## 5.13 Packet Slicing on the 7280 Switch

This feature removes unwanted or unneeded bytes from a packet at a configurable byte position (offset). This approach is beneficial when the data of interest is situated within the headers or early in the packet payload. This action reduces the volume of the monitoring stream, particularly in cases where payload data is not necessary.

Another use case for packet slicing (slice action) can be removing payload data to ensure compliance with the captured traffic.

Within the DMF fabric, two types of slice-managed services (packet slicing service) now exist. These types are distinguished based on whether installing the service on a service node or an interface of a supported switch. The scope of this document is limited to the slice-managed service configured on a switch. The managed service interface is the switch interface used to configure this service.

All DMF 8.4 compatible 7280 switches support this feature. Use the **show switch all property** command to check which switch in DMF fabric supports this feature. The feature is supported if the **Min Truncate Offset** and **Max Truncate Offset** properties have a non-zero value.

```
# show switch all property
# Switch Min Truncate Offset ... Max Truncate Offset
-|-----|-----| ... |-----|
1 7280 100 ... 9236
2 core1 ...
```



**Note:** The CLI output example above is truncated for illustrative purposes. The actual output will differ.

## 5.13.1 Using the CLI to Configure Packet Slicing - 7280 Switch

Use the following steps to configure **slice** managed service on a switch.

1. Create a managed service using the **managed-service service name** command.
2. Add **slice** action with **packet-start** anchor and an **offset** value between the supported range as reported by the **show switch all property** command.
3. Configure the service interface under the **config-managed-srv** submode using the service-interface switch **switch-name interface-name** command as shown in the following example.

```
> enable
# config
(config)# managed-service slice-action-7280-J2-J2C
(config-managed-srv)# 1 slice packet-start 101
(config-managed-srv)# service-interface switch 7280-J2-J2C Ethernet10/1
```

1. This feature requires the service interface to be in MAC loopback mode.
4. To set the service interface in MAC loopback mode, navigate to the **config-switch-if** submode and configure using the **loopback-mode mac** command as shown in the following example..

```
(config)# switch 7280-J2-J2C
(config-switch)# interface Ethernet10/1
(config-switch-if)# loopback-mode mac
```

2. Once a managed service for slice action has been created, any policy can use it.
5. Enter the **config-policy** submode, and chain the managed service using the **use-managed-service service same sequence sequence** command.

```
(config)# policy timestamping-policy
(config-policy)# use-managed-service slice-action-7280-J2-J2C sequence 1
```

Key points to consider while configuring the **slice** action on a supported switch:

1. Only the **packet-start** anchor is supported.
2. Offset should be within the Min/Max truncate size bounds reported by the **show switch all property** command. If the configured value is beyond the bound, then we choose the closest value of the range.

For example, if the user configures the offset as 64, and the min truncate offset reported by switch properties is 100, then the offset used is 100. If the configured offset is 10,000 and the max truncate offset reported by the switch properties is 9236, then the offset used is 9236.

3. A configured offset for slice managed service includes FCS when programmed on a switch interface and means an offset of 100 will result in a packet size of 96 bytes (accounting for 4-byte FCS).
4. Configuring an offset below 17 is not allowed.
5. The same service interface cannot chain multiple managed services.
6. The **insert-original-packet-length** option is not applicable for switch-based slice managed service.

### CLI Show Commands

Use the **show policy policy name** command to see the runtime state of a policy using the slice managed service. The command shows the service interface information and stats.

```
Controller# show policy packet-slicing-policy
Policy Name      : packet-slicing-policy
Config Status    : active - forward
Runtime Status   : installed
Detailed Status  : installed - installed to forward
Priority         : 100
Overlap Priority  : 0
# of switches with filter interfaces : 1
# of switches with delivery interfaces : 1
# of switches with service interfaces : 1
# of filter interfaces : 1
# of delivery interfaces : 1
# of core interfaces : 0
# of services : 1
# of pre service interfaces : 1
# of post service interfaces : 1
```



```

Push VLAN           : 1
Post Match Filter Traffic : -
Total Delivery Rate  : -
Total Pre Service Rate : -
Total Post Service Rate : -
Overlapping Policies : none
Component Policies  : none
Runtime Service Names : packet-slicing-7280
Installed Time       : 2023-08-09 19:00:40 UTC
Installed Duration    : 1 hour, 17 minutes
~ Match Rules ~
# Rule
-|-----|
1 1 match any

~~~~~ Filter Interface(s) ~~~~~
# DMF IF Switch IF Name      State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 fl      7280  Ethernet2/1 up   rx    0         0      0      -      2023-08-09 19:00:40.305000 UTC

~~~~~ Delivery Interface(s) ~~~~~
# DMF IF Switch IF Name      State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 dl      7280  Ethernet3/1 up   tx    0         0      0      -      2023-08-09 19:00:40.306000 UTC

~~~~~ Service Interface(s) ~~~~~
# Service name      Role Switch IF Name      State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 packet-slicing-7280 pre 7280  Ethernet10/1 up   tx    0         0      0      -      2023-08-09 19:00:40.305000 UTC
2 packet-slicing-7280 post 7280  Ethernet10/1 up   rx    0         0      0      -      2023-08-09 19:00:40.306000 UTC

~ Core Interface(s) ~
None.

~ Failed Path(s) ~
None.

```

Use the **show managed-services** command to view the status of all the managed services, including the packet-slicing managed service on a switch.

```

Controller# show managed-services
~~~~~ Managed-services ~~~~~
# Service Name      Switch Switch Interface Installed Max Post-Service BW Max Pre-Service BW Total Post-Service BW Total Pre-
Service BW
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 packet-slicing-7280 7280  Ethernet10/1   True      400Gbps      400Gbps      80bps 80bps

~~~~~ Actions of Service Names ~~~~~
# Service Name      Sequence Service Action Slice Anchor Insert original packet length Slice Offset
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 packet-slicing-7280 1          slice      packet-start False      101

```

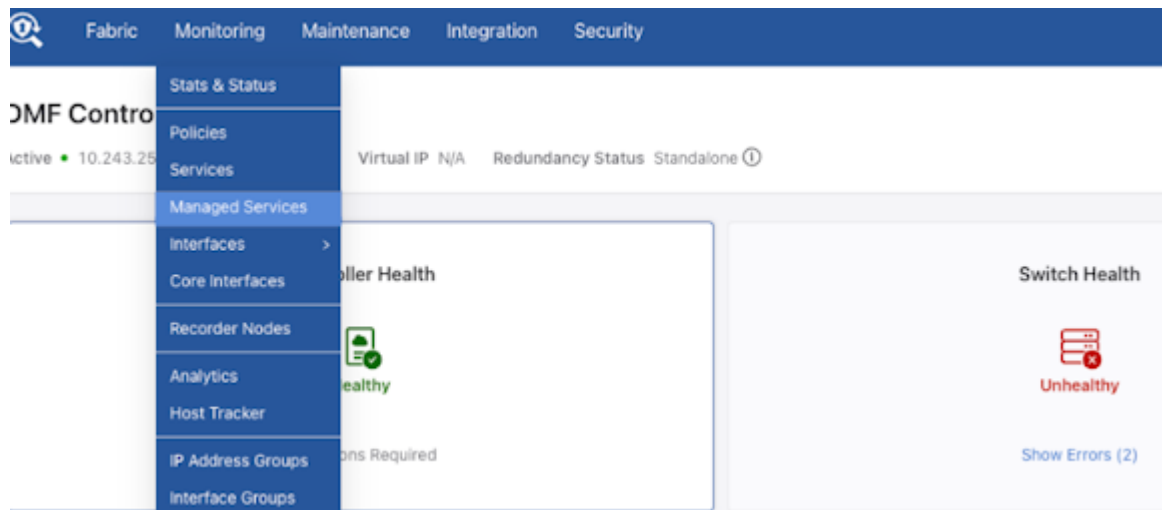
### 5.13.2 Using the GUI to Configure Packet Slicing - 7280 Switch

Perform the following steps to configure or edit a managed service.

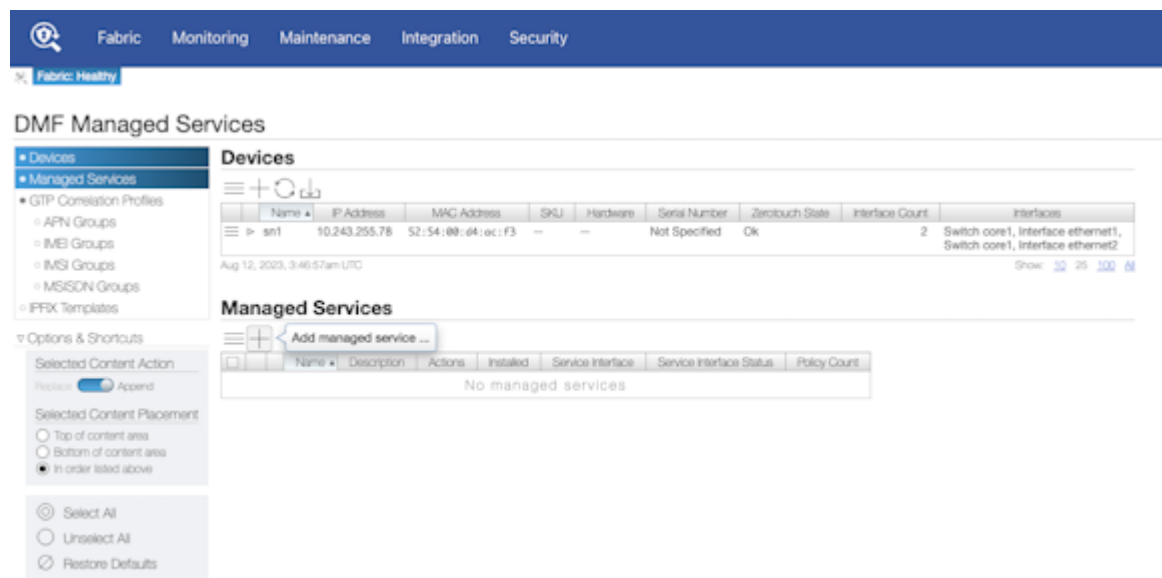
#### Managed Service Configuration

1. To configure or edit a managed service, navigate to the **DMF Managed Services** page from the **Monitoring** menu and click **Managed Services**.

**Figure 5-19: Managed Services**



**Figure 5-20: DMF Managed Services Add managed service**



2. Configure a managed service interface on a switch that supports packet slicing. Make sure to deselect the **Show Managed Device Switches Only** checkbox.

**Figure 5-21: Create Managed Service**

**Create Managed Service**

1. Info ✓ 2. Actions ✓ 3. Post-Service Match ✓

Name:

Description:

Switch:  ☐ Show Managed Device Switches Only  
Optional: use to filter interface choices below ☐ Show Connected Switches Only

Interface:   
Choose from a list of interfaces not already configured for DMF

BACK NEXT RESET CANCEL SAVE

3. Configure a new managed service action using **Add Managed service action**. The action chain supports only one action when configuring packet slicing on a switch.

**Figure 5-22: Add Managed service action**

**Create Managed Service**

1. Info ✓ 2. Actions ✓ 3. Post-Service Match ✓

Click + and - to include or exclude actions.

+ Add Managed service action ...

No actions

BACK NEXT RESET CANCEL SAVE

4. Use **Action > Slice** with **Anchor > Packet Start** for configuring the packet slicing managed service on a switch.

**Figure 5-23: Configure Managed Service Action**

**Configure Managed Service Action**

Sequence:  Click + and - to include or exclude actions.

Action:

Insert Original Packet Length: No ☒ Yes ☐  
Use to insert original pre-slice packet length into additional header

Anchor:  Offset:

Diagram: Packet Start (blue), L3 Header (green), L4 Header (blue), L4 Payload (brown)

BACK NEXT CANCEL APPEND

- Click **Append** to continue. The slice action appears in the **Managed Services** page.

**Figure 5-24: Slice Action Added**

**Managed Services**

Name	Description	Actions	Installed	Service Interface	Service Interface Status	Policy Count
slice-action-7280-J2-J2C	-	Slice	✓	7280-J2-J2C / Ethernet10/1	✓ Up	0

Aug 12, 2023, 3:55:11am UTC Show: 10 25 100 All

## 1. Interface Loopback Configuration

The managed service interface used for slice action must be in MAC loopback mode.

- Configure the loopback mode in the **Fabric > Interfaces** page by clicking on the **configuration icon** of the interface.

**Figure 5-25: Interfaces**

**Interfaces**

Fabric: Healthy

Filter table rows FILTER

Switch	Switch DPID	Interface Name	Status	Tunnel	LAG	LAG Parent	Device Name	Speed
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet9/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet8/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet7/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet6/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet5/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet4/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet3/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet24/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet23/1	✗ Link Down	-	-	-	-	0 bit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet22/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet21/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet20/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet2/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet19/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet18/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet17/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet16/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet15/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet14/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet13/1	✓ Up	-	-	-	-	400 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet12/1	✓ Up	-	-	-	-	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet11/1	✓ Up	-	-	-	-	400 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet10/1	✓ Up	-	-	-	7280-J2-J2C	100 Gbit/s
7280-J2-J2C	00:00:98:5d:82:84:ca:b5	Ethernet1/1	✗ Link Down	-	-	-	-	0 bit/s
core1	00:00:52:54:00:ec:e9:c9	ethernet9	✗ Link Down	-	-	-	-	10 Gbit/s

Aug 14, 2023, 7:59:17pm UTC



**Note:** The image above has been edited for documentation purposes. The actual output will differ.

7. Enable the toggle for **MAC Loopback Mode** (set the toggle to **Yes**).

**Figure 5-26: Edit Interface**

**Edit Interface**

1. Port ✓ 7280-J2-J2C / Ethernet10/1

2. Traffic ✓

3. DMF ✓

Description

Admin Status: Down Up

Enable Optics: Only When Up Always

Disable Transmitting Packets: No Yes

MAC Loopback Mode: No Yes

Force Link Up: No Yes

Use for Management Traffic: No Yes

Buttons: BACK NEXT RESET CANCEL SAVE

8. After all configuration changes are done **Save** the changes.
2. **Policy Configuration**
9. Create a new policy from the **DMF Policies** page.

**Figure 5-27: DMF Policies Page**

**DMF Policies**

Filter table rows

Buttons: FILTER

No policies

10. Add the previously configured packet slicing managed service.

Figure 5-28: Create Policy

The screenshot shows the 'Create Policy' page with the following details:

- Policy Information:**
  - Name: packet-slicing-policy
  - Description: Input value
- Scheduling:**
  - Start Policy: Automatically (selected), Set Time, Set Delay
  - Run Policy: Always (selected), For Duration
- Additional Advanced Options:**
  - Actions: Forward
  - Priority: 100
  - Push VLAN: All
  - Root Switch: No Switch Configured
  - Enable PTP Timestamping: Off

11. Select **Managed Service** under the **Service Type** option and the previously configured slice managed service.

Figure 5-29: Service Type - Managed Service

The screenshot shows the 'Edit Policy' page with the 'Services' tab selected. The table below shows the assigned services:

Sequence ID	Service Type	Service	Optional	Backup Service	Use as Delivery Service
1	Managed Service	slice-action-7280-J2-J2C	No	Backup Service Name	No

12. The slice managed service (packet-slicing-policy) appears in the DMF Policies page.

Figure 5-30: DMF Policy Configured

The screenshot shows the 'DMF Policies' page with the following table:

Name	Description	Runtime Status	Config Status	Type	Action	Active	Priority	Push VLAN	Filter Interfaces	Delivery Interfaces	Flow Count
packet-slicing-policy		Installed	active - forward	Configured	Forward	Yes	100	1	f1 - 7280-J2-J2C / Ethernet2/1	e1 - 7280-J2-J2C / Ethernet2/1	4



**Note:** The image above has been edited for documentation purposes. The actual output will differ.

### 5.13.3 Troubleshooting Packet Slicing

The **show switch all property** command provides upper and lower bounds of packet slicing action's offset. If bounds are present, the feature is supported; otherwise, the switch does not support the packet slicing feature.

The **show fabric errors managed-service-error** command provides information when DMF fails to install a configured packet slicing managed service on a switch.

The following are some of the failure cases:

1. The managed service interface is down.
2. More than one action is configured on a managed service interface of the switch.
3. Managed service interface on a switch is neither a physical interface nor a LAG port.
4. A non-slice managed service is configured on a managed service interface of a switch.
5. The switch does not support packet slicing managed service, and its interface is configured with slice action.
6. Slice action configured on a switch interface is not using a packet-start anchor.
7. The managed service interface is not in MAC loopback mode.

Use the following commands to troubleshoot packet slicing issues.

```
Controller# show fabric errors managed-service-error
Managed Service related error
# Error      Service Name
-----
1 Pre-service interface 7280-Ethernet10/1-to-managed-service on switch 7280 is inactive; Service interface Ethernet10/1 on switch 7280 is down packet-slicing-7280
2 Post-service interface 7280-Ethernet10/1-to-managed-service on switch 7280 is inactive; Service interface Ethernet10/1 on switch 7280 is down packet-slicing-7280
```

The **show switch switch name interface interface name dmf-stats** command provides Rx and Tx rate information for the managed service interface.

```
Controller# show switch 7280 interface Ethernet10/1 dmf-stats
# Switch DPID Name      State Rx Rate Pkt Rate Peak Rate Peak Pkt Rate TX Rate Pkt Rate Peak Rate Peak Pkt Rate Pkt Drop Rate
1 7280      Ethernet10/1 down -      0      128bps      0 -      0      128bps      0 0
```

The **show switch switch name interface interface name stats** command provides Rx and Tx counter information for the managed service interface.

```
Controller# show switch 7280 interface Ethernet10/1 stats
# Name      Rx Pkts Rx Bytes Rx Drop Tx Pkts Tx Bytes Tx Drop
1 Ethernet10/1 22      843477 0      5140      845937 0
```

#### Considerations

1. Managed service action chaining is not supported when a switch interface is used as a managed service interface.
2. When configured for a supported switch, the managed service interface for slice action can only be a physical interface or a LAG.
3. When using packet slicing managed service, packets ingressing on the managed service interface are not counted in the ingress interface counters, which affects the output of the **show switch switch name interface interface name stats** and **show switch switch name interface interface name dmf-stats** commands. This issue does not impact byte counters; all byte counters will show the original packet size, not the truncated size.

## 5.14 Session-slice Action

The session-slice action tracks the state of a TCP session (distinguished by its source IP address, source port, destination IP address and destination port), and counts the number of packets sent in both directions (client-to-server and server-to-client). After recognizing the session, the action transmits a user-configured number of packets to the tool node.

A session is usually identified by tracking the packets of the three-way TCP handshake that establishes the session. However, observing the three-way handshake is not necessary, since a new session will be created for any TCP packet that does not match an existing session.

Once a TCP session has been recognized and the session-slice action is applied, the service node tracks packets in both directions, and drops them after the counts in both directions meet a threshold configured by the user on the controller.



**Note:** The count of packets in one direction may exceed the user-configured threshold, because fewer packets have arrived in the other direction. Counts in both directions must be greater than or equal to the threshold before packets are dropped.

A maximum of **512K** IPv4 and **512K** IPv6 sessions can be tracked and sliced simultaneously per service-node interface.

### GUI Configuration

Figure 5-31: Create Managed Service: Slice Action

The screenshot shows the 'Create Managed Service' GUI. The '2. Action' tab is selected, and the 'Action' dropdown is set to 'Slice'. The 'Insert Original Packet Length' section has a 'No' button selected. The 'Anchor' dropdown is set to 'Word'. The 'Offset' is set to '-12, word'. The 'L2 Header Options' section has 'Don't Add L2 Header' selected. The 'Customize L2 Header' section has 'EtherType' set to 'Select By Name', 'Dst. MAC' set to 'MAC Address', and 'Src. MAC' set to 'MAC Address'. A diagram on the right shows a packet structure with layers: Packet Start, L3 Header, L4 Header, and L4 Payload.

This page provides the option to configure the number of packets the service node accounts for.

### CLI Configuration

Use the **session-slice** keyword to enable TCP the session-slice service action.



The following example allows the service node to account for a new TCP session followed with 10 packets. Additional packets received from this TCP session are dropped.

```
! managed-service
managed-service SESSION-SLICE
!
1 session-slice
slice-after 10
```

## 5.15 Timestamp Action

The timestamp service action identifies and timestamps every packet it receives with the time that the packet is received by the service node for matching traffic.

### GUI Configuration

**Figure 5-32: Create Managed Service: Timestamp Action**

The screenshot shows the 'Create Managed Service' window with the 'Timestamp' action selected. The '1. Info' tab is active, showing the service name 'SESSION-SLICE' and the action 'Timestamp'. The '2. Action' tab is also visible, showing a list of interfaces for the service. The '3. Post-Service Match' tab is also visible, showing a list of match rules for the service.

Interface Count	Interfaces
10	Switch supaspine-ss7016-1, interface ethernet15/1, interface ethernet15/2, interface ethernet15/3, interface ethernet15/4, interface ethernet15/5, interface ethernet15/6, interface ethernet15/7, interface ethernet15/8, interface ethernet15/9, interface ethernet15/10, interface ethernet15/11, interface ethernet15/12, interface ethernet15/13, interface ethernet15/14, interface ethernet15/15, interface ethernet15/16, interface ethernet15/17, interface ethernet15/18, interface ethernet15/19, interface ethernet15/20
10	Switch supaspine-ss7016-1, interface ethernet15/1, interface ethernet15/2, interface ethernet15/3, interface ethernet15/4, interface ethernet15/5, interface ethernet15/6, interface ethernet15/7, interface ethernet15/8, interface ethernet15/9, interface ethernet15/10, interface ethernet15/11, interface ethernet15/12, interface ethernet15/13, interface ethernet15/14, interface ethernet15/15, interface ethernet15/16, interface ethernet15/17, interface ethernet15/18, interface ethernet15/19, interface ethernet15/20

Name	Description	Action	Action Data	Anchor/Offset	Insert Origin	Packet Length	Installed	State
TS-FILTER-ACTIVE	Application	1 match amazon, 2 match ebay, 3 match facebook, 4 match google, 5 match microsoft, 6 match netflix, 7 match office365, 8 match onepoint, 9 match skype, 10 match skype					✓	active
TS-FILTER-BACKUP	Application	1 match amazon, 2 match ebay, 3 match facebook, 4 match google, 5 match microsoft, 6 match netflix, 7 match office365, 8 match onepoint, 9 match skype, 10 match skype					✓	active

Buttons: BACK, NEXT, RESET, CANCEL, SAVE

### CLI Configuration

```
! managed-service
managed-service MS-TIMESTAMP-1
1 timestamp
service-interface switch CORE-SWITCH-1 ethernet15/3
```

## 5.16 UDP-replication Action

The UDP-replication service action copies UDP messages, such as syslog or NetFlow messages, and sends the copied packets to a new destination IP address.

A rate limit must be configured when enabling UDP replication. When upgrading from a version of BMF before release **6.3.1**, the UDP-replication configuration is not applied until a rate limit is applied to the delivery interface.

The following is an example of applying a rate limit to a delivery interface used for UDP replication:

```
CONTROLLER-1(config)# switch DMF-DELIVERY-SWITCH-1
CONTROLLER-1(config-switch)# interface ethernet1
CONTROLLER-1(config-switch-if)# role delivery interface-name udp-delivery-1
CONTROLLER-1(config-switch-if)# rate-limit 256000
```



**Note:** No other service action can be applied after a UDP-replication service action.

### GUI Configuration

The screenshot shows the 'Create Managed Service' window with the '2. Action' tab selected. The 'Action' dropdown is set to 'UDP Replication'. The 'Input Packet Destination IP' is set to '- IPv4 Address -'. The 'Delivery Interface' is a dropdown menu. The 'Output Packet Destination IPs' section shows '0 items'. The 'L2 Header Options' section has three radio buttons: 'Don't Add L2 Header' (selected), 'Use Original (Stripped) L2 Header', and 'Customize L2 Header'. The 'EtherType' is set to '- Select By Name -' and the 'Dst. MAC' is set to '- MAC Address -'. The 'Src. MAC' is set to '- MAC Address -'. The 'Active Timeout' is set to '10s'.

You can use the UDP-replication service to copy UDP traffic, such as syslog messages or NetFlow packets, and send the copied packets to a new destination IP address. This lets you send traffic to more destination syslog servers or NetFlow collectors than would otherwise be allowed.

Enable the checkbox for the destination for the copied output, or click the provision control (+) and add the IP address in the dialog that appears.

**Figure 5-33: Configure Output Packet Destination IP**

For the header-strip service action only, configure the policy rules for matching traffic after the header-strip service action has been applied. After completing pages 1-4, click **Append** and enable the checkbox to apply the policy.

Click **Save** to save the managed service.

### CLI Configuration

When you enter the `1 udp-replicate` command and identify the name of the configuration, the submode changes to the **config-managed-srv-udp-replicate** submode, where you can view and configure a specific UDP-replication configuration.

```
controller-1(config)# managed-service MS-UDP-REPLICATE-1
controller-1(config-managed-srv)# 1 udp-replicate DELIVERY-INTF-TO-COLLECTOR
controller-1(config-managed-srv-udp-replicate)#
```

From this submode, define the destination address of the packets to copy and the destination address where you want to send the copied packets.

```
controller-1(config-managed-srv-udp-replicate)# in-dst-ip 10.1.1.1
controller-1(config-managed-srv-udp-replicate)# out-dst-ip 10.1.2.1
```

## 5.17 Redundancy of Managed Services in Same DMF Policy

In this method, users can use a second managed service as a backup service in the same DMF policy. The backup service is activated only when the primary service becomes unavailable for any reason. The backup service can be on the same service node and/or core switch, or on a different service node and/or core switch.



**Note:** Transitioning from active to backup managed service require reprogramming of switches and associated managed appliances. This reprogramming, done seamlessly, will result in a slight loss of traffic.

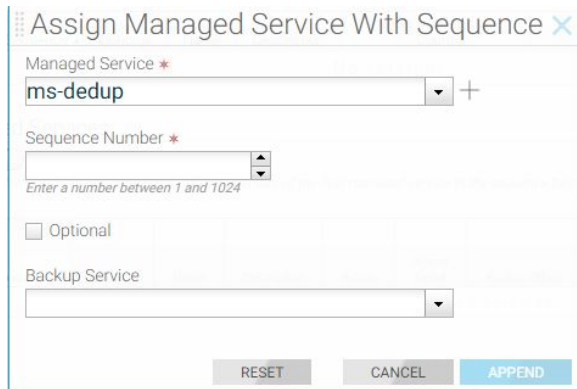
### 5.17.1 Using the GUI to Configure a Backup Managed Service

To assign a managed service as a backup service in a DMF policy, complete the following steps:

1. Select **Monitoring > Policies** and click the Provision control (+) to create a new policy.

2. Configure the policy as required. when you get to the **Services** section, click the Provision control (+) at the top of the **Managed Services** table.

**Figure 5-34: Policy with Backup Managed Service**



3. Select the primary managed service from the **Managed Service** selection list.
4. Select the backup service from the **Backup Service** selection list and click **Append**.

### 5.17.2 Using the CLI to Configure a Backup Managed Service

To implement backup managed services, complete the following steps:

1. Identify the first managed service.

```
managed-service MS-SLICE-1
1 slice l3-header-start 20
service-interface switch CORE-SWITCH-1 lag1
```

2. Identify the second managed service.

```
managed-service MS-SLICE-2
1 slice l3-header-start 20
service-interface switch CORE-SWITCH-1 lag2
```

3. Configure the policy where you refer to the backup managed service.

```
policy SLICE-PACKETS
action forward
delivery-interface TOOL-PORT-1
filter-interface TAP-PORT-1
use-managed-service MS-SLICE-1 sequence 1 backup-managed-service MS-SLICE-2
1 match ip
```

## 5.18 Application Identification

Application Identification monitors applications identified from packets taken from filter interfaces and sent through the fabric by sending IPFIX reports to a collector. Filter by forwarding or dropping packets from specific applications before sending them to the tools.



**Note:** Application identification is supported on R640 Service Nodes (DCA-DM-SC and DCA-DM-SC2).

### 5.18.1 Using the CLI to Configure Application Identification

Configure this feature through the controller in managed services.

There are two application identification services to configure:

- `app-id`
- `app-id-filter`

### Configuration Steps of `app-id-filter` and `app-id` Combined

You can follow the configuration steps described in the services above to configure `app-id-filter` and `app-id` together. However, in this case, `app-id` should use a higher **seq num** than `app-id-filter`. Thus, the traffic is processed through the `app-id-filter` policy first, then through `app-id`. This behavior can be helpful to monitor certain types of traffic.

**Note:** The two drawbacks of this configuration are `app-id` dropping the original traffic, and this type of service chaining can cause a performance hit and high memory utilization.

## 5.18.2 Using the CLI to Configure `app-id`

Perform the following steps to configure `app-id`:

1. Create a managed service. Enter the service interface.
2. Choose the `app-id` managed service with this command: `<seq num> app-id`.  
The above command should enter the `app-id` submode, which supports two configuration parameters: `collector` and `13-delivery-interface`. Both are required.
3. To configure the IP address of the IPFIX collector, enter the following command: `collector ip-address`.  
The UDP port and MTU parameters are optional: the default values are 4739 and 1500 respectively.
4. To configure the delivery interface, enter the command: `13-delivery-interface delivery interface name`.

Below is an example of `app-id` configuration that sends IPFIX application records to the collector (analytics node) at IP address 192.168.1.1 over the configured delivery interface named `app-to-analytics`:

```
managed-service ms
service-interface switch core1 ethernet2
!
1 app-id
collector 192.168.1.1
13-delivery-interface app-to-analytics
```

With this configuration in place check the analytics node for application reports and visualizations. For instance, a flow is classified internally with the following tuple: `ip, tcp, http, google, google_maps`. As a consequence, the analytics node displays the most specific app ID for this flow as `google_maps` under `appName`.

## 5.18.3 Using the CLI to Configure `app-id-filter`

Perform the following steps to configure `app-id-filter`:

1. Create a managed service. Enter the service interface.
2. Choose the `app-id` managed service with this command: `<seq num> app-id-filter`.  
The above command should enter the `app-id-filter` submode, which supports two configuration parameters: `app` and `filter-mode`. `app` is required while `filter-mode` is optional and has a default value of `forward`.
3. To configure the application name, enter the following command: `app application name`.  
Tip: Press the **tab** key after entering the `app` keyword to see all possible application names. Type in a partial name and press **tab** to see all possible choices to auto-complete the name. The entered application name must match a name in this list of app names. A service node must be connected to the controller

for this list to appear. Any number of apps can be entered one at a time with the app application name command. An example of a (partial) list of names is shown below:

**Figure 5-35: Command Example**

```
dmf-dogfood-controller-1(config-managed-srv-dpi-filter)# app goog
google          google_desktop    google_picasa    google_toolbar
google_accounts google_docs      google_play      google_translate
google_ads      google_earth     google_play_music google_trusted_store
google_analytics google_gen       google_plus      google_weblight
google_appengine google_groups    google_safebrowsing googlebot
google_cache    google_inbox     google_skymap
google_calendar google_localguides google_spaces
google_classroom google_maps      google_sprayscape
google_code     google_photos   google_tags
```

4. The **filter-mode** parameter supports two modes: forward and drop. Enter **filter-mode forward** to allow the packets to be forwarded based on the applications that have been configured. Enter **filter-mode drop** to drop these packets.

An example of app-id-filter configuration that drops all facebook and google packets:

```
! managed-service
managed-service ms
service-interface switch core1 ethernet2
!
1 app-id-filter
app facebook
app google
filter-mode drop
```



**CAUTION:** The app-id-filter configuration filters based on flows. For example, if a session is internally identified with the following tuple: ip, tcp, http, google, google\_maps, then adding any of these parameters to the filter list permits or drops all the packets matching such parameters after classification is determined (e.g., adding tcp to the filter list permits or blocks packets from the aforementioned 5-tuple flow as well as all other tcp flows). Use caution when filtering using lower layer protocols and apps.

#### 5.18.4 Using the GUI to Configure app-id and app-id-filter

The new managed service actions “App ID” and “App ID Filter” are added in the Managed Service workflow. Perform the following steps to complete the configuration.

1. Navigate to the **Monitoring > Managed Services** page. Click the table action **+** icon button to add a new managed service.

**Figure 5-36: DMF Managed Services**

**DMF Managed Services**

**Devices**

Name	IP Address	MAC Address	SKU	Hardware	Serial Number	ZeroTouch State	Interface Count	Interfaces
DMF-SN-R740-1	78:ac:44:8a:59:b8					Reloading	0	

Jun 1, 2023, 3:46:51am UTC

**Managed Services**

Name	Description	Actions	Installed	Service Interface	Service Interface Status	Policy Count
GEN-DAPPER-ACTIVE	TCP Analysis		✗	Tomahawk2-SuperSpine-es7816-1 / ethernet13/2	✓ Up	1
GEN-DAPPER-BACKUP	TCP Analysis		✗	Tomahawk2-SuperSpine-es7816-1 / ethernet19/2	✓ Up	0
GEN-NETFLOW-ACTIVE	NetFlow		✗	Tomahawk2-SuperSpine-es7816-1 / ethernet13/1	✓ Up	1
GEN-NETFLOW-BACKUP	NetFlow		✗	Tomahawk2-SuperSpine-es7816-1 / ethernet19/1	✓ Up	0

Jun 1, 2023, 3:46:51am UTC

2. In the **Info** step, configure the **Name**, **Switch**, and **Interface** inputs.

**Figure 5-37: Info**

**Create Managed Service**

1. Info

2. Actions

3. Post-Service Match

Name

Description

Switch

Optional: use to filter interface choices below

Interface

Choose from a list of interfaces not already configured for DMF

BACK NEXT RESET CANCEL SAVE

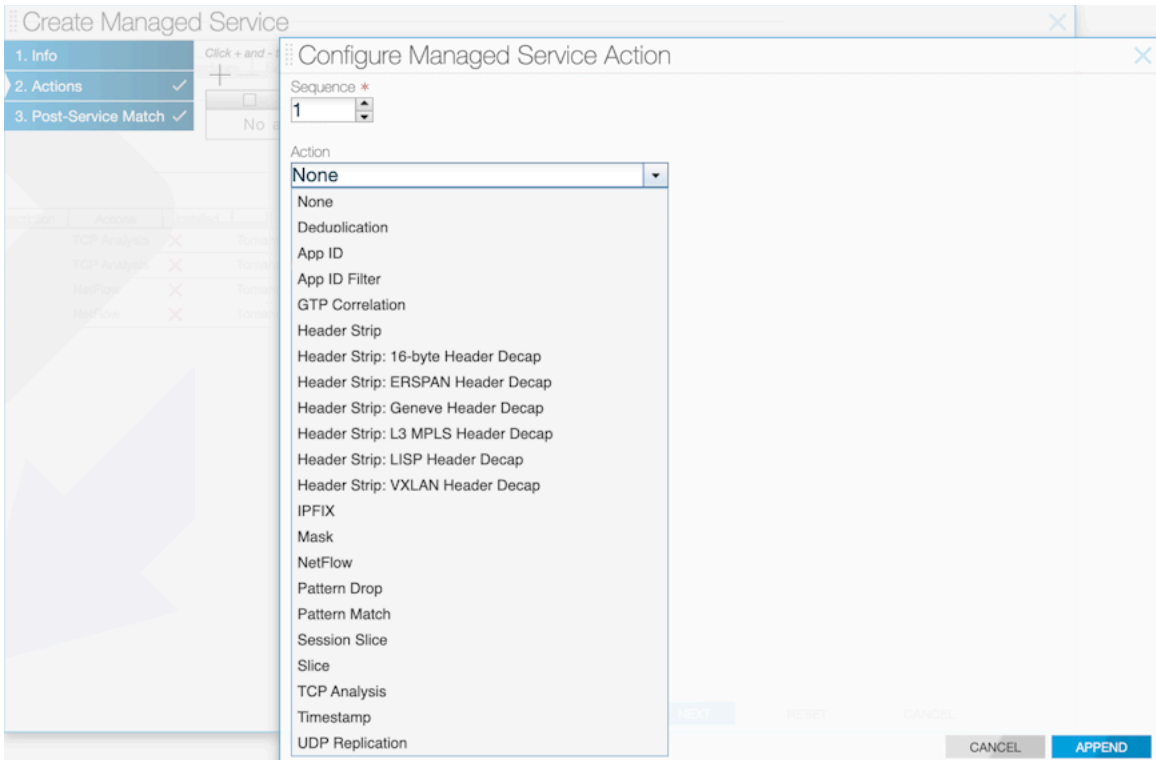
3. In the **Actions** step, click the **+** icon to add a new managed service action.

**Figure 5-38: Add App ID Action**



4. To Add the **App ID** Action, select **App ID** from the action selection input:

**Figure 5-39: Select App ID**





5. Fill in the **Delivery Interface**, **Collector IP**, **UDP Port** and **MTU** inputs and click **Append** to include the action in the managed service:

**Figure 5-40: Delivery Interface**

Configure Managed Service Action

Sequence \*  
1

Action  
App ID

Delivery Interface  
+ -

Collector IP  
- IP Address -

UDP Port \*  
4739

MTU \*  
1500

BACK NEXT RESET CANCEL APPEND

6. To Add the **App ID Filter** Action, select **App ID Filter** from the action selection input:

**Figure 5-41: Select App ID Filter**

Create Managed Service

1. Info  
2. Actions ✓  
3. Post-Service Match ✓

Click + and - to add or remove actions

Configure Managed Service Action

Sequence \*  
1

Action  
None

- None
- Deduplication
- App ID
- App ID Filter
- GTP Correlation
- Header Strip
- Header Strip: 16-byte Header Decap
- Header Strip: ERSPAN Header Decap
- Header Strip: Geneve Header Decap
- Header Strip: L3 MPLS Header Decap
- Header Strip: LISP Header Decap
- Header Strip: VXLAN Header Decap
- IPFIX

7. Select the **Filter** input as Forward or Drop action:

**Figure 5-42: Select Filter Input**

Configure Managed Service Action

Sequence \*  
1

Action  
App ID Filter

Filter  
[Dropdown]

Custom App Names  
Click + and - to include or exclude App Names.

Qosmos App Names  
Click + and - to include or exclude App Names.

+ -  
[Table with 1 row: Name, No apps]

+ -  
[Table with 1 row: Name, No apps]

When saving, the app names from the two tables will be merged together, and any duplicate app names will be removed.

BACK [DISABLED] RESET [DISABLED] CANCEL [DISABLED]

CANCEL APPEND

8. Use the **Qosmos App Names** section to add app names from Qosmos.
- a. Click the **+** button to open a modal pane to add an app name.
  - b. All the app names are listed in a table. Use the text search to filter out app names. Select the checkbox for app names to include and click **Append Selected**.
  - c. Repeat the above step to add more app names, as necessary.

**Figure 5-43: Associate App Names**

Associate App Names

google FILTER [X] [Copy]

1 2 3 ... 5 <>

Name
Filtered Results
<input checked="" type="checkbox"/> google
<input checked="" type="checkbox"/> google_accounts
<input checked="" type="checkbox"/> google_ads
<input checked="" type="checkbox"/> google_analytics
<input checked="" type="checkbox"/> google_api
<input type="checkbox"/> google_appengine
<input type="checkbox"/> google_blog
<input checked="" type="checkbox"/> google_books
<input type="checkbox"/> google_cache
<input checked="" type="checkbox"/> google_calendar

When saving, the app names from the two tables will be merged together, and any duplicate app names will be removed.

Show: 10 25 (1 - 10 / 44) 1 2 3 ... 5 <>

CANCEL APPEND SELECTED

9. The selected app names are now listed. Use the - icon button to remove any app names, if necessary:

**Figure 5-44: Application Names**

Configure Managed Service Action

Sequence \*  
1

Action  
App ID Filter

Filter

Custom App Names  
Click + and - to include or exclude App Names.

Qosmos App Names  
Click + and - to include or exclude App Names.

When saving, the app names from the two tables will be merged together, and any duplicate app names will be removed.

BACK CANCEL APPEND

If the required app names are not present in the **Qosmos App Names**, add a custom app name under the **Custom App Names** section. This name has to exist in the namespace of Qosmos app names above or will not be filtered.

10. Click the + button to open a modal pane to add an app name. Enter a custom app name and click **Append**.

**Figure 5-45: Configure Custom Name**

Configure Custom App Name

Name  
my\_app

CANCEL APPEND

11. The custom app names are now listed. Use the - icon button to remove any app names, if necessary.

**Figure 5-46: Custom App Names**

Configure Managed Service Action

Sequence \*  
1

Action  
App ID Filter

Filter

Custom App Names  
Click + and - to include or exclude App Names.

	Name
<input type="checkbox"/>	my_app

Show: 10 25 100 All

Qosmos App Names  
Click + and - to include or exclude App Names.

	Name
<input type="checkbox"/>	google
<input type="checkbox"/>	google_accounts
<input type="checkbox"/>	google_ads
<input type="checkbox"/>	google_analytics
<input type="checkbox"/>	google_api
<input type="checkbox"/>	google_books
<input type="checkbox"/>	google_calendar

(1 - 7 / 7)

When saving, the app names from the two tables will be merged together, and any duplicate app names will be removed.

BACK CANCEL APPEND

12. Click the **Append** button at the bottom to add the action to the managed service and click **Save** to save the managed service.

## 5.18.5 Dynamic Signature Updates (Beta Version)

This beta feature allows the app-id and app-id-filter services to classify newly supported applications at runtime rather than waiting for an update in the next DMF release. Perform such runtime service updates during a maintenance cycle. There can be issues with backward compatibility if attempting to revert to a too-old bundle. You should adopt only supported versions. In the controller's CLI, perform the following steps as recommended:

1. Remove all policies containing app-id or app-id-filter. Remove the app-id and app-id-filter managed services from the policies using the command: **no use-managed-service** in policy config.

Arista Networks recommends this step to avoid errors and service node reboots during the update process. A warning message is printed right before confirming a push. Proceeding without this step may work but is not recommended as there is a risk of service node reboots.



**Note:** In the command displayed below, the update file is a specific file provided by Arista Networks.

2. To pull the signature file onto the controller node, use the command:

```
C1(config)# app-id pull-signature-file user@host:path to file.tar.gz
Password:
file.tar.gz          5.47MB 1.63MBps 00:03
```

3. Fetch and validate the file using the command:

```
C1(config)# app-id fetch-signature-file file://file.tar.gz
Fetch successful.
Checksum      : abcdefgh12345
Fetch time    : 2023-08-02 22:20:49.422000 UTC
Filename      : file.tar.gz
```

4. To see which files are currently saved on the controller node after the fetch operation is successful, use the following command:

```
C1(config)# app-id list-signature-files
# Signature-file      Checksum      Fetch time
-|-----|-----|-----|
1 file.tar.gz      abcdefgh12345    2023-08-02 22:20:49.422000 UTC
```



**Note:** Only the files listed by this command can be pushed to service nodes.

5. Push the file from the controller to the service nodes, use the following command:

```
C1(config)# app-id push-signature-file file.tar.gz
App ID update: WARNING: This push will affect all service nodes
App ID update: Remove policies configured with app-id or app-id-filter
before continuing to avoid errors
App ID update: Signature file: file.tar.gz
App ID update: Push app ID signatures to all Service Nodes? Update ("y" or
"yes" to continue): yes
Push successful.

Checksum      : abcdefgh12345
Fetch time    : 2023-08-02 22:20:49.422000 UTC
Filename      : file.tar.gz
Sn push time  : 2023-08-02 22:21:49.422000 UTC
```

6. Add the **app-id** and **app-id-filter** managed services back to the policies.  
As a result of adding back app-id, service nodes are now able to identify and report new applications to the analytics node.  
After adding back app-id-filter, new application names should appear in the app-id-filter controller app list. To test this, enter **app-id-filter** submode and press **tab** to see the full list of applications. New identified applications should appear in this list.
7. To delete a signature file from the controller, use the command below.



**Note:** You are only allowed to delete a signature file that is not actively in use by any service node, which needs to keep a working file in case of issues—attempting to delete an active file causes the command to fail.

```
C1(config)# app-id delete-signature-file file.tar.gz
Delete successful for file: file.tar.gz
```

### Useful Information

The fetch and delete operations are synced with standby controllers as follows:

- **fetch:** when a fetch is successful on the active controller, it invokes the fetch rpc on the standby controller by providing a signed HTTP URL as the source. This URL points to an internal REST API that provides the recently fetched signature file.
- **delete:** the active controller invokes the delete rpc call on the standby controllers.

The controller stores the signature files in this location: `/var/lib/capture/appidsignatureupdate`.

On a service node, files are overwritten and always contain the full set of applications.



**Note:** An analytics node cannot display these applications in the current version.

This step is only for informational purposes:

- The bundle version on the service node may be checked by entering the **show service-node app-id-bundle-version** command in the service node CLI, as shown below.

#### Figure 5-47: Before Update

```
[SD-appid84x-Hormigueros-SN1(config)# show service-node app-id-bundle-version
Name : bundle_version
Data : 1.640.2-24 (build date Mar  2 2023)
```

#### Figure 5-47: After Update

```
[SD-appid84x-Hormigueros-SN1(config)# show service-node app-id-bundle-version
Name : bundle_version
Data : 1.650.2-22 (build date Apr 27 2023)
```

## 5.18.6 CLI Show Commands

In the service node CLI use the following **show** command:

```
> show service-node app-id-bundle-version
```

This command shows the version of the bundle being used. An instance of app-id or app-id-filter must be configured or an error message is displayed.

```
dmf-service-node-1> show service-node app-id-bundle-version
Name : bundle_version
data : 1.640.2-24 (build date Mar 2 2023)
```

## 5.18.7 Syslog Messages

Syslog messages for configuring the app-id and app-id-filter services are included in a service node's syslog through journalctl.

A service node syslog registers events for the app-id add/modify/delete actions.

These events contain the keywords dpi and dpi-filter which correspond to app-id and app-id-filter.

For example:

Adding dpi for port, Modifying dpi for port, Deleting dpi for port,

Adding dpi filter for port, Modifying dpi filter for port, Deleting dpi filter for port,

App *appname* does not exist - An invalid app name was entered.

Addition/modification/deletion of app names in an app-id-filter managed-service in the controller node's CLI influences the policy refresh activity and these events are registered in `floodlight.log`.

## 5.18.8 Scale

Max concurrent sessions are currently set to permit less than 200k active contexts per core. Surpassing this threshold may cause some flows not to be processed. The new flows will not be identified or filtered if this is surpassed. Entries for inactive flows time out over a few minutes and a few seconds after the session ends.

Heavy application traffic load degrades performance.

## 5.18.9 Troubleshooting and Considerations

### Troubleshooting

- If IPFIX reports do not appear on an analytics node or collector, ensure the UDP port is correctly configured and check if the analytics node is receiving traffic.
- If the app-id-filter's app list does not appear, ensure a service node is connected using 'show service-node' on the controller.
- Be aware that a flow may contain other ids and protocols when using app-id-filter. For example, the specific application for a flow may be google\_maps, but there may be protocols or broader applications under it, such as SSH, HTTP, or Google. Adding google\_maps will filter this flow. However, adding SSH will also filter this flow. Therefore, adding any of these to the filter list will cause packets of this flow to be forwarded/dropped.
- During a dynamic signature update, if service node reboots occur, it is most likely because some policies are still configured with app-id or app-id-filter at the time of the update. If these are not removed, they may cause service nodes to reboot. To avoid this problem, remove those policies before an update and perform an update during a maintenance window.

### Considerations

- If using a drop filter, a small amount of packets may slip through the filter before determining an application ID for a flow. When using a forward filter, a small amount of packets may not be forwarded. A small amount is estimated to be between 1 and 6 packets at the beginning of a flow.
- If using a drop filter, add the 'unknown' app ID to the filter list to drop any traffic that was unable to be identified if these packets are unwanted.
- The controller must be connected to a service node for the app-id-filter's app list to appear. If the list does not appear and you don't remember the application names, use the app-id to send reports to the analytics node, and the application names seen there can be used to configure an app-id-filter. The name must match exactly.
- For app-id, a report is only generated for a flow once that flow has been fully classified. Therefore, the number of reported applications may not match the total number of flows. These reports are sent once enough applications are identified on the service node. If many applications are identified, the reports will be sent quickly. However, if only a few applications are identified, these reports may be sent every 10 seconds.
- For app-id, there may be some instances where multiple reports are generated for the same 5-tuple.
- The app-id and app-id-filter services are more resource-intensive than other services. Combining them in a service chain or configuring many instances of them may lead to degradation in performance.

## 5.19 Redundancy of Managed Services Using Two DMF Policies

In this method, users can use a second policy with a second managed service to provide redundancy. The idea here is to duplicate the policies, but assign a lower policy priority to the second DMF policy. In this case, the backup policy (and by extension, the backup service) will always be active but will not receive any relevant traffic until the primary policy goes down. This method provides true redundancy at the policy, service-node, and core switch levels, but it uses additional network and node resources.

An example is provided below:

```
! managed-service
managed-service MS-SLICE-1
1 slice l3-header-start 20
service-interface switch CORE-SWITCH-1 lag1
!
managed-service MS-SLICE-2
1 slice l3-header-start 20
```

```

service-interface switch CORE-SWITCH-1 lag2
! policy
policy ACTIVE-POLICY
priority 101
action forward
delivery-interface TOOL-PORT-1
filter-interface TAP-PORT-1
use-managed-service MS-SLICE-1 sequence 1
1 match ip
!
policy BACKUP-POLICY
priority 100
action forward
delivery-interface TOOL-PORT-1
filter-interface TAP-PORT-1
use-managed-service MS-SLICE-2 sequence 1
1 match ip

```

## 5.20 Cloud Services Filtering

The DANZ Monitoring Fabric supports the filtering of traffic to specific services hosted in the public cloud. Filtered traffic is redirected to customer tools. This is achieved by reading the source and destination IP address of specific flows, identifying the Autonomous System number they belong to, tagging the flows with their respective AS numbers, and redirecting the flows to customer tools for consumption.

The following is the list of services supported:

- amazon: traffic with src/dst IP belonging to Amazon
- ebay: traffic with src/dst IP belonging to eBay
- facebook: traffic with src/dst IP belonging to FaceBook
- google: traffic with src/dst IP belonging to Google
- microsoft: traffic with src/dst IP belonging to Microsoft
- netflix: traffic with src/dst IP belonging to Netflix
- office365: traffic for Microsoft Office365
- sharepoint: traffic for Microsoft Sharepoint
- skype: traffic for Microsoft Skype
- twitter: traffic with src/dst IP belonging to Twitter
- default :traffic not matching other rules in this service. Supported types are **match** or **drop**.

The option **drop** instructs the DMF Service Node to drop packets matching the configured application.

The option **match** instructs the DMF Service Node to deliver packets to the delivery interfaces connected to the customer tool.

A default drop action is auto-applied as the last rule, except when the last rule is configured as **match default**. It instructs the DMF Service Node to drop packets if either of the following conditions are met:

- Source IP address or destination IP address of stream doesn't belong to any AS number.
- Source IP address or destination IP address of stream is associated with an AS number, but has no specific action set.

### 5.20.1 Cloud Services Filtering Configuration

#### Managed Service Configuration

```

Controller(config)# managed-service <name>
Controller(config-managed-srv)#

```



## Service Action Configuration

```
Controller(config-managed-srv)# 1 app-filter
Controller(config-managed-srv-appfilter)#
```

## Filter Rules Configuration

```
Controller(config-managed-srv-appfilter)# 1 drop sharepoint
Controller(config-managed-srv-appfilter)# 2 match google
Controller(config-managed-srv-appfilter)# show this
! managed-service
managed-service sf3
service-interface switch CORE-SWITCH-1 ethernet13/1
!
1 service- app-filter
1 drop sharepoint
2 match google
```

A policy having a managed service with app-filter as the managed service, but with no matches specified will fail to install. The example below shows a policy incomplete-policy having failed due to the absence of a Match/Drop rule in the managed service incomplete-managed-service.

```
Controller(config)# show running-config managed-service incomplete-managed-
service
! managed-service
managed-service incomplete-managed-service
1 app-filter
Controller(config)# show running-config policy R730-sf3
! policy
policy incomplete-policy
action forward
delivery-interface TOOL-PORT-1
filter-interface TAP-PORT-1
use-managed-service incomplete-managed-service sequence 1
1 match any
```

```
Controller(config-managed-srv-appfilter)# show policy incomplete-policy
Policy Name : incomplete-policy
Config Status : active - forward
Runtime Status : one or more required service down
Detailed Status : one or more required service down - installed to
forward
Priority : 100
Overlap Priority : 0
```

## 5.21 Multiple Services Per Service Node Interface

Starting with the release of *DANZ Monitoring Fabric 8.1.0*, the service-node capability is augmented to support more than one service action per service-node interface. Though this feature is economical in terms of per-interface cost, it could cause packets to be dropped in a high-volume traffic environment. Use this feature judiciously.

An example is shown below:

```
controller-1# show running-config managed-service Test
! managed-service
managed-service Test
service-interface switch CORE-SWITCH-1 ethernet13/1
```

```

1 dedup full-packet window 2
2 mask BIGSWITCH
3 slice l4-payload-start 0
!
4 netflow an-collector
collector 10.106.6.15 udp-port 2055 mtu 1500

```

This feature replaces the **service-action** command with sequential numbers. The allowed range of sequence numbers is **1 -20000**. In the above example, the sequence numbering impacts the order in which the managed services influence the traffic.



**Note:** On upgrade to **DANZ Monitoring Fabric release 8.1.0** and later, the service-action CLI will be automatically replaced with sequence number(s).

Specific managed service statistics can be viewed via the following CLI command:

```

controller-lf show managed-service-device R740 stats Test

```

Stats															
#	Service	Node	Interface	Name	Function	Service Name	Rx packets	Rx bytes	Rx Bit Rate	Applied packets	Applied bytes	Applied Bit Rate	Tx packets	Tx bytes	Tx Bit Rate
1			sn19	dedup	Test		12984912	10273331330	-	12984912	10273331330	-	12984912	10273331330	-
2			sn19	netflow	Test		12984704	597296384	-	12984704	597296384	-	74	101180	-
3			sn19	mask	Test		12984880	10273305423	-	11728176	10104821837	-	12984880	10273305423	-
4			sn19	slice	Test		12984944	10273357891	-	12984944	10273357891	-	12984944	597307424	-

The above information can also be viewed in the DMF GUI at **Monitoring > Managed Services > Devices > Service Stats**.

Service Stats

Interface Name	Service Name	Table Name	Load	RX Packet Count	RX Byte Count	RX Bit Rate	TX Packet Count	TX Byte Count	TX Bit Rate	Applied Packets	Applied Bytes	Applied Bit Rate
sn19	Test	dedup	No	12,984,912	10,273,331,330	0 bit/s	12,984,912	10,273,331,330	0 bit/s	12,984,912	10,273,331,330	0 bit/s
sn19	Test	netflow	No	12,984,704	597,296,384	0 bit/s	74	101,180	0 bit/s	12,984,704	597,296,384	0 bit/s
sn19	Test	mask	No	12,984,880	10,273,305,423	0 bit/s	12,984,880	10,273,305,423	0 bit/s	11,728,176	10,104,821,837	0 bit/s
sn19	Test	slice	No	12,984,944	10,273,357,891	0 bit/s	12,984,944	597,307,424	0 bit/s	12,984,944	10,273,357,891	0 bit/s



**Note:** Following limitations apply for this mode of configuration:

- The NetFlow/IPFIX-action configuration should not be followed by the timestamp service action.
- The UDP-replication action configuration should be the last service in the sequence.
- The header-stripping service with post-service-match rule configured should not be followed by the NetFlow, IPFIX, udp-replication, timestamp and TCP-analysis services.

## Using the DMF Recorder Node

---

This chapter describes how to configure the DMF Recorder Node to record packets from DMF filter interfaces. For related information, refer to the following:

- Installing and configuring DMF Recorder Node: ***DANZ Monitoring Fabric Deployment Guide***
- Integrating DMF Recorder Node with Analytics: ***Arista Analytics User Guide***
- Stenographer Queries: [Stenographer Reference for DMF Recorder Node](#)
- DMF Recorder Node REST API: [DMF Recorder Node REST APIs](#)

### 6.1 Overview

The DMF Recorder Node is integrated with the DANZ Monitoring Fabric for single-pane-of-glass monitoring. A single DMF controller can manage multiple Recorder Nodes, delivering packets for recording through Out-of-Band policies. The DMF controller also provides central APIs for packet queries across one or multiple recorder nodes and for viewing errors, warnings, statistics, and the status of connected recorder nodes.

A DMF out-of-band policy directs matching packets to be recorded to one or more recorder nodes. A recorder node interface identifies the switch and port used to attach the recorder node to the fabric. A DMF policy treats these as delivery interfaces and adds them to the policy so that flows matching the policy are delivered to the specified recorder node interfaces.

### 6.2 Configuration Summary

At a high level, three steps are required for using the recorder node.

**Step 1:** Define a recorder node.

**Step 2:** Define a DMF policy to select the traffic to forward to the recorder node.

**Step 3:** View and analyze the recorded traffic.

The recorder node configuration on the DMF controller includes the following:

- **Name:** Each recorder node requires a name that is unique among recorder nodes in the connected fabric. If the name is removed, all configuration for the given recorder node is removed.
- **Management MAC address:** Each recorder node must have a management MAC address that is unique in the connected fabric.
- **Packet removal policy:** This defines the behavior when the recorder node disks reach capacity. The default policy causes the earliest recorded packets to be overwritten by the most recent packets. The other option is to stop recording and wait until space is available.
- **Record enable or Record disable:** Recording of packets is enabled by default, but it can be enabled or disabled for a specific recorder node.
- **Static auth tokens:** Static auth tokens are pushed to each recorder node as an alternative form of authentication in headless mode, when the DMF controller is unreachable, or by third-party applications that do not have or do not need DMF controller credentials.
- **Controller auth token:** The recorder node treats the controller as an ordinary client and requires it to present valid credentials in the form of an authentication token. The DMF controller authentication token is automatically generated but can be reset upon request.

- **Pre-buffer:** This buffer, which is defined in minutes, is used for proactive network monitoring without recording and retaining unnecessary packets. Once the buffer is full, the oldest packets are deleted.
- **Maximum disk utilization:** This defines the maximum disk utilization in terms of a percentage between 5% and 95%. When the configured utilization is reached, the packet removal policy is enforced. The default maximum disk utilization is 95%.
- **Maximum packet age:** This defines the maximum age in minutes of any packet in the recorder node. It can be used in combination with the packet removal policy to control when packets are deleted based on age rather than disk utilization alone. When not set, the maximum packet age is not enforced and packets are kept until the maximum disk utilization is reached.

## 6.3 Indexing Configuration

The recorder node indexing configuration defines the fields that can be used to query packets on the recorder node. By default, all indexing fields are enabled in the indexing configuration. You can selectively disable indexing fields you do not wish to use in recorder node queries.

Disabling indexing fields has two advantages. First, it reduces the index space required for each packet recorded. Second, it improves query performance by reducing unnecessary overhead. It is recommended that unnecessary indexing fields be disabled.

The recorder node supports the following indexing fields:

- MAC Source
- MAC Destination
- VLAN 1: Outer VLAN ID
- VLAN 2: Inner/Middle VLAN ID
- VLAN 3: Innermost VLAN ID
- IPv4 Source
- IPv4 Destination
- IPv6 Source
- IPv6 Destination
- IP protocol
- Port Source
- Port Destination
- MPLS
- Community ID
- MetaWatch Device ID
- MetaWatch Port ID



**Note:** The Outer VLAN ID indexing field must be enabled in order to query the recorder node using a DMF policy name or a DMF filter interface name.

To understand how indexing configuration can be leveraged to your advantage, consider the following examples:

**Example 1:** To query packets based on applications defined by unique transport ports, disable all indexing fields except source and destination transport ports. This results in only transport ports being saved as meta data for each packet recorded. This greatly reduces per-packet index space consumption and also increases the speed of recorder-node queries.

However, you will not be able to effectively query on any other indexing field because that meta data was not saved when the packets were recorded.

**Example 2:** The recorder node supports community ID indexing, which is a hash of IP addresses, IP protocol, and transport ports that can be used to identify a flow of interest. If the recorder node use case is to query based on community ID, it might be redundant to index on IPv4 source and destination addresses, IPv6 source and destination addresses, IP protocol, and transport port source and destination addresses.

## 6.4 Pre-buffer Configuration and Events

The recorder node pre-buffer is a circular buffer in which packets to be recorded are received. When enabled, the pre-buffer feature allows for retention of the packets received by the recorder node for a specified length of time prior to an event that triggers recording of buffered and future packets to disk. In the absence of an event, the recorder node will record into this buffer, deleting the oldest packets in the buffer when the buffer reaches capacity. When a recorder node event is triggered, the packets in the pre-buffer are saved to disk, and the packets received from the time of the event trigger to the time of the event termination are saved directly to disk upon termination of the event, received packets are received and retained in the pre-buffer until the next event. By default, the pre-buffer feature is disabled, indicated by a value of zero minutes.

For example, if you configure the pre-buffer to thirty minutes, up to thirty minutes of packets will be received by the buffer. When you trigger an event, the packets currently in the buffer are recorded to disk, and packets newly received by the recorder node bypass the buffer and are written directly to disk until the event is terminated. When you terminate the event, the pre-buffer resets, accumulating received packets for up to the defined thirty-minute pre-buffer size.

The packets associated with an event can be queried, replayed, or analyzed using any type of recorder node query. Each triggered event is identified by a unique, user-supplied name, which can be used in the query to reference packets recorded in the pre-buffer prior to and during the event itself.

## 6.5 Using an Authentication Token

When using a DMF controller authentication token, the recorder node treats the DMF controller as an ordinary client and requires it to present valid credentials either in the form of an HTTP basic username and password or an authentication token.

Static authentication tokens are pushed to each recorder node as an alternative form of authentication in headless mode, when the DMF controller is unreachable, or by third-party applications that do not have or do not need controller credentials.

## 6.6 Using the GUI to Add a Recorder Device

To configure a recorder node or update the configuration of an existing recorder node, complete the following steps:

1. Select **Monitoring > Recorder Nodes** from the main menu bar of the DMF GUI.

The system displays the page shown below.

**Figure 6-1: Recorder Nodes**

The screenshot displays the 'Recorder Nodes' page within the DANZ Monitoring Fabric 8.0.0 interface. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. A status indicator shows 'Fabric: Healthy'.

**Recorder Nodes**

**Devices**

Filter table rows

Name	MAC	IP Address	SKU	Hardware	Serial Number	Connected	ZeroTouch State	Recording	Disk Full Policy	Max Packet Age	Pre Buffer	Max Disk Utilization	Parse MetaWatch Trailer	Indexing	Alerts
mt	92:14:00:eb:24:fb	10.243.255.4	--	--	Not Specified	✓	Ok	Off	--	--	--	--	--	X MAC Source X MAC Destination X VLAN 1 X VLAN 2 X VLAN 3 <a href="#">and 11 More...</a>	--

Oct 13, 2023, 3:08:10am UTC

Show: 10 25 100 All

**Interfaces**

Filter table rows

Name	Description	Switch Name	Interface Name
No recorder node interfaces			

**Global Query Configuration**

Maximum Size of Query Result 256.00 MB

**Query History**

Filter table rows

Recorder Node	Stenographer Query	Query Type	Start Time	Query Wait Duration	Query Running Duration	Total Duration	Error Message
No query histories							

**Active Queries**

Stenographer Query Query Type Invoked via Recorder Nodes Status Percent Complete Current Value Max Value

No active queries

**Query**

Submit Reset Clear Inputs Clear Results

Query Type Window

Fetch oldest and latest timestamp

© Arista Networks, Inc. DANZ Monitoring Fabric 8.0.0

- To add a new recorder node, click the provision control (+) at the top of the **Recorder Nodes Devices** table.

**Figure 6-2: Provision Recorder Node**

**Provision Recorder Node**

**1. Info** ✓

Name \*

**2. Indexing** ✓

**3. Network** ✓

MAC Address

- MAC Address -

Drop-down includes connected switches without a fabric role and addresses from failed ZTN requests. Choose from the drop-down or enter a new value expected to connect in the future. When a switch with the entered MAC connects, this configuration will be applied to it.

Interface Name

**4. Storage** ✓

Recording

Off ☒ On

Disk Full Policy

Rolling FIFO ☒ Stop and Wait ?

Pre Buffer

minutes

Duration to record into a pre-buffer until an event occurs

Parse MetaWatch Trailer \* ☒ Error Message

Off

Max Packet Age

minutes

The maximum age of a recorded packet in minutes. Packets older than this age will be deleted automatically.

Max Disk Utilization

%

The maximum allowed utilization on the index and packet disks. The packet removal policy will be enforced at this limit.

BACK NEXT RESET CANCEL SAVE

- Complete the following required fields:
  - Assign a name to the recorder node.
  - Set the MAC address of the recorder node. Obtain the MAC address from the chassis ID of the connected device, using the **Fabric > Connected Devices** option.
- Configure the following options as needed:
  - Recording:** Recording is enabled by default. To disable recording on the recorder node, move the **Recording** slide to **Off**. When recording is enabled, the recorder node records the matching traffic directed from the filter interface defined in a DMF policy.
  - Disk Full Policy:** Change the **Disk Full Policy** to **Stop and Wait** if required. The default packet removal policy is **Rolling FIFO** (First In First Out), which means the oldest packets will be deleted to make room for newer packets. This occurs only when the recorder node disks are full. The alternative removal policy is **Stop and Wait**, which causes the recorder node to stop recording when the disks are full and wait until disk space becomes available. Disk space can be made available by leveraging the recorder node delete operation to remove all or selected time ranges of recorded packets.
  - Max Packet Age:** Change the **Max Packet Age** to set the maximum number of minutes that recorded packets will be kept on the recorder node. Packets recorded are discarded after the specified number of minutes. This defines the maximum age in minutes of any packet in the recorder node. It can be used in combination with the **Disk Full Policy** to control when packets are deleted based on age rather than disk utilization alone. When unset, **Max Packet Age** is not enforced.
  - Pre-Buffer:** Assign the number of minutes the recorder node pre-buffer allows for windowed retention of packets received by the recorder node for a specified length of time. By default, the **Pre-Buffer** is set to zero minutes (disabled). With a nonzero **Pre-Buffer** setting, when you trigger a recorder event, any packets in the pre-buffer are saved to disk, and any packets received by the recorder after the trigger are saved directly to disk. When you terminate an ongoing recorder event, a new pre-buffer is established in preparation for the next event.
  - Max Disk Utilization:** Specify the maximum utilization allowed on the index and packet disks. The **Disk Full Policy** will be enforced at this limit. If left unset, then the disks space will be used to capacity.

- **Parse MetaWatch Trailer:** Determine when the MetaWatch trailer should be parsed.
  - **Off:** When set to **Off**, the recorder node will not parse the MetaWatch trailer, even if it is present in incoming packets.
  - **Auto:** When set to **Auto**, the recorder node will look for a valid timestamp in the last 12 bytes of the packet. If it matches the system timestamp closely enough, the trailer will be parsed by the recorder node.
  - **Force:** When set to **Force**, recorder node will assume the last 12 bytes of packet is a MetaWatch trailer and parse it, even if it did not find a valid timestamp.
- 5. Configure either local or external storage for the node using the appropriate procedure from below.

## 6.6.1 Configuring a Node to Use Local Storage

1. To configure a node to use local storage, use the following steps:
1. **Network:** To use local storage, set the **Auxiliary NIC Configuration** to default (**No**) as shown in the figure below.

**Figure 6-3: Network Provisioning**

The screenshot shows the 'Provision Recorder Node' window with the 'Network' step selected. The 'Auxiliary NIC Configuration' is set to 'No'. The window includes a sidebar with steps 1-4, a main content area with a description and a toggle switch, and a bottom bar with buttons for BACK, NEXT, RESET, CANCEL, and SAVE.

2. **Storage:** To use local storage, set the **Index Disk Configuration** and **Packet Disk Configuration** to default (**No**) as shown in the figure below.

**Figure 6-4: Configure to Use Local Storage**

The screenshot shows the 'Provision Recorder Node' window with the 'Storage' step selected. Both the 'Index Disk Configuration' and 'Packet Disk Configuration' are set to 'No'. The window includes a sidebar with steps 1-4, a main content area with descriptions and toggle switches for both configurations, and a bottom bar with buttons for BACK, NEXT, RESET, CANCEL, and SAVE.


3. Click **Save** to add the recorder node configuration to the controller.



## 6.6.2 Configuring a Node to Use External Storage

To store packets on external storage with an NFS mount, the auxiliary interface of the recorder node has to be connected to the network and subnet where the NFS storage is located as displayed in the figure below.

**Figure 6-5: Topology to Use External Storage**



**Note:** Volume for index and packet on the NFS storage should be created first. Refer to vendor-specific NFS storage documentation about how to create the volume (or path).

To configure a recorder node for external NFS storage, update the configuration of an existing recorder node or add a new node with the following steps:



**Note:** For **release 7.2**, only Isilon NFS storage is supported.

1. **Network:** For external NFS storage, such as Isilon, the auxiliary interface of the recorder node should be connected to a network and subnet which has reachability to Isilon NFS storage. Set the **Auxiliary NIC Configuration** slide to **YES** and assign an IP address to the auxiliary interface as shown in the figure

below. Ensure the IP address for the auxiliary interface is not in the same subnet as the recorder node management IP address.

**Figure 6-6: Provision External Storage**

Provision Recorder Node

1. Info ✓

2. Indexing ✓

3. Network ✓

4. Storage ✓

Auxiliary NIC Configuration

No ☐ Yes ☐

IP Address \*

10.111.10.10 / 24

Interface Name

256.00 MB

BACK NEXT RESET CANCEL SAVE

2. Storage: To specify the location of the external NFS storage, configure the following options:
  - **Index Disk Configuration** and **Packet Disk Configuration** are disabled by default (slide set to **NO**). Set the slide for both **Index** and **Packet Disk Configuration** to **YES**.
  - **NFS Server [Index and Packet Disk Configuration]**: assign the IP address or host name for the NFS Server (e.g., Isilon Smart Connect host name).
  - **Transport Port of NFS Service [Index and Packet Disk Configuration]**: if no value is specified, default will be used (2049). Specify a value for this if the NFS storage has been configured to use something other than the default.
  - **Transport Port of Mounted Service [Index and Packet Disk Configuration]**: if no value is specified, the default will be used. Specify a value for this if the NFS storage mounted service has been configured to use something other than default.
  - **Index Disk Configuration**- Volume and path on the NFS server where the index will be stored.

- **Packet Disk Configuration**- Volume and path on the NFS server where packets will be stored.

**Figure 6-7: Provision External Storage**

**Provision Recorder Node**

1. Info ✓ Configure how and where we store packets/metadata

2. Indexing ✓ Index Disk

3. Network ✓ Index Disk Configuration

4. Storage ✓

Index Disk Configuration

No ☐ Yes ☒

NFS Server \* bsn-isilon-pool ?

Transport Port of NFS Service 1 - 65535

Transport Port of mountd Service 1 - 65535  
The mountd service is assumed to run on the same host as the NFS server.

Volume \* /index\_volume

Packet Disk

Packet Disk Configuration

No ☐ Yes ☒

NFS Server \* bsn-isilon-pool ?

Transport Port of NFS Service 1 - 65535

Transport Port of mountd Service 1 - 65535  
The mountd service is assumed to run on the same host as the NFS server.

Volume \* /packet\_volume

BACK NEXT RESET CANCEL SAVE

3. Click **Save** to add the recorder node configuration to the controller.



**Note:** If the configuration of a previously added packet recorder is edited to use external storage from local storage or vice versa, then the packet recorder must be rebooted.

## 6.7 Configuring a Recorder Node Interface

To record packets to a recorder node using a DMF policy, configure a DMF Recorder Node interface that defines the switch and interface in the monitoring fabric where the recorder node is connected. The DMF Recorder Node interface is referenced by name in the DMF policy as the destination for traffic matched by the policy. To configure a DMF Recorder Node interface, complete the following steps:

1. Click the provision control (+) at the top of the **Recorder Node Interfaces** table. The system displays the following page:

**Figure 6-8: Create DMF Recorder Node Interface**

Create Recorder Node InterfaceX

Name \*

Switch

Optional: use to filter interface choices below

Interface

- Interface -

Choose from a list of interfaces not already configured for DMF

Description

RESET CANCEL SAVE


2. Assign a name for the DMF Recorder Node interface in the **Name** field.
3. Select the switch containing the interface that connects the recorder node to the monitoring fabric.
4. Select the interface that connects the recorder node to the monitoring fabric.
5. (Optional) Type information about the interface in the **Description** field.
6. Click **Save** to add the configuration to the DMF controller.

## 6.8 Using the GUI to Assign a Recorder Interface to a Policy

To forward traffic to a recorder node, include one or more recorder node interfaces as a delivery interface in a DMF policy.

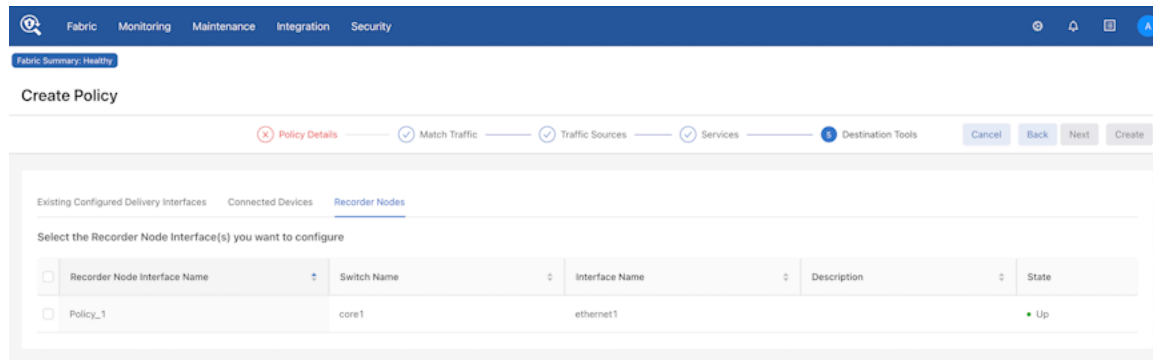
When you create a new policy or edit an existing policy, you can select the recorder node interfaces from the **Monitoring > Policies** dialog, as shown in the following screen.

**Figure 6-9: DMF Policies**



To create a Recorder Node interface, proceed to the **Monitoring > Recorder Nodes** page and click the **+** in the Interface section. While creating a policy, from the **Destination Tools > Recorder Nodes** page, select the Recorder Node interface previously configured following the procedure mentioned above.

**Figure 6-10: Recorder Node - Create Policy**



**Note:** The Recorder Node interface can only be selected and not created in the create policy dialogue.

## 6.9 Using the GUI to Define a Recorder Query

The recorder node records all the packets received on a filter interface that match the criteria defined in a DMF policy. Recorded packets can be recalled from or analyzed on the recorder node using a variety of queries. Use the options in the recorder node **Query** section to create a query and submit it to the recorder node for processing. The following queries are supported:

- **Window:** Retrieves the timestamps of the oldest and most recent packets recorded on the recorder.
- **Size:** Provides the number of packets and their aggregate size in bytes that match the filter criteria specified.
- **Application:** Performs deep packet inspection to identify applications communicating with the packets recorded and that match the filter criteria specified.
- **Packet-data:** Retrieves all the packets that match the filter criteria specified.
- **Packet-object:** The packet object query extracts unencrypted HTTP objects from packets matching the given stenographer filter.
- **HTTP, HTTP Request, and HTTP Stat:** Analyzes HTTP packets, extracting request URLs, response codes, and statistics.
- **DNS:** Analyzes any DNS packets, extracting query and response meta data.
- **Replay:** Replays selected packets and transmits them to the specified delivery interface.
- **IPv4:** Identifies and dissects distinct IPv4 flows.
- **IPv6:** Identifies and dissects distinct IPv6 flows.
- **TCP:** Identifies and dissects distinct TCP flows.
- **TCP Flow Health:** Analyzes TCP flows for information such as maximum RTT, retransmissions, throughput, etc.
- **UDP:** Identifies and dissects distinct UDP flows.
- **Hosts:** Identifies all the unique hosts that match the filter criteria specified.
- **RTP Stream:** Characterizes the performance of Real Time Protocol streaming packets.

After making a selection from the **Query Type** list, the system displays additional fields that can be used to filter the retrieved results, as shown below:

**Figure 6-11: Packet Recorder Node Query**

Active Queries

Stenographer Query

Query Type

Invoked via

Recorder Nodes

Status

Perce

No active queries

Query

▶

Submit

⊘

Reset

✕

Clear Inputs

✕

Clear Results

Query Type \*

Size

Fetch size and packet count

Relative Time

Absolute Time

Last \*

10

minute(s)

Any IP

Directional IP

IP

- IP Address -

Src Port

Dst Port

IP Protocol

- Select By Name -

- Dec -

Community ID

- Community ID -

Src Mac

- MAC Address -

Dst Mac

- MAC Address -

VLAN

Filter Interfaces

+

- None -

Policies

+

- None -

MetaWatch Device ID

MetaWatch Port ID

Deduplication

No

Yes

Dedup. Time Window

ms


Use the following options to specify the packets to include in the query:

- **Relative Time:** A time range relative to the current time in which look for packets.
- **Absolute Time:** A specific time range in which to look for packets.
- **Any IP:** Include packets with the specified IP address in the IP header (either source or destination).
- **Directional IP:** Include packets with the specified source and/or destination IP address in the IP header.
- **Src Port:** Include packets with the specified protocol port number in the Src Port field in the IP header.
- **Dst Port:** Include packets with the specified protocol port number in the Dst Port field in the IP header.

- **IP Protocol:** Select the IP protocol from the selection list or specify the numeric identifier of the protocol.
- **Community ID:** Select packets with a specific BRO community ID string.
- **Src Mac:** Select packets with a specific source MAC address.
- **Dst Mac:** Select packets with a specific destination MAC address.
- **VLAN:** Select packets with a specific VLAN ID.
- **Filter Interfaces:** Click the provision (+) control and, in the dialog that appears, enable the checkbox for one or more filter interfaces to which the query should be restricted. To add interfaces to the dialog, click the provision (+) control on the dialog and select the interfaces from the list that is displayed.
- **Policies:** Click the provision (+) control and, in the dialog that appears, enable the checkbox for one or more policies to which the query should be restricted. To add policies to the dialog, click the provision (+) control on the dialog and select the policies from the list that is displayed.
- **Max Bytes:** This option is only available for packet queries. Specify the maximum number of bytes returned by a packet query in a PCAP file.
- **Max Packets:** This option is only available for packet queries. Specify the maximum number of packets returned by a packet query in a PCAP file.
- **MetaWatch Device ID:** Filter packets with the specified MetaWatch device ID.
- **MetaWatch Port ID:** Filter packets with the specified MetaWatch port ID.

Alternatively, Global Query Configuration can be used to set the byte limit on packet query results.

**Figure 6-12: Global Query Configuration**



## 6.10 Viewing Query History

You can view the queries that have been submitted to the recorder node using the GUI or CLI.

To use the GUI to view the query history, select **Monitoring > Recorder Nodes** and scroll down to the **Query History** section.

**Figure 6-13: Monitoring > Recorder Nodes > Query History**

Packet Recorder	Stenographer Query	Query Type	Start Time	Status
HW-PR-1	after 180m ago	Size	Yesterday 9:58:10pm Pacific Daylight Time	✓
HW-PR-2	after 180m ago	Size	Yesterday 9:58:10pm Pacific Daylight Time	✓
HW-PR-2	after 120m ago	Size	Yesterday 9:58:00pm Pacific Daylight Time	✓
HW-PR-1	after 120m ago	Size	Yesterday 9:58:00pm Pacific Daylight Time	✓
HW-PR-1	after 120m ago	Application	Yesterday 9:53:34pm Pacific Daylight Time	✓
HW-PR-2	after 120m ago	Application	Yesterday 9:53:34pm Pacific Daylight Time	✗
HW-PR-2	after 120m ago	Application	Yesterday 9:49:39pm Pacific Daylight Time	✗
HW-PR-1	after 120m ago	Application	Yesterday 9:49:39pm Pacific Daylight Time	✓
HW-PR-1	after 60m ago	Application	Yesterday 9:49:14pm Pacific Daylight Time	✗
HW-PR-2	after 60m ago	Application	Yesterday 9:49:14pm Pacific Daylight Time	✗

The **Query History** section displays the queries submitted to each recorder node and the status of the query.

To download the query results, select **Download Results** from the **Menu** control for a specific query. To export the query history, click the **Export** control at the top of the table (highlighted in the figure above, to the right of the **Refresh** control).

To display query history using the CLI, enter the following command:

```
controller-1> show recorder-node query-history
```

#	Packet Recorder	Query	Type	Start	Duration
1	HW-PR-2	after 10m ago	analysis-hosts	2019-03-20 09:52:38.021000 PDT	3428
2	HW-PR-1	after 10m ago	analysis-hosts	2019-03-20 09:52:38.021000 PDT	3428

## 6.11 Using the CLI to Manage the DMF Recorder Node

### 6.11.1 Basic Configuration

To perform basic recorder node configuration, complete the following steps:

1. Assign a name to the recorder node device.

```
controller-1(config)# recorder-node device rn-alias
```

2. Set the MAC address of the recorder node.

```
controller-1(config-recorder-node)# mac 18:66:da:fb:6d:b4
```

If the management MAC is unknown, it can be determined from the chassis ID of connected devices.

```
controller-1> show connected-devices packet-recorder
```

#	Switch	IF Name	DMF name	SPAN?	Device Name	Device Description	Chassis ID	Port ID	Port Description	Management Address	Protocol
1	bt-1b9-1	ethernet50		False	recorder-1	DMF Recorder Node, SN HLEZYH2	18:66:da:fb:6d:b4	3c:fd:fe:1f:0f:82	enpl80e0f1	10.4.100.200	LLDP

3. Define the recorder node interface name.

```
controller-1(config)# recorder-fabric interface Intf-alias
controller-1(config-pkt-rec-intf)#
```

Any alphanumeric identifier can be assigned for the name of the recorder node interface, which changes the submode to **config-pkt-rec-intf**, where an optional description can be provided. This submode allows you to specify the switch and interface where the recorder node is connected.

4. Provide an optional description and identify the switch interface connected to the recorder node.

```
controller-1(config-pkt-rec-intf)# description 'Delivery point for recorder-node'
controller-1(config-pkt-rec-intf)# recorder-interface switch Switch-z9100 ethernet37
```

5. (Optional) **Recording**: Recording is enabled by default. To disable recording, enter the following commands:

```
controller-1(config)# recorder-node device rn-alias
controller-1(config-recorder-node)# no record
```

6. (Optional) **Disk Full Policy**: By default, **Disk Full Policy** is set to **rolling-fifo**, which means oldest packets will be deleted to make room for newer packets when recorder node disks are full. This configuration can be changed to **stop-and-wait**, which will allow the recorder node to stop recording until disk space becomes available. Enter the commands below to configure **Disk Full Policy** to **stop-and-wait**.

```
controller-1(config)# recorder-node device rn-alias
controller-1(config-recorder-node)# when-disk-full stop-and-wait
```

7. (Optional) **Max Packet Age**: This defines the maximum age in minutes of any packet in the recorder node. By default, **Max Packet Age** is not set, which means no limit is enforced. When **Max Packet Age** is set, packets recorded on the recorder node will be discarded after the specified number of minutes. To



set the maximum number of minutes that recorded packets will be kept on the recorder node, enter the following commands:

```
controller-1(config)# recorder-node device rn-alias
controller-1(config-recorder-node)# max-packet-age 30
```

This sets the maximum time to keep recorded packets to **30** minutes.



**Note: Max Packet Age** can be used in combination with the packet removal policy to control when packets are deleted based on age rather than disk utilization alone.

8. (Optional) **Max Disk Utilization:** This defines the maximum disk utilization in terms of a percentage between 5% and 95%. When this utilization is reached, the **Disk Full Policy** (**rolling-fifo** or **stop-and-wait**) is enforced. If unset, the default maximum disk utilization is **95%**; however, it can be configured using the following commands:

```
controller-1(config)# recorder-node device rn-alias
controller-1(config-recorder-node)# max-disk-utilization 80
```

9. (Optional) Disable any indexing configuration fields that will not be used in subsequent recorder node queries. All indexing fields are enabled by default. To disable a specific indexing option, enter the following commands from the **config-recorder-node-indexing** submode. To re-enable a disabled option, enter the command without the **no** prefix.

Use the following command enter the recorder node indexing submode:

```
controller-1(config-recorder-node)# indexing
controller-1(config-recorder-node-indexing)#
```

Then use the following commands to disable any fields that will not be used in subsequent queries:

- Disable MAC Source indexing: **no mac-src**
- Disable MAC Destination indexing: **no mac-dst**
- Disable outer VLAN ID indexing: **no vlan-1**
- Disable inner/middle VLAN ID indexing: **no vlan-2**
- Disable innermost VLAN ID indexing: **no vlan-3**
- Disable IPv4 Source indexing: **no ipv4-src**
- Disable IPv4 Destination indexing: **no ipv4-dst**
- Disable IPv6 Source indexing: **no ipv6-src**
- Disable IPv6 Destination indexing: **no ipv6-dst**
- Disable IP Protocol indexing: **no ip-proto**
- Disable Port Source indexing: **no port-src**
- Disable Port Destination indexing: **no port-dst**
- Disable MPLS indexing: **no mpls**
- Disable Community ID indexing: **no community-id**
- Disable MetaWatch Device ID: **no mw-device-id**
- Disable MetaWatch Port ID: **no mw-port-id**

For example, the following command disables indexing for the destination MAC address:

```
controller-1(config-recorder-node-indexing)# no mac-dst
```

10. Identify the recorder node interface by name in an out-of-band policy.

```
controller-1(config)# policy RecorderNodePolicy
controller-1(config-policy)# use-recorder-fabric-interface intf-1
controller-1(config-policy)#
```

11. Configure the DMF policy to identify the traffic to send to the recorder node.

```
controller-1(config-policy)# 1 match any
controller-1(config-policy)# # filter-interface FilterInterface1
controller-1(config-policy)# # action forward
```

This example forwards all traffic received in the monitoring fabric on filter interface **FilterInterface1** to the recorder node interface. The following is the running-config for this example configuration:

```
recorder-fabric interface intf-1
description 'Delivery point for recorder-node'
recorder-interface switch 00:00:70:72:cf:c7:cd:7d ethernet37
policy RecorderNodePolicy
action forward
filter-interface FilterInterface1
use-recorder-fabric intf-1
1 match any
```

### 6.11.2 Authentication Token Configuration

Static authentication tokens are pushed to each recorder node as an alternative form of authentication in headless mode, when the DMF controller is unreachable, or by third-party applications that do not have or do not need DMF controller credentials in order to query the recorder node.

To configure the recorder node with a static authentication token, use the following commands:

```
controller-1(config)# recorder-node auth token mytoken
Auth : mytoken
Token : some_secret_string <--- secret plaintext token displayed once here
controller-1(config)# show running-config recorder-node auth token
! recorder-node
recorder-node auth token mytoken $2a$12$ctw4PvsPySXrmMLYA.Mnyus9DpQ/bydGWD4LEhNL6xhPpkKNLzqWS <---hashed token shows in running
config
```

The DMF controller uses its own hidden authentication token to query the recorder node. To regenerate the controller authentication token, use the following command:

```
controller-1(config)# recorder-node auth generate-controller-token
```

### 6.11.3 Configuring the Pre-buffer

To enable the pre-buffer or change the time allocated, enter the following commands:

```
controller-1(config)# recorder-node device <name>
controller-1(config-recorder-node)# pre-buffer <minutes>
```

Replace **name** with the name of the recorder node. Replace **minutes** with the number of minutes to allocate to the pre-buffer.

### 6.11.4 Triggering a Recorder Node Event

To trigger an event for a specific recorder node, enter the following command from enable mode:

```
controller-1# trigger recorder-node <name> event <event-name>
```

Replace **name** with the name of the recorder node and replace **event-name** with the name to assign to the current event.

## 6.11.5 Terminating a Recorder Node Event

To terminate a recorder node event, use the following command:

```
controller-1# terminate recorder-node <name> event <event-name>
```

Replace **name** with the name of the recorder node and replace **event-name** with the name of the recorder node event to terminate.

## 6.11.6 Viewing Recorder Node Events

To view recorder node events, enter the following command from enable mode:

```
controller-1# show recorder-node events
# Packet Recorder Time                                     Event
-----|-----|-----|-----|-----|-----|
1 pkt-rec-740      2018-02-06 16:21:37.289000 UTC Pre-buffer event my-event1 complete. Duration 3 minute(s)
2 pkt-rec-740      2018-02-06 20:23:59.758000 UTC Pre-buffer event event2 complete. Duration 73 minute(s)
3 pkt-rec-740      2018-02-07 22:39:15.036000 UTC Pre-buffer event event-02-7/event3 complete. Duration 183 minute(s)
4 pkt-rec-740      2018-02-07 22:40:15.856000 UTC Pre-buffer event event5 triggered
5 pkt-rec-740      2018-02-07 22:40:16.125000 UTC Pre-buffer event event4/event-02-7 complete. Duration 1 minute(s)
6 pkt-rec-740      2018-02-22 06:53:10.216000 UTC Pre-buffer event triggered
```

## 6.12 Using the CLI to Run Recorder Node Queries



**Note:** The DMF controller prompt is displayed immediately after entering a query or replay request, but the query continues in the background. If you try to enter another replay or query command before the previous command is completed, an error message is displayed.

### 6.12.1 Packet Replay

To replay the packets recorded by a recorder node, enter the **replay recorder-node** command from enable mode.

```
controller-1# replay recorder-node <name> to-delivery <interface> filter
<stenographer-query>
[realtime | replay-rate <bps> ]
```

The following are the options available with this command.

- **name:** Specify the recorder node for which you wish to replay the recorded packets from.
- **interface:** The name of the DMF delivery interface to which the packets should be delivered.
- **stenographer-query:** The filter used to look up desired packets.
- (Optional) **real-time:** Replay the packets at the original rate recorded by the specified recorder node. The absence of this parameter will result in a replay up to the line rate of the recorder node interface.
- (Optional) **replay-rate bps:** Specify the number of bits per second to be used for replaying the packets recorded by the specified recorder node. The absence of this parameter will result in a replay up to the line rate of the recorder node interface.

The following command shows an example of a replay command using the **to-delivery** option.

```
controller-1# replay recorder-node packet-rec-740 to-delivery eth26-del filter
'after 1m ago'
controller-1#
Replay policy details:
controller-1# show policy-flow | grep replay
1 __replay_131809296636625 packet-as5710-2 (00:00:70:72:cf:c7:cd:7d) 0 0 6400 1
in-port 47 apply: name=__replay_131809296636625 output: max-length=65535,
port=26
```

## 6.12.2 Packet Data Query

You can use a packet query to search the packets recorded by a specific recorder node. The operation uses a Stenographer query string to filter only the interesting traffic. The query returns a URL that can be used to download and analyze the packets using Wireshark or other packet-analysis tools.

From enable mode, enter the **query recorder-node** command.

```
switch # query recorder-node <name> packet-data filter <stenographer-query>
```

The following is the meaning of each parameter:

- **name:** Identify the recorder instance.
- **packet-data filter stenographer-query:** Look up only the packets that match the specified Stenographer query.

The following example shows the results returned:

```
switch # query packet-recorder bq-bmf-packet-recorder-1 packet-data filter "after 1m ago and src host 8.8.8.8"
----- Packet Query Results -----
Individual URL(s) : /pcap/_packet_recorder_/bq-bmf-packet-recorder-1-2018-10-19-08-19-59-40b3dd8e.
pcap
- Error(s) -
None.
Past packet queries can be referenced on the controller using the show packet-capture files policy
__packet_recorder__ command. Any HTTP client can be used to download the files using the URLs indicated.
switch# show packet-capture files policy __packet_recorder__ | head
----- Packet Capture Files for Selected Policy -----
# File Name                                         File Size Last Modified URL
-----
1 coalesced-all-2018-09-20-13-13-33-5f7cc9ea.pcap 642KB 2018-09-20 13:13:33 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-20-13-13-33-5f7cc9ea.pcap
2 coalesced-all-2018-09-20-16-29-16-d0f90c5.pcap 106KB 2018-09-20 16:29:16 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-20-16-29-16-d0f90c5.pcap
3 coalesced-all-2018-09-21-10-13-48-526682b1.pcap 2.34MB 2018-09-21 10:13:48 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-21-10-13-48-526682b1.pcap
4 coalesced-all-2018-09-21-10-13-08-23d8266e.pcap 3.36MB 2018-09-21 10:13:08 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-21-10-13-08-23d8266e.pcap
5 coalesced-all-2018-09-21-10-45-53-b94fb5b.pcap 3.22MB 2018-09-21 10:45:53 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-21-10-45-53-b94fb5b.pcap
6 coalesced-all-2018-09-29-12-53-28-1aa5887f.pcap 276B 2018-09-29 12:53:28 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-09-29-12-53-28-1aa5887f.pcap
7 coalesced-all-2018-10-03-08-17-34-34e259e6.pcap 11.4MB 2018-10-03 08:17:35 PDT https://10.100.6.15/pcap/_packet_recorder_/coalesced-all-2018-10-03-08-17-34-34e259e6.pcap
```

## 6.12.3 Packet Object Query

The packet object query extracts unencrypted HTTP objects from packets matching the given stenographer filter. To run a packet object query, run the following query command:

```
switch# query recorder-node bmf-integrations-pr-1 packet-object filter 'after 5m ago'
```

The following example shows the results returned:

```
switch# query recorder-node bmf-integrations-pr-1 packet-object filter 'after 1m ago'
----- Packet Object Query Results -----
Coalesced URL : /pcap/_packet_recorder_/coalesced-bmf-2022-11-21-14-27-56-67a73ea9.tgz
Individual URL(s) : /pcap/_packet_recorder_/bmf-integrations-pr-1-2022-11-21-14-27-55-598f5ae7.tgz
```

Untar the folder to extract the HTTP objects.

## 6.12.4 Size Query

You can use a size query to analyze the number of packets and the total size of the packets recorded by a specific recorder node. The operation uses a Stenographer query string to filter only the interesting traffic.

To run a size query, enter the **query recorder-node** command from enable mode.

```
# query recorder-node <name> size filter <stenographer_query>
```

The following is the meaning of each parameter:

- **name:** Identify the recorder node.
- **size filter stenographer-query:** Analyze only the packets that match the specified Stenographer query.

The following example shows the results returned:

```
switch# query recorder-node <hq-bmf-packet-recorder-1> size filter "after 1m
ago and src host 8.8.8.8"
~ Summary Query Results ~
# Packets : 66
Size      : 7.64KB
~ Error(s) ~
None.
```

### 6.12.5 Window Query

You can use a window query to analyze the oldest available packet and most recent available packet recorded by a specific recorder node.

To run a window query, enter the **query recorder-node** command from enable mode.

```
switch# query recorder-node <name> window
```

The following is the meaning of each parameter:

- **name**: Identify the recorder node.

The following example shows the results returned:

```
switch# query recorder-node hq-bmf-packet-recorder-1 window
~~~~~ Window Query Results ~~~~~
Oldest Packet Available : 2020-07-30 05:01:08 PDT
Newest Packet Available : 2020-10-19 08:14:21 PDT
~ Error(s) ~
None.
```

### 6.12.6 Stopping a Query

You can use the abort **recorder-node** command to stop the current query running on the specified recorder. From enable mode, enter the following command:

```
controller-1# abort recorder-node <name> filter <string>
```

Replace **name** with the name of the recorder node, and use the **filter** keyword to identify the specific filter used to submit the query. If the specific query being run is unknown, an empty-string filter of "" can be used to terminate any running query.

```
controller-1# abort recorder-node hq-bmf-packet-recorder-1 filter ""
Abort any request with the specified filter? This cannot be undone. enter
"yes" (or "y") to continue:
yes
Result : Success
~ Error(s) ~
None.
```

## 6.13 Using RBAC to Manage Access to the DMF Recorder Node

You can use Role-Based Access Control (RBAC) to manage access to the DMF Recorder Node by associating a recorder node with an RBAC group.

To restrict access for a specific recorder to a specific RBAC group, use the CLI or GUI as described below.

### 6.13.1 RBAC Configuration Using the CLI

1. Identify the group to which you want to associate the recorder node.

Enter the following command from config mode on the active DMF controller:

```
controller-1(config)# group test
controller-1(config-group)#
```

2. Associate one or more recorder nodes with the group.

Enter the following CLI command from the **config-group** submode:

```
controller-1(config-group)# associate recorder-node <device-name>
```

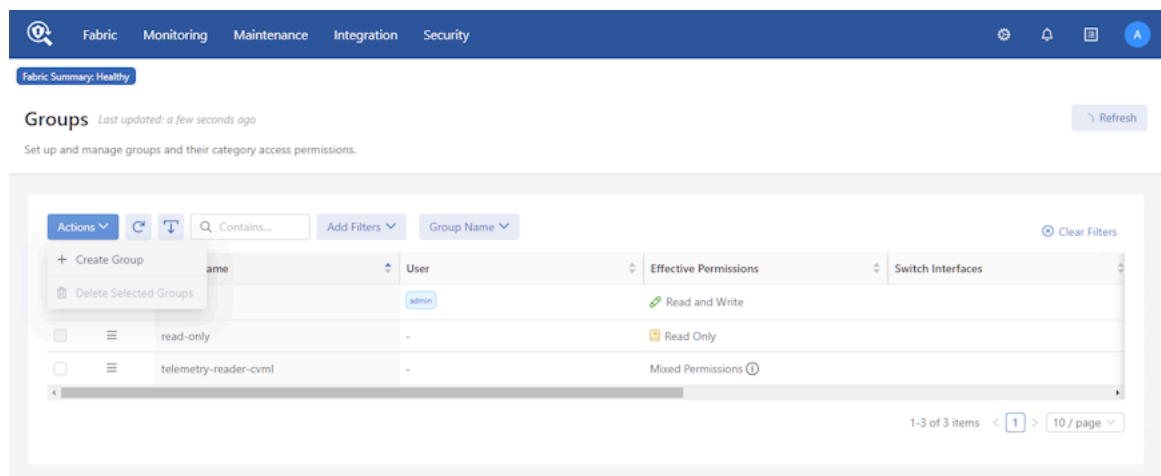
Replace **device-name** name with the name of the recorder node, as in the following example:

```
controller-1(config-group)# associate recorder-node HW-PR-1
```

### 6.13.2 RBAC Configuration Using the GUI

1. Select **Security > Groups**, and select **Edit** from the **Actions** and click **+ Create Group**.

**Figure 6-14: Create Security Group**



## 2. Enter a **Group Name**.

**Figure 6-15: Create Group**

Create Group

This group will manage the categories based on permissions.

\* Group Name

admin1

**Category Access**

Categories have a hierarchy. If no permission is assigned for a category, it will be inherited from the parent.

▼ Default

Default  
General configuration and states.

Default/Sensitive  
Sensitive data included in "General configuration and states."

AAA

App Log

Sysops

Time

Mixed Permissions

Does Not Elevate

Inherit

**Role Based Access Control**

Configure fine-grained access to specific Policies, Services, Switch Interfaces and Recorder Nodes.

Policy 0 Policies configured

Service 0 Services configured

Switch Interface 0 Switch Interfaces configured

Recorder Node 0 Recorder Nodes configured

Add Policy

Add Service

Add Switch Interface

Add Recorder Node

Cancel Create

## 3. Under the **Role Based Access Control** section select **Add Recorder Node**.

## 4. Select the **Recorder Node** from the selection list, and assign the permissions required.

- **Read:** The user can view recorded packets.
- **Use:** The user can define and run queries.
- **Configure:** The user can configure packet recorder instances and interfaces.

- **Export:** The user can export packets to a different device.

**Figure 6-16: Associate Recorder Node**

5. Click **Create**.

## 6.14 Using the CLI to View Information About a Recorder Node

This section describes how to monitor and troubleshoot recorder node status and operation. The recorder node stores packets on the main hard disk and the indices on the SSD volumes.

### 6.14.1 Viewing the Recorder Node Interface

To view information about the recorder node interface information, use the following command:

```
controller-1(config)# show topology recorder-node
# DMF IF      Switch IF  Name      State Speed  Rate Limit
-|-----|-----|-----|-----|-----|-----|
1 RecNode-Intf Arista7050 ethernet1 up    25Gbps -
```

### 6.14.2 Viewing Recorder Node Operation

```
controller-1# show recorder-node device packet-rec-740 interfaces stats
Packet Recorder Name Rx Pkts      Rx BytesRx Drop  Rx Errors Tx Pkts  Tx Bytes  Tx Drop Tx Errors
-----|-----|-----|-----|-----|-----|-----|-----|
packet-rec-740 pri1 2640908588614 172081747460802 84204084 0 24630503 3053932660 0 0
```

Information about a recorder node interface used as a delivery port in a DMF out-of-band policy is displayed in a list. Recorder node interfaces are listed as dynamically-added delivery interfaces.

```
Ctrl-2(config)# show policy PR-policy
Policy Name      : PR-policy
Config Status    : active - forward
Runtime Status   : installed
Detailed Status  : installed - installed to forward
Priority         : 100
Overlap Priority  : 0
# of switches with filter interfaces : 1
# of switches with delivery interfaces : 1
# of switches with service interfaces : 0
# of filter interfaces : 1
# of delivery interfaces : 1
# of core interfaces : 0
# of services : 0
# of pre service interfaces : 0
# of post service interfaces : 0
Push VLAN        : 1
Post Match Filter Traffic : 1.51Gbps
```



```

Total Delivery Rate      : 1.51Gbps
Total Pre Service Rate  : -
Total Post Service Rate : -
Overlapping Policies     : none
Component Policies      : none
Installed Time           : 2023-09-22 12:16:55 UTC
Installed Duration       : 3 days, 4 hours
~ Match Rules ~
# Rule
-|-----|
1 1 match any

~~~~~ Filter Interface(s) ~~~~~
# DMF IP      Switch      IF Name      State Dir Packets      Bytes      Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 Lab-traffic Arista-7050SX3-T3X5 ethernet7 up    rx  97831460642 51981008309480 382563  1.51Gbps 2023-09-22 12:16:55.738000 UTC

~~~~~ Delivery Interface(s) ~~~~~
# DMF IP      Switch      IF Name      State Dir Packets      Bytes      Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 FR-intf Arista-7050SX3-T3X5 ethernet35 up    tx  97831460642 51981008309480 382563  1.51Gbps 2023-09-22 12:16:55.738000 UTC

~ Service Interface(s) ~
None.

~ Core Interface(s) ~
None.

~ Failed Path(s) ~
None.
Ctrl-2(config)#

```

### 6.14.3 Viewing Errors and Warnings

The following table lists the errors and warnings that may be displayed by a recorder node. In the CLI, these errors and warnings can be displayed by entering the following commands:

- `show fabric errors`
- `show fabric warnings`
- `show recorder-node errors`
- `show recorder-node warnings`

### Table 3: Errors and Warnings

Type	Condition	Cause	Resolution
Error	Recorder Node (RN) management link down	RN has not received controller LLDP	Wait 30s if the recorder node is newly configured. Verify it is not connected to a switch port that is a DMF interface.
Error	RN fabric link down	Controller has not received RN LLDP	Wait 30s if recorder node is newly configured. Check it is online otherwise.
Warning	Disk/RAID health degraded	Possible hardware degradation	Investigate specific warning reported. Could be temperature issue. Possibly replace indicated disk soon.
Warning	Low disk space	Packet or index disk space has risen above threshold	Prepare for disk full soon
Warning	Disk full	Packet or index disk space is full. Packets are being dropped or rotated depending on removal policy.	Do nothing if removal policy is rolling-FIFO. Consider erasing packets to free up space otherwise.
Warning	Recorder misconfiguration on a DMF interface	A recorder node has been detected in the fabric on a switch interface that is configured as a filter or delivery interface.	Remove the conflicting interface configuration, or re-cable the recorder node to a switch interface not defined as a filter or delivery interface.

# 6.15 Using the GUI to view Recorder Node Statistics

Recorder node statistics can be viewed by clicking on the recorder node alias from the **Monitoring > Recorder Nodes** page.

Figure 6-17: List of Connected Recorder Nodes

Fabric Monitoring Maintenance Integration Security

Fabric: Healthy

Recorder Nodes

▽ Devices

≡ + − ↺ ↴

Filter table rows

	Name	MAC	IP Address	SKU	Hardware	Serial Number	Connected	Zerotouch State	Recording	Disk Full Policy
<input type="checkbox"/>	m1	52:54:00:eb:26:f8	10.243.255.4	—	—	Not Specified	✓	Ok	🔴 Off	—

Oct 13, 2023, 4:18:10am UTC

Click a **Recorder Node** to display the available recorder node statistics. All statistics are disabled/ hidden by default.

**Figure 6-18: Available Recorder Node Statistics**

The screenshot displays the configuration page for a Recorder Node named 'rn1'. The interface includes a top navigation bar with tabs for Fabric, Monitoring, Maintenance, Integration, and Security. Below the navigation bar, a status indicator shows 'Fabric: Healthy'. The main content area is divided into a left sidebar with a navigation menu and a main panel with configuration details.

**Navigation Menu (Left Sidebar):**

- Summary (selected)
- Alerts
- Events
- Interfaces
- Health
  - CPU
  - Memory
  - Storage
  - File Descriptors
  - Mount
  - Stenographer
  - Recording
- Options & Shortcuts
  - Selected Content Action: Replace (selected) / Append
  - Selected Content Placement:
    - Top of content area
    - Bottom of content area
    - In order listed above (selected)
  - Select All
  - Unselect All
  - Restore Defaults

**Main Panel Configuration:**

**Summary**

**Configuration**

- MAC: 52:54:00:eb:26:f8
- Recording: Off (toggle switch) / On
- Disk Full Policy: Rolling FIFO (toggle switch) / Stop and Wait
- Max Packet Age: —
- Pre Buffer: 0 minutes
- Max Disk Utilization: —
- Parse MetaWatch Trailer: —

**Fabric Connection**

- Connected: ✓

**Network**

- Auxiliary NIC: ✗

**Storage: Index Disk**

- Configured: ✗

**Storage: Packet Disk**

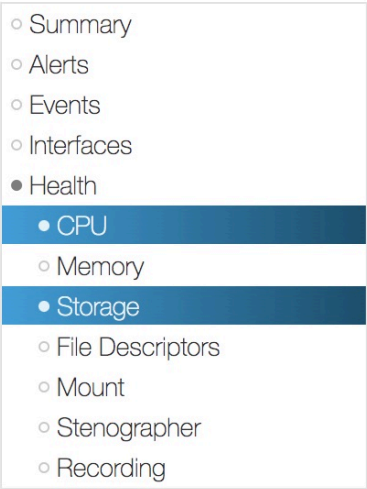
- Configured: ✗

**Indexing**

- MAC Source: ✗
- MAC Destination: ✗
- VLAN 1: ✗
- VLAN 2: ✗
- VLAN 3: ✗
- IPv4 Source: ✗
- IPv4 Destination: ✗
- IPv6 Source: ✗
- IPv6 Destination: ✗
- IP Protocol: ✗
- Port Source: ✗
- Port Destination: ✗
- MPLS: ✗
- Community ID: ✗
- MetaWatch Device ID: ✗
- MetaWatch Port ID: ✗

Statistics can be enabled/viewed by simply clicking on them. Selected statistics are highlighted in blue.

Figure 6-19: Selected Recorder Node Statistics



The recorder node shows health statistics for the following:

**CPU:** CPU health displays the compute resource utilization of the recorder node.

Figure 6-20: Recorder Node CPU Health Statistics

**CPU**

Collection Time Jun 18, 2020, 12:15:33pm Pacific Daylight Time

Cores

		Utilization						
Core Number	Name	User	Low Priority User	Kernel	I/O	H Interrupt	S Interrupt	Idle ▲
14	record-13	9.43%	—	1.11%	—	—	—	88.89%
15	record-14	8.43%	—	1.81%	—	—	—	89.16%
33	query-1	—	—	0.5%	—	—	—	99.5%
34	query-2	0.5%	—	—	—	—	—	99.5%
35	query-3	—	—	0.5%	—	—	—	99.5%

Jun 18, 2020, 7:15:41pm UTC Show: 10 25 100 All (1 - 5 / 40) 1 2 3 ... 8 <>

**Memory:** Memory related stats are displayed, such as total memory, used, free, available, etc.

Figure 6-21: Recorder Node Memory Statistics


**Memory**

Collection Time Jun 18, 2020, 12:15:34pm Pacific Daylight Time

Total	251.13 GB
Used	109.83 GB
Free	62.00 GB
Shared	3.00 MB
Buffer	48.00 MB
Cache	79.25 GB
Available	196.97 GB


**Storage:** Storage health displays the storage utilization percentage along with total and available capacity of Index and Packet virtual disks.

**Figure 6-22: Recorder Node Storage Statistics**



**Time-based Disk Utilization Statistics:** Time-based Disk Utilization Statistics provides an estimated time period until the Index and Packet virtual disks reach full storage capacity. This estimate is calculated based on data points (incoming data rate) collected periodically from recorder node for a certain time duration. Note that if the collected data points are insufficient to calculate the disk-full estimate, it will show inaccurate. However, once a sufficient number of data points are collected, the estimate will be calculated and displayed automatically.

**Figure 6-23: Time-based Disk Utilization Statistics**



**Virtual Disks:** Virtual Disks health stats displays the Index and Packet virtual disks size, state, health and RAID level configuration.

**Figure 6-24: Recorder Node Virtual Disk Details**

Virtual Disks

	Name ▲	Size	State	Block Health	RAID Level	Physical Disk Count
▶	index	6.98 TB	Optimal	✓ Good	RAID 0	2
▶	packet	142.55 TB	Optimal	✓ Good	RAID 0	16

Jun 18, 2020, 7:15:41pm UTC Show: [10](#) [25](#) [100](#) [All](#) (1 - 2 / 2)

Click on the drop-down arrow next to the virtual disk name to obtain information regarding participating physical disks, such as slot numbers, type, size, state, temperature and Dell's Self Monitoring Analysis and Report Technology (SMART) stats, such as errors and failures, if any.

**Figure 6-25: Recorder Node Virtual Disk Statistics**

Virtual Disks

	Name ▲	Size	State	Block Health	RAID Level	Physical Disk Count	
▼	index	6.98 TB	Optimal	✓ Good	RAID 0	2	
Physical Disks							
					SMART		
Slot Number ▲	Type	Size	State	Temperature (°C)	Predictive Failure Count	Media Error Count	Other Error Count
12	SSD	3.49 TB	Online, Spun Up	<div><div></div></div> 44	0	0	0
13	SSD	3.49 TB	Online, Spun Up	<div><div></div></div> 45	0	0	0
Jun 18, 2020, 7:40:21pm UTC					Show: <a href="#">10</a> <a href="#">25</a> <a href="#">100</a> <a href="#">All</a>		(1 - 2 / 2)

▼ packet

142.55 TB

Optimal

✓ Good

RAID 0

16

Physical Disks

					SMART		
Slot Number ▲	Type	Size	State	Temperature (°C)	Predictive Failure Count	Media Error Count	Other Error Count
0	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
1	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
2	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
3	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
4	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
5	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
6	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0
7	HDD	8.91 TB	Online, Spun Up	<div><div></div></div> 26	0	0	0

**File Descriptors:** the **File Descriptor** section displays the following:

- **File Descriptors (current):** Current number of files open in the entire system.
- **Max System File Descriptors:** Highest number of open files allowed on the entire system.
- **Max Stenographer File Descriptors:** Highest number of open files allowed for Stenographer application.

**Figure 6-26: Recorder Node File Descriptors Statistics**

#### File Descriptors

Collection Time Jun 18, 2020, 12:15:30pm Pacific Daylight Time

File Descriptors (current) 74,752

Max System File Descriptors 26,284,554

Max Stenographer File Descriptors 10,000,000

**Mount:** Mount section displays the Index and Packet disk mount information, such as volume name, mount point, file system type and mount health.

**Figure 6-27: Recorder Node Mount Information**

**Mount**

Collection Time Jun 18, 2020, 12:48:51pm Pacific Daylight Time

Index Mounts

Volume ▲	Mount Point	File System Type	Mount Point Health
/dev/sdb1	/idx	xfs	healthy

Jun 18, 2020, 7:49:04pm UTC

Packet Mounts

Volume ▲	Mount Point	File System Type	Mount Point Health
/dev/sda1	/pkt	xfs	healthy


Jun 18, 2020, 7:49:04pm UTC

**Stenographer:** Stenographer Statistics are displayed as follows:

**Figure 6-28: Recorder Node Stenographer Statistics**

**Stenographer**

Collection Time Jun 18, 2020, 12:48:54pm Pacific Daylight Time

Initialized 

Tracked Files 36,392

Cached Files 72,784

Max Cached Files 3,000,000

- **Initialized:** Displays the Stenographer application running state. A green check mark indicates that the application was initialized successfully. When the Stenographer application is starting up, a red x mark is expected. During this time, recording and querying is disallowed.
- **Tracked Files:** Tracked files are the total number of files stored under each CPU instance thread.
- **Cached Files:** Cached files are the number of files that are open and have a file descriptor.
- **Max Cached Files:** Maximum cached files is the total number of files that are allowed to be open.

These numbers are further divided and displayed for each recording thread and can be viewed in the **Recording Threads** table:

**Figure 6-29: Recorder Node Max Cached Files Statistics**

Recording Threads

1 2 3 4 <>

Instance ▲	Tracked Files	Cached Files	Max Cached Files
0	1,188	2,376	187,500
1	1,186	2,372	187,500
2	1,195	2,390	187,500
3	1,262	2,524	187,500
4	3,574	7,148	187,500

Jun 18, 2020, 7:49:04pm UTC Show: 10 25 100 All (1 - 5 / 16) 1 2 3 4 <>

**Recording:** Recording stats displays packet stats, such as dropped packets, total packets and collection time for each CPU core.

**Figure 6-30: Recorder Node Statistics**

Recording Threads

CPU Core ▲	Dropped Packets	Total Packets	Collection Start Time
1	0	965,277	Yesterday, 1:40:53am Pacific Daylight Time
2	0	923,565	Yesterday, 1:40:53am Pacific Daylight Time
3	0	970,454	Yesterday, 1:40:53am Pacific Daylight Time
4	0	918,498	Yesterday, 1:40:53am Pacific Daylight Time
5	0	968,702	Yesterday, 1:40:54am Pacific Daylight Time

Jun 18, 2020, 8:02:41pm UTC Show: [10](#) [25](#) [100](#) [All](#) (1 - 5 / 16) 1 [2](#) [3](#) [4](#) <>

The following displays packet size distribution stats.

**Figure 6-31: Recorder Node Packet Size Distribution Statistics**

Frames under 64 bytes 0  
Frames of 64 bytes 0  
Frames between 65 and 127 bytes 5,024,888  
Frames between 128 and 255 bytes 75,931,017,000  
Frames between 256 and 511 bytes 714,644,788  
Frames between 512 and 1,023 bytes 181,532,956  
Frames between 1,024 and 1,522 bytes 179,379,007  
Frames between 1,523 and 9,522 bytes 0  
Frames over 9,522 bytes 0

The following displays interface errors, such as CRC errors, frame length errors and back pressure errors:

**Figure 6-32: Recorder Node Interface Errors**

Collection Time Jun 18, 2020, 1:02:29pm Pacific Daylight Time  
CRC Errors 0  
Frame Length Errors 0  
Back Pressure Errors 0

## 6.16 Changing the Recorder Node Default Configuration

Configuration settings are automatically downloaded to the recorder node from the DMF controller, which eliminates the need for box-by-box configuration. However, you can override the default configuration for a recorder node from the **config-recorder-node** submode for any recorder node.



**Note:** In the current release, these options are available only from the CLI, and are not included in the DMF controller GUI.

To change the CLI mode to **config-recorder-node**, enter the following command from config mode on the active DMF controller:

```
controller-1(config)# recorder-node device <instance>
```



Replace **instance** with the alias you want to use for the recorder node. This alias is associated with the MAC hardware address, using the **mac** command.

Use any of the following commands from **config-recorder-node** submode to override the default configuration for the associated recorder node:

- **banner**: Set the recorder node pre-login banner message
- **mac**: Configure the MAC address for the recorder node

Additionally, the below configurations can be overridden to use values specific to the recorder node or can also be used in a merge-mode along with the configuration inherited from the DMF controller:

- **ntp**: Configure recorder node to override default timezone and NTP parameters.
- **snmp-server**: Configure recorder node SNMP parameters and traps.
- **logging**: Enable recorder node logging to controller.
- **tacacs**: Set TACACS defaults, server IP address(es), timeouts and keys.

The following commands can be used, from the **config-recorder-node** submode, to change the default configuration on the recorder node:

- **ntp override-global**: Override global time configuration with recorder node time configuration.
- **snmp-server override-global**: Override global SNMP configuration with recorder node SNMP configuration.
- **snmp-server trap override-global**: Override global SNMP trap configuration with recorder node SNMP trap configuration.
- **logging override-global**: Override global logging configuration with packet recorder logging configuration.
- **tacacs override-global**: Override global TACACS configuration with recorder node TACACS configuration.

To configure the recorder node to work in a merge mode by merging its specific configuration with that of the DMF controller, execute the following commands in the **config-recorder-node** submode:

- **ntp merge-global**: Merge global time configuration with recorder node time configuration.
- **snmp-server merge-global**: Merge global SNMP configuration with recorder node SNMP configuration.
- **snmp-server trap merge-global**: Merge global SNMP trap configuration with recorder node SNMP trap configuration.
- **logging merge-global**: Merge global logging configuration with recorder node logging configuration.


TACACS configuration does not have a merge option. It can either be inherited completely from the DMF controller or overridden to use only the recorder node specific configuration.

## 6.17 Large PCAP Queries

To run large PCAP queries to the recorder node, access the recorder node via a web browser. This allows you to run packet queries directly to the recorder node without specifying the maximum byte or packet limit for the PCAP file (which is required if the query is executed from the DMF Controller).

To access the recorder node directly, use the URL **https://RecorderNodeIP** in a web browser, as shown below:

**Figure 6-33: URL to Recorder Node**



The following page will be displayed:

**Figure 6-34: Recorder Node Page**

Recorder Node IP Address:   
BMF Controller Username:   
BMF Controller Password:   
Stenographer Query Filter:   
Stenographer Query ID:   
Save pcap as:

- **Recorder Node IP Address:** Enter the IP address of the target recorder node.
- **DMF Controller Username:** Provide the DMF controller username.
- **DMF Controller Password:** Provide the password for authentication.
- **Stenographer Query Filter:** The query filter can be used to filter the query results to look for specific packets. For example, to search for packets with a source IP address of **10.0.0.145** in the last 10 minutes, use the following filter:

```
after 10m ago and src host 10.0.0.145
```

- **Stenographer Query ID:** Starting in DMF 8.0, a Universally Unique Identifier (UUID) is required to run queries. To generate a UUID, run the following command on any Linux machine and use the result as the Stenographer query ID:

```
$ uuidgen  
b01308db-65f2-4d7c-b884-bb908d111400
```

- **Save pcap as:** Provide the file name to be used for this PCAP query result.
- **Submit Request:** Click on **Submit Request**. This will send a query to the specified recorder node, and it will save the PCAP file with the provided file name to the default download location for the browser.

## 6.18 Recorder Node Management Migration L3ZTN

After the first boot (initial configuration) is completed, the recorder node can be removed from the old controller, and it can be pointed to a new controller via the CLI in the case of a Layer-3 topology mode.



**Note:** In order for appliances to connect to the DMF controller in Layer-3 Zero Touch Network (L3ZTN) mode, the DMF controller deployment mode must be configured as **pre-configure**.

To migrate management to a new controller, follow the steps below:

1. Remove the recorder node and switch from the old controller using the commands below:

```
controller-1(config)# no recorder-node device <RecNode>
```

```
controller-1(config)# no switch <Arista7050>
```

2. Add the switch to the new controller.
3. SSH to the recorder node and configure the new controller IP using the **zerotouch 13ztn controller-ip** command:

```
controller-1(config)# zerotouch 13ztn controller-ip 10.2.0.151
```

4. After pointing the recorder node to use the new controller, reboot the recorder node.
5. Once the recorder node is back online, the DMF controller should receive the ZTN request.

```

controller-1(config)# show zerotouch request
# Request-history      Ip address Action      Timestamp      Result      Message
-----
1 24:6e:96:78:58:b4 (Bell) 10.106.8.7 switch-light-manifest 2021-05-26 18:25:05.149000 UTC unable-to-service ZTN is not allowed for this device: No application configuration for device mac 24:6e:96:78:58:b4

```

6. After the DMF controller has received a ZTN request from the recorder node, it can be added to the DMF controller running-configuration using the below command:

```

controller-1(config)# recorder-node device RecNode
controller-1(config-recorder-node)# mac 24:6e:96:78:58:b4

```

7. The recorder node should now be added to the new DMF controller. It can be verified using the command below:

```

controller-1(config)# show zerotouch request
# Request-history      Ip address Action      Timestamp      Result      Message
-----
1 24:6e:96:78:58:b4 (Bell) 10.106.8.7 switch-light-manifest 2021-05-26 18:25:05.149000 UTC unable-to-service ZTN is not allowed for this device: No application configuration for device mac 24:6e:96:78:58:b4

```

## 6.19 Recorder Node CLI

The following commands are available from the recorder node:

The **show version** command can be used to view the version and image information that recorder node is running on.

```

RecNode(config)# show version
Controller Version : DMF Recorder Node 8.1.0 (bigswitch/enable/dmf-8.1.x #5)
RecNode(config)#

```

The **show controllers** command can be used to view the connected DMF controllers to the recorder node. Note that if the recorder node is connected to a DMF controller cluster, then all the cluster nodes should be listed in the command output:

```

RecNode(config)# show controllers
controller      Role    State    Aux
-----
tcp://10.106.8.2:6653 master  connected 0
tcp://10.106.8.3:6653 slave   connected 0
tcp://10.106.8.3:6653 slave   connected 1
tcp://10.106.8.3:6653 slave   connected 2
tcp://10.106.8.2:6653 master  connected 1
tcp://10.106.8.2:6653 master  connected 2
RecNode(config)#

```

## 6.20 Multiple Queries

The GUI can be used to run multiple recorder node queries

To run queries on recorded packets by the recorder node, navigate to the **Monitoring > Recorder Nodes** page.

Under the **Query** section, click on the **Query Type** drop-down to select the type of analysis that you would like to run on the recorded packets as shown below:

**Active Queries**

Stenographer Query	Query Type	Invoked via	Recorder Nodes	Status	Percent Complete	Current Value	Max Value
No active queries							

**Query**

▶ Submit
⊘ Reset
✕ Clear Inputs
✕ Clear Results

---

Query Type \* **TCP**

- Window
- Size
- Application
- Packet Data**
- Packet Object
- HTTP
- HTTP Request
- HTTP Stat
- DNS
- Replay
- IPv4
- IPv6
- TCP**
- TCP Flow Health
- UDP
- Hosts
- RTP Stream
- Delete

Last \* **Packet Data**

IP **Directional IP**

Src Port

Dst Port

IP Protocol **- Dec -**

Community ID

Src Mac

Dst Mac

VLAN

Filter Interfaces + - None -

Policies + - None -

MetaWatch Device ID


MetaWatch Port ID


Deduplication No ☐ Yes ☒


Dedup. Time Window  ms


After selecting the query type, you can use filters to limit or narrow the search to obtain specific results. Providing specific filters also helps to complete the query analysis faster. In the following example, the query result for the TCP query type will return the results for IP address **10.240.30.24** for the past **10** minutes.

▼ Query

  
Submit

  
Reset

  
Clear Inputs

  
Clear Results

Query Type \* **TCP** ▼  
*Fetch TCP conversation analysis*

Relative Time

Absolute Time

Last \* **10** ▼ minute(s) ▼

Any IP

Directional IP

IP **10.120.30.24** / **32** ▼

**Figure 6-35:**

After entering the desired filters, click on the **Submit** button. The **Progress** dialog will be displayed, showing the **Elapsed Time** and **Progress** percentage of the running query:

Progress

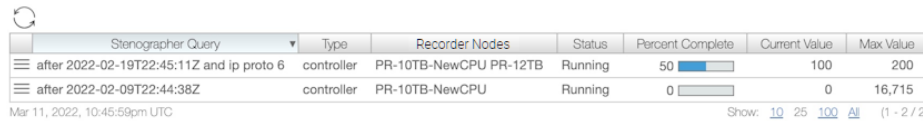
Querying ...

✗ Elapsed Time: 1 min, 59 secs

Progress: 0%

While a query is in progress, another query can be initiated from a new DMF controller web session. The query progress can be viewed under the **Active Queries** section:

▼ Active Queries



Stenographer Query	Type	Recorder Nodes	Status	Percent Complete	Current Value	Max Value
after 2022-02-19T22:45:11Z and ip proto 6	controller	PR-10TB-NewCPU PR-12TB	Running	50	100	200
after 2022-02-09T22:44:38Z	controller	PR-10TB-NewCPU	Running	0	0	16,715

Mar 11, 2022, 10:45:59pm UTC Show: 10 25 100 All (1 - 2 / 2)

## 6.21 Ability to Deduplicate Packets - Query from Recorder Node

For Recorder Node queries, the recorded packets matching a specified query filter may contain duplicates when packet recording occurs at several different TAPs within the same network; i.e., as a packet moves through the network, it may be recorded multiple times. The dedup feature removes duplicate packets from the query results. By eliminating redundant information, packet deduplication improves query results' clarity, accuracy, and conciseness. Additionally, the dedup feature significantly reduces the size of query results obtained from packet query types.

### 6.21.1 Using the CLI to Deduplicate Packets

In the DMF controller CLI, packet deduplication is available for the packet data, packet object, size, and replay query types. Deduplication is turned off by default for these queries. To enable deduplication, "dedup" must be added to the end of the query command after all optional values have been selected (if any).

The following are command examples of enabling deduplication.

Enabling deduplication for a size query:

```
controller# query recorder-node rn size filter "before 5s ago" dedup
```

Enabling deduplication for a packet data query specifying a limit for the size of the PCAP file returned in bytes:

```
controller# query recorder-node rn packet-data filter "before 5s ago" limit-bytes 2000 dedup
```

Enabling deduplication for a replay query:

```
controller# replay recorder-node rn to-delivery dintf filter "before 5s ago" dedup
```

Enabling deduplication for a replay query specifying the replay rate:

```
controller# replay recorder-node rn to-delivery dintf filter "before 5s ago" replay-rate 100 dedup
```

A time window (in milliseconds) can also be specified for deduplication. The time window defines the time required between timestamps of identical packets to no longer be considered duplicates of each other. For example, for a time window of 200 ms, two identical packets with timestamps that are 200 ms (or less) apart are duplicates of each other. In contrast, if the two identical packets had timestamps more than 200 ms apart, they would not be duplicates of each other.

The time window must be an integer between 0 and 999 (inclusive) with a default time window of 200 ms when deduplication is enabled and no set time window value.

---

To configure a time window value, **dedup-window** must be added after **dedup** and followed by an integer value for the time window.

```
controller# query recorder-node rn size filter "before 5s ago" dedup dedup-  
window 150
```

### 6.21.2 Using the GUI to Deduplicate Packets

In the DMF controller GUI, packet deduplication is available for the packet data, packet object, size, replay, application, and analysis query types. Deduplication is not enabled by default for these queries. To enable deduplication perform the following steps:

1. Set the toggle switch deduplication to **Yes** in the query submission window.
2. Specify an optional time window (in milliseconds) as required by entering an integer between 0 and 999 (inclusive) into the **Deduplication Time Window** field. The time window will default to 200 ms if no time window value is set.
3. Click **Submit** to continue.







**Note:** If a time window value is specified, but deduplication is not toggled, packet deduplication will not occur.

The following is an example of enabling deduplication for a size query specifying a time window value.

**Figure 6-36: Query**

▽ Query

Submit
Reset
Clear Inputs
Clear Results

Query Type \* **Size**  
Fetch size and packet count

Relative Time Absolute Time

Last \* **10** minute(s)

Any IP Directional IP

IP **- IP Address -** / **- Prefix -**

Src Port

Dst Port

IP Protocol **- Select By Name -** **- Dec -**

Community ID **- Community ID -**

Src Mac **- MAC Address -**

Dst Mac **- MAC Address -**

VLAN

Filter Interfaces **+ - None -**

Policies **+ - None -**

MetaWatch Device ID

MetaWatch Port ID

Deduplication No ☒ Yes ?

Dedup. Time Window **200** ms ?

### 6.21.3 Limitations

Expect a query with packet deduplication enabled to take longer to complete than with packet deduplication disabled. Hence, packet deduplication, by default, is disabled.

The maximum time window value permitted is 999 ms to ensure that TCP retransmissions are not regarded as duplicates, assuming that the receive timeout value for TCP retransmissions (of any kind) is at least 1 second. If the receive timeout value is less than 1 second (particularly, exactly 999 ms or less), then it is possible for TCP retransmissions to be regarded as duplicates when the time window value used is larger than the receive timeout value.

Due to memory constraints, removing some duplicates may not occur as expected. This scenario is likely to occur if a substantial amount of packets match the query filter, which all have timestamps within the specified time window from each other. We refer to this scenario as the query having exceeded the packet window capacity. To mitigate this from occurring, decrease the time window value or use a more specific query filter to reduce the number of packets matching the query filter at a given time.



## Link Aggregation

---

This chapter describes how to configure link aggregation groups between switches, between switches and tools, or between switches and taps.

### 7.1 Configuring Link Aggregation

Link aggregation lets you combine multiple LAN links and/or cables in parallel. Link aggregation provides a high level of redundancy and higher transmission speed.



**Note:** When connecting a Link Aggregation Group (LAG) to a DMF Service Node appliance, member links can be connected to multiple DMF Service Node appliances with data ports of the same speed.

DMF provides a configurable method of hashing for load distribution among the members of a LAG. The enhanced hashing algorithm automatically assigns the best hashing type for the switch and traffic. This setting also lets you manually select the packet types and fields used for load distribution among the members of a port-channel interface. For the supported switch platforms, enhanced mode and symmetric hashing are enabled by default. With symmetric hashing, bidirectional traffic between two hosts going out on a port channel is distributed on the same member port.

The default hashing option uses the best available packet header field that applies to each packet and that is supported by the switch. These fields can include the following:

- IPv4
- IPv6
- MPLS (disabled by default)
- L2GRE packet

If none of these headers can be used, Layer-2 header fields (source MAC address, destination MAC address, VLAN-ID, and ethertype) are used to distribute traffic among the LAG member interfaces. Hashing on the following packet header fields is enabled by default:

- `hash l2 dst-mac eth-type src-mac vlan-id`
- `hash ipv4 dst-ip src-ip`
- `hash ipv6 dst-ip src-ip`
- `hash l2gre inner-l3 dst-ip src-ip`
- `hash symmetric`



**Note:** VN-tagged packets and QinQ packets are treated as L2 packets and Layer-2 headers are used to distribute traffic among LAG member interfaces for these packets:

#### 7.1.1 Using the GUI to Configure Link Aggregation Groups

To view, manage, or create Link Aggregation Groups (LAGs) in the monitoring fabric, complete the following steps:

1. Select **Fabric > LAGs** from the main menu.

**Figure 7-1: Link Aggregation Groups**

Link Aggregation Groups (LAGs)

Switch	Name	Hash Type	State Flags	Members	Tx Packet Count	Tx Byte Count	Tx Packet Rate	Tx Bit Rate	Tx Peak Bit Rate	Tx Peak Packet Rate	Tx Peak Byte Time	Tx Peak Packet Time	Rx Packet Count
No lags													

2. To create a new LAG, click the provision control (+) at the upper left corner of the table.

**Figure 7-2: Create LAG**

Create LAG

Name \*

Switch \*

CANCEL SAVE

3. Enter a name for the LAG.

4. Select the switch where you want to create the LAG.

**Figure 7-3: Create LAG**

**Create LAG**

Name \*  
MY-LAG

Switch \*  
T3X5-Leaf-7050sx3-1 (00:00:d4:af:f7:4)  
This switch has LAG Enhanced Hash settings ?

Member Interfaces

T3X5-Leaf-7050sx3-1

- ☒ ethernet35
- ☒ ethernet37

T3X5-Leaf-7050sx3-1

ethernet1	ethernet2
ethernet3	ethernet4
ethernet5	ethernet6
ethernet7	ethernet8
ethernet9	ethernet10
ethernet11	ethernet12
ethernet13	ethernet14
ethernet15	ethernet16
ethernet17	ethernet18
ethernet19	ethernet20
ethernet21	ethernet22
ethernet23	ethernet24
ethernet25	ethernet26
ethernet27	ethernet28
ethernet29	ethernet30
ethernet31	ethernet32
ethernet33	ethernet34
ethernet35	ethernet36
ethernet37	ethernet38
ethernet39	ethernet40
ethernet41	ethernet42
ethernet43	ethernet44
ethernet45	ethernet46
ethernet47	ethernet48
ethernet49	ethernet50
ethernet51	ethernet52
ethernet53/1	ethernet53/2
ethernet53/3	ethernet53/4

CANCEL SAVE

5. Select the interfaces to include in the LAG and click **Submit**.

### 7.1.2 Using the CLI to Configure Link Aggregation Groups

1. Use the **lag-interface** command to enter the config-switch-lag-if submode, where you can define the LAG member interfaces and specify the type of load distribution (hashing) to use for the LAG.
2. Use the **member** command to add an interface to a LAG. Enter this command for each interface you want to add to the LAG. To remove an interface, use the **no member** version of the command.

For example, the following commands add two interfaces to a LAG named **my-lag**.

```
controller-1(config)# switch DMF-FILTER-SWITCH-1
controller-1(config-switch)# lag-interface mylag
controller-1(config-switch-lag-if)# member ethernet13
controller-1(config-switch-lag-if)# member ethernet14
```

3. To configure multiple delivery interfaces as a LAG, complete the following steps:
  - a. Assign a name to the LAG and enter **config-switch-lag-if** submode.

```
controller-1(config)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch)# lag-interface lag1
controller-1(config-switch-lag-if)#
```

- b. Assign members to the LAG.

```
controller-1(config-switch-lag-if)# member ethernet39
controller-1(config-switch-lag-if)# member ethernet40
```

- c. To view the configured LAGs, enter the **show lag** command, as in the following example:

```
controller-1> show lag
```

### 7.1.3 Configuring Hashing Fields

To configure the hashing fields manually via the CLI, use the **lag-enhanced-hash** command to enter **config-switch-hash** mode as in the following example:

```
controller1(config)# switch DMF-FILTER-SWITCH-1
controller1(config-switch)# lag-enhanced-hash
controller1(config-switch-hash)#
```

The hash commands have the following syntax:

- To hash on GTP fields, pick one of the following options:

```
controller1(config-switch-hash)# hash gtp
header-first-byte Configure fields to identify GTP traffic
port-match Configure UDP tunnel port match entry
```

- To hash on IPv4 fields, pick one of the following options:

```
controller1(config-switch-hash)# hash ipv4
<cr>
dst-ip Destination IPv4 address (optional)
l4-dst-port TCP/UDP destination port (optional)
l4-src-port TCP/UDP source port (optional)
protocol IP protocol (optional)
src-ip Source IPv4 address (optional)
vlan-id Vlan Id (optional)
```

- To hash on IPv6 fields, pick one of the following options

```
controller1(config-switch-hash)# hash ipv6
<cr>
dst-ip Collapsed destination IPv6 address (optional)
l4-dst-port TCP/UDP destination port (optional)
l4-src-port TCP/UDP source port (optional)
nxt-hdr Next Header (optional)
src-ip Collapsed source IPv6 address (optional)
vlan-id Vlan Id (optional)
```

- To hash on Layer-2 fields, pick one of the following options:

```
controller1(config-switch-hash)# hash l2
dst-mac Destination xMAC address
eth-type Ethernet Type
src-mac Source MAC address
vlan-id Vlan Id
```

- To hash on L2GRE fields, pick one of the following options:

```
controller1(config-switch-hash)# hash l2gre
inner-l2 Use inner L2 fields for hash computation (optional)
inner-l3 Use inner L3 fields for hash computation (optional)
```

- To hash on MPLS labels, pick one of the following options:

```
controller1(config-switch-hash)# hash mpls
<cr>
label-1 Lower 16 bits of MPLS label 1 (optional)
label-2 Lower 16 bits of MPLS label 2 (optional)
label-3 Lower 16 bits of MPLS label 3 (optional)
label-hi-bits Higher 4 bits of MPLS Labels 1,2 and 3 (optional)
```

- To manually configure the hash seeds:

```
controller1(config-switch)# hash seeds
<First hash seed> Configure seed1 for hash computation
controller1(config-switch-hash)# hash seeds 3809
<cr>
<Second hash seed> Configure seed2 for hash computation (optional)
controller1(config-switch-hash)# hash seeds 3809 90901
<cr>
```

- To enable/disable symmetric hashing

```
controller1(config-switch-hash)# hash symmetric
<cr>
disable Disable symmetric hashing
enable Enable symmetric hashing
```

## 7.1.4 L2 GRE Key Hashing

The L2 GRE Key-based hashing feature introduced in DMF 8.4 allows the L2 GRE packets to hash based on the L2 GRE (Tunnel) Key on Core DMF switches.

Previously, L2 GRE payload-based hashing (InnerL2 or InnerL3) applied only to L2 GRE packets terminated at DMF delivery or filter switches. If a user wanted to hash L2 GRE packets transiting a DMF core switch, the L2 GRE payload-based hashing across port-channel interfaces would not have been functional as the L2 GRE tunnel was not terminating on the core DMF switch.

With the L2 GRE Key-based hashing feature, users can now hash L2 GRE packets based on the L2 GRE Key on core DMF switches.



**Note:** The L2 GRE Key-based hashing feature applies to switches running SWL OS and does not apply to switches running EOS.

### CLI Configuration

L2 GRE Key-based hashing is supported only for the IPv4-based packets with L2 GRE payload. This feature does NOT support the IPv6 packets with L2 GRE payloads.

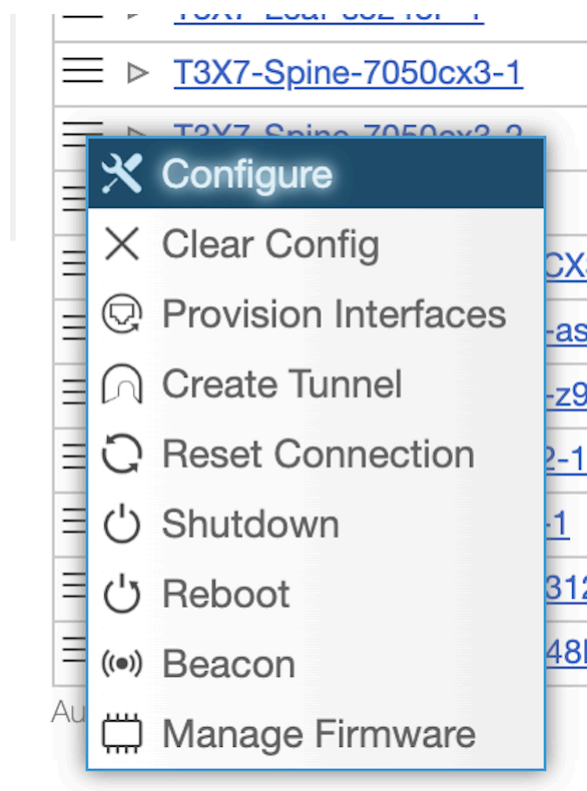
Enable the L2 GRE Key hashing by setting the **l2-gre-key** parameter as shown in the following example.

```
Controller-Active# show running-config switch DMF-SWITCH-1
! switch
switch DMF-SWITCH-1
mac c0:d6:82:17:fd:5a
!
lag-enhanced-hash
hash ipv4 l2-gre-key
hash symmetric disable
```

### UI Configuration

1. Configure the L2 GRE Key Hashing in the UI for a switch in the **Fabric > Switches** page using the table row menu action **Configure** option.

**Figure 7-4: Fabric Switch Configure Menu**



2. Enable the **L2 GRE Key** for the **IPv4** packets in the **LAG Enhanced Hash** step.



**Note:** The L2 GRE Key is unsupported for IPv6 and VxLAN Inner L3.

**Figure 7-5: Configure Switch L2 GRE Key**

- Click the **Submit** button to save the configuration.

### CLI Commands

Use the following CLI commands to verify settings and to troubleshoot any issues that may arise.

```
# show lag-enhanced-hash
```

While logged into a switch use the following commands to troubleshoot this feature.

```
root@DMF-SWITCH-1:~# ofad-ctl gt PORT_CHANNEL_ENHANCED_HASH_FIELD
Hash Field Configs:
-----
Symmetric Hashing:
Disabled
L2GRE Key Hashing:
Enabled
L2 Fields:
IPv4 Fields:
DSTL4 SRCL4
IPv6 Fields:
MPLS Fields:
```

```
L2GRE L2 Fields:
L2GRE L3 Fields:
VXLAN L2 Fields:
VXLAN L3 Fields:
```

```
root@DMF-SWITCH-1:~# ofad-ctl bshell getreg RTAG7_HASH_CONTROL_L2GRE_MASK_A
RTAG7_HASH_CONTROL_L2GRE_MASK_A.ipipe0[1][0x6a001900]=0xffffffff
: <L2GRE_TUNNEL_GRE_KEY_MASK_A=0xffffffff>
```

```
root@DMF-SWITCH-1:~# ofad-ctl bshell getreg RTAG7_HASH_CONTROL_L2GRE_MASK_B
RTAG7_HASH_CONTROL_L2GRE_MASK_B.ipipe0[1][0x6a001a00]=0xffffffff
: <L2GRE_TUNNEL_GRE_KEY_MASK_B=0xffffffff>
```

```
root@s5248f-1:~# ofad-ctl bshell getreg RTAG7_L2GRE_PAYLOAD_L2_HASH_FIELD_BMAP
RTAG7_L2GRE_PAYLOAD_L2_HASH_FIELD_BMAP.ipipe0[1][0x6a001b00]=0: <
L2GRE_PAYLOAD_L2_BITMAP_B=0,L2GRE_PAYLOAD_L2_BITMAP_A=0>
```

```
root@s5248f-1:~# ofad-ctl bshell getreg RTAG7_L2GRE_PAYLOAD_L3_HASH_FIELD_BMAP
RTAG7_L2GRE_PAYLOAD_L3_HASH_FIELD_BMAP.ipipe0[1][0x6a001c00]=0: <
L2GRE_PAYLOAD_L3_BITMAP_B=0,L2GRE_PAYLOAD_L3_BITMAP_A=0>
root@DMF-SWITCH-1:~#
```



**Note:** The L2GRE\_KEY offset is the same as the SRCL4 and DSTL4 offset in hardware. Hence, the hardware requires setting SRCL4 and DSTL4 hash fields and the L2GRE\_KEY hash field to hash the packets using the L2GRE\_KEY.

## 7.1.5 VxLAN Hashing

VxLAN hashing enables hashing on a VxLAN payload, including hashing on the Inner L3 Source IP, Inner L3 Destination IP, Inner L2 Source MAC, and inner L2 Destination MAC. This only applies to terminated cases.

Symmetric hashing works with VxLAN packet Inner L3 Source IP/Destination IP, Inner L4 Source Port/Destination Port, and Outer L3 Source IP/Destination IP.



**Note:** VxLAN hashing applies to switches running SWL OS.

### CLI Configuration

VxLAN hashing includes hashing on L2 and L3 and the setting of at least one parameter enabled under the switch construct on Controller CLI:

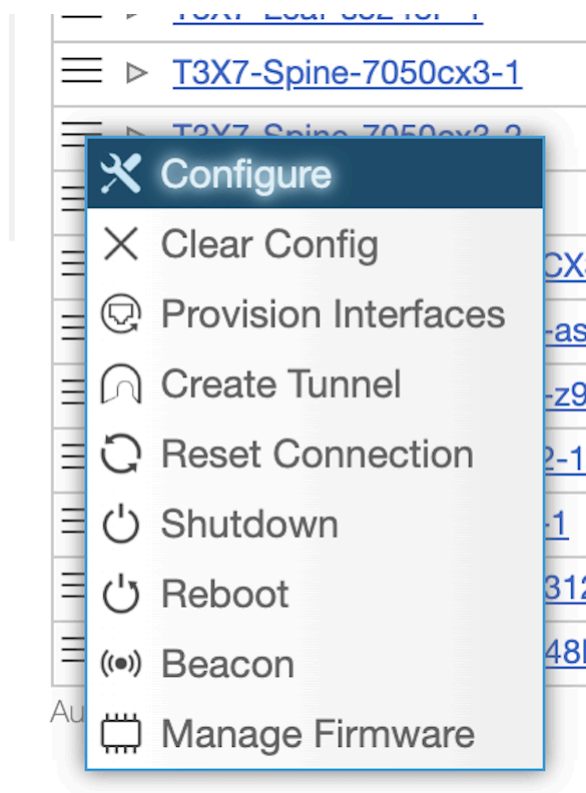
```
# lag-enhanced-hash
hash vxlan inner-l2 dst-mac
hash vxlan inner-l3 dst-ip
```

### UI Configuration



1. Configure the VxLAN Hashing in the UI for a switch in the **Fabric > Switches** page using the table row menu action **Configure** option.

**Figure 7-6: Fabric Switch Configure Menu**



2. In the **LAG Enhanced Hash** step, configure the following fields depending on your requirements:
  - **L2 VxLAN Inner L2 fields**
  - **VxLAN Inner L3 fields**



**Note:**

- L2 GRE Key is not supported for VxLAN hash fields.
- Cannot simultaneously specify enhanced hash for L2 GRE Inner L2 and Inner L3.

- Cannot simultaneously specify enhanced hash for VxLAN Inner L2 and Inner L3.

**Figure 7-7: Configure Switch LAG Enhanced Hash**

3. Click the **Submit** button to save the configuration.

### CLI Commands

Use the following CLI commands to verify settings and to troubleshoot any issues that may arise.

```
# show lag-enhanced-hash
```

Use the following commands to troubleshoot this feature. For example, when the hashing happens on VxLAN payload inner L3 Src IP.

```
root@DMF-SWITCH-1:~# root@mrv1:~# ofad-ctl gt PORT_CHANNEL_ENHANCED_HASH_FIELD
Hash Field Configs:
-----
Symmetric Hashing:
Disabled
L2 Fields:
IPv4 Fields:
IPv6 Fields:
MPLS Fields:
L2GRE L2 Fields:
L2GRE L3 Fields:
VXLAN L2 Fields:
```

```

VXLAN L3 Fields:
IP4SRC_LO IP4SRC_HI

root@mrv1:~# ofad-ctl bshell getreg RTAG7_HASH_CONTROL_4
RTAG7_HASH_CONTROL_4.ipipe0[1][0x6a000700]=3:
<VXLAN_PAYLOAD_HASH_SELECT_B=1,VXLAN_PAYLOAD_HASH_SELECT_A=1,DISABLE_H
ASH_VXLAN_B=0,DISABLE_HASH_VXLAN_A=0>

root@mrv1:~# ofad-ctl bshell getreg RTAG7_VXLAN_PAYLOAD_L2_HASH_FIELD_BMAP
RTAG7_VXLAN_PAYLOAD_L2_HASH_FIELD_BMAP.ipipe0[1][0x6a001d00]=0:<
VXLAN_PAYLOAD_L2_BITMAP_B=0,VXLAN_PAYLOAD_L2_BITMAP_A=0>

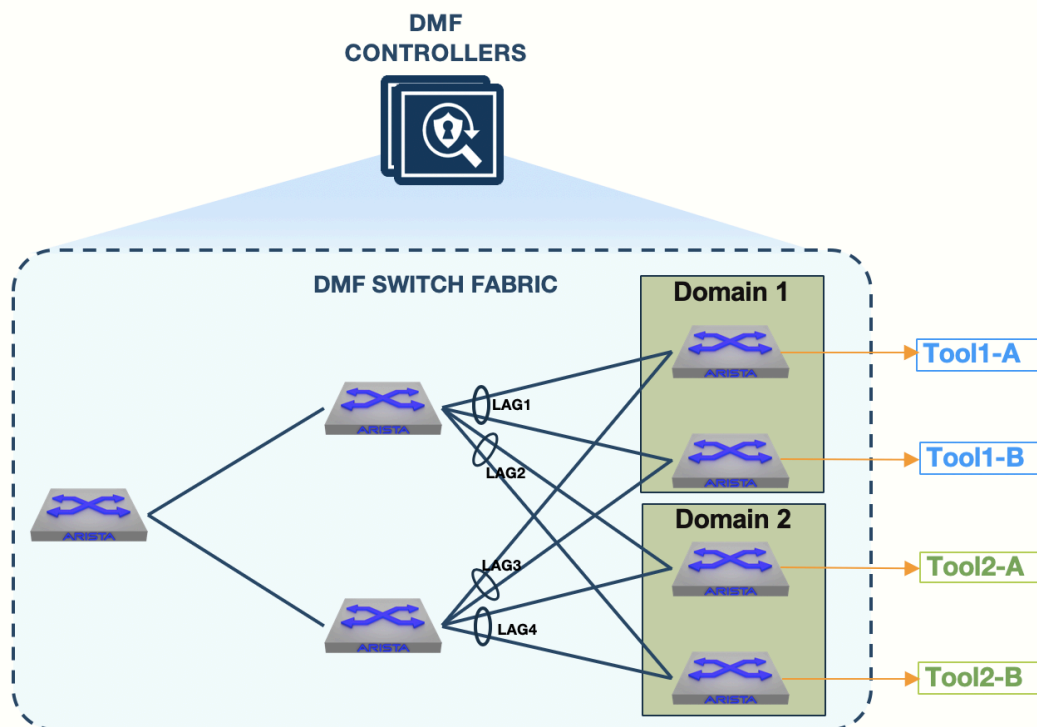
root@mrv1:~# ofad-ctl bshell getreg RTAG7_VXLAN_PAYLOAD_L3_HASH_FIELD_BMAP
RTAG7_VXLAN_PAYLOAD_L3_HASH_FIELD_BMAP.ipipe0[1][0x6a001e00]=0x1800c00
: <VXLAN_PAYLOAD_L3_BITMAP_B=0xc00,VXLAN_PAYLOAD_L3_BITMAP_A=0xc00>

```

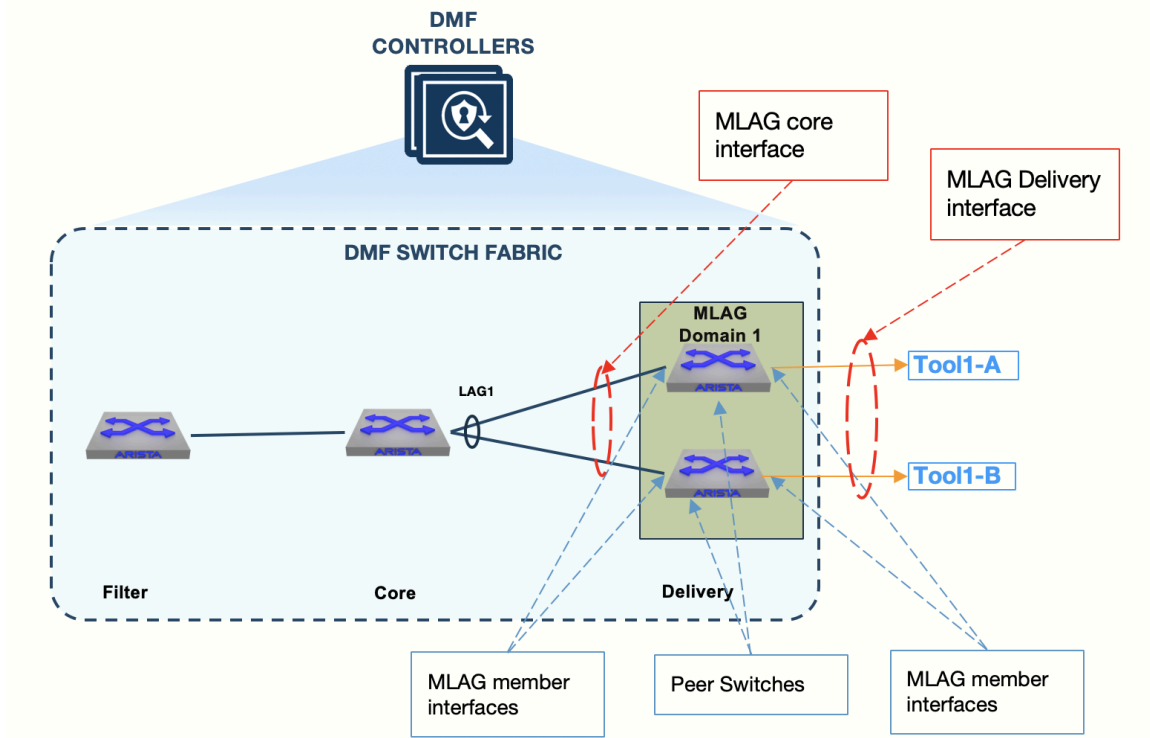
## 7.2 Pseudo Multi-Chassis Link Aggregation

Currently in DMF, we support Link Aggregation Groups (LAGs) that allow 2 or more physical interfaces on the same DMF switch to be aggregated into 1 logical interface to increase the aggregate bandwidth and provide link redundancy against link failure. This feature works well if all the tools are connected to the same DMF delivery switch, which is typically the case when customer tools are co-located in the same physical location. However, in cases where tools are located in different data centers or physical locations, where a single DMF switch cannot connect to all the tools, load-balancing across two DMF delivery switches is required.

A pseudo Multi-chassis Link Aggregation Group (MLAG) provides delivery switch redundancy where each delivery switch is connected to a multi-homed tool. With MLAG, traffic is hashed on the upstream DMF switch across two active-active links toward the delivery switches. In case one of the switches fails, the traffic will be switched over to the healthy switch.



## 7.3 MLAG Components



- **MLAG Domain:** An MLAG domain is a logical grouping of two delivery switches that will participate in an MLAG.
- **Peer Switch:** Member switches added into the MLAG domain.
- **MLAG Interface:** An MLAG interface, configured under the MLAG domain, is a logical binding of two physical interfaces or LAG interfaces, one from each peer switch.
- **Core MLAG Link:** A fabric-facing MLAG link. A core switch LAG interface, whose members connect to the two peer switches participating in the MLAG domain.
- **Delivery MLAG Link:** An MLAG interface that is assigned the delivery interface role. This interface is used in a policy as a delivery interface.
- **MLAG Member Interface:** A physical interface or a LAG interface added into an MLAG interface.
- **DMF Policy:** A user-configured DMF policy that contains at least one MLAG delivery interface.
- **Dynamic MLAG Domain Policy:** Dynamically configured policies that follow the naming convention `_mlag_<DMF- policy>_<DeliverySwitch>`. For one user-configured MLAG policy, a policy that uses at least one MLAG delivery interface, two dynamic MLAG domain policies are created, one for each peer switch.

## 7.4 MLAG Limitations

- An MLAG domain cannot have more than two switches.
- A switch can only be a part of one MLAG domain.
- An MLAG interface can only have two member interfaces.
- An MLAG interface can only have one interface (physical interface or LAG interface) from each peer switch.
- Tunnel interfaces are not supported as members in MLAG interface configuration.

## 7.5 Configuring an MLAG via the CLI

To configure an MLAG, use the following steps:

1. Configure an MLAG domain by specifying an alias, and add peer switches that will be participating in the MLAG.

```
Controller-1(config)# mlag-domain MLAG-Domain1
Controller-1(config-mlag-domain)# peer-switch DeliverySwitch-1
Controller-1(config-mlag-domain)# peer-switch DeliverySwitch-2
```

2. Configure the core MLAG interface.

```
Controller-1(config-mlag-domain)# mlag-interface MLAG-Core-Intf
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-1
interface ethernet50
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-2
interface ethernet50
```

The above MLAG interface configuration selects one physical interface from each peer switch added into the MLAG domain. This MLAG interface is fabric-facing, which means that **ethernet50** of **DeliverySwitch-1** and **ethernet50** of **DeliverySwitch-2** are connected to the DMF core switch, where traffic hashing is performed.

3. Configure the core LAG interface, a LAG interface on the core switch. The members of the LAG interface are connected to the peer switches in the MLAG domain. This configuration ensures that the traffic will be hashed toward the two connected delivery switches.

```
Controller-1(config)# switch CoreSwitch-1
Controller-1(config-switch)# lag-interface Core-LAG
Controller-1(config-switch-lag-if)# member ethernet10
Controller-1(config-switch-lag-if)# member ethernet20
```

4. Configure the delivery MLAG interface by specifying an interface alias and selecting one member from each delivery switch.

```
Controller-1(config-mlag-domain-if)# mlag-interface MLAG-Del-Intf
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-1
interface ethernet1
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-2
interface ethernet1
Controller-1(config-mlag-domain-if)# role delivery interface-name MLAG-Tool-1
```

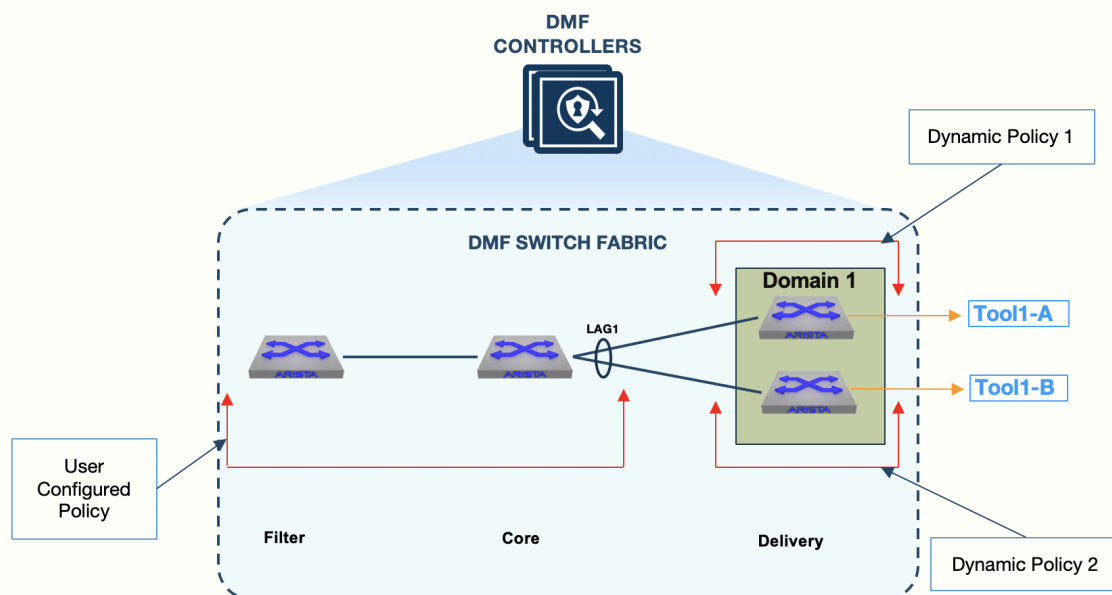
The above MLAG interface configuration selects one physical interface from each peer switch added into the MLAG domain. The members of this MLAG interface, **ethernet1** of **DeliverySwitch-1** and **ethernet1** of **DeliverySwitch-2**, are connected to multi-homed tools. Note that unlike the core MLAG interface, the delivery MLAG interface is assigned the delivery role and its interface name is configured, so that it can be used in DMF policies as a delivery interface.

5. Configure a DMF policy by following the procedure shown below:

```
Controller-1(config)# policy Policy-1
Controller-1(config-policy)# action forward
Controller-1(config-policy)# 1 match any
Controller-1(config-policy)# filter-interface Filter-1
Controller-1(config-policy)# delivery-interface MLAG-Tool-1
```

The above policy is configured using the **MLAG-Tool-1** interface configured in **Step 4**. Configuring the policy to use an MLAG delivery interface will result in two dynamic policies, one for each peer switch. Refer to the following topology for the policy breakdown.

**Figure 7-8: MLAG Policy Breakdown**



As seen in the topology above:

- The user-configured policy delivers traffic from the filter switch to the core switch LAG interface.
- Dynamic Policy 1 delivers traffic to delivery **switch 1**.
- Dynamic Policy 2 delivers traffic to delivery **switch 2**.

The following output displays the three policies as configured on the DMF controller:

```
Controller-1(config)# show policy
# Policy Name      Action  Runtime Status Type      Priority Overlap Priority Push VLAN Filter BW Delivery BW Post Match Filter Traffic Delivery Traffic Services
-----
1 MLAG-policy      forward installed      Configured 100      0          3          25Gbps 80Gbps 1.22Gbps 1.23Gbps
2 _mlag_Policy-1_DeliverySwitch-1 forward installed      Dynamic 100      1          2          40Gbps 1Gbps 612Mbps 612Mbps
3 _mlag_Policy-1_DeliverySwitch-2 forward installed      Dynamic 100      1          1          40Gbps 1Gbps 613Mbps 613Mbps
Controller-1(config)#
```

Below are the details for each policy:

#### Policy: **Policy-1** Interfaces

- Filter Interface(s) section lists the filter interface configured for the policy, **Policy-1**.
- Core Interface(s) section lists the interfaces that connect the filter switch and the core switch selected for the policy.
- MLAG Core Interface(s) section displays the core LAG interface that hashes the traffic towards the peer switches.
- MLAG Delivery Interface(s) section lists the delivery MLAG interface members.

#### Policy: **\_mlag\_Policy-1\_DeliverySwitch-1** Interfaces

- Filter Interfaces(s) section lists the dynamically configured interface name on **DeliverySwitch1** to which the core switch is connected.
- MLAG Delivery Interface(s) section lists the delivery MLAG interface member on **DeliverySwitch1**.

#### Policy: **\_mlag\_Policy-1\_DeliverySwitch-2** Interfaces

- Filter Interfaces(s) section lists the dynamically configured interface name on **DeliverySwitch2** to which the core switch is connected.

- MLAG Delivery Interface(s) section lists the delivery MLAG interface member on **DeliverySwitch2**.

## 7.6 MLAG Link Discovery

Link Layer Discovery Protocol (LLDP) is used to discover MLAG links. When the DMF controller receives an LLDP message, it looks for the switch and interface names. If the switch is a part of an MLAG domain, and the reported interface corresponds to the MLAG interface, then it is classified as an MLAG link.

```
Controller-1(config)# show link all link-type mlag-member
-----
# Active State Src switch          Src IF Name Dst switch          Dst IF Name Link Type   Since
-----
1 active      CoreSwitch-1  ethernet10  DeliverySwitch-1  ethernet50  mlag-member  2022-11-11 21:54:28 UTC
2 active      CoreSwitch-1  ethernet20  DeliverySwitch-2  ethernet50  mlag-member  2022-11-11 21:54:28 UTC
3 active      DeliverySwitch-1 ethernet50  CoreSwitch-1      ethernet10  mlag-member  2022-11-11 21:54:28 UTC
4 active      DeliverySwitch-2 ethernet50  CoreSwitch-1      ethernet20  mlag-member  2022-11-11 21:54:13 UTC
Controller-1(config)#
```

## 7.7 Configure MLAG via GUI

To configure an MLAG domain from the GUI, go to the **Fabric > MLAGs** tab.

Figure 7-9: Fabric -> MLAGs page

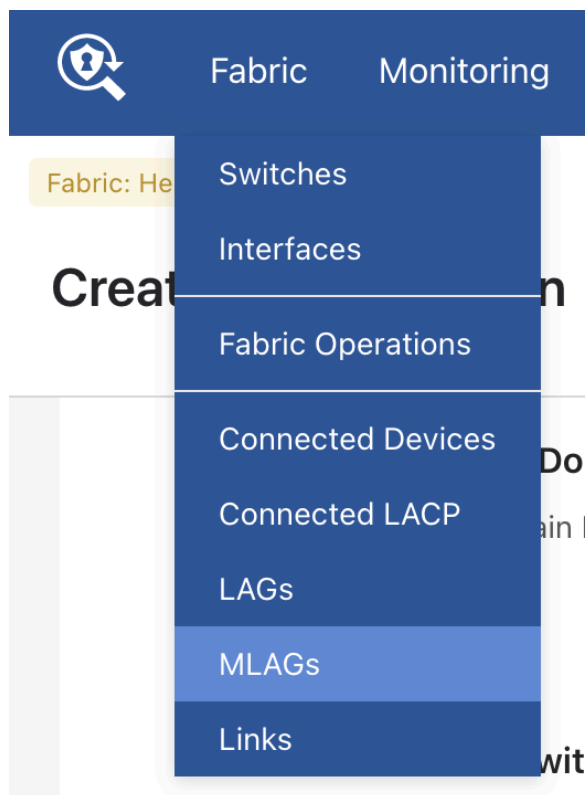



Figure 7-10: Create MLAG page



Click on **Create MLAG Domain** and enter the following:

- **Domain Name:** Enter the MLAG domain alias.
  - **Peer Switch 1:** From the drop-down, select the first switch that will be participating in the MLAG domain.
  - **Peer Switch 2:** From the drop-down, select the second switch that will be participating in the MLAG domain.
- **MLAG Interfaces:** Enter an alias for the fabric-facing MLAG interface. This interface connects the core switch to the peer switches in the MLAG domain.
  - **Peer Switch 1 and Peer Switch 2:** After selecting peer switches under the domain name, the peer switches under the MLAG interface will automatically be selected.
  - **Interface 1:** Select the member interface that connects the core switch to **DeliverySwitch-1**
  - **Interface 2:** Select the member interface that connects the core switch to **DeliverySwitch-2**.
- **MLAG Delivery Interfaces:** Enter an alias for each MLAG delivery interface.
  - **DMF Interface Name:** Enter the DMF interface name for the MLAG delivery interface. This alias will be used to identify the delivery interface while configuring the DMF policy.
  - **Strip VLAN on Egress:** Select the strip VLAN configuration for the MLAG delivery interface
  - **Peer Switch 1 and Peer Switch 2:** These will be automatically selected based on the peer switches selected under the domain name.
  - **Interface 1:** Select the member interface on Peer **Switch 1**.
  - **Interface 2:** Select the member interface on Peer **Switch 2**.

Click on the **Create** button to save the above configuration.

**Figure 7-11: MLAG Domain State**

Multi-Chassis LAGs (MLAGs) Refresh

View and configure MLAG domains and interfaces

Delete MLAG Domain(s) Edit MLAG Domain Create MLAG Domain				
Contains...	Add Filters	Domain Name	Clear current filters	
Domain Name	Peer Switches	MLAG Interfaces	MLAG Delivery Interfaces	Status
+ Domain1	DMF-D1 and DMF-D2	2	1	Active

1-1 of 1 items < 1 > 10 / page

**Figure 7-12: MLAG Domain Expanded View**

Multi-Chassis LAGs (MLAGs) Refresh

View and configure MLAG domains and interfaces

Delete MLAG Domain(s) Edit MLAG Domain Create MLAG Domain				
Contains...	Add Filters	Domain Name	Clear current filters	
Domain Name	Peer Switches	MLAG Interfaces	MLAG Delivery Interfaces	Status
- Domain1	DMF-D1 and DMF-D2	2	1	Active

**MLAG Interfaces**

Name	Member Interfaces	Status
Intf-core-del	DMF-D1 / ethernet51 DMF-D2 / ethernet51	Active

**MLAG Delivery Interfaces**

Name	Dmf Name	Role	Member Interfaces	Strip Vlan on Egress	Status
Intf-del-tool	MLAG-Tool-1	Delivery	DMF-D1 / ethernet23 DMF-D2 / ethernet11	-	Active

The above screenshot displays the MLAG domain status. Click on the **+** button to expand the MLAG domain configuration and status of each MLAG interface.



## 7.8 Create MLAG Policy from GUI

To configure an MLAG policy, go to **Monitoring > Policies** page. Click on the **+** to add a new policy, and configure all four of the below attributes:

**Figure 7-13: Configure MLAG Policy**

Create Policy

Advanced Options ▾

1 Policy Details 2 Match Traffic 3 Traffic Sources 4 Destination Tools

- **Policy Details:** Assign a unique name to the policy.
- **Match Traffic:** Specify the traffic rule for the policy.
- **Traffic Source:** Select the filter interface (traffic source) for the policy.
- **Destination Tools:** Select an MLAG delivery interface for the policy using the DMF interface name **MLAG-Tool-1**.

## 7.9 Viewing Policy Statistics in the GUI

Once the MLAG policy is configured, it can be viewed under **Monitoring > Policies** along with the dynamic policies created as part of the MLAG policy.

**Figure 7-14: MLAG Policy**

	Name	Description	Runtime Status	Config Status	Type	Action	Active	Priority	Push VLAN
<input type="checkbox"/>	► <a href="#">_mlag_Policy-1_DM-F-D1</a>	—	✓ installed	✓ active - forward	Dynamic	NA	✓ Yes	100	10
<input type="checkbox"/>	► <a href="#">_mlag_Policy-1_DM-F-D2</a>	—	✓ installed	✓ active - forward	Dynamic	NA	✓ Yes	100	12
<input type="checkbox"/>	► <a href="#">Policy-1</a>	—	✓ installed	✓ active - forward	Configured	→ Forward	✓ Yes	100	11

To view the policy statistics, click on the MLAG policy menu → Monitor Interfaces. The following dialog box will open:

**Figure 7-15: Policy Interface Statistics**

Configured → Forward ✓ Yes 100

POLICY: Policy-1

☒ Bytes/Bits
 ☐ Received (Rx)

☐ Packets
 ☐ Transmitted (Tx)

☐ Dropped Packets
 ☒ Rx + Tx

☐ Zero Counters

DMF Interface	Switch/Interface	Byte Counter	Bit Rate
<b>F</b> Filter-1	DMF-F1/ethernet5	152,580,482,150	973.2 Mbit/s
<b>C</b> NA	DMF-F1/ethernet49	439,234,946,644	2.3 Gbit/s
<b>C</b> NA	core2-7050-cx3-32S/ethernet3	439,018,233,233	2.3 Gbit/s
<b>MD</b> MLAG-Tool-1	DMF-D2/ethernet11	83,486,008,527	496.4 Mbit/s
<b>MD</b> MLAG-Tool-1	DMF-D1/ethernet23	69,027,820,642	477.2 Mbit/s
<b>O</b> NA	core2-7050-cx3-32S/MLAG-Core2-LAG	153,316,313,500	978.5 Mbit/s

**Table 4: DMF Policy Statistics**

DMF Interfaces	Description
[F] - Filter DMF Interface: Filter-1	<b>Filter-1</b> is the filter interface selected for the policy.
[C] - Core DMF Interface: N/A	Link between the filter switch and the core switch (core switch interface).
[C] - Core DMF Interface: N/A	Link between the filter switch and the core switch (filter switch interface).
[O] - Other DMF Interface: N/A	LAG interface configured on the core switch.
[MD] - MLAG Delivery DMF Interface: MLAG-Tool-1	MLAG delivery interface member 1.
[MD] - MLAG Delivery DMF Interface: MLAG-Tool-1	MLAG Delivery interface member 2.

There are no DMF interface names associated with core interfaces and core switch LAG interfaces; therefore, non-filter and delivery interfaces show N/A for DMF interface names.

## 7.10 Viewing MLAG Links in the GUI

To view the MLAG links, go to **Fabric > Links > MLAG Member Links** tab.

**Figure 7-16: MLAG Member Links**

Links Last updated: a few seconds ago

Switch Links

LAG Links

MLAG Member Links

History

Contains...

Add Filters

Dest. Interface Number

Source Interface Number

Source Switch Name	Source Interface Name	Source Interface Number	Dest. Switch Name	Dest. Interface Name
DMF-D1	ethernet51	51	core2-7050-cx3-32S	ethernet7
DMF-D2	ethernet51	51	core2-7050-cx3-32S	ethernet8
core2-7050-cx3-32S	ethernet7	7	DMF-D1	ethernet51
core2-7050-cx3-32S	ethernet8	8	DMF-D2	ethernet51

The above screenshot shows the MLAG links established between the core switch and the peer switches that are part of the MLAG domain. The links are discovered via LLDP message exchange.

## 7.11 Using LAG Interfaces as Members in MLAG Interfaces

MLAG interface members can be physical interfaces or LAG interfaces to increase bandwidth. To add a LAG member to an MLAG interface, use the following procedure:

1. Configure the LAG interface on Peer Switch 1.

```
Controller-1(config)# switch DeliverySwitch-1
Controller-1(config-switch)# lag-interface LAG-peer-switch-1
Controller-1(config-switch-lag-if)# member ethernet11
Controller-1(config-switch-lag-if)# member ethernet12
```

2. Configure the LAG interface on Peer Switch 2.

```
Controller-1(config)# switch DeliverySwitch-2
Controller-1(config-switch)# lag-interface LAG-peer-switch-2
Controller-1(config-switch-lag-if)# member ethernet11
Controller-1(config-switch-lag-if)# member ethernet12
```

3. Add the configured LAG interfaces as members into the MLAG interface.

```
Controller-1(config)# mlag-domain Domain1
Controller-1(config-mlag-domain)# mlag-interface MLAG-LAG-Del-Intf
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-1
interface LAG-peer-switch-1
Controller-1(config-mlag-domain-if)# member switch DeliverySwitch-2
interface LAG-peer-switch-2
Controller-1(config-mlag-domain-if)# role delivery interface-name MLAG-LAG-Tool-1
```

4. Configure the DMF policy using the delivery interface MLAG-LAG-Tool-1.

```
Controller-1(config)# policy Policy-1
Controller-1(config-policy)# action forward
Controller-1(config-policy)# 1 match any
Controller-1(config-policy)# filter-interface Filter-1
Controller-1(config-policy)# delivery-interface MLAG-LAG-Tool-1
```



**Note:** Traffic will not hash toward tools if the core switch LAG has the same number of member interfaces as the LAG on the peer switches of MLAG delivery.

**Workaround:**

- Ensure that the hash fields on the core switch and MLAG peer switches are different, OR
- Ensure that the number of member interfaces in the LAG interface configured on the core switch is different from the number of members in the LAG interface configured on peer switches.

## 7.12 Overlapping Policies in LAGs

An overlapping policy is dynamically configured if two configured policies share at least one filter interface, and at least one of the delivery interfaces is different.

When two DMF policies that are configured to use an MLAG interface as a delivery interface overlap, the following policies are created:

1. **Policy-1** is configured to use the filter interface **Filter-1**, and the delivery interface **MLAG-Tool-1**.
2. **Policy-2** is configured to use the filter interface **Filter-1**, and the delivery interface **MLAG-Tool-2**.
3. The above two policies will result in an overlapping policy. An overlapping policy will be configured following the naming convention **\_Policy-1\_o\_Policy-2**.

Once the overlapping policy is calculated for the two user configured policies, two dynamic policies will be configured: one policy for each peer switch in the MLAG domain and one for each of the three policies listed above.

**Table 5: Dynamic Policies**

MLAG Dynamic Policy	Parent Policy	Delivery Switch/Peer switch
<i>_mlag_Policy-1_DeliverySwitch-1</i>	<i>Policy-1</i>	<i>DeliverySwitch-1</i>
<i>_mlag_Policy-1_DeliverySwitch-2</i>	<i>Policy-1</i>	<i>DeliverySwitch-2</i>
<i>_mlag_Policy-2_DeliverySwitch-1</i>	<i>Policy-2</i>	<i>DeliverySwitch-1</i>
<i>_mlag_Policy-2_DeliverySwitch-2</i>	<i>Policy-2</i>	<i>DeliverySwitch-2</i>
<i>_mlag_Policy-1_o_Policy-2_DeliverySwitch-1</i>	<i>_Policy-1_o_Policy-2</i>	<i>DeliverySwitch-1</i>
<i>_mlag_Policy-1_o_Policy-2_DeliverySwitch-2</i>	<i>_Policy-1_o_Policy-2</i>	<i>DeliverySwitch-2</i>

The following policies, **Policy-1** and **Policy-2**, share the same filter interface, **Filter-1**, but they are configured to use different delivery interfaces, **MLAG-Tool-1** and **MLAG-Tool-2**. No priority is configured; therefore, these policies will be using the same default priority.

### Policy-1 Configuration

```
policy Policy-1
action forward
delivery-interface MLAG-Tool-1
filter-interface Filter-1
1 match ip src-ip 200.200.0.0 255.255.255.0
```

### Policy-2 Configuration

```
policy Policy-2
action forward
delivery-interface MLAG-Tool-2
filter-interface Filter-1
1 match ip dst-ip 100.100.0.0 255.255.255.0
```

The above two policies will result in an overlapping policy.

```
Controller-1(config)# show policy
# Policy Name      Action Runtime Status Type      Priority Overlap Priy Push VLAN Filter BW Delivery BW Post Match Filt Traff Del Traffic Services
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 Policy-1         forward installed          Configured      100          0          1    25Gbps    80Gbps    314Mbps    315Mbps
2 Policy-2         forward installed          Configured      100          0          3    25Gbps    80Gbps    314Mbps    315Mbps
3 _Policy-1_o_Policy-2 forward installed          Dynamic         100          1          5    25Gbps    80Gbps    314Mbps    315Mbps
```

- **Policy-1:** User-configured policy to forward packets matching source IP **200.200.0.0/24** to **MLAG-Tool-1**.
- **Policy-2:** User-configured policy to forward packets matching destination IP **100.100.0.0/24** to **MLAG-Tool-2**.
- **\_Policy-1\_o\_Policy-2:** A dynamically configured overlapping policy with higher Overlap Priority to ensure that if a packet matches rules from both the policies (source IP of **200.200.0.1** and destination IP of **100.100.0.1**) it will be forwarded to both **MLAG-Tool-1** and **MLAG-Tool-2**.

Following are the dynamic policies configured for each delivery switch in the MLAG Domain.

### Policy-1 Dynamic Policies

```
# Policy Name      Action Runtime Status Type      Priority Overlap Priority Push VLAN Filter BW Delivery BW Post Match Fil Traff Del Traffic Services
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
4 _mlag_Policy-1_DeliverySwitch-1 forward installed          Dynamic         100          1          2    40Gbps    1Gbps    -
5 _mlag_Policy-1_DeliverySwitch-2 forward installed          Dynamic         100          1          6    40Gbps    1Gbps    -
```

- **\_mlag\_Policy-1\_DeliverySwitch-1:** MLAG dynamic policy for **Policy-1** for **DeliverySwitch-1**.
- **\_mlag\_Policy-1\_DeliverySwitch-2:** MLAG dynamic policy for **Policy-1** for **DeliverySwitch-2**.

### Policy-2 Dynamic Policies

#	Policy Name	Action	Runtime Status	Type	Priority	Overlap	Pri	Push	VLAN	Filter	BW	Delivery	BW	Post Match	Filter	Traffic	Del	Traffic	Services
6	_mlag_Policy-2_DeliverySwitch-1	forward	installed	Dynamic	100		1		9	40Gbps		10Gbps						-	
7	_mlag_Policy-2_DeliverySwitch-2	forward	installed	Dynamic	100		1		7	40Gbps		10Gbps						-	

- **\_mlag\_Policy-2\_DeliverySwitch-1**: MLAG dynamic policy for **Policy-2** for **DeliverySwitch-1**.
- **\_mlag\_Policy-2\_DeliverySwitch-2**: MLAG dynamic policy for **Policy-2** for **DeliverySwitch-2**.

### **\_Policy-1\_o\_Policy-2 Dynamic Policies**

The following policies have higher Overlap Priority as compared to the rest of the policies. This is to prioritize the overlapping traffic to be forwarded to both **DeliverySwitch-1** and **DeliverySwitch-2**.

#	Policy Name	Action	Runtime Status	Type	Priority	Overlap	Priority	Push	VLAN	Filter	BW	Delivery	BW	Post Match	Filter	Traffic	Delivery	Traffic	Services
8	_mlag_Policy-1_o_Policy-2_DeliverySwitch-1	forward	installed	Dynamic	100		3		4	40Gbps		11Gbps				1578Mbps		3158Mbps	
9	_mlag_Policy-1_o_Policy-2_DeliverySwitch-2	forward	installed	Dynamic	100		3		8	40Gbps		11Gbps				1578Mbps		3158Mbps	

- **\_mlag Policy-1\_o\_Policy-2\_DeliverySwitch-1**: MLAG dynamic policy for overlapping policy for **DeliverySwitch-1**.
- **\_mlag Policy-1\_o\_Policy-2\_DeliverySwitch-2**: MLAG dynamic policy for overlapping policy for **DeliverySwitch-2**.

## Tunneling Between Data Centers

---

This chapter describes how to establish Generic Routing Encapsulation (GRE) and/or Virtual Extensible LAN (VxLAN) tunnels between DMF switches in different locations or between a DMF switch and a third-party device.

### 8.1 Understanding Tunneling

DMF can forward traffic between two DMF switches controlled by the same controller over a tunnel. This feature can be used to extend a DMF deployment across multiple data centers or branch offices over networks connected by Layer-3 networks. This feature supports centralization or distribution of tools and taps across multiple locations when they cannot be cabled directly.



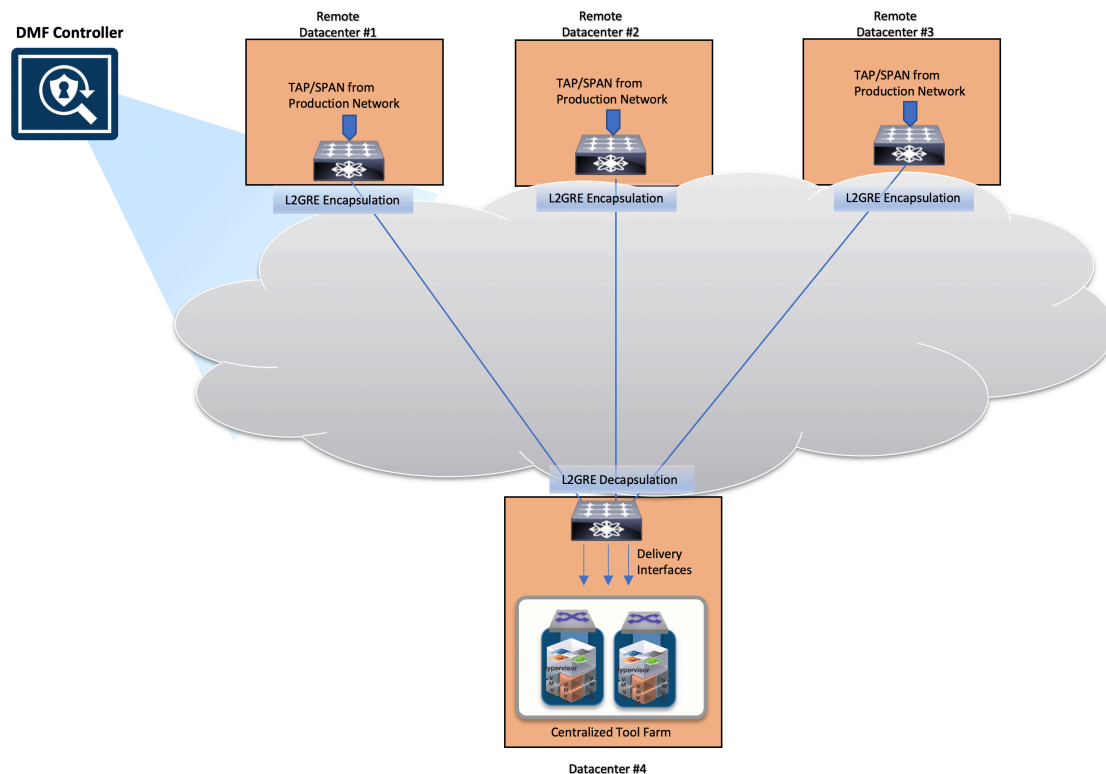
**Note:** For a list of the switches that support tunneling, refer to the ***DANZ Monitoring Fabric 8.4 Hardware Compatibility List***. The ***DANZ Monitoring Fabric 8.4 Verified Scale Guide*** indicates the number of tunnels supported by each supported switch (Verified Scalability Values/encap tunnels/decap tunnels).

When enabling tunneling between DMF switches, keep the following in mind:

- Connect switch ports in the main data center and the remote location to the appropriate WAN routers, and ping each interface to make sure that IP connectivity is established.
- Create tunnel endpoints and configure the tunnel attributes on each end of the tunnel.
- The **CRC Check** option must be enabled if tunneling is enabled, which it is by default. If CRC checking has been disabled, re-enable it before configuring a tunnel.
- In case of GRE tunnels, the optional **gre-key-decap** value on the receiving end must match the GRE key value of the sender. You can set multiple values on the same tunnel to decapsulate traffic with different keys.
- Multiple tunnels can be initiated from a single switch. For each tunnel (transmit-only or bidirectional), a separate encap-loopback-interface must be configured.

- The **loopback-mode** must be set to mac on the **encap-loopback-interface**.

**Figure 8-1: Connecting DMF Switches Using a Layer-2 GRE Tunnel**



**Note:** For EOS switches running DMF 8.4: L2GRE tunneling is supported on Arista 7280R3 switches only and subject to the following limitations:

- L2GRE tunnels are not supported on DMF 7280R and 7280R2 switches.
- Wildcard tunnel decapsulation is not supported.
- DSCP configuration is not supported.
- Traffic steering for traffic arriving on an L2GRE tunnel will only allow for matching based on inner src/dst IP, IP protocol, and inner L4 src/dst port.
- Packets may only be redirected to a single L2GRE tunnel.
- Packets may not be load-balanced across multiple L2GRE tunnels.
- Only IPv4 underlays in the default VRF are supported.
- Matching on inner IPv6 headers may not be supported.
- The maximum number of tunnels on EOS Jericho switches is 32.
- There is no bi-directional tunnel support. The parent/uplink router-facing interface is used for either encapsulation or decapsulation, but not simultaneously.
- When using tunnel-as-a-filter, there is no inner L3/L4 matching support immediately after decapsulation in the same switch pass. Using a loopback may potentially work around this limitation.
- VxLAN tunnels are currently NOT supported on 7280 switches.

## 8.2 Encapsulation Type

DMF supports VxLAN tunnel type along with Level-2 Generic Routing Encapsulation (L2GRE). Tunnel type is a per-switch configuration, which means that the switch pipeline can either be set to VxLAN or L2GRE. Once the switch pipeline is set, all tunnels configured on the switch will use the same tunnel type.

The encapsulation type can be configured in the GUI while adding a new switch into the DMF controller as shown in the figure below:

**Provision Switch**

1. Info ✓ 2. Clock ✓ 3. SNMP ✓ 4. SNMP Traps ✓ 5. Logging ✓ 6. TACACS ✓ 7. sFlow ✓ 8. LAG Enhanced Hash ✓

Name \*

MAC Address

Description

Admin Status \* ☒ Down ☒ Up

Management Interface \*

Tunnel Encapsulation Type \*

Password

The encapsulation type can be edited for an existing switch from the **Fabric > Switches > Configure Switch** page as shown in the figure below:

**Switches**

Filter table rows

Name	Connected	Quarantine Reason	Admin Status
Switch-1	✓	—	✓ Up
Switch-2	✗	—	✓ Up

Context Menu:

- Configure
- Clear Config
- Provision Interfaces

The encapsulation type can also be configured or edited from the CLI in configuration mode:

```
Ctrl-1(config)# switch Switch-1
Ctrl-1(config-switch)# tunnel-type
gre Select GRE as the tunnel type of the switch. (default selection)
vxlan Select VxLAN as the tunnel type of the switch.
```

The switch pipeline mode can be viewed from the CLI using the following command:

```
Ctrl-1(config)# show switch
# Switch Name      IP Address          State      Pipeline Mode
-----
1 Switch-1         fe80::d6af:f7ff:fe99:e2b0%9 connected 13-14-offset-match-push-vlan-vxlan
2 Switch-2         fe80::e6f0:4ff:fe69:6aee%9 connected 13-14-offset-match-push-vlan
3 Switch-3         fe80::e6f0:4ff:fe78:1ffe%9 connected 13-14-offset-match-push-vlan-vxlan
Ctrl-1(config)#
```



In the above CLI output, **Switch-1** and **Switch-3** are using the VxLAN tunnel type as seen in the Pipeline Mode column. **Switch-2** is using the L2GRE tunnel type.

## 8.3 Using Tunnels in Policies

Tunnels can be used as a core link, as a filter interface, or as a delivery interface. The most common use case is for linking multiple sites, in which case the tunnel is treated as a core link. If used as a core link, DMF automatically discovers the link as if it was a physical link and determines connectivity (link-state) in the same manner. If the tunnel goes down for any reason, DMF treats the failure as it would a physical link failure.

Another typical use case for the tunnel is as A filter interface to decapsulate L2 GRE/VxLAN tunneled production traffic, or a tunnel initiated by another DMF instance managed by a different DMF controller. Tunnel endpoint can also be used as a delivery interface to encapsulate filtered monitoring traffic to send to analysis tools or to send to another DANZ Monitoring Fabric managed by a different DMF controller.



**Note:** By default, sFlow and other Arista Analytics meta data cannot be generated for decapsulated L2 GRE/VxLAN tunneled production traffic on a tunnel interface configured as a filter interface. In order to generate this meta data, create a policy with a filter interface as a tunnel interface and send the decapsulated traffic to a MAC loopback port configured in a filter-and-delivery role. Now create a second policy with the filter interface as the MAC loopback port and the delivery interface going to the tools. The sFlow and meta data will now be generated for the decapsulated tunnel traffic.

## 8.4 Using the GUI to Configure a GRE Tunnel

To configure a VxLAN tunnel, complete the following steps:

1. Select **Fabric > Switch**.
2. On the **Switches** page, click the **Menu** control to the left of the switch or interface to include in the tunnel and select **Create Tunnel**.

Alternatively, tunnels can be configured from the **Fabric > Interfaces** page by clicking on the top **Menu Control > Create Tunnel** option. The system displays the dialog as shown in the figure below:

**Figure 8-2: Configure VxLAN Tunnel**

3. Complete the fields on this page, as described below.

- **Switch:** From the drop down, select the DMF switch.
- **Encapsulation Type:** The encapsulation type will automatically be selected based on the pipeline mode of the selected switch.
- **Name:** Name of the tunnel, beginning with the word `tunnel`.
- **Rate Limit** (Optional): Packets entering the tunnel can be rate limited to restrict the bandwidth usage of the tunnel. This can help ensure that a WAN link is not saturated with monitoring traffic being tunneled between sites. This setting is applicable on the tunnel encapsulation side.
- **Direction:** Direction can be bidirectional, transmit-only or receive-only. For bidirectional tunnels, the tunnel direction should be bidirectional on both ends. For uni-directional tunnels from remote to main datacenter, the tunnel direction is transmit-only on the remote datacenter switch and receive-only on the main data center switch.
- **Local IP:** Local IP address and subnet mask in CIDR format (/nn).
- **Gateway IP:** IP address of the default (next-hop) gateway.
- **Remote IP:** This is the IPv4 address of the other end (remote end) of the tunnel.
- **Parent Interface:** Physical port or port-channel interface associated with the tunnel. This is the destination interface for the tunnel.
- **Loopback Interface:** A physical interface on each switch with a transmit-only or a bidirectional tunnel endpoint. This physical interface will be used for tunnel encapsulation purposes, and cannot be used for any other DMF purpose, such as a filter, delivery, service, or core interface.
- **DSCP** (Optional): Mark the tunnel traffic with the specified DSCP value.

4. After configuring the appropriate options, click **Submit**.



**Note:** This procedure should be configured on both switches, at each end of the tunnel. The Auto VLAN mode must be set to **Push Per Policy** or **Push Per Filter Interface**.

## 8.5 Using the CLI to Configure a GRE Tunnel

To configure a GRE tunnel using the CLI, complete the following steps:

1. Connect switch ports (on remote and main datacenter) to their respective WAN routers and ensure that they can communicate via IP.
2. Enable tunneling on the DMF network by entering the following command from config mode:

```
controller-1(config)# tunneling
Tunneling is an Arista Licensed feature. Please ensure that you have
purchased the license
for tunneling before using this feature. enter "yes" (or "y") to continue:
yes
controller-1(config)#
```

3. Configure the MAC loopback mode, as shown in the following example:

```
controller-1(config)# switch DMF-CORE-SWITCH
controller-1(config-switch)# interface ethernet7
controller-1(config-switch-if)# loopback-mode mac
```

4. Create tunnel endpoints.

The following CLI example configures a bi-directional tunnel from **remote-dc1-filter-sw** to **main-dc-delivery-sw**:

```
!
switch remote-dc1-filter-sw
gre-tunnel-interface tunnel1
remote-ip 192.168.200.50
gre-key-decap 4097 === 4097 is the VPN key used for the tunnel ID
```

```
parent-interface ethernet6
local-ip 192.168.100.50 mask 255.255.255.0 gateway-ip 192.168.100.1
direction bidirectional encaps-loopback-interface ethernet38
!
switch main-dc-delivery-sw
gre-tunnel-interface tunnel1
remote-ip 192.168.100.50
gre-key-decap 4097 === 4097 is the VPN key used for the tunnel ID
parent-interface ethernet5
local-ip 192.168.200.50 mask 255.255.255.0 gateway-ip 192.168.200.1
direction bidirectional encaps-loopback-interface ethernet3
```

The following CLI example configures a uni-directional tunnel from **remote-dc1-filter-sw** to **main-dc-delivery-sw**:

```
!
switch remote-dc1-filter-sw
gre-tunnel-interface tunnel1
remote-ip 192.168.200.50
gre-key-decap 4097 === 4097 is the VPN key used for the tunnel ID
interface parent-interface ethernet6
local-ip 192.168.100.50 mask 255.255.255.0 gateway-ip 192.168.100.1
direction transmit-only encaps-loopback-interface ethernet38
!
switch main-dc-delivery-sw
gre-tunnel-interface tunnel1
remote-ip 192.168.100.50
gre-key-decap 4097 === 4097 is the VPN key used for the tunnel ID
parent-interface ethernet5
local-ip 192.168.200.50 mask 255.255.255.0 gateway-ip 192.168.200.1
direction receive-only
```

### 8.5.1 Using the CLI to Rate Limit the Packets on a GRE Tunnel

Packets entering the GRE tunnel can be rate limited to limit bandwidth usage by the tunnel and help ensure that a WAN link is not saturated with monitoring traffic being tunneled between sites. This setting is applicable on the tunnel encapsulation side.



**Note:** The minimum recommended value for rate limiting on the tunnel interface is **25** kbps. If you set a value below this, the switch will still set the rate limit value to **25** kbps.

```
switch DMF-CORE-SWITCH-1
gre-tunnel-interface tunnel1
direction bidirectional encaps-loopback-interface ethernet10
-----example truncated-----
interface ethernet10
rate-limit 1000
```

### 8.5.2 Using the CLI to View GRE Tunnel Interfaces

All CLI **show** commands for regular interfaces can be used for GRE tunnel interfaces, and the **show running-config** command can be used to view the configuration of tunnel interfaces.

Enter the **show tunnel** command to see the configuration parameters as well as runtime state for a tunnel interface.

```
controller-1# show tunnel
# Switch DPID      Tunnel Name Tunnel Status Direction Src IP      Dst IP      Parent Name Loopback Name
-----
1 DMF-CORE-SWITCH-1 tunnel1     ESTABLISHED_TUNNEL bidirectional 198.82.215.1 216.47.143.1 ethernet5:1 ethernet6
2 DMF-CORE-SWITCH-2 tunnel1     ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
3 DMF-CORE-SWITCH-2 tunnel2     ESTABLISHED_TUNNEL bidirectional 192.168.43.1 192.168.42.1 ethernet11:4 ethernet17
4 DMF-CORE-SWITCH-3 tunnel2     ESTABLISHED_TUNNEL bidirectional 192.168.42.1 192.168.43.1 ethernet6 ethernet33

controller-1# show tunnel switch DMF-CORE-SWITCH-2
# Switch DPID      Tunnel Name Tunnel Status Direction Src IP      Dst IP      Parent Name Loopback Name
```

```

1 DMF-CORE-SWITCH-2      tunnel1  ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
2 DMF-CORE-SWITCH-2      tunnel2  ESTABLISHED_TUNNEL bidirectional 192.168.43.1 192.168.42.1 ethernet11:4 ethernet17

controller-1# show tunnel switch DMF-CORE-SWITCH-2 tunnel1
# Switch DPID Tunnel Name Tunnel Status Direction Src IP Dst IP Parent Name Loopback Name
1 DMF-CORE-SWITCH-2 tunnel1 ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
controller-1#

```

## 8.6 Using the GUI to Configure a VxLAN Tunnel

To configure a VxLAN tunnel using the GUI, complete the following steps:

1. Select **Fabric > Switch**.
2. On the **Switches** page, click the **Menu** control to the left of the switch or interface to include in the tunnel and select **Create Tunnel**.

Alternatively, tunnels can be configured from the **Fabric > Interfaces** page by clicking on the top **Menu Control > Create Tunnel** option. The system displays the dialog as shown in the figure below:

**Figure 8-3: Configure VxLAN Tunnel**

3. Complete the fields on this page, as described below:
  - **Switch:** From the drop down, select the DMF switch.
  - **Encapsulation Type:** Encapsulation type will automatically be selected based on the pipeline mode of the selected switch.
  - **Name:** Name of the tunnel, beginning with the word **tunnel1**.
  - **Rate Limit** (Optional): Packets entering the tunnel can be rate limited to restrict bandwidth usage of the tunnel. This can help ensure that a WAN link is not saturated with monitoring traffic being tunneled between sites. This setting is applicable on the tunnel encap side.
  - **Direction:** bidirectional, transmit-only or receive-only. For bidirectional tunnels, tunnel-direction should be bidirectional on both ends. For uni-directional tunnels from remote to main datacenter, tunnel-direction is transmit only on the remote datacenter switch and the tunnel-direction is receive-only on the main data center switch.
  - **Local IP:** Local IP address and subnet mask in CIDR format (/nn).
  - **Gateway IP:** IP address of the default (next-hop) gateway.
  - **Remote IP:** This is the IPv4 address of the other end (remote end) of the tunnel.

- **Parent Interface:** Physical port or port-channel interface associated with the tunnel. This is the destination interface for the tunnel.
  - **Loopback Interface:** A physical interface on each switch with a transmit-only or a bidirectional tunnel endpoint. This physical interface will be used for tunnel encapsulation purposes, and cannot be used for any other DMF purpose, such as a filter, delivery, service, or core interface.
  - **DSCP (Optional):** Mark the tunnel traffic with the specified DSCP value.
4. After configuring the appropriate options, click **Submit**.



**Note:** This procedure should be configured on both switches, at each end of the tunnel. The Auto VLAN mode must be set to **Push Per Policy Or Push Per Filter Interface**.

## 8.7 Using the CLI to Configure a VxLAN Tunnel

To configure a VxLAN tunnel using the CLI, complete the following steps:

1. Connect switch ports (on remote and main datacenter) to their respective WAN routers and ensure that they can communicate via IP.
2. Enable tunneling on the DMF network by entering the following command from config mode:

```
controller-1(config)# tunneling
Tunneling is an Arista Licensed feature. Please ensure that you have
purchased the license
for tunneling before using this feature. enter "yes" (or "y") to continue:
yes
controller-1(config)#
```

3. Configure the MAC loopback mode, as shown in the following example:

```
controller-1(config)# switch DMF-CORE-SWITCH
controller-1(config-switch)# interface ethernet7
controller-1(config-switch-if)# loopback-mode mac
```

4. Create tunnel endpoints.

The following CLI example configures a bi-directional tunnel from *remote-dc1-filter-sw* to *main-dc-delivery-sw*:

```
!
switch remote-dc1-filter-sw
vxlan-tunnel-interface tunnel1
remote-ip 192.168.200.50
parent-interface ethernet6
local-ip 192.168.100.50 mask 255.255.255.0 gateway-ip 192.168.100.1
direction bidirectional encap-loopback-interface ethernet38
!
switch main-dc-delivery-sw
vxlan-tunnel-interface tunnel1
remote-ip 192.168.100.50
parent-interface ethernet5
local-ip 192.168.200.50 mask 255.255.255.0 gateway-ip 192.168.200.1
direction bidirectional encap-loopback-interface ethernet3
```

The following CLI example configures a uni-directional tunnel from *remote-dc1-filter-sw* to *main-dc-delivery-sw*:

```
!
switch remote-dc1-filter-sw
vxlan-tunnel-interface tunnel1
remote-ip 192.168.200.50
interface parent-interface ethernet6
```

```

local-ip 192.168.100.50 mask 255.255.255.0 gateway-ip 192.168.100.1
direction transmit-only encap-loopback-interface ethernet38
!
switch main-dc-delivery-sw
vxlan-tunnel-interface tunnell1
remote-ip 192.168.100.50
parent-interface ethernet5
local-ip 192.168.200.50 mask 255.255.255.0 gateway-ip 192.168.200.1
direction receive-only

```

### 8.7.1 Using the CLI to Rate Limit the Packets on a VxLAN Tunnel

Packets entering the VxLAN tunnel can be rate limited to limit bandwidth usage by the tunnel and help ensure that a WAN link is not saturated with monitoring traffic being tunneled between sites. This setting is applicable on the tunnel encapsulation side.



**Note:** The minimum recommended value for rate limiting on the tunnel interface is **25 kbps**. If you set a value below this, the switch will still set the rate limit value to **25 kbps**.

```

switch DMF-CORE-SWITCH-1
vxlan-tunnel-interface tunnell1
direction bidirectional encap-loopback-interface ethernet10
<snip>
interface ethernet10
rate-limit 1000

```

### 8.7.2 Using the CLI to View VxLAN Tunnel Interfaces

All CLI **show** commands for regular interfaces can also be used for tunnel interfaces. The **show running-config** command can be used to display the configuration of tunnel interfaces.

Enter the **show tunnel** command to see the configuration parameters as well as runtime state for a VxLAN tunnel interface.

```

controller-1# show tunnel
# Switch DPID Tunnel Name Tunnel Status Direction Src IP Dst IP Parent Name Loopback
Name
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 DMF-CORE-SWITCH-1 tunnell1 ESTABLISHED_TUNNEL bidirectional 198.82.215.1 216.47.143.1 ethernet5:1 ethernet6
2 DMF-CORE-SWITCH-2 tunnell1 ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
3 DMF-CORE-SWITCH-2 tunnel2 ESTABLISHED_TUNNEL bidirectional 192.168.43.1 192.168.42.1 ethernet11:4 ethernet17
4 DMF-CORE-SWITCH-3 tunnel2 ESTABLISHED_TUNNEL bidirectional 192.168.42.1 192.168.43.1 ethernet6 ethernet33
controller-1#
controller-1# show tunnel switch DMF-CORE-SWITCH-2
# Switch DPID Tunnel Name Tunnel Status Direction Src IP Dst IP Parent Name Loopback Name
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 DMF-CORE-SWITCH-2 tunnell1 ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
2 DMF-CORE-SWITCH-2 tunnel2 ESTABLISHED_TUNNEL bidirectional 192.168.43.1 192.168.42.1 ethernet11:4 ethernet17
controller-1#
controller-1# show tunnel switch DMF-CORE-SWITCH-2 tunnell1
# Switch DPID Tunnel Name Tunnel Status Direction Src IP Dst IP Parent Name Loopback Name
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 DMF-CORE-SWITCH-2 tunnell1 ESTABLISHED_TUNNEL bidirectional 216.47.143.1 198.82.215.1 ethernet11:3 ethernet5
controller-1#

```

## 8.8 Viewing or Modifying Existing Tunnels

To view or modify the configuration of an existing tunnel, use the **Fabric > Interfaces** option. To view the tunnel configuration, expand the interface. The tunnel configuration is displayed, as shown in the following figure:

**Figure 8-4: Tunnel Interfaces**

The screenshot shows the Arista DMF web interface. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', 'Security', and 'Profile'. The 'Interfaces' section is active, displaying a table of tunnel configurations. The table has columns for Switch, Switch DPID, Interface Name, Status, Tunnel, Device Name, Speed, and DMF Config. The selected interface, 'ARISTA-TUN-DEL', is expanded to show its detailed configuration. The configuration includes: Interface Name (ARISTA-TUN-DEL), Source IP (19.1.1.1/30), Status (Up), Loopback Interface (ethernet34), Destination IP (18.1.1.1), Admin Status (Up), Parent Interface (ethernet53/3), Gateway IP (19.1.1.2), Link Status (Up), Direction (transmit-only), GRE Key Decap (12345), Encapsulation Type (GRE), DSCP (—), and Rate Limit (8 kbit/s). The 'Core Links' section below shows no links.

The expanded row displays the status and other properties of the tunnel configured for the selected interface. Use the **Menu** control and select **Configure Tunnel** to modify the tunnel configuration. Select **Delete Tunnel** to remove the tunnel.

## 8.9 Using a Tunnel with User-defined Offsets

With an L2-GRE or VxLAN tunnel, matching traffic on a user-defined offset results in dropping interesting traffic. The offset calculation is thrown off by the tunnel header, and traffic that you want to select may be dropped. This behavior is due to the way switch hardware calculates the anchor and offset with respect to incoming packets. When the core link is a tunnel, the anchor and offset calculation is different when encapsulating packets compared to when decapsulating.

There are two ways to work around this issue:

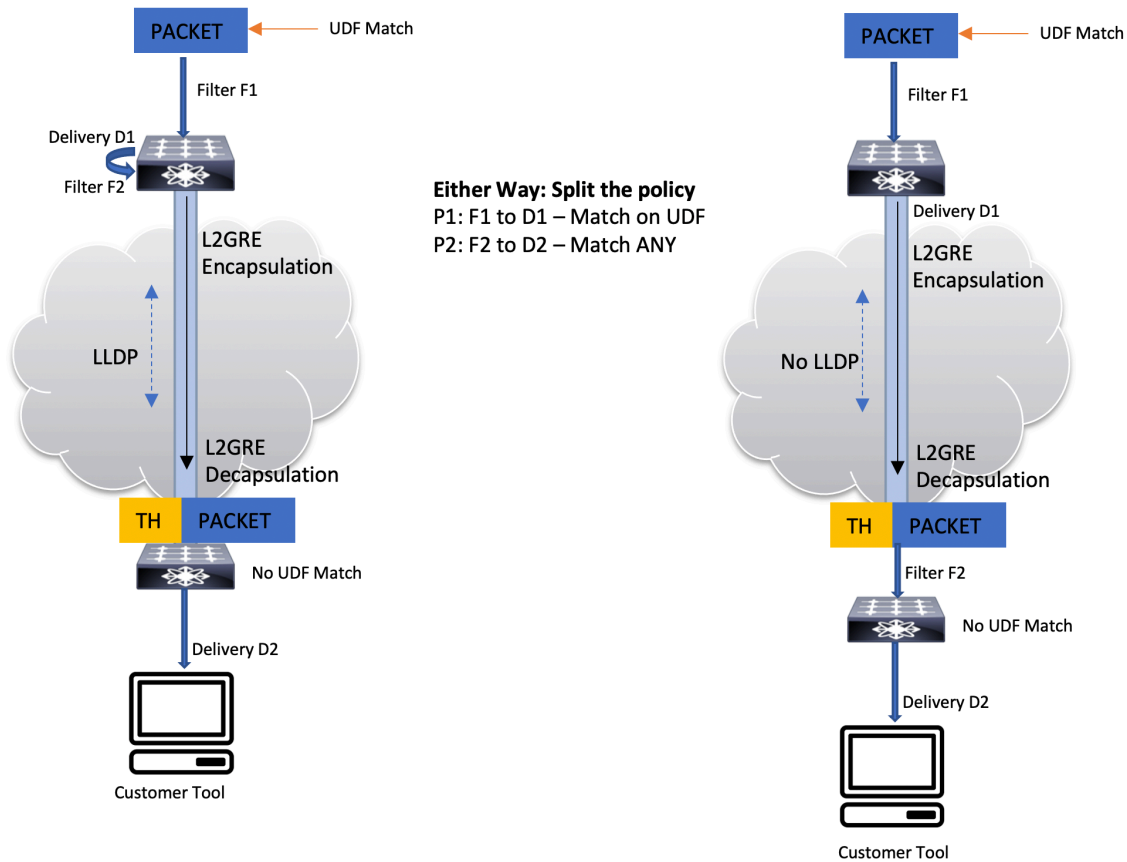
- Avoid matching on user-defined offsets on tunnel interfaces
- Avoid using a tunnel as a core link when matching on a user-defined offset

You can avoid matching on user-defined offsets when the ingress filter interface is a tunnel by filtering on the user-defined offset before the traffic enters a tunnel used as a filter interface. This preserves the LLDP messaging on the core tunnel link, but it requires an extra physical loopback interface on the encapsulating switch. Figure below illustrates both of these workarounds. In either case, a UDF match is applied to the ingress traffic on filter interface F. For example, the policy might apply a match at offset 20 after the start of the L4 header. In both workarounds, the policy has been split into two policies:

**P1: F to D1**, match on user-defined offset **P2: F1 to D**, match any.

In the example on the left, the ingress interface on the decapsulating switch, which is included in a core tunnel link, no longer has to calculate the user-defined offset. This solution preserves LLDP messages on the tunnel link, but requires an extra loopback interface.

**Figure 8-5: Using Tunnels with User-Defined Offsets**



In the example on the right, the tunnel endpoints are configured as filter and delivery interfaces. This solution avoids using the tunnel as a core link and does not require an extra physical loopback interface. However, LLDP updates are lost on the tunnel link.



## Integrating vCenter with the DANZ Monitoring Fabric

This chapter describes how to integrate VMware vCenter with the DANZ Monitoring Fabric and monitor Virtual Machines (VM) in the vCenter.

### 9.1 Overview

The DANZ Monitoring Fabric allows you to integrate with and monitor VMs in a VMware vCenter cluster. After integrating a vCenter with the DMF fabric, you can use DMF policies to select different types of traffic from specific VMs and apply managed services, such as deduplication or header slicing, to the selected traffic.

Currently, the following versions of VMware vCenter are supported for monitoring:


- vCenter Server 6.5.0
- vCenter Server 6.7.0
- vCenter Server 7.0.0
- vCenter Server 8.0.0

The DANZ Monitoring Fabric provides two options to monitor a VMware vCenter cluster:

- **Monitoring using span ports:** Using this method, you can monitor VMware vCenter clustering using a separate monitoring network. The advantage of this configuration is that it has no impact on the production network, and has a minimal effect on compute node CPU performance. However, in this configuration each compute node must have a spare NIC that can be used for monitoring traffic.

The following figure illustrates the topology used for local SPAN configuration:

**Figure 9-1: Mirroring on a Separate SPAN Physical NIC (SPAN)**



- **Monitoring using ERSPAN/L2GRE tunnels:** You can also use Remote SPAN (ERSPAN) to monitor VMs running on the ESX hosts within a vCenter instance integrated with DMF. This lets you monitor traffic to and from VMs anywhere in the network and does not require a dedicated physical interface card on the

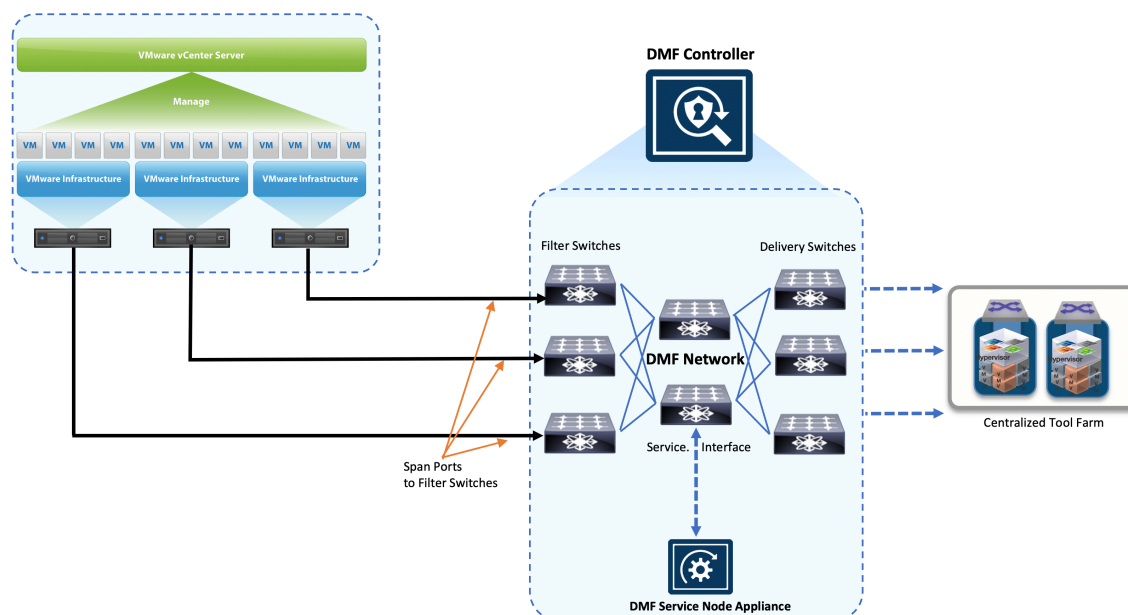
ESX host. However, ERSPAN can potentially affect network performance, especially if you are monitoring VMs connected to the DMF controller over WAN links or production networks with high utilization.

## 9.2 Using SPAN to Monitor VMs

This section describes the configuration required to integrate the DMF controller with one or more vCenter instances and to monitor traffic from VMs connected to the VMware vCenter after integration.

The following figure illustrates the topology required for integrating a vCenter instance with the monitoring fabric and for delivering the traffic selected by DMF policies to specified delivery ports connected to different monitoring tools.

**Figure 9-2: VMware vCenter Integration and VM Monitoring**



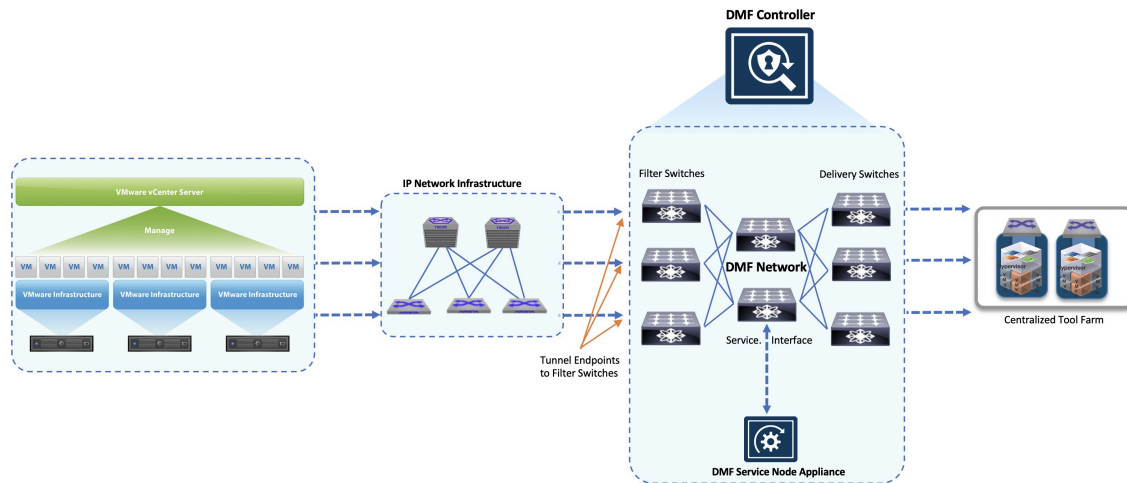
When integrated with vCenter, the DMF controller uses Link Layer Discovery Protocol (LLDP) to automatically identify the available filter interfaces that are connected to the vCenter instance.

## 9.3 Using ERSPAN to Monitor VMs

You can use Remote SPAN (ERSPAN) to monitor VMs running on the ESX hosts within a VMware vCenter instance integrated with the DANZ Monitoring Fabric. This lets you monitor traffic to and from VMs anywhere in the network and does not require a dedicated physical interface card on the ESX host. However, ERSPAN

can potentially affect network performance, especially if you are monitoring VMs connected to the DMF controller over WAN links or production networks with high utilization.

**Figure 9-3: Using ERSPAN to Monitor VMs**



The procedure for deploying ERSPAN is similar to SPAN but requires an additional step to define the tunnel endpoints used on the DMF network to terminate the ERSPAN session.

### 9.3.1 Configuration Summary for vCenter Integration

The following procedure summarizes the high-level steps required to integrate the vCenter and monitor traffic to or from selected VMs:

1. (For ERSPAN only) Define the tunnel endpoint.

Identify a fabric interface connected to the vCenter instance for the tunnel endpoint by entering the **tunnel-endpoint** command in config mode. To define the tunnel endpoint, refer to the **Defining a Tunnel Endpoint** section.

2. Provide the vCenter address and credentials.

The vSphere extension on the DMF controller discovers an inventory of VMs and the associated details for each VM.

3. Select the VMs to monitor on the DMF controller.

The DMF controller uses APIs to invoke the vSphere vCenter instance.

vSphere calls the DVS to create a SPAN session. The preferred option is to SPAN on a separate physical NIC. However, you can also use ERSPAN by tunneling to the remote interface.

4. Create policies in DMF to filter, replicate, process, and redirect traffic to tools.

When using tunnels with ERSPAN, DMF terminates the tunnels using the specified tunnel endpoint. A DMF policy for monitoring VM traffic using a SPAN session must include the required information regarding the vCenter configuration. All match conditions, including User-Defined ofFsets (UDFs) are supported.

The policy used for selecting VM traffic to monitor is similar to other DMF policies except that the filtering interfaces are orchestrated automatically (filter interfaces are auto-discovered cannot be specified manually). All managed-service actions are supported.

## 9.4 Defining a Tunnel Endpoint

You can predefine the tunnel endpoints to be used for creating tunnels when monitoring VMware vCenter traffic using either the GUI or the CLI.

## GUI Procedure

To manage tunnel endpoints in the GUI, select **Monitoring > Tunnel Endpoints**.

**Figure 9-4: Monitoring > Tunnel Endpoints**

Name	Switch	Interface	Gateway	IP Address	Mask
ERSPAN	superspine-as7816-1	ethernet7	172.27.1.2	172.27.1.1	255.255.255.0

This page lists the tunnel endpoints that are already configured and provides information about each endpoint.

To create a new tunnel endpoint, click the provision (+) control at the upper left corner of the **Tunnel Endpoints** table.

**Figure 9-5: Create Tunnel Endpoint**

Create Tunnel Endpoint

Name \*

Switch

Optional: use to filter interface choices below

Interface \*

Choose from a list of interfaces not already configured for DMF

Gateway \*

IP Address \*

Mask \*

Subnet Mask

RESET CANCEL SAVE

To create the tunnel endpoint, complete the following information on this dialog and click **Save**:

- **Name:** Type a descriptive name for the endpoint.
- **Switch:** Select the DMF switch from the selection list where the endpoint interface is configured.
- **Interface:** Select the interface from the selection list to use for the endpoint.
- **Gateway:** Type the address of the default gateway.
- **IP Address:** Type the IP address of the endpoint.
- **Mask:** Type the subnet mask for the endpoint.

## CLI Procedure

To configure a tunnel endpoint using the CLI, enter the **tunnel-endpoint** command from config mode using the following syntax:

```
controller-1(config)# tunnel-endpoint <name> switch <switch> <interface> ip-  
address <address> mask  
<mask> gateway <address>
```

For example, the following command defines **ethernet24** on **F-SWITCH-1** as a tunnel endpoint named **OSEP1**:

```
controller-1(config)# tunnel-endpoint ERSpan switch CORE-SWITCH ethernet7 ip-  
address 172.27.1.1  
mask 255.255.255.0 gateway 172.27.1.2
```

The IP address assigned to this endpoint is **172.27.1.1**, and the next hop address for connecting to the vCenter via ERSpan is **172.27.1.2**.

## 9.5 Wildcard Tunnels for VMware vCenter Monitoring

The current implementation of VMware vCenter creates one tunnel interface from every ESXi host to DMF. Using a wildcard tunnel on DMF for VMware vCenter reduces the number of tunnels created.

### Platform Compatibility

This feature is only compatible with switches that support wildcard tunneling. The following DMF supported switch models currently support wildcard tunnels:

- DCS-7050SX3-48YC8
- DCS-7050CX3-32S
- DCS-7050SX3-48YC8
- DCS-7050SX3-96YC8
- DCS-7050SX3-48YC12
- DCS-7260CX3-64E
- DCS-7260CX3-64
- DCS-7050TX3-48C8
- DCS-7050SX3-48C8
- Dell S5212F-ON
- Dell S5248F-ON
- Dell Z9264F-ON
- Dell S4112F-ON
- Dell S5232F-ON

### 9.5.1 Configuration

Configure wildcard tunnels using the CLI or the GUI.

### 9.5.2 Using the CLI to create Wildcard Tunnels

The CLI construct `wildcard-tunnels` is available as a configuration option when configuring a VMware vCenter in DMF, as shown below:

---

**Table 6: Commands**

<b>cluster</b>	Configure tunnel-endpoint for cluster
<b>default-tunnel-endpoint</b>	Configure tunnel endpoints
<b>description</b>	Provide a description for this vCenter
<b>hashed-password</b>	Set the vCenter password (to log into vCenter)
<b>host-name</b>	Set the vCenter host name
<b>mirror-type</b>	Set the vCenter vm monitoring mode
<b>mirrored-packet-length</b>	Set the mirrored packet length
<b>password</b>	Set the vCenter password (to log into vCenter)
<b>sampling-rate</b>	Set the packet sampling rate
<b>user-name</b>	Set the vCenter user name (to log into vCenter)
<b>vm-monitoring</b>	Enter <code>vm-monitoring config</code> submode
<b>wildcard-tunnels</b>	Enable wildcard tunnels

Enable wildcard tunnels by setting the above leaf parameter, as shown in the following example of vCenter configuration on the controller node.

```
Controller-C1(config)# vcenter VC1
Controller-C1(config-vcenter)# wildcard-tunnels
Controller-C1(config-vcenter)# show this
! vcenter
vcenter VC1
wildcard-tunnels
Controller-C1(config-vcenter)#
```

Similarly, disable wildcard tunnels by issuing the **no** command as shown below:

```
Controller-C1(config-vcenter)# show this
! vcenter
vcenter VC1
wildcard-tunnels
Controller-C1(config-vcenter)# no wildcard-tunnels
Controller-C1(config-vcenter)# show this
! vcenter
vcenter VC1
Controller-C1(config-vcenter)#
```

### Show Commands

There is no specific show command for wildcard tunnels, but it can be checked in the vCenter running config. In addition, the **show tunnels** command shows the tunnels created for the selected vCenter configuration with a wildcarded remote IP address.

### Troubleshooting

Verify error and warnings are clear using the **show fabric errors** and **show fabric warnings** commands. The **show tunnels** command displays tunnels created based on the vCenter configuration on the controller with a wildcarded remote IP address. To display tunnels programmed on the switch use the **show switch <name> table gre-tunnel** command.

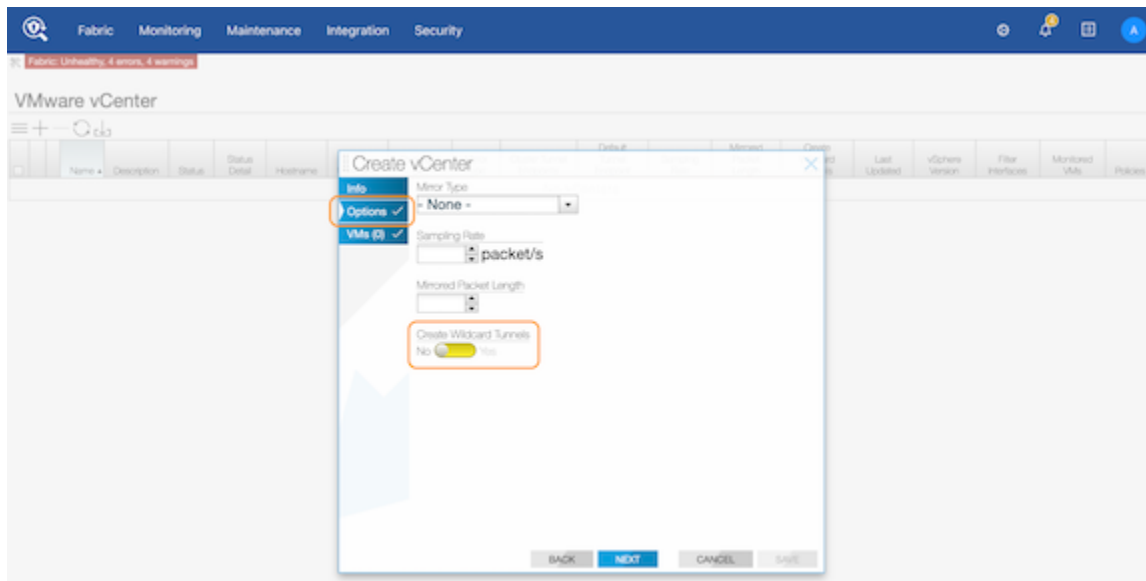
### 9.5.3 Using the GUI to create Wildcard Tunnels

Use the DMF GUI to create wildcard tunnels as outlined below.

Navigate to the **Integration > VMware vCenter** page.

As part of the **Options** step of the **Add/Edit vCenter** workflow, enable wildcard tunnels using the **Create Wildcard Tunnels** toggle input. By default, the feature is disabled, as shown below:

**Figure 9-6: VMware vCenter Create vCenter Options**



### 9.5.4 Limitations

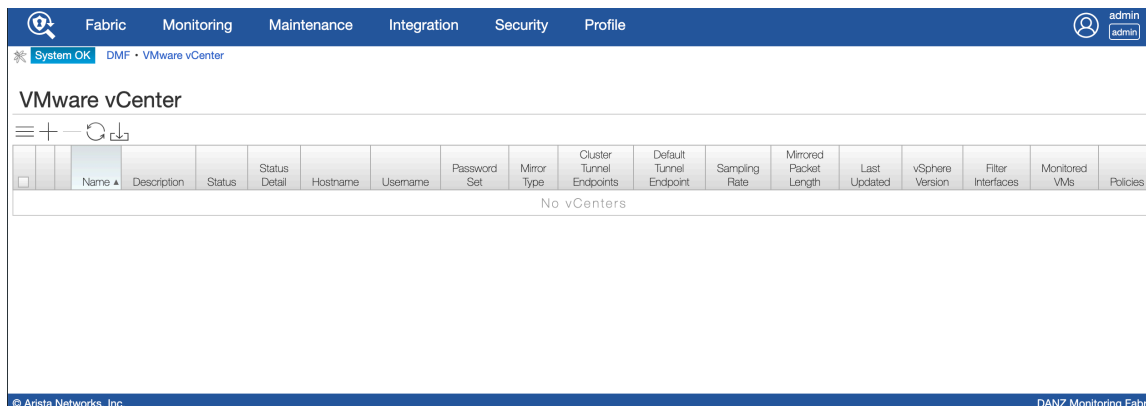
Wildcard tunnels are only supported on select Broadcom® switch ASICs, ensure your switch model supports this feature before configuring it for vCenter.

Please refer to the Platform Compatibility section for more information.

## 9.6 Using the GUI to Integrate a vCenter Instance

To integrate a vCenter instance with DMF to begin monitoring VMs, select **Integration > vCenter** from the **DMF** menu bar.

**Figure 9-7: Integration > vCenter**



This page displays information about the vCenter instances that have been integrated with DMF. To add a vCenter instance for integration with DMF, complete the following steps:

1. Click the provision control (+) at the top of the table.

**Figure 9-8: Create vCenter: Info**

The screenshot shows a web form titled "Create vCenter" with a sidebar on the left containing three tabs: "Info" (selected), "Options", and "VMs (0)". The "Info" tab is active, showing a form with the following fields and controls:

- Name:** A text input field containing "my-vcenter".
- Description:** A large text area for adding a description.
- Host Name:** A text input field with a note below it: "Up to 255 characters in length".
- Username:** A text input field with a note below it: "Up to 255 characters in length".
- Password:** A text input field containing "- Password -". Below it is a checkbox labeled "Show Password" and a note: "Up to 255 characters in length".

At the bottom of the form, there are four buttons: "BACK", "NEXT" (highlighted in blue), "CANCEL", and "SAVE".

2. Type an alphanumeric identifier for the vCenter instance, and (optionally) add a description in the fields provided.
3. Identify the vCenter host name to be integrated.
4. Enter the vCenter username and password for authenticating to the vCenter instance.

These credentials are used by the DMF controller when communicating with the vCenter host.



5. Click **Next**.**Figure 9-9: Create vCenter: Options (page 2)**

**Create vCenter**

Info ✓ Options ✓ VMs (0) ✓

Mirror Type: **Encapsulated Remote**

Cluster Tunnel Endpoints

Cluster	Tunnel Endpoint

No cluster tunnel endpoints

Default Tunnel Endpoint: **- Tunnel Endpoint -**

Configuration of a default tunnel endpoint is only allowed when the mirror type is set to encapsulated remote.

Sampling Rate: **packet/s**

Mirrored Packet Length: **1024**

BACK NEXT CANCEL SAVE

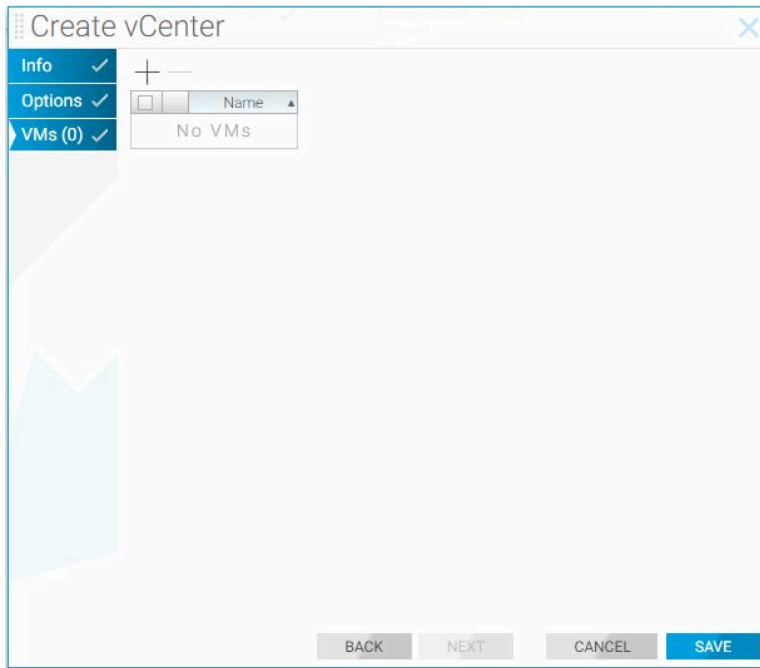
This page lets you define the mirror type as SPAN or ERSPAN. When you select ERSPAN, the following additional fields are provided to complete the ERSPAN configuration:

- **Cluster Tunnel Endpoints** (optional)
- **Default Tunnel Endpoint** (required)
- **Sampling Rate** (optional)
- **Mirrored Packet Length** (optional)

Use **Cluster Tunnel Endpoints** to specify a common tunnel endpoint for all the ESXi hosts in the cluster. Use **Default Tunnel Endpoint** to specify a common tunnel endpoint for all the ESXi hosts regardless of the cluster. When both cluster and default tunnel endpoints are configured, all hosts in clusters form tunnels using the cluster-specific configuration, and all the other hosts that are not a part of any cluster use the default configuration to form tunnels.

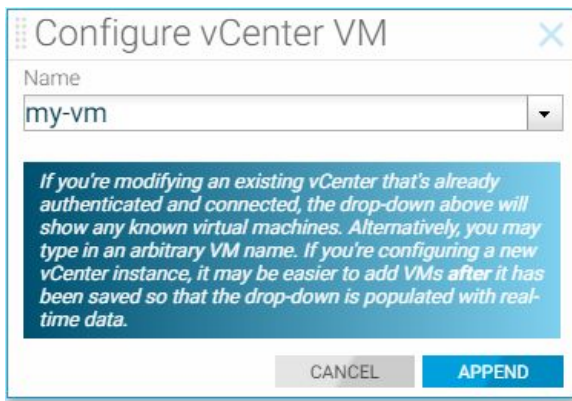
6. Click **Next**.

**Figure 9-10: Create vCenter/VMs**



7. To add a VM for monitoring, click the provision control (+).

**Figure 9-11: Configure vCenter VM**



You can select VMs from the selection list after the vCenter has been integrated and the VMs have been discovered, or you can manually add the VM hostname.

8. After identifying the VM to monitor, click **Append**.
9. On the **VMs** of the **Create vCenter** dialog, click **Save**.

### 9.6.1 Using a vCenter Instance as the Traffic Source in a DMF Policy

To identify a vCenter instance integrated with the DMF controller as the source of traffic for a DMF policy, click the **vCenters** tab on the **Feeds** page of the **Create Policy** dialog.

**Figure 9-12: Identifying Managed Service Feeds**

**Create Policy**

The following settings may affect the availability of some configuration options.

Switching Mode **Full Match**

☐ All Filter Interfaces

Filter Interfaces Filter Interface Groups OpenStacks **vCenters**

**vCenters**

Click + and - to include or exclude vCenters.

	Name	Description	Status	Status Detail	Hostname	Username	Password Set	Mirror Type	Cluster Tunnel Endpoints	Default Tunnel Endpoint	Sample Rate
No vCenters											

BACK NEXT RESET CANCEL SAVE

Select one or more instances from the table and click **Append Selected**.

To add an existing instance to the table, click the provision control (+) at the top of the table and complete the fields provided by the **Associate vCenter** dialog.

To define a new instance, click the provision control (+) at the top of the **Associate vCenters** table and complete the fields provided by the **Create vCenters** dialog. You can also use the **Integration > vCenters** option to complete the same configuration.

## 9.7 Using the CLI to Integrate a vCenter Instance

To configure the DMF controller for monitoring VMs on a locally connected vCenter instance, complete the following steps:

1. Add the vCenter instance details by entering the following commands:

```
controller-1(config)# vcenter vc-1
controller-1(config-vcenter)# host-name 10.8.23.70
controller-1(config-vcenter)# password 094e470e2a121e060804
```

```
controller-1(config-vcenter)# user-name root
```

2. Specify the mirror type by entering the following commands:

```
controller-1(config-vcenter)# mirror-type span | erspan  
controller-1(config-vcenter)# sampling-rate 60  
controller-1(config-vcenter)# mirrored-packet-length 60
```

The **sampling-rate** and **mirrored-packet-length** commands are optional.

3. For ERSPAN mirroring, a tunnel endpoint configuration is needed. Use the **cluster** command to specify a common tunnel endpoint for all the ESXi hosts in the cluster. Use the **default-tunnel-endpoint** command to specify a common tunnel endpoint for all the ESXi hosts regardless of the cluster. When both the **cluster** and **default-tunnel-endpoint** commands are used, all hosts in clusters form tunnels using the cluster-specific configuration and all the other hosts that are not a part of any cluster use the default configuration to form tunnels.

```
controller-1(config-vcenter)# default-tunnel-endpoint VCEP1  
controller-1(config-vcenter)# cluster <cluster-name> tunnel-endpoint  
                                  <tunnel-endpoint-name>
```

Using the **tab** auto-complete feature with the cluster suggests existing cluster names associated with the vCenter.

4. Add the VMs you want to monitor by entering the following commands:

```
controller-1(config-vcenter)# vm-monitoring  
controller-1(config-vcenter-vm-monitoring)# vm vm-2001  
controller-1(config-vcenter-vm-monitoring)# vm vm-2002
```

5. To view the vCenter configuration, enter the **show running-config vcenter** command as in the following example:

```
controller-1# show running-config vcenter  
! vcenter  
vcenter vc-1  
hashed-password 752a3a3211040e0200090409090611  
host-name 10.8.23.70  
mirror-type span  
mirrored-packet-length 60  
sampling-rate 60  
user-name administrator@vsphere.local  
!  
vm-monitoring  
vm vm-2001  
vm vm-2002
```

6. Configure the policies specifying the match rules and delivery interfaces.

```
controller-1(config)# policy dmf-policy-with-vcenter  
controller-1(config-policy)# action forward  
controller-1(config-policy)# vcenter vc-1  
controller-1(config-policy)# 1 match any  
controller-1(config-policy)# delivery-interface TOOL-PORT-03
```

The filter interfaces are automatically learned using LLDP. All the hosts belonging to vc-1 that have physical connections to DMF switches become the filter interfaces. If new connections are made later on (or existing connections are changed), policies will be recomputed with the new interfaces.

7. To view the automatically assigned filter interfaces, enter the **show running-config policy** command, as in the following example:

```
controller-1# show running-config policy dmf-policy-with-vcenter
```

```

! policy
policy dmf-policy-with-vcenter
action forward
delivery-interface TOOL-PORT-03
filter-interface vc-filter-1 origination vc-10-9-19-7--filter-interface
filter-interface vc-filter-3 origination vc-10-9-19-7--filter-interface
filter-vcenter vc-1
1 match any

```

8. To configure the DMF controller for monitoring VMs on a second locally connected vCenter instance, add the VMs you want to monitor and configure the DMF policies to specify the match rules and delivery interfaces.

```

(config)# vcenter vc-2
(config-vcenter)# host-name 10.8.23.71
(config-vcenter)# password 094e470e2a121e060804
(config-vcenter)# user-name root
(config-vcenter)# mirror-type span
(config-vcenter)# sampling-rate 60
(config-vcenter)# mirrored-packet-length 60
(config-vcenter)# vm-monitoring
(config-vcenter-vm-monitor)# vm vm-1001
(config-vcenter-vm-monitor)# vm vm-1002

```

9. Configure the policy for the second vCenter instance.

```

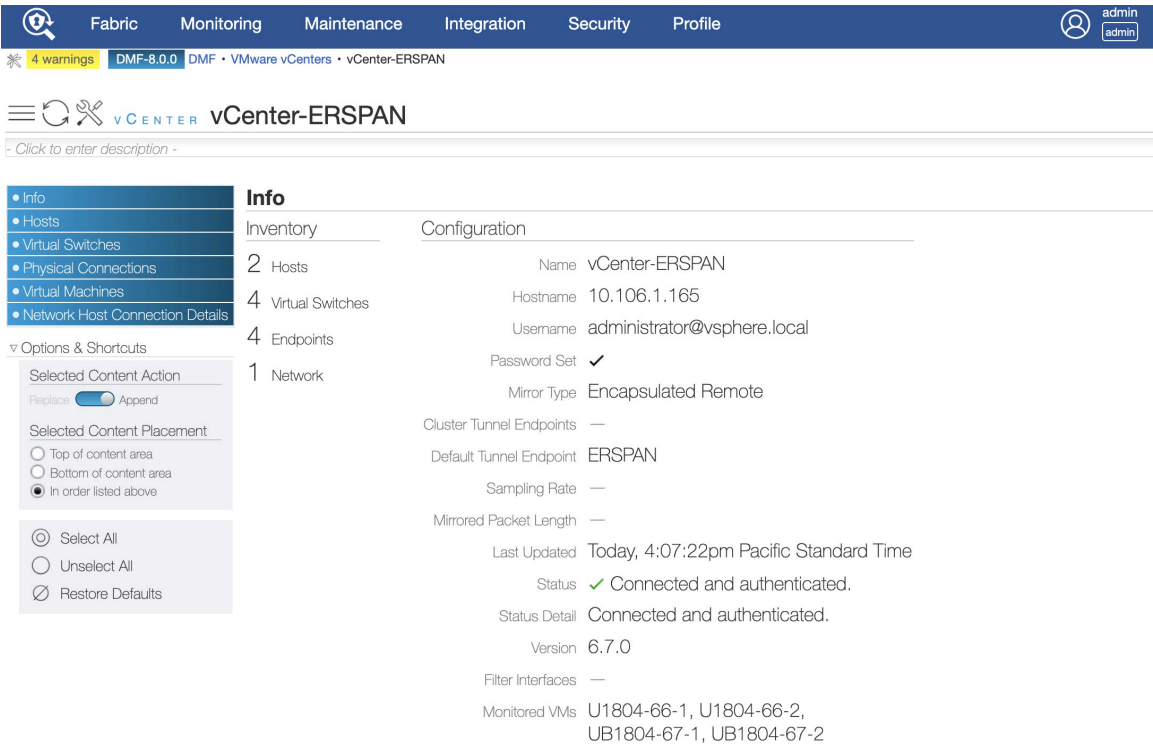
(config)# policy dmf-policy-with-vcenter-2
(config-policy)# vcenter vc-2
(config-policy)# 1 match any
(config-policy)# delivery-interface TOOL-PORT-02

```

# 9.8 Using the GUI to View vCenter Configuration

After integrating a vCenter instance, click the link in the **vCenter** table to view vCenter activity.

Figure 9-13: VMware vCenter Configuration



This page displays information about the configuration of the vCenter instance. To view information about vCenter resources, scroll down to the following sections:

- Hosts
- Virtual Switches
- Physical Connections
- Virtual Machines

- Network Host Connection Details

### Figure 9-14: Hosts, Virtual Switches, and Physical Connections

Hosts									
Filter table rows								FILTER	
	Inventory Host Name	Host DNS Name	Product Name	Hardware Model	CPU Usage (%)	Memory Usage (%)	Virtual Switches	State Summary	State Description
>	10.90.30.41	localhost.gigawatch.com	VMware ESXi 6.7.0 build-8169922	PowerEdge R320	75 <div><div></div></div>	92 <div><div></div></div>	2	[duplicate dns name]	[WARNING] duplicate-dns-name
>	10.90.30.60	app-animesh-esxi-4.gigawatch.com	VMware ESXi 6.5.0 build-8294253	PowerEdge R210 II	0 <div><div></div></div>	32 <div><div></div></div>	2	—	—
>	10.90.30.61	localhost.gigawatch.com	VMware ESXi 6.7.0 build-9214924	—	4 <div><div></div></div>	13 <div><div></div></div>	2	[duplicate dns name]	[WARNING] duplicate-dns-name

## Virtual Switches

Filter table rows

FILTER

	Name	Inventory Host Name	Host DNS Name	Status	NICs	Link Discovery Protocol	Link Discovery Direction	Virtual Switch Type
▶	VC22-V65-0Switch-1	10.9.30.60	app-animesh-esxi-4.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic3 is not connected</div> <div>[WARNING] vmnic2 is not connected</div>	2	LLDP	both	dvS
▶	VC22-V67-0Switch-1	10.9.30.61	localhost.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic3 is not connected</div> <div>[WARNING] vmnic2 is not connected</div>	2	LLDP	both	dvS
▶	VC22-V67-0Switch-1	10.9.30.41	localhost.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic3 is not connected</div> <div>[WARNING] vmnic2 is not connected</div>	2	LLDP	both	dvS
▶	vSwitch0	10.9.30.60	app-animesh-esxi-4.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic0 is not connected</div> <div>[WARNING] Link Discovery advertisement not configured</div>	1	CDP	listen	vSS
▶	vSwitch0	10.9.30.61	localhost.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic0 is not connected</div> <div>[WARNING] Link Discovery advertisement not configured</div>	1	CDP	listen	vSS
▶	vSwitch0	10.9.30.41	localhost.qa.bigswitch.com	<div></div> <div>[WARNING] vmnic0 is not connected</div> <div>[WARNING] Link Discovery advertisement not configured</div>	1	CDP	listen	vSS

Oct 24, 2018, 3:11:20am UTC

Show: 10 25 100 1 / 6 / 8

## Physical Connections

Inventory Host Name	Host DNS Name	Virtual Switch	NIC Team Group	NIC Name	Connected Interface	Status
10.9.30.41	localhost.qa.bigswitch.com	VC22-V67-DSwitch-1	—	vmnic3	—	✓ Connected
10.9.30.41	localhost.qa.bigswitch.com	VC22-V67-DSwitch-1	—	vmnic2	—	✓ Connected
10.9.30.41	localhost.qa.bigswitch.com	vSwitch0	—	vmnic0	—	✓ Connected
10.9.30.60	app-animesh-esxi-4.qa.bigswitch.c...	vSwitch0	—	vmnic0	—	✓ Connected

### Figure 9-15: Virtual Machines and Network Host Connection Details

Virtual Machines																
Filter table rows										FILTER						
	vSphere Host Name	VM Interface Name	Endpoint name	Power State	Monitored In Tg Tab	Tunnel Endnode	Status	VLAN	Inventry Host Net	Physical Connection	Virtual Switch Type	Endpoint type	MAC Address	IP Addresses	Default Gateway	Nat Load Balance Policy
	BMP-700 Analytics	VM adapter 1	localhost.qa.bigwatches.com	Powered On	—	—	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected [WARNING] VM Adapter not connected [WARNING] vmnic1 is not connected <a href="#">go 2 More...</a>	2000	10.9.30.81	vmnic3 vmnic2	dvsw virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.111/23,	—	IP Hash	—
	vcs2-vss-u1404-1	Network adapter 2	app-armeh-exst-4.qa.bigwatches.com	Powered On	—	—	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected [WARNING] vmnic1 is not connected [WARNING] Link Discovery advertisement not configured	Untagged	10.9.30.60	vmnic0 vmnic3 vmnic2	dvs virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.111/23,	192.168.23.3/24, 192.168.23.10/24, fdd:2bb:5af:fefb:d95a/f9	—	Originating Port ID
	vcs2-wc2-u1404-1	Network adapter 2	localhost.qa.bigwatches.com	Powered On	✓	VC22	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected [WARNING] Link Discovery advertisement not configured	Untagged	10.9.30.41	vmnic3 vmnic2	dvs virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.115/23,	192.168.23.3/24, fdd:2bb:5af:fefb:d95a/f9	—	Originating Port ID
	vcs2-wc2-u1404-2	Network adapter 2	localhost.qa.bigwatches.com	Powered On	✓	VC22	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected [WARNING] vmnic1 is not connected [WARNING] Link Discovery advertisement not configured	Untagged	10.9.30.61	vmnic0 vmnic3 vmnic2	dvs virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.115/23,	192.168.24.2/24, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.115/23,	—	Originating Port ID
	vcs2-wc2-u1604-1	Network adapter 2	localhost.qa.bigwatches.com	Powered On	✓	VC22	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected [WARNING] VMware tools not correctly installed	Untagged	10.9.30.41	vmnic3 vmnic2	dvs virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.115/23,	192.168.23.10/24, fdd:2bb:5af:fefb:d95a/f9	—	IP Hash
	vcs2-wc2-u1604-2	Network adapter 2	localhost.qa.bigwatches.com	Powered On	✓	VC22	[WARNING] vmnic3 is not connected [WARNING] vmnic2 is not connected	Untagged	10.9.30.61	vmnic3 vmnic2	dvs virtual machine	virtual machine	00:10:1d:a8:f6:b2:0e, fdd:2bb:5af:fefb:d95a/f9, 10.9.30.115/23,	192.168.23.10/24, fdd:2bb:5af:fefb:d95a/f9	—	IP Hash

Oct 24, 2018, 3:11:21am UTC

Show 10 20 100 1000

## Network Host Connection Details

Filter table rows

Filter

×

URL	Destination Host Name	Virtual Switch	VM Count	NIC	Connection Status	Security	View Logs
Unmapped	10.0.30.61, 10.0.30.41	VC22 VSI Switch-1	2	vmnic3	—	—	—
Unmapped	10.0.30.61, 10.0.30.41	VC22 VSI Switch-1	2	vmnic2	—	—	—
Unmapped	10.0.30.60	VC22 VSI Switch-1	2	vmnic3	—	—	—
Unmapped	10.0.30.61, 10.0.30.41	VC22 VSI Switch-1	1	vmnic3	—	—	—
Unmapped	10.0.30.61, 10.0.30.60, 10.0.30.61	VC22 VSI Switch-1	2	vmnic3	—	—	—
2009	10.0.30.61, 10.0.30.41	VC22 VSI Switch-1	0	vmnic3	—	—	—
2009	10.0.30.61, 10.0.30.41	VC22 VSI Switch-1	0	vmnic2	—	—	—

02/24/2018, 3:11 PM UTC

Show

25

25

0

## 9.9 Using the CLI to View vCenter Configuration

To view the vCenter configuration in the CLI, use the **show vcenter** command, as in the following examples:

controller-1# show vcenter							
#	vCenter Name	vCenter Host Name or IP	Last vCenter Update	Time	Detail State	vSphere Version	
1	vc-10-9-0-75	10.9.0.75	2017-09-09	18:02:35.980000 PDT	Connected and authenticated.	6.5.0	
2	vc-10-9-0-76	10.9.0.76	2017-09-09	18:02:36.488000 PDT	Connected and authenticated.	6.5.0	
3	vc-10-9-0-77	10.9.0.77	2017-09-09	18:02:35.908000 PDT	Connected and authenticated.	6.0.0	
4	vc-10-9-0-78	10.9.0.78	2017-09-09	18:02:33.507000 PDT	Connected and authenticated.	6.5.0	
5	vc-10-9-0-79	10.9.0.79	2017-09-09	18:02:32.248000 PDT	Connected and authenticated.	6.5.0	
6	vc-10-9-0-80	10.9.0.80	2017-09-09	18:02:32.625000 PDT	Connected and authenticated.	6.0.0	
7	vc-10-9-0-81	10.9.0.81	2017-09-09	18:02:34.672000 PDT	Connected and authenticated.	6.0.0	
8	vc-10-9-0-82	10.9.0.82	2017-09-09	18:02:33.008000 PDT	Connected and authenticated.	6.0.0	
9	vc-10-9-0-83	10.9.0.83	2017-09-09	18:02:30.011000 PDT	Connected and authenticated.	6.0.0	

```

10 vc-10-9-0-84 10.9.0.84 2017-09-09 18:02:33.024000 PDT Connected and authenticated. 6.5.0
11 vc-10-9-0-85 10.9.0.85 2017-09-09 18:02:34.827000 PDT Connected and authenticated. 6.0.0
12 vc-10-9-0-86 10.9.0.86 2017-09-09 18:02:35.164000 PDT Connected and authenticated. 6.0.0
13 vc-10-9-0-87 10.9.0.87 2017-09-09 18:02:38.042000 PDT Connected and authenticated. 6.5.0
14 vc-10-9-0-88 10.9.0.88 2017-09-09 18:02:37.212000 PDT Connected and authenticated. 6.0.0
15 vc-10-9-0-89 10.9.0.89 2017-09-09 18:02:33.436000 PDT Connected and authenticated. 6.5.0
controller-1#

controller-1# show vcenter vc-10-9-0-75
# vCenter Name vCenter Host Name or IP Last vCenter Update Time Detail State vSphere Version
--|-----|-----|-----|-----|-----|-----|
1 vc-10-9-0-75 10.9.0.75 2017-09-09 18:02:44.698000 PDT Connected and authenticated. 6.5.0
controller-1#

controller-1# show vcenter vc-10-9-0-75 detail
vCenter Name : vc-10-9-0-75
vCenter Host Name or IP : 10.9.0.75
Last vCenter Update Time : 2017-09-09 18:02:49.463000 PDT
Detail State : Connected and authenticated.
vSphere Version : 6.5.0
controller-1#

controller-1# show vcenter vc-10-9-0-75 error
vCenter Name : vc-10-9-0-75
vCenter Host Name or IP : 10.9.0.75
State : connected
Detail State : Connected and authenticated.
Detailed Error Info :
controller-1#

```

## 9.10 Minimum Permissions for Non-admin Users

For a non-admin user to add, remove, edit, or monitor a vCenter via the DANZ Monitoring Fabric, the privilege level assigned to the non-admin user should be VSPAN operation. To assign VSPAN operation privileges to a user, do the following:

1. From the vCenter GUI, navigate to **Menu > Administration**.
2. Once on the page, click on the **Users and Groups** link in the navigation bar on the left.

**Figure 9-16: Users and Groups**

The screenshot shows the vSphere Client interface. The left navigation pane is expanded to 'Administration' > 'Users and Groups'. The main content area is titled 'Users and Groups' and shows a list of users in the 'localos' domain. The table below lists the users:

Username	First Name	Last Name	Email	Description	Locked	Disabled	Domain
sshd	sshd	PrivSep			No	No	localos
eam	eam				No	No	localos
root	root				No	No	localos
imagebuilder	imagebuilder				No	No	localos
vapiEndpoint	vapiEndpoint				No	No	localos
pod	pod				No	No	localos
cm	cm				No	No	localos

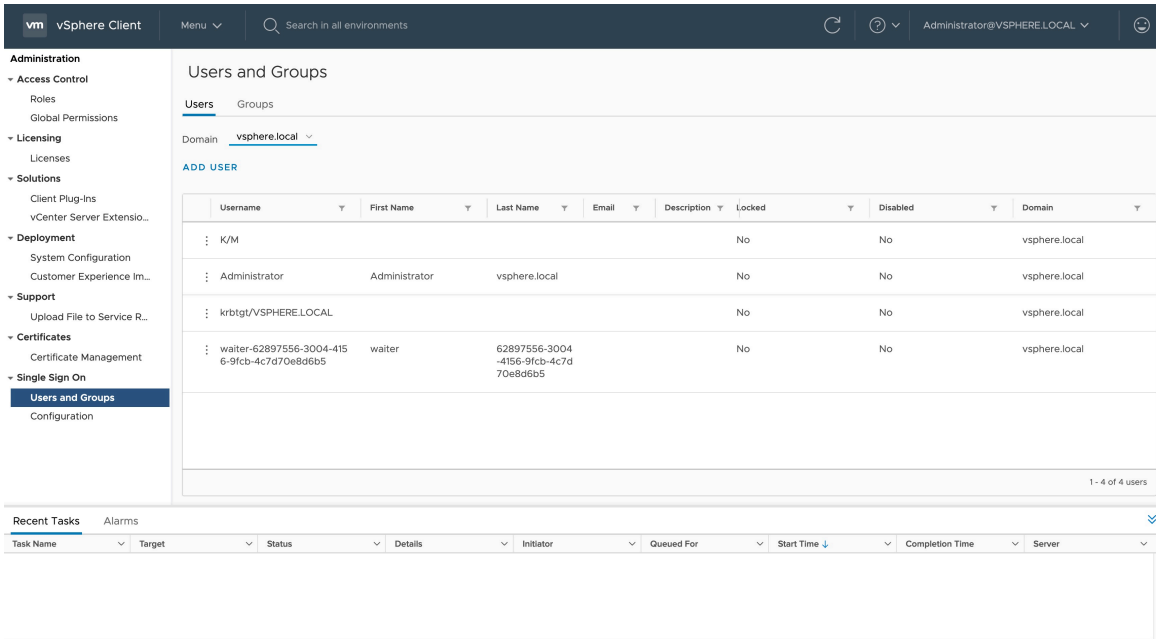
At the bottom of the interface, there is a 'Recent Tasks' section with a table for task management:

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server



3. Click on the **Users** tab and make sure the appropriate domain is selected (in this case, the domain is `vsphere.local`).

Figure 9-17: Domain Selection



4. Next click on the **ADD USER** link and create the desired user. (In the example below, we create a user called **dmf-alice**.)

Figure 9-18: Add a New User

Add User

Username \*

dmf-alice

Password \*

.....

Confirm Password \*

.....

First Name

Alice

Last Name

D

Email

alice@arista.com

Description

CANCEL

ADD

5. Verify that the newly created user is on the **Users and Groups** page.

**Figure 9-19: Verify User Has Been Created**

The screenshot shows the vSphere Client interface with the 'Users and Groups' page selected. The left navigation pane shows 'Administration' > 'Access Control' > 'Users and Groups'. The main content area displays a table of users for the 'vsphere.local' domain. The table has columns for Username, First Name, Last Name, Email, Description, Locked, Disabled, and Domain. The users listed are K/M, Administrator, krbtgt/VSHERE.LOCAL, waiter, and dmfl-alice. The 'waiter' user is highlighted.

Username	First Name	Last Name	Email	Description	Locked	Disabled	Domain
K/M					No	No	vsphere.local
Administrator	Administrator	vsphere.local			No	No	vsphere.local
krbtgt/VSHERE.LOCAL					No	No	vsphere.local
waiter-62897556-3004-4156-9fcb-4c7d70e8d6b5	waiter	62897556-3004-4156-9fcb-4c7d70e8d6b5			No	No	vsphere.local
dmfl-alice	Alice	D	alice@arist a.com		No	No	vsphere.local

6. Once the desired user has been created, create and assign a role to this user. Click on **Roles** under **Access Control** in the navigation bar on the left. Next click on the **+** sign to add a new role.

**Figure 9-20: Add a New Role**

The screenshot shows the vSphere Client interface with the 'Roles' page selected. The left navigation pane shows 'Administration' > 'Access Control' > 'Roles'. The main content area displays a list of roles for the '10.106.1164' provider. The roles listed are Administrator, Read-only, No access, AutoUpdateUser, Content library administrator (sample), Datastore consumer (sample), Network administrator (sample), No cryptography administrator, Resource pool administrator (sample), Tagging Admin, Virtual Machine console user, Virtual machine power user (sample), Virtual machine user (sample), VMware Consolidated Backup user (sample), and vSphere Client Solution User. The 'Administrator' role is highlighted.

Role	Description
Administrator	Full access rights
Read-only	
No access	
AutoUpdateUser	
Content library administrator (sample)	
Datastore consumer (sample)	
Network administrator (sample)	
No cryptography administrator	
Resource pool administrator (sample)	
Tagging Admin	
Virtual Machine console user	
Virtual machine power user (sample)	
Virtual machine user (sample)	
VMware Consolidated Backup user (sample)	
vSphere Client Solution User	

7. In the **New Role** pop up dialog, select **Distributed Switch** from the left and then scroll down to find and select **VSPAN operation** as the role. Click **Next** and give the new role a new name. (In the example below, we named the new role **monitor-dmf**.) Click **Finish** to create the new role.

**Figure 9-21: Select Role Type**

New Role

Alarms  
AutoDeploy  
Certificate Management  
Certificates  
Cns  
Content Library  
Cryptographic operations  
Datacenter  
Datastore  
Datastore cluster  
Distributed switch  
ESX Agent Manager  
Extension  
External stats provider  
Folder  
Global  
Health update provider  
Host

☐ All Distributed switch Privileges

☐ IPFIX operation  
☐ Modify  
☐ Move  
☐ Network I/O control operation  
☐ Policy operation  
☐ Port configuration operation  
☐ Port setting operation  
☒ VSPAN operation

All | Selected | Unselected

CANCEL BACK NEXT

**Figure 9-22: Save New Role**

New Role

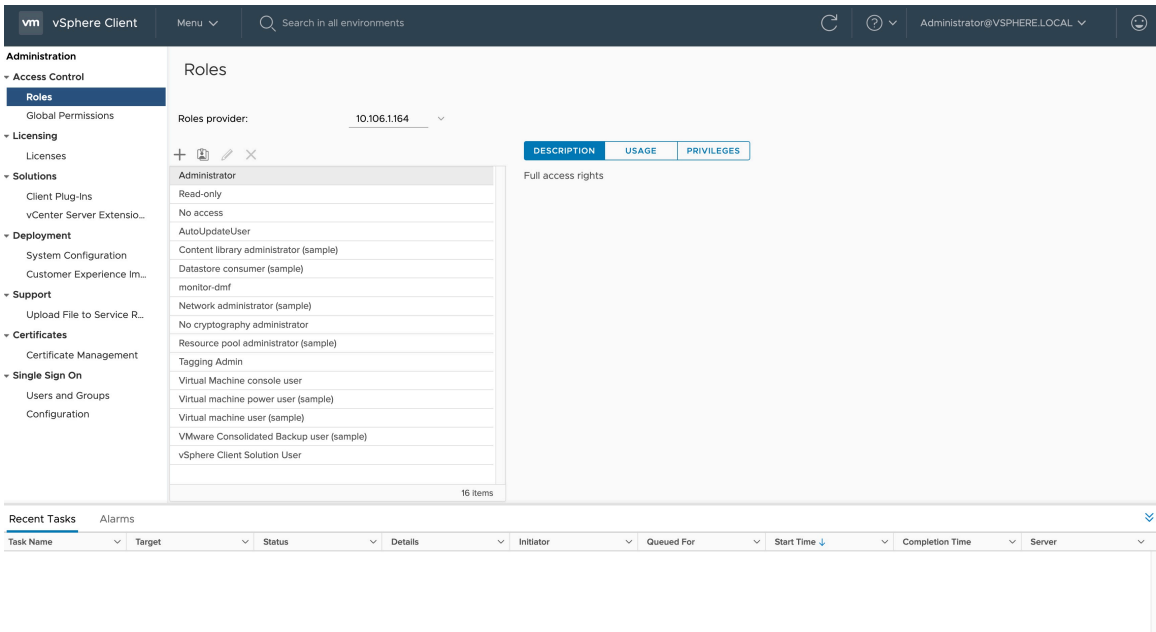
Role name

Description

CANCEL BACK FINISH

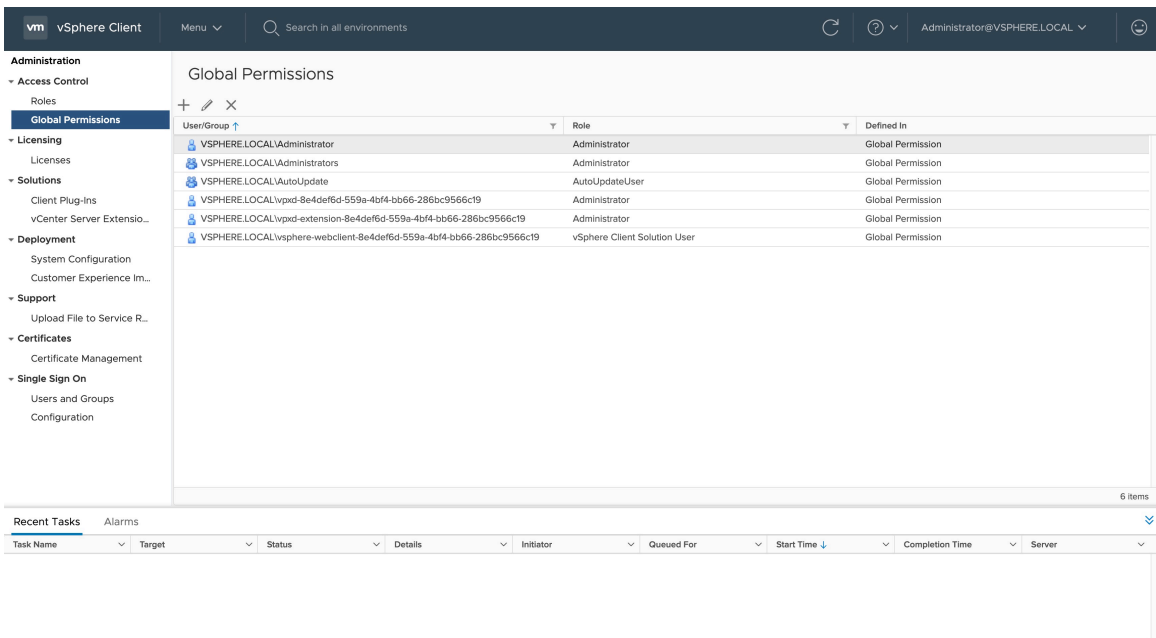
8. Verify that the new role has been created on the **Roles** page.

**Figure 9-23: Verify New Role Has Been Created**



9. To assign the new role to the new user, click the **Global Permissions** link in the navigation bar on the left. Next click on the + sign to assign the new role.

**Figure 9-24: Global Permissions**



10. In the **Add Permission** dialog, type the newly created username and select the newly created role as shown in the figure below.



**Note:** Do not forget to check the **Propagate to children** checkbox.

**Figure 9-25: Assign Role to User**

Add Permission

Global Permission Root

✕

User

vsphere.local

▼

Q dmf-alice

Role

monitor-dmf

▼

☒

Propagate to children

CANCEL

OK

11. Verify that the newly created role has been assigned to the newly created user.

**Figure 9-26: Verify Role Assignment to User**

vm vSphere Client

Menu

Search in all environments

Administrator@VSPHERE.LOCAL

Administration

Access Control

Roles

Global Permissions

Licensing

Solutions

Deployment

Support

Certificates

Single Sign On

Global Permissions

User/Group

Role

Defined In

VSPHERE.LOCAL\Administrator

Administrator

Global Permission

VSPHERE.LOCAL\Administrators

Administrator

Global Permission

VSPHERE.LOCAL\AutoUpdate

AutoUpdateUser

Global Permission

VSPHERE.LOCAL\dmf-alice

monitor-dmf

Global Permission

VSPHERE.LOCAL\vpdx-8e4def6d-559a-4bf4-bb66-286bc9566c19

Administrator

Global Permission

VSPHERE.LOCAL\vpdx-extension-8e4def6d-559a-4bf4-bb66-286bc9566c19

Administrator

Global Permission

VSPHERE.LOCAL\vsphere-webclient-8e4def6d-559a-4bf4-bb66-286bc9566c19

vSphere Client Solution User

Global Permission

7 items

Recent Tasks

Alarms

Task Name

Target

Status

Details

Initiator

Queued For

Start Time

Completion Time

Server

## Advanced Fabric Settings

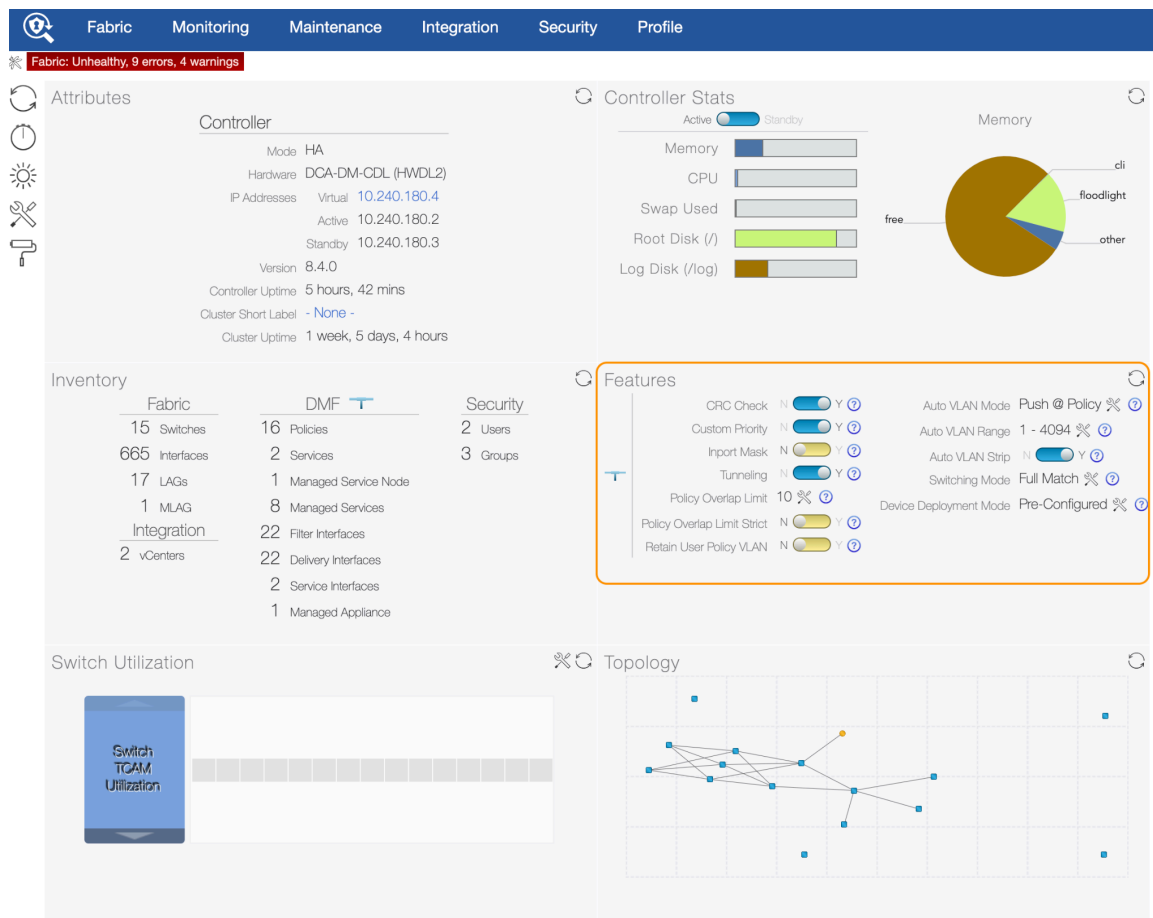
This chapter describes fabric-wide configuration options that are required in advanced use cases for deploying DMF policies.

### 10.1 Configuring Advanced Fabric Settings

#### Overview

Prior to the DMF 8.4 release, the fabric-wide settings, specifically the Features section as shown below were available on the home page after logging in.

**Figure 10-1: DMF Legacy Page (pre 8.4)**



In DMF 8.4, a newly designed Dashboard replaces the former home page. The **Features** section is now the new **DMF Features** page. To navigate to the DMF Features Page, click on the **gear icon** in the navigation bar.

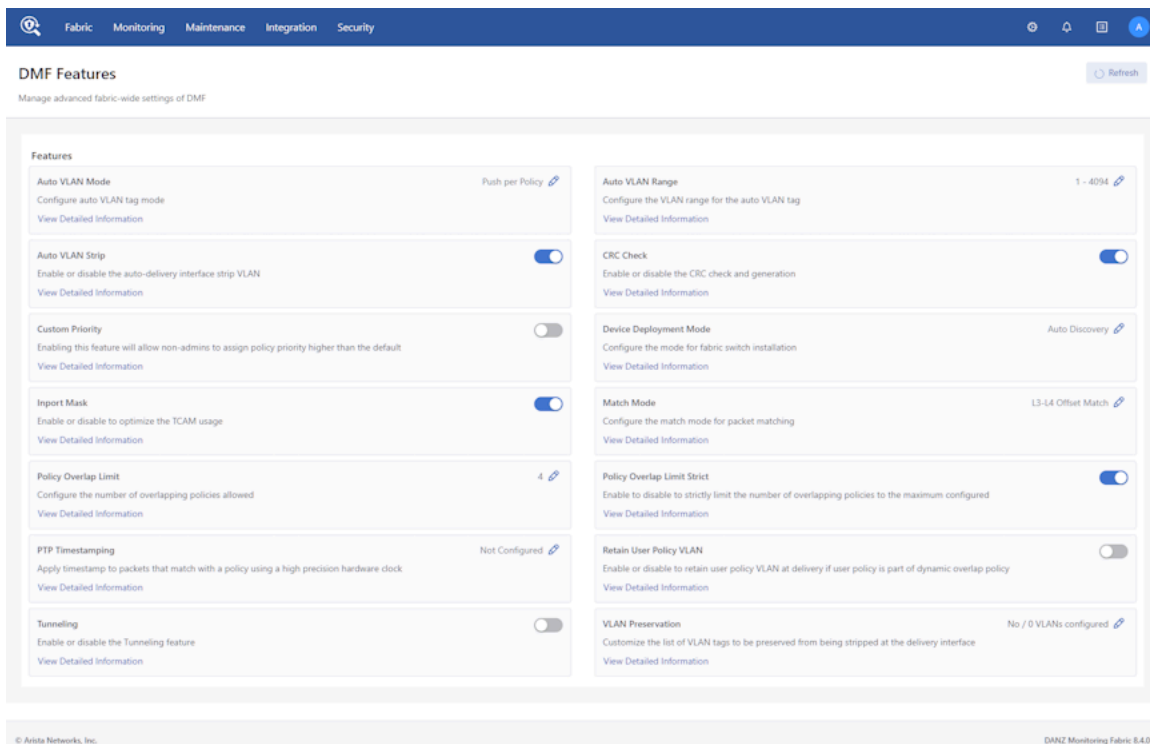
**Figure 10-2: Gear Icon**



## Page Layout


All fabric-wide configuration settings required in advanced use cases for deploying DMF policies appear in the new DMF Features Page.

**Figure 10-3: DMF Features Page**



Each card on the page corresponds to a feature set.

**Figure 10-4: Feature Set Card Example**



---

The UI displays the following:

- Feature Title
- A brief description
- View / Hide detailed information link
- Current Setting
- Edit Link - Use the **Edit** configuration button (**pencil icon**) to change the value.

The fabric-wide options used with DMF policies include the following:

**Table 7: Feature Set**

Auto VLAN Mode	Auto VLAN Range
Auto VLAN Strip	CRC Check
Custom Priority	Device Deployment Mode
Inport Mask	Match Mode
Policy Overlap Limit	Policy Overlap Limit Strict
PTP Timestamping	Retain User Policy VLAN
Tunneling	VLAN Preservation

## 10.2 Managing VLAN Tags in the Monitoring Fabric

VLAN tags are often used by analysis tools to identify the filter interface where the traffic was received. How VLAN IDs are assigned to traffic depends on which auto-VLAN mode is enabled. The system automatically assigns the VLAN ID from a configurable range of VLAN IDs, which by default is from **1** to **4094**. Available auto-VLAN modes behave as follows:

- **push-per-policy (default)**: Automatically adds a unique VLAN ID to all traffic selected by a specific policy. This setting enables tag-based forwarding.
- **push-per-filter**: Automatically adds a unique VLAN ID from the default auto-VLAN range (**1-4094**) to each filter interface. A custom VLAN range can be specified using the **auto-vlan-range** command. Any VLAN ID not in the auto-VLAN range can be manually assigned to a filter interface.

The VLAN ID assigned to policies or filter interfaces remains the same after controller reboot or failover. However, it changes if the policy is removed and added back again. Also, when the VLAN range is changed, existing assignments are discarded and new assignments are made.

The push-per-filter feature preserves the original VLAN tag, but if the packet already has two VLAN tags the outer VLAN tag is rewritten with the assigned VLAN ID.

The following table summarizes how VLAN tagging occurs with the different auto-VLAN modes:



Table 8: VLAN Tagging Across VLAN Modes

Traffic with VLAN tag type	push-per-policy Mode (Applies to all supported switches)	push-per-filter Mode (Applies to all supported switches)
Untagged	Pushes a single tag	Pushes a single tag
Single tag	Pushes an outer (second) tag	Pushes an other (second) tag
Two tags	Pushes an outer (third) tag. Except on T3 based switches it rewrites the outer tag. Due to this outer customer VLAN is replaced by DMF policy VLAN.	Rewrites the outer tag. Due to this outer customer VLAN is replaced by DMF filter VLAN.



**Note:** When you enable push-per-policy, the `auto-delivery-interface-vlan-strip` feature is enabled if it was disabled prior to enabling push-per-policy. When you enable push-per-filter, the global delivery strip option is not enabled if it was previously disabled.

The following table summarizes how different auto-VLAN modes affect the applications and services supported.



**Note:** Matching on untagged packets cannot be applied to DMF policies when in push-per-policy mode.

Table 9: Auto-VLAN Mode Comparison

Auto-VLAN Mode	Supported Platform	TCAM Optimization in the Core	L2 GRE Tunnels Support	Q-in-Q Packets Preserve Both Original Tags	Support DMF Service Node Services	Manual Tag to Filter Interface
Push-per- policy (default)	All	Yes	Yes	Yes	All	Policy tag overwrites manual
Push-per- filter	All	No	Yes	No	All	Configuration not allowed



**Note:** Tunneling is supported with full-match or offset-match modes but not with l3-l4-match mode.

Tag-based forwarding, which improves traffic forwarding and reduces TCAM utilization on the monitoring fabric switches, is enabled only when you choose the push-per-policy option.

When the mode is push-per-filter, the VLAN that is getting pushed or rewritten can be displayed using the `show interface-names` command as shown below:

```
controller-1> show interface-names
~~~~~ Filter Interface(s) ~~~~~
#  DMF IF      Switch IF      Name      Dir State Speed  VLAN Tag Analytics Ip address Connected Device
--  -
1  TAP-PORT-eth1  FILTER-SW1    ethernet1  rx  up   10Gbps  5      True
2  TAP-PORT-eth10  FILTER-SW1    ethernet10 rx  up   10Gbps  10     True
3  TAP-PORT-eth12  FILTER-SW1    ethernet12 rx  up   10Gbps  11     True
4  TAP-PORT-eth14  FILTER-SW1    ethernet14 rx  up   10Gbps  12     True
5  TAP-PORT-eth16  FILTER-SW1    ethernet16 rx  up   10Gbps  13     True
6  TAP-PORT-eth18  FILTER-SW1    ethernet18 rx  up   10Gbps  14     True
7  TAP-PORT-eth20  FILTER-SW1    ethernet20 rx  up   10Gbps  16     True
8  TAP-PORT-eth22  FILTER-SW1    ethernet22 rx  up   10Gbps  17     True
```

## 10.3 Auto VLAN Mode

VLAN tags are also often used by analysis tools to identify the filter interface where the traffic was received. How VLAN IDs are assigned to traffic depends on which of Auto VLAN modes is enabled. The system automatically assigns the VLAN ID from a configurable range of VLAN IDs, which by default is from 1 to 4094.

- Push per Policy (default): Automatically adds a unique VLAN ID to all traffic selected by a specific policy. This setting enables tag-based forwarding.
- Push per Filter: Automatically adds a unique VLAN ID, from the default auto-vlan-range (1-4094) to each filter interfaces. New vlan range can be specified using the auto-vlan-range command. VLAN ID not in auto-vlan-range can be used to manually assign to filter interfaces.

The following table summarizes how VLAN tagging occurs with the different Auto VLAN modes.

Traffic with VLAN tag type	push-per-policy Mode (Applies to all supported switches)	push-per-filter Mode (Applies to all supported switches)
Untagged	Pushes a single tag	Pushes a single tag
Single tag	Pushes an outer (second) tag	Pushes an other (second) tag
Two tags	Pushes an outer (third) tag. Except on T3 based switches it rewrites the outer tag. Due to this outer customer VLAN is replaced by DMF policy VLAN.	Rewrites the outer tag. Due to this outer customer VLAN is replaced by DMF filter VLAN.



**Note:** When users enable push-per-policy, the auto-delivery-interface-vlan-strip feature is enabled if it was disabled prior to enabling push-per-policy.

When users enable push-per-filter feature, the global delivery strip option is not enabled if it was previously disabled.

The following table summarizes how different Auto VLAN modes affect the applications and services supported.



**Note:** Matching on untagged packets cannot be applied to DMF policies when in push-per-policy mode.

Auto-VLAN Mode	Supported Platform	TCAM Optimization in the Core	L2 GRE Tunnels Support	Q-in-Q Packets Preserve Both Original Tags	Supported DMF Service Node Services	Manual Tag to Filter Interface
Push per Policy (default)	All	Yes	Yes	Yes	All	Policy tag overwrites manual
Push per Filter	All	No	Yes	No	All	Configuration not allowed

Tag-based forwarding, which improves traffic forwarding and reduces TCAM utilization on the monitoring fabric switches, is enabled only when users choose the push per policy option.

Use the CLI or the GUI to configure Auto VLAN Mode as described in the following topics.

- [Configuring Auto VLAN Mode using the CLI](#)
- [Configuring Auto VLAN Mode using the GUI](#)

### 10.3.1 Configuring Auto VLAN Mode using the CLI

To set the auto VLAN mode, complete the following steps:

1. When the auto VLAN mode is set to push-per-filter, you can define the range of VLAN IDs to be automatically assigned by entering the following command from config mode:

```
auto-vlan-range vlan-min <start> vlan-max <end>
```

Replace **start** and **end** with the first and last VLAN ID in the range to use for assignment. For example, the following command assigns a range of 100 VLAN IDs from **3994** to **4094**:

```
controller-1(config)# auto-vlan-range vlan-min 3994 vlan-max 4094
```

2. Select the VLAN mode using the following command from config mode:

```
auto-vlan-mode command { push-per-filter | push-per-policy }
```

The effect of these options is summarized in the [Managing VLAN Tags in the Monitoring Fabric](#) section.

For example, the following command adds a unique outer VLAN tag to each packet received on each filter interface:

```
controller-1(config)# auto-vlan-mode push-per-filter
Switching to auto vlan mode would cause policies to be re-installed. Enter
"yes" (or "y")
to continue: y
```

3. To display the configured VLAN mode, enter the **show fabric** command, as in the following example:

```
controller-1# show fabric
~~~~~ Aggregate Network State ~~~~~
Number of switches                : 5
Inport masking                    : True
Start time                        : 2018-11-02 23:42:29.183000 UTC
Number of unmanaged services      : 0
Filter efficiency                  : 1:1
Number of switches with service interfaces : 0
Total delivery traffic (bps)      : 411Kbps
Number of managed service instances : 0
Number of service interfaces      : 0
Match mode                        : full-match
Number of delivery interfaces     : 13
Max pre-service BW (bps)         : -
Auto VLAN mode                    : push-per-filter
Number of switches with delivery interfaces : 4
Number of managed devices        : 1
Uptime                            : 2 days, 19 hours
Total ingress traffic (bps)       : 550Kbps
Max overlap policies (0=disable)  : 10
Auto Delivery Interface Strip VLAN : False
Number of core interfaces        : 219
Max filter BW (bps)              : 184Gbps
Number of switches with filter interfaces : 5
State                            : Enabled
Max delivery BW (bps)            : 53Gbps
Total pre-service traffic (bps)   : -
Track hosts                      : True
Number of filter interfaces       : 23
Number of active policies        : 3
Number of policies               : 25
-----output truncated-----
```

4. To display the VLAN IDs assigned to each policy, enter the **show policy** command, as in the following example:

```

controller-1> show policy
# Policy Name Action Runtime Status Type Priority Overlap Priority
┌→Push VLAN Filter BW Delivery BW Post Match Filter Traffic Delivery Traffic
Services
┌|-----|-----|-----|-----|-----
└|-----|-----|-----|-----|-----
┌→-----|-----|-----|-----|-----
└-----|-----|-----|-----|-----
┌→-----|
1 GENERATE-NETFLOW-RECORDS forward installed Configured 100 0 4
┌→ 100Gbps 10Gbps - - DMF-OOB-NETFLOWSERVICE
2 P1 forward inactive Configured 100 0 1
┌→ - 1Gbps - -
3 P2 forward inactive Configured 100 0 3
┌→ - 10Gbps - -
4 TAP-WINDOWS10-NETWORK forward inactive Configured 100 0 2
┌→ 21Gbps 1Gbps - -
5 TIMESTAMP-INCOMING-PACKETS forward inactive Configured 100 0 5
┌→ - 100Gbps - - DMF-OOB-TIMESTAMPINGSERVICE
controller -1>)#

```




**Note:** The strip VLAN option, when enabled, removes the outer VLAN tag, including the VLAN ID applied by any rewrite VLAN option.

## 10.3.2 Configuring Auto VLAN Mode using the GUI

### Auto VLAN Mode

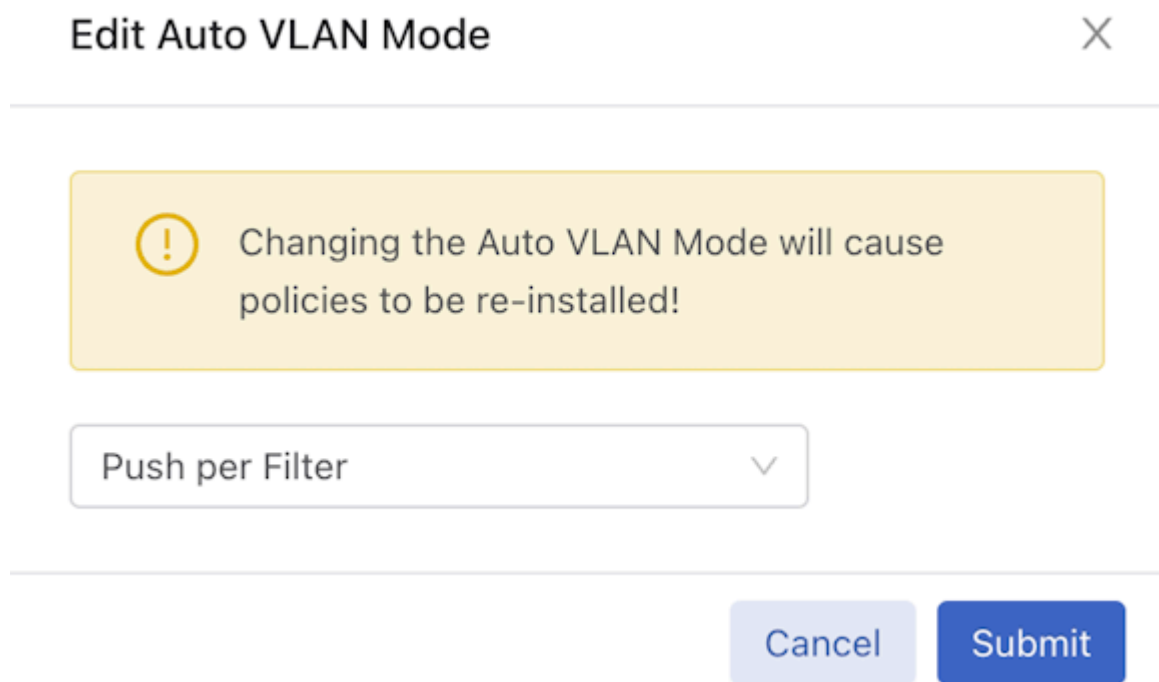
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 10-5: Auto VLAN Mode Config**



2. A confirmation edit dialogue window appears displaying the corresponding prompt message.

**Figure 10-6: Edit VLAN Mode**



The dialog box is titled "Edit Auto VLAN Mode" with a close button (X) in the top right corner. Below the title bar is a yellow warning box containing a yellow circle with an exclamation mark and the text "Changing the Auto VLAN Mode will cause policies to be re-installed!". Below the warning box is a dropdown menu with the text "Push per Filter" and a downward arrow. At the bottom right of the dialog are two buttons: "Cancel" (light blue) and "Submit" (dark blue).

3. To configure different modes, click the drop-down arrow to open the menu.


**Figure 10-7: Drop-down Example**



A single dropdown menu with the text "Push per Filter" and a downward arrow.

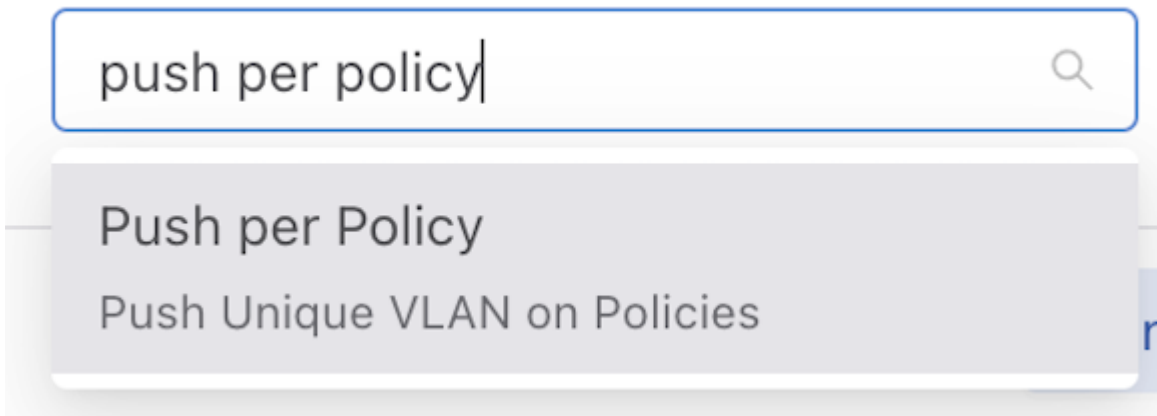
4. From the drop-down menu, select and click on the desired mode.

**Figure 10-8: Push Per Policy**



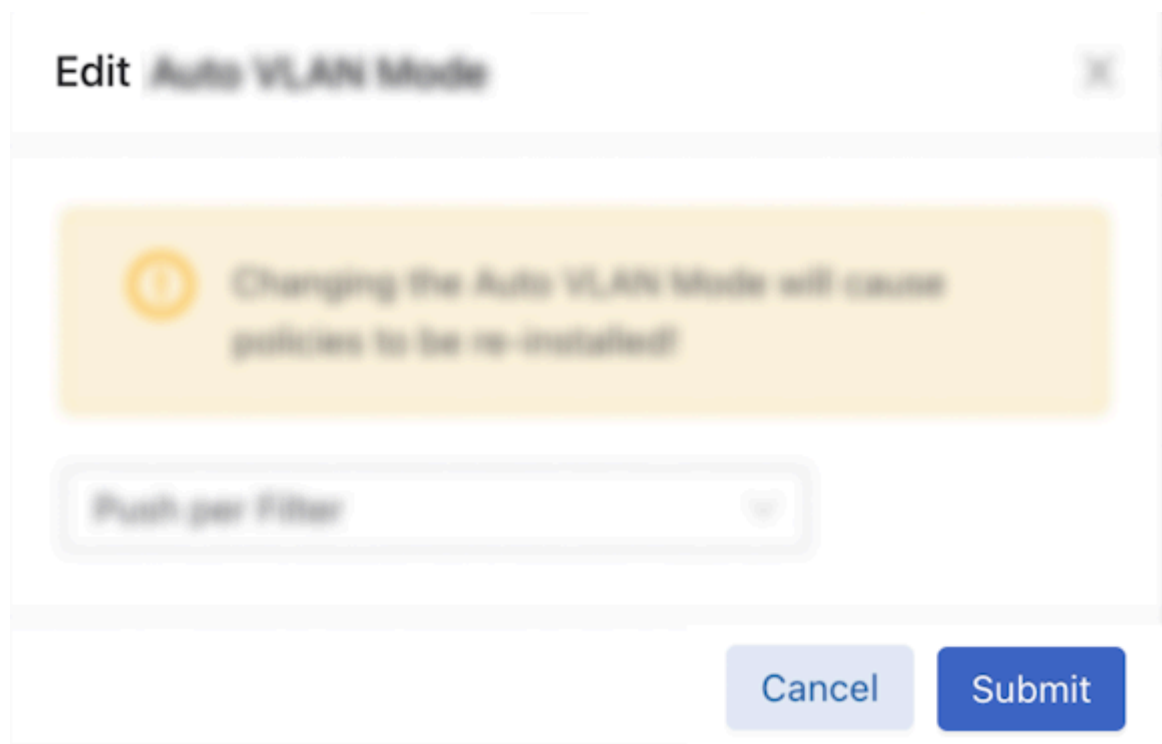
5. Alternatively, you can directly input the desired mode name in the input area.

**Figure 10-9: Push Per Policy**



- Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.

**Figure 10-10: Submit Button**



- The current configuration status displays next to the edit button after successfully setting the configuration.

**Figure 10-11: Current Configuration Status**



The following feature sets work in the same manner as the **Auto VLAN Mode** feature described above.

- **Device Deployment Mode**
- **Match Mode**

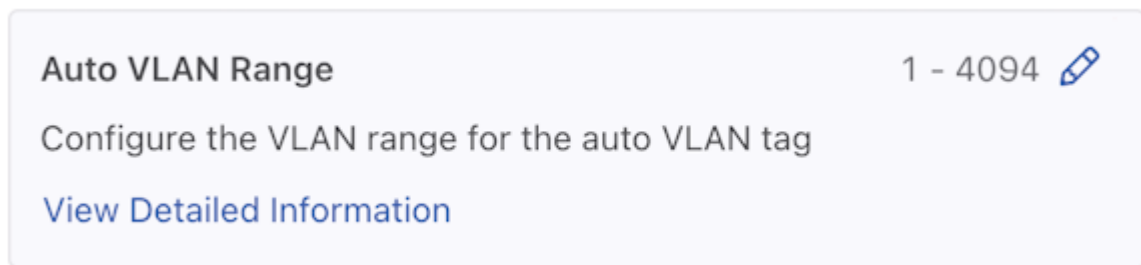
## 10.4 Auto VLAN Range

### Auto VLAN Range

The range of automatically generated VLANs only applies when setting Auto VLAN Mode to push-per-filter. VLANs are picked from the range 1 - 4094 when not set.

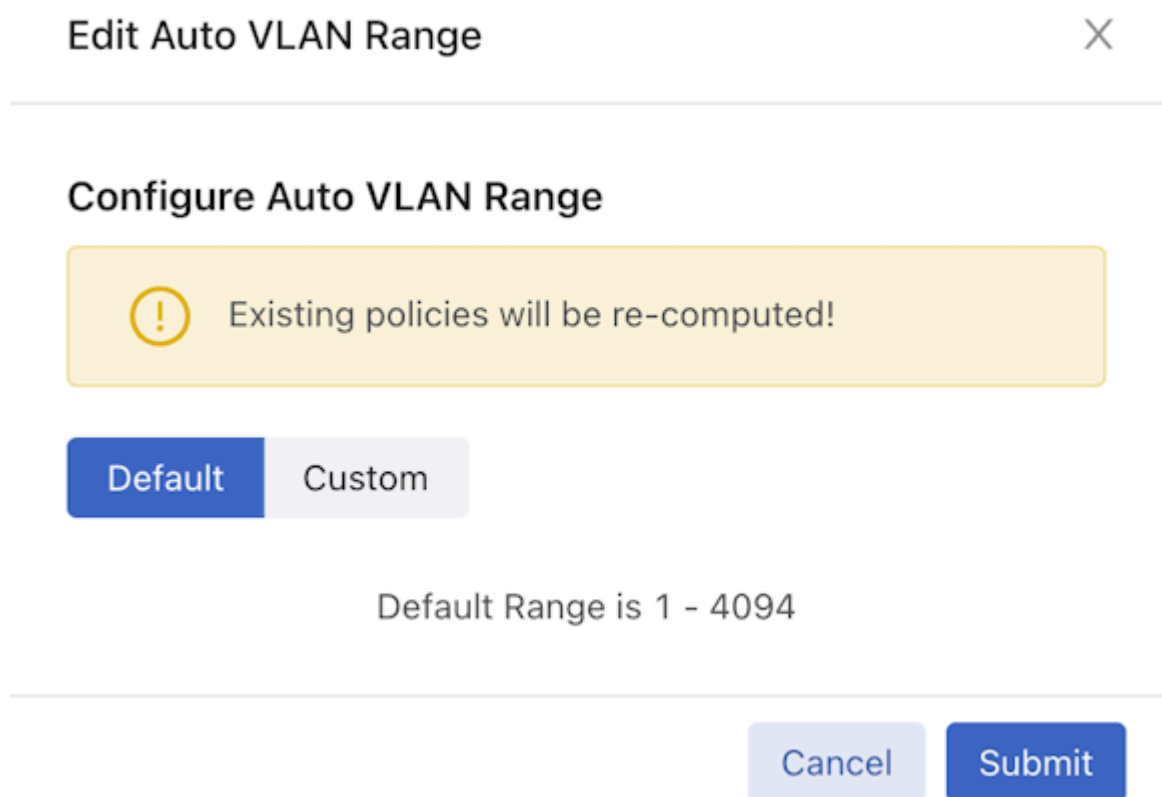
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 10-12: Edit Auto VLAN Range**



2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. The Auto VLAN Range defaults to 1 - 4094.

**Figure 10-13: Edit Auto VLAN Range**





- Click on the **Custom** button to configure the custom range.

**Figure 10-14: Custom Button**

- Adjust range value (minimum value: 1, maximum value: 4094). There are three ways to adjust the value of a range:
  - Directly enter the desired value in the input area, with the left side representing the minimum value of the range and the right side representing the maximum value.
  - Adjust the value by dragging the **slider** using a mouse. The left knob represents the minimum value of the range, while the right knob represents the maximum value.
  - Use the up and down arrow buttons in the input area to adjust the value accordingly. Pressing the up arrow increments the value by 1, while pressing the down arrow decrements it by 1.
- Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.
- After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 10-15: Configuration Change Success**

### Configuring Auto VLAN Range using the CLI

To set the Auto VLAN Range, complete the following:

```
auto-vlan-range vlan-min start vlan-max end
```

To set the Auto VLAN Range, complete the following steps. Replace **start** and **end** with the first and last VLAN ID in the desired range. For example, the following command assigns a range of 100 VLAN IDs from 3994 to 4094:

```
controller-1(config)# auto-vlan-range vlan-min 3994 vlan-max 4094
```

## 10.5 Auto VLAN Strip

The strip VLAN option removes the outer VLAN tag before the packet is forwarded to a delivery interface. If the packet has two VLAN tags, only the outer tag is removed. If it has no VLAN tag, the packet is not modified. Users can remove the VLAN ID on traffic forwarded to a specific delivery interface, globally for all delivery interfaces. The strip VLAN option removes any VLAN ID applied by the rewrite VLAN option.

---

The strip vlan option removes the VLAN ID on traffic forwarded to delivery interface. The following are the two methods available:

- Remove VLAN IDs fabric-wide for all delivery interfaces. This method removes only the VLAN tag added by DMF Fabric.
- On specific delivery interfaces. This method has four options
  - Keep all tags intact. This will preserve VLAN tag added by DMF Fabric and other tags in the traffic using strip-no-vlan option during delivery interface configuration.
  - Remove only the outer VLAN tag added by the DANZ Monitoring Fabric using strip-one-vlan option during delivery interface configuration.
  - Remove only the second (inner) tag. This will preserve VLAN (outer) tag added by DMF Fabric and remove second (inner) tag in the traffic using strip-second-vlan option during delivery interface configuration.
  - Remove two tags. Removes the outer VLAN tag added by DMF fabric and inner vlan tag in the traffic using strip-two-vlan option during delivery interface configuration.



**Note:** The strip vlan command for a specific delivery interface overrides fabric-wide strip vlan option.

By default, the VLAN ID is stripped, when it is added by DMF as a result of enabling the following options.

- Push per Policy
- Push per Filter
- Rewrite VLAN under filter-interfaces

Tagging and stripping VLANs as they ingress and egress DMF differs based on whether the switch is a Trident 3 based switch or not.

Use the CLI or the GUI to configure Auto VLAN Strip as described in the following topics.

- [Auto VLAN Strip using the CLI](#)
- [Auto VLAN Strip using the GUI](#)

## 10.5.1 Auto VLAN Strip using the CLI

The strip VLAN option removes the outer VLAN tag before the packet is forwarded to a delivery interface. If the packet has two VLAN tags, only the outer tag is removed. If it has no VLAN tag, the packet is not modified. You can remove the VLAN ID on traffic forwarded to a specific delivery interface or globally for all delivery interfaces. The strip VLAN option removes any VLAN ID applied by the rewrite VLAN option.

The following are the two methods available:

- **Remove VLAN IDs fabric-wide for all delivery interfaces:** This method removes only the VLAN tag added by the DMF Fabric.
- **Remove VLAN IDs only on specific delivery interfaces:** This method has four options:
  - Keep all tags intact. This will preserve VLAN tag added by DMF Fabric and other tags in the traffic using **strip-no-vlan** option during delivery interface configuration.
  - Remove only the outer VLAN tag added by the DANZ Monitoring Fabric using the **strip-one-vlan** option during delivery interface configuration.
  - Remove only the second (inner) tag. This will preserve VLAN (outer) tag added by DMF and remove the second (inner) tag in the traffic using the **strip-second-vlan** option during delivery interface configuration.
  - Remove two tags. Removes the outer VLAN tag added by DMF fabric and the inner VLAN tag in the traffic using the **strip-two-vlan** option during delivery interface configuration.



**Note:** The **strip vlan** command for a specific delivery interface overrides the **fabric-wide strip VLAN** option.

By default, the VLAN ID is stripped when it is added by DMF as a result of enabling the following options:

- push-per-policy
- push-per-filter
- rewrite vlan under filter-interfaces

To see the current **auto-delivery-interface-vlan-strip** configuration, enter the following command:

```
controller-1> show running-config feature details
! deployment-mode
deployment-mode pre-configured
! auto-delivery-interface-vlan-strip
auto-delivery-interface-vlan-strip
! auto-vlan-mode
auto-vlan-mode push-per-policy
! auto-vlan-range
auto-vlan-range vlan-min 3200 vlan-max 4094
! crc
crc
! match-mode
match-mode full-match
! tunneling
tunneling
! allow-custom-priority
allow-custom-priority
! inport-mask
no inport-mask
! overlap-limit-strict
no overlap-limit-strict
! overlap-policy-limit
overlap-policy-limit 10
! packet-capture
packet-capture retention-days 7
```

To see the current **auto-delivery-interface-vlan-strip** state, enter the following command:

```
controller-1> show fabric
~~~~~ Aggregate Network State ~~~~~
Number of switches : 5
Inport masking : True
Start time : 2018-10-16 22:30:03.345000 UTC
Number of unmanaged services : 0
Filter efficiency : 3005:1
Number of switches with service interfaces : 0
Total delivery traffic (bps) : 232bps
Number of managed service instances : 0
Number of service interfaces : 0
Match mode : l3-l4-match
Number of delivery interfaces : 24
Max pre-service BW (bps) : -
Auto VLAN mode : push-per-policy
Number of switches with delivery interfaces : 5
Number of managed devices : 1
Uptime : 21 hours, 53 minutes
Total ingress traffic (bps) : 697Kbps
Max overlap policies (0=disable) : 10
Auto Delivery Interface Strip VLAN : True
```

To disable this global command, enter the following command:

```
controller-1(config-switch-if)# no auto-delivery-interface-vlan-strip
```

---

The delivery interface level command to strip the VLAN overrides the global **auto-delivery-interface-vlan-strip** command. For example, when global VLAN stripping is disabled or when you want to override the default strip option on a delivery interface use the below options:

When you want to strip the VLAN added by DMF fabric on a specific delivery interface, use the following command:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-one-vlan
```

When global VLAN stripping is enabled, it strips only the outer VLAN ID. If you want to remove outer VLAN ID that was added by DMF as well as the inner VLAN ID, enter the following command:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-two-vlan
```

To strips only the inner VLAN ID and preserve the outer VLAN ID that was added by DMF, use the following command:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-second-vlan
```

To preserve the VLAN tag added by DMF and other tags in the traffic, use the following command:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-no-vlan
```



**Note:** For all modes VLAN stripping is supported at both the global and delivery interface levels. The **rewrite-per-policy** and **rewrite-per-filter** options have been removed in **BMF Release 6.0** because the related use cases are now supported by the **push-per-policy** and **push-per-filter** options.

The syntax for the strip VLAN ID feature is as follows:

```
controller-1(config-switch-if)# role delivery interface-name <name> [strip-no-vlan | strip-onevlan | strip-second-vlan | strip-two-vlan]
```

You can use the option to leave all VLAN tags intact, to remove the outermost tag, to remove the second (inner) tag, or to remove the outermost two tags.

By default, VLAN stripping is enabled and it removes the outer VLAN added by DMF.

To preserve the outer VLAN tag, enter the **strip-no-vlan** command, as in the following example, which preserves the outer VLAN ID for traffic forwarded to delivery interface **TOOL-PORT-1**:

```
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-no-vlan
```

When global VLAN stripping is disabled, the following commands remove the outer VLAN tag, added by DMF, on packets transmitted to the specific delivery interface **ethernet20** on **DMF-DELIVERY-SWITCH-1**:

```
controller-1(config)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch)# interface ethernet20
controller-1(config-switch-if)# role delivery interface-name TOOL-PORT-1 strip-one-vlan
```

To restore the default configuration, which is to strip the VLAN IDs from traffic to every delivery interface, enter the following command:

```
controller-1(config)# auto-delivery-interface-vlan-strip
This would enable auto delivery interface strip VLAN feature.
Existing policies will be re-computed. Enter "yes" (or "y") to continue: yes
```

As mentioned earlier, tagging and stripping VLANs as they ingress and egress DMF differs based on whether the switch uses a Trident 3 chipset. The following scenarios show how DMF behaves in different VLAN modes with various knobs set.

### Scenario 1

- VLAN mode: **Push per Policy**
- Filter interface on any switch except a Trident 3 switch
- Delivery interface on any switch
- Global VLAN stripping is **enabled**

**Table 10: Behavior of Traffic as It Egresses with Different Strip Options on a Delivery Interface**

VLAN tag type	No Configuration	strip-no-VLAN	strip-one-VLAN	strip-second-VLAN	strip-two-VLAN
	DMF policy VLAN is stripped automatically on delivery interface using default global strip VLAN added by DMF	DMF policy VLAN and customer VLAN preserved	Strips the outermost VLAN that is DMF policy VLAN	DMF policy VLAN is preserved and outermost customer VLAN is removed	Strip two VLANs, DMF policy VLAN and customer outer VLAN removed
Untagged	Packets exit DMF as untagged packets	Packets exit DMF as singly tagged packets. VLAN in the packet is DMF policy VLAN.	Packets exit DMF as untagged packets.	Packets exit DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packets exit DMF as untagged traffic.
Singly Tagged	Packets exit DMF as single-tagged traffic with customer VLAN.	Packets exit DMF as doubly tagged packets. Outer VLAN in the packet is DMF policy VLAN.	Packets exit DMF as single-tagged traffic with customer VLAN.	Packets exit DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packets exit DMF as untagged traffic.
Doubly Tagged	Packet exits DMF as doubly tagged traffic. Both VLANs are customer VLANs.	Packet exits DMF as triple-tagged packets. Outermost VLAN in the packet is the DMF policy VLAN.	Packet exits DMF as doubly tagged traffic. Both VLANs are customer VLANs.	Packet exits DMF as double-tagged packets. Outer VLAN is DMF policy VLAN, inner VLAN is inner customer VLAN in the original packet.	Packet exits DMF as singly tagged traffic. VLAN in the packet is the inner customer VLAN.

## Scenario 2

- VLAN Mode: **Push per Policy**
- Filter interface on any switch except a Trident 3 switch
- Delivery interface on any switch
- Global VLAN strip is **disabled**

**Table 11: Behavior of Traffic as It Egresses with Different Strip Options on a Delivery Interface**

VLAN tag type	No Configuration	strip-no-VLAN	strip-one-VLAN	strip-second-VLAN	strip-two-VLANs
	DMF policy VLAN and customer VLAN are preserved	DMF policy VLAN and customer VLAN are preserved	Strips only the outermost VLAN that is DMF policy VLAN	DMF policy VLAN is preserved and outer most customer VLAN is removed	Strip two VLANs, DMF policy VLAN and customer outer VLAN removed
Untagged	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as untagged packets.	Packet exits DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as untagged traffic.
Singly Tagged	Packet exits DMF as doubly tagged packets. Outer VLAN in packet is DMF policy VLAN and inner VLAN is customer outer VLAN.	Packet exits DMF as doubly tagged packets. Outer VLAN in the packet is DMF policy VLAN.	Packet exits DMF as single-tagged traffic with customer VLAN.	Packet exits DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packets exits DMF as untagged traffic.
Doubly Tagged	Packet exits DMF as triple-tagged packets. Outermost VLAN in the packet is the DMF policy VLAN.	Packet exits DMF as triple-tagged packets. Outermost VLAN in the packet is the DMF policy VLAN.	Packet exits DMF as doubly tagged traffic. Both VLANs are customer VLANs.	Packet exits DMF as doubly tagged packets. Outer VLAN is DMF policy VLAN, inner VLAN is inner customer VLAN in the original packet.	Packet exits DMF as singly tagged traffic. VLAN in the packets is the inner customer VLAN.

## Scenario 3

- VLAN Mode - **Push per Policy**
- Filter interface on a Trident 3 switch
- Delivery interface on any switch
- Global VLAN strip is **enabled**

**Table 12: Behavior of traffic as it egresses with different strip options on a delivery interface**

VLAN tag type	No Configuration	strip-no-VLAN	strip-one-VLAN	strip-second-VLAN	strip-two-VLAN
	DMF policy VLAN is stripped automatically on delivery interface using default global strip VLAN added by DMF	DMF policy VLAN and customer VLAN preserved	Strips the outermost VLAN that is DMF policy VLAN	DMF policy VLAN is preserved and outermost customer VLAN is removed	Strip two VLANs , DMF policy VLAN and customer outer VLAN removed
Untagged	Packet exits DMF as untagged packets.	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as untagged packets.	Packet exits DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as untagged traffic.
Singly Tagged	Packet exits DMF as single-tagged traffic with customer VLAN.	Packet exits DMF as doubly tagged packets. Outer VLAN in the packet is DMF policy VLAN.	Packet exits DMF as single tagged traffic with customer VLAN.	Packet exits DMF as single-tagged traffic. VLAN in the packet is DMF policy VLAN.	Packet exits DMF as untagged traffic.
Doubly Tagged	Packet exits DMF as singly tagged traffic. VLAN in the packet is the inner customer VLAN	Packet exits DMF as doubly tagged traffic. Outer customer VLAN is replaced by DMF policy VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the packet is the inner customer VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the packet is the DMF policy VLAN.	Packet exits DMF as untagged traffic.

**Scenario 5**

- VLAN Mode - ***Push per Filter***
- Filter interface on any switch
- Delivery interface on any switch
- Global VLAN strip is **enabled**

**Table 13: Behavior of Traffic as It Egresses with Different Strip Options on a Delivery Interface**

VLAN tag type	No Configuration	strip-no-VLAN	strip-one-VLAN	strip-second-VLAN	strip-two-VLAN
	DMF filter VLAN is stripped automatically on delivery interface using global strip VLAN added by DMF.	DMF filter VLAN and customer VLAN preserved.	Strips the outermost VLAN that is DMF filter VLAN.	DMF filter VLAN is preserved and outermost customer VLAN is removed.	Strip two VLANs, DMF filter interface VLAN and customer outer VLAN removed.
Untagged	Packet exits DMF as untagged packets.	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged packets.	Packet exits DMF as single tagged traffic. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged traffic.
Singly Tagged	Packet exits DMF as singly tagged traffic. VLAN in the packet is the customer VLAN.	Packet exits DMF as doubly tagged packets. Outer VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the packet is the customer VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged traffic.
Doubly Tagged	Packet exits DMF as singly tagged traffic. VLAN in the policy is the inner customer VLAN.	Packet exits DMF as doubly tagged traffic. Outer customer VLAN is replaced by DMF filter interface VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the policy is the inner customer VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the policy is the DMF filter interface VLAN.	Packet exits DMF as untagged traffic.

#### Scenario 6

- VLAN Mode - **Push per Filter**
- Filter interface on any switch
- Delivery interface on any switch
- Global VLAN strip is **disabled**



**Table 14: Behavior of Traffic as It Egresses with Different Strip Options on a Delivery Interface**

VLAN tag type	No Configuration	strip-no-VLAN	strip-one-VLAN	strip-second-VLAN	strip-two-VLAN
	DMF filter VLAN is stripped automatically on delivery interface using global strip VLAN added by DMF.	DMF filter VLAN and customer VLAN preserved.	Strips the outermost VLAN that is DMF filter VLAN.	DMF filter VLAN is preserved and outermost customer VLAN is removed.	Strip two VLANs, DMF filter interface VLAN and customer outer VLAN removed.
Untagged	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as singly tagged packets. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged packets.	Packet exits DMF as singly tagged traffic. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged traffic.
Singly Tagged	Packet exits DMF as doubly tagged traffic. Outer VLAN in the packet is DMF filter VLAN and inner VLAN is the customer VLAN.	Packet exits DMF as doubly tagged packets. Outer VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as single tagged traffic. VLAN in the packet is the customer VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the packet is DMF filter interface VLAN.	Packet exits DMF as untagged traffic.
Doubly Tagged	Packet exits DMF as doubly tagged traffic. Outer customer VLAN is replaced by DMF filter interface VLAN.	Packet exits DMF as doubly tagged traffic. Outer customer VLAN is replaced by DMF filter interface VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the policy is the inner customer VLAN.	Packet exits DMF as singly tagged traffic. VLAN in the policy is the DMF filter interface VLAN.	Packet exits DMF as untagged traffic.

## 10.5.2 Auto VLAN Strip using the GUI

### Auto VLAN Strip

1. A **toggle button** controls the configuration of this feature. Locate the corresponding card and click the **toggle** switch.

**Figure 10-16: Toggle Switch**

2. A confirm window pops up, displaying the corresponding prompt message. Click the **Enable** button to confirm the configuration changes or the **Cancel** button to cancel the configuration. Conversely, to disable the configuration, click **Disable**.

**Figure 10-17: Confirm / Enable**



3. Review any warning messages that appear in the confirmation window during the configuration process.

**Figure 10-18: Warning Message - Changing**



The following feature sets work in the same manner as the **Auto VLAN Strip** feature described above.

- **CRC Check**
- **Custom Priority**
- **Inport Mask**
- **Policy Overlap Limit Strict**
- **Retain User Policy VLAN**
- **Tunneling**

## 10.6 CRC Check

If the Switch CRC option is enabled, which is the default, each DMF switch drops incoming packets which enter the fabric with a CRC error. The switch generates a new CRC if the incoming packet was modified using one of the options that modify the original CRC checksum, which include the push VLAN, rewrite VLAN, strip VLAN, and L2 GRE tunnel options.



**Note:** The Switch CRC option must be enabled to use the DMF tunneling feature.

If the Switch CRC option is disabled, DMF switches do not check the CRC of incoming packets and do not drop packets with CRC errors. Also, switches do not generate a new CRC if the packet is modified. This mode is useful, if packets with CRC errors need to be delivered to a destination tool unmodified for analysis. If users disable the Switch CRC option, make sure that the destination tool does not drop packets having CRC errors. Also recognize that CRC errors will be caused by modification of packets by DMF options so that these CRC errors are not mistaken for CRC errors from the traffic source.



**Note:** When the Switch CRC option is disabled, packets going to the Service Node or Recorder Node are dropped because a new CRC is not calculated when push-per-policy or push-per-filter adds a VLAN tag.

Enable and disable CRC Check using the steps described in the following topics.

- [CRC Check using the CLI](#)
- [CRC Check using the GUI](#)

## 10.6.1 CRC Check using the CLI

If the Switch CRC option is enabled, which is the default, each DMF switch drops incoming packets which enter the fabric with a CRC error. The switch generates a new CRC if the incoming packet was modified using one of the options that modify the original CRC checksum, which include the push VLAN, rewrite VLAN, strip VLAN, and L2 GRE tunnel options.



**Note:** The Switch CRC option must be enabled to use the DMF tunneling feature.

If the Switch CRC option is disabled, DMF switches do not check the CRC of incoming packets and do not drop packets with CRC errors. Also, switches do not generate a new CRC if the packet is modified. This mode is useful if packets with CRC errors need to be delivered to a destination tool unmodified for analysis. If you disable the Switch CRC option, make sure that the destination tool does not drop packets having CRC errors. Also recognize that CRC errors will be caused by modification of packets by DMF options so that these CRC errors are not mistaken for CRC errors from the traffic source.

To disable the Switch CRC option, enter the following command from config mode:

```
controller-1(config)# no crc
Disabling CRC mode may cause problems to tunnel interface. Enter "yes" (or "y")
to continue: y
```

To re-enable the Switch CRC option after it has been disabled, enter the following command from config mode:

```
controller-1(config)# crc
Enabling CRC mode would cause packets with crc error dropped. Enter "yes" (or
"y") to continue: y
```



**Tip:** To enable/disable the CRC through the GUI, refer to the chapter, [Check CRC using the GUI](#).




**Note:** When the Switch CRC option is disabled, packets going to the service node or recorder node are dropped because a new CRC is not calculated when **push-per-policy** or **push-per-filter** adds a VLAN tag.

## 10.6.2 CRC Check using the GUI

From the DMF Features page, proceed to the CRC Check feature card and perform the following steps to enable the feature.

1. Select the **CRC Check** card.

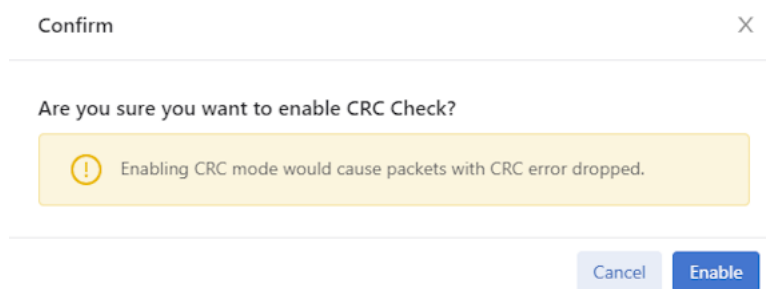
**Figure 10-19: CRC Check Disabled**



2. Toggle the CRC Check switch to **On**.

3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

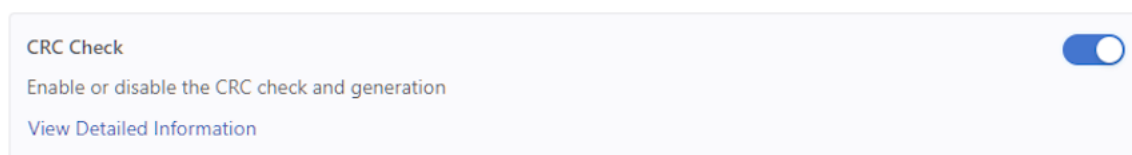
**Figure 10-20: Enable CRC Check**



A confirmation dialog box titled "Confirm" with a close button (X) in the top right corner. The main text asks, "Are you sure you want to enable CRC Check?". Below this is a yellow warning box containing an exclamation mark icon and the text: "Enabling CRC mode would cause packets with CRC error dropped." At the bottom right, there are two buttons: "Cancel" and "Enable".

4. CRC Check is running.

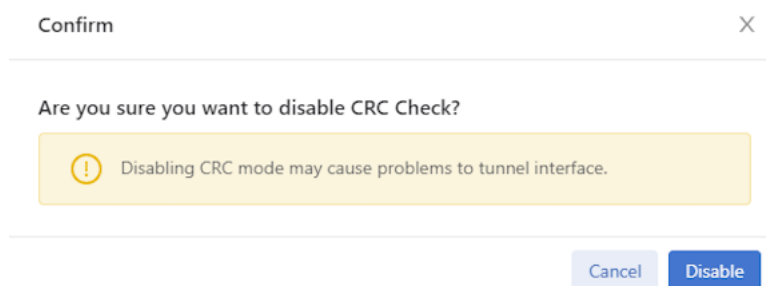
**Figure 10-21: CRC Check Enabled**



A feature card for "CRC Check". It has a title "CRC Check" and a toggle switch that is currently turned on (blue). Below the title is the text "Enable or disable the CRC check and generation". At the bottom left, there is a link "View Detailed Information".

5. To disable the feature, toggle the CRC Check switch to **Off**. Click **Disable** and confirm.

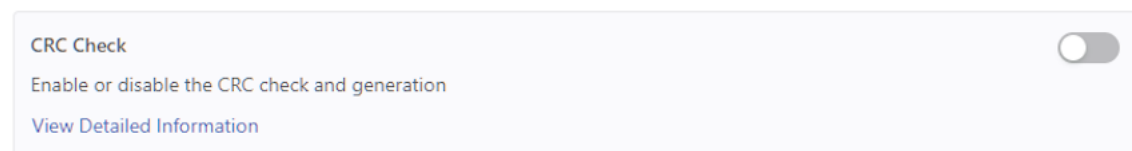
**Figure 10-22: Disable CRC Check**



A confirmation dialog box titled "Confirm" with a close button (X) in the top right corner. The main text asks, "Are you sure you want to disable CRC Check?". Below this is a yellow warning box containing an exclamation mark icon and the text: "Disabling CRC mode may cause problems to tunnel interface." At the bottom right, there are two buttons: "Cancel" and "Disable".

The feature card updates with the status.

**Figure 10-23: CRC Check Disabled**



A feature card for "CRC Check". It has a title "CRC Check" and a toggle switch that is currently turned off (grey). Below the title is the text "Enable or disable the CRC check and generation". At the bottom left, there is a link "View Detailed Information".

## 10.7 Custom Priority

When custom priorities are allowed, non-admin users may assign policy priorities between 0 and 100 (the default value). However, when custom priorities are not allowed, the default priority 100 will be automatically assigned to non-admin users' policies.

Enable and disable Custom Priority using the steps described in the following topics.

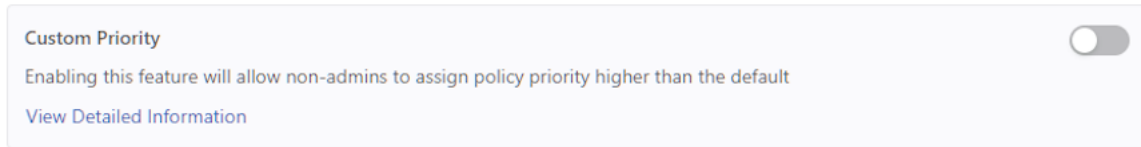
- [Custom Priority using the CLI](#)
- [Custom Priority using the GUI](#)

## 10.7.1 Configuring Custom Priority using the GUI

From the DMF Features page, proceed to the Custom Priority feature card and perform the following steps to enable the feature.


1. Select the **Custom Priority** card.

**Figure 10-24: Custom Priority Disabled**



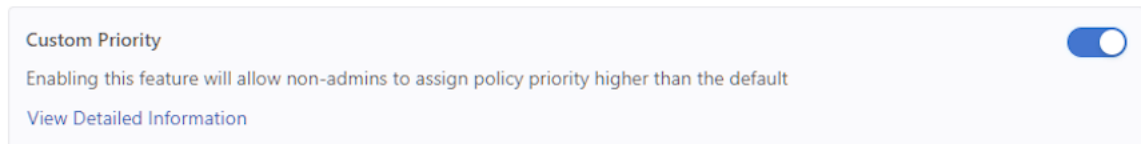
2. Toggle the Custom Priority switch to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 10-25: Enable Custom Priority**




4. Custom Priority is running.

**Figure 10-26: Custom Priority Enabled**




5. To disable the feature, toggle the **Custom Priority** switch to **Off**. Click **Disable** and confirm.

**Figure 10-27: Disable Custom Priority**



The feature card updates with the status.

**Figure 10-28: Custom Priority Disabled**



## 10.7.2 Configuring Custom Priority using the CLI

To enable the Custom Priority, enter the following command:

```
controller-1(config)# allow-custom-priority
```

To disable the Custom Priority, enter the following command:

```
controller-1(config)# no allow-custom-priority
```

## 10.8 Device Deployment Mode

Complete the fabric switch installation in one of the following two modes:

- **Layer 2 Zero Touch Fabric (L2ZTF, Auto-discovery switch provisioning mode)**

In this mode, which is the default, Switch ONIE software automatically discovers the controller via IPv6 local link addresses and downloads and installs the appropriate Switch Light OS image from the controller. This installation method requires all the fabric switches and the DMF controller to be in the same Layer 2 network (IP subnet). Also, suppose the fabric switches need IPv4 addresses to communicate with SNMP or other external services. In that case, users must configure IPAM, which provides the controller with a range of IPv4 addresses to allocate to the fabric switches.

- **Layer 3 Zero Touch Fabric (L3ZTF, Preconfigured switch provisioning mode)**

When fabric switches are in a different Layer 2 network from the controller, log in to each switch individually to configure network information and download the ZTF installer. Subsequently, the switch automatically downloads Switch Light OS from the controller. This mode requires communication between the controller and the fabric switches to occur using IPv4 addresses, and no IPAM configuration is required.

The following table summarizes the requirements for installation using each mode:

Requirement	Layer 2 mode	Layer 3 mode
Any switch in a different subnet from the controller	No	Yes
IPAM configuration for SNMP and other IPv4 services	Yes	No
IP address assignment	IPv4 or IPv6	IPv4
Refer to this section (in User Guide)	Using L2 ZTF (Auto-Discovery) Provisioning Mode	Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode


All the fabric switches in a single fabric must be installed using the same mode. If users have any fabric switches in a different IP subnet than the controller, users must use Layer 3 mode for installing all the switches, even those in the same Layer 2 network as the controller. Installing switches in mixed mode, with some switches using ZTF in the same Layer 2 network as the controller, while other switches in a different subnet are installed manually or using DHCP is unsupported.

### 10.8.1 Configuring Device Deployment Mode using the GUI

From the DMF Features page, proceed to the Device Deployment Mode feature card and perform the following steps to manage the feature.


1. Select the **Device Deployment Mode** card.

**Figure 10-29: Device Deployment Mode - Auto Discovery**



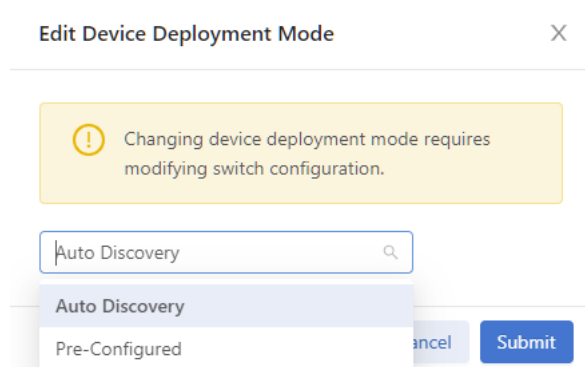
2. Enter the edit mode using the **pencil icon**.

**Figure 10-30: Configure Device Deployment Mode**




3. Change the switching mode as required using the drop-down menu. The default mode is **Auto Discovery**.

**Figure 10-31: Device Deployment Mode Options**




**Figure 10-31: Device Deployment Mode - Pre-Configured Option**



4. Click **Submit** and confirm the operation when prompted.

5. The Device Deployment Mode status updates.

**Figure 10-33: Device Deployment Mode - Status Update**



## 10.8.2 Configuring Device Deployment Mode using the CLI

Device Deployment Mode has two options, select desired option either auto-discovery or pre-configured as shown below:

```
controller-1(config)# deployment-mode auto-discovery
```

Changing device deployment mode requires modifying switch configuration. Enter "yes" (or "y") to continue: **y**

```
controller-1(config)# deployment-mode pre-configured
```

Changing device deployment mode requires modifying switch configuration. Enter "yes" (or "y") to continue: **y**

## 10.9 Inport Mask

Enable and disable Inport Mask using the steps described in the following topics.


- [Configure Inport Mask using the CLI](#)
- [Configure Inport Mask using the GUI](#)

### 10.9.1 InPort Mask using the CLI

DANZ Monitoring Fabric implements multiple flow optimizations to reduce the number of flows that are programmed in the DMF switch TCAM space. This feature enables effective usage of TCAM space, and it is on by default.

When this feature is off, TCAM rules are applied for each ingress port belonging to the same policy. For example, in the following topology, if a policy were configured with **10** match rules and filter-interface as **F1** and **F2**, then **20** (**10** for **F1** and **10** for **F2**) TCAM rows were consumed.

**Figure 10-34: Simple Inport Mask Optimization**




With inport mask optimization, only **10** rules are consumed. This feature optimizes TCAM usage at every level (filer, core, delivery) in the DMF network.



Consider the more complex topology illustrated below:

**Figure 10-35: Complex Inport Mask Optimization**



In this topology, if a policy has  $N$  rules, without in-port optimization, the policy will consume  $3N$  at **Switch 1**,  $3N$  at **Switch 2**, and  $2N$  at **Switch 3**. With the in-port optimization feature enabled, the policy consumes only  $N$  rules at each switch.

However, this feature loses a level of granularity in the statistics available because there is only one set of flow mods for multiple filter ports per switch. Statistics without this feature are maintained per filter port per policy.

With inport optimization enabled, the statistics are combined for all input ports sharing rules on that switch. You can obtain filter port statistics for different flow mods for each filter port. However, this requires disabling inport optimization, which is enabled by default.

To disable the inport optimization feature, enter the following command from config mode:


```
controller-1(config)# controller-1(config)# no inport-mask
```

## 10.9.2 Inport Mask using the GUI

From the DMF Features page, proceed to the Inport Mask feature card and perform the following steps to enable the feature.

1. Select the **Inport Mask** card.

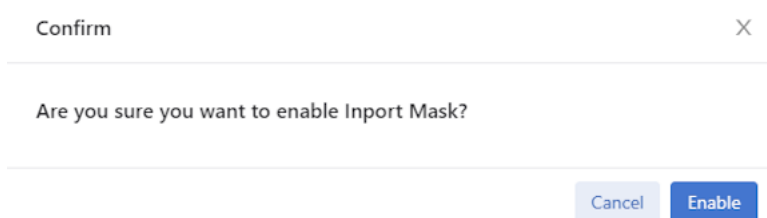
**Figure 10-36: Inport Mask Disabled**



2. Toggle the Inport Mask switch to **On**.

3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

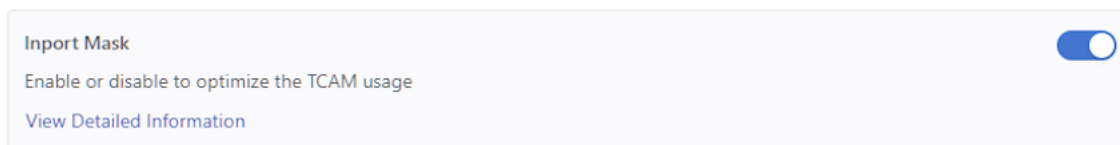
**Figure 10-37: Enable Inport Mask**



A confirmation dialog box with a title bar containing 'Confirm' and a close button 'X'. The main text asks 'Are you sure you want to enable Inport Mask?'. At the bottom, there are two buttons: 'Cancel' and 'Enable'.

4. Inport Mask is running.

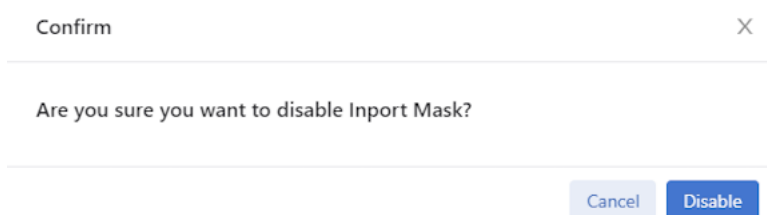
**Figure 10-38: Inport Mask Enabled**



A feature card for 'Inport Mask'. It has a title 'Inport Mask' and a toggle switch that is turned on (blue). Below the title, it says 'Enable or disable to optimize the TCAM usage'. At the bottom, there is a link 'View Detailed Information'.

5. To disable the feature, toggle the Inport Mask switch to **Off**. Click **Disable** and confirm.

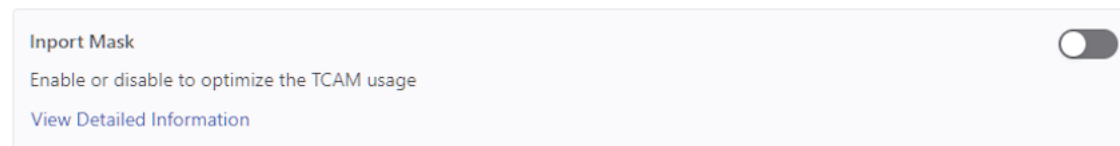
**Figure 10-39: Disable Inport Mask**



A confirmation dialog box with a title bar containing 'Confirm' and a close button 'X'. The main text asks 'Are you sure you want to disable Inport Mask?'. At the bottom, there are two buttons: 'Cancel' and 'Disable'.

The feature card updates with the status.

**Figure 10-40: Inport Mask Disabled**



A feature card for 'Inport Mask'. It has a title 'Inport Mask' and a toggle switch that is turned off (grey). Below the title, it says 'Enable or disable to optimize the TCAM usage'. At the bottom, there is a link 'View Detailed Information'.

## 10.10 Match Mode

Switches have finite hardware resources available for packet matching on aggregated traffic streams. This resource allocation is relatively static and must be configured in advance. The DANZ Monitoring Fabric supports three allocation schemes, which are referred to as switching (match) modes:

- **L3-L4 mode (default mode):** With L3-L4 mode, fields other than **src-mac** and **dst-mac** can be used for specifying policies. If there are no policies using **src-mac** or **dst-mac**, using the L3-L4 mode allows a larger number of match rules per switch.
- **Full-match mode:** With full-match mode, all matching fields, including **src-mac** and **dst-mac**, can be used while specifying policies.
- **L3-L4 Offset mode:** L3-L4 offset mode allows matching beyond the L4 header up to 128 bytes from the beginning of the packet. The number of matches per switch in this mode is the same as in full-match mode. As with L3-L4 mode, matches using **src-mac** and **dst-mac** are not permitted.



**Note:** Changing switching modes causes all fabric switches to disconnect and reconnect with the controller. Also, all existing policies will be reinstalled. The switching mode applies to all DMF switches in the DANZ Monitoring Fabric. You can switch between modes, but any match rules that are incompatible with the new mode will fail.

### 10.10.1 Setting the Match Mode Using the CLI

To use the CLI to set the **match** mode, enter the following command:

```
controller-1(config)# match-mode {full-match | l3-l4-match | l3-l4-offset-match}
```

For example, the following command sets the **match** mode to **full-match** mode:


```
controller-1(config)# match-mode full-match
```

### 10.10.2 Setting the Match Mode Using the GUI

From the DMF Features page, proceed to the Match Mode feature card and perform the following steps to enable the feature.


1. Select the **Match Mode** card.

**Figure 10-41: L3-L4 Match Mode**




2. Enter the edit mode using the **pencil icon**.

**Figure 10-42: Configure Switching Mode**



3. Change the switching mode as required using the drop-down menu. The default mode is **L3-L4 Match**.

**Figure 10-43: L3-L4 Match Options**



4. Click **Submit** and confirm the operation when prompted.



**Note:** If the existing configuration of the monitoring fabric is incompatible with the specified switching mode, an error message is displayed.

## 10.11 Retain User Policy VLAN

Enable and disable Retain User Policy VLAN using the steps described in the following topics.

- [Retain User Policy VLAN using the CLI](#)
- [Retain User Policy VLAN using the GUI](#)

### 10.11.1 Retain User Policy VLAN using the CLI

This feature will send traffic to a delivery interface with the user policy VLAN tag instead of the overlap dynamic policy VLAN tag for traffic matching the dynamic overlap policy only. This feature is supported only in **push-per-policy** mode. For example, policy **P1** with filter interface **F1** and delivery interface **D1**, and policy **P2** with filter interface **F1** and delivery interface **D2**, and overlap dynamic policy **P1\_o\_P2** is created when the overlap policy condition is met. In this case, the overlap dynamic policy is created with filter interface **F1** and delivery interfaces **D1** and **D2**. The user policy **P1** assigns a VLAN (**VLAN 10**) and **P2** assigns a VLAN (**VLAN 20**) when it is created, and the overlap policy also assigns a VLAN (**VLAN 30**) when it is dynamically created. When this feature is enabled, traffic forwarded to **D1** will have a policy VLAN tag of **P1 (VLAN 10)** and **D2** will have a policy VLAN tag of policy **P2 (VLAN 20)**. When this feature is disabled, traffic forwarded to **D1** and **D2** will have the dynamic overlap policy VLAN tag (**VLAN 30**). By default, this feature is disabled.

#### Feature Limitations:

- A overlap dynamic policy will fail when the overlap policy has filter (**F1**) and delivery interface (**D1**) on the same switch (**switch A**) and another delivery interface (**D2**) on another switch (**switch B**).
- Post-to-delivery dynamic policy will fail when it has a filter interface (**F1**) and a delivery interface (**D1**) on the same switch (**switch A**) and another delivery interface (**D2**) on another switch (**switch B**).
- Overlap policies may be reinstalled when a fabric port goes up or down when this feature is enabled.
- Double-tagged VLAN traffic is not supported and will be dropped at the delivery interface.
- Tunnel interfaces are not supported with this feature.
- Only IPv4 traffic is supported; other non-IPv4 traffic will be dropped at the delivery interface.
- Delivery interfaces with IP addresses (L3 delivery interfaces) are not supported.
- This feature is not supported on EOS switches (Arista 7280 switches).
- Delivery interface statistics may not be accurate when displayed using the **sh policy** command. This will happen when policy **P1** has **F1, D1, D2** and policy **P2** has **F1, D2**. In this case, overlap policy

**P1** o **P2** will be created with delivery interfaces **D1**, **D2**. Since **D2** is in both policies **P1** and **P2**, overlap traffic will be forwarded to **D2** with both the **P1** policy VLAN and the **P2** policy VLAN. The **sh policy <policy\_name>** command will not show this doubling of traffic on delivery interface **D2**. Delivery interface statistics will show this extra traffic forwarded out of the delivery interface.

To enable this feature, enter the following command:

```
controller-1(config)# retain-user-policy-vlan
This will enable retain-user-policy-vlan feature. Non-IP packets will be
dropped at delivery. Enter
"yes" (or "y") to continue: yes
```

To see the current Retain Policy VLAN configuration, enter the following command:

```
controller-1> show fabric
~~~~~ Aggregate Network State ~~~~~
Number of switches : 14
Import masking : True
Number of unmanaged services : 0
Number of switches with service interfaces : 0
Match mode : 13-14-offset-match
Number of switches with delivery interfaces : 11
Filter efficiency : 1:1
Uptime : 4 days, 8 hours
Max overlap policies (0=disable) : 10
Auto Delivery Interface Strip VLAN : True
Number of core interfaces : 134
State : Enabled
Max delivery BW (bps) : 2.18Tbps
Health : unhealthy
Track hosts : True
Number of filter interfaces : 70
Number of policies : 101
Start time : 2022-02-28 16:18:01.807000 UTC
Number of delivery interfaces : 104
Retain User Policy Vlan : True
```

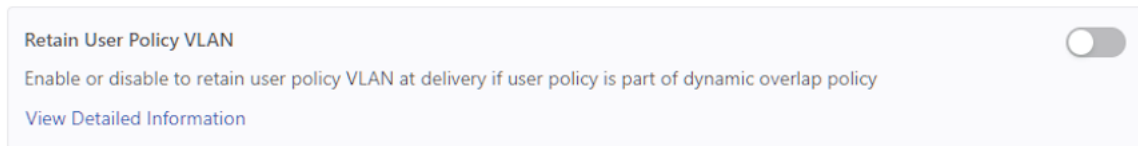
This feature can be used with the **strip-second-vlan** option during delivery interface configuration to preserve the outer DMF fabric policy VLAN, strip the inner VLAN of traffic forwarded to a tool, or the **strip-no-vlan** option during delivery interface configuration.

## 10.11.2 Retain User Policy VLAN using the GUI

From the DMF Features page, proceed to the Retain User Policy VLAN feature card and perform the following steps to enable the feature.

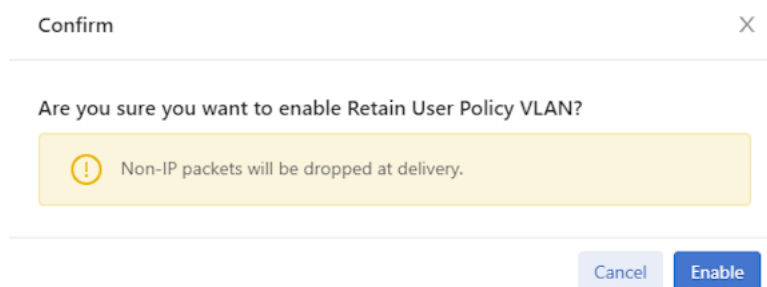
1. Select the **Retain User Policy VLAN** card.

**Figure 10-44: Retain User Policy VLAN Disabled**



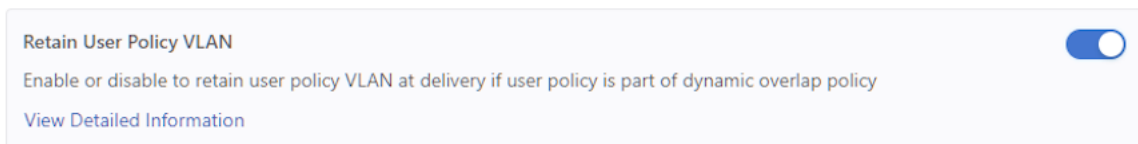
2. Toggle the Retain User Policy VLAN switch to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 10-45: Enable Retain User Policy VLAN**



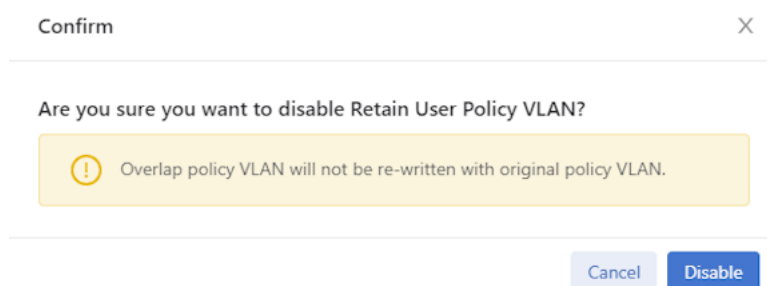
4. Retain User Policy VLAN is running.

**Figure 10-46: Retain User Policy VLAN Enabled**



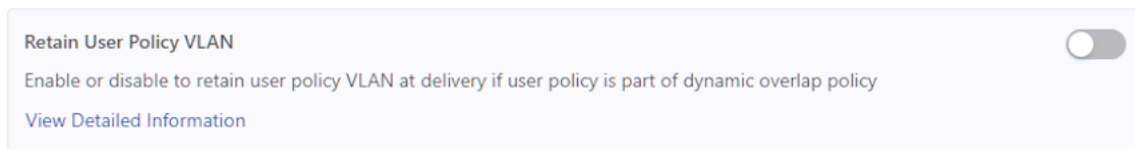
5. To disable the feature, toggle the Retain User Policy VLAN switch to **Off**. Click **Disable** and confirm.

**Figure 10-47: Disable Retain User Policy VLAN**



The feature card updates with the status.

**Figure 10-48: Retain User Policy VLAN Disabled**



## 10.12 Tunneling

For more information about Tunneling please refer to the [Understanding Tunneling](#) section.

Enable and disable Tunneling using the steps described in the following topics.


- [Configure Tunneling using the GUI](#)
- [Configure Tunneling using the CLI](#)

### 10.12.1 Configuring Tunneling using the GUI

From the DMF Features page, proceed to the Tunneling feature card and perform the following steps to enable the feature.

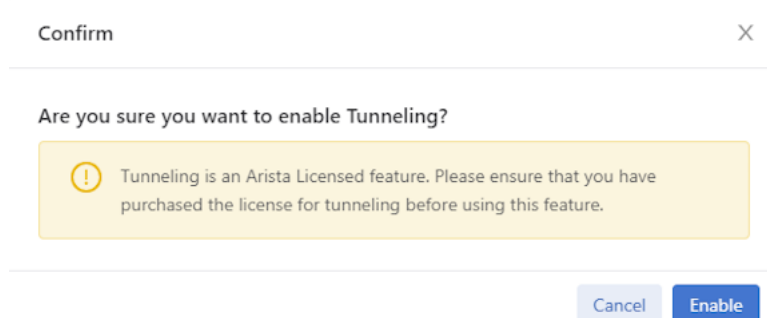
1. Select the **Tunneling** card.

**Figure 10-49: Tunneling Disabled**



2. Toggle the Tunneling switch to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

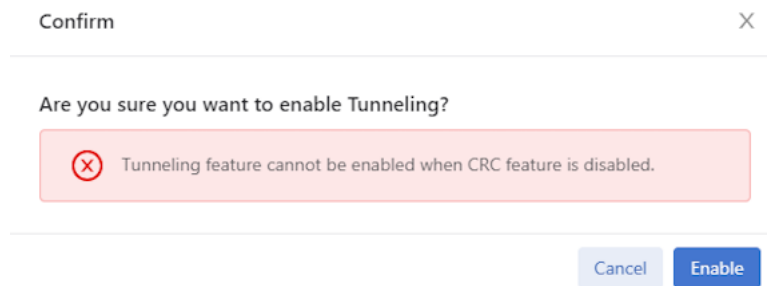
**Figure 10-50: Enable Tunneling**



**Note:** CRC Check must be running before attempting to enable Tunneling. An error message displays if CRC Check is not enabled. Proceeding to click Enable results in a validation error


message. Refer to the [CRC Check](#) section for more information on configuring the CRC Check feature.

**Figure 10-51: CRC Check Warning Message**



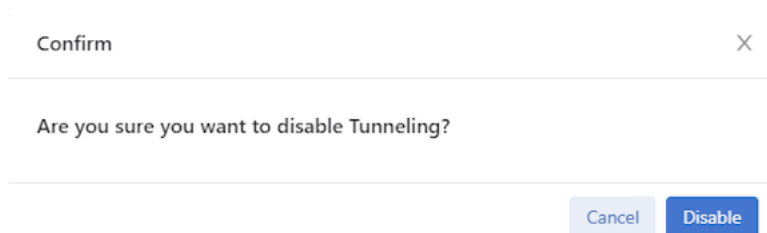
4. Tunneling VLAN is running.

**Figure 10-52: Tunneling Enabled**




5. To disable the feature, toggle the Tunneling switch to **Off**. Click **Disable** and confirm.

**Figure 10-53: Disable Tunneling**



The feature card updates with the status.

**Figure 10-54: Tunneling VLAN Disabled**



## 10.12.2 Configuring Tunneling using the CLI

To enable the Tunneling, enter the following command:

```
controller-1(config)# tunneling

Tunneling is an Arista Licensed feature.
Please ensure that you have purchased the license for tunneling before using
this feature.
Enter "yes" (or "y") to continue: y
controller-1(config)#
```



To disable the Tunneling, enter the following command:

```
controller-1(config)# no tunneling
This would disable tunneling feature? Enter "yes" (or "y") to continue: y
controller-1(config)#
```

## 10.13 VLAN Preservation

In DMF, meta data is appended to the packets forwarded by the fabric to a tool attached to a delivery interface. This meta data is encoded primarily in the outer VLAN tag of the packets.

By default (using the auto-delivery-strip feature), this outer VLAN tag is always removed on egress upon delivery to a tool.

The VLAN preservation feature introduces a choice to selectively preserve a packet's outer VLAN tag instead of stripping or preserving all of it.

VLAN preservation works in both push-per-filter and push-per-policy mode for auto-assigned and user-configured VLANs.



**Note:** VLAN preservation applies to switches running SWL OS and does not apply to switches running EOS.

This functionality only supports 2000 VLAN ID and port combinations per switch.

Support for VLAN preservation is on select Broadcom® switch ASICs. Ensure your switch model supports this feature before attempting to configure it.

### 10.13.1 Using the CLI to Configure VLAN Preservation

VLAN preservation can be configured at two levels: global and local. A local configuration can override the global configuration.

#### Global Configuration

Enable VLAN preservation globally using the `vlan-preservation` command from the config submode to apply a global configuration.

```
(config)# vlan-preservation
```

Two options exist while in the `config-vlan-preservation` submode:

- `preserve-user-configured-vlans`
- `preserve-vlans`

Use the help function to list the options by entering a ? (question mark).

```
(config-vlan-preservation)# ?
Commands:
preserve-user-configured-vlans    Preserve all user-configured VLANs for all
delivery interfaces
preserve-vlan                    Configure VLAN ID to preserve for all
delivery interfaces
```

Use the **preserve-user-configured-vlans** option to preserve all user-configured VLANs. The packets with the user-configured VLANs will have their fabric-applied VLAN tags preserved even after leaving the respective delivery interface.

```
(config-vlan-preservation)# preserve-user-configured-vlans
```

Use the **preserve-vlan** option to specify and preserve a particular VLAN ID. Any VLAN ID may be provided. In the following example, the packets with VLAN ID 100 or 200 will have their fabric-applied VLAN tags preserved upon delivery to the tool.

```
(config-vlan-preservation)# preserve-vlan 100  
(config-vlan-preservation)# preserve-vlan 200
```

## Local Configuration

This feature applies to delivery and both-filter-and-delivery interface roles.

Fabric-applied VLAN tag preservation can be enabled locally on each delivery interface as an alternative to the global VLAN preservation configuration. To enable this functionality locally, enter the following configuration submode using the **if-vlan-preservation** command to specify either one of the two available options. Use the help function to list the options by entering a ? (question mark).

```
(config-switch-if)# if-vlan-preservation  
(config-switch-if-vlan-preservation)# ?  
Commands:  
preserve-user-configured-vlans      Preserve all user-configured VLANs for all  
delivery interfaces  
preserve-vlan                        Configure VLAN ID to preserve for all  
delivery interfaces
```

Use the **preserve-user-configured-vlans** option to preserve all user-configured VLAN IDs in **push-per-policy** or **push-per-filter** mode on a selected delivery interface. All packets egressing such delivery interface will have their user-configured fabric VLAN tags preserved.

```
(config-switch-if-vlan-preservation)# preserve-user-configured-vlans
```

Use the **preserve-vlan** option to specify and preserve a particular VLAN ID. For example, if any packets with VLAN ID 100 or 300 egress the selected delivery interface, VLAN IDs 100 and 300 will be preserved.

```
(config-switch-if-vlan-preservation)# preserve-vlan 100  
(config-switch-if-vlan-preservation)# preserve-vlan 300
```



**Note:** Any local vlan-preservation configuration overrides the global configuration for the selected interfaces by default.

On an MLAG delivery interface, the local configuration follows the same model, as shown below.

```
(config-mlag-domain-if)# if-vlan-preservation member role  
(config-mlag-domain-if)# if-vlan-preservation  
(config-mlag-domain-if-vlan-preservation)# preserve-user-configured-vlans  
preserve-vlan
```

To disable selective VLAN preservation for a particular delivery or both-filter-and-delivery interface, use the following command to disable the feature's global and local configuration for the selected interface:

```
(config-switch-if)# role delivery interface-name del  
<cr>                no-analytics                strip-no-vlan                strip-second-  
vlan  
ip-address          no-vlan-preservation  strip-one-vlan                strip-two-vla  
n  
(config-switch-if)# role delivery interface-name del no-vlan-preservation
```

## CLI Show Commands

The following show command displays the name of the device on which VLAN preservation is enabled and the information about which VLAN is preserved on specific selected ports. Use the data in this table primarily for debugging purposes.

```
# show switch all table vlan-preserve
# Vlan-preserve Device name Entry key
-|-----|-----|-----|
1 0          delivery1  VlanVid(0x64), Port(6)
2 0          filter1   VlanVid(0x64), Port(6)
3 0          core1     VlanVid(0x64), Port(6)
```

## 10.13.2 Using the GUI to Configure VLAN Preservation

VLAN preservation can be configured at two levels: global and local. A local configuration can override the global configuration. Follow the steps outlined below to configure a Global Configuration (steps 1 - 4), Local Configuration (steps 5-7) or an MLAG Delivery Interface configuration within a MLAG domain (step 8).

### 1. Global Configuration

- To view or edit the global configuration, navigate to the DMF Features page by clicking the gear icon in the top right of the navigation bar.

Figure 10-55: DMF 8.4 Menu Bar



The DMF Feature page is new in DMF release 8.4. It allows for managing fabric-wide settings for DMF.

- Scroll to the **VLAN Preservation** card.

Figure 10-56: DMF Features Page

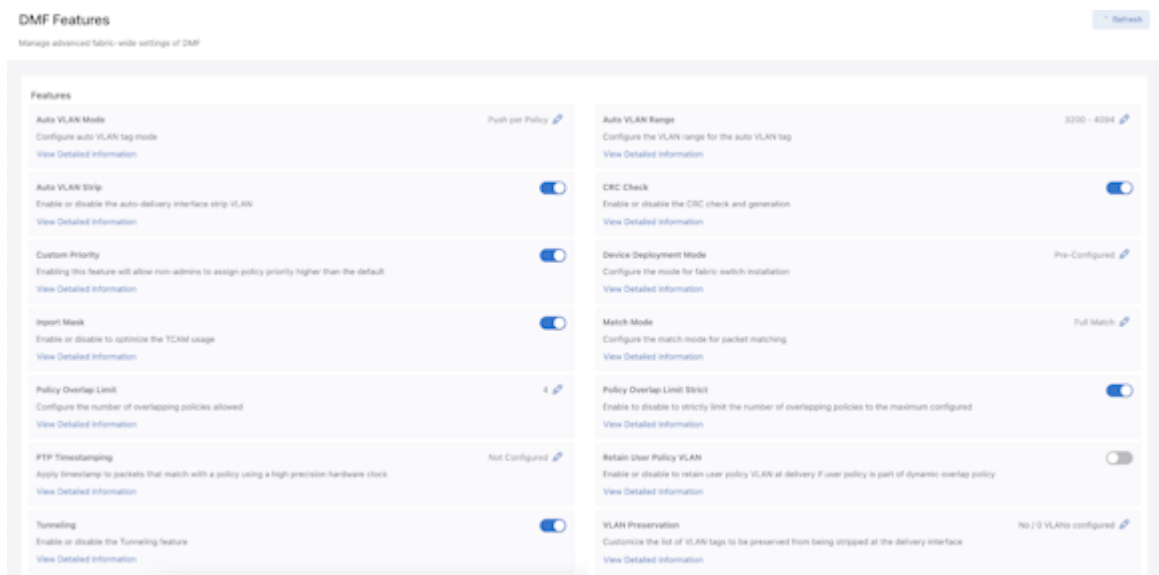
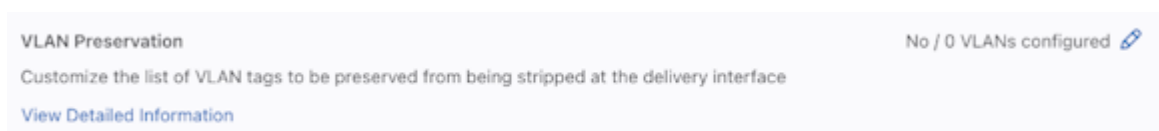


Figure 10-57: VLAN Preservation Card



3. Click the **Edit** button (pencil icon) to configure or modify the global VLAN Preservation feature settings.

**Figure 10-58: Edit VLAN Preservation Configuration**

Edit VLAN Preservation Configuration

X

Preserve User Configured VLANs

VLAN Preserve

⊕ Add VLAN

VLAN

Single

Value

#

VLAN

Range

From

#

To

#

Cancel

Submit

The edit modal screen has two input sections:

- Toggle **on** or off the **Preserve User Configured VLANs**.
  - Enter the parameters for **VLAN Preserve** using the following functions:
    - Use the **Add VLAN** button to add VLAN IDs.
    - Select the **Single** VLAN type drop-down to add a single VLAN ID.
    - Select the **Range** VLAN type drop-down to add a continuous VLAN ID range.
    - Use the **Trash** button (delete) to delete a single VLAN ID or a VLAN ID range.
4. Click the **Submit** button to save the configuration.
2. **Local Configuration**

- The VLAN Preservation configuration can be applied per-delivery interface while configuring or editing a delivery or filter-and-delivery interface in the **Monitoring Interfaces** page and **Monitoring Interfaces > Delivery Interfaces** page.

**Figure 10-59: Monitoring Interfaces Delivery Interface Create Interface**

- The following inputs are available for the local feature configuration:
  - Toggle the **Preserve User Configured VLANs** button to on. Use this option to preserve all user-configured VLAN IDs in push-per-policy or push-per-filter mode on a selected delivery interface. The packets with the user-configured VLANs will have their fabric-applied VLAN tags preserved even after leaving the respective delivery interface.
  - VLAN Preserve**. Use the + and - icon buttons to respectively add and remove VLAN IDs.
  - Toggle the **VLAN Preservation** button to **on**. Disabling this option will ignore this feature configuration given globally/locally for this delivery interface. This is enabled by default.
- Click the **Save** button to save the configuration.
- VLAN Preservation for MLAG Delivery Interfaces**

- Configure VLAN preservation for MLAG delivery interfaces using **Fabric > MLAGs** page while configuring a MLAG Domain toggling the **VLAN Preservation** and **Preserve User Configured VLANs** switches to **on** (as required).

**Figure 10-60: Create MLAG Domain**

**Figure 10-61: MLAG VLAN Preservation & Preserve User Configured VLANs**

### 10.13.3 Troubleshooting

Use the following commands to troubleshoot the scenario in which a tool attached to a delivery interface expects a packet with a preserved VLAN tag, but instead, there is no tag attached to it; double-check the following.

- Verify the running config and review if the VLAN preservation configuration is enabled for that VLAN ID and on that delivery interface.

```
(config-vlan-preservation)# show running-config | grep "preserve"
! vlan-preservation
vlan-preservation
preserve-vlan 100
```

- Verify the **show switch switch-name table vlan-preserve** output. It displays the ports and VLAN ID combinations that are enabled.

```
(config-policy)# show switch core1 table vlan-preserve
# Vlan-preserve Device name Entry key
-|-----|-----|-----|
1 0          core1      VlanVid(0x64), Port(6)
```

3. The same configuration can be verified from a switch (e.g., core1) by using the command below:

```
root@core1:~# ofad-ctl gt vlan_preserve
VLAN PRESERVE TABLE:
-----
VLAN: 100 Port: 6 PortClass: 6
```

4. Verify if a switch has any associated preserve VLAN warnings among the fabric warnings like so:

```
(config-vlan-preservation)# show fabric warnings | grep "preserve
1 delivery1 (00:00:52:54:00:85:ca:51) Switch 00:00:52:54:00:85:ca:51 cannot
preserve VLANs for some interfaces due to resource exhaustion.
```

If you find any preserve VLAN fabric warnings, please contact the [Arista Support Team](#) for assistance.

## 10.14 Reuse of Policy VLANs

From **DMF Release 8.2** policies will be allowed to reuse policy VLANs for policies in different switch islands. A switch island is an isolated fabric managed by a single pair of controllers; there is no dataplane connection between fabrics in different switch islands. For example, with a single controller pair managing six switches (**switch1**, **switch2**, **switch3**, **switch4**, **switch5**, and **switch6**), you can create two fabrics with three switches each (**switch1**, **switch2**, **switch3** in one switch island and **switch4**, **switch5**, and **switch6** in another switch island), as long as there is no dataplane connection between switches in the different switch islands.

There is no command needed to enable this feature. As long as the above condition is met, you can create policies in each switch island with the same policy VLAN tag.

In the above-mentioned condition, you can assign the same policy VLAN to two policies in different switch islands using the **push-vlan <vlan-tag>** command under policy configuration. For example, policy **P1** in **switch island 1** assigned push-vlan **10** and policy **P2** in **switch island 2** assigned the same vlan tag **10** using push-vlan **10** under policy configuration.

When two switch islands are connected by a dataplane link, they become one switch island. In that case two policies cannot use same policy vlan tag, so one of the policies (**P1** or **P2**) will become inactive.

## 10.15 Rewriting the VLAN ID for a Filter Interface

If a destination tool is shared with multiple filter interfaces, you can use the VLAN identifier assigned by the rewrite VLAN option to identify the ingress filter interface for specific packets. To use the rewrite VLAN option for this purpose, assign a unique VLAN identifier to each filter interface. This VLAN ID should be outside the auto-VLAN range.



**Note:** You can not enable the rewrite VLAN feature on filter interfaces in push-per-policy mode. If you try, a validation error is displayed. This feature is available only in the push-per-filter mode.

The following commands change the VLAN tag on packets received on the interface ethernet10 on f-switch1 to 100. The **role** command in this example also assigns the alias **TAP-PORT-1** to Ethernet interface **10**.

```
controller-1(config)# switch f-switch1
controller-1(config-switch-if)# interface ethernet10
controller-1(config-switch-if)# role filter interface-name TAP-PORT-1 rewrite
vlan 100
```

The rewrite VLAN option overwrites the original VLAN frame tag, if it was already tagged, and this changes the CRC checksum so it no longer matches the modified packet. The switch CRC option, which is enabled by default, rewrites the CRC after the frame has been modified so that a CRC error does not occur.



**Note:** Starting with **BMF Release 7.1.0**, you can simultaneously rewrite the VLAN ID and MAC address. This lets you to use VLAN rewriting to isolate traffic while using MAC rewriting to forward traffic to specific VMs.

## 10.16 Reusing Filter Interface VLAN IDs

A DMF fabric comprises groups of switches, known as islands, connected over the data plane. There are no data plane connections between switches in different islands. When Push-Per-Filter forwarding is enabled, monitored traffic is forwarded within an island using the VLAN ID associated with a Filter Interface. These VLAN IDs are configurable. Previously, the only recommended configuration was for these VLAN IDs to be globally unique.

This feature adds official support for associating the same VLAN ID with multiple Filter Interfaces as long as they are *in different islands*. This feature provides more flexibility when duplicating Filter Interface configurations across islands and helps prevent using all available VLAN IDs.

Note that within each island, VLAN IDs must still be unique, which means that Filter Interfaces in the same group of switches cannot have the same ID. Suppose you try to reuse the same VLAN ID within an island. In that case, a fabric error is generated, and only the first Filter Interface (as sorted alphanumerically by DMF name) remains in use.

### Configuration

This feature does not require any special configuration beyond the existing Filter Interface configuration workflow.

### Troubleshooting

A fabric error occurs if the same VLAN ID is configured more than once in the same island. The error message includes the Filter Interface name, the switch name, and the VLAN ID that is not unique. When encountering this error, pick a different non-conflicting VLAN ID.

Filter Interface invalid VLAN errors can be displayed in the CLI using the following command:

```
>show fabric errors filter-interface-invalid-vlan
~~~~~ Invalid Filter Interface VLAN(s) ~~~~~
# DMF Name    IF Name    Switch                               Rewrite VLAN Details
-|-----|-----|-----|-----|-----|-----|
1 filter1-f1 ethernet2 filter1 (00:00:52:54:00:4b:c9:bc) 1 The configured rewrite VLAN
1 for filter interface filter1-f1 is not unique within its fabric.
```

It is helpful to know all of the switches in an island. The following command lists all of the islands (referred to in this command as **switch clusters**) and their switch members:

```
>show debug switch-cluster
# Member
-|-----|
1 core1, filter1
```

It can also be helpful to know how the switches within an island are interconnected. Use the following command to display all the links between the switches:

```
>show link all
~~~~~ Links ~~~~~
# Active State Src switch Src IF Name Dst switch Dst IF Name Link Type Since
-|-----|-----|-----|-----|-----|-----|
1 active      filter1    ethernet1  core1    ethernet1  normal  2023-05-24 22:31:39 UTC
2 active      core1     ethernet1  filter1  ethernet1  normal  2023-05-24 22:31:40 UTC
```



### Considerations

- VLAN IDs must be unique within an island. Filter Interfaces in the same island with the same VLAN ID are not supported.
- This feature only applies to manually configured Filter Interface VLAN IDs. VLAN IDs that are automatically assigned are still unique across the entire fabric.

## 10.17 Using Push-per-filter Mode

The push-per-filter mode setting does not enable tag-based forwarding. Each filter interface is automatically assigned a VLAN ID; the default range is from **1** to **4094**. To change the range, use the **auto-vlan-range** command.

You can manually assign a VLAN that is not included in the currently defined range to a filter interface.

To manually assign a VLAN to a filter interface in **push-per-filter** mode, complete the following steps:

1. Change the auto-vlan-range from the default (**1-4094**) to a limited range, as in the following example:

```
controller-1(config)# auto-vlan-range vlan-min 1 vlan-max 1000
```

The example above configures the auto-VLAN feature to use VLAN IDs from **1** to **1000**.

2. Assign a VLAN ID to the filter interface that is not in the range assigned to the auto-VLAN feature.

```
controller-1(config)# role filter interface-name TAP-1 rewrite vlan 1001
```

## 10.18 Tag-based Forwarding

The DMF controller configures each switch with forwarding paths based on the most efficient links between the incoming filter interface and the delivery interface, which is connected to analysis tools. The TCAM capacity of the fabric switches may limit the number of policies that you can configure. The controller can also use VLAN tag-based forwarding, which reduces the TCAM resources required to implement a policy.

Tag-based forwarding is automatically enabled when the auto-VLAN Mode is push-per-policy, which is the default. This configuration improves traffic forwarding within the monitoring fabric. The assigned VLAN tags are used for forwarding traffic to the correct delivery interface, which saves TCAM space. This feature is especially useful when using switches based on the Tomahawk chipset because these switches have higher throughput but reduced TCAM space.

## 10.19 Prefix Optimizations

A policy can match with a large number of IPv4 or IPv6 addresses. These matches can be configured explicitly on each of the match rules or the match rules can use an address group. With prefix optimization, which is based on IPv4, IPv6, and TCP ports, DMF uses efficient masking algorithms to minimize the number of flow entries in hardware.

**Example 1:** Optimize same mask addresses.

```
controller-1(config)# policy ip-addr-optimization
controller-1(config-policy)# action forward
controller-1(config-policy)# delivery-interface TOOL-PORT-1
controller-1(config-policy)# filter-interface TAP-PORT-1
controller-1(config-policy)# 10 match ip dst-ip 1.1.1.0 255.255.255.255
controller-1(config-policy)# 11 match ip dst-ip 1.1.1.1 255.255.255.255
controller-1(config-policy)# 12 match ip dst-ip 1.1.1.2 255.255.255.255
```

```
controller-1(config-policy)# 13 match ip dst-ip 1.1.1.3 255.255.255.255
controller-1(config-policy)# show policy ip-addr-optimization optimized-match
Optimized Matches :
10 ether-type 2048 dst-ip 1.1.1.0 255.255.255.252
```

**Example 2:** Generic prefix optimization. In this case, if there is a generic prefix, all the specific addresses are not programmed in TCAM.

```
controller-1(config)# policy ip-addr-optimization
controller-1(config-policy)# action forward
controller-1(config-policy)# delivery-interface TOOL-PORT-1
controller-1(config-policy)# filter-interface TAP-PORT-1
controller-1(config-policy)# 10 match ip dst-ip 1.1.1.0 255.255.255.255
controller-1(config-policy)# 11 match ip dst-ip 1.1.1.1 255.255.255.255
controller-1(config-policy)# 12 match ip dst-ip 1.1.1.2 255.255.255.255
controller-1(config-policy)# 13 match ip dst-ip 1.1.1.3 255.255.255.255
controller-1(config-policy)# 100 match ip dst-ip 1.1.0.0 255.255.0.0
controller-1(config-policy)# show policy ip-addr-optimization optimized-match
Optimized Matches :
100 ether-type 2048 dst-ip 1.1.0.0 255.255.0.0
```

**Example 3:** IPv6 prefix optimization. In this case, if there is a generic prefix, the specific addresses are not programmed in the TCAM.

```
controller-1(config)# policy ip-addr-optimization
controller-1(config-policy)# 25 match ip6 src-ip 2001::100:100:100:0
      FFFF:FFFF:FFFF::0:0
controller-1(config-policy)# 30 match ip6 src-ip 2001::100:100:100:0
      FFFF:FFFF::0
controller-1(config-policy)# show policy ip-addr-optimization optimized-match
Optimized Matches :
30 ether-type 34525 src-ip 2001::100:100:100:0 FFFF:FFFF::0
```

**Example 4:** Different subnet prefix optimization. In this case, addresses belonging to different subnets are optimized.

```
controller-1(config)# policy ip-addr-optimization
controller-1(config-policy)# 10 match ip dst-ip 2.1.0.0 255.255.0.0
controller-1(config-policy)# 11 match ip dst-ip 3.1.0.0 255.255.0.0
controller-1(config-policy)# show policy ip-addr-optimization optimized-match
Optimized Matches : 10 ether-type 2048 dst-ip 2.1.0.0 254.255.0.0
```

## 10.20 Switch Dual Management Port

### 10.20.1 Overview

When a DMF switch disconnects from controller, the switch is taken out of the fabric causing service interruptions. The dual management feature provides a solution to the problem by providing physical redundancy of the switch-to-controller management connection. This is achieved by allocating a switch data path port to be bonded with its existing management interface and thereby acting as a standby management interface. Hence, it eliminates a single-point failure in the management connectivity between the switch and the controller.

Once an interface on a switch is configured for management, this configuration persists across reboots and upgrades until the management configuration is explicitly disabled on the controller.

You can configure an interface for dual management using either the CLI or the GUI.



**Note:** Along with the configuration on the controller detailed below, dual management requires a physical connection in the same subnet as the primary management link from the data port to a management switch.

## 10.20.2 Configuring Dual Management Using the CLI

1. From config mode, specify the switch which needs to be configured with dual management as in the following example:

```
Controller-1(config)# switch DMF-SWITCH-1
Controller-1(config-switch)#
```

The CLI changes to the config-switch submenu, which lets you configure the specified switch.

2. From **config-switch** mode, enter the **interface** command to specify the interface which needs to be configured as the standby management interface:

```
Controller-1(config-switch)# interface ethernet40
Controller-1(config-switch-if)#
```

The CLI changes to the config-switch-if submenu, which lets you configure the specified interface.

3. From **config-switch-if** mode, enter the **management** command to specify the role for the interface:

```
Controller-1(config-switch-if)# management
Controller-1(config-switch-if)#
```

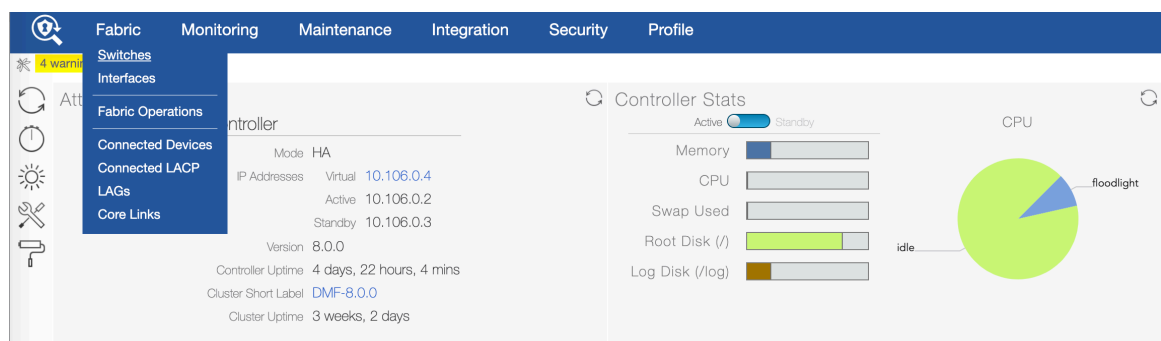


**Note:** When an interface is assigned to management, no other interface-specific commands are honored for that interface (e.g., **shut- down**, **role**, **speed**, etc.).

## 10.20.3 Configuring Dual Management Using the GUI

1. Select **Fabric > Switches** from the main menu.

**Figure 10-62: Controller GUI Showing Fabric Menu List**



2. Click on the name of the switch that needs to be configured with dual management.

**Figure 10-63: Controller GUI Showing Inventory of Switches**

**Switches**

Name	Connected	Admin Status	Management Interface
leaf-7050sx3-1	✓	✓ Up	—
leaf-7050sx3-2	✓	✓ Up	—
leaf-7050sx3-3	✓	✓ Up	—
leaf-as5812T-1	✓	✓ Up	—
leaf-as5812X-1	✓	✓ Up	—
leaf-as7312XS-1	✓	✓ Up	—
leaf-s4048-1	✓	✓ Up	—
leaf-s4048T-1	✓	✓ Up	—
leaf-s4112F-1	✓	✓ Up	—
leaf-s4112T-1	✓	✓ Up	—

3. Click on the **Interfaces** tab.

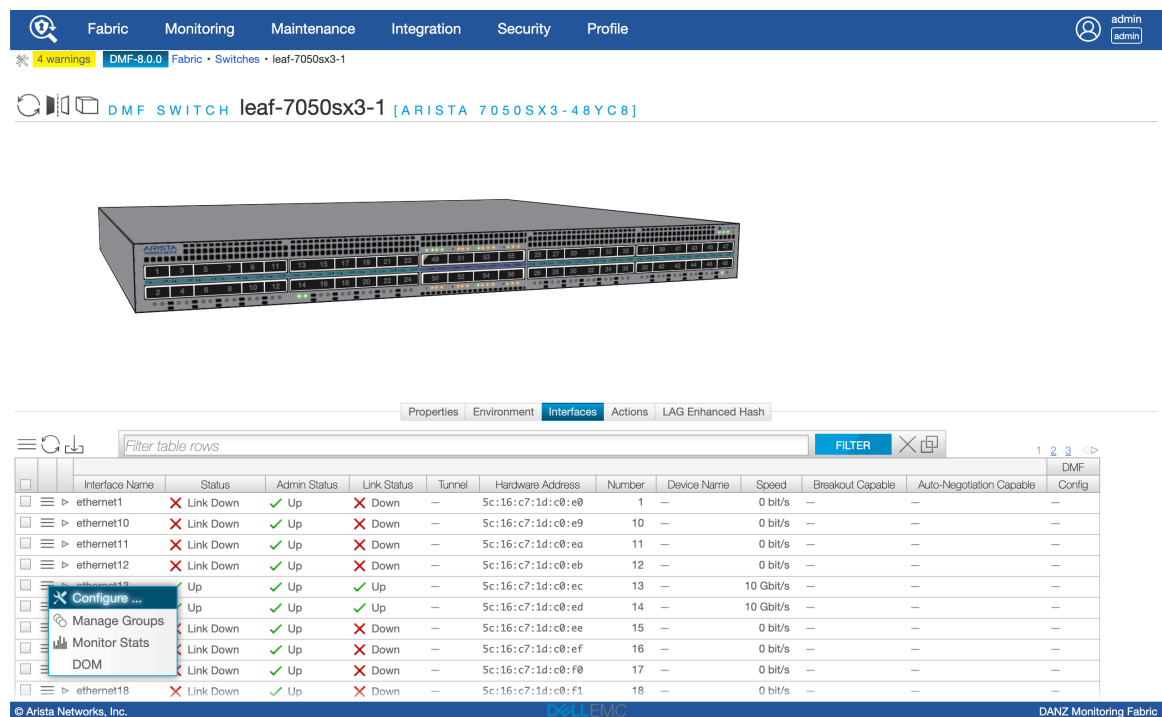
**Figure 10-64: Controller GUI Showing Switch Interfaces**

**DMF SWITCH leaf-7050sx3-1 [ARISTA 7050SX3-48YC8]**

Interface Name	Status	Admin Status	Link Status	Tunnel	Hardware Address	Number	Device Name	Speed	Breakout Capable	Auto-Negotiation Capable	DMF Config
ethernet1	✗ Link Down	✓ Up	✗ Down	—	5c:16:c7:1d:c0:e0	1	—	0 bit/s	—	—	—
ethernet10	✗ Link Down	✓ Up	✗ Down	—	5c:16:c7:1d:c0:e9	10	—	0 bit/s	—	—	—
ethernet11	✗ Link Down	✓ Up	✗ Down	—	5c:16:c7:1d:c0:ea	11	—	0 bit/s	—	—	—
ethernet12	✗ Link Down	✓ Up	✗ Down	—	5c:16:c7:1d:c0:eb	12	—	0 bit/s	—	—	—
ethernet13	✓ Up	✓ Up	✓ Up	—	5c:16:c7:1d:c0:ec	13	—	10 Gbit/s	—	—	—

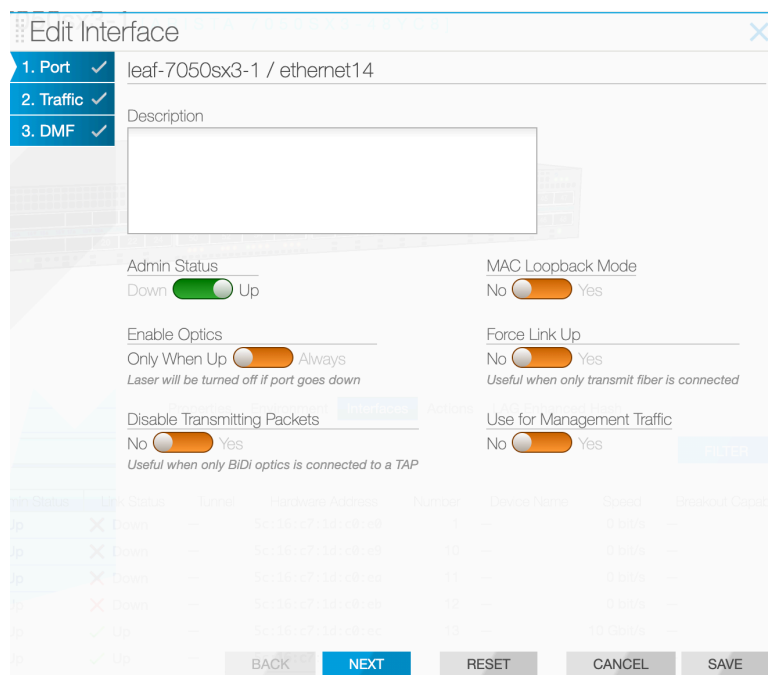
- Identify the interface to be configured as the standby management interface.

**Figure 10-65: Controller GUI Showing Configure Knob**



- Click on the button to the left of the identified interface, then click on **Configure**.
- Set **Use for Management Traffic** to **Yes**. This configures the interface to standby management role.

**Figure 10-66: Controller GUI Showing Interface Settings**



## 10.20.4 Management Interface Selection

By default, the dedicated management interface serves as the management port, with the front panel data port acting as backup only when the management interface is unavailable:

- When the dedicated management interface fails, the front panel data port becomes active as the management port.
- When the dedicated management interface comes back, it becomes the active management port again.

When the management network is undependable, this can lead to switch disconnects. Management interface selection allows you to choose what happens when the management interface comes back up after a failover. This selection can be made using the GUI or the CLI.

### Using the GUI

**Figure 10-67: Controller GUI Showing Dual Management Settings**

The screenshot shows the 'Configure Switch' GUI. On the left is a sidebar with a list of configuration sections: 1. Info, 2. Clock, 3. SNMP, 4. SNMP Traps, 5. Logging, 6. TACACS, 7. sFlow, and 8. LAG Enhanced Hash. The 'Info' section is selected. The main area shows the switch configuration for 'leaf-7050sx3-2'. The MAC Address is 'd4:af:f7:43:24:5c'. The 'Management Interface' dropdown menu is open, showing three options: 'Prefer Dedicated Management Interface' (selected), 'Prefer Dedicated Management Interface', and 'Prefer Current Interface'. Below the dropdown, a message states: 'management traffic will be changed to the dedicated management interface even if the front panel interface is still up.' The 'Admin Status' is 'Up' and the 'Change Password' toggle is 'No'.

### Using the CLI

```
Controller-1(config)# switch DMF-SWITCH-1
Controller-1(config-switch)#management-interface-selection ?
prefer-current-interface Set management interface selection algorithm
prefer-dedicated-management-interface Set management interface selection
algorithm (default selection)
Controller-1(config-switch)#
```

If you select Prefer Dedicated Management Interface (the default), when the dedicated management interface goes down, the front panel data port becomes the active management port for the switch. When

the dedicated management port comes back up, the dedicated management port becomes the active management port again and the front panel data port is put in admin down state.

If you select Prefer Current Interface, when the dedicated management interface goes down, the front panel data port still becomes the active management port for the switch. However, when the dedicated management port comes back up, the front panel data port continues to be the active management port.

### 10.20.5 Switch Fabric Management Redundancy Status

To check the status of all switches configured with dual management as well as the interface which is being actively used for management, enter the following command in the CLI:

```
Controller-1# show switch all mgmt-stats
```

### 10.20.6 Additional Notes

- A maximum of one data-plane interface on a switch can be configured as a standby management interface.
- The switch management interface **ma1** is a bond interface, having **oma1** as the primary slave and the data plane interface as the secondary slave.
- The bandwidth of the data-plane interface is limited regardless of the physical speed of the interface. Arista recommends immediate remediation when the **oma1** link fails.

## 10.21 Controller Lockdown

Controller lockdown mode, when enabled, disallows user configuration such as policy configuration, inline configuration, and rebooting of fabric components, and disables data path event processing. If there is any change in the data path, it will not be processed.

The primary use case for this feature is planned management switch upgrade. During a planned management switch upgrade, DMF switches are disconnected from the controller, and DMF policies are reprogrammed and disrupt traffic forwarding to tools. Enabling this feature before starting a management switch upgrade will not disrupt the existing DMF policies when DMF switches disconnects from controller, thereby forwarding traffic to the tools.



**Note:** DMF policies are reprogrammed when the switches reconnect to the DMF fabric when controller lockdown mode is disabled after management switch upgrade is completed. Controller lockdown mode is a special operation and should not be enabled for a prolonged period.

- Operations such as switch reboot, controller reboot, controller failover, controller upgrade, policy configuration, etc are disabled when controller lockdown mode is enabled.
- The command to enable controller lockdown mode, **system control-plane-lockdown enable**, is not saved to running config. Hence, controller lockdown mode is disabled after controller power down/up. When failover happens with a redundant controller configured, the new active controller will be in controller lockdown mode but may not have all policy information.
- In controller lockdown mode, copying running config to a snapshot will not include the **system control-plane-lockdown enable** command.
- When this feature is enabled, the CLI prompt will start with the prefix **LOCKDOWN**.
- Link up/down and other events that happen during controller lockdown mode are processed after controller lockdown mode is disabled.
- All the events handled by the switch are processed in controller lockdown mode. For example, if one member of a LAG fails, traffic is hashed to other members automatically in controller lockdown mode. Likewise, all switch-handled events related to inline are processed in controller lockdown mode.

Use the below commands to enable controller lockdown mode. Only an admin user can enable or disable this feature.

```
Controller# configure
Controller(config)# system control-plane-lockdown enable
Enabling control-plane-lockdown may cause service interruption. Do you want to
continue ("y" or "yes
" to continue):yes
LOCKDOWN Controller(config)#
```

To disable controller lockdown mode, use the command below:

```
LOCKDOWN Controller(config)# system control-plane-lockdown disable
Disabling control-plane-lockdown will bring the fabric to normal operation.
This may cause some
service interruption during the transition. Do you want to continue ("y" or
"yes" to continue):
yes
Controller(config)#
```



## Advanced Policy Configuration

This chapter describes advanced features and use cases for DMF policies.

### 11.1 Advanced Match Rules

Optional parameters of a match rule (such as **src-ip**, **dst-ip**, **src-port**, **dst-port**) must be listed in a specific order. To determine the permitted order for optional keywords, use the tab key to display completion options. If the keywords in a match rule are not entered in the correct order, the following message is displayed:

```
Error: Unexpected additional arguments ...
```

#### 11.1.1 Match Fields and Criteria

The following summarizes the different match criteria available:

- **src-ip**, **dst-ip**, **src-mac**, and **dst-mac** are maskable. If the mask for **src-ip**, **dst-ip**, **src-mac**, or **dst-mac** is not specified, it is assumed to be an exact match.
- For **src-ip** and **dst-ip**, the mask can be specified in either CIDR notation (for example /24) or dotted-decimal notation (for example, **255.255.255.0**).
- For **src-ip** and **dst-ip**, the mask must be contiguous. For example, a mask of **255.0.0.255** or **0.0.255.255** is not supported.
- For tcp, the **tcp-flags** option lets you match on the following TCP flags: **URG**, **ACK**, **PSH**, **RST**, **SYN**, and **FIN**.

The following match combinations are not allowed in the same match rule in the same DMF policy.

- **src-ip-range** and **dst-ip-range**
- **src-ip address group** and **dst-ip address group**
- **ip-range** and **ip address group**

You can match on user-defined L3/L4 offsets instead of matching on these criteria. However, you cannot use both methods of matching packets in the same DANZ Monitoring Fabric. Switching between these match modes may cause policies defined under the previous mode to fail.

Match rules can be applied to the following fields in the packet header:

```
dscp-value Match on DSCP value. Value range is 0..63
dst-ip Match dst ip
dst-port Match dst port
is-fragment Match if the packet is IP fragmented
is-not-fragment Match if the packet is not IP fragmented
l3-offset Match on l3 offset
l4-offset Match on l4 offset
range-dst-ip Match dst-ip range
range-dst-port Match dst port range
range-src-ip Match src-ip range
range-src-port Match src port range
src-ip Match src ip
src-port Match src port
untagged Untagged (no vlan tag)
vlan-id Match vlan-id
vlan-id-range Match vlan-id range
```

```
<ip-proto> IP Protocol
```



**Warning:** Matching on untagged packets cannot be applied to DMF policies when in *push-per-policy* mode.

If a policy match rule has more than one field, a logical AND is used. For example, the following rule matches if the packet has **src-ip 1.1.1.1 AND dst-ip 2.2.2.2**:

```
1 match ip src-ip 1.1.1.1 255.255.255.255 dst-ip 2.2.2.2 255.255.255.255
```

If these fields are configured in two different match rules, a logical OR is used. For example, the following matches if packet has **src-ip 1.1.1.1 OR dst-ip 2.2.2.2**:

```
1 match ip src-ip 1.1.1.1 255.255.255.255
2 match ip dst-ip 2.2.2.2 255.255.255.255
```

A match rule with the any keyword matches all traffic entering the filter interfaces in a policy:

```
controller-1(config)# policy dmf-policy-1
controller-1(config-policy)# 10 match any
```

The following commands match on the **TCP SYN** and **SYN ACK** flags:

```
1 match tcp tcp-flags 2 2
2 match tcp tcp-flags 18 18
```



**Note:** In the DMF GUI, when you configure a match on TCP flags, the current GUI workflow sets the hex value of the TCP flags for the mask attribute as well. If you have configured a different value for the tcp-flags and tcp-flags-mask attributes in a rule via the DMF CLI, editing the rule in the GUI will override the tcp-flags-mask.

## 11.1.2 Match-except Rules

The following summarizes match-except rules with examples:

- They allow a policy to permit packets that meet the match criteria, except packets that match the value specified using the **except** command.
- Match-except only supports IPv4 source-IP and IPv4 destination-IP match fields.

```
Permit src-ip network, except ip-address
1 match ip src-ip 172.16.0.0/16 except-src-ip 172.16.0.1
Permit dst-ip network, except subnet
1 match ip dst-ip 172.16.0.0/16 except-dst-ip 172.16.128.0/17
```

- In a rule, the **except** condition can only be used with either src-ip or dst-ip, but not with src-ip and dst-ip together.

```
Except being used with src-ip
1 match icmp src-ip 172.16.0.0/16 except-src-ip 172.16.0.1 dst-ip
  172.16.0.0/16
```

```
Except being used with dst-ip
1 match icmp src-ip 224.248.0.0/24 dst-ip 172.16.0.0/16 except-dst-ip
  172.16.0.0/18
```

- Except-src-ip or except-dst-ip can only be used after a match for src-ip or dst-ip, respectively.

```
Incorrect match rule
1 match icmp except-src-ip 192.168.1.10
Correct match rule
```

```
1 match icmp src-ip 192.168.1.0/24 except-src-ip 192.168.1.10
```

- In a match rule, only one IP address, or one subnet (range of IP addresses) can be used with the **except** command.

```
Deny a subnet
1 match ip dst-ip 172.16.0.0/16 except-dst-ip 172.16.0.0/18
Deny an IP address
1 match ip dst-ip 172.16.0.0/16 except-dst-ip 172.16.0.1
```

### 11.1.3 Matching with IPv6 Addresses

The value of the EtherType field determines whether the src-ip field to match is IPv4 or IPv6. The DMF controller displays an error if there is a mismatch between the EtherType and the format of the IP address.

DMF supports IPv6 address/mask matching, either on src-IP or dst-IP. Optionally, UDP/TCP ports can be used with the IPv6 address/mask match. You must specify an address/mask or a group; ranges are not supported for IPv6 addresses.



**Note:** Match rules containing both MAC addresses and IPv6 addresses are not accepted and cause a validation error.

- The preferred IPv6 address representation is as follows: **xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**, where each **x** is a hexadecimal digit representing **4** bits.
- IPv6 addresses range from **0000:0000:0000:0000:0000:0000:0000:0000** to **ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff**.

In addition to this preferred format, IPv6 addresses may be specified in two other shortened formats:

- **Omit Leading Zeros:** Specify IPv6 addresses by omitting leading zeros. For example, IPv6 address **1050:0000:0000:0000:0005:0600:300c:326b** may be written as **1050:0:0:0:5:600:300c:326b**.
- **Double Colon:** Specify IPv6 addresses by using double colons (::) in place of a series of zeros. For example, IPv6 address **ff06:0:0:0:0:0:c3** may be written as **ff06::c3**. Double colons may be used only once in an IP address.

DMF does not support IPv4 address embedded IPv6 address format. For example, neither **0:0:0:0:0:101.45.75.219** nor **::101.45.75.219** can be used. Both IPv4 and IPv6 mask must be CIDR format. For example, **FFFF:FFFF:FFFF:FFFF:0:0:0:0** is valid in DMF, but **FFFF:0:0:FFFF:FFFF:0:0:0** is not a valid mask. Both the colon-separated hexadecimal representation and the CIDR-style mask format are supported. The following example shows the correct format for IPv6 addresses and subnet masks:

```
controller-1(config)# policy dmf-ipv6-policy
controller-1(config-policy)# 10 match ip6 src-ip 2001::0 ffff:ffff:fff
f:ffff:0:0:0:0
controller-1(config-policy)# 11 match ip6 dst-ip 2001:db8:122:344::/64
controller-1(config-policy)# filter-interface all
controller-1(config-policy)# action drop
```

### 11.1.4 Port and VLAN Range Matches

DMF policy supports matching on source and destination port ranges with optimized hardware resource utilization. Efficient masking algorithms are used to minimize the number of flow entries in hardware for each VLAN range. For example, a VLAN range of **12-99** uses only five flows in hardware.



**Note:** Use the untagged keyword to match traffic without a VLAN tag.

You must provide the IP protocol information when using source and destination port ranges, which are supported for IPv4 and IPv6 for TCP and UDP. These keywords have the following options:

- **range-dst-ip:** Match dst-ip range.
- **range-dst-port:** Match dst port range.

- **range-src-ip**: Match src-ip range.
- **range-src-port**: Match src port range.

Either **src-port-range** or **dst-port-range** or both can be specified in each match rule, as shown in the following example:

```
controller-1(config)# policy ip-port-range-policy
controller-1(config-policy)# 10 match tcp range-src-port 10 100
controller-1(config-policy)# 15 match udp range-dst-port 300 400
controller-1(config-policy)# 20 match tcp range-src-port 10 2000 range-dst-port
400 800
controller-1(config-policy)# 30 match tcp6 range-src-port 8 20
controller-1(config-policy)# 40 match tcp6 range-src-ip 1:2:3:4::/64 range-src-
port 10 300
controller-1(config-policy)# filter-interface all
controller-1(config-policy)# delivery-interface all
controller-1(config-policy)# action forward
```

DMF policy also supports matches for VLAN ID range with optimized hardware resource utilization. A VLAN ID range can be combined with a source port or destination port range, but you cannot use all three ranges in a single match. The following example shows a valid use of the VLAN ID range option;

```
controller-1(config)# policy vlan-range-policy
controller-1(config-policy)# 10 match mac vlan-id-range 30 400
controller-1(config-policy)# 20 match full ether-type ip protocol 6 vlan-id-
range 1000 3000 srcip 1.
1.1.1 255.255.255.255 src-port-range 100 500
```

To determine the number of flow entries required for a range, use the **optimized-match** option, as shown in the following example:

```
controller-1(config-policy)# show running-config policy
! policy
policy vlan-range-policy
action forward
delivery-interface TOOL-PORT-1
filter-interface TAP-PORT-1
10 match mac vlan-id-range 12 99
controller-1(config-policy)# show policy vlan-range-policy optimized-match
Optimized Matches :
10 vlan-min 12 vlan-max 15
10 vlan-min 16 vlan-max 31
10 vlan-min 32 vlan-max 63
10 vlan-min 64 vlan-max 95
10 vlan-min 96 vlan-max 99
```

### 11.1.5 User-defined Filters

Starting with **BMF Release 7.1.0**, up to eight two-byte user-defined offsets are allowed on each switch. To view the currently defined offsets, select **Monitoring > User Defined Offsets**.



**Note:** In order for a user-defined filter to work accurately, the DMF controller must be in push-per-policy mode.

If you select the **User Defined Offsets** option when the L3-L4 Offset Match switching mode is not enabled, the system displays a message asking you to enable the correct match mode.

After enabling the L3-L4 Offset Match mode, when you select **Monitoring > User Defined Offsets** the system displays a table that lists the offsets currently defined.



**Note:** Matching on a user-defined offset is not recommended when forwarding traffic to a tunnel, because some packets may be dropped.

Each offset match has the following four components:

- **Anchor:** Specified from where user can define the matching criteria. There are 3 options to choose from: a) L3-start: Start of layer 3 header. b) L4-start: Start of layer 4 header. c) Packet-start: Start of the packet from layer 2 header.
- **Offset:** The number of bytes from the specified anchor.
- **Length:** The number of bytes to be matched. User can match 2 or 4 bytes.
- **Value:** The value to be matched of the specified length, which can be specified in hexadecimal, decimal, or IPv4 format.
- **Mask:** The value that is ANDed with the match value.



**Note:** DMF allows users to combine up to four 4-byte user-defined offsets or up to eight 2-byte offsets to match up to sixteen bytes in the same match condition. In this case, the multiple offset matching conditions in a single match statement will be considered ANDed. For example, to match on eight bytes, in a single match condition, define two user-defined offsets and configure two rules in an AND fashion so that the first rule matches on the first four bytes and the second rule matches on the remaining four bytes.

Each switch can be configured with a maximum of eight different offsets matching two bytes each, which can be used in a single policy, or in any combination in different policies. In the example below, the policy matches on a value of **0x00001000** at offset **40** from the start of the L3-header and a value of **0x00002000** at offset **64** from the start of the L4-header.

```
controller-1(config-policy)# 1 match udp dst-port 2152 l3-offset 40 length 4
value 0x00001000 mask
0xffffffff l4-offset 64 length 4 value 0x00002000 mask 0xffffffff
```

To display the values configured in the user-defined-offset table, enter the **show user-defined-offset** command.

```
controller-1# show user-defined-offset
# Switch      Slot Anchor  Offset Length Policy
-----
1 DMF-FILTER-SW1 0      14-start 64      2      DMF-UDF-TEST-1, _DMF-UDF-TEST-1_o_SAVE-TO-RECORDER-NODE
2 DMF-FILTER-SW1 1      14-start 66      2      DMF-UDF-TEST-1, _DMF-UDF-TEST-1_o_SAVE-TO-RECORDER-NODE
3 DMF-FILTER-SW1 2      13-start 40      2      DMF-UDF-TEST-1, _DMF-UDF-TEST-1_o_SAVE-TO-RECORDER-NODE
4 DMF-FILTER-SW1 3      13-start 42      2      DMF-UDF-TEST-1, _DMF-UDF-TEST-1_o_SAVE-TO-RECORDER-NODE
controller-1#
```

User-defined filtering is supported on Trident 3 switches. The following are the UDF limitations on a Trident 3 switch in comparison to a non-Trident 3 switch:

UDF Features	Non-Trident 3	Trident 3
Total UDF Length	16 bytes	12 bytes
Minimum Chunk Size	2 bytes	2 bytes
Packet Start (Layer 2 Anchor)	8 offsets	2 offsets
Layer 3 Anchor	8 offsets	6 offsets
Layer 4 Anchor	8 offsets	6 offsets
Layer 2 Offset Range	0 - 126 bytes	0 - 62 bytes
Layer 3 Offset Range	0 - 114 bytes	0 - 112 bytes
Layer 4 Offset Range	0 - 96 bytes	0 - 112 bytes



**Note:** For a complete list of supported switches and their corresponding Network ASIC types (Trident 3, Trident 2 etc), please refer to the ***DMF Hardware Compatibility List***.

## 11.2 Using the Filter and Delivery Role with MAC Loopback for a Two-stage Policy

The Filter and Delivery role can be used with a MAC (software) loopback to support monitoring as a service. This option uses a two-stage policy to replicate the incoming feed from one or more filter interfaces and send it to multiple intermediate interfaces (one per end customer or organization).

The end customer can define policies on the intermediate interface for forwarding to customer-specific tools. These intermediate interfaces must also be assigned the Filter and Delivery role enabled with the MAC loopback option. This method eliminates the need for a physical loopback cable and a second interface, which simplifies the deployment of monitoring as a service.

When multiple user-defined policies with overlapping rules select traffic from the same filter interfaces for forwarding to different delivery interfaces, overlapping policies are automatically generated to replicate the requisite traffic to the delivery interfaces. The number of overlapping policies increases exponentially with the number of user-defined policies.


The total number of policies in the fabric is limited by switch hardware limits. The mechanism of using a Filter and Delivery role with a MAC loopback can also be useful for eliminating scale and operational issues seen with overlapping policies.

To configure an interface with the Filter and Delivery role and enable the MAC (software) loopback option, use the `loopback-mode mac` command to assign an unused interface as a loopback. The physical interface is enabled by this command without requiring a physical connection to the interface. You can use a software loopback interface for copying traffic in any scenario where you might use a physical loopback.

You can also assign the Filter and Delivery role to a software loopback interface, which lets you use a single interface for copying traffic to multiple destination interfaces. When you assign this role to an interface in loopback mode, you can then use the interface as a delivery interface in relation to the original filter interface, and as a filter interface in relation to the final destination interface.

The following figure illustrates the physical configuration for a switch that uses four software loopback interfaces to copy traffic from a single filter interface to four different tools:

**Figure 11-1: Using Software Loopback Interfaces to Avoid Overlapping Policies**



This configuration is used to copy different types of traffic from a single filter interface (**F1**) to four different delivery interfaces (**D1** to **D4**). Assigning the Filter and Delivery role to the software loopback interfaces (**LFD1** through **LFD4**) lets you do this with four physical interfaces. Physical loopbacks would require twice as many interfaces.

## 11.2.1 Using the GUI To Configure a Filter and Delivery Interface with MAC Loopback

To configure an interface with the Filter and Delivery role and enable the MAC (software) loopback option in the GUI, complete the following steps:

1. Display the available interfaces by selecting **Fabric > Interfaces**.  
The system displays the **Interfaces** page, which lists the interfaces connected to the DMF fabric.
2. Click the **Menu** control for the interface you want to use and select **Configure** from the pull-down menu.

The system displays the following dialog:

**Figure 11-2: Fabric > Interfaces > Edit Interface > Port**

**Edit Interface**

1. Port ☒ leaf-7050sx3-1 / ethernet13

2. Traffic ☒ Description

3. DMF ☒

Admin Status: ☒ Up ☐ Down

Enable Optics: ☒ Only When Up ☐ Always ☐ Down  
*Laser will be turned off if port goes down*

Disable Transmitting Packets: ☒ No ☐ Yes  
*Useful when only BiDi optics is connected to a TAP*

MAC Loopback Mode: ☒ No ☐ Yes

Force Link Up: ☒ No ☐ Yes  
*Useful when only transmit fiber is connected*

Use for Management Traffic: ☒ No ☐ Yes

BACK NEXT RESET CANCEL SAVE

3. (Optional) Type a description for the interface.
4. Enable the **MAC Loopback Mode** slider.

**5. Click Next.**

**Figure 11-3: Fabric > Interfaces > Edit Interface > Traffic**

## Edit Interface

1. Port ✓

2. Traffic ✓

3. DMF ✓

Forward Error Correction

☒ Use Switch Default
 ☐ Enable
 ☐ Disable

Auto-Negotiation

☒ Use Switch Default
 ☐ Enable
 ☐ Disable

Breakout

Default

Breakout

4x1G

4x10G

4x25G

Speed

Default

1G

10G

25G

40G

100G

Rate Limit

☐ 0 kbps

BACK

NEXT

RESET

CANCEL

SAVE



6. (Optional) Configure **Rate Limiting**, if required, and click **Next**.

**Figure 11-4: Fabric > Interfaces > Edit Interface > DMF**

7. Enable the **Filter and Delivery** radio button.  
You can optionally enable the **Rewrite VLAN** option.



**Note:** The rewrite VLAN ID feature cannot be used with tunneling.

8. Click **Save** to complete and save the configuration.

## 11.2.2 Using the CLI To Configure a Filter and Delivery Interface with MAC Loopback

The CLI interface configuration for copying traffic to multiple delivery ports is shown in the following example:

```
switch DMF-FILTER-SWITCH-1
admin hashed-password
$6$5niT1gPm$Jc24qOMF.hxNPI20DvnKaFZKYD61Io59IMp3O4xIdwVTu2hx0s8Djpvz9xXAXXndiSkKe5jH.9PKoHHrWviS10
mac 70:72:cf:dc:99:5c
interface ethernet1
role filter interface-name TAP-PORT-1
interface ethernet13
role delivery interface-name TOOL-PORT-1
interface ethernet15
role delivery interface-name TOOL-PORT-1
interface ethernet17
role delivery interface-name TOOL-PORT-3
interface ethernet19
role delivery interface-name TOOL-PORT-4
interface ethernet25
loopback-mode mac
role both-filter-and-delivery interface-name LOOPBACK-PORT-1
interface ethernet26
```

```

loopback-mode mac
role both-filter-and-delivery interface-name LOOPBACK-PORT-2
interface ethernet27
loopback-mode mac
role both-filter-and-delivery interface-name LOOPBACK-PORT-3
interface ethernet28
loopback-mode mac
role both-filter-and-delivery interface-name LOOPBACK-PORT-4

```

The following example shows how five policies can be used to implement this use case without creating overlapping policies. Without the use of the loopback interfaces to copy the traffic to separate filter interfaces, sixteen overlapping policies would be created.

```

! policy
policy TAP-NETWORK-1
action forward
delivery-interface LOOPBACK-PORT-1
delivery-interface LOOPBACK-PORT-2
delivery-interface LOOPBACK-PORT-3
delivery-interface LOOPBACK-PORT-4
filter-interface TAP-PORT-1
1 match any
!
policy DUPLICATED-TRAFFIC-1
action forward
delivery-interface TOOL-PORT-1
filter-interface LOOPBACK-PORT-1
1 match ip src-ip 100.1.1.1 255.255.255.252
!
policy DUPLICATED-TRAFFIC-2
action forward
delivery-interface TOOL-PORT-2
filter-interface LOOPBACK-PORT-2
1 match ip dst-ip 100.1.1.1 255.255.255.252
!
policy DUPLICATED-TRAFFIC-3
action forward
delivery-interface TOOL-PORT-3
filter-interface LOOPBACK-PORT-3
1 match tcp src-port 1234
!
policy DUPLICATED-TRAFFIC-4
action forward
delivery-interface TOOL-PORT-4
filter-interface LOOPBACK-PORT-4
1 match tcp dst-port 80

```

You can use **show policy** command to verify the policy configuration.

## 11.3 Rate Limiting Traffic to Delivery Interfaces

You can limit the traffic rate on a delivery interface, which can be a regular interface, a port channel, a tunnel interface, or a loopback interface.

For information about using rate limiting on tunnels, refer to the [Using the CLI to Rate Limit the Packets on a VxLAN Tunnel](#) section.

Use kbps to configure the rate-limit for regular delivery interface. Arista recommends configuring the rate limit in multiples of 64 kbps.

### 11.3.1 Rate Limiting Using the GUI

To use the GUI to set the rate limit for an interface, complete the following steps:

1. Select **Fabric > Interfaces**.
2. Click the **Menu** control for a specific interface and select **Configure**.
3. Click **Next** or select **Traffic** to display the **Traffic** page on the **Edit Interface** dialog.

**Figure 11-5: Setting the Rate Limit for an Interface**

4. Enable the **Rate Limit** checkbox.
5. Use the number spinner to set the number of Kbps to which the traffic should be limited.
6. Click **Save**.

### 11.3.2 Rate Limiting Using the CLI

#### CLI Procedure

The following example applies a rate limit of **10 Mb/s** to the delivery interface **tobcotDelivery**:

```
CONTROLLER-1(config)#switch DMF-DELIVERY-SWITCH-1
CONTROLLER-1(config-switch)# interface ethernet1
CONTROLLER-1(config-switch-if)# role delivery interface-name TOOL-PORT-1
CONTROLLER-1(config-switch-if)# rate-limit 10240
```

To view the configuration, enter the **show this** command, as in the following example:

```
CONTROLLER-1(config-switch-if)# show this
! switch
switch DMF-DELIVERY-SWITCH-1
!
interface ethernet1
rate-limit 10000
role delivery interface-name TOOL-PORT-1
CONTROLLER-1 (config-switch-if)#
```

To rate limit a port channel, configure the rate limit for each of the member interfaces. If the port channel has two member interfaces, configure individual rate limits for each member interface.

```
lag-interface lag1
hash-type 13
```

```

member ethernet43
member ethernet45
interface ethernet43
rate-limit 10000 <----- set the rate-limit to 10 Mbps
interface ethernet45
rate-limit 128000 <----- set the rate-limit to 128 Mbps

```

To display the configured rate limit, use the **show topology** and **show interface-names** commands, as in the following examples:



**Note:** In the current release, for LAG and tunnel interfaces, the Rate Limit column does not show the configured value.

```

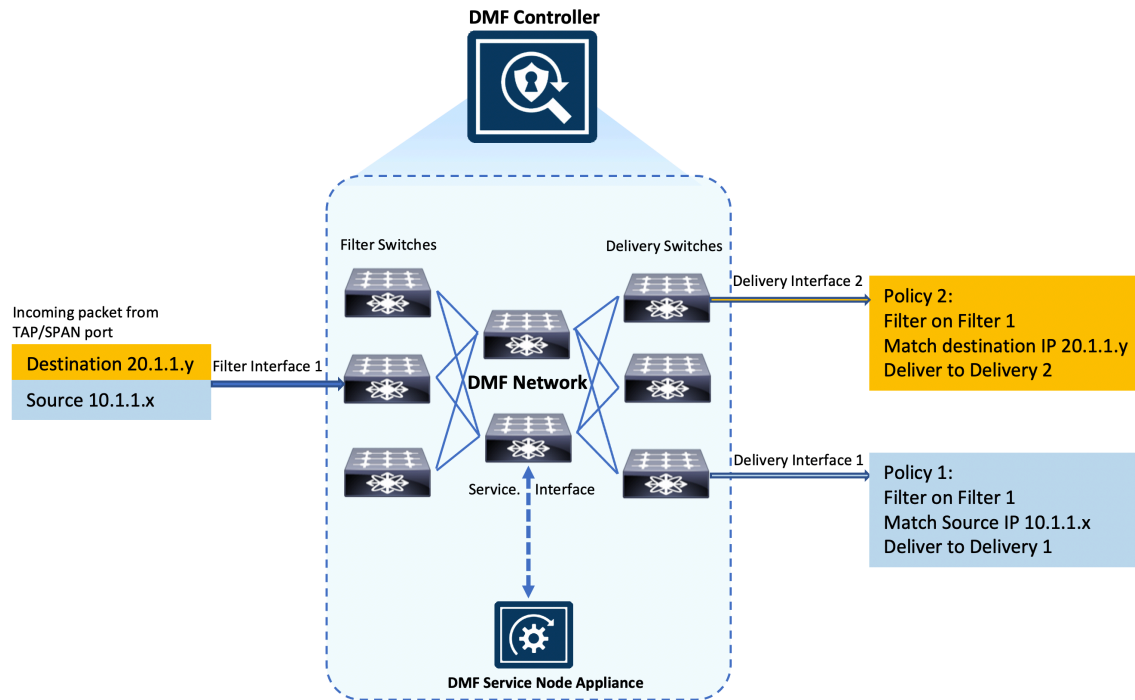
controller-1# show topology
----- Delivery Interface(s) -----
# IP Switch      IP Name      State      Speed Rate  Limit
-----|-----|-----|-----|-----|
1 to-switch-scale-2 dmf-switch-2 ethernet18 up 10Gbps -
2 tunnelSwitch dmf-switch-2 tunnelSwitch up - -
3 tohostDelivery dmf-switch-1 ethernet1 up 1Gbps 250Mbps
4 delivery-loop1 dmf-switch-1 ethernet15 up 10Gbps -
5 test-filter-delivery dmf-switch-1 ethernet18 up 10Gbps -
controller-1# show interface-names
----- Filter Interface(s) -----
# IP Switch      IP Name      Dir State Speed  VLAN Tag Analytics Ip address Connected Device
-----|-----|-----|-----|-----|
1 LOOPBACK-FILTER-DELIVERY-1 DMP-CORE-SWITCH-1 ethernet17 rx up 100Gbps 0 True
2 LOOPBACK-FILTER-DELIVERY-2 DMP-CORE-SWITCH-1 ethernet18 rx up 100Gbps 0 True
3 ve-filter-17 DMP-FILTER-SWITCH-1 ethernet7 rx up 10Gbps 0 True anet-vsw-host1.qa.bigswitch.com port 24 on dvSwitch WIN10-DP (etherswitch)
4 ve-filter-18 DMP-FILTER-SWITCH-1 ethernet8 rx up 10Gbps 0 True anet-vsw-host2.qa.bigswitch.com port 22 on dvSwitch WIN10-DP (etherswitch)
5 ve-filter-19 DMP-FILTER-SWITCH-1 ethernet9 rx up 1Gbps 0 True anet-vsw-host3.qa.bigswitch.com port 20 on dvSwitch WIN10-DP (etherswitch)
----- Delivery Interface(s) -----
# IP Switch      IP Name      Dir State Speed  Rate Limit Strip Pwly Vlan Strip Customer Vlan IP Address ARP Resolved Connected Device
-----|-----|-----|-----|-----|
1 LOOPBACK-FILTER-DELIVERY-1 DMP-CORE-SWITCH-1 ethernet17 tx up 100Gbps - True False
2 LOOPBACK-FILTER-DELIVERY-2 DMP-CORE-SWITCH-1 ethernet18 tx up 100Gbps - True False
3 NETFLIX-DELIVERY-PORT DMP-DELIVERY-SWITCH-1 ethernet4 tx up 10Gbps 10Mbps True False 172.43.75.1 True app-arista-5.qa.bigswitch.com Connected to Eth4 on DMP-DELIVERY-SWITCH /DMP DogFood
4 TOOL-PORT-01 DMP-DELIVERY-SWITCH-1 ethernet1 tx up 1Gbps - True False
5 TOOL-PORT-02 DMP-DELIVERY-SWITCH-1 ethernet15 tx up 10Gbps - True False
6 TOOL-PORT-03 DMP-DELIVERY-SWITCH-1 ethernet16 tx up 10Gbps - True False
----- Service Interface(s) -----
# IP Switch      IP Name      Dir State Speed  Connected Device
-----|-----|-----|-----|-----|
1 DMP-CORE-SWITCH-1-ethernet1311-to-managed-service DMP-CORE-SWITCH-1 ethernet1311 tx up 10Gbps DMP-DOGFOD-HP-IN-1 1
2 DMP-CORE-SWITCH-1-ethernet1311-to-managed-service DMP-CORE-SWITCH-1 ethernet1311 rx up 10Gbps DMP-DOGFOD-HP-IN-1 1
3 DMP-CORE-SWITCH-1-ethernet1312-to-managed-service DMP-CORE-SWITCH-1 ethernet1312 tx up 10Gbps DMP-DOGFOD-HP-IN-1 2
4 DMP-CORE-SWITCH-1-ethernet1312-to-managed-service DMP-CORE-SWITCH-1 ethernet1312 rx up 10Gbps DMP-DOGFOD-HP-IN-1 2
5 post-serv-intf-1 DMP-CORE-SWITCH-1 ethernet22 rx up 100Gbps
6 pre-serv-intf-1 DMP-CORE-SWITCH-1 ethernet21 tx up 100Gbps
controller-1#

```

## 11.4 Configuring Overlapping Policies

When two or more policies have one or more filter ports in common, the match rules in these policies may intersect. If the priorities are different, the policy with the higher priority takes effect. However, if the policies have the same priority, the policies overlap as illustrated in the figure below:

**Figure 11-6: Overlapping Policies**



In the policy illustrated, packets received on interface **Filter 1** with the source-IP address **10.1.1.x/24** are delivered to **D1**. In a separate policy, with the same priority, packets received at **Filter 1** with the destination IP address **20.1.1.y/24** are delivered to **D2**. With both these policies applied, when a packet arrives at **F1** with a source IP address **10.1.1.5/24** and a destination IP address of **20.1.1.5/24**, the packets are copied and forwarded to both **D1** and **D2**. This behavior is caused by the DMF policy overlap feature, which is enabled by default.

DMF manages overlapping policies automatically by copying packets received on the same filter interface that match multiple rules, but which the policy forwards to different delivery interfaces.

Two policies are said to be overlapping when all of the following conditions are met:

- At least one delivery interface is different.
- At least one filter interface is shared.
- Match rules across policies intersect, which occurs under these conditions:
  - The match rules match on the same field but a different value OR both policies have the same configured priority (or same default priority).
  - The match rules match on completely different fields.



**Note:** Automatically created dynamic policies will be visible in the **show policy** command. However, they will not be visible in the **running config**, nor can they be deleted manually.

When overlapping policies are detected, by default DMF performs the following operations:

- Creates a new dynamic policy that aggregates the policy actions.
- Assigns policy names, using this dynamic policy naming convention: **\_<policy1>\_o\_<policy2>\_**

- Adds match combinations and configuration as appropriate.
- Assigns a slightly higher priority to the new aggregated policy so that it overrules the overlapping policies, which as a result only applies to traffic that does not match the new aggregated policy. An incremental value of .1 is added to the original policy priority. For example, if the original policies have a priority of **100**, the dynamic policy priority is **101**.



**Note:** When changing the configurable parameters in an existing DMF out-of-band policy, any counters associated with the policy, including service-node-managed services counters, are reset to zero.

Starting with **BMF Release 7.1.0**, the **overlap-limit-strict** command, which is enabled by default, strictly limits the number of overlapping policies to the maximum configured using the **overlap-policy-limit** command. For example, when the maximum number of overlapping policies is set to four (the default) and you create a fifth policy using the same filter interface, the operation fails with a validation error. To disable strict enforcement, use the **no overlap-limit-strict** command.



**Note:** After upgrading from **BMF Release 7.0.x**, the **overlap-strict-limit** command is disabled and must be manually enabled to enforce configurable policy limits.

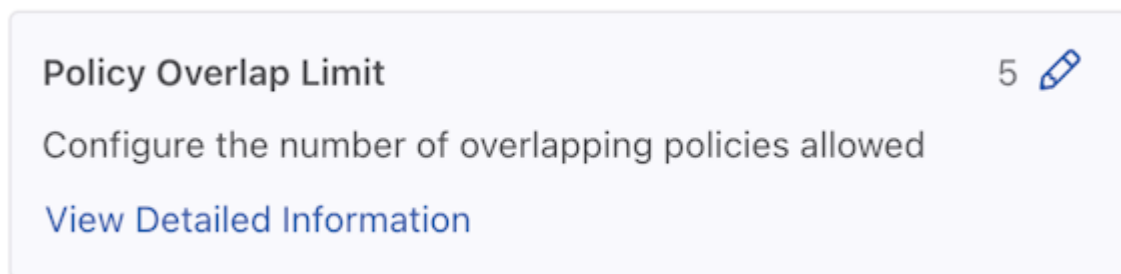
### 11.4.1 Configuring the Policy Overlap Limit Using the GUI

Policy Overlap Limit

Perform the following steps to configure the Policy Overlap Limit.

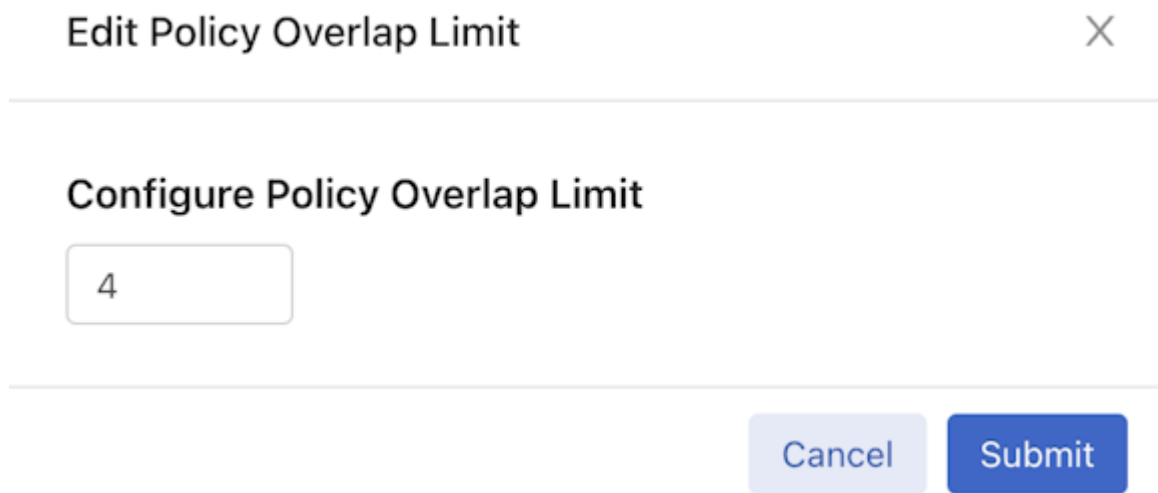
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 11-7: Policy Overlap Limit**



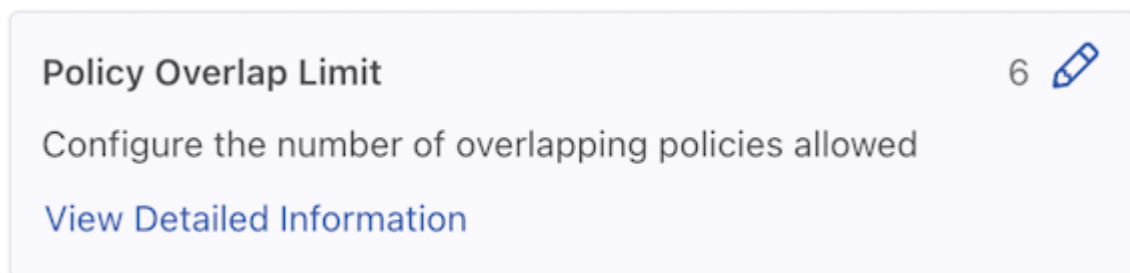
2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. By default, the Policy Overlap Limit is 4.

**Figure 11-8: Edit Policy Overlap Limit**



3. Adjust the Value (minimum value: 0, maximum value: 10). There are two ways to adjust the value:
  - Directly enter the desired value in the input area.
  - Use the up and down arrow buttons in the input area to adjust the value accordingly. Pressing the up arrow increments the value by 1, while pressing the down arrow decrements it by 1.
4. Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.
5. After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 11-9: Policy Overlap Limit Change Success**



### 11.4.2 Configuring the Overlapping Policy Limit Using the CLI

By default, the number of overlapping policies allowed is **four**. The maximum number that you can configure for overlapping policies is **ten**. To disable the overlapping policy feature, set the overlap policy limit to **zero**.

To change the default limit for overlapping policies, use the following command:

```
controller-1(config)# overlap-policy-limit <integer>
```

Replace **integer** with the maximum number of overlapping policies to support fabric-wide.

For example, the following command sets the number of overlapping policies supported to the maximum value (**10**):

```
controller-1(config)# overlap-policy-limit 10
```

The following command disables the overlapping policies feature:

```
controller-1(config)# overlap-policy-limit 0
```



**Note:** If you set the Policy Overlap Limit to **zero**, be careful that your policies do not overlap. If active policies overlap after disabling this feature, the forwarding result may be unpredictable.

### 11.4.3 Using the CLI to View Overlapping Policies

To view statistics for dynamic (overlapping) policies, enter the **show policy** command. If an overlapping policy is listed in the output, the parent policies are identified, as in the following example:

```
controller-1(config-policy)# show policy
# Policy Name Config Status Runtime Status Action Type Priority Overlap Priority Rewrite VLAN Filter BW Delivery BW Services
-----
1 p2_o_p1 active and forwarding installed forward Dynamic 100 1 0 - - -
2 p1 active and forwarding installed forward Configured 100 0 0 - - -
3 p2 active and forwarding installed forward Configured 100 0 0 - - -
```

In this example:

- **show overlap \_P1\_O\_P2**, lists component policies: source **P1**, **P2**.
- **show P1**, lists dynamic policies: **overlap \_P1\_O\_P2**.

To view the details for a specific overlapping policy, append the policy name to the **show policy** command, as in the following example:

```
controller-1(config-policy)# show policy _p1_o_p2
Policy Name : _p1_o_p2
Config Status : active and forwarding
Runtime Status : installed
Detailed Status : installed - installed to forward
Action : forward
Priority : 100
Overlap Priority : 1
Description : runtime policy
# of switches with filter interfaces : 1
# of switches with delivery interfaces : 1
# of switches with service interfaces : 0
# of filter interfaces : 1
# of delivery interfaces : 2
# of core interfaces : 4
# of services : 0
# of pre service interfaces : 0
# of post service interfaces : 0
Rewrite VLAN : 0
Total Ingress Rate : -
Total Delivery Rate : -
Total Pre Service Rate : -
Total Post Service Rate : -
Overlapping Policies : none
Component Policies : p2, p1,
Failed Overlap Policy Exceeding Max Rules :
Rewrite valid? : False
Service Names :
Overlap Matches :
1 ether-type 2048 src-ip 10.1.1.1 255.255.255.0 dst-ip 20.1.1.1 255.255.255.0
Strip VLAN : False
Delivery Bandwidth : 20 Gbps
explicitly-scheduled : False
Filter Bandwidth : 10 Gbps
Type : Dynamic
~ Match Rules ~
None.
~~~~~ Filter Interface(s) ~~~~~
# IF Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate
-----
1 f1 filter-sw-1 s11-eth1 up rx 0 0 0 -
~~~~~ Delivery Interface(s) ~~~~~
# IF Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate
-----
1 d1 filter-sw-2 s12-eth1 up tx 0 0 0 -
2 d2 filter-sw-2 s12-eth2 up tx 0 0 0 -
~ Service(s) ~
None.
~~~~~ Core Interface(s) ~~~~~
# Switch IF State Dir Packets Bytes Pkt Rate Bit Rate
-----
1 filter-sw-1 s11-eth3 up tx 0 0 0 -
2 core-sw-2 s10-eth1 up rx 0 0 0 -
3 core-sw-2 s10-eth2 up tx 0 0 0 -
~ Failed Path(s) ~
None.
~~~~~ Event History ~~~~~
# Time Event Detail
-----
1 2014-08-05 22:22:27 start forward pending installation - installed to forward
2 2014-08-05 22:22:27 installation complete installed - installed to forward
```

### 11.4.4 Configuring the Policy Overlap Limit Strict using the GUI

The Policy Overlap Limit Strict option, **enabled** by default, strictly limits the number of overlapping policies to the maximum configured. For example, when the maximum number of overlapping policies is set to 4 (the default) and users create a fifth policy using the same filter interface, the operation fails with a validation error.



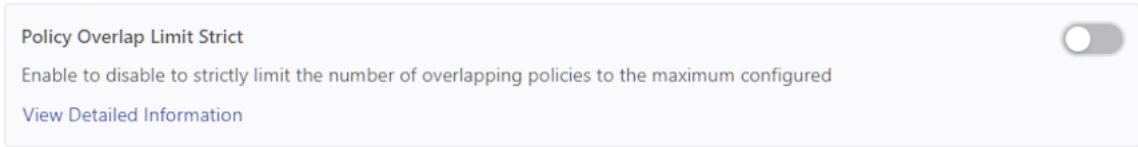
From the DMF Features page, proceed to the Configuring the Policy Overlap Limit Strict feature card.

1. Select the **Configuring the Policy Overlap Limit Strict** card.




**Note:** The Policy Overlap Limit Strict option is **enabled** by default. The following steps guide if the Policy Overlap Limit Strict option is disabled.

**Figure 11-10: Policy Overlap Limit Strict Disabled**



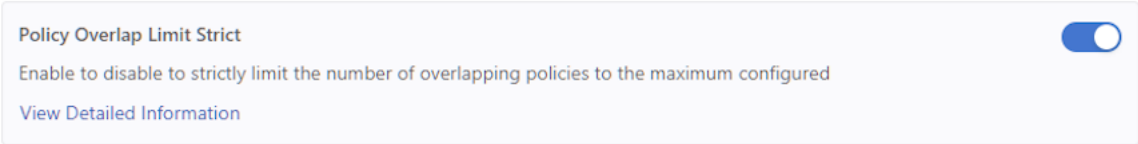
2. Toggle the Configuring the Policy Overlap Limit Strict switch to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 11-11: Enable Policy Overlap Limit Strict**



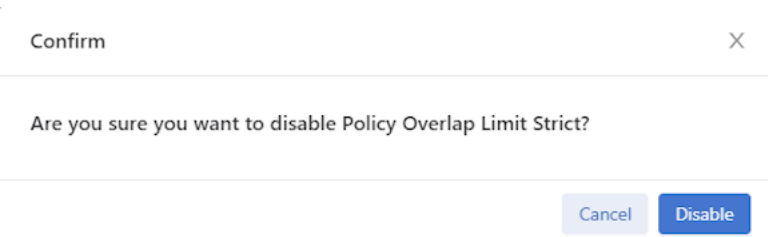
4. Retain Configuring the Policy Overlap Limit Strict is running.

**Figure 11-12: Policy Overlap Limit Strict Enabled**



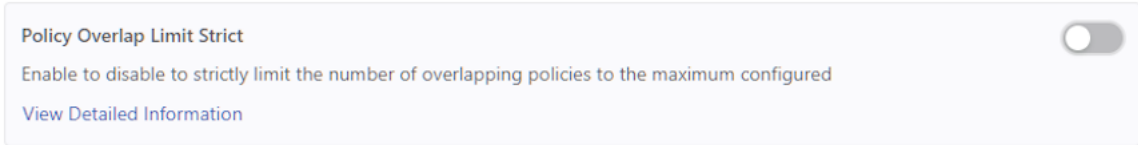
5. To disable the feature, toggle the Configuring the Policy Overlap Limit Strict switch to **Off**. Click **Disable** and confirm.

**Figure 11-13: Disable Policy Overlap Limit Strict**



The feature card updates with the status.

**Figure 11-14: Policy Overlap Limit Strict Disabled**



## 11.4.5 Configuring the Policy Overlap Limit Strict using the CLI

The Policy Overlap Limit Strict option, **enabled** by default, strictly limits the number of overlapping policies to the maximum configured. For example, when the maximum number of overlapping policies is set to 4 (the default) and users create a fifth policy using the same filter interface, the operation fails with a validation error.

Use the following commands to disable or enable the Policy Overlap Limit Strict feature using the CLI.

```
controller-1(config)# no overlap-limit-strict
```

```
controller-1(config)# overlap-limit-strict
```

## 11.5 Viewing Information about Policies

Installing and activating overlapping policies may take more than a minute, depending on the number of overlapping policies and the number of rules in each policy.

### 11.5.1 Viewing Policy Flows

The **show policy-flow** command lists all the flows installed by the DMF application on the switches in the monitoring fabric. The following is the command syntax:

```
show policy-flow [<policy_name>]
```

Flows are sorted on a per-policy basis. Each flow entry includes the policy name for which it is configured. The packet and byte count is associated with each flow entry, as shown in the following example:

```
controller-1# show policy-flow _P1_o_P2
# Policy Name Switch
1 _P1_o_P2 DMF-CORE-SWITCH-1 (00:00:cc:37:ab:a0:90:71) 0 0 6401 1 in-port 16,vlan-vid 7 apply: name=_P1_o_P2 output: max-length=65535, port=15
2 _P1_o_P2 DMF-CORE-SWITCH-1 (00:00:cc:37:ab:a0:90:71) 0 0 6401 1 in-port 16,eth-type ipv6,vlan-vid 7 apply: name=_P1_o_P2 output: max-length=65535, port=15
3 _P1_o_P2 DMF-DELIVERY-SWITCH-1 (00:00:cc:37:ab:60:d4:74) 0 0 6401 1 in-port 49,eth-type ipv6,vlan-vid 7 apply: name=_P1_o_P2 output: max-length=65535, port=15
4 _P1_o_P2 DMF-DELIVERY-SWITCH-1 (00:00:cc:37:ab:60:d4:74) 0 0 6401 1 in-port 49,vlan-vid 7 apply: name=_P1_o_P2 output: max-length=65535, port=15
-----output truncated-----
```

### 11.5.2 Viewing Packets Dropped by Policies

The **drops** option displays the current value of the transmit drop packet counters at the filter, delivery, and core interfaces for the specified policy, as shown in the following example:

```
controller-1# show policy p1 drops
# IF Switch IF Name state speed Xmit Drops Pkt Count Xmit Drops Pkt Rate Rx Drops Pkt Count Rx Drops Pkt Rate
1 f1 00:00:00:00:00:00:00:0c s12-eth1 up 10 Gbps 0 0 0 0 0
# IF Switch IF Name state speed Xmit Drops Pkt Count Xmit Drops Pkt Rate Rx Drops Pkt Count Rx Drops Pkt Rate
1 d1 00:00:00:00:00:00:00:0c s12-eth2 up 10 Gbps 0 0 0 0 0
~ Core Interface(s) Drops ~
None.
~ Service Interface(s) Drops ~
None.
```

## 11.6 Using Rule Groups

You can use an IP address group in multiple policies, referring to the group by name in match rules. If no subnet mask is provided in the address group, it is assumed to be an exact match. For example, for an IPv4 address group, no mask is interpreted as a mask of **/32**. For an IPv6 address group, no mask is interpreted as **/128**.


You can identify only a single IP address group for a specific policy match rule. Address lists with both **src-ip** and **dst-ip** options cannot be used in the same match rule.

## 11.6.1 Using the GUI to Configure Rule Groups

To create an interface group from the **Monitoring > Interfaces** table, complete the following steps:

1. Select the **Monitoring > Rule Groups** option.

**Figure 11-15: Creating Rule Groups**



2. On the **Rule Groups** table, click on the **+** sign to create a new rule group.
3. In the pop-up menu, enter your preferred name for the rule group and, optionally, a description.

**Figure 11-16: Creating Rule Groups: Enter a Rule Group Name and Description**

The screenshot shows the 'Create Rule Group' pop-up form. It has two tabs: '1. Info' and '2. Rules'. The '1. Info' tab is active. It contains a 'Name' field with the value 'TestGroup' and a 'Description' field with the value 'Test Rule Group'. At the bottom of the form, there are buttons for 'BACK', 'NEXT', 'RESET', 'CANCEL', and 'SAVE'. The 'NEXT' button is highlighted in blue.

4. Click **NEXT** to add specific rules to your rule group.

5. In this section of the pop-up, you can add predefined rules by clicking on the options provided. In the example below, we add a rule to match all IPv4 traffic by clicking on **IPv4**.

**Figure 11-17: Creating Rule Groups: Add a Predefined Rule to the Rule Group**

Create Rule Group

1. Info ✓

2. Rules ✓

Match Rule Shortcuts

AnyARPIPv4ICMP SCTPTCPUDP IPv6 SCTP6 TCP6 UDP6

HTTP SourceHTTP DestSMTP SourceSMTP DestSMTP(SSL) Source

SMTP(SSL) DestHTTPS SourceHTTPS DestSSH SourceSSH Dest

Rules

+ —

Sequence ▲

Any Traffic

Ethertype

≡

1

—

2048 (IPv4)

Show: 1025100All

BACK

NEXT

RESET

CANCEL

SAVE

- As an alternative to the previous step, you can add your own custom rules by clicking the + sign under **Rules** and adding necessary fields in the new pop-up screen.

**Figure 11-18: Creating Rule Groups: Add Custom Rules to the Rule Group**

- Complete the dialog that appears to assign a descriptive name to the rule group.
- You can now add this rule group to DMF policies as a match condition.

## 11.6.2 Using the CLI to Configure Interface Groups

The following example shows how to configure two interface groups: a filter interface group **TAP-PORT-GRP** and a delivery interface group **TOOL-PORT-GRP**.

```
controller-1(config-switch)# filter-interface-group TAP-PORT-GRP
controller-1(config-filter-interface-group)# filter-interface TAP-PORT-1
controller-1(config-filter-interface-group)# filter-interface TAP-PORT-2
controller-1(config-switch)# delivery-interface-group TOOL-PORT-GRP
controller-1(config-delivery-interface-group)# delivery-interface TOOL-PORT-1
controller-1(config-delivery-interface-group)# delivery-interface TOOL-PORT-2
```

To view information about the interface groups in the DMF fabric, enter the **show filter-interface-group** command, as in the following examples:

- Filter Interface Groups**

```
controller-1(config-filter-interface-group)# show filter-interface-group
! show filter-interface-group TAP-PORT-GRP
# Name Big Tap IF Name Switch IF Name Direction Speed State VLAN Tag
-----|-----|-----|-----|-----|-----|-----|-----|
1 TAP-PORT-GRP TAP-PORT-1 DMF-CORE-SWITCH-1 ethernet17 rx 100Gbps up 0
2 TAP-PORT-GRP TAP-PORT-2 DMF-CORE-SWITCH-1 ethernet18 rx 100Gbps up 0
controller-1(config-filter-interface-group)#
```

- Delivery Interface Groups**

```
controller-1(config-filter-interface-group)# show delivery-interface-group
! show delivery-interface-group DELIVERY-PORT-GRP
# Name Big Tap IF Name Switch IF Name Direction Speed Ratelimit State Strip Forwarding Vlan
-----|-----|-----|-----|-----|-----|-----|-----|
1 TOOL-PORT-GRP TOOL-PORT-1 DMF-DELIVERY-SWITCH-1 ethernet15 tx 10Gbps up True
2 TOOL-PORT-GRP TOOL-PORT-2 DMF-DELIVERY-SWITCH-1 ethernet16 tx 10Gbps up True
```

## 11.7 PTP Timestamping

DMF rewrites the source MAC address of packets that match a policy with a 48-bit timestamp value sourced from a high-precision hardware clock.

- A switch with filter interface should be connected to a PTP network with a dedicated interface for the Precision Time Protocol. With a valid PTP interface, the switch will be configured in boundary clock mode and can sync the hardware clock with an available Grandmaster clock.
- Once a policy is configured to use timestamping, any packet matching on this policy will get its source MAC address rewritten with a timestamp value. Same holds true for any overlapping policy that carries traffic belonging to a user policy with timestamp enabled.
- Following options are available to configure a switch in boundary mode:
  - **domain:** Value for dataplane PTP domain (0-255) (optional)
  - **Priority1:** Value of priority1 dataplane PTP (0-255) (optional)
  - **Source IPv4 Address:** Used to restamp PTP messages from a switch to the endpoints (optional)
  - **Source IPv6 Address:** Used to restamp PTP messages from a switch to the endpoints (optional)
- Following options are available to configure an interface with role “ptp”:
  - **Announce Interval:** Set ptp announce interval between messages (-3,4). Default is 1 (optional)
  - **Delay Request Interval:** Set ptp delay request interval between messages (-7,8). Default is 5 (optional)
  - **Sync Message Interval:** Set ptp sync message interval between messages (-7,3). Default is 0 (optional)
  - **PTP Vlan:** VLANs used for Trunk or Access mode of operation for a ptp interface
- A policy should have enabled timestamping and have its filter interfaces on a switch with valid PTP config to get its packets timestamped.

### 11.7.1 Platform Compatibility

The timestamping feature is supported on 7280R3 switches.

Use the show switch all property command to check which switch in DMF fabric supports timestamping. If the following properties exist in the output, the feature is supported:

- ptp-timestamp-cap-replace-smac
- ptp-timestamp-cap-header-48bit
- ptp-timestamp-cap-flow-based

```
# show switch all property
# Switch
1 S1 (00:00:2c:dd:e9:96:2b:ff)
2 S2 (00:00:cc:1a:a3:91:a7:6c)
2 S3 (00:00:cc:1a:a3:c0:94:3e)
```

...	PTP Timestamp Supported Capabilities
...	ptp-timestamp-cap-replace-smac,
...	ptp-timestamp-cap-header-64bit,
...	ptp-timestamp-cap-header-48bit,
...	ptp-timestamp-cap-flow-based,
...	ptp-timestamp-cap-add-header-after-12



**Note:** The CLI output example above is truncated for illustrative purposes. The actual output will differ.

## 11.7.2 Configuration

The following three sections describe the configuration for PTP and timestamping:

- [PTP Switch Configuration](#)
- [PTP Interface Configuration](#)
- [Policy Configuration for Timestamping](#)



**Note:** Configuring the PTP domain for the distribution of PTP-based time sync in the network is beyond the scope of this document. We assume the network exists with a PTP domain and all necessary components.

## 11.7.3 Configuring PTP Timestamping using the CLI

Configure the switch at a global level under the **config** submode in the CLI or for each switch under the **config-switch** submode. Irrespective of the place, it has the following options:

1. **Domain:** Set the data plane PTP domain. The default value is 0. Valid values are [0 to 255] inclusive.
2. **Priority1:** Set the value of priority1 data plane PTP. The default value is 128. Valid values are [0 to 255] inclusive.
3. **Source-ipv4-address:** This is the source IPv4 address used to restamp PTP messages from this switch to the endpoints. Some master clock devices do not accept default source IP (0.0.0.0), it can be configured to sync with such devices. The default is 0.0.0.0
4. **Source-ipv6-address:** This is the source IPv6 address used to restamp PTP messages from this switch to the endpoints. Some master clock devices do not accept default source IP (::/0), it can be configured to sync with such devices. The default is ::/0

All fields are optional, and default values are selected if not configured by the user.

### Global Configuration

The global configuration is a central place to provide a common switch config for PTP. It only takes effect after creating a **ptp-interface** for a switch. Under the **config** submode, provide PTP switch properties using the following commands:

```
> enable
# config
(config)# ptp priority1 0 domain 1 source-ipv4-address 1.1.1.1
```

### Local Configuration

The local configuration provides a local PTP configuration or overrides a global PTP config for a selected switch. Select the switch using the command **switch switch name**. PTP switch config (local or global) only takes effect after creating a **ptp-interface** for a switch. Under the **config-switch** submode, provide local PTP switch properties using the following commands:

```
(config)# switch eos
(config-switch)# ptp priority1 1 domain 2
```

## 11.7.4 Configuring PTP Timestamping using the GUI

### Global Configuration

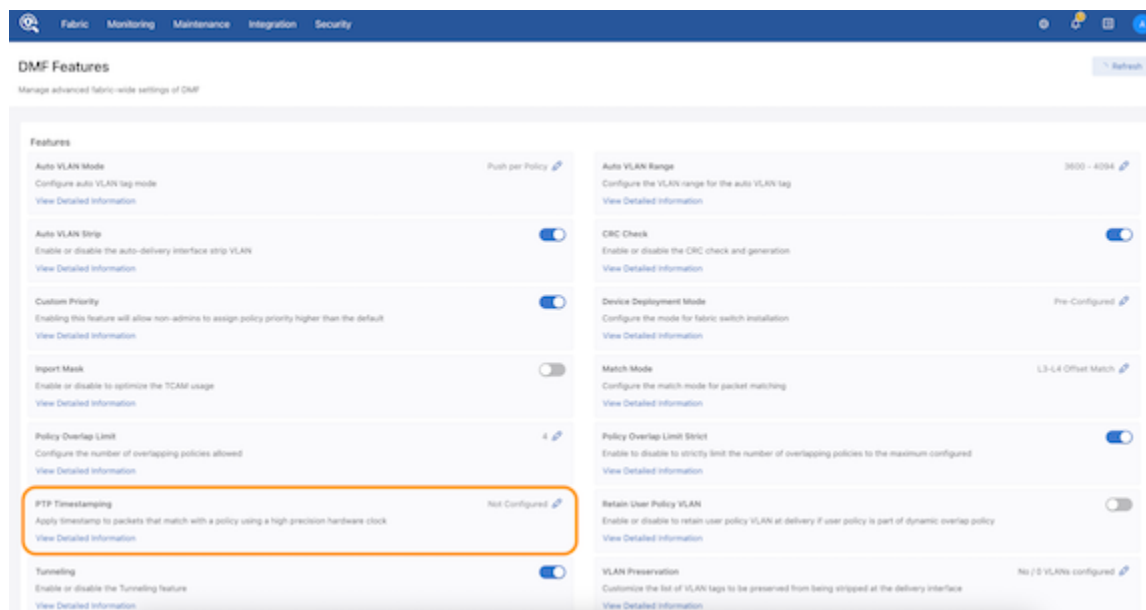
To view or edit the global PTP configuration, navigate to the DMF Features page by clicking the gear icon.

**Figure 11-19: DMF Menu Gear Icon**



The DMF Feature page is new in DMF release 8.4. It provides fabric-wide settings management for DMF. Scroll to the PTP Timestamping card and click the edit button (pencil icon) to configure or modify the global PTP Timestamping settings.

**Figure 11-20: DMF Features Page**



**Figure 11-21: Edit PTP**

Edit PTP

×

Domain ?

Priority 1 ?

Source IPv4 Address ?

Source IPv6 Address ?

Cancel

Submit

## Local Configuration



Provide a local PTP configuration for the switch or override global PTP configuration for a selected switch while configuring or editing a switch configuration (under the PTP step) using the **Monitoring > Switches** page.

**Figure 11-22: Configure Switch**

**Configure Switch**

1. Info ✓

2. PTP ✓

3. Clock ✓

4. SNMP ✓

5. SNMP Traps ✓

6. Logging ✓

7. TACACS ✓

8. sFlow ✓

9. LAG Enhanced Hash ✓

Domain: [0-255] Set the dataplane PTP domain (0-255)

Priority1: [0-255] Set the value of priority1 dataplane PTP (0-255)

Source IPv4 Address: [IPv4 Address -] Used to restamp PTP messages from this switch to the endpoints

Source IPv6 Address: [IPv6 Address -] Used to restamp PTP messages from this switch to the endpoints

Date	Time	Status
Aug 8, 2023	3:05:04pm UTC	3h 27m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 27m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24
Aug 8, 2023	3:05:04pm UTC	3h 28m 10.24

### PTP Interface Configuration

Configure a PTP Interface on the **Monitoring > Interfaces** page.

**Figure 11-23: Create Interface**

**Create Interface**

1. Interface

2. Configure

The following settings may affect the availability of some configuration options.

Auto VLAN Mode: push-per-policy

Strip Forwarding VLANs: ✓ Enabled

Name	Status	Speed	Unit	Description	Action
ethernet0/24	Up	10 Gb/s			
ethernet0/25	Up	10 Gb/s			
ethernet0/26	Up	10 Gb/s			
ethernet0/27	Up	10 Gb/s			
ethernet0/28	Up	10 Gb/s			
ethernet0/29	Up	10 Gb/s			
ethernet0/30	Up	10 Gb/s			
ethernet0/31	Up	10 Gb/s			
ethernet0/32	Up	10 Gb/s			
ethernet0/33	Up	10 Gb/s			
ethernet0/34	Up	10 Gb/s			
ethernet0/35	Up	10 Gb/s			
ethernet0/36	Up	10 Gb/s			
ethernet0/37	Up	10 Gb/s			
ethernet0/38	Up	10 Gb/s			
ethernet0/39	Up	10 Gb/s			
ethernet0/40	Up	10 Gb/s			
ethernet0/41	Up	10 Gb/s			
ethernet0/42	Up	10 Gb/s			
ethernet0/43	Up	10 Gb/s			
ethernet0/44	Up	10 Gb/s			
ethernet0/45	Up	10 Gb/s			
ethernet0/46	Up	10 Gb/s			
ethernet0/47	Up	10 Gb/s			
ethernet0/48	Up	10 Gb/s			
ethernet0/49	Up	10 Gb/s			
ethernet0/50	Up	10 Gb/s			
ethernet0/51	Up	10 Gb/s			
ethernet0/52	Up	10 Gb/s			
ethernet0/53	Up	10 Gb/s			
ethernet0/54	Up	10 Gb/s			
ethernet0/55	Up	10 Gb/s			
ethernet0/56	Up	10 Gb/s			
ethernet0/57	Up	10 Gb/s			
ethernet0/58	Up	10 Gb/s			
ethernet0/59	Up	10 Gb/s			
ethernet0/60	Up	10 Gb/s			
ethernet0/61	Up	10 Gb/s			
ethernet0/62	Up	10 Gb/s			
ethernet0/63	Up	10 Gb/s			
ethernet0/64	Up	10 Gb/s			
ethernet0/65	Up	10 Gb/s			
ethernet0/66	Up	10 Gb/s			
ethernet0/67	Up	10 Gb/s			
ethernet0/68	Up	10 Gb/s			
ethernet0/69	Up	10 Gb/s			
ethernet0/70	Up	10 Gb/s			
ethernet0/71	Up	10 Gb/s			
ethernet0/72	Up	10 Gb/s			
ethernet0/73	Up	10 Gb/s			
ethernet0/74	Up	10 Gb/s			
ethernet0/75	Up	10 Gb/s			
ethernet0/76	Up	10 Gb/s			
ethernet0/77	Up	10 Gb/s			
ethernet0/78	Up	10 Gb/s			
ethernet0/79	Up	10 Gb/s			
ethernet0/80	Up	10 Gb/s			
ethernet0/81	Up	10 Gb/s			
ethernet0/82	Up	10 Gb/s			
ethernet0/83	Up	10 Gb/s			
ethernet0/84	Up	10 Gb/s			
ethernet0/85	Up	10 Gb/s			
ethernet0/86	Up	10 Gb/s			
ethernet0/87	Up	10 Gb/s			
ethernet0/88	Up	10 Gb/s			
ethernet0/89	Up	10 Gb/s			
ethernet0/90	Up	10 Gb/s			
ethernet0/91	Up	10 Gb/s			
ethernet0/92	Up	10 Gb/s			
ethernet0/93	Up	10 Gb/s			
ethernet0/94	Up	10 Gb/s			
ethernet0/95	Up	10 Gb/s			
ethernet0/96	Up	10 Gb/s			
ethernet0/97	Up	10 Gb/s			
ethernet0/98	Up	10 Gb/s			
ethernet0/99	Up	10 Gb/s			

Interface Name: \*

Switchport Mode: Access

Announce Interval: seconds

Delay Request Interval: seconds

Sync Message Interval: seconds

PTP VLANs: + -

VLAN: No VLANs

### Timestamping Policy Configuration

DMF supports flow-based timestamping. This function requires programming a policy to match relevant traffic and enable timestamping for the matched traffic. In the Create/Edit Policy workflow (on the **Monitoring > Policies** page), use the Enable PTP Timestamping toggle to enable or disable timestamping.

**Figure 11-24: Create Timestamping Policy**

## 11.7.5 PTP Interface Configuration

A switch that syncs its hardware clock using PTP requires a physical front panel interface of the switch to be configured as a PTP interface. This interface is solely responsible for communication with the master clock and has no other purpose.

To configure the PTP interface, select an interface on the switch, as illustrated in the following command.

```
(config-switch) # interface Ethernet6/1
```

Use the **role** command to assign a **ptp** role and interface name and select **switchport-mode** for the specified interface.

```
(config-switch-if) # role ptp interface-name ptp1 access-mode announce-interval 1 delay-request-interval 1 sync-message-interval 1
```

Switchport is required to configure a PTP interface. The options for switchport mode are:

- **trunk-mode**
- **access-mode**
- **routed-mode**

The switchport mode configuration for a PTP interface is necessary to match the PTP master switch's interface configuration. The master switch could be configured to communicate PTP messages with or without a vlan tag. Use the trunk-mode with appropriate ptp vlan when the neighbor is configured for the same. If the neighbor's interface is in switch-port access mode or routed mode, we can use either of these to match it on the filter switch.

Other fields are optional; a default value is selected when no configuration is provided.

Optional fields:

- **announce-interval**: Set PTP to announce interval between messages [-3,4]. The default value is 1.
- **delay-request-interval**: Set PTP delay request interval between messages [-7,8]. The default value is 5.
- **sync-message-interval**: Set PTP sync message interval between messages (-7,3). The default value is 0.

Depending on the switchport mode selected for this interface, you can provide VLANs that will be associated with the selected **ptp-interface** using the following commands:

```
(config-switch-if)# ptp vlan 1
(config-switch-if)# ptp vlan 2
```

In routed switchport mode, we ignore the configured VLANs. In access switchport mode, the first VLAN is used for programming while ignoring the rest. In trunk switchport mode, all configured VLANs are programmed into the switch.

## 11.7.6 Policy Configuration for Timestamping

DMF supports flow-based timestamping. This function requires programming a policy to match relevant traffic and enable timestamping for the matched traffic.

Create a policy using the command **policy *policy name***.

Under **config-policy** submode, enable timestamping using the following command:

```
(config-policy)# use-timestamping
```

## 11.7.7 Using the CLI Show Commands

### PTP State Show Commands

Use the **show switch *switch name* ptp info| masters | interface | local-clock** command to obtain the PTP state of the selected switch.

The **show switch *switch name* ptp info** command summarizes the switch's PTP state and the PTP interfaces' status.

```
Controller# show switch eos ptp info
PTP Mode: Boundary Clock
PTP Profile: Default ( IEEE1588 )
Clock Identity: 0x2c:dd:e9:ff:ff:96:2b:ff
Grandmaster Clock Identity: 0x44:a8:42:ff:fe:34:fd:7e
Number of slave ports: 1
Number of master ports: 1
Slave port: Ethernet1
Offset From Master (nanoseconds): -128
Mean Path Delay (nanoseconds): 71
Steps Removed: 2
Skew (estimated local-to-master clock frequency ratio): 1.0000080070748882
Last Sync Time: 00:52:44 UTC Aug 09 2023
Current PTP System Time: 00:52:44 UTC Aug 09 2023
Interface      State      Transport  Delay
Mechanism
-----
Et1 Slave      ipv4       e2e
Et47Master     ipv4       e2e
```

The **show switch *switch name* ptp master** command provides information about the PTP master and grandmaster clocks.

```
Controller# show switch eos ptp master
Parent Clock:
Parent Clock Identity: 0x28:99:3a:ff:ff:21:81:d3
Parent Port Number: 10
Parent IP Address: N/A
Parent Two Step Flag: True
```

---

```
Observed Parent Offset (log variance): N/A
Observed Parent Clock Phase Change Rate: N/A
```

```
Grandmaster Clock:
Grandmaster Clock Identity: 0x44:a8:42:ff:fe:34:fd:7e
Grandmaster Clock Quality:
Class: 127
Accuracy: 0xfe
OffsetScaledLogVariance: 0x7060
Priority1: 120
Priority2: 128
```

The **show switch switch name ptp interface interface name** command provides PTP interface configuration and state on the device.

```
Controller# show switch eos ptp interface Ethernet1
Ethernet1
Interface Ethernet1
PTP: Enabled
Port state: Slave
Sync interval: 1.0 seconds
Announce interval: 2.0 seconds
Announce interval timeout multiplier: 3
Delay mechanism: end to end
Delay request message interval: 2.0 seconds
Transport mode: ipv4
Announce messages sent: 3
Announce messages received: 371
Sync messages sent: 4
Sync messages received: 739
Follow up messages sent: 3
Follow up messages received: 739
Delay request messages sent: 371
Delay request messages received: 0
Delay response messages sent: 0
Delay response messages received: 371
Peer delay request messages sent: 0
Peer delay request messages received: 0
Peer delay response messages sent: 0
Peer delay response messages received: 0
Peer delay response follow up messages sent: 0
Peer delay response follow up messages received: 0
Management messages sent: 0
Management messages received: 0
Signaling messages sent: 0
Signaling messages received: 0
```

The **show switch switch name ptp local-clock** command provides PTP local clock information.

```
Controller# show switch eos ptp local-clock
PTP Mode: Boundary Clock
Clock Identity: 0x2c:dd:e9:ff:ff:96:2b:ff
Clock Domain: 0
Number of PTP ports: 56
Priority1: 128
Priority2: 128
Clock Quality:
Class: 248
Accuracy: 0x30
OffsetScaledLogVariance: 0xffff
Offset From Master (nanoseconds): -146
Mean Path Delay: 83 nanoseconds
Steps Removed: 2
```

```
Skew: 1.0000081185368557
Last Sync Time: 01:01:41 UTC Aug 09 2023
Current PTP System Time: 01:01:41 UTC Aug 09 2023
```

### Policy State Show Commands

Use the **show policy** command to see if timestamping is enabled for any policy.

```
> show policy
# Policy Name Action Runtime Status Type Priority Overlap
Priority Push VLAN Filter BW Delivery BW Post Match Filter Traffic Delivery
Traffic Services Installed Time Installed Duration Ptp Timestamping
-|-----|-----|-----|-----|-----|-----|-----|-----|
---|-----|-----|-----|-----|-----|-----|-----|
1 pl unspecified action inactive Configured 100 0
1 - - - - -
True
```

## 11.7.8 Configuration Validation Messages

In **push-per-policy** mode, a validation exception occurs if a policy uses NetFlow managed-service with **records-per-interface** option and the same policy also uses timestamping. The following message appears:

```
Validation failed: Policy policy1 cannot have timestamping enabled along with
header modifying netflow service.
Netflow service netflow1 is configured with records-per-interface in push-per-
policy mode
```

In **push-per-policy** mode, a validation exception occurs if a policy uses the ipfix managed-service (using a template with **records-per-dmf-interface** key) and the same policy also uses timestamping. The following message appears:

```
Validation failed: Policy policy1 cannot have timestamping enabled along with
header modifying ipfix service.
Ipfix service ipfix1 is configured with records-per-dmf-interface in push-per-
policy mode
```

Only unicast source-ipv4-address or source-ipv6-address are allowed in the switch PTP config.

Examples of invalid ipv6 addresses: "ff02::1", "ff02::1a", "ff02::d", "ff02::5"

```
Validation failed: Source IPv6 address must be a unicast address
```

Examples of invalid ipv4 addresses: "239.10.10.10", "239.255.255.255", "255.255.255.255"

```
Validation failed: Source IPv4 address must be a unicast address
```

## 11.7.9 Troubleshooting

A policy programmed to use timestamping can fail for the following reasons:

1. The filter switch does not support syncing its hardware clock using PTP.
2. The PTP interface is not configured, or the interface is inactive.
3. PTP switch config or PTP interface config is invalid/ incomplete.
4. The PTP interface is configured on a logical port (Lag / Tunnel).

Reasons for failure will be available in the runtime state of the policy and viewed using the **show policy *policy name*** command.

As the Platform Compatibility Section describes, use the **show switch all properties** command to confirm a switch supports the feature.

### 11.7.10 Limitations

The source MAC address of the user packet is re-written with a 48-bit timestamp value on the filter switch. This action can exhibit the following behavior changes or limitations:

1. Dedup managed service will not work as expected. A high-precision timestamp can be different for duplicate packet matching on two different filter interfaces. Thus, the dedup managed service will consider this duplicate packet to be different in the L2 header. To circumvent this limitation, use an anchor/offset in the dedup managed-service config to ignore the source MAC address.
2. Any Decap managed service except for **decap-13-mps** will remove the timestamp information header.
3. The user source MAC address is lost and cannot be recovered when using this feature.
4. The **rewrite-dst-mac** feature cannot be used on the filter interface that is part of the policy using the timestamping feature.
5. In push-per-filter mode, if a user has src-mac match condition as part of their policy config, the traffic will not be forwarded as expected and can get dropped at the core switch.
6. The **in-port masking** feature will be disabled for a policy using PTP timestamping.
7. Logical ports (Lag/Tunnel) as PTP interfaces are not allowed.

# Stenographer Reference for DMF Recorder Node

---

This appendix provides information about composing Stenographer queries and submitting them through REST API.

## A.1 Stenographer Query Syntax

The DMF Recorder Node accepts Stenographer queries using a syntax based on the Berkeley Packet Filter (BPF) syntax. If a malformed BPF string is entered, the recorder node will respond with an error. The entire BPF grammar is not supported, but query strings can be composed using the predicates listed in the following table:

**Table 15: Table 1: Supported Stenographer BPF Query Strings**

BPF Predicate	Value	Description
before <b>value</b>	time	string before the specified time
before <b>value</b> m ago	duration	before <i>value</i> minutes ago
before <b>value</b> h ago	duration	before <i>value</i> hours ago
before <b>value</b> d ago	duration	before <i>value</i> days ago
before <b>value</b> w ago	duration	before <i>value</i> weeks ago
after <b>value</b>	time string	after the specified time
after <b>value</b> m ago	duration	after <i>value</i> minutes ago
after <b>value</b> h ago	duration	after <i>value</i> hours ago
vlan <b>value</b>	VLAN ID	match the specified VLAN tag (outer, inner, or inner inner)
outer vlan <b>value</b>	VLAN ID	match the specified outer VLAN tag
inner vlan <b>value</b>	VLAN ID	match the specified inner VLAN tag (or middle tag of triple-tagged packets)
inner vlan <b>value</b>	VLAN ID	match the specified innermost VLAN tag of triple-tagged packets
src mac <b>value</b>	MAC address	match the specified MAC address in typical colon-delimited form (e.g. 11:22:33:44:55)
dst mac <b>value</b>	MAC address	match the specified MAC address in typical colon-delimited form (e.g. 11:22:33:44:55)
mpls <b>value</b>	MPLS label	match the specified MPLS label
src host <b>value</b>	IPv4/v6 address	match the specified source address exactly
dst host <b>value</b>	IPv4/v6 address	match the specified destination address exactly
src net <b>value</b>	IPv4/v6 address	match the specified source address with an optional CIDR mask. All octets of address must be specified, e.g. good → 1.2.3.0/24, bad → 1.2.3/24
src net <b>value</b> mask <b>value</b>	IPv4/v6 address	match the specified source address with masked with the specified address



BPF Predicate	Value	Description
dst net <b>value</b>	IPv4/v6 address	match the specified destination address with an optional CIDR mask. All octets of address must be specified, e.g. good → 1.2.3.0/24, bad → 1.2.3/24
dst net value mask <b>value</b>	IPv4/v6 address	match the specified destination address with masked with the specified address
ip proto <b>value</b>	protocol number	match the specified IP protocol number
icmp		match ICMP packets (shortcut for "ip proto 1")
tcp		match TCP packets (shortcut for "ip proto 6")
udp		match UDP packets (shortcut for "ip proto 17")
src port <b>value</b>	transport port number	match the specified transport port number
dst port <b>value</b>	transport port number	match the specified transport port number
cid <b>value</b>	Community ID	match the provided community ID in standard version:base-64 encoded form (e.g. 1:hO+sN4H+MG5MY/8hIrXPqc4ZQz0=)
policy <b>value</b>	DMF policy name	match the forwarding VLAN(s) of the specified DMF policy. Only supported through the DMF controller. Not supported when using a Recorder Node REST API directly.
filter-interface <b>value</b>	DMF filter interface name	match the forwarding VLAN of the specified filter interface. Only supported through the DMF controller. Not supported when using a Recorder Node REST API directly.
event <b>value</b>	Recorder Node event name	match the time range of the specified event. Only supported through the DMF controller. Not supported when using a Recorder Node REST API directly.
and		logical "and"

BPF Predicate	Value	Description
&&		logical “and”
or		logical “or”
		logical “or”
(		begin grouping
)		end grouping

## A.2 Example Stenographer Queries



**Note:** It is recommended that a specific time range always be included in each query.

After two hours ago but before one hour ago, search for all packets to or from Google DNS (8.8.8.8).

```
(after 2h ago and before 1h ago) and (src host 8.8.8.8 or dst host 8.8.8.8)
```

In the last twenty-four hours, search for all SSH (TCP port 22) packets destined to IP 10.4.100.200.



**Note:** This will not match any SSH packets from 10.4.100.200.

```
after 24h ago and dst host 10.4.100.200 and tcp and src port 22
```

Within the last five minutes, search for all packets to or from 10.1.1.100. And, in the five minutes before that, search for all packets to or from 10.1.100.101.

```
(after 5m ago and (src host 10.1.1.100 or dst host 10.1.1.100)) or (after 10m ago and before 5m ago and (src host 10.1.1.101 or dst host 10.1.1.101))
```

Within the timespan of event `abc` and within the last hour, search for all SSH (TCP port 22) packets destined to IP 1.2.3.4.

```
(event abc or after 1h ago) and dst host 1.2.3.4 and tcp and dst port 22
```

Within the timespan defined by the intersection of events `abc` and `def`, search for all packets sent from any IP in subnet 1.2.3.0/24 that were seen on filter interface `xyz`.

```
(event abc and event def) and filter-interface xyz and src net 1.2.3.0/24
.. note::
To use the filter-interface predicate, the DMF controller must be in the push-per-filter Auto VLAN mode.
```

Within the last five minutes, search for all packets sent from IP 1.2.3.4 that were sent to the DMF Recorder Node using DMF policy `abc`.

```
after 5m ago and policy abc and src host 1.2.3.4
.. note::
To use the policy predicate the DMF controller must be in the push-per-policy or push-per-filter Auto VLAN mode. When in push-per-policy auto-vlan-mode, the policy's forwarding tag will be queried. When in push-per-filter mode, the forwarding tags of the filter interfaces used in
```

```
the policy are queried.
```

Within the last five minutes, search for all packets with any VLAN tag 100.

```
after 5m ago and vlan 100
```

Within the last five minutes, search for all packets with an outer VLAN tag 100.

```
after 5m ago and outer vlan 100
```

Within the last five minutes, search for all packets with an inner (or middle) VLAN tag 100.

```
after 5m ago and inner vlan 100
```

Within the last five minutes, search for all triple-tagged packets with innermost VLAN tag 100.

```
after 5m ago and inner inner vlan 100
```

Within the last five minutes, search for packets belonging to a flow with community ID of 1:hO+sN4H+MG5MY/8hIrXPqc4ZQz0=.

```
after 5m ago and cid 1:hO+sN4H+MG5MY/8hIrXPqc4ZQz0=
```

This matches packets in each direction of the flow, if applicable.

Within the last five minutes, search for all L2 broadcast packets originating from MAC address 11:22:33:44:55:66.

```
after 5m ago and src mac 11:22:33:44:55:66 and dst mac ff:ff:ff:ff:ff:ff
```

# DMF Recorder Node REST API

---

The REST server is available over HTTPS on the default port (**443**) using either of the two authentication methods supported:

- HTTP basic: The client presents a valid username and password for the controller with which the recorder is connected. The DMF Recorder Node verifies at the DMF controller if the provided username and password is valid and has sufficient privileges to use the Recorder REST API.
- Authentication tokens: Revocable authentication tokens are accepted by the DMF Recorder Node REST API as an alternative to HTTP basic. Valid authentication tokens are configured in the controller and pushed down to the DMF Recorder Node using a gentable. Any client with a valid authentication token will be permitted to query the DMF Recorder Node REST API without real time consultation with the controller.

Some APIs accept a Stenographer query string as input or return a Stenographer query string as output. A Stenographer query string is a BPF-like syntax for defining the scope of a query. Packets that match this scope are included in the query result or operation. For details about the Stenographer query syntax supported by the DMF Recorder Node, refer to [Stenographer Reference for DMF Recorder Node](#). The DMF Recorder Node provides a REST API so that clients can look up packets and meta data. The REST server runs securely (HTTPS) on TCP port **443**.

## B.1 Authentication

Clients must either authenticate using a valid DANZ Monitoring Fabric controller username and password over HTTP Basic or with an authentication token that has been configured on the DMF controller specifically for DMF Recorder Node REST API authentication.

### B.1.1 Basic HTTP Authentication

A valid DMF controller username and password can be used to authenticate with a DMF Recorder Node over its REST API. The recorder node delegates authentication to the DMF controller. If the username and password provided are valid, the recorder node proceeds to authorize the user for the recorder node REST endpoint being invoked.

In the following example, a query is made using the HTTP Basic method of authentication:

```
$ curl https://1.2.3.4/query/window -u admin:12345 -k | python -m json.tool
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 79 100 79 0 0 143 0 --:--:-- --:--:-- --:--:-- 143
{
  "begin": "2019-01-23 15:15:23 +0000 UTC",
  "end": "2019-02-04 17:39:52 +0000 UTC"
}
```

In this example, the recorder node IP address is **1.2.3.4**. The username on the DMF controller is **admin** and the password is **12345**.

### B.1.2 Authentication with an Authentication Token

The use of an authentication token is primarily designed for third-party applications and automation scripts where creating an account or storing a username and password is not desirable. This method can also allow access to a DMF Recorder Node if the management network connection to the controller is disrupted.

To create an authentication token, login to the DMF controller associated with the recorder node, then complete the following steps:

1. Change to config mode on the active DMF controller.

```
controller-1# configure
controller-1(config)#
```

2. Define the authentication token using a unique name.

```
controller-1(config)# recorder-node auth token my-token
Auth : my-token
Token : the-secret-token
```



**Note:** This name does not need to be secret. This example uses the name my-token.

The controller generates a secret token for you (in this example, the-secret-token). Treat this token as private. Anyone who presents it to the DMF Recorder Node can use the DMF Recorder Node REST APIs.



**Note:** Only the non-reversible hash of this token is stored on the DMF Recorder Node and controller. There is no way to recover the token if it is lost. (See below for how to revoke the token in the event it is lost or compromised.)

The controller stores the token hash and the name assigned, which can be seen by entering the **show running-config recorder-node** command, as in the following example:

```
controller-1(config)# show running-config recorder-node auth token
! recorder-node
recorder-node auth token my-token $2a$12$pxm62t15rMD8c4vSrzu6X.DTjeoBmRUw
ZvTkvnXatsZ8TFb4PxanC
```

If the token is lost or compromised, remove it from the controller and the controller will fail any attempt to authenticate to the recorder using the token.

```
controller-1(config)# no recorder-node auth token my-token
controller-1(config)# show running-config recorder-node auth token
controller-1(config)#
```

The following example shows a query using the authentication token method. The authentication token is defined in the HTTP request as the value of the cookie header.

```
$ curl https://1.2.3.4/query/inventory/window --header "Cookie:plaintext-
secret-auth-token" -k |
python
-m json.tool
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 79 100 79 0 0 83 0 --:--:-- --:--:-- --:--:-- 83
{
  "begin": "2019-01-23 15:15:23 +0000 UTC",
  "end": "2019-02-04 17:50:46 +0000 UTC"
}
```

In this example, the DMF Recorder Node IP address is **1.2.3.4**. The authentication token has already been generated on the DMF controller associated with the recorder node and is included in the cookie header as **plaintext-secret-auth-token**.

You must include the plaintext authentication token and not the token hash, which is saved in the controller running configuration. If you do not know the plaintext token, revoke access for the token and generate a new one. Make note of the plaintext value displayed after the token has been generated.

## B.2 DMF Recorder Node API Headers

The supported REST API HTTP header entries are listed in the following table:

**Table 16: DMF Recorder Node REST API HTTP Headers**

Header	Value Type	Description
Steno-Limit-Bytes:value	integer	max number of bytes to accept in a packet query response
Steno-Limit-Packets:value	integer	max number of packets to accept in a packet query response
Cookie:value	string	auth token to use in lieu of HTTP basic auth

## B.3 DMF Recorder Node REST APIs

The supported DMF Recorder Node REST APIs are listed below:

### B.3.1 /ready

- Description: Is the DMF Recorder Node able to accept queries? Return payload indicates progress towards start up completion.
- HTTP Method: GET
- Request Payload:
- Return MIME Type:
- Return Payload:

```
{
  "current-value": <int>,
  "max-value": <int>,
  "percent-complete": <float>
}
```

- Return Status Code:
  - 200, ready
  - 503, not ready

### B.3.2 /query/window

- Description: Get timestamp of oldest and newest packet available for query.
- HTTP Method: GET
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "begin" : <RFC-3339>,
  "end" : <RFC-3339>
}
```

- Return Status Code:

- 200, success
- 400, input error
- 500, internal error
- 503, not ready

### B.3.3 **/query/size**

- Description: Get count and aggregate size of packets matching provided filter.
- HTTP Method: POST
- Request Payload: Stenographer query string
- Return MIME Type: application/json
- Return Payload:

```
{
  "packet-count" : <int>,
  "aggregate-size" : <int>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.4 **/query/application**

- Description: Perform DPI on packets matching provided filter. DPI is performed using nDPI.
- HTTP Method: POST
- Request Payload: Stenographer query string
- Return MIME Type: application/json
- Return Payload: Defined by nDPI
- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.5 **/query/packet**

- Description: Download pcap of packets matching provided filter.
- HTTP Method: POST
- Request Payload: Stenographer query string
- Return MIME Type: application/vnd.tcpdump.pcap
- Return Payload: .pcap file
- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

---

### B.3.6 `/query/analysis[filter="<stenographer-query-string>"][[type = "<analysis-type>"]`

- Description: Perform an analysis on the packets matching the stenographer query string. Supported values for *analysis-type* are:
  - analysis\_http\_tree
  - analysis\_http\_stat
  - analysis\_http\_req\_tree
  - analysis\_http\_srv\_tree
  - analysis\_dns\_tree
  - analysis\_hosts
  - analysis\_conv\_ipv4
  - analysis\_conv\_ipv6
  - analysis\_conv\_tcp
  - analysis\_conv\_udp
  - analysis\_rtp\_streams
  - analysis\_sip\_stat
  - analysis\_conv\_sip
  - analysis\_tcp\_packets
  - analysis\_tcp\_flow\_health
- HTTP Method: GET
- Request Payload:
- Return MIME Type: application/json
- Return Payload: Determined by the analysis type selected.
- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.7 `/query/replay/request[filter="<stenographer-query-string>"][[real-time="<boolean>"]`

- Description: Asynchronously request packets matching filter be replayed into the monitoring fabric. Replay is performed using tcp replay.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "id" : <int>,
  "message": <string>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready



### B.3.8 `/query/replay/request[filter="<stenographer-query-string>"] [speed-mbps="<int>"]`

- Description: Asynchronously request packets matching filter be replayed into the monitoring fabric. Replay is performed using tcp replay.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "id" : <int>,
  "message": <string>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.9 `/query/replay/done`

- Description: Check the status of a replay matching the provided ID. Message contains replay result from tcp replay.
- HTTP Method: POST
- Request Payload: Replay ID
- Return MIME Type: application/json
- Return Payload:

```
{
  "id" : <int>,
  "done" : <boolean>,
  "message": <string>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 404, replay ID unknown
  - 406, replay not done
  - 500, internal error
  - 503, not ready

### B.3.10 `/erase/packet[filter="<stenographer-query-string>"]`

- Description: Erase packets matching the provided filter. Note that any packet not matching the filter but in the same packet file of a packet matching the filter will also be deleted.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "bytes-erased" : <int>,
  "message" : <string>
}
```

---

```
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.11 **/event/update[trigger="<boolean>"][name="<string>"][pre-buffer-minutes="<int>"]**

- Description: Trigger or terminate the named event. Set pre-buffer-minutes to 0 to use the available pre-buffer.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{  
  "message" : <string>,  
  "event-queued" : <boolean>  
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.12 **/abort/query**

- Description: Terminate a particular query defined by the provided Stenographer query string.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{  
  "message" : <string>  
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.13 **/abort-all/query**

- Description: Terminate all running queries.
- HTTP Method: POST
- Request Payload:
- Return MIME Type: application/json

- Return Payload:

```
{
  "message" : <string>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.14 /queries

- Description: Determine the currently running queries, enumerated by the Stenographer query string of the query.
- HTTP Method: GET
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "queries" : [
    <stenographer-query-string>, ...
  ]
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.15 /status/query

- Description: Determine how far a given query has progressed. This can be used to estimate the time remaining to run the query.
- HTTP Method: GET
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "query" : <stenographer-query-string>, "current-value" : <int>,
  "max-value" : <int>, "percent-complete" : <float>
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

### B.3.16 /status/all

- Description: Determine how far all queries have progressed. This can be used to estimate the time remaining to run the queries.
- HTTP Method: GET
- Request Payload:
- Return MIME Type: application/json
- Return Payload:

```
{
  "queries" : [
    {
      "query" : <stenographer-query-string>,
      "current-value" : <int>,
      "max-value" : <int>,
      "percent-complete" : <float>
    },
    ...
  ]
}
```

- Return Status Code:
  - 200, success
  - 400, input error
  - 500, internal error
  - 503, not ready

## Configuring Third-party Services

### C.1 Services in the DANZ Monitoring Fabric

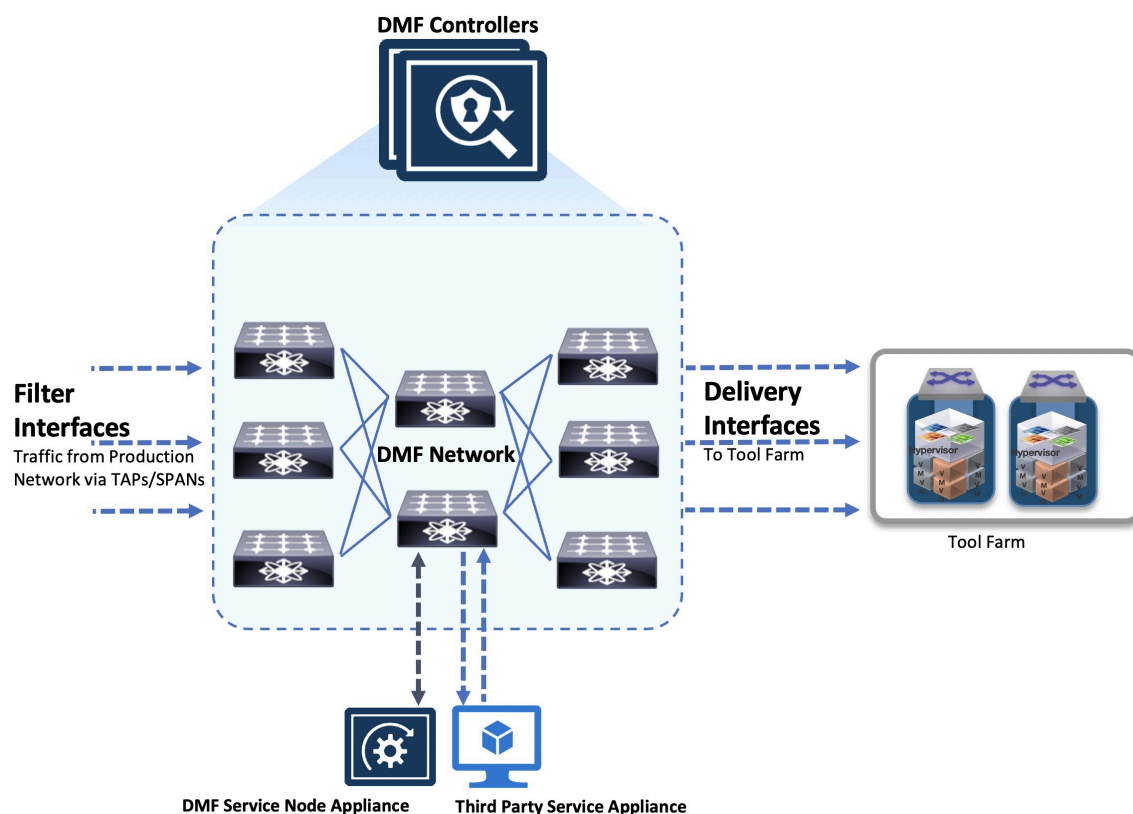
Services in the DANZ Monitoring Fabric refer to packet modification operations provided by third-party network packet brokers (NPBs), referred to as service nodes. Services can include various operations that refine or modify the data stream that is delivered to analysis tools.

Each service instance is assigned a numeric identifier because multiple services can be specified for a given policy. Services are applied in sequential order; a service with a lower sequence number is applied first.

Service nodes are optional sets of devices that process interesting traffic before it is forwarded to the delivery ports specified by the policy. Example services include time-stamping packets, packet slicing, or payload obfuscation. To configure a service node:

- Create all the pre-service and post-service interfaces that will be used with the service.
- Create a service node and add pre-service and post-service interfaces using the DMF interface names.

**Figure C-1: Using Services with a Policy**



In the figure above, the time-stamping service is applied first, followed by the packet-slicing service. The illustration shows the CLI commands used to associate the service with a specific policy. For the illustrated policy, the packet path is as follows:

1. Filter interface (**F3**)
2. Time-stamping service node (pre-service and post-service interfaces)
3. (optional) Packet-slicing service node (pre-service and post-service interfaces)
4. Delivery-interface (**D2**)

Once a service has been included in a policy, it is only optional if specifically defined as optional. If not defined as optional in the policy, packet forwarding does not occur when the service is not available. For example, if the packet-slicing service is configured as optional and if a pre-service or post-service interface assigned to that service node is down, then the service is skipped and the packets are delivered to the **D2** delivery interface after the time-stamping service is completed. However, if at least one pre-service and post-service interface is not available for the time-stamping service, packets are not forwarded to the delivery interfaces by this policy.

Configure all the service interfaces before you create a service definition that uses them.



**Note:** Before defining a service, first create the service interface names. Otherwise, the service might enter an inconsistent state. If that happens, delete the service definition, create the interfaces, then re-create the service definition. Alternatively, re-create the service definition without the nonexistent interfaces.

A DMF service can have multiple pre-service and post-service interfaces. A Link Access Group (LAG) can be used either as a pre-service or a post-service interface.



**Note:** For any DMF service, Arista strongly recommends configuring the post-service and pre-service interfaces on the same switch.

### C.1.1 Using the GUI to Configure a DMF Unmanaged Service

To create a DMF unmanaged service, perform the following steps:

1. Select **Monitoring > Services**.

The system displays the following table:

**Figure C-2: DMF Unmanaged Service**

Name	Description	Pre-Service Interface Count	Post-Service Interface Count	Policy Count	Total Inbound Bit Rate	Total Outbound Bit Rate	Max Inbound Bandwidth	Max Outbound Bandwidth
THIRD-PARTY-SERVICE-1		1	1	0			10 Gbit/s	10 Gbit/s

This table lists the services configured for the DANZ Monitoring Fabric and lets you add, delete, or modify existing services.

2. To create a new service, click the provision control (+) at the upper-left corner of the table.

The system displays the following dialog:

**Figure C-3: Create Service Dialog: Info**

Create Service

1. Info ✓

2. Pre-Service Interfaces ✓

3. Post-Service Interfaces ✓

Name \*  
THIRD-PARTY-SERVICE-1

Description

BACK NEXT RESET CANCEL SAVE

3. Type a unique name for the service and optional text description, then click **Next**.

The system displays the following dialog:

**Figure C-4: Create Service Dialog: Pre-service Interfaces**

Create Service

1. Info ✓

2. Pre-Service Interfaces ✓

3. Post-Service Interfaces ✓

Click + and - to include or exclude interfaces.

+ -

	Name ▲	Switch DPID	Switch Name	Interface	Description	Groups	State
No interfaces							

BACK NEXT RESET CANCEL SAVE

This table lists the interfaces assigned as pre-service interfaces for the current service.

4. To add a pre-service interface, click the provision control (+) at the top of the table.

The system displays the following dialog:

**Figure C-5: Select Pre-service Interfaces**

	Name	Switch DPID	Switch Name	Interface	Description	Groups
<input type="checkbox"/>	POST-SERVICE-INTF-1	00:00:cc:37:ab:be:11:20	leaf-as5812X-1	ethernet16	—	—
<input type="checkbox"/>	PRE-SERVICE-INTF-1	00:00:cc:37:ab:be:11:20	leaf-as5812X-1	ethernet15	—	—

This table lists the interfaces available for assignment as pre-service interfaces. To configure a new interface, click the provision control (+) at the top of the table. The system displays a dialog for adding a service interface.

5. Enable the checkbox for one or more interfaces to assign as a pre-service interface for the current service and click **Append Selected**.
6. On page two of the **Create Service Interface** dialog, click **Next**.

The system displays the following dialog:

**Figure C-6: Create Service Dialog: Post-service Interfaces**

Click + and - to include or exclude interfaces.

	Name	Switch DPID	Switch Name	Interface	Description	Groups	State
No interfaces							

This table lists the interfaces assigned as post-service interfaces for the current service.

7. To add a post-service interface, click the provision control (+) at the top of the table.



The system displays the following dialog.

**Figure C-7: Select Post-service Interfaces**

	Name	Switch DPID	Switch Name	Interface	Description	Groups
<input type="checkbox"/>	POST-SERVICE-INTF-1	00:00:cc:37:ab:be:11:20	leaf-as5812X-1	ethernet16	—	—
<input type="checkbox"/>	PRE-SERVICE-INTF-1	00:00:cc:37:ab:be:11:20	leaf-as5812X-1	ethernet15	—	—

Show: 10

BACK CANCEL APPEND SELECTED

This table lists the interfaces available for assignment as post-service interfaces.

To configure a new interface, click the provision control (+) at the top of the table. The system displays a dialog for adding a service interface, as described in the **Configuring DMF Unmanaged Services** section.

8. Enable the checkbox for one or more interfaces to assign as a post-service interface for the current service and click **Append Selected**.
9. Click **Save** on page three of the **Create Service Dialog**.

### C.1.2 Using the CLI to Configure a DMF Unmanaged Service

In the DANZ Monitoring Fabric, third-party tools that provide packet manipulation services, such as time stamping and packet slicing, are referred to as DMF Unmanaged Services. These are optional devices that process traffic from filter interfaces before being forwarded to delivery interfaces.



**Note:** Once a service is added to a policy, it is no longer optional unless you specifically define it as optional. If not defined as optional, packets are not forwarded by the policy if the service is not available.

To configure an unmanaged service using the CLI, complete the following steps:

1. Create one or more pre-service interfaces for delivering traffic to the NPB, as in the following example.

```
controller-1(config-switch-if) # switch DMF-CORE-SWITCH
controller-1(config-switch-if) # interface s9-eth1
controller-1(config-switch-if) # role service interface-name pre-serv-intf-1
```

2. Create one or more post-service interfaces for receiving traffic from the NPB, as in the following example:

```
controller-1(config-switch-if) # interface s9-eth2
controller-1(config-switch-if) # role service interface-name post-serv-intf-1
```

3. Create a service node and add at least one pre-service and at least one post-service interface using the DMF interface names, as in the following example:

```
controller-1(config)#controller-1(config)# unmanaged-service THIRD-PARTY-SERVICE-1
controller-1(config-unmanaged-srv)# description "this is a third-party unmanaged service"
controller-1(config-unmanaged-srv)# pre-service PRE-SERVICE-INTF-1
controller-1(config-unmanaged-srv)# post-service POST-SERVICE-INTF-1
```

To list the services that have been configured in the DMF fabric, enter the **show unmanaged-services** command, as in the following example:

```
controller-1# show unmanaged-service
# Service Name      Max from service bandwidth bps Max to service bandwidth bps Total from service bps Total to service bps
-----
1 THIRD-PARTY-SERVICE-1 10Gbps
~~~~~ Post-groups of Service Names ~~~~~
# Service Name      Dmf name
-----
1 THIRD-PARTY-SERVICE-1 POST-SERVICE-INTF-1
~~~~~ Pre-groups of Service Names ~~~~~
# Service Name      Dmf name
-----
1 THIRD-PARTY-SERVICE-1 PRE-SERVICE-INTF-1
```


To display information about a service, specify the service name, as in the following example:

```
controller-1 # show unmanaged-service THIRD-PARTY-SERVICE-1
# Service Name      Max from service bandwidth bps Max to service bandwidth bps Total from service bps Total to service bps
-----
1 THIRD-PARTY-SERVICE-1 10Gbps
~~~~~ Post-groups of Service Names ~~~~~
# Service Name      Dmf name
-----
1 THIRD-PARTY-SERVICE-1 POST-SERVICE-INTF-1
~~~~~ Pre-groups of Service Names ~~~~~
# Service Name      Dmf name
-----
1 THIRD-PARTY-SERVICE-1 PRE-SERVICE-INTF-1
```

## C.2 Service Insertion and Chaining in a DMF Policy

To configure a DMF policy that uses services provided by an NPB, add the **use-service** command to the policy. Services can be configured in series, called chaining, as shown below:

**Figure C-8: Service Insertion and Chaining**



Because multiple services can be specified in a given policy, specify the sequence number of each service instance so the services are applied in order for the policy traffic. A lower sequence number means the service is applied first.

To configure a DMF out-of-band policy that uses services provided by an NPB, use the **use-service** command from the config-policy submode to add the service to the policy.

The following are the configuration commands for implementing the illustrated example:

```
controller-1(config)# policy DMF-POLICY-1
controller-1(config-policy)# use-service UMS-DEDUPLICATE-1 sequence 100
controller-1(config-policy)# use-service UMS-TIMESTAMP-1 sequence 101
```

In this example, the packet deduplication service is applied first, followed by time stamping. If all the pre-service interfaces or all the post-service interfaces for the packet-slicing service nodes are down, then this service is skipped if it is configured as optional. In this example, the time-stamping service is applied before the packet deduplication service, and the packet deduplication service is configured as optional.

```
controller-1(config)# policy DMF-POLICY-1
controller-1(config-policy)# use-service UMS-TIMESTAMP-1 sequence 100
controller-1(config-policy)# use-service UMS-DEDUPLICATE-1 sequence 101
optional
.. note::
If a service is inserted, the policy can only become active and begin
forwarding when at
least one delivery port is reachable from all the post-service interfaces
defined for the service.
```

To display the run time services being applied, enter the **show policy** command from any mode.

## References

---

### D.1 Related Documents

The following documentation is available for ***DANZ Monitoring Fabric***:

- *DANZ Monitoring Fabric Release Notes*
- *DANZ Monitoring Fabric User Guide*
- *DANZ Monitoring Fabric Deployment Guide*
- *DANZ Monitoring Fabric Hardware Compatibility List*
- *DANZ Monitoring Fabric Hardware Guide*
- *DANZ Monitoring Fabric Verified Scale Guide*
- *DANZ Monitoring Fabric REST API Guide*
- *DANZ Monitoring Fabric SNMP MIB Reference Guide*