

ARISTA

User Guide

DANZ Monitoring Fabric

Version 8.6



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
www.arista.com/en/	support@arista.com	sales@arista.com

© Copyright 2024 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 is subject to the Arista Networks Terms of Use Policy, available at www.arista.com/en/terms-of-use. Use of marks belonging to other parties is for informational purposes only.

Contents

Chapter 1: Introduction to DMF and Overview.....	1
1.1 Out-of-band Monitoring with DANZ Monitoring Fabric.....	2
1.2 Using the DANZ Monitoring Fabric CLI.....	4
1.3 Using the DANZ Monitoring Fabric GUI.....	7
1.3.1 DMF Features Page.....	11
1.3.2 Dashboard Layout.....	27
1.3.3 Controller Health.....	28
1.3.4 Switch Health.....	33
1.3.5 Policy Health.....	42
1.3.6 Smart Node Health.....	44
1.3.7 Empty State.....	47
Chapter 2: Managing DMF Switches and Interfaces.....	49
2.1 Overriding the Default Configuration for a Switch.....	49
2.1.1 CLI Configuration.....	49
2.2 DMF Interfaces.....	49
2.2.1 Using the GUI to Configure a DMF Filter or Delivery Interface.....	50
2.2.2 Using the CLI to Configure a DANZ Filter or Delivery Interface.....	52
2.2.3 Using the CLI to Identify a Filter Interface using Destination MAC Rewrite.....	54
2.2.4 Using the GUI to Identify a Filter Interface using Destination MAC Rewrite.....	56
2.2.5 Forward Slashes in Interface Names.....	61
2.3 Using Interface Groups.....	62
2.3.1 Using the GUI to Configure Interface Groups.....	62
2.3.2 Using the CLI to Configure Interface Groups.....	63
2.4 Switch Light CLI Operational Commands.....	64
Chapter 3: Managing DMF Policies.....	68
3.1 Overview.....	68
3.2 DMF Policies Page.....	69
3.3 Policy Elements.....	86
3.4 Configuring a Policy.....	88
3.4.1 Configure a Policy Using the GUI.....	88
3.4.2 Configure a Policy Using the CLI.....	120
3.4.3 Define Out-of-band Match Rules Using the CLI.....	122
3.4.4 Stop, Start, and Schedule a Policy Using the CLI.....	122
3.4.5 Clear a Policy Using the CLI.....	123
3.4.6 View Policies Using the CLI.....	124
3.4.7 Rename a Policy Using the CLI.....	125
3.5 Using the Packet Capture Action in a Policy.....	126
3.6 Define a Policy with a Packet Capture Action Using the CLI.....	127
Chapter 4: Viewing Information about Monitoring Fabric and Production Networks.....	129
4.1 Monitoring DMF Interfaces.....	129
4.1.1 Using the GUI to Monitor DMF Interfaces.....	129
4.1.2 Viewing Oversubscription Statistics.....	129

4.1.3 Using the CLI to Monitor Interface Configuration.....	130
4.2 Viewing Devices Connected to the Monitoring Fabric.....	132
4.2.1 Using the GUI to View Fabric-Connected Devices.....	132
4.2.2 Using the CLI to View Switch Configuration.....	133
4.3 Viewing Information about a Connected Production Network.....	134
4.3.1 Using the CLI to View Connected Devices and LAGs.....	134

Chapter 5: Using the DMF Service Node Appliance.....136

5.1 Overview.....	136
5.2 Changing the Service Node Default Configuration.....	136
5.3 Using SNMP to Monitor DPDK Service Node Interfaces.....	137
5.4 Configuring Managed Services.....	137
5.4.1 Using the GUI to Define a Managed Service.....	138
5.4.2 Using the CLI to Define a Managed Service.....	139
5.4.3 Monitoring Managed Services.....	141
5.5 Deduplication Action.....	141
5.6 Header Strip Action.....	143
5.6.1 Configuring the Post-service Match.....	145
5.7 IPFIX and Netflow Actions.....	145
5.7.1 IPFIX Template.....	146
5.7.2 IPFIX Keys.....	146
5.7.3 IPFIX Fields.....	147
5.7.4 Active and Inactive Timers.....	147
5.7.5 Example Flowset.....	147
5.7.6 Using the GUI to Define an IPFIX Template.....	148
5.7.7 Using the CLI to Define an IPFIX Template.....	151
5.7.8 Using the GUI to Define an IPFIX Service Action.....	151
5.7.9 Using the CLI to Define an IPFIX Service Action.....	152
5.7.10 Records Per Interface Netflow using DST-MAC Rewrite.....	152
5.8 Packet-masking Action.....	154
5.9 Arista Analytics Node Capability.....	154
5.9.1 Configuring the Arista Analytics Node Using the GUI.....	155
5.9.2 Configuring the Arista Analytics Node Using the CLI.....	156
5.10 Pattern-drop Action.....	158
5.11 Pattern-match Action.....	159
5.12 Slice Action.....	160
5.13 Packet Slicing on the 7280 Switch.....	162
5.13.1 Using the CLI to Configure Packet Slicing - 7280 Switch.....	162
5.13.2 Using the GUI to Configure Packet Slicing - 7820 Switch.....	164
5.13.3 Troubleshooting Packet Slicing.....	170
5.14 VXLAN Stripping on the 7280R3 Switch.....	171
5.14.1 Using the CLI to Configure VXLAN Header Stripping.....	172
5.14.2 Using the GUI to Configure VXLAN Header Stripping.....	174
5.14.3 Syslog Messages.....	178
5.14.4 Troubleshooting.....	178
5.14.5 Limitations.....	179
5.15 Session Slicing for TCP and UDP Sessions.....	179
5.15.1 Using the CLI to Configure Session Slicing.....	179
5.15.2 Using the GUI to Configure Session Slicing.....	180
5.16 Timestamp Action.....	185
5.17 UDP-replication Action.....	186
5.18 Redundancy of Managed Services in Same DMF Policy.....	188
5.18.1 Using the GUI to Configure a Backup Managed Service.....	188
5.18.2 Using the CLI to Configure a Backup Managed Service.....	188
5.19 Application Identification.....	189

5.19.1 Using the CLI to Configure Application Identification.....	189
5.19.2 Using the CLI to Configure app-id.....	189
5.19.3 Using the CLI to Configure app-id-filter.....	190
5.19.4 Using the CLI to Configure app-id and app-id-filter Combined.....	192
5.19.5 Using the GUI to Configure app-id and app-id-filter.....	192
5.19.6 Dynamic Signature Updates (Beta Version).....	197
5.19.7 CLI Show Commands.....	199
5.19.8 Syslog Messages.....	200
5.19.9 Scale.....	201
5.19.10 Troubleshooting.....	201
5.19.11 Limitations.....	201
5.20 Redundancy of Managed Services Using Two DMF Policies.....	202
5.21 Cloud Services Filtering.....	203
5.21.1 Cloud Services Filtering Configuration.....	203
5.22 Multiple Services Per Service Node Interface.....	204
5.23 Sample Service.....	205
5.23.1 Sample Service CLI Configuration.....	206
5.23.2 Sample Service GUI Configuration.....	207
5.23.3 Troubleshooting Sample Service.....	208
5.23.4 Limitations.....	209
5.24 Flow Diff Latency and Drop Analysis.....	209
5.24.1 Configure Flow Diff Latency and Drop Analysis Using the CLI.....	209
5.24.2 Configuring Tap Points.....	210
5.24.3 Configuring Policy.....	211
5.24.4 Show Commands.....	212
5.24.5 Syslog Messages.....	213
5.24.6 Troubleshooting.....	213
5.24.7 Limitations.....	214
5.24.8 Resources.....	215
5.25 Service Node Management Migration L3ZTN.....	215
Chapter 6: Using the DMF Recorder Node.....	217
6.1 Overview.....	217
6.2 Configuration Summary.....	217
6.3 Indexing Configuration.....	218
6.4 Pre-buffer Configuration and Events.....	219
6.5 Using an Authentication Token.....	219
6.6 Using the GUI to Add a Recorder Device.....	219
6.6.1 Configuring a Node to Use Local Storage.....	223
6.6.2 Configuring a Node to Use External Storage.....	224
6.7 Configuring a Recorder Node Interface.....	227
6.8 Using the GUI to Assign a Recorder Interface to a Policy.....	228
6.9 Using the GUI to Define a Recorder Query.....	231
6.10 Viewing Query History.....	233
6.11 Using the CLI to Manage the DMF Recorder Node.....	234
6.11.1 Basic Configuration.....	234
6.11.2 Authentication Token Configuration.....	236
6.11.3 Configuring the Pre-buffer.....	236
6.11.4 Triggering a Recorder Node Event.....	237
6.11.5 Terminating a Recorder Node Event.....	237
6.11.6 Viewing Recorder Node Events.....	237
6.12 Using the CLI to Run Recorder Node Queries.....	237
6.12.1 Packet Replay.....	237
6.12.2 Packet Data Query.....	238
6.12.3 Packet Object Query.....	238

6.12.4 Size Query.....	239
6.12.5 Window Query.....	239
6.12.6 Stopping a Query.....	239
6.13 Using RBAC to Manage Access to the DMF Recorder Node.....	240
6.13.1 RBAC Configuration Using the CLI.....	240
6.13.2 RBAC Configuration Using the GUI.....	240
6.14 Using the CLI to View Information About a Recorder Node.....	243
6.14.1 Viewing the Recorder Node Interface.....	243
6.14.2 Viewing Recorder Node Operation.....	243
6.14.3 Viewing Errors and Warnings.....	244
6.15 Using the GUI to View Recorder Node Statistics.....	245
6.16 Changing the Recorder Node Default Configuration.....	251
6.17 Large PCAP Queries.....	252
6.18 Recorder Node Management Migration L3ZTN.....	253
6.19 Recorder Node CLI.....	254
6.20 Multiple Queries.....	254
6.21 Ability to Deduplicate Packets - Query from Recorder Node.....	256
6.21.1 Using the CLI to Deduplicate Packets.....	256
6.21.2 Using the GUI to Deduplicate Packets.....	257
6.21.3 Limitations.....	258
6.22 Enabling Egress sFlow® on Recorder Node Interfaces.....	259
6.22.1 Using the CLI to Enable Egress sFlow.....	259
6.22.2 Using the GUI to Enable Egress sFlow.....	259
6.22.3 Troubleshooting Egress sFlow Configurations.....	261
6.22.4 Guidelines and Limitations for Enabling Egress sFlow.....	262
6.23 Recorder Node Recording State API.....	262
Chapter 7: Link Aggregation.....	266
7.1 Configuring Link Aggregation.....	266
7.1.1 Using the CLI to Configure Link Aggregation Groups.....	266
7.1.2 Using the GUI to Configure Link Aggregation Groups.....	267
7.1.3 Minimum Link and Activation Configuration for LAG Interfaces.....	272
7.1.4 Tunnel Endpoint for vCenter Integration.....	274
7.1.5 Configuring Hashing Fields.....	274
7.1.6 L2 GRE Key Hashing.....	276
7.1.7 VXLAN Hashing.....	279
7.1.8 VXLAN LAG Hashing on DCS7280 Platforms.....	282
7.1.9 MLAG RTAG7 Hash Computation for Unbalanced Load Balancing.....	283
7.2 Pseudo Multi-Chassis Link Aggregation.....	286
7.3 MLAG Components.....	287
7.4 MLAG Limitations.....	288
7.5 Configuring an MLAG via the CLI.....	289
7.6 MLAG Link Discovery.....	291
7.7 Configure MLAG via GUI.....	291
7.8 Create MLAG Policy from GUI.....	293
7.9 Viewing Policy Statistics in the GUI.....	295
7.10 Viewing MLAG Links in the GUI.....	297
7.11 Using LAG Interfaces as Members in MLAG Interfaces.....	297
7.12 Overlapping Policies in LAGs.....	298
Chapter 8: Tunneling Between Data Centers.....	300
8.1 Understanding Tunneling.....	300
8.2 Encapsulation Type.....	301
8.3 Using Tunnels in Policies.....	303

8.4 Using the GUI to Configure a GRE Tunnel.....	303
8.5 Using the CLI to Configure a GRE Tunnel.....	305
8.5.1 Using the CLI to Rate Limit the Packets on a GRE Tunnel.....	306
8.5.2 Using the CLI to View GRE Tunnel Interfaces.....	306
8.6 Using the GUI to Configure a VXLAN Tunnel.....	306
8.7 Using the CLI to Configure a VXLAN Tunnel.....	308
8.7.1 Using the CLI to Rate Limit the Packets on a VXLAN Tunnel.....	309
8.7.2 Using the CLI to View VXLAN Tunnel Interfaces.....	309
8.8 Viewing or Modifying Existing Tunnels.....	309
8.9 Using a Tunnel with User-defined Offsets.....	310
8.10 Wildcard Tunnels on SAND Platforms.....	311
8.10.1 Configuring Wildcard Tunnels Using the CLI.....	311
8.10.2 Configuring Wildcard Tunnels Using the GUI.....	312
Chapter 9: Integrating vCenter with DMF.....	316
9.1 Overview.....	316
9.2 Using SPAN to Monitor VMs.....	317
9.3 Using ERSPAN to Monitor VMs.....	317
9.3.1 Configuration Summary for vCenter Integration.....	318
9.4 Defining a Tunnel Endpoint.....	318
9.5 Using the GUI to Integrate a vCenter Instance.....	320
9.5.1 Using a vCenter Instance as the Traffic Source in a DMF Policy.....	323
9.6 Using the CLI to Integrate a vCenter Instance.....	328
9.6.1 VMs using ERSPAN on a Locally Connected vCenter Instance.....	328
9.6.2 VMs using SPAN on a Locally Connected vCenter Instance.....	330
9.6.3 VMs on a Second Locally Connected vCenter Instance.....	331
9.7 Using the GUI to View vCenter Configuration.....	332
9.8 Using the CLI to View vCenter Configuration.....	333
9.9 Integrating vCenter with DMF using Mirror Stack.....	334
9.9.1 vCenter Configuration.....	334
9.9.2 Configuring DMF.....	339
9.9.3 Troubleshooting.....	341
9.9.4 Limitations.....	341
9.9.5 Resources.....	342
9.10 Wildcard Tunnels for VMware vCenter Monitoring.....	342
9.10.1 Configuration.....	342
9.10.2 Using the CLI to Create Wildcard Tunnels.....	342
9.10.3 Using the GUI to Create Wildcard Tunnels.....	344
9.10.4 Limitations.....	344
9.11 Minimum Permissions for Non-admin Users.....	344
9.12 Monitor VMware vCenter Traffic by VM Names.....	350
9.12.1 Using the CLI to Monitor vCenter Traffic by VM Names.....	351
9.12.2 Using the GUI to Monitor vCenter Traffic by VM Names.....	352
9.12.3 Troubleshooting.....	357
9.12.4 Limitations.....	358
Chapter 10: CloudVision DMF Integration.....	359
10.1 Overview.....	359
10.2 Compatibility Requirements.....	360
10.3 CloudVision DMF Integration using the CLI.....	360
10.4 CloudVision DMF Integration using the GUI.....	365
10.5 Limitations.....	369
10.6 Troubleshooting.....	370
10.7 Resources.....	370

Chapter 11: Advanced Fabric Settings.....	371
11.1 Configuring Advanced Fabric Settings.....	371
11.2 Managing VLAN Tags in the Monitoring Fabric.....	373
11.3 Auto VLAN Mode.....	374
11.3.1 Configuring Auto VLAN Mode using the CLI.....	375
11.3.2 Configuring Auto VLAN Mode using the GUI.....	377
11.4 Auto VLAN Range.....	380
11.5 Auto VLAN Strip.....	382
11.5.1 Auto VLAN Strip using the CLI.....	383
11.5.2 Auto VLAN Strip using the GUI.....	391
11.6 CRC Check.....	393
11.6.1 CRC Check using the CLI.....	393
11.6.2 CRC Check using the GUI.....	394
11.7 Custom Priority.....	395
11.7.1 Configuring Custom Priority using the GUI.....	395
11.7.2 Configuring Custom Priority using the CLI.....	396
11.8 Device Deployment Mode.....	396
11.8.1 Configuring Device Deployment Mode using the GUI.....	397
11.8.2 Configuring Device Deployment Mode using the CLI.....	399
11.9 Inport Mask.....	399
11.9.1 InPort Mask using the CLI.....	399
11.9.2 Inport Mask using the GUI.....	400
11.10 Match Mode.....	401
11.10.1 Setting the Match Mode Using the CLI.....	402
11.10.2 Setting the Match Mode Using the GUI.....	402
11.11 Retain User Policy VLAN.....	403
11.11.1 Retain User Policy VLAN using the CLI.....	403
11.11.2 Retain User Policy VLAN using the GUI.....	404
11.12 Tunneling.....	405
11.12.1 Configuring Tunneling using the GUI.....	406
11.12.2 Configuring Tunneling using the CLI.....	407
11.13 VLAN Preservation.....	407
11.13.1 Using the CLI to Configure VLAN Preservation.....	408
11.13.2 Using the GUI to Configure VLAN Preservation.....	409
11.13.3 Troubleshooting.....	413
11.14 Reuse of Policy VLANs.....	414
11.15 Rewriting the VLAN ID for a Filter Interface.....	415
11.16 Reusing Filter Interface VLAN IDs.....	415
11.17 Using Push-per-filter Mode.....	417
11.18 Tag-based Forwarding.....	417
11.19 Policy Rule Optimization.....	417
11.19.1 Prefix Optimization.....	417
11.19.2 Transport Port Range and VLAN Range Optimization.....	419
11.20 Switch Dual Management Port.....	420
11.20.1 Overview.....	420
11.20.2 Configuring Dual Management Using the CLI.....	420
11.20.3 Configuring Dual Management Using the GUI.....	420
11.20.4 Management Interface Selection Using the GUI.....	424
11.20.5 Management Interface Selection Using the CLI.....	426
11.20.6 Switch Fabric Management Redundancy Status.....	427
11.20.7 Additional Notes.....	427
11.21 Controller Lockdown.....	427
11.22 CPU Queue Stats and Debug Counters.....	428
11.22.1 Troubleshooting.....	430

11.22.2 Limitations.....	430
11.23 Egress Filtering.....	430
11.23.1 Configuring Egress Filtering using the CLI.....	431
11.23.2 Configuring Egress Filtering using the GUI.....	434
11.23.3 Syslog Messages.....	437
11.23.4 Troubleshooting.....	437
11.23.5 Limitations.....	438
Chapter 12: Advanced Policy Configuration.....	439
12.1 Advanced Match Rules.....	439
12.1.1 Match Fields and Criteria.....	439
12.1.2 Match-except Rules.....	440
12.1.3 Matching with IPv6 Addresses.....	441
12.1.4 Port and VLAN Range Matches.....	442
12.1.5 User Defined Filters.....	443
12.2 Using the Filter and Delivery Role with MAC Loopback for a Two-stage Policy.....	444
12.2.1 Using the GUI To Configure a Filter and Delivery Interface with MAC Loopback.....	445
12.2.2 Using the CLI To Configure a Filter and Delivery Interface with MAC Loopback.....	447
12.3 Rate Limiting Traffic to Delivery Interfaces.....	448
12.3.1 Rate Limiting Using the GUI.....	449
12.3.2 Rate Limiting Using the CLI.....	449
12.4 Configuring Overlapping Policies.....	450
12.4.1 Configuring the Policy Overlap Limit Using the GUI.....	452
12.4.2 Configuring the Overlapping Policy Limit Using the CLI.....	453
12.4.3 Using the CLI to View Overlapping Policies.....	454
12.4.4 Configuring the Policy Overlap Limit Strict using the GUI.....	455
12.4.5 Configuring the Policy Overlap Limit Strict using the CLI.....	456
12.4.6 Exclude Inactive Policies from Overlap Limit Calculation.....	456
12.5 Viewing Information about Policies.....	458
12.5.1 Viewing Policy Flows.....	458
12.5.2 Viewing Packets Dropped by Policies.....	458
12.6 Using Rule Groups.....	458
12.6.1 Using the GUI to Configure Rule Groups.....	459
12.6.2 Using the CLI to Configure Interface Groups.....	460
12.7 PTP Timestamping.....	461
12.7.1 Platform Compatibility.....	462
12.7.2 Configuration.....	462
12.7.3 Configuring PTP Timestamping using the CLI.....	462
12.7.4 Configuring PTP Timestamping using the GUI.....	463
12.7.5 PTP Interface Configuration.....	466
12.7.6 Policy Configuration for Timestamping.....	467
12.7.7 L2GRE Encapsulation of Packets with Arista Timestamp Headers.....	467
12.7.8 Using the CLI Show Commands.....	467
12.7.9 Configuration Validation Messages.....	470
12.7.10 Troubleshooting.....	470
12.7.11 Limitations.....	470
Appendix A: Stenographer Reference for DMF Recorder Node.....	472
A.1 Stenographer Query Syntax.....	472
A.2 Example Stenographer Queries.....	475
Appendix B: DMF Recorder Node REST API.....	477
B.1 Authentication.....	477

B.1.1 Basic HTTP Authentication.....	477
B.1.2 Authentication with an Authentication Token.....	477
B.2 DMF Recorder Node API Headers.....	479
B.3 DMF Recorder Node REST APIs.....	479
B.3.1 Ready.....	479
B.3.2 Query Window.....	480
B.3.3 Query Size.....	480
B.3.4 Query Application.....	480
B.3.5 Query Packet.....	481
B.3.6 Query Analysis Filter Stenographer Analysis Type.....	481
B.3.7 Query Replay Request Stenographer Real-time Boolean.....	482
B.3.8 Query Replay Request Filter Stenographer Integer.....	482
B.3.9 Query Replay Done.....	482
B.3.10 Erase Packet Filter Stenographer Query String.....	483
B.3.11 Event Update Trigger Boolean Name Integer.....	483
B.3.12 Abort Query.....	484
B.3.13 Abort All Query.....	484
B.3.14 Queries.....	485
B.3.15 Status Query.....	485
B.3.16 Status All.....	485
Appendix C: DMF Controller in Microsoft Azure.....	487
C.1 Configuration.....	487
C.2 Syslog Messages.....	490
C.3 Troubleshooting.....	491
C.4 Limitations.....	491
C.5 Resources.....	491
C.6 Diagrams.....	492
Appendix D: DMF Controllers in Google Cloud VMware Engine.....	493
D.1 Firstboot Parameters.....	498
D.2 Post Installation Verification.....	500
D.3 Troubleshooting.....	501
D.4 Limitations.....	501
Appendix E: Telemetry Collector.....	502
Appendix F: DMF GUI REST API Inspector.....	506
Appendix G: Configuring Third-party Services.....	514
G.1 Services in the DANZ Monitoring Fabric.....	514
G.1.1 Using the GUI to Configure a DMF Unmanaged Service.....	515
G.1.2 Using the CLI to Configure a DMF Unmanaged Service.....	517
G.2 Service Insertion and Chaining in a DMF Policy.....	518
Appendix H: References.....	520
H.1 Related Documents.....	520

Introduction to DMF and Overview

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

DANZ Monitoring Fabric (DMF) is a cloud-first Network Packet Broker (NPB) that 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.

1.1 Out-of-band Monitoring with DANZ Monitoring Fabric

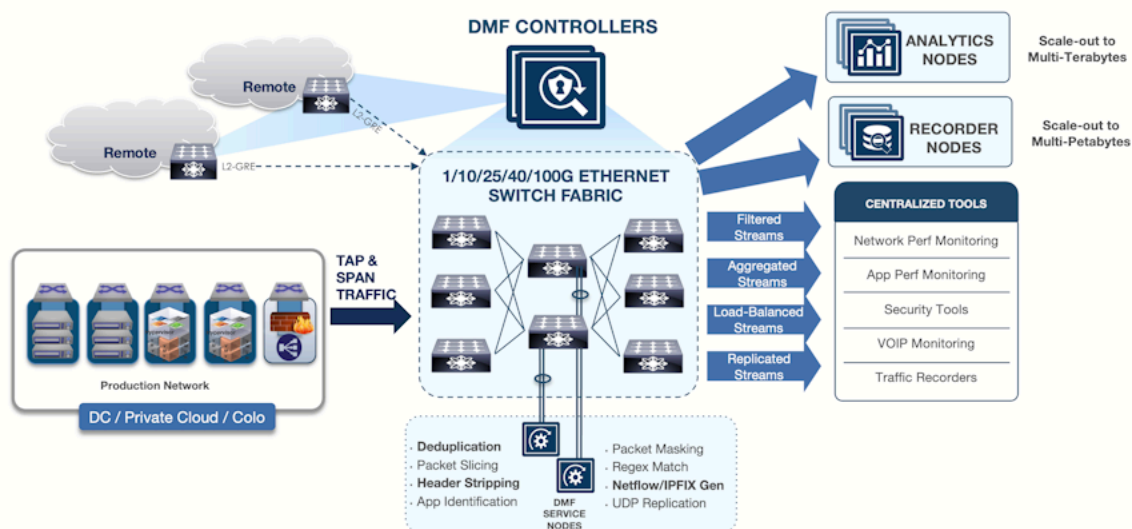
As data center networks move toward 40/100G designs, cloud computing, hyper scale data analytics, and 5G mobile services, traffic monitoring must transition to next-generation designs. To manage the modern data center, much network traffic must be copied and aggregated from TAP or SPAN ports and forwarded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

DANZ Monitoring Fabric (DMF) uses high-performance open-networking switches to deliver an open, production-grade, and scalable monitoring solution based on Software Defined Networking (SDN) technology. 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, inspired by hyper scale networking designs, consists of the following components:

- HA pair of SDN-enabled DMF Controllers (VMs or hardware appliances), enable simplified and centralized configuration, monitoring, and troubleshooting.
- Arista Networks SDN-enabled Switch Light OS is a production-grade, ONIE-deployable, lightweight OS that runs on DMF Ethernet switches.
- Open Ethernet switches (white box/brite box) use the same merchant silicon ASICs used by most incumbent switch vendors and are widely deployed in production data center networks. These switches ship with an 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 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) are x86-based appliances connected to the DMF and are managed via the DMF Controller to provide petabyte packet recording, querying, and replay functions.
- Analytics Nodes (optional) are 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 needing a 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: Achieving a significant reduction in the total cost of ownership by using open Ethernet switches in combination with industry-standard x86 servers, optimized usage of tools, and SDN-enabled operations and automation.

1.2 Using the DANZ Monitoring Fabric CLI

Before connecting to the DANZ Monitoring Fabric (DMF) Controller, ensure 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: Make all fabric switch configuration changes using 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.

Mode and sub-modes divide CLI commands, 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 significantly affect system configuration and 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                                Print remaining arguments
echo                                Enter enable mode
enable                               Exit submode
exit                                Show help
help                                Show commands recently executed
history                              Logout
logout                              Prefix existing commands to delete item
no                                  Send echo messages
ping                                Send echo messages
ping6                               Configure user profile
profile                             Reauthenticate
reauth                              Manage CLI sessions settings
set                                  Generate diagnostic data bundle for technical support
show                                Manage CLI sessions settings
support                             Show documentation on topic
terminal                            Upload diagnostic data bundle for technical support
topic                               Show output of other commands
upload                              Identify the current authenticated account
watch                              Show workflow documentation
whoami
workflow
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 any command's permitted values or keywords, 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 8.6 Deployment Guide***.

1.3 Using the DANZ Monitoring Fabric GUI


The DANZ Monitoring Fabric (DMF) Graphical User Interface (GUI) performs similar operations to the CLI using a graphic user interface instead of text commands and options. The DMF GUI is compatible 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 configured during deployment. Using the VIP ensures that the user connects to the current active Controller, regardless of any failover that may have occurred.

Use the active Controller for all configuration operations and to obtain reliable information when monitoring DMF. The standby Controller is provided only for redundancy if the active Controller becomes unavailable. Do not perform any configuration using the standby Controller, and any 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



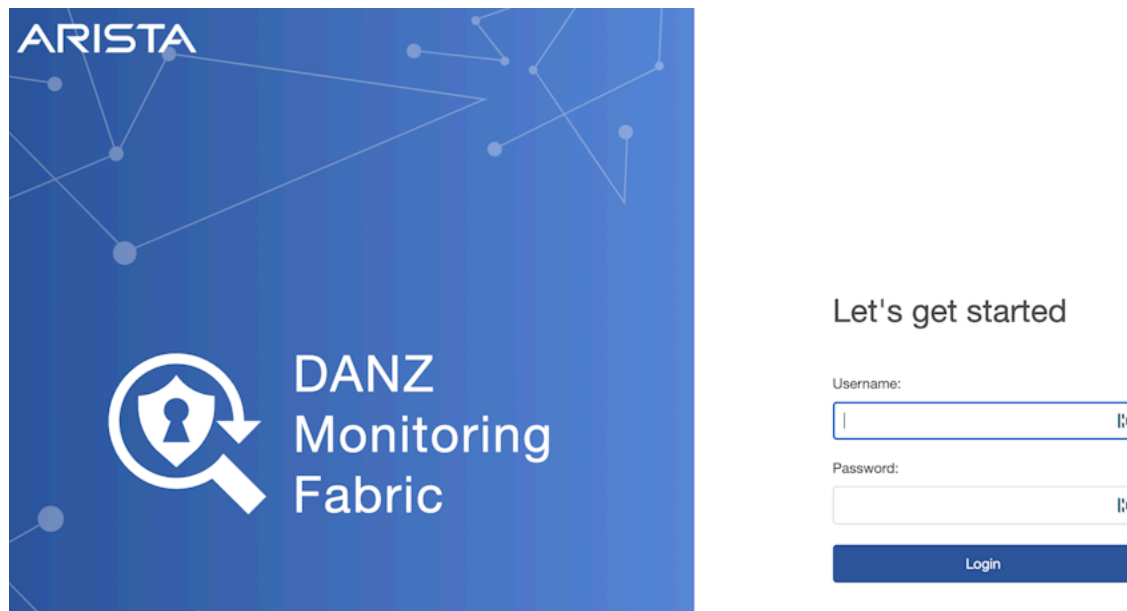
Connecting to the Controller for the first time may result in a security exception prompt (message) because the Controller HTTPS server uses an unknown (self-signed) certificate authority.



Note: When using Internet Explorer, the login attempt may fail if the system time is different than the Controller time. To remedy this, ensure the system used to log in to the Controller is synchronized with the Controller.

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

Figure 1-3: DANZ Monitoring Fabric GUI Login Prompt



Use the admin username and password 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 cannot change the configuration.

Figure 1-4: DANZ Monitoring Fabric GUI Main Menu



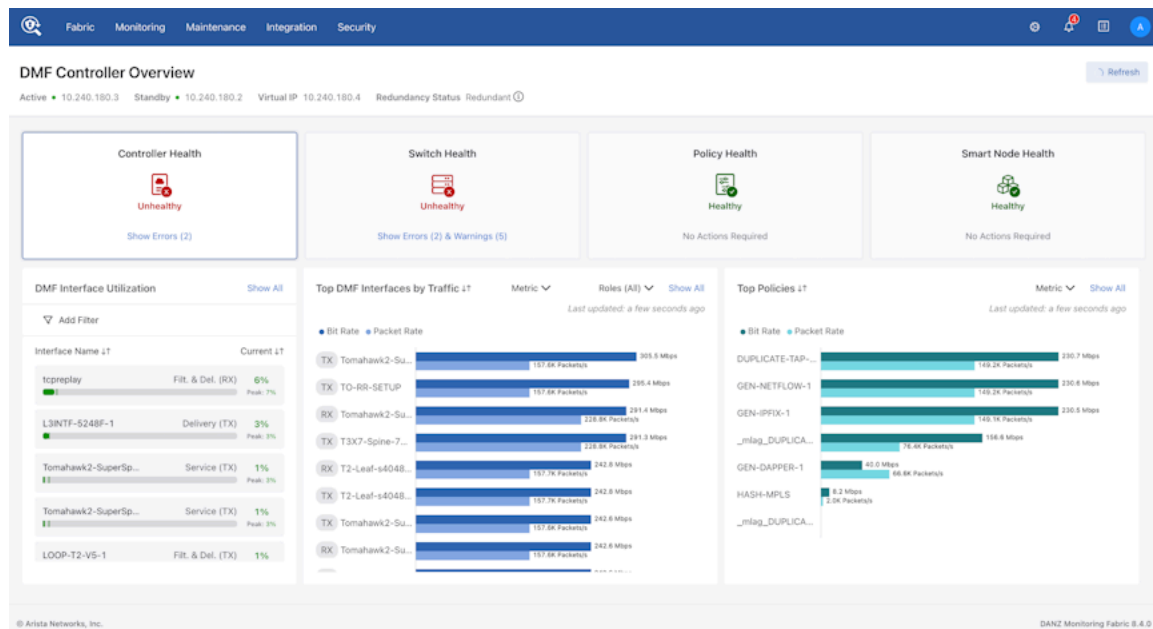
When logging in to the DMF GUI, a landing page appears. This page shows the **DMF Controller Overview**, dashboard, and a menu bar (pictured above) with sub-menus containing options for setting up DMF and monitoring network activity. The menu bar includes the following sub-menus:

- **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 the integration of vCenter instances to allow monitoring traffic using DMF.
- **Security:** manage administrative access.
- A profile page that displays or edits user preferences, the ability to change the password or sign out.

* sFlow[®] is a registered trademark of Inmon Corp.

The newly designed dashboard 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.

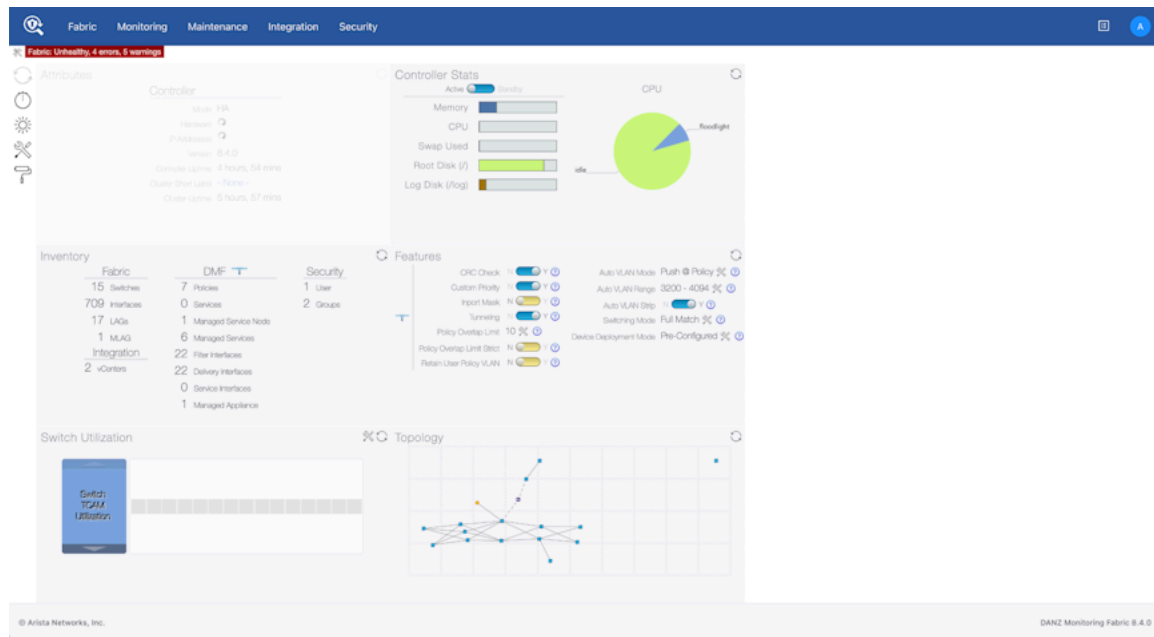
Four cards control the type of content displayed on the main section of the page. The cards are:

- 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



1.3.1 DMF Features Page

To navigate the **DMF Features Page**, click on the **gear icon** in the navigation bar.

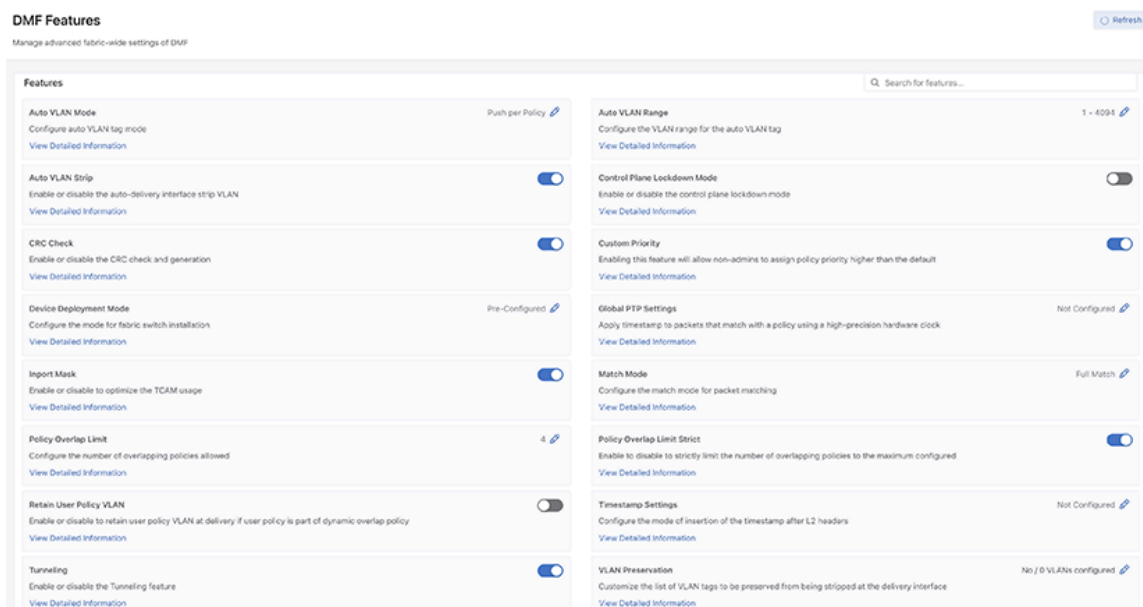
Figure 1-7: Gear Icon



Page Layout

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

Figure 1-8: 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	Control Plane Lockdown Mode
CRC Check	Custom Priority
Device Deployment Mode	Global PTP Settings
Inport Mask	Match Mode
Policy Overlap Limit	Policy Overlap Limit Strict
Retain User Policy VLAN	Timestamp Settings
Tunneling	VLAN Preservation

Each card on the page corresponds to a feature set.

Figure 1-9: Feature Set Card




The UI displays the following:

- Feature Title
- A brief description
- View / Hide Detailed Information
- Current Setting
- Edit Link - Use the **Edit** 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 1-10: View Detailed Information

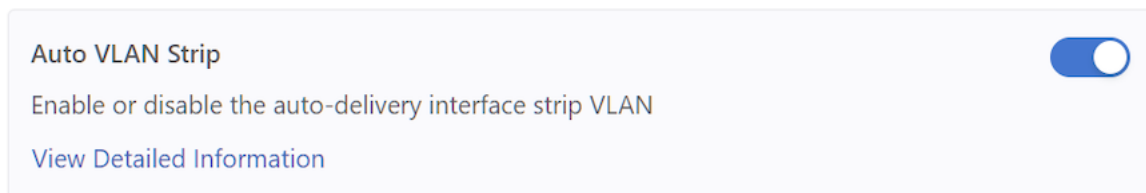


Feature Settings

Auto VLAN Strip

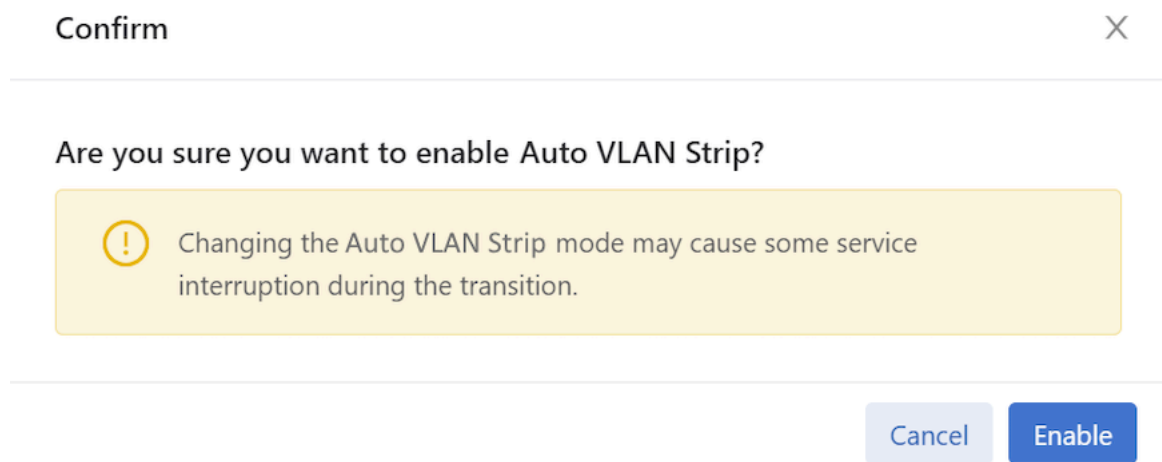
1. A **toggle button** controls the configuration of this feature. Locate the corresponding card and move the **toggle** button.

Figure 1-11: Toggle Switch



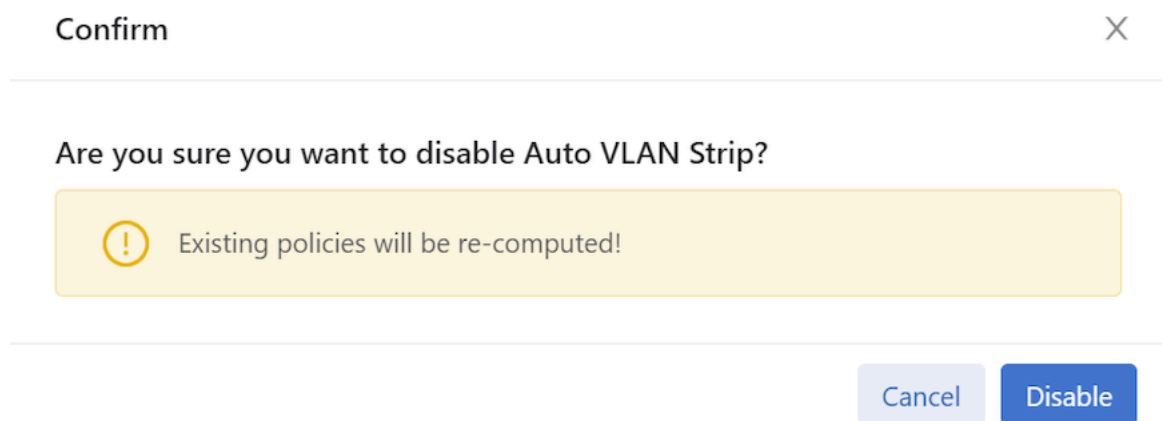
2. A confirm window pops up, displaying the corresponding prompt message. Select the **Enable** button to confirm the configuration changes or the **Cancel** button to cancel the configuration. Conversely, to disable the configuration, select **Disable**.

Figure 1-12: Confirm Enable



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 Auto VLAN Strip?". Below this is a yellow warning box containing a yellow exclamation mark icon and the text: "Changing the Auto VLAN Strip mode may cause some service interruption during the transition." At the bottom right, there are two buttons: "Cancel" (light blue) and "Enable" (dark blue).

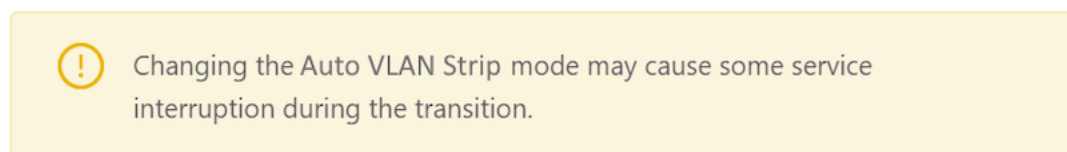
Figure 1-13: Confirm Disable



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 Auto VLAN Strip?". Below this is a yellow warning box containing a yellow exclamation mark icon and the text: "Existing policies will be re-computed!". At the bottom right, there are two buttons: "Cancel" (light blue) and "Disable" (dark blue).

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

Figure 1-14: Warning Message - Changing



A yellow warning box containing a yellow exclamation mark icon and the text: "Changing the Auto VLAN Strip mode may cause some service interruption during the transition."


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



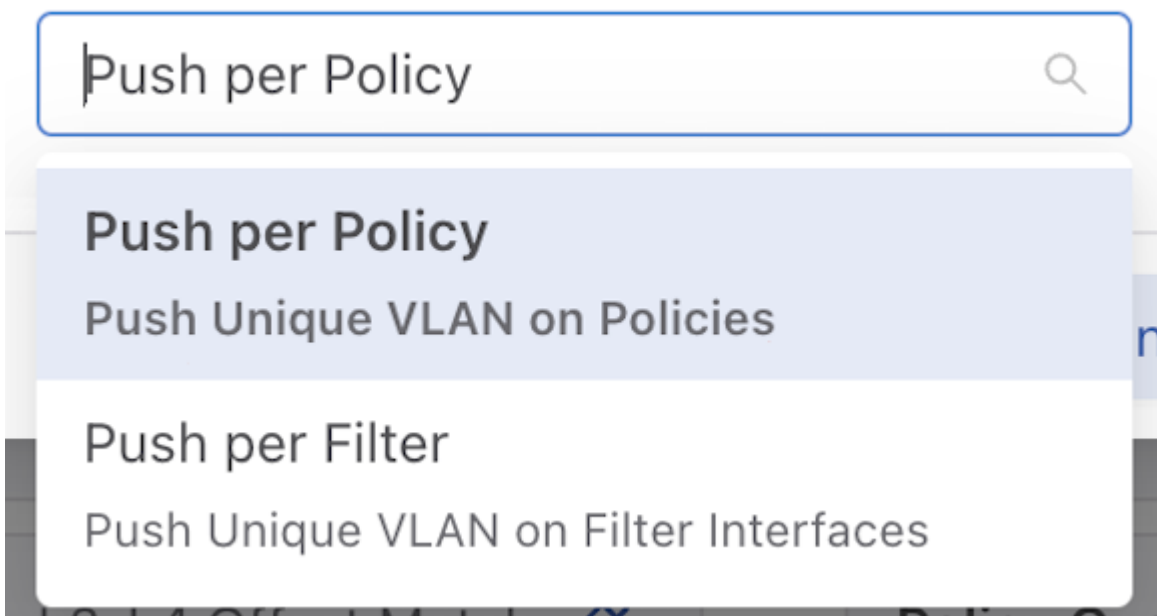
3. To configure different modes, click the drop-down arrow to open the menu.

Figure 1-17: Drop-down Example



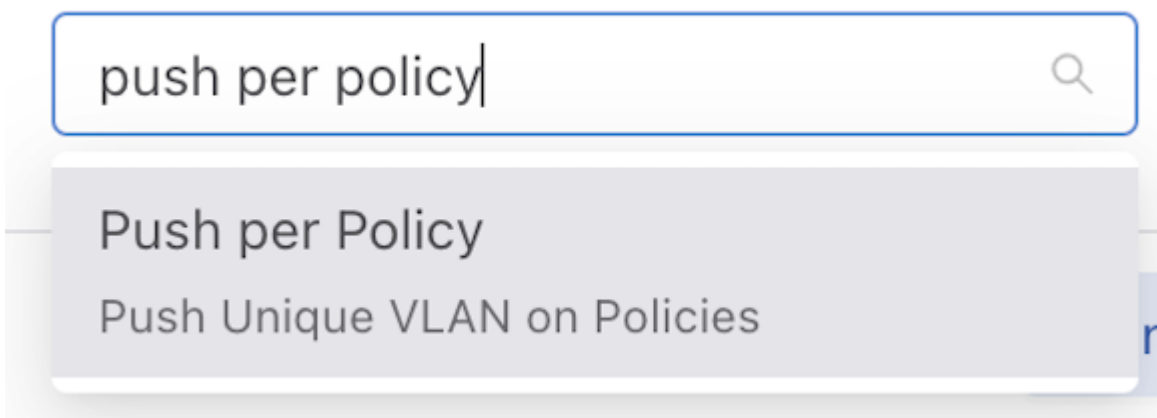
4. From the drop-down menu, select and click on the desired mode.

Figure 1-18: Push Per Policy



5. Alternatively, enter 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

Edit Auto VLAN Mode

! Changing the Auto VLAN Mode will cause policies to be re-installed!

Push per Filter

Cancel Submit

- After successfully setting the configuration, the current configuration status displays next to the edit button.

Figure 1-21: Current Configuration Status

Auto VLAN Mode Push per Policy

Configure auto VLAN tag mode

[View Detailed Information](#)

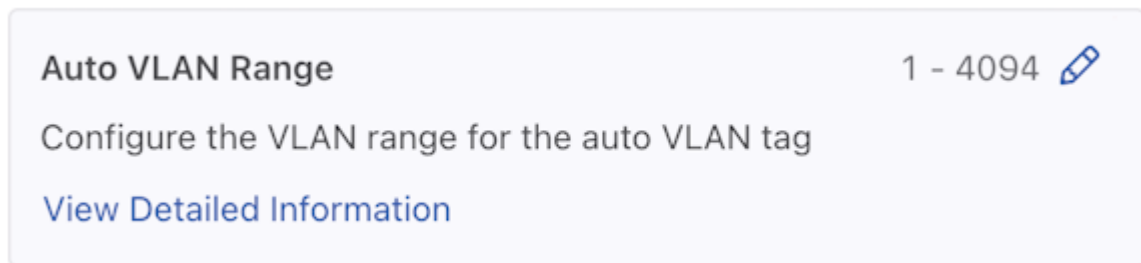
The following feature sets work in the same manner as the **Auto VLAN Mode** feature described above.

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

Auto VLAN Range

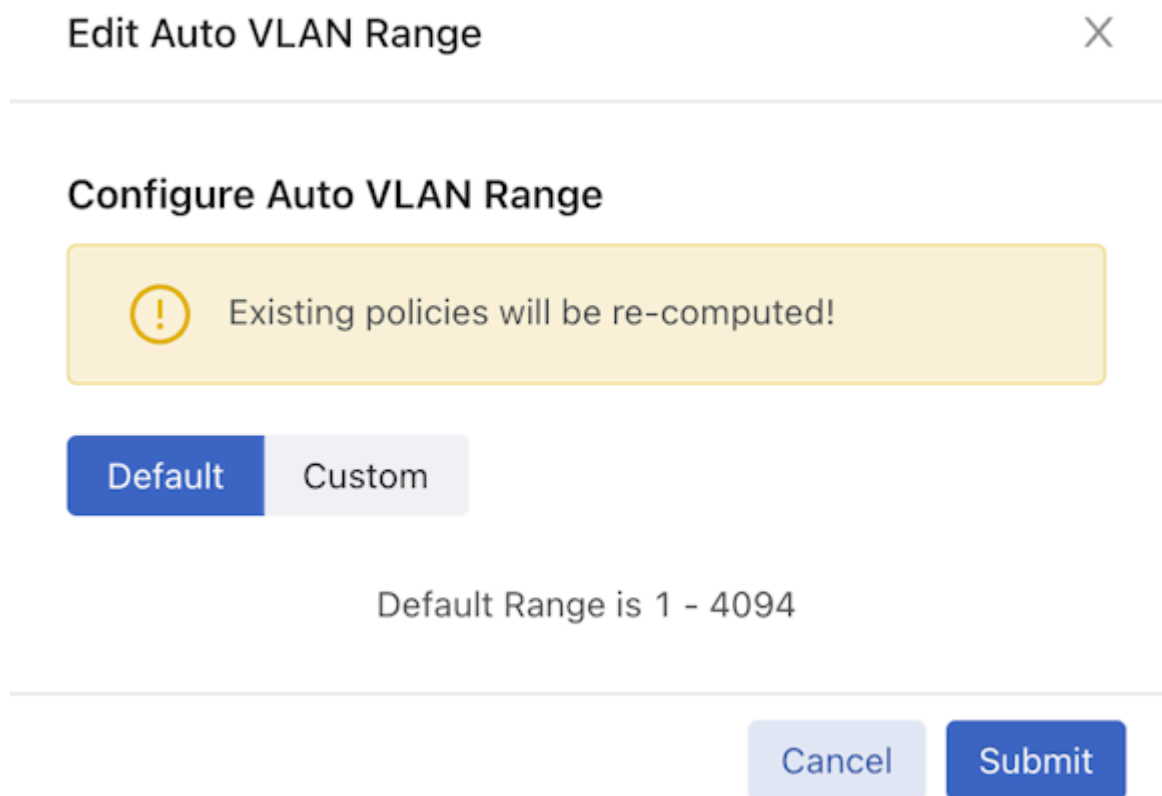
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

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



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

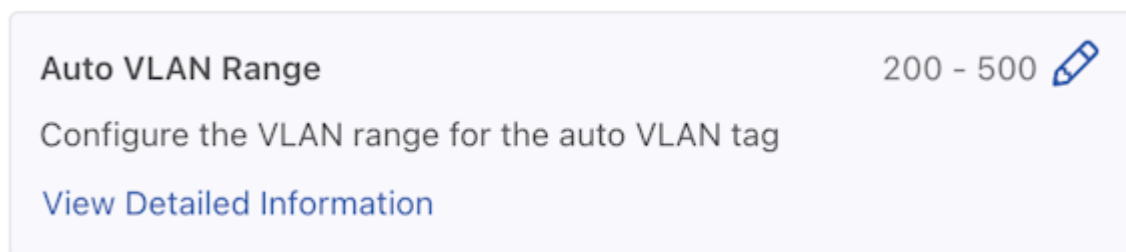
Figure 1-24: Custom Button



The image shows a configuration interface with two tabs: 'Default' and 'Custom'. The 'Custom' tab is selected and highlighted in blue. Below the tabs, there is a range configuration section. It consists of a left input box containing the number '1', a horizontal blue slider with a circular knob in the middle, and a right input box containing the number '4094'.

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

Figure 1-25: Configuration Change Success

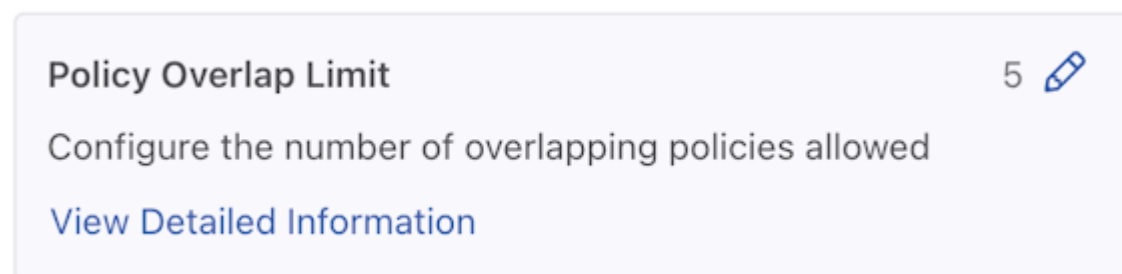


The image shows a light blue confirmation card. At the top, it says 'Auto VLAN Range' on the left and '200 - 500' followed by a pencil icon on the right. Below this, it says 'Configure the VLAN range for the auto VLAN tag'. At the bottom, there is a blue link that says 'View Detailed Information'.

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**.

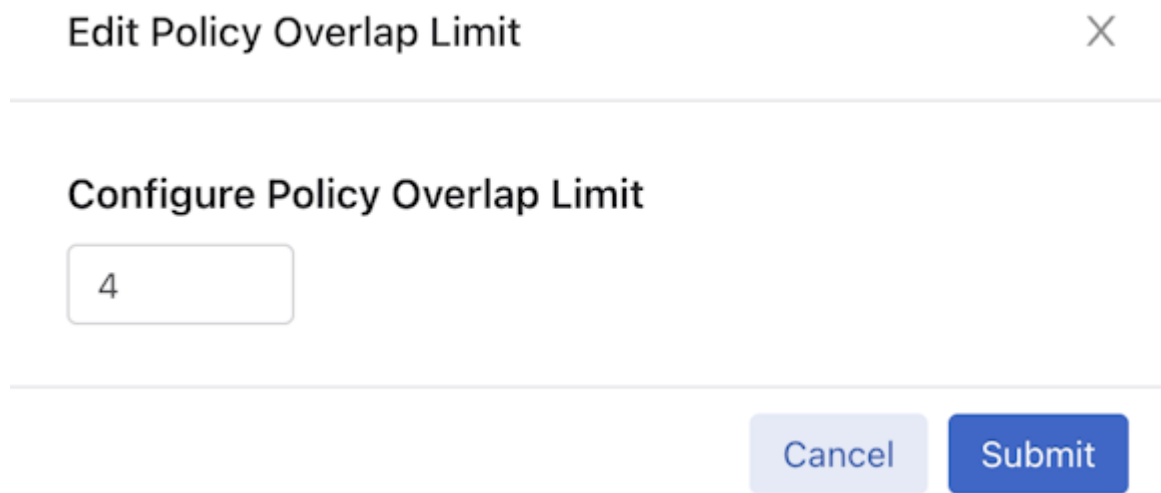
Figure 1-26: Policy Overlap Limit



The image shows a light blue configuration card. At the top, it says 'Policy Overlap Limit' on the left and '5' followed by a pencil icon on the right. Below this, it says 'Configure the number of overlapping policies allowed'. At the bottom, there is a blue link that says 'View Detailed Information'.

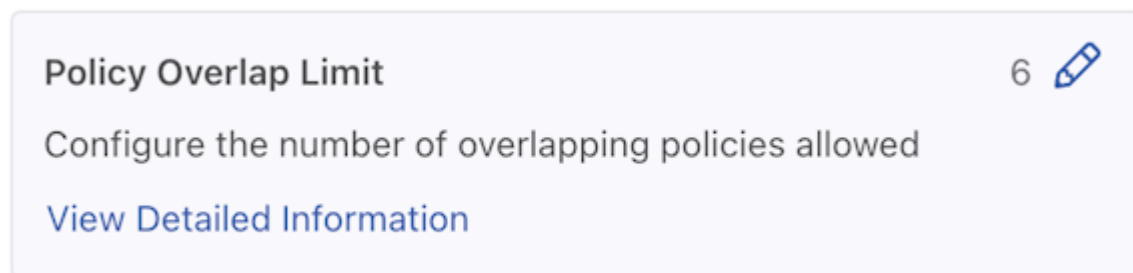
2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. By default, the Policy Overlap Limit is 4.

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.

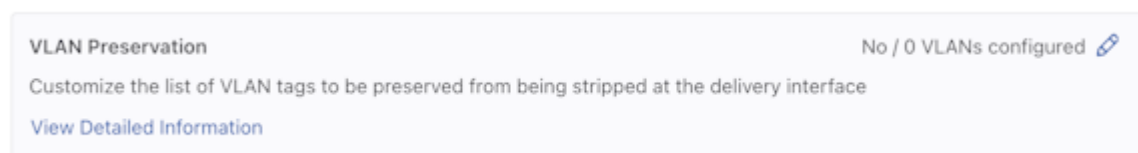
Figure 1-28: Policy Overlap Limit Change Success



VLAN Preservation

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

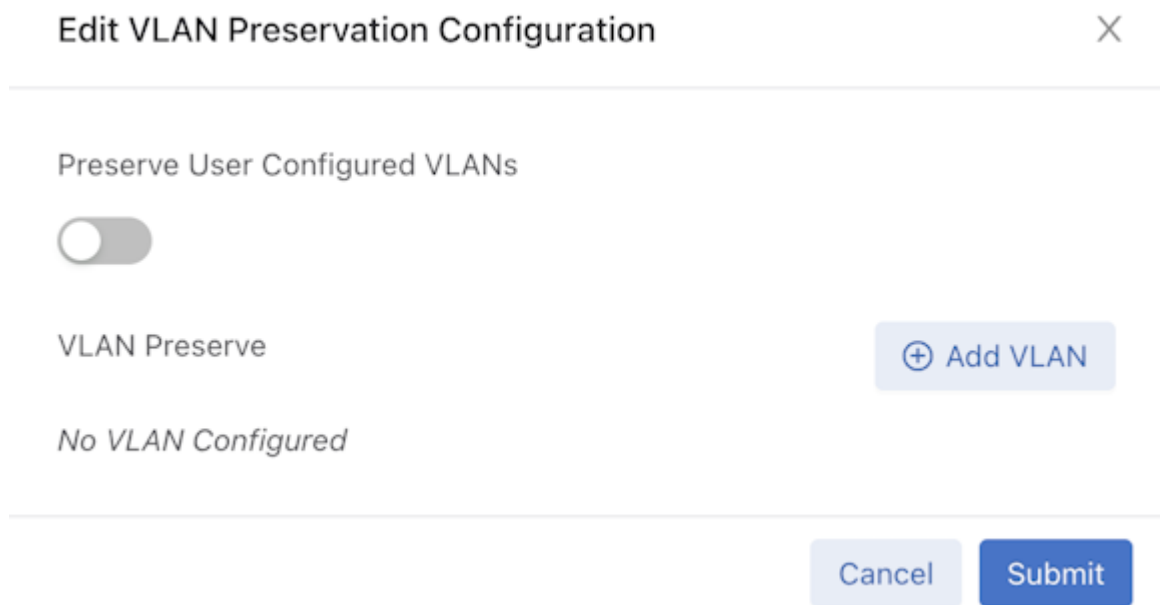
Figure 1-29: VLAN Preservation Feature Set



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 X

Preserve User Configured VLANs

☐

VLAN Preserve + Add VLAN

No VLAN Configured

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

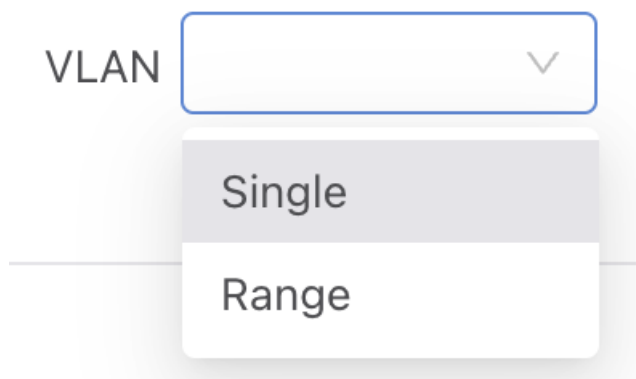


Preserve VLAN X

+ 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 Example X

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

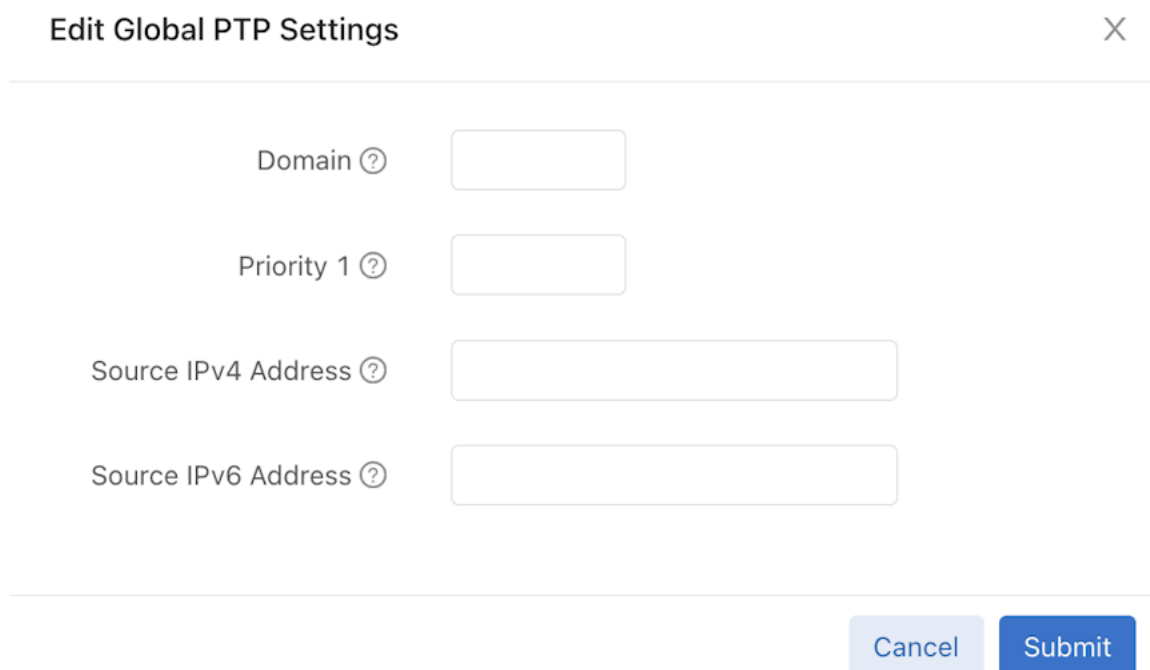
Global PTP Settings

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: Global PTP Settings

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 Settings



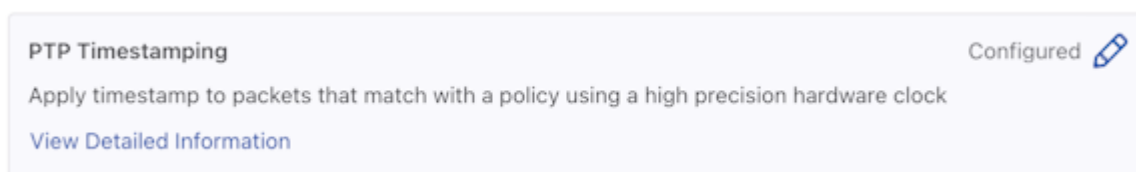
The dialog box is titled "Edit Global PTP Settings" with a close button (X) in the top right corner. It contains four input fields, each with a question mark icon for help:

- 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 card displays the configuration status for "PTP Timestamping". It includes a description, a status indicator, and a link for more details.

PTP Timestamping Configured 

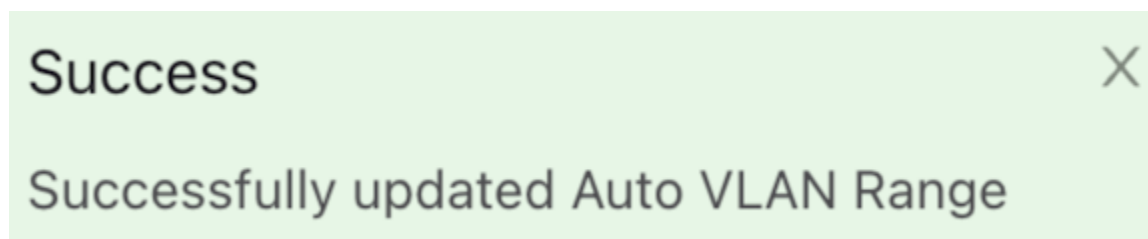
Apply timestamp to packets that match with a policy using a high precision hardware clock

[View Detailed Information](#)

Feature Setting Notification Message

Successfully configuring a feature results in a success notification message pop-up with specific details.

Figure 1-39: Success Message



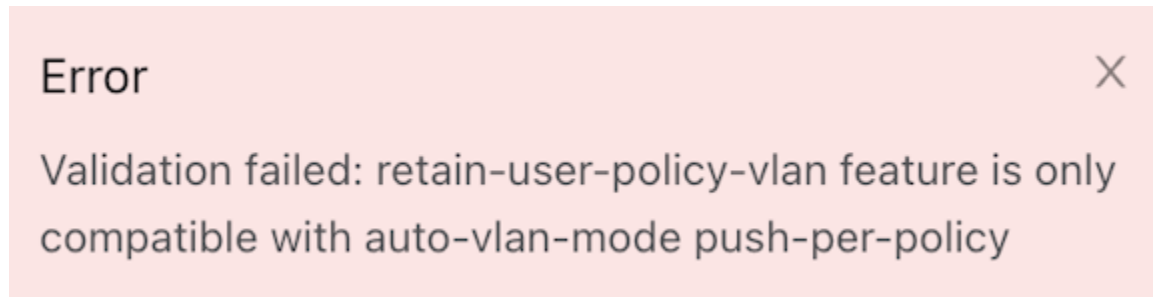
The success message pop-up has a green background and a close button (X) in the top right corner. It displays the following text:

Success

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




Control Plane Lockdown Mode

Enable or disable the Control Plane Lockdown Mode feature.

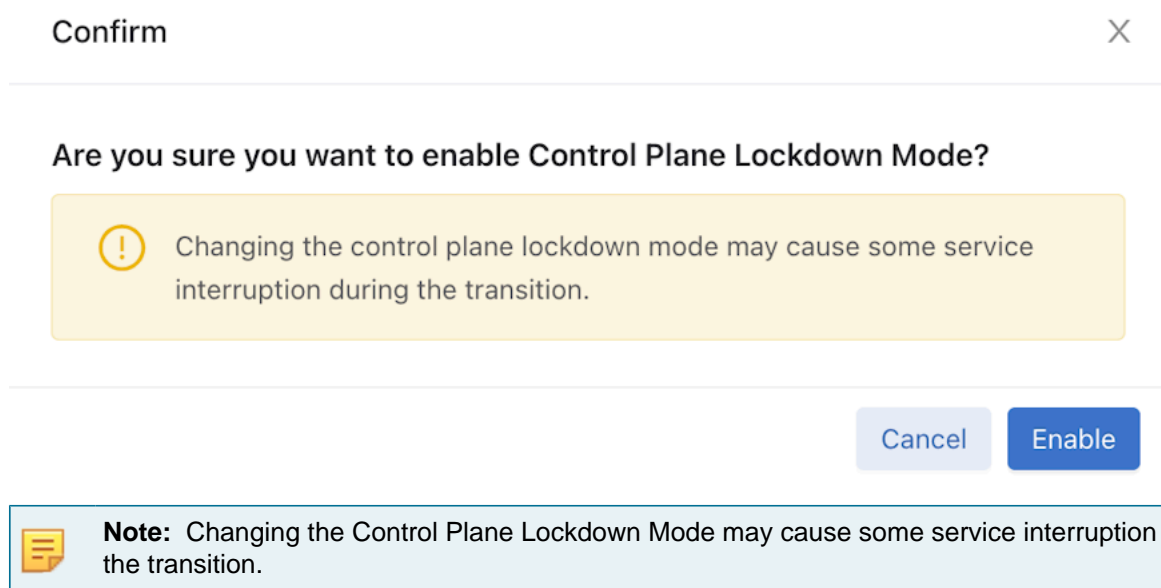
1. A **toggle button** controls the configuration of this feature. Locate the corresponding card and click the **toggle** switch.

Figure 1-41: Control Plane Lockdown Mode




2. Click the **Enable** button to enable Control Plane Lockdown Mode or the **Cancel** button to discard the changes.

Figure 1-42: Enable Control Plane Lockdown Mode



3. On enabling Control Plane Lockdown Mode, a success notification message pops up with specific details.


Figure 1-43: Success Message



Timestamp Settings

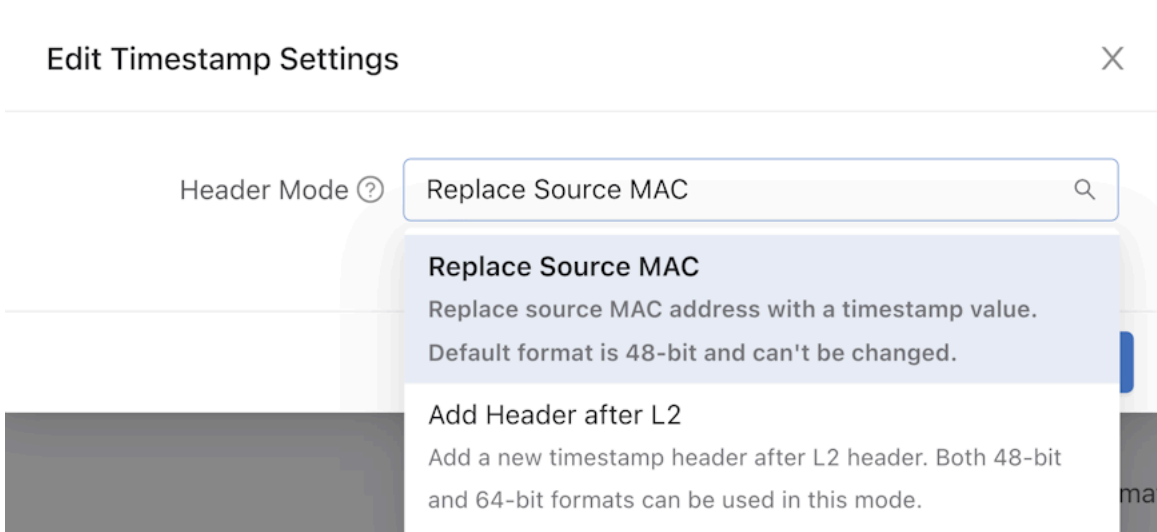
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

Figure 1-44: Timestamp Settings



2. To configure different header modes, click on the drop-down arrow. There are two ways to edit the timestamp settings - **Replace Source MAC** or **Add Header after L2**.

Figure 1-45: Edit Timestamp Settings



- For **Add Header after L2 Mode**, choose the header format as 48-bit or 64-bit.

Figure 1-46: Add Header after L2

Header Mode ⓘ Add Header after L2

* Header Format ⓘ 48-bit

48-bit

64-bit

Cancel Submit

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

Figure 1-47: Submit Timestamp Changes

Header Mode ⓘ Add Header after L2

* Header Format ⓘ 48-bit

Cancel Submit

- Select **Replace Source MAC Mode** and Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.

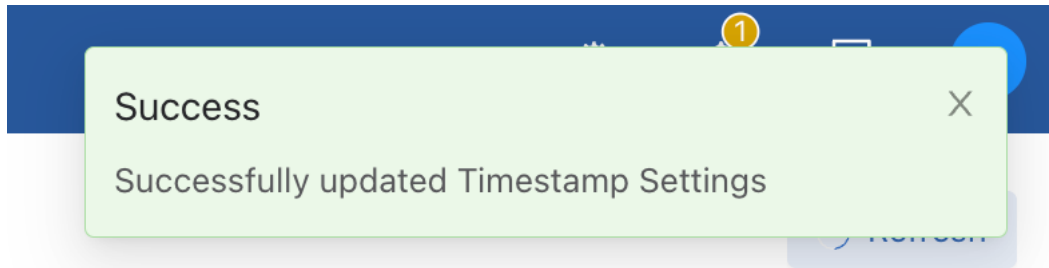
Figure 1-48: Replace Source MAC

Header Mode ⓘ Replace Source MAC

Cancel Submit

-
6. On enabling Timestamp Settings, a success notification message pops up with specific details.

Figure 1-49: Success Message




1.3.2 Dashboard Layout

The dashboard data displays four tabs:

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

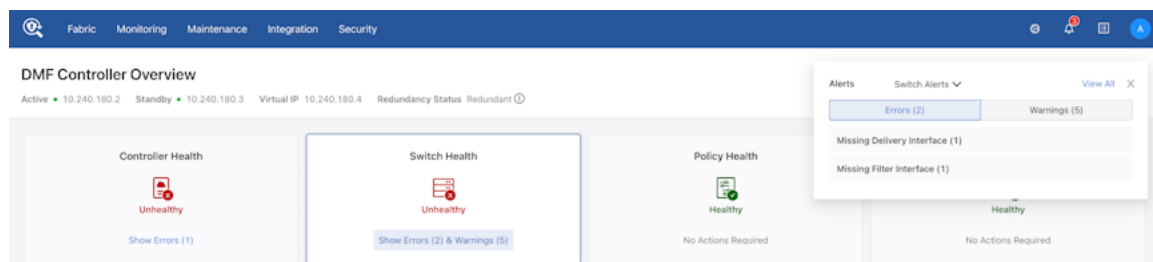
Each tab has health indicators for that category, and accessing the tab displays the relevant data below.

Figure 1-50: DANZ Monitoring Fabric (DMF) Controller Tabs



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

Figure 1-51: 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 1-52: Notification Bell



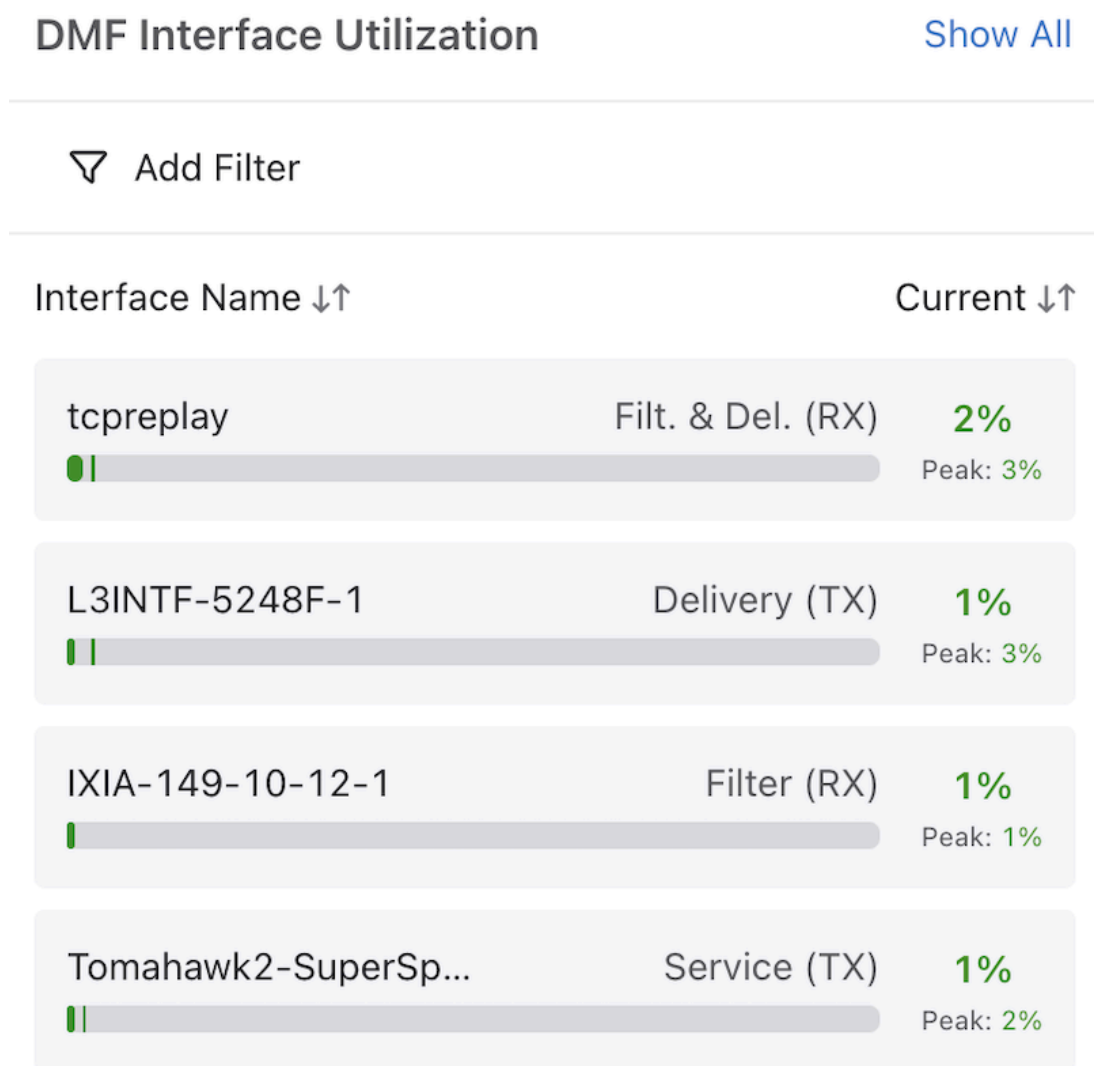
1.3.3 Controller Health

DANZ Monitoring Fabric (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 1-53: 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:

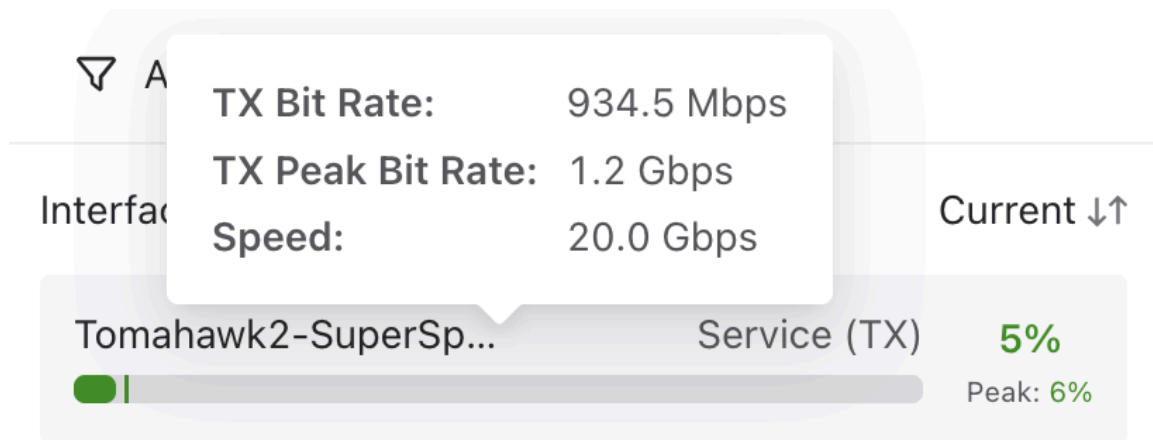
- Red means the utilization percentage is greater than 95%.
- Yellow means the utilization percentage is greater than 70%.
- Green means the utilization percentage is less than 70%.

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 leads to the DMF Interfaces page.

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

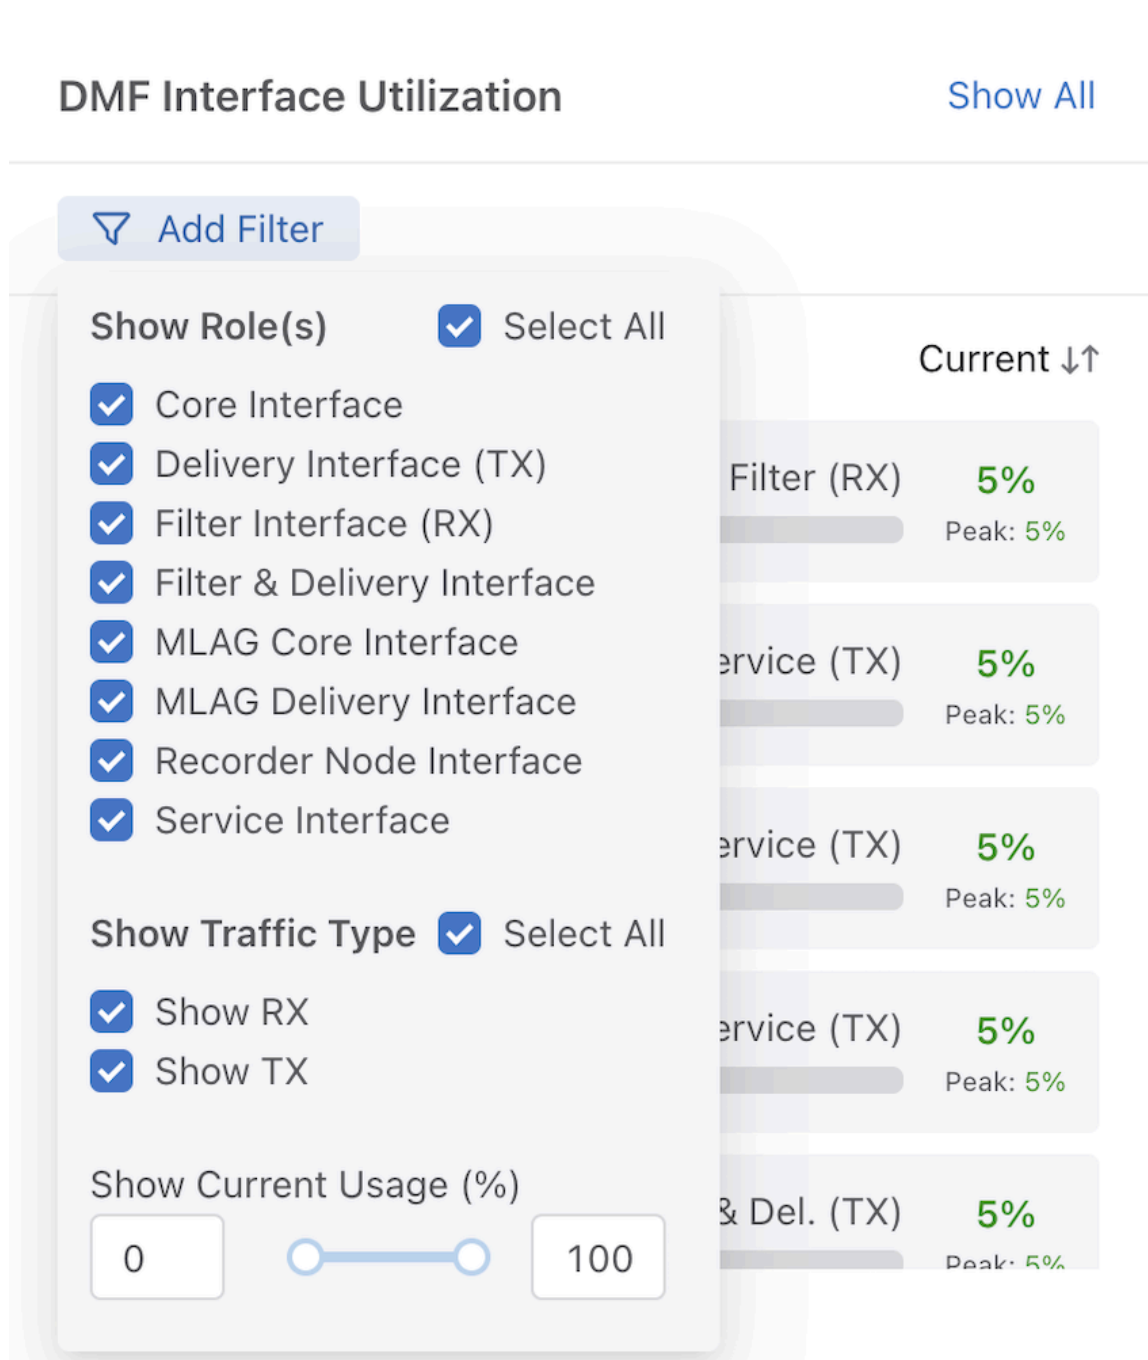
Figure 1-54: DMF Interface Utilization Hover Details



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 1-55: Sort Roles



Top DMF Interfaces by Traffic

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

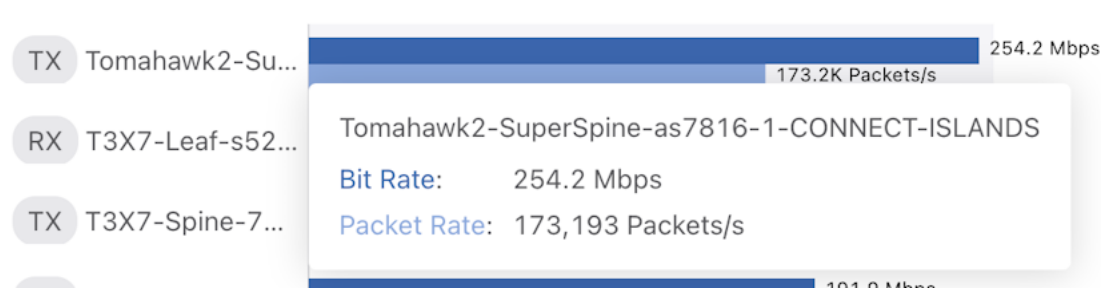
Figure 1-56: Top DMF Interfaces by Traffic



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

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


Figure 1-57: 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 1-58: DMF Top Policies

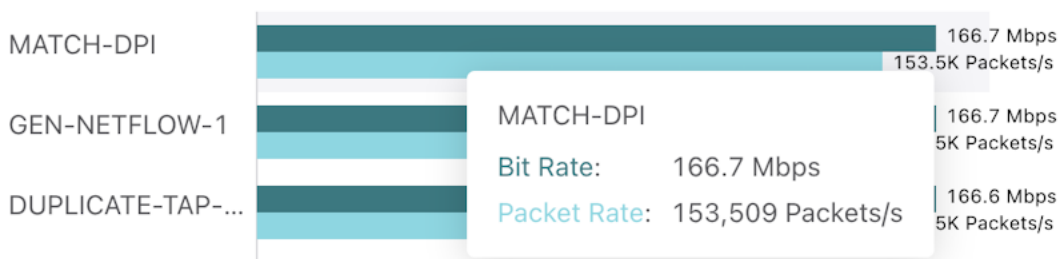


For each policy, the bar chart shows the following:

- 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 1-59: DMF Policies Hover Details



Sort policies by Bit Rate and filter by Metric. By default, DMF sorts the policies in descending order of bit rate.

The **Show All** button leads to the Policies page.

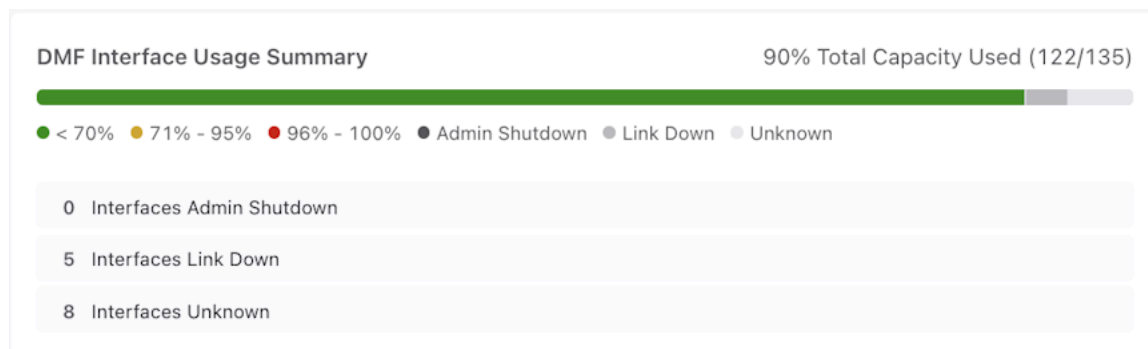
1.3.4 Switch Health

Interface Usage Summary

This widget displays the usage statistics for all DANZ Monitoring Fabric (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 1-60: DMF Interface Usage Summary

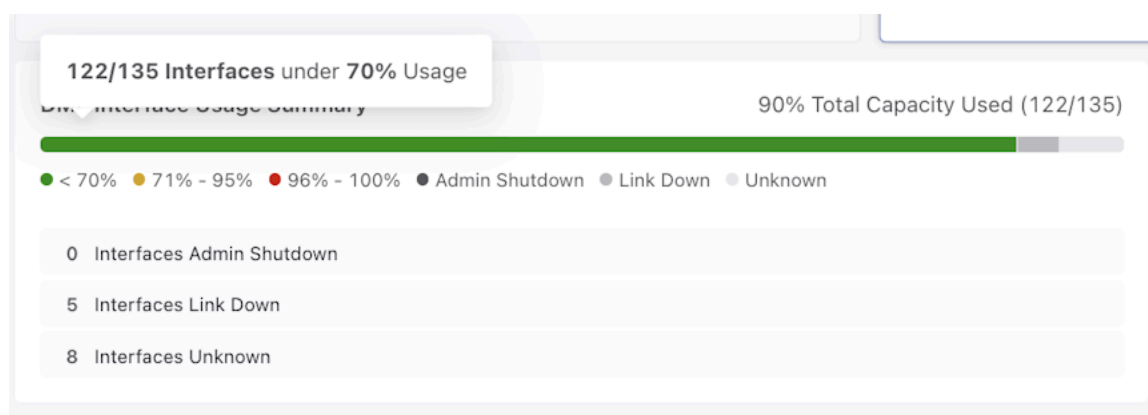


There are three other categories for DMF Interfaces with no traffic. These appear beneath the Usage Bar:

- 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 appears in the respective usage bar.

Figure 1-61: DMF Interface Usage Hover Details

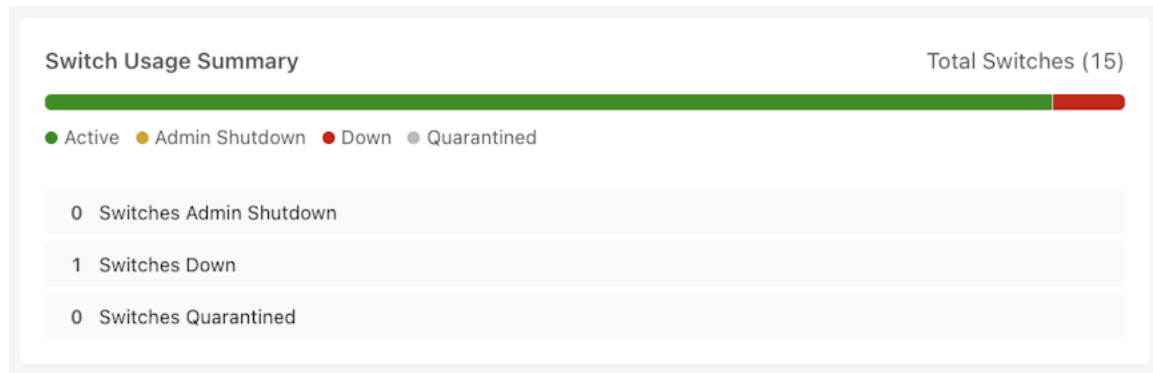


Switch Usage Summary

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

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

Figure 1-62: DMF Switch Usage Summary



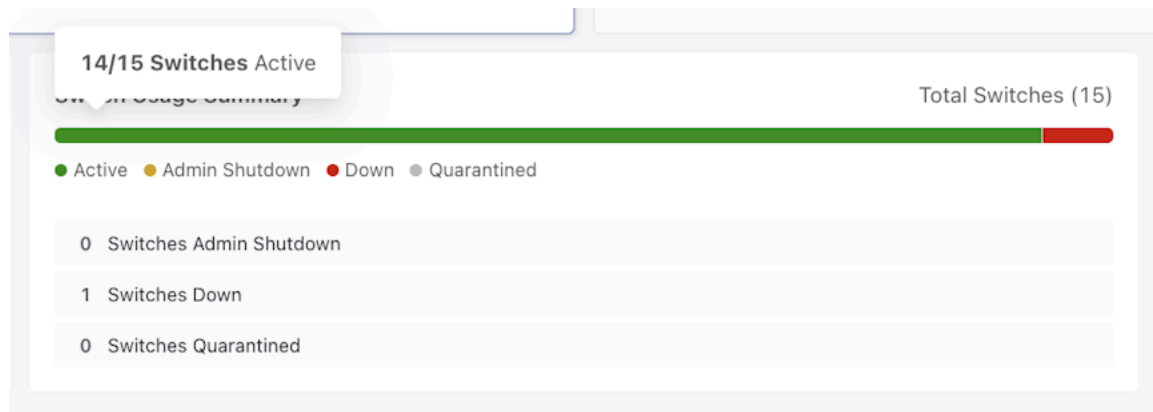
The total number of switches is displayed.

Three list items display the number of:

- Switches Admin Shutdown
- Switches Down
- Switches Quarantined

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

Figure 1-63: DMF Switch Usage Hover Details

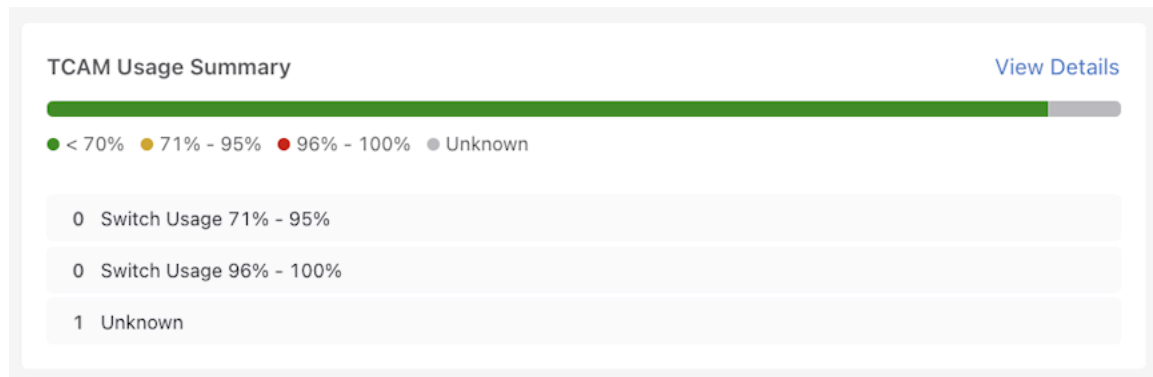


TCAM Usage Summary

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 1-64: DMF TCAM Usage Summary



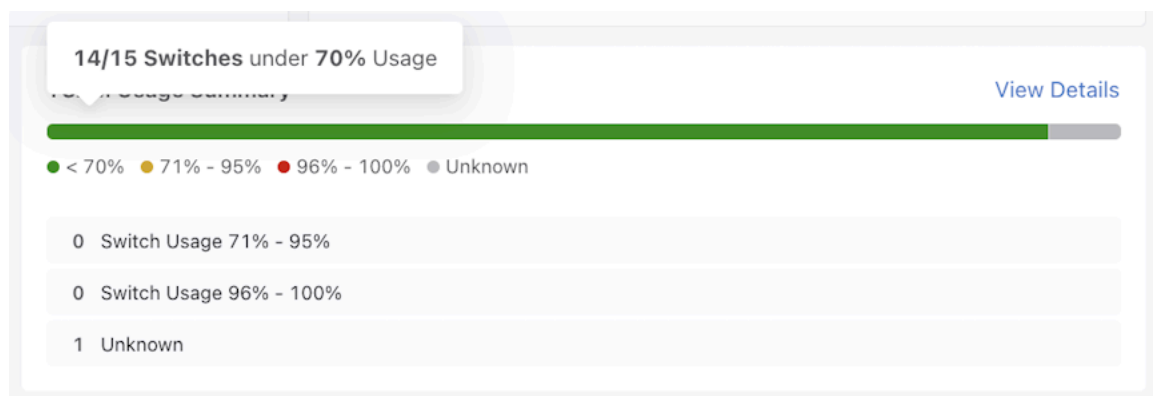
The **View Details** link leads to the TCAM Utilization tab of the Switches page.

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 appears in the respective usage bar.

Figure 1-65: 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

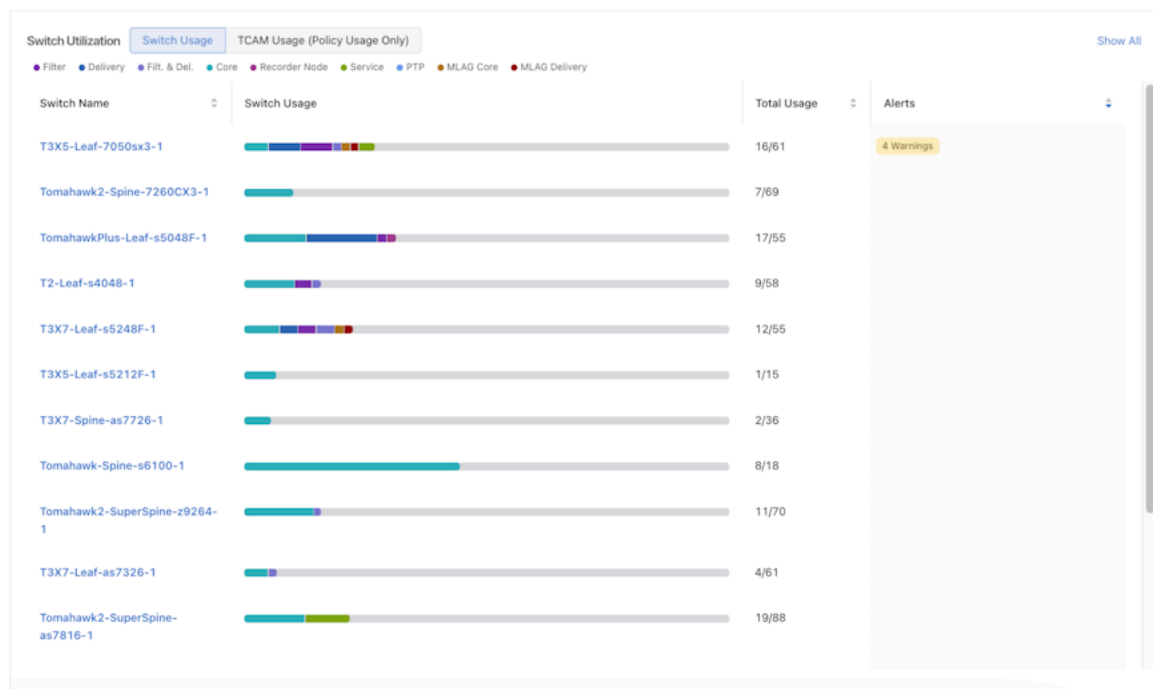
This widget contains two tabs:

- 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 1-66: 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 and 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, which ensures the switches with the highest number of alerts are initially at the top.

On hover, the number of each interface appears.

Figure 1-67: DMF Switch Usage Filter Interfaces Hover Details

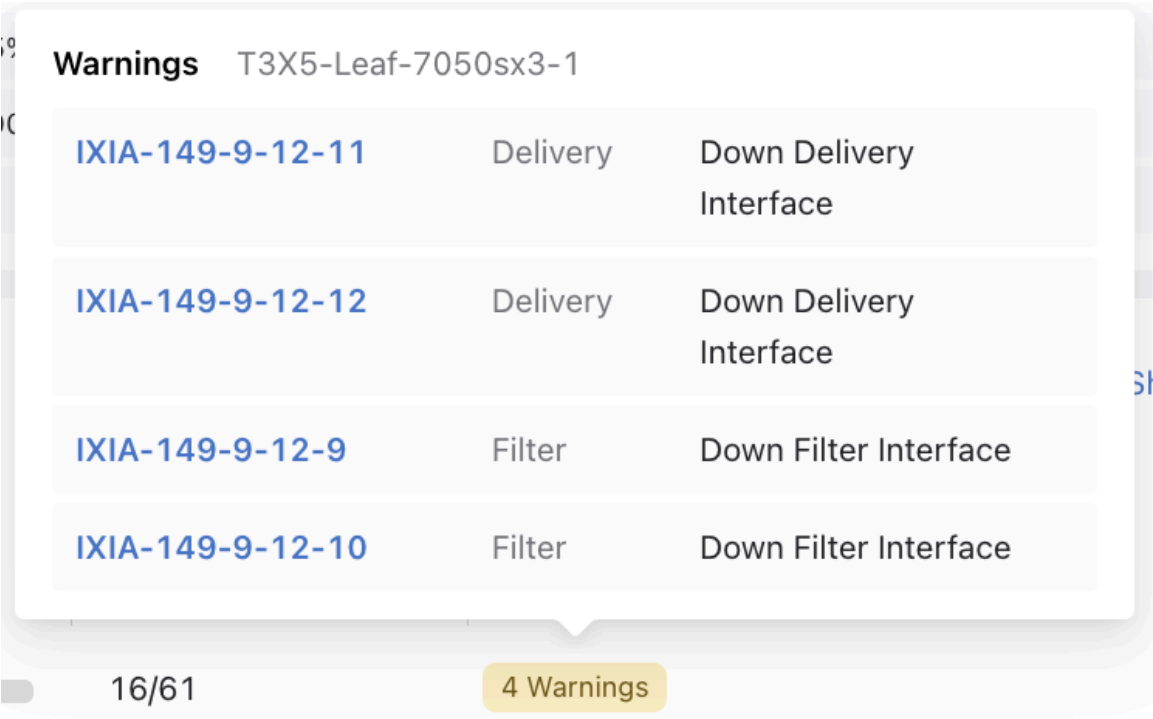


While hovering over the warnings or alerts badge, a table appears and displays Warnings for the yellow badge and Errors for the red badge, and it will also show the switch name.

Each row of the table contains the following:

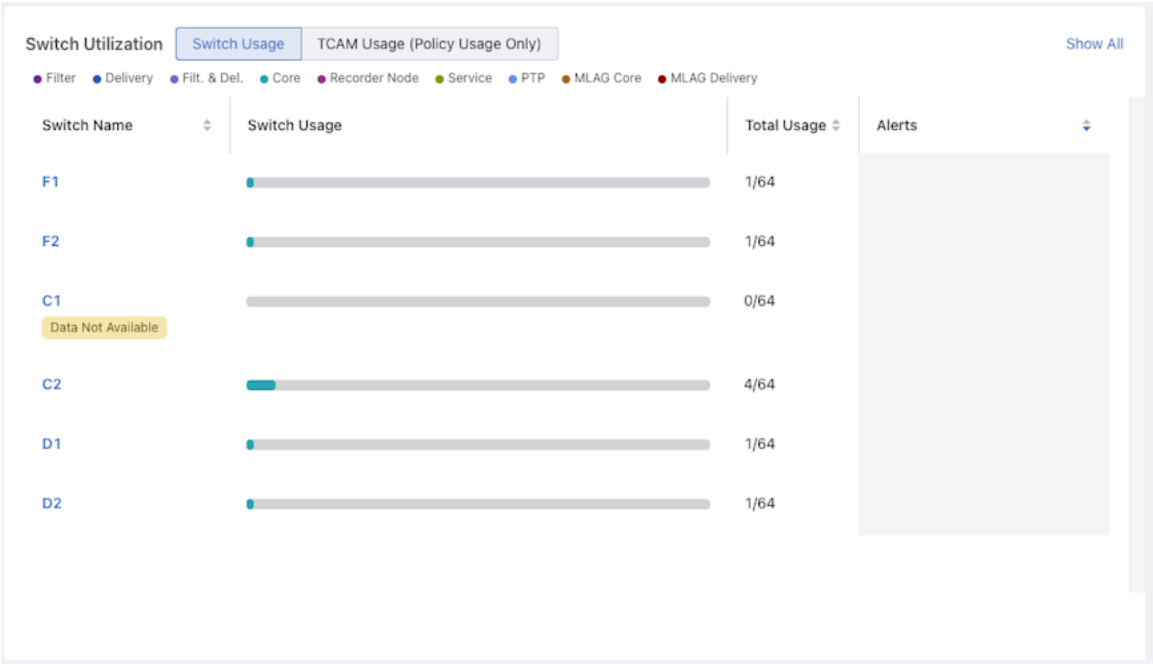
- Interface name (includes a link to Interfaces/[INTERFACE-NAME] page)
- Interface role
- Alert type (e.g., Down Delivery Interface)

Figure 1-68: DMF Interface Warnings



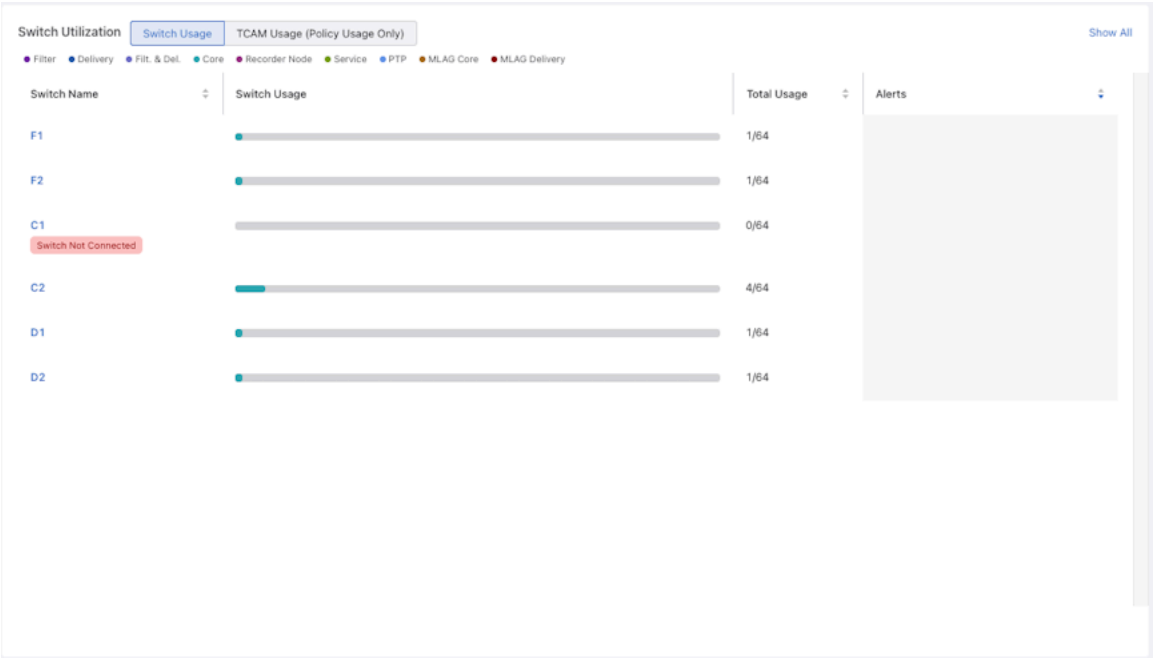
When data is unavailable for a switch (C1), there will be a yellow badge under the Switch Name that says **Data Not Available**. The Switch Usage column will have an empty usage bar, the Total Usage Column will show 0 (zero) for the number of currently used interfaces, and the Alerts column will be empty.

Figure 1-69: 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.

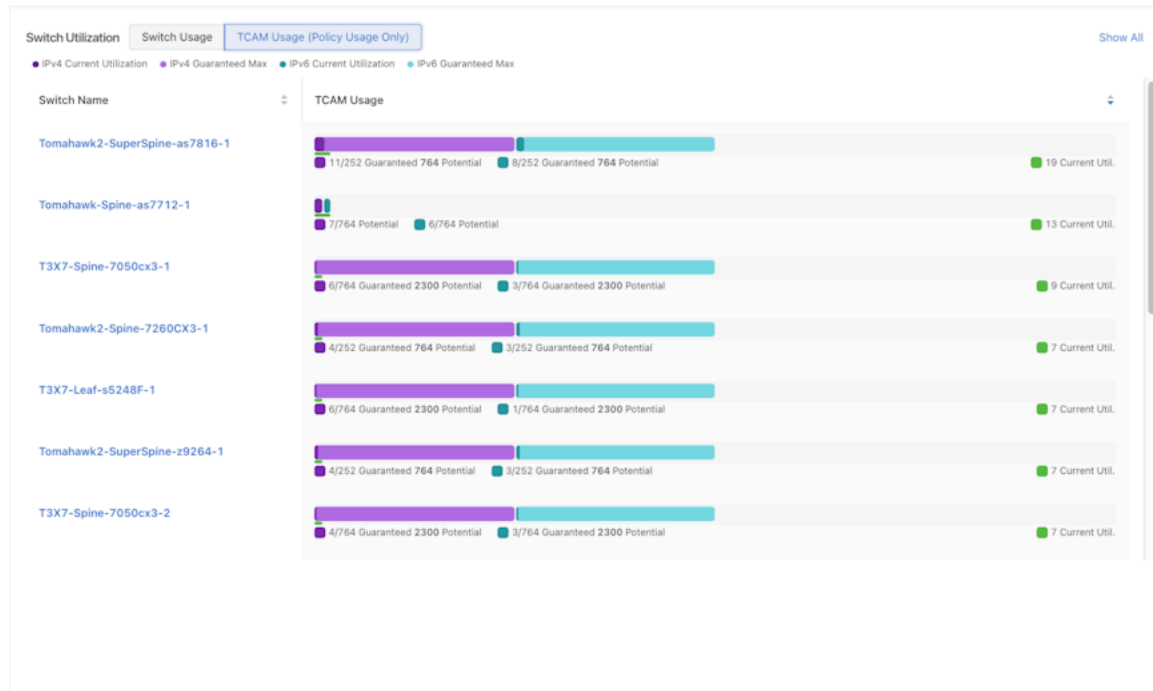
Figure 1-70: DMF Switch Usage - Switch Not Connected



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 1-71: 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 and Delivery
- MLAG Core
- MLAG Delivery
- Recorder Node
- Service

Figure 1-72: 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 leads 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 1-73: DMF Sort Interfaces by Bit Rate




1.3.5 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 1-74: DANZ Monitoring Fabric (DMF) Policies Usage by Traffic



On hover, similar information displays.

Sort policies by Bit Rate.

The **Show All** button leads to the Policies page.

Active Interfaces by Policy

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

Figure 1-75: 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-8-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 leads to the DMF Interfaces page.

Sort the table by each column; DMF sorts the items in descending bit rate order by default.

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

Figure 1-76: DMF Active Interfaces by Policy - Roles

Active Interfaces by Policy

DMF Interface Name	Role	Policy Name(s)	Bit Rate	Packet Rate
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	RX: 129.63K Packets/s TX: 129.63K Packets/s
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 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 1-77: DMF Active Interfaces by Policy - Interfaces

Active Interfaces by Policy

DMF Interface Name	Role	Policy Name(s)	Bit Rate	Packet Rate
_MLAG-DELIVERY-SX3-55248_T3X7-Leaf-s5248F-1_ethernet32	MLag Delivery	DUPLICATE-TAP-FLOWS _mlag_DUPLICATE-TAP-FLOWS_T3X7-Leaf-s5248F-1	357.8 Mbps	cond/s
TO-RR-SETUP	Delivery	DUPLICATE-TAP-FLOWS	209.0 Mbps	
V5-FEED	Filter	DUPLICATE-TAP-FLOWS	188.5 Mbps	
LOOP-TOM2-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS	RX: 188.1 Mb	ets/s
LOOP-T3X5-V5-1	Filter And Delivery	DUPLICATE-TAP-FLOWS GEN-NETFLOW-1	RX: 181.8 Mb	ets/s

Roles (All) Interfaces Show All

Select Interface(s)

Search

☒ V5-FEED
☒ TO-EOS-TASWL-1
☒ LOOP-T2-V5-1
☒ LOOP-T3X7-V5-1
☒ L3INTF-S248F-1
☒ LOOP-TOM2-V5-1
☒ LOOP-TX07-2
☒ LOOP-TOM-V5-1
☒ LOOP-T3X5-V5-1
☒ TO-RR-SETUP
☒ IXIA-149-10-12-1
☒ _MLAG-DELIVERY-SX3-55248_T3X5-Leaf-7050sx3-1_ethernet33
☒ _MLAG-DELIVERY-SX3-55248_T3X7-Leaf-s5248F-1_ethernet32
☒ Tomahawk-7-Gen4Firmware-7816-1-M5-1-85-6-1-in-manual-1-service

1.3.6 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 status of the Recorder Node recording configuration, either **Yes** or **No**.
- 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 1-78: 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
PB-NashRV		...08	Yes	Unknown	Unknown	Unknown			Errors (1)
RMA-Chrial	10.24...	...28	Yes	Index: 12% Packet: 17%	Index: 0% Packet: 0%	Index: Good Packet: Good	RMA-Chrial-intf	OK	

On hover, various details appear 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 detailed information when hovering.

Figure 1-79: Example - Index Disk Storage




Figure 1-80: Example - Index Backup Disk Storage




Figure 1-81: Example - Recorder Node Fabric Interface





Figure 1-82: Example - Errors



The **View All** link leads 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 1-83: DANZ Monitoring Fabric (DMF) Service Nodes



Hovering over the Service Node Interface Load column displays:

- Interface Name
- Service Name
- Action

Figure 1-84: DMF Service Nodes Hover Details

The screenshot shows a web interface with two main sections: Recorder Nodes and Service Nodes. A hover tooltip is displayed over the 'Service Node Interface Load' column of the Service Nodes table.

Recorder Nodes		
Recorder Node Name	IP Address	MAC Address
DMF-RACK258-RN-1	10.2.1.1	08:00:27:00:12:34

Service Nodes	
Service Node Name	IP Address
DMF-SN-R740-1	10.2.1.2

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
sni10	MS-NETFLOW-ACTIVE-1	NetFlow

Legend: Moderate (5) Low (1) Ok

The **View All** link leads to the **Service Node** page.

Analytics Node

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

- IP Address: The configured Analytics Node IP address.
 - 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 1-85: DMF Analytics Node

The screenshot shows the 'Analytics Node' table with a 'View Details' link in the top right corner.

IP Address	Redis Status	Replicated Redis Status
10.2.1.1	Healthy Last updated: 5 seconds ago	Healthy Last updated: 5 seconds ago

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

Refreshing Data

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

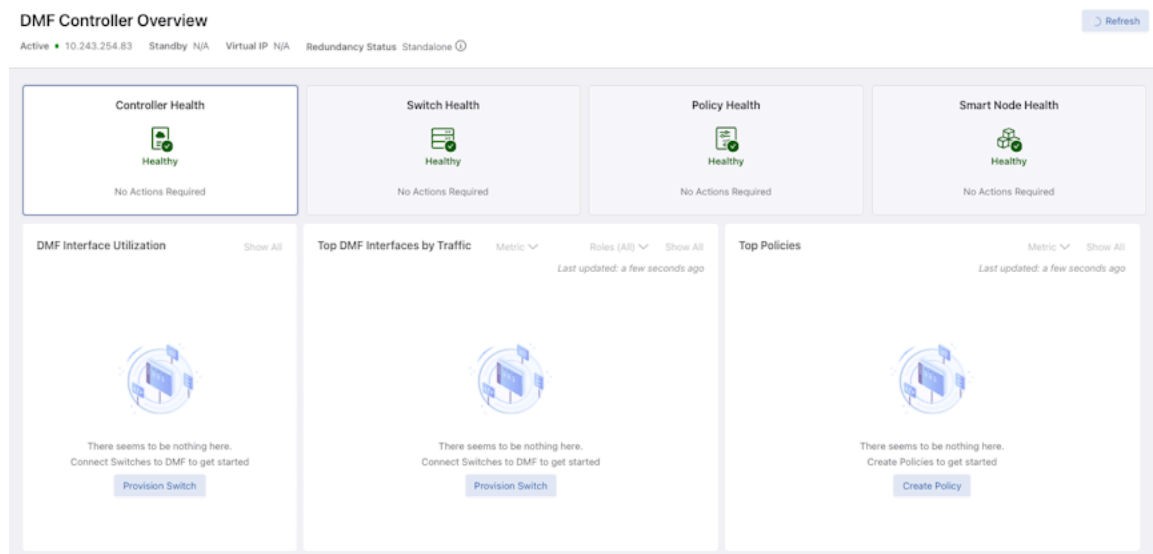
Manually refresh dashboard data using the **Refresh** button.

1.3.7 Empty State

When there are no provisioned switches, DANZ Monitoring Fabric (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.

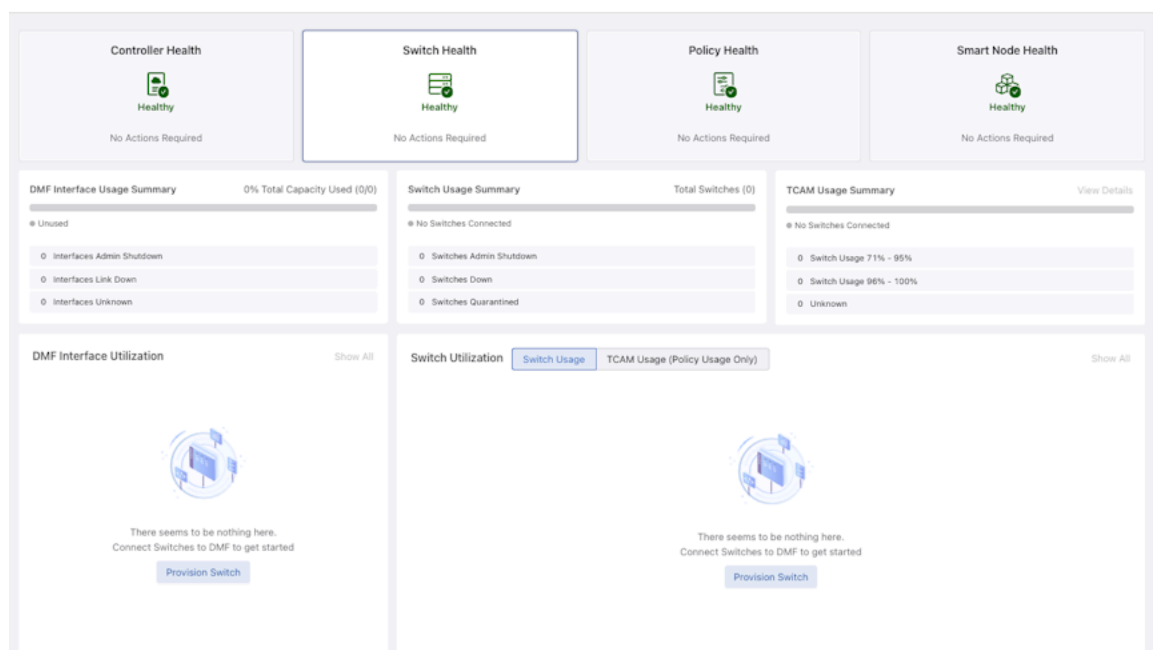
Figure 1-86: DMF Controller Overview - Empty State



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

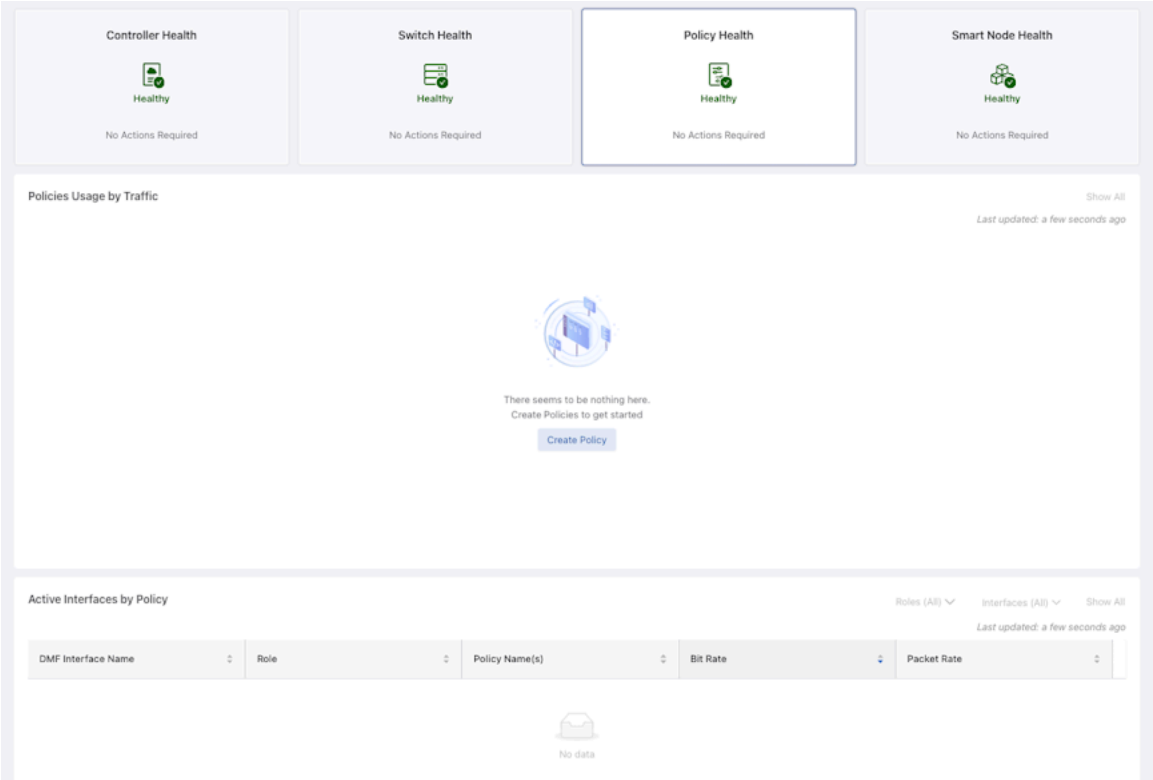
Use the **Create Policy** button to go to the **Create Policy** page.

Figure 1-87: DMF Switch Health - Empty State



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

Figure 1-88: 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 (DMF) switches and interfaces.

2.1 Overriding the Default Configuration for a Switch

By default, each switch inherits its configuration from the DANZ Monitoring Fabric (DMF) Controller. Use the following pages 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 the **switch-name** with the alias previously 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, configure the specific switch and override the default configuration pushed from the DANZ Monitoring Fabric (DMF) Controller to the switch.

The **DANZ Monitoring Fabric 8.6 Deployment Guide** provides detailed instructions on overriding the switch's default configuration..

* sFlow[®] is a registered trademark of Inmon Corp.

2.2 DMF Interfaces

To monitor traffic, 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, monitoring, and compliance tools. These include 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. When placing the port in loopback mode, use a filter and delivery interface to send outgoing traffic back into the switch for further processing. To reduce cost, use a filter and delivery interface when transmit and receive cables are connected to two separate devices.
- **Service interfaces:** interfaces connected to third-party services or network packet brokers, including any interface that sends or receives traffic to or from an NPB.

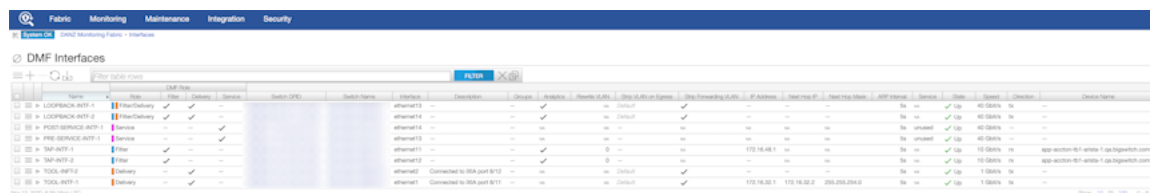
In addition, interfaces connected to managed service nodes and DANZ recorder nodes can be referenced in the configuration directly without assigning a role explicitly. Also, Inter-Switch Links (ISLs), which interconnect DANZ monitoring switches, are automatically detected and referred to as core interfaces.

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

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

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

Figure 2-1: DMF Interfaces



Name	Role	Status	Switch Name	Interface	Description	Direction	Source IP	Destination IP	Source Port	Destination Port	Protocol	Service	Group	Source	Destination	Service Name
LOOPBACK-INT-1	Filter	✓	ethernet1	ethernet1		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
LOOPBACK-INT-2	Filter	✓	ethernet2	ethernet2		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
PROD-SERVICE-INT-1	Service	✓	ethernet3	ethernet3		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
PROD-SERVICE-INT-2	Service	✓	ethernet4	ethernet4		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
TAP-INT-1	Filter	✓	ethernet5	ethernet5		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
TAP-INT-2	Filter	✓	ethernet6	ethernet6		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
TOTAL-INT-1	Delivery	✓	ethernet7	ethernet7	Connected to NPB port 811	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
TOTAL-INT-2	Delivery	✓	ethernet8	ethernet8	Connected to NPB port 811	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

- Click the provision control (+) in 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 (1a)

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 already configured interface.
- 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 the selection.

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

Figure 2-3: Create Interface: Filter

The screenshot shows the 'Create Interface' dialog box with the 'Filter' role selected. The dialog has a title bar with 'Create Interface' and a close button. Below the title bar are two tabs: '1. Interface' and '2. Configure', both with checkmarks. An orange banner at the top right states: 'The following settings may affect the availability of some configuration options.' Below this banner, there are two rows of settings: 'Auto VLAN Mode' set to 'push-per-policy' and 'Strip Forwarding VLANs' set to 'Enabled' with a checkmark. The 'Filter' role is selected with a radio button, and other roles (Delivery, Filter and Delivery, Service) are unselected. The 'Interface Name' field contains 'filter-1'. The 'Rewrite VLAN' section has the text 'Controller will automatically push VLAN.' and a checkmark. The 'Analytics' section has a slider set to 'Enabled'. The 'IP Address' field is empty and has a placeholder '- IPv4 Address -' and a note 'For receiving IP datagram traffic'. At the bottom are buttons for 'BACK', 'NEXT', 'RESET', 'CANCEL', and '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*.
- Optionally, enable the Rewrite VLAN option for a filter or a filter and delivery interface.
- Enable the Rewrite VLAN option when configuring a Filter interface 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, perform the following steps:

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

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



Note: Identify the switch using the alias if configured. The CLI changes to the config-switch submode to configure the specified switch.

- 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. After identifying the interface, the CLI changes to the config-switch-if mode to configure the specified interface.

- From the **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 a single delivery interface as per policy definition. These drops are accounted for as micro bursts occur at the 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 the wire but with a bigger inter-frame gap). This means the ingress may oversubscribe the egress due to the 30G to 25G traffic ratio. For example, if each ingress port receives one packet at the 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 promptly and accumulate inside the egress TX queue. If this pattern continues, the egress queue eventually drops packets due to the TX buffer becoming full. Therefore, expect this behavior in the case of many-to-one forwarding. After reconfiguring and using only one 25G ingress port to one 25G egress port, there is no TX drop problem.



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 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.

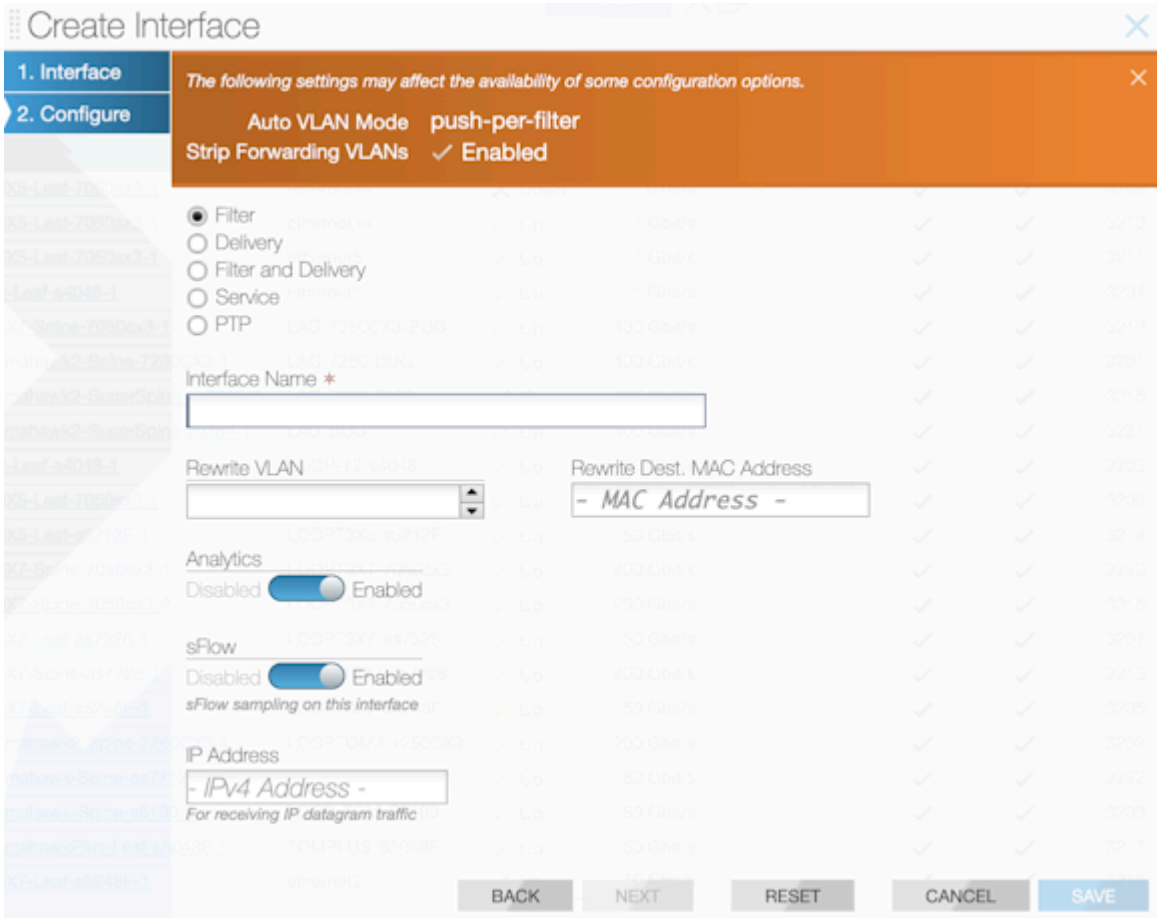
Workflow One

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

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



- 4. Click **Save** to continue.

Edit Interface

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

Figure 2-6: DANZ Monitoring Fabric (DMF) Interfaces - Edit

FabricMonitoringMaintenanceIntegrationSecurity

Fabric: Healthy, 4 warnings

DMF Interfaces

is_filter:true

		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	✓	✓	—	—

Edit ...

Delete ...

Monitor Stats

Clear Stats

Clear Peak Stats

Manage RBAC Group Membership

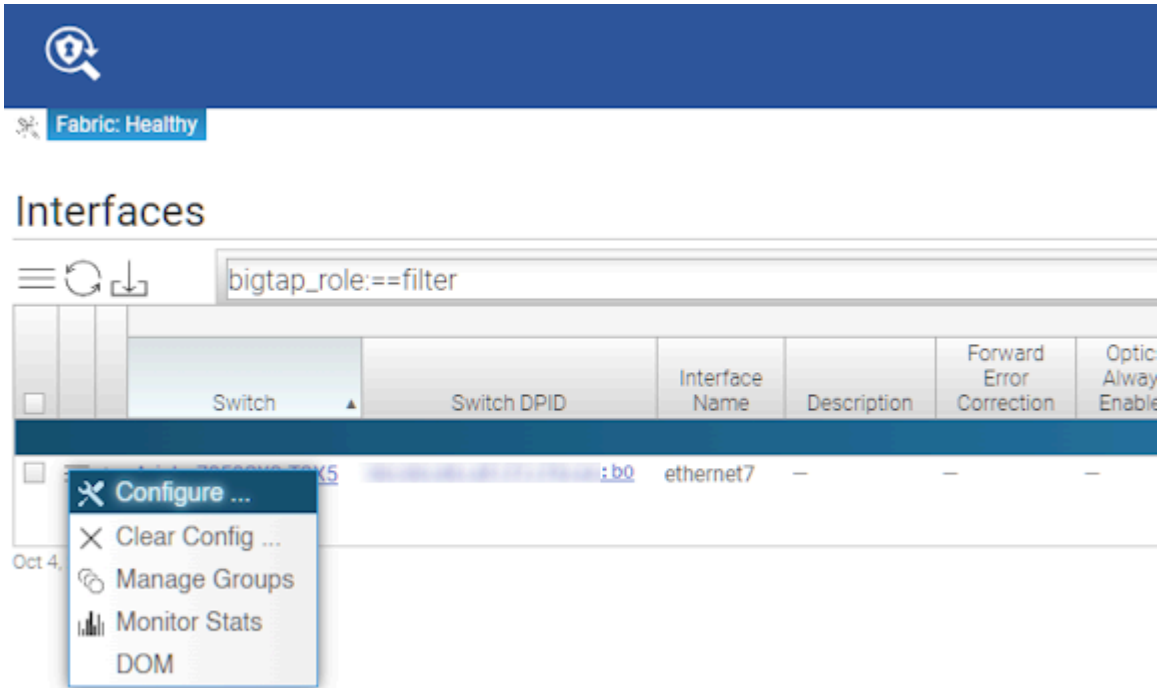
Workflow Two

When using the **Fabric > Interfaces** page, use the following workflow.

6. In the **Configure** step, use the **Rewrite Dest. MAC Address** input to configure the MAC address to override.

- 7. Select the row menu of the switch interface associated with the filter interface 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

Edit Interface

1. Port ✓

2. Traffic ✓

3. DMF ✓

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

Auto VLAN Mode push-per-policy

Auto Strip VLAN ✓ Enabled

Name and role cannot be changed on interfaces used by policies or services

☐ None

☒ Filter

☐ Delivery

☐ Filter and Delivery

☐ Service

☐ PTP

Interface Name *

IXIA-149-9_12-9

Rewrite VLAN

Controller will automatically push 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

9. Click **Save** to continue.

2.2.5 Forward Slashes in Interface Names

DANZ Monitoring Fabric (DMF) supports using forward slashes (/) in interface names to aid in managing interfaces in the DMF fabric. For example, when:

- Defining the SPAN device name and port numbers which generally contain a forward slash (eth2/1/1) in the name for easy port identification.
- Using separate SPAN sessions for Tx and Rx traffic, when there are multiple links from a device to a filter switch.

The following is the comprehensive list of DMF interfaces supporting the use of a forward slash:

- filter interface
- filter interface group
- delivery interface
- delivery interface group
- filter-and-delivery interface
- PTP interface
- managed service interface
- unmanaged service interface
- recorder node interface
- MLAG interface
- LAG interface
- GRE tunnel interface
- VXLAN tunnel interface



Note: An interface name cannot start with a forward slash. However, multiple forward slashes are allowed while adhering to the maximum allowed length limitation.

Configuration

The configuration of DMF filter interfaces remains unchanged. This feature relaxes the existing naming convention by allowing a forward slash to be a part of the name.

The following are several examples:

For switch interfaces, for any of the roles: **both-filter-and-delivery**, **delivery**, **filter**, **ptp**, **service**

```
dmf-controller-1(config-switch-if)# role role interface-name a/b/c
```

For filter and delivery interface groups:

```
dmf-controller-1(config)# filter-interface-group a/b/c
dmf-controller-1(config)# delivery-interface-group a/b/c
```

Adding interfaces or interface groups to a policy:

```
dmf-controller-1(config-policy)# filter-interface f1/a/b
dmf-controller-1(config-policy)# filter-interface-group f/a/b
dmf-controller-1(config-policy)# delivery-interface d1/a/b
dmf-controller-1(config-policy)# delivery-interface-group d/a/b
```

Recorder Node interface:

```
dmf-controller-1(config)# recorder-fabric interface a/b/c
```

For a managed service:

```
dmf-controller-1(config-managed-srv-flow-diff)# l3-delivery-interface a/b/c
```

MLAG interface:

```
dmf-controller-1(config-mlag-domain)# mlag-interface a/b/c  
dmf-controller-1(config-mlag-domain-if)# role delivery interface-name a/b/c
```

LAG interface:

```
dmf-controller-1(config-switch)# lag-interface a/b/c
```

GRE tunnel interface:

```
dmf-controller-1(config-switch)# gre-tunnel-interface a/b/c
```

VXLAN tunnel interface:

```
dmf-controller-1(config-switch)# vxlan-tunnel-interface a/b/c
```

Show Commands

There are no new **show** commands. The existing **show running-config** and **show this** commands for the configurations mentioned earlier should display the interface names without any issue.

2.3 Using Interface Groups

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 identify every interface to which the policy applies explicitly.

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, in an IPv4 address group, the absence of a mask implies a mask of /32. For an IPv6 address group, the absence of a mask implies a mask of /128.

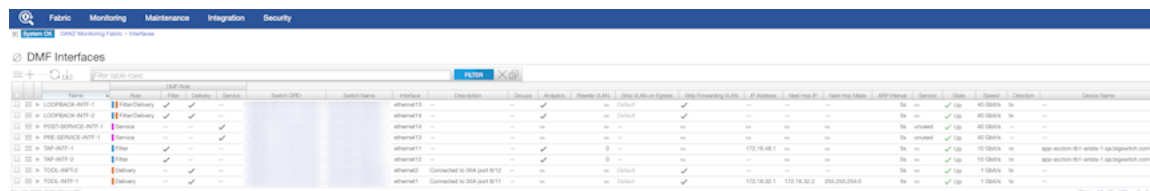
Identify only a single IP address group for a specific policy match rule. Address lists with both **src-ip** and **dst-ip** options cannot exist in the same match rule.

2.3.1 Using the GUI to Configure Interface Groups

To create an interface group from the **Monitoring > Interfaces** table, perform 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 checkboxes 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: Optionally, define an interface group using the **Monitoring > Interface > Groups** UI.

2.3.2 Using the CLI to Configure Interface Groups

The following example illustrates 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)#
```

2.4 Switch Light CLI Operational Commands

As a result of upgrading the Debian distribution to Bookworm, the original Python CLI (based on python2) was removed, as the interaction with the DANZ Monitoring Fabric (DMF) is performed mainly from the Controller.

However, several user operations involve some of the commands used on the switch. These commands are implemented in the new CLI (based on python3) in Switch Light in the Bookworm Debian distribution.

The Zero-Trust Network (ZTN) Security CLI is the default shell when logged into the switch.



Note: The following commands are only available on Dell and Arista Switch platforms running Switch Light OS.

Operational Commands

After connecting to the switch and from the DMF Controller, use the **debug admin** command to enter the switch admin CLI from the ZTN CLI.

Enter the **exit** command to leave the switch admin CLI, as illustrated in the following example.

```
DMF-CONTROLLER# connect switch dmf-sw-7050sx3-1
Switch Light OS SWL-OS-DMF-8.6.x(0), 2024-05-16.08:26-17f56f6
Linux dmf-sw-7050sx3-1 4.19.296-OpenNetworkLinux #1 SMP Thu May 16 08:35:25 UTC
2024 x86_64
Last login: Tue May 21 10:39:05 2024 from 10.240.141.151

Switch Light ZTN Manual Configuration. Type help or ? to list commands.

(ztn-config) debug admin
(admin)
(admin) exit

(ztn-config)
```

Help

The following commands are available under the admin shell.

```
(admin) help

Documented commands (type help <topic>):
=====
EOF copy exit help ping ping6 quit reboot reload show
```

Ping

Use the **ping** command to test a host's accessibility using its IPV4 address.

```
(admin) ping 10.240.141.151
PING 10.240.141.151 (10.240.141.151) 56(84) bytes of data.
64 bytes from 10.240.141.151: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 10.240.141.151: icmp_seq=2 ttl=64 time=0.206 ms
64 bytes from 10.240.141.151: icmp_seq=3 ttl=64 time=0.221 ms
64 bytes from 10.240.141.151: icmp_seq=4 ttl=64 time=0.161 ms

--- 10.240.141.151 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3049ms
```

```
rtt min/avg/max/mdev = 0.161/0.206/0.238/0.028 ms
```

Ping6

Use the **ping6** command to test a host's accessibility using its IPV6 address.

```
(admin) ping6 fe80::3673:5aff:febf:9dec
PING fe80::3673:5aff:febf:9dec(fe80::3673:5aff:febf:9dec) 56 data bytes
64 bytes from fe80::3673:5aff:febf:9dec%ma1: icmp_seq=1 ttl=64 time=0.490 ms
64 bytes from fe80::3673:5aff:febf:9dec%ma1: icmp_seq=2 ttl=64 time=0.232 ms
64 bytes from fe80::3673:5aff:febf:9dec%ma1: icmp_seq=3 ttl=64 time=0.218 ms
64 bytes from fe80::3673:5aff:febf:9dec%ma1: icmp_seq=4 ttl=64 time=0.238 ms

--- fe80::3673:5aff:febf:9dec ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3069ms
rtt min/avg/max/mdev = 0.218/0.294/0.490/0.113 ms
```

Copy Tech-Support Data

Use the **copy tech-support data** command to collect the switch support bundle. The tech support command executed from the Controller collects support bundles from all the switches in the fabric.

To collect tech support data for an individual switch, only run the **copy tech-support data** command from that switch. Perform this action when the switch isn't accessible from the Controller or when collecting a support bundle only for that switch.

```
(admin) copy tech-support data
Writing /mnt/onl/data/tech-support_240521104337.txt.gz...
tech-support_240521104337.txt.gz created in /mnt/onl/data/
(admin)
```

Show Commands

The following **show** commands are available under the admin shell.

Show Clock

```
(admin) show clock
Tue May 21 10:46:20 2024
(admin)
```

Show NTP

```
(admin) show ntp
remote                               refid                               st t when poll reach
delay    offset    jitter
=====
10.130.234.12      .STEP.      16 u    - 1024    0
0.0000    0.0000    0.0002
*time2.google.com  .GOOG.      1 u    649 1024   377
40.0469    1.6957    0.4640
+40.119.6.228      25.66.230.1 3 u    601 1024   377
50.6083   -0.5323    7.8069
(admin)
```

Show Controller

The **show controller** command displays all the configured Controllers and their connection status and role.

```
(admin) show controller
IP:Port          Proto State          Role          #Aux
10.240.141.151:6653 tcp  CONNECTED      ACTIVE         2
10.240.189.233:6653 tcp  CONNECTED      STANDBY        2
127.0.0.1:6653   tcp  CONNECTED      STANDBY        1
(admin)
```

Use the **history** and **statistics** options in the **show controller** command to obtain additional information.

```
(admin) show controller history | statistics
(admin)
```

The **show controller history** command displays the history of controller-to-switch connections and disconnections.

```
(admin) show controller history
Mon May 20 15:53:42 2024 tcp:127.0.0.1:6653:0 - Connected
Mon May 20 15:53:43 2024 tcp:127.0.0.1:6653:1 - Connected
Mon May 20 15:54:46 2024 tcp:127.0.0.1:6653:1 - Disconnected
Mon May 20 15:54:46 2024 tcp:127.0.0.1:6653:0 - Disconnected
Mon May 20 08:57:07 2024 tcp:127.0.0.1:6653:0 - Connected
Mon May 20 08:57:07 2024 tcp:127.0.0.1:6653:1 - Connected
Mon May 20 08:57:07 2024 tcp:10.240.141.151:6653:0 - Connected
Mon May 20 08:57:07 2024 tcp:10.240.141.151:6653:1 - Connected
Mon May 20 08:57:07 2024 tcp:10.240.141.151:6653:2 - Connected
Mon May 20 11:16:07 2024 tcp:10.240.189.233:6653:0 - Connected
Mon May 20 11:16:19 2024 tcp:10.240.189.233:6653:1 - Connected
Mon May 20 11:16:19 2024 tcp:10.240.189.233:6653:2 - Connected
(admin)
```

The **show controller statistics** command displays connection statistics, including keep-alive timeout, timeout threshold count, and other important information, as shown in the following example.

```
(admin) show controller statistics
Connection statistics report
Outstanding async op count from previous connections: 0
Stats for connection tcp:10.240.141.151:6653:0:
  Id: 131072
  Auxiliary Id: 0
  Controller Id: 0
  State: Connected
  Keepalive timeout: 2000 ms
  Threshold: 3
  Outstanding Echo Count: 0
  Tx Echo Count: 46438
  Messages in, current connection: 52887
  Cumulative messages in: 52887
  Messages out, current connection: 52961
  Cumulative messages out: 52961
  Dropped outgoing messages: 0
  Outstanding Async Operations: 0
Stats for connection tcp:10.240.189.233:6653:0:
  Id: 112066561
  Auxiliary Id: 0
  Controller Id: 1
  State: Connected
  Keepalive timeout: 2000 ms
```



```

Threshold: 3
Outstanding Echo Count: 0
Tx Echo Count: 42269
Messages in, current connection: 43108
Cumulative messages in: 43108
Messages out, current connection: 43114
Cumulative messages out: 43114
Dropped outgoing messages: 0
Outstanding Async Operations: 0
(admin)

```

The **show log** command displays log messages from the Syslog file.

```

(admin) show log
2024-05-20T15:53:04+00:00 localhost syslog-ng[3787]: NOTICE syslog-ng starting
up; version='3.38.1'
2024-05-20T15:52:51+00:00 localhost kernel: NOTICE Linux version 4.19.296-Open
NetworkLinux (bsn@sbs3) (gcc version 12.2.0 (Debian 12.2.0-14)) #1 SMP Thu May
16 08:35:25 UTC 2024
2024-05-20T15:52:51+00:00 localhost kernel: INFO Command line: reboot=p
acpi=on Aboot=Aboot-norcal6-6.1.10-14653765 platform=magpie sid=Calpella
console=ttyS0,9600n8 tsc=reliable pcie_ports=native pti=off reassign_pref
mem amd_iommu=off onl_mnt=/dev/mmcblk0p1 quiet=1 onl_platform=x86-64-arista
-7050sx3-48yc12-r0 onl_sku=DCS-7050SX3-48YC12
2024-05-20T15:52:51+00:00 localhost kernel: INFO BIOS-provided physical RAM
map:
...
...
2024-05-21T10:40:31-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO Starting
logrotate.service - Rotate log files...
2024-05-21T10:40:32-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO logrotate.service:
Deactivated successfully.
2024-05-21T10:40:32-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO Finished
logrotate.service - Rotate log files.
2024-05-21T10:45:32-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO Starting
logrotate.service - Rotate log files...
2024-05-21T10:45:33-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO logrotate.service:
Deactivated successfully.
2024-05-21T10:45:33-07:00 dmf-sw-7050sx3-1 systemd[1]: INFO Finished
logrotate.service - Rotate log files.
(admin)

```

Reboot and Reload

Use the **reboot** and **reload** commands to either reboot or reload the switch, as needed.

```

(admin) help reboot

Reboot the switch.

(admin) help reload

Reload the switch.

(admin) reboot
Proceed with reboot [confirm]?

(admin) reload
Proceed with reload [confirm]?
(admin)

```

Managing DMF Policies

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

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 the destination tool(s) (delivery interfaces). The DANZ Monitoring Fabric (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. Users can also send traffic to be processed by a managed service, such as time stamping, slicing, or deduplication, on a DMF service node. Forward the output from the service node to the appropriate tool for analysis.

While policies can be simple, they can also be more complicated when optimizing hardware resources, such as switching TCAM space. Also, DMF provides different switching modes to optimize policies based on use cases and switch capabilities. Arista Networks recommends planning the switching mode before configuring policies in a production deployment.

For further information, refer to the chapter [Advanced Policy Configuration](#).

3.2 DMF Policies Page

Overview

While retaining all information from the previous version, the new policy page features a new layout and design and offers additional functionality for easier viewing, monitoring, and troubleshooting of policies.


Figure 3-1: DMF Policies

Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces
_PLC_test_anthony	1 Interface	-	-	-	-	2 Interfaces
APP-FILTER-DROP	1 Interface	-	-	1 Service	-	1 Interface
DUPLICATE-DC-TAPS	1 Interface	-	-	-	-	8 Interfaces
ERSPAN-DC-1-CLUSTER-1	2 Interfaces	-	1 vCenter	-	-	-
ERSPAN-DC-1-CLUSTER-2	1 Interface	-	1 vCenter	-	-	-
ERSPAN-DC-2-CLUSTER-1	1 Interface	-	1 vCenter	-	-	-
ERSPAN-TO-DMF-DC-EMPTY	-	-	1 vCenter	-	-	-
GENERATE-DAPPER	1 Interface	-	-	1 Service	-	-
GENERATE-IPFIX	1 Interface	-	-	1 Service	-	-
GENERATE-IPFIX	1 Interface	-	-	1 Service	-	-

Header Action Items

- Refresh Button

Figure 3-2: Refresh Button



The page refreshes every 60 seconds automatically. Click the **Refresh** button to manually refresh the page.

- Create Policy Button


Figure 3-3: Create Policy Button



Click the **+ Create Policy** button to open the policy creation page.

- Clear Stats Button

Figure 3-4: Clear Stats Button



Click the **Clear Stats** button to clear all DMF interface's runtime stats.

Quick Filters

- Show Quick Filters Button

Figure 3-5: Show Quick Filters



By default, the feature is toggled **on** and displays four quick filter options. When toggled **off**, the four quick filters are no longer displayed.

Figure 3-6: Four Filter Options



Four quick filter cards display the policy counts that meet the filter criteria and the filter name. The quick filter cards support multi-select.

- Radio Buttons

Figure 3-7: Table View / Interface View



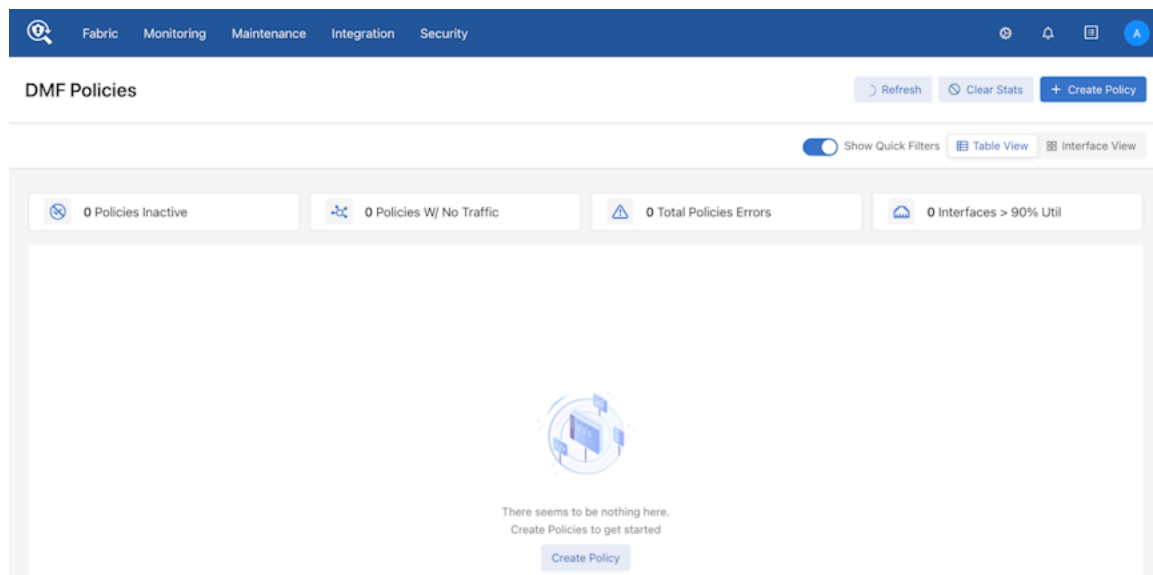
Switch page views between **Table View** and **Interface View**. Refer to the **Table View** and **Interface View** sections below for more information.

Table View

The table view is the default landing view of the **Policies** Page.

The page displays an empty table with the Create Policy button when no configured policies exist.

Figure 3-8: DMF Policies



Conversely, when configured policies exist, the table view displays the list of policies.


Figure 3-9: List of Policies

Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces
_mlag_MLAG-POLICY_T3X5-Leaf-s5212F-1	1 Interface	-	-	-	-	1 Interface
_mlag_MLAG-POLICY_T3X7-Leaf-s5248F-1	1 Interface	-	-	-	-	1 Interface
APP-FILTER-DROP	1 Interface	-	-	1 Service	-	1 Interface
DUPLICATE-DC-TAPS	1 Interface	-	-	-	-	12 Interfaces
ERSPAN-DC-1-CLUSTER-1	2 Interfaces	-	1 vCenter	-	-	-
ERSPAN-DC-1-CLUSTER-2	1 Interface	-	1 vCenter	-	-	-
ERSPAN-DC-2-CLUSTER-1	1 Interface	-	1 vCenter	-	-	-
FRSPAN-TN-						

Action Buttons

Several buttons in the policy table provide quick access to corresponding functionality. These are:

Figure 3-10: Action Buttons



Delete Button

- Disabled by default (when no policies are selected).
- Enabled when one or more policies are selected.
- Used to delete selected policies.

Edit Button

- Disabled by default (when no policies are selected).
- Enabled only when a policy is selected.
- Navigate to the editing workflow (the new policy edit workflow).

Duplicate Button

- Disabled by default (when no policy is elected).
- Enabled only when one policy is selected.
- Navigate to the create policy workflow (the new policy create workflow) with an empty name input field while retaining the same settings from the selected policy.

Table View Filters

Figure 3-11: Filter Views

Filter

Policy Filter(s)

☐ Select All

☐ Show Inactive Policies

☐ Show Policies w/ No Traffic

☐ Show Policies w/ Errors

☐ Show Overlapping Policies

DMF Interface Name(s)

Select

Click the **Filter** button to open the filter menu.

Policy Filter(s)

- There are four quick policy filters. The first three filters overlap with the quick filters; thus, enabling or disabling them will trigger changes to the quick filter button.

DMF Interface Name(s)

- Filters out policies by DMF interfaces that are selected from the drop-down list.
- Searchable
- Allows multiple selections applying OR logic.

Policies Table

Figure 3-12: Policy Table

<input type="checkbox"/>	Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces	
<input type="checkbox"/>	APP-FILTER-DROP	1 Interface	-	-	1 Service	-	1 Interface	
<input type="checkbox"/>	DUPLICATE-DC-TAPS	1 Interface	-	-	-	-	8 Interfaces	
<input type="checkbox"/>	ERSPAN-DC-1-CLUSTER-1	2 Interfaces	-	1 vCenter	-	-	-	
<input type="checkbox"/>	ERSPAN-DC-1-CLUSTER-2	1 Interface	-	1 vCenter	-	-	-	
<input type="checkbox"/>	ERSPAN-DC-2-CLUSTER-1	-	-	1 vCenter	-	-	-	
<input type="checkbox"/>	ERSPAN-TO-DMF-DC-EMPTY	-	-	1 vCenter	-	-	-	
<input type="checkbox"/>	GENERATE-DAPPER	-	-	-	1 Service	-	-	
<input type="checkbox"/>	GENERATE-IPFIX	1 Interface	-	-	1 Service	-	-	
<input type="checkbox"/>	IDENTIFY-APP	1 Interface	-	-	1 Service	-	-	
<input type="checkbox"/>	RECONSTRUCT-APP	1 Interface	-	-	1 Service	-	-	

The **Policy** table displays all policies; each column shows the number of interfaces and services corresponding to that policy.

Figure 3-13: Search



Table Search

The **Policy** table supports search functionality. Click the **magnifying glass** icon in the last column to activate the search input fields and search the results by the context of each column.

Table Search: The Policy table supports search functionality. Click on the **magnifying glass** icon in the last column to activate the search input fields. Search results by the context of each column.

Figure 3-14: Table Search

Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces	
Type to Search	Type to Search	Type to Search	Type to Search	Type to Search	Type to Search	Type to Search	

Figure 3-15: Expand Icon




Expand Policy + Icon

Hidden for an unconfigured policy. Click the expand + icon to view the policy's interfaces and services information.

Figure 3-16: Expanded View Example

Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces	
APP-FILTER-DROP	1 Interface LAG-T3X7-7050cx3-2 534.1 Mbps of 200.0 Gbps	-	-	1 Service MS-APP-FILTER-ACTIVE	-	1 Interface LOOP-DEL-2 101.2 kbps of 25.0 Gbps	

Figure 3-17: Expand Group



Interfaces Group Expand + Icon

For policies configured with an interface group, an expand + icon with group displays by default. Click on the group expand + icon to view the detailed information on the interfaces belonging to that group.

Figure 3-18: Filter Interface Details

Policy Name	Filter Interfaces
TOI-test2	1 Interface filter1 0.0 bps of 10.0 Gbps
	Group group1

Policy Name Tooltip

Hovering over policy names displays the tooltips for the policy, including Configuration / Runtime / Details state.

Figure 3-19: Tooltip

ERSPAN-DC-1-CLUSTER-1	2 Interfaces	Config Status: active - drop
		Runtime Status: installed to drop
		Detailed Status: installed to drop


Policy Error Icon

Figure 3-20: Policy Error Icon



Policies with errors will display this icon after the policy name.

Figure 3-21: Error with Policy Name



Clicking the **error** icon will display an error window with detailed information.

Figure 3-22: Detailed Error Information



Checkbox

Figure 3-23: Checkbox



Disabled for unconfigured policies. Use the **checkbox** to select a policy and the applicable function buttons (described above) as required.

Table Interaction


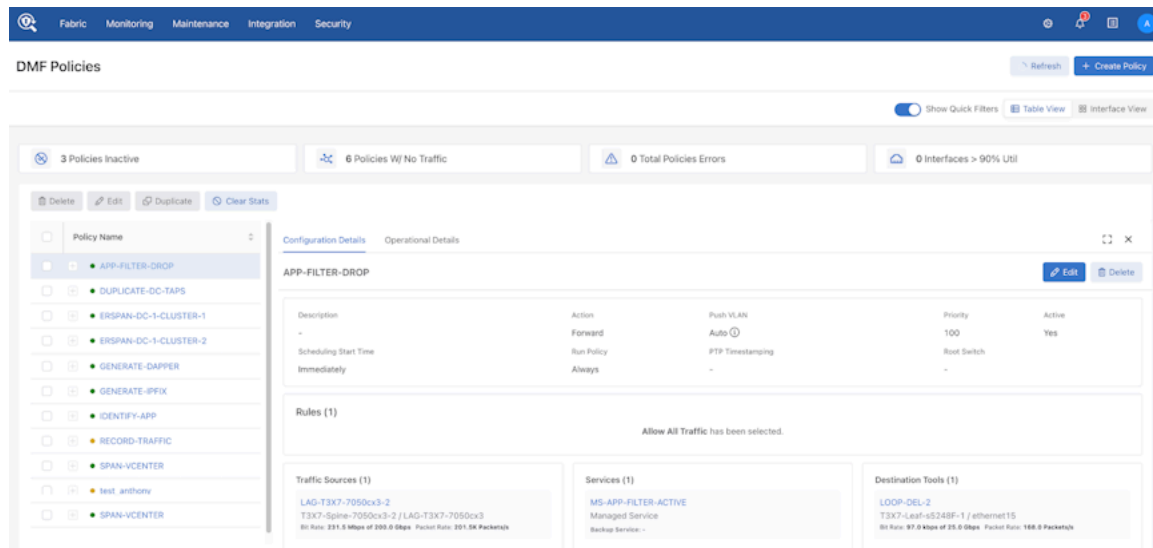
- All columns support sorting.
- Clicking on a **policy name** opens the policy table split view. The table on the left displays the policy names. The table on the right provides two tabs showing **Configuration** and **Operational Details**.
- Select a policy name from the **Policy Name** list to view its configuration or operational details in the split view.
- Use the  icon to view the information in full-screen mode or the **X** icon to close the split view and return to the table view.

Figure 3-24: DMF Policies

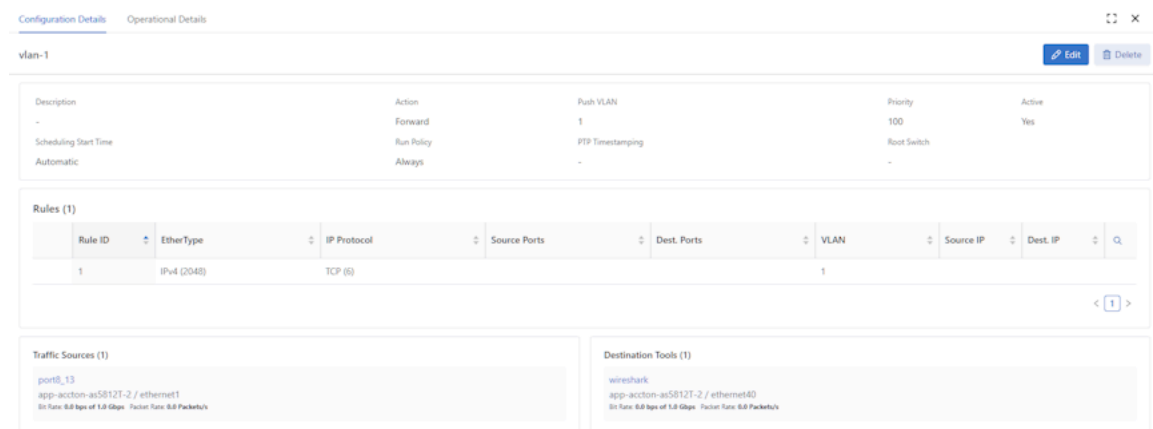


Configuration Details

Access the **Configuration Details** tab by selecting a policy in either **Table View** or **Interface View**. This tab displays all of the configured settings for the selected policy.

The top row of the **Configuration Details** tab displays the selected policy name and an **Edit** and **Delete** button. The **Edit** button opens the **Edit Policy** configuration page with policy information prefilled, and the **Delete** button opens a confirmation dialog window before deleting a policy. The default **Table View** opens after deleting a policy.

Figure 3-25: Configuration Details



The second component of the Configuration Details is the Quick Facts box. This component displays the Description, Action, Push VLAN, Priority, Active, Scheduling Start Time, Policy Run Duration, PTP Timestamping, and Root Switch values.

- **Description:** An info icon shows the entire description in a tooltip.
- **Action:** Forward, Drop, Capture, or None.
- **Active:** Policy active status, Yes or No.

Scheduling Start Time: Either **Automatically** or the DateTime it is scheduled to start, in terms of the current Time Zone configured on the DMF. When setting DateTime to **Now** during policy creation, the time of creation will be the Scheduling start time.

- **Automatic:** The policy will always run. There's no expiration.
- **Now:** The policy starts from now, and duration and packet expiration may apply. The policy runs from now with no expiration.

Figure 3-26: Start Time

Scheduling Start Time
Dec 13 2023, 19:16:31 PM UTC

- **Run Policy:** The duration the policy should run. The default value is **Always**. Set a time limit (i.e., 4 hours) or a packet limit (i.e., 1,000 packets) The tooltip explains that the policy will stop running when reaching either of the limits.

The third component is the **Rules Table**, which displays all **Match Traffic** rules configured for the policy. The default value is **Allow All Traffic**. Optionally, configure **Deny All Traffic**.

Figure 3-27: Allow All Traffic





Figure 3-28: Deny All Traffic



When configuring custom rules, the **Rules Table** is displayed. The table is horizontally scrollable, and each column is searchable and sortable. The **Edit Policy** feature provides rule management, including **Edit**, **Add**, and **Delete** functionality.

Figure 3-29: Rules Table

Rules (4)										
Rule ID	EtherType	IP Protocol	Source Ports	Dest. Ports	VLAN	Source IP	Dest. IP			
1	IPv4 (2048)									
2	IPv4 (2048)	TCP (6)								
3	IPv4 (2048)	UDP (17)								
4	IPv6 (34525)									

The next component is the Interface Info Columns.

Figure 3-30: Information Columns

<p>Traffic Sources (1)</p> <p>V5-FEED TomahawkPlus-Leaf-s5048F-1 / V5-FEED Bit Rate: 239.6 Mbps of 20.0 Gbps Packet Rate: 271.3K Packets/s</p>	<p>Destination Tools (8)</p> <p>LAG-T2-s4048 T2-Leaf-s4048-1 / LAG-T2-s4048 Bit Rate: 240.3 Mbps of 20.0 Gbps Packet Rate: 272.1K Packets/s</p> <p>LAG-T3X5-s5212F T3X5-Leaf-s5212F-1 / LAG-T3X5-s5212F Bit Rate: 239.7 Mbps of 50.0 Gbps Packet Rate: 271.4K Packets/s</p> <p>LAG-T3X7-7050cx3-1 T3X7-Spine-7050cx3-1 / LAG-T3X7-7050cx3 Bit Rate: 239.5 Mbps of 200.0 Gbps Packet Rate: 271.1K Packets/s</p> <p>LAG-T3X7-7050cx3-2 T3X7-Spine-7050cx3-2 / LAG-T3X7-7050cx3 Bit Rate: 239.6 Mbps of 200.0 Gbps Packet Rate: 271.4K Packets/s</p> <p>LAG-T3X7-as7326 T3X7-Leaf-as7326-1 / LAG-T3X7-as7326 Bit Rate: 0.0 bps of 50.0 Gbps Packet Rate: 0.0 Packets/s</p>
---	--

There are three primary columns: **Traffic Sources**, **Services**, and **Destination Tools**.

- The Traffic Sources column includes Filter Interfaces, vCenters, and CloudVision Portal associated with the policy.
- The Services column includes Managed Services and Services associated with the policy.
- The Destination Tools column includes Delivery interfaces and RN Fabric Interfaces associated with the policy.

These columns display the DMF Interface name in the **interface card**, and the name includes a link to the Interfaces page. The switch name and physical interface name appear in this format: **SWITCH-NAME / INTERFACE-NAME** under the DMF interface name. The bit rate and packet rate operational state data appear for each interface. Each column is only displayed if the policy has one or more interfaces of that type.

Figure 3-31: Traffic Sources

<p>Traffic Sources (1)</p> <p>LAG-T3X7-7050cx3-1 T3X7-Spine-7050cx3-1 / LAG-T3X7-7050cx3 Bit Rate: 173.8 Mbps of 200.0 Gbps Packet Rate: 205.3K Packets/s</p>	<p>Services (1)</p> <p>MS-APP-ID-ACTIVE Managed Service MS-APP-ID-BACKUP</p>
--	---

The services column renders for all policies that have at least one service. The service name appears for each card, which contains a link to either the Services or Managed Services page. Under the service name, the service type (Managed Service or Service) appears if the service has a backup name that also appears.

Figure 3-32: Managed Service

<p>CloudVision (1)</p> <p>cvp689 cal433: 3 Entries</p>	<p>Filter Interfaces (1)</p> <p>cvp--i2gre--cvp689--cal433--TUNNEL2--dmf--C2--cvp-212b6632 C2 / cvp-212b6632 Bit Rate: 0.0 bps of 0.0 bps Packet Rate: 0.0 Packets/s</p>	<p>Destination Tools (1)</p> <p>n D1 / ethernet4 Bit Rate: 0.0 bps of 10.0 Gbps Packet Rate: 0.0 Packets/s</p>
---	---	---

There is a special case for policies that have CloudVision port mirroring sessions. To differentiate the CloudVision source interfaces from the auto-generated DMF filter interfaces, DMF creates two columns: **CloudVision** and **Filter Interfaces**.

The cards in the CloudVision column show the connected CloudVision portal and the number of port mirroring sessions for each device in the CloudVision portal. Filter Interfaces and vCenters are now in the Filter Interfaces column. There are no differences between the Services and Destination Tools columns.

The last component only displays for policies with CloudVision port mirroring sessions.

Figure 3-33: Port Mirroring Sessions

cvp689 Port Mirroring Session Entries (3)							
ID	Device	Source Interface	Monitor Type	Tunnel Source	Tunnel Endpoint	SPAN Interface	Direction
1	cal433	Ethernet1	GRE	1.1.1.1	TUNNEL2		Bidirectional
2	cal433	Ethernet2	GRE	1.1.1.1	TUNNEL2		Bidirectional
3	cal433	Ethernet3	GRE	1.1.1.1	TUNNEL2		Bidirectional

The **Port Mirroring Session Entries** table shows all configured Port Mirroring Sessions for a CloudVision portal. The **Device**, **Source Interface**, **Monitor Type**, **Tunnel Source**, **Tunnel Endpoint**, **SPAN Interface**, and **Direction** columns display the same values configured in the Port Mirroring Table in the Add Traffic Sources component in the Create Policy flow. Each column is sortable.

For more information on the configuration flow for CloudVision port mirroring, please refer to the documentation in the **Create Policy** section.

Operational Details

Clicking on the **Operational Details Tab** navigates to the Operational Details view.

Figure 3-34: Operational Details

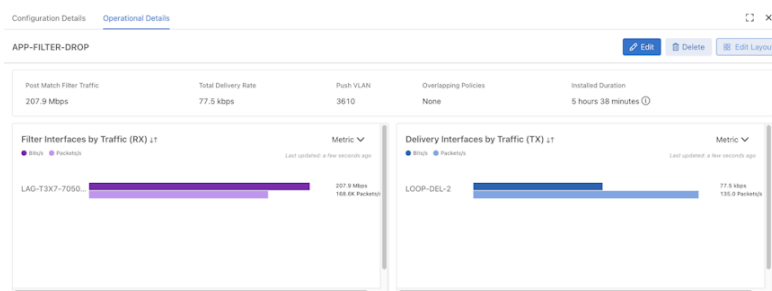



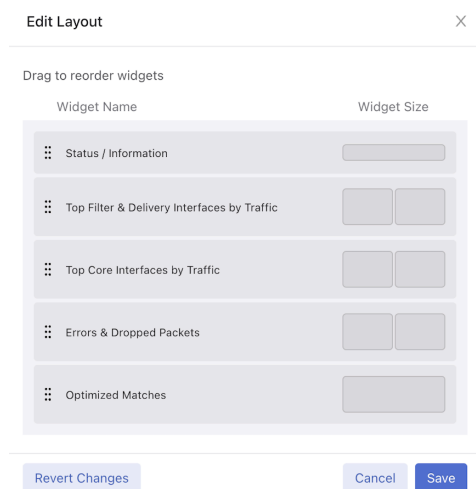
Figure 3-35: Action Buttons



- **Edit:** Clicking the Edit button opens the Editing Policy window for making changes to the policy.
- **Delete:** Clicking the Delete button deletes the policy.

- **Edit Layout:** Clicking the Edit Layout button opens the editing layout window. Move the widgets by dragging the components in order of user preference. Click the **Save** button to save the changes. DMF preserves the order of the widgets when the same user logs back in.

Figure 3-36: Edit Layout



Widgets

Status / Information

Status and information include basic operational information about the policy.


Figure 3-37: Operational Information

Post Match Filter Traffic 264.4 Mbps	Total Delivery Rate 258.1 kbps	Push VLAN 3610	Overlapping Policies None	Installed Duration 5 hours 45 minutes ⓘ
---	-----------------------------------	-------------------	------------------------------	--

Installed Duration

Hover over the info icon to see the installed time in the UTC time zone.

Figure 3-38: Install Time



Top Filter and Delivery Interfaces by Traffic

Figure 3-39: Top Filter and Delivery Interfaces by Traffic

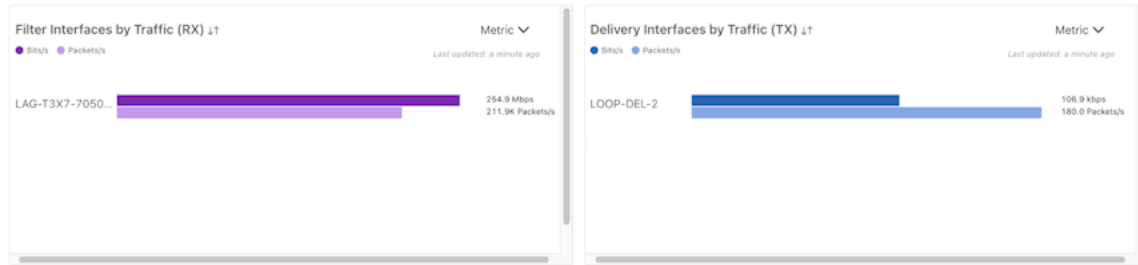




Figure 3-40: Select Metric



Click the **Metric Drop-down** menu and choose the metrics to display in the chart. Only the selected metrics appear in the **Badge**, **Labels**, and **Bar Chart**.

- **Badge:** Colored dots and text indicate the content represented by different bars in the bar chart.
- **Interface Name**

Figure 3-41: Labels




Hover over the **interface name** to see the full name in the tooltips.

- **Labels:** Display the number and unit corresponding to the bar.
- **Bar Chart:** Displays the numerical value of traffic.
- **Empty State**
 - Display title, last updated time, and disabled metric drop-down.
 - The Edit Policy button opens the edit policy window.

Top Core Interfaces by Traffic


Figure 3-42: Top Core Interfaces by Traffic



The **Core Interfaces by Traffic** chart is similar to **Filter Interfaces / Delivery Interfaces by Traffic** charts, which have **Metric Drop-down**, **Badge**, **Interface Name**, **Labels**, and **Bar Charts** with similar functionality.

Errors & Dropped Packets

Figure 3-43: Errors



The **Errors** chart is similar to **Filter Interfaces / Delivery Interfaces by Traffic** charts, which have **Metric Drop-down**, **Badge**, **Interface Name**, **Labels**, and **Bar Charts** with similar functionality. Hovering over the bar displays all error counts and rate information.

Figure 3-44: Packets Dropped



The **Dropped Packets** chart is similar to **Filter Interfaces / Delivery Interfaces by Traffic** charts, which have **Badge**, **Interface Name**, **Labels**, and **Bar Charts** with similar functionality. Hovering over the bar displays all packet dropped counts and rate information.

Optimized Matches

Displays optimized match rules.

Figure 3-45: Optimized Matches

Optimized Matches Last updated: a few seconds ago

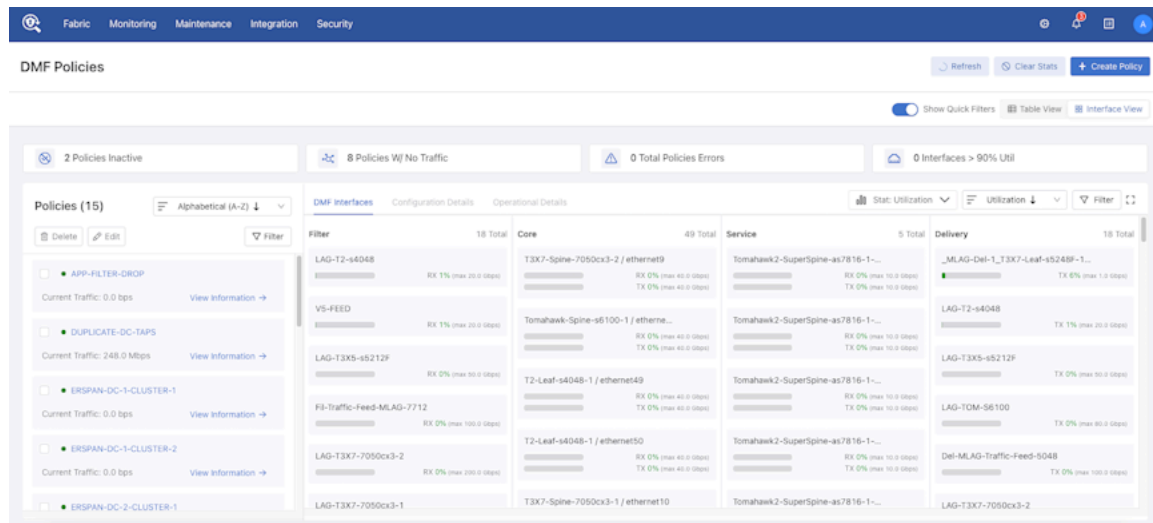
Rule ID	Rule	
1	any	
1	ether-type 34525	

< 1 >

Interface View


As a new feature of the **DMF Policies** page, the **Interface** view offers an alternative way to view policies, allowing for an intuitive visualization of all policies-related interfaces.

Figure 3-46: Interface View




Policies Column

Figure 3-47: Policies Column



- A Policies header displaying count. The column shows the total count when no filters are applied, or the filtered policies count in the format of **x Associated**.
- The drop-down menu enables data sorting using multiple attributes.
- The **Delete** button deletes the selected policies.
- The **Edit** button opens the selected policy in edit mode.
- The **Filter** drop-down is similar to the table view filters but without an interface filtering option.

Figure 3-48: Filter

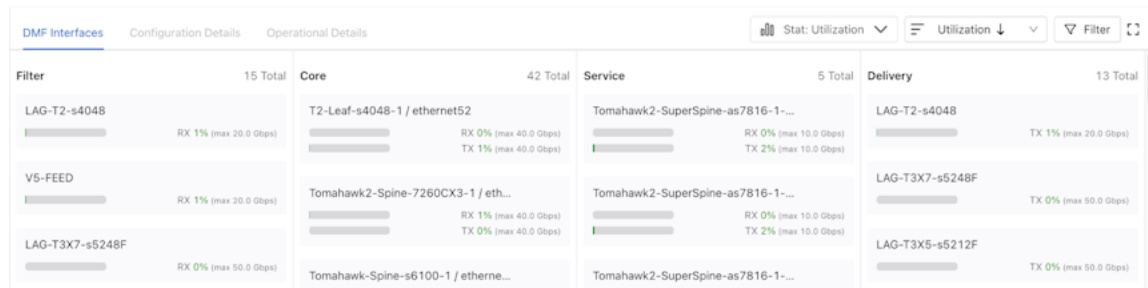


- A list of policies with quick facts and user interactions.
- The **checkbox** enables policy selection for deletion and editing.
- Badges with different colors indicate policy run time status.
- Policy name with tooltip on hover displaying configuration, runtime, and detailed status.
- **Current Traffic** display in bps.
- Clicking the **View Information** button highlights the policy:
 - Only shows the interfaces associated with the selected policy in the **DMF Interfaces** tab.
 - Enable **Configuration Details** and **Operational Details**.

- Clicking on an active policy card deselects the previously selected policy:
 - De-emphasizes the policy and resets card styles and tabs accessibility.
 - Reveals all the interfaces in **DMF Interfaces**.
 - Interface card highlights in the **DMF Interfaces** tab can co-exist, leading to a more granular search.

DMF Interfaces

Figure 3-49: DMF Interfaces



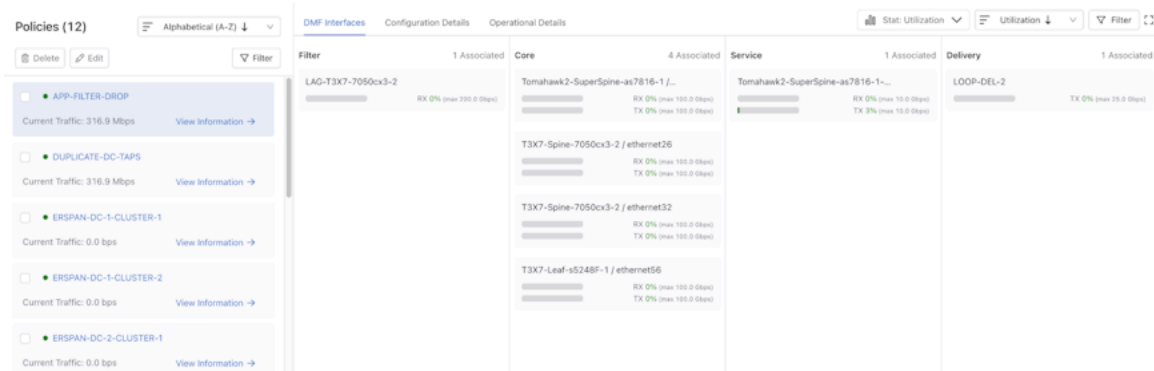
- Active tab by default
- Header Row**
 - Stat selector:** Choose between Utilization, Bit Rate, and Packet Rate to display in the subsequent interface info cards.
 - Sorter selector:** Choose between Utilization and interface name to sort the interfaces in ascending or descending order.
 - Filter drop-down:**
 - Utilization range filter
 - Switch name selector
 - DMF interface name selector
- Interface Column**
 - Header:** Specifies interface category and count, showing **X Associated** when filters apply and **X Total** otherwise.
 - Interface Information Card
 - Interface name
 - Stat**
 - Utilization
 - Bit Rate
 - Packet Rate
 - Text:** Display detailed information about the selected stat of the current interface.

Interaction

- Selecting one policy card:

The selected policy card highlights and filters interfaces to only those configured to the policy and hides interfaces not configured in the selected policy.

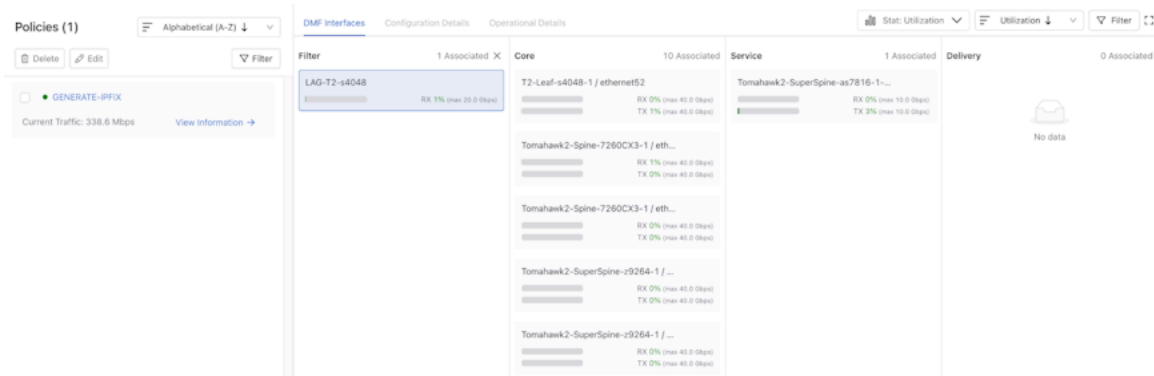
Figure 3-50: Policy Card



- Selecting one interface card:

The selected interface card highlights and filters policies to only those configured to the interface and hides interfaces not configured in the filtered policies mentioned above.

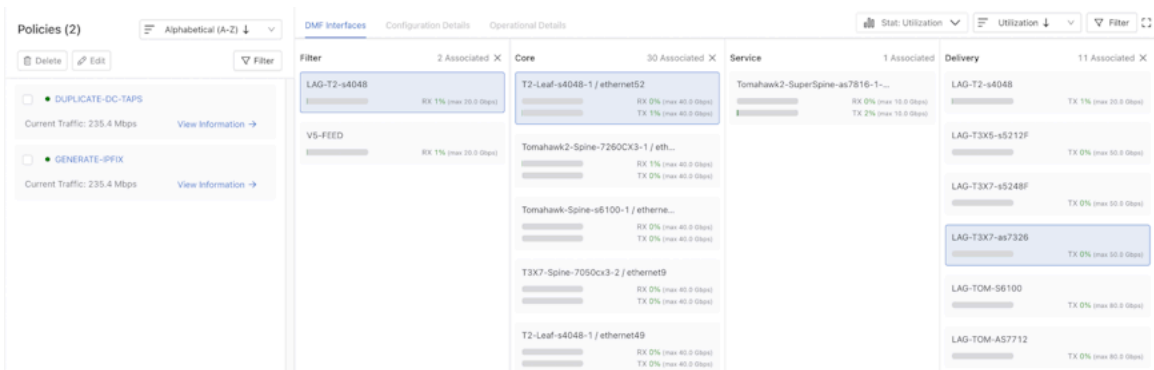
Figure 3-51: Single Interface Card



- Selecting multiple interface cards (any columns):

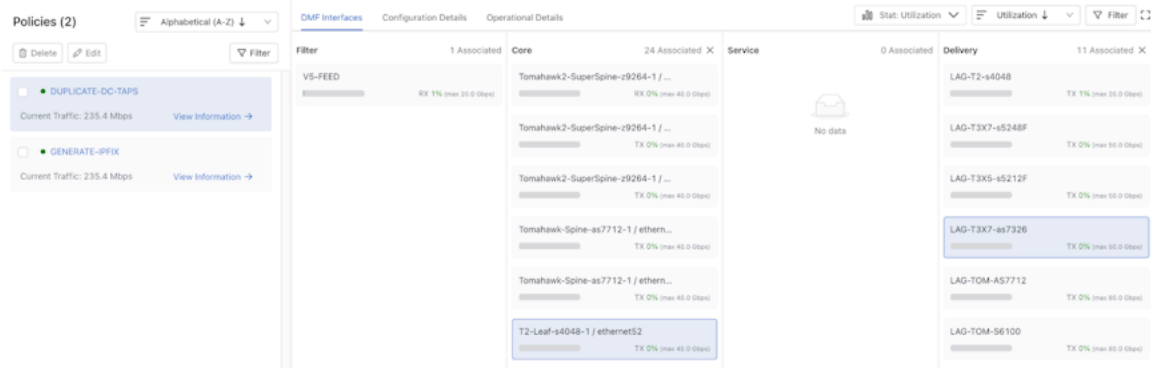
The selected interface cards highlight and filter policies to only those configured on the selected interfaces and hide interfaces not configured in the filtered policies mentioned above.

Figure 3-52: Multiple Interface Cards



Highlighted policy and interface cards can co-exist, leading to a more granular search.

Figure 3-53: Policy and Interface Cards



Configuration Details

The GUI is similar to **Table View > Configuration Details**. Please refer to the [Configuration Details](#) section.

Operational Details

The GUI is similar to **Table View > Operational Details**. Please refer to the [Operational Details](#) section.

3.3 Policy Elements

Each policy includes the following configuration elements:

- **Filter interfaces:** these identify the ingress ports for analyzing the traffic for this policy. Choose individual filter interfaces or one or more filter interface groups. Select the **Select All Filter Interfaces** option, intended for small-scale deployments.
- **Delivery interfaces:** these identify the egress ports for analyzing the traffic as part of this policy. Choose individual delivery interfaces or one or more delivery interface groups. Like filter interfaces, a **Select All Delivery Interfaces** option is available for small deployments.
- **Action:** identifies the policy action 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. 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. Use Drop to measure the bandwidth of matching traffic without forwarding it to a tool. 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. Select at least one or more filter interfaces. A policy with a capture action can only run for a short period. For continuous packet capture, use the DANZ Monitoring Fabric (DMF) recorder node. Refer to the chapter [Using the DMF Recorder Node](#) for details.



Note: The policy will not be installed if an action is not selected.

- **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 or the match rule is **Deny All Traffic**, the policy is not installed. One policy can specify multiple match rules, differentiating each rule by its rule number.



Note: The rule numbers do not define the order in which the rules will be installed or processed. The numbering allows a 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 using the **active** or **inactive** sub-command from the config-policy sub-model. By default, a policy is active when initially configured.
- **Priority** (optional): unless a user specifies, all policies have a priority of **100**. When sharing filter/ingress ports across policies, a policy with a higher priority will get access to matching traffic first. Traffic 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. Typically, this VLAN ID is in the range of **1-4094** and auto-increments by **1**. However, to specific policy with a specific VLAN ID, first define a smaller VLAN range using the command **auto-vlan-range** and then pick a VLAN outside 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, the policy does not move back. In some cases, this causes traffic overload. One way to overcome this problem is to specify a root

switch in each policy. The policy is rerouted through other switches when the root switch goes down. When the root switch comes back, DMF reroutes the 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 affiliated with a filter interface forwards to all the delivery interfaces defined in the policy.

Except for a capture action policy, a policy runs indefinitely once activated. 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 automatically deactivates.

Note:



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

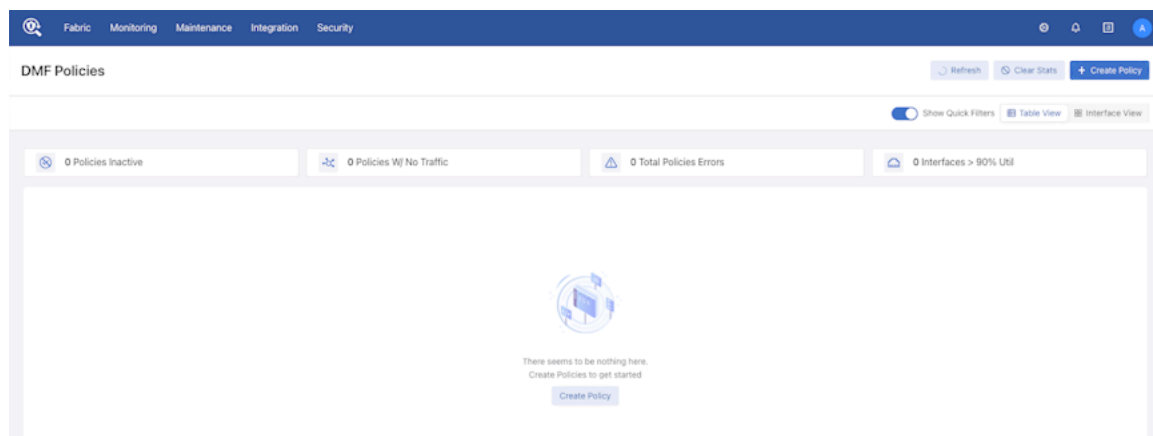
3.4 Configuring a Policy

3.4.1 Configure a Policy Using the GUI

DANZ Monitoring Fabric (DMF) introduces a newly designed **Create Policy** configuration workflow, replacing the former workflow page.

There are two possible entry points for creating a policy. The first is via the **Create Policy** button continuously displayed on the top-right corner of the **DMF Policies** page, or the second is via the **Create Policy** button, which appears on the central panel of the same page when no configured policies exist.

Figure 3-54: DMF Policies



Clicking the **Create Policy** button opens the new **Policy Creation** configuration page, which supports moving, minimizing, expanding, collapsing, and closing the window using the respective icons in the menu bar.


Figure 3-55: Create Policy


The screenshot shows the 'Create Policy' window with the following elements:

- Title Bar:** 'Create Policy' with window control icons (move, minimize, maximize, close).
- Policy Details:** A section with a toggle for 'Enable Advanced Options' and a 'Collapse' button.
- Form Fields:**
 - Name:** Input field with 'Input value'.
 - Description:** Input field with 'Input value'.
 - Action:** Dropdown menu with 'Forward' selected.
 - Push VLAN:** Input field with '##'.
 - Priority:** Input field with '100'.
 - Active:** Toggle switch.
- Port Selection:** A tab labeled 'Match Traffic' with a 'Stat: Bit Rate' dropdown.
- Traffic Sources:** A section with a 'Create' button and a '+ Add Port(s)' button.
- Services:** A section with a '+ Add Service(s)' button.
- Destination Tools:** A section with a 'Create' button and a '+ Add Port(s)' button.
- Footer:** A message 'Match Traffic: Allow All Traffic (Default) has been selected' and 'Cancel' and 'Create Policy' buttons.

Figure 3-56: UI Controls

Create Policy



Move: Click (and hold) any part of the title section of the window or the  icon to drag and reposition as required. Moving the window in full-size mode is not possible.

Expand: Click the  icon to enlarge the window.

Minimize: Click the  icon to minimize the window and the  icon to return to the standard view.

Proceed to the following sections for create and manage policies.

3.4.1.1 Create a New Policy

Create a New Policy

To create a new Policy, complete the required fields in the **Policy Details** section and configure settings under the **Port Selection** tab (optional) and the **Match Traffic** tab (optional). Please refer to the [Policy Details](#), [Port Selection Tab](#), and [Match Traffic Tab](#) sections for more detailed information on configuring settings.

Once configured, click the **Create Policy** button on the bottom-right corner to save the changes and finish the policy creation.

Figure 3-57: Create Policy

Match Traffic: **Allow All Traffic (Default)** has been selected Cancel Create Policy

Policy Details

Figure 3-58: Policy Details

Policy Details ☐ Enable Advanced Options ⌵ Collapse

* Name	Description	Action	Push VLAN	* Priority ?	Active
<input type="text" value="Input value"/>	<input type="text" value="Input value"/>	Forward ⌵	<input type="text" value="##"/>	<input type="text" value="100"/>	<input checked="" type="checkbox"/>

Enter the primary information for the policy:

- **Policy Name** (must be unique)
- **Description**
- **Policy Action:** Capture, Drop, Forward (default)



Note: The Destination Tools column is not available when Drop and Capture actions are selected.

- **Push VLAN**
- **Priority:** By default, set to 100 if no value is specified.
- **Active:** By default, set to enabled.
- **Advanced Options:** By default, disabled.

When **Advanced Options** is enabled, the following configuration settings are available:

Figure 3-59: Advanced Options

Scheduling Run Time ? PTP Timestamping Root Switch ?

⌵ ⌵ ☐ ⌵

- **Scheduling:** There are four options:
 - **Automatic:** The policy runs indefinitely.
 - **Now:** The policy starts running immediately; use **Run Time** to determine when the policy should stop.

- **Set Time:** Set a specific date and time to start the policy.

Figure 3-60: Scheduling

<< < Dec 2023 > >>									
Su	Mo	Tu	We	Th	Fr	Sa	13	40	02
26	27	28	29	30	1	2	14	41	03
3	4	5	6	7	8	9	15	42	04
10	11	12	13	14	15	16	16	43	05
17	18	19	20	21	22	23	17	44	06
24	25	26	27	28	29	30	18	45	07
31	1	2	3	4	5	6	19	46	08
							20	47	09

Now OK

- **Set Delay:** Start the policy using relative time options.

Figure 3-61: Set Delay

Set Delay

1 minute
10 minutes
1 hour
2 hours
12 hours

hour(s)
 min(s)

*Delay starts when policy is saved Ok

- **Run Time:** There are two options:
 - **Always:** (default).

- **For Duration:** Selecting **For Duration** allows using **Time Input** to set the time number and the **Unit** selector to set the time unit. Select the checkbox to use **Packet Input** and enter the required packet number (1000, by default).

Figure 3-62: Run Time



The image shows a configuration panel titled "Run Time". It contains three main input fields in the first row: a dropdown menu currently showing "For Duration", a text input field containing the number "0", and another dropdown menu currently showing "second(s)". Below these, there is a checkbox that is currently unchecked, followed by the text "Or until", a text input field containing "1000", and the text "packet(s) are delivered". At the bottom of this section, the text "(whichever comes first)" is displayed.


- **PTP Timestamping:** Disabled by default.
- **Root Switch:** By default, set to a locked state. Click the lock icon to unlock and select a root switch.

Additional Controls

Figure 3-63: Collapse



Figure 3-64: Show



- **Collapse and Show:** Visually hide or unhide the basic policy configurations to manage the view of the other configuration fields.

3.4.1.2 Traffic Sources

The **Traffic Sources** column displays the associated traffic sources in the policy.

Figure 3-65: Traffic Sources

The screenshot shows the 'Create Policy' dialog box with the 'Port Selection' tab active. The 'Traffic Sources' section is expanded, showing three groups of traffic sources. Each group has a list of interfaces with their descriptions and bit rates. The 'Match Traffic' section shows 'Allow All Traffic (Default)' selected. The 'Services' and 'Destination Tools' sections are empty.

To add Sources, click on the **Add Port(s)** button. The page allows adding **Filter Interfaces** and **Groups**, **vCenters**, or **CloudVision Portals**.



Note: The left column has three multiple groups. Select the corresponding type of traffic source to view the available selections. After making all desired selections, confirm them using the **Add N Sources** button.

Figure 3-66: Add Sources

The screenshot shows the 'Create Policy' dialog box with the 'Port Selection' tab active. The 'Traffic Sources' section is expanded, showing a list of vCenters and CloudVision Portals. The 'Selected Traffic Sources' section shows the selected sources: LAG-T3X7-as7726, T3X7-Spine-as7726-1, and vCenters (2).

Interfaces can be searched by the available information in the interface tiles using the search bar. Clicking the



icon reveals sorting and filtering options using **Display Data**, which includes:

- **Sort** - By default, DMF sorts the data in descending Bit Rate order. Optionally, sort the data by ascending Bit Rate order or alphabetically.
- **Bit Rate** (default), **Utilization** percentage, or **Packet Rate**
- **Switch Name**
- **Interface Name(s)**

Figure 3-67: Traffic Sources Display Data

The screenshot displays the 'Traffic Sources' management interface. On the left, a list of traffic sources is shown, each with a green dot icon, a name, a description, and a bandwidth value. A modal window titled 'Display Data' is open in the foreground, allowing users to filter and sort the data.

Traffic Sources

^ Filter Interfaces/Groups

☐ Select All Filter Interfaces

Search

- LAG-TOM-AS7712
Tomahawk-Spin... LAG-TOM-A
252.5 Mbps
- LAG-T3X7-7050cx3-2
T3X7-Spine-70... LAG-T3X7-
251.2 Mbps c
- LAG-T3X7-s5248F
T3X7-Leaf-s52... LAG-T3X7-
250.9 Mbps
- LAG-TOM2-AS7816
Tomahawk2-Sup...LAG-TOM2-
239.4 Mbps c
- LAG-T3X7-7050cx3-1
T3X7-Spine-70... LAG-T3X7-
239.1 Mbps c
- LAG-T3X7-as7326

Selected Traffic Sources

Display Data

☒ Bit Rate ☐ Util% ☐ Packet Rate

Sort

- ☒ Bit Rate (Descending)
- ☐ Bit Rate (Ascending)
- ☐ Alphabetical (A-Z)
- ☐ Alphabetical (Z-A)

Filter

Bit Rate (Mbps)

0 20000

Switch Name

Select

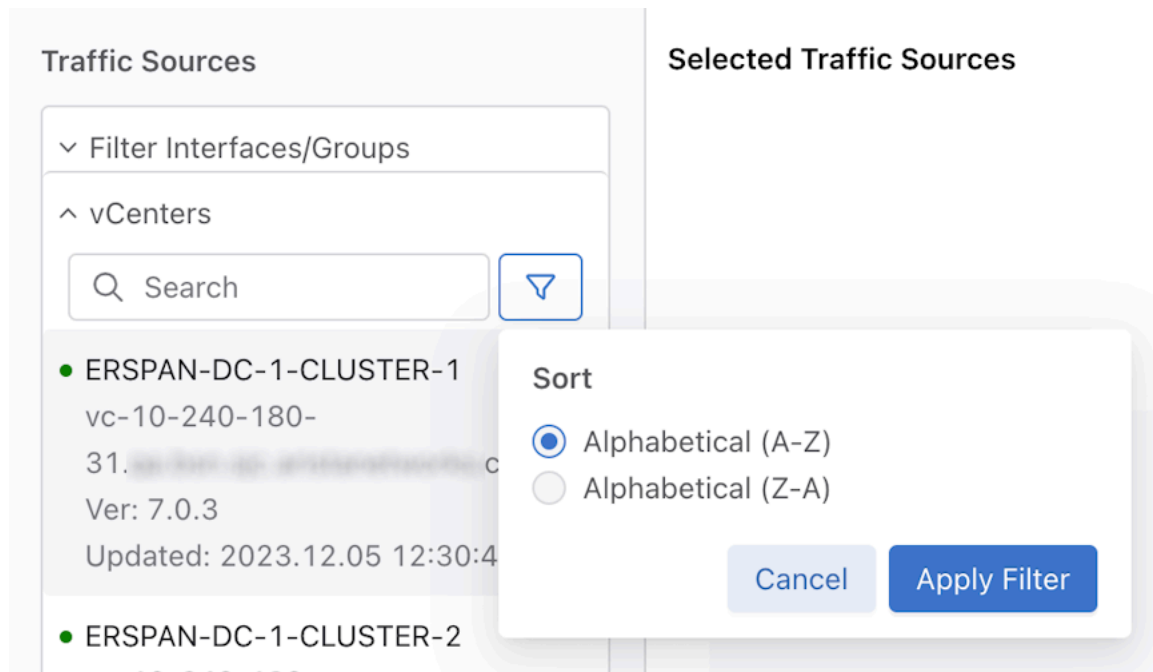
Interface Name(s)

Select

Cancel Apply Filter

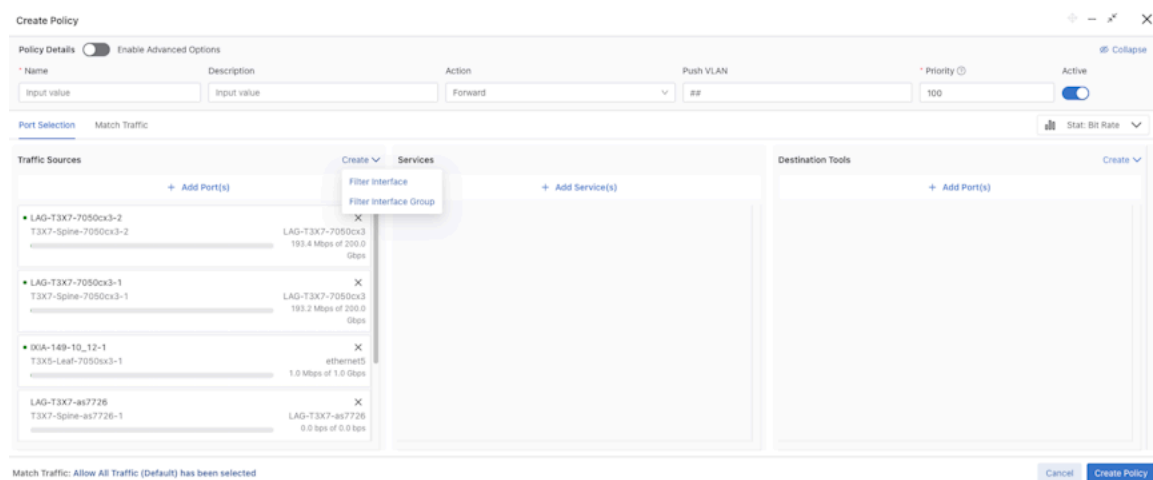
DMF sorts vCenters and CloudVision Portals alphabetically (A-Z, by default).

Figure 3-68: Sort Traffic Sources



Suppose a Filter Interface has not been created yet, the **Create** button has two selections: create **Filter Interfaces** and **Filter Interface Groups**.

Figure 3-69: Filter Interfaces / Filter Interface Groups



Clicking the **Create Filter Interface** button opens a form to configure a **Filter Interface**. Enter the required settings to configure the new **Filter Interface**.

Figure 3-70: Configure Filter Interface

The screenshot shows the 'Create Policy' window. On the left, under 'Connected Devices', there is a search bar and a list of devices including 'dmf-rack257-arista-1' and 'T2-Leaf-s4048-1'. On the right, the 'Create New Filter Interface' form is displayed. It includes fields for 'Filter Interface Name', 'Switch Name', 'Interface Name', 'Description (Optional)', 'IP Address', and 'Rewrite Dest. MAC Address'. The 'Create and Select' button is highlighted in blue.

Alternatively, the left column allows the selection of an existing connected device to pre-populate the **Switch Name** and **Interface Name** fields and to configure a **Filter Interface** based on a connected device. Click the **Create and Select** button to create the **Filter Interface** and associate it with the current policy.

Figure 3-71: Associate Filter Interface

The screenshot shows the 'Create Policy' window. On the left, under 'Connected Devices', there is a search bar and a list of devices. On the right, the 'Create New Filter Interface' form is displayed. The 'Connected Device Selected' section shows 'T2-Leaf-s4048-1'. The 'Create and Select' button is highlighted in blue.

To create multiple **Filter Interface(s)**, click the **Create another** button to create an interface using the current configuration. This action clears the form to allow the creation of an additional **Filter Interface**.

Figure 3-72: Add Multiple Filter Interfaces

The screenshot shows the 'Create Policy' dialog with two main panels. The left panel, titled 'Connected Devices', contains a search bar and a list of network devices including 'dmf-rack257-arista-1' and 'dmf-rack237-4x4200-1'. The right panel, titled 'Created DMF Interfaces', shows a 'test_daniel' interface. Below this is a 'Create New Filter Interface' form with fields for 'Filter Interface Name', 'Switch Name', 'Interface Name', 'Description (Optional)', 'IP Address', and 'Rewrite Dest. MAC Address'. There are also checkboxes for 'Analytics' and 'sFlow'.



Note: The **Select (n)** interface button associates all created **Filter Interfaces** to the current policy.

Click the **Create Filter Interface Group** button to create a group of filter interfaces.

Figure 3-73: Create Filter Interface Group

The screenshot shows the 'Create Policy' dialog with two main panels. The left panel, titled 'Filter Interfaces', contains a search bar and a list of network interfaces including 'IXIA-149-10-12-1', 'LAG-T2-s4048', and 'LAG-T3X5-65212F'. The right panel, titled 'Create Filter Interface Group', shows a 'Filter Interface Group Name' field and a large area with a 'Select or drag interfaces to begin' prompt.

Select one or more filter interfaces to create a **Filter Interface Group**.

Figure 3-74: Add Filter Interfaces

The screenshot shows the 'Create Policy' dialog with two main sections. On the left, the 'Filter Interfaces' section contains a list of available interfaces with search and filter options. On the right, the 'Create Filter Interface Group' section shows the 'Filter Interface Group Name' set to 'test_group' and a list of 'Selected Interfaces' including 'IXIA-149-10_12-3' and 'LAG-T3X5-s5212F'.

Filter Interfaces	Selected Interfaces
IXIA-149-10_12-1 T3X5-Leaf-7050cx3-1 1.0 Mbps of 1.0 Gbps ethernet5	IXIA-149-10_12-3 Switch Name T3X7-Leaf-s5248F-1 Interface Name ethernet1 Description -
IXIA-149-10_12-3 T3X7-Leaf-s5248F-1 0.0 bps of 1.0 Gbps ethernet1	LAG-T3X5-s5212F Switch Name T3X5-Leaf-s5212F-1 Interface Name LAG-T3X5-s5212F Description -
LAG-T2-s4048 T2-Leaf-s4048-1 261.0 Mbps of 20.0 Gbps LAG-T2-s4048	
LAG-T3X5-s5212F T3X5-Leaf-s5212F-1 255.8 Mbps of 50.0 Gbps LAG-T3X5-s5212F	
LAG-T3X7-7050cx3-1 T3X7-Spine-7050cx3-1 296.7 Mbps of 200.0 Gbps LAG-T3X7-7050cx3	
LAG-T3X7-7050cx3-2 T3X7-Spine-7050cx3-2 262.6 Mbps of 200.0 Gbps LAG-T3X7-7050cx3	
LAG-T3X7-as7326	

Click the **Create Group** button to create the **Filter Interface Group** and associate the group with the current policy.

Figure 3-75: Create Group

The screenshot shows the 'Create Policy' dialog with the 'Port Selection' and 'Match Traffic' sections. The 'Port Selection' section shows a list of 'Traffic Sources' including 'test_daniel' and 'test_group'. The 'Match Traffic' section shows a list of 'Services' and 'Destination Tools'. The 'Match Traffic' section also includes a 'Match Traffic' dropdown menu with 'Allow All Traffic (Default)' selected.

Traffic Sources	Services	Destination Tools
test_daniel T2-Leaf-s4048-1 0.0 bps of 0.0 bps ethernet48		
test_group Interface Group		

Expand the group tile to view interfaces within an **Interface Group**.

Figure 3-76: Expand Details

The screenshot shows the 'Create Policy' dialog with the 'Policy Details' tab selected. The 'Traffic Sources' section is expanded, showing a list of traffic sources. The first source is 'test_daniel' (T2-Leaf-s4049-1) with a description 'ethernet49' and a rate of '0.0 tps of 0.0 Gbps'. The second source is 'test_group' (Interface Group) with a description 'ethernet1' and a rate of '0.0 tps of 1.0 Gbps'. The 'test_group' is expanded, showing two sub-items: '00A-149-10_12-3' (T3X7-Leaf-s5248F-1) with a description 'ethernet1' and a rate of '0.0 tps of 1.0 Gbps', and 'LAG-T3X5-s5212F' (T3X5-Leaf-s5212F-1) with a description 'LAG-T3X5-s5212F' and a rate of '196.5 Mbps of 50.0 Gbps'. The 'Match Traffic' section shows 'Allow All Traffic (Default)' selected. The 'Create Policy' button is visible at the bottom right.



Note: Clicking the **x** icon on the top right of each tile disassociates the Filter Interface from the current policy. Clicking the **Undo** button restores the association.

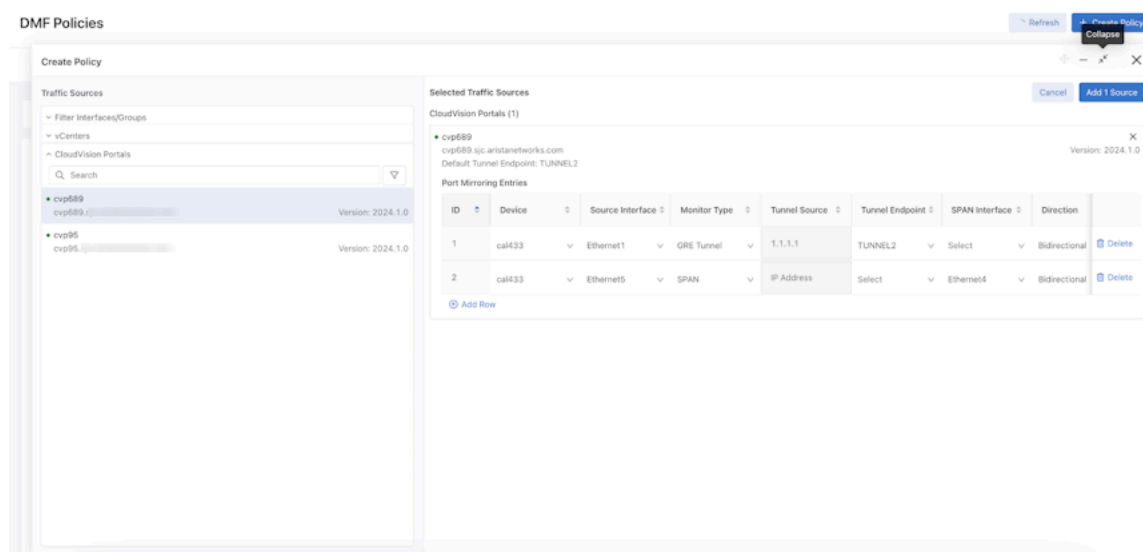
Figure 3-77: Disassociate Filter Interface

The screenshot shows the 'Create Policy' dialog with the 'Policy Details' tab selected. The 'Traffic Sources' section is expanded, showing a list of traffic sources. The first source, 'test_daniel', is now greyed out and has a 'Deleted test_daniel' message above it with an 'Undo' button. The second source is 'test_group' (Interface Group) with a description 'ethernet1' and a rate of '0.0 tps of 1.0 Gbps'. The 'test_group' is expanded, showing two sub-items: '00A-149-10_12-3' (T3X7-Leaf-s5248F-1) with a description 'ethernet1' and a rate of '0.0 tps of 1.0 Gbps', and 'LAG-T3X5-s5212F' (T3X5-Leaf-s5212F-1) with a description 'LAG-T3X5-s5212F' and a rate of '196.4 Mbps of 50.0 Gbps'. The 'Match Traffic' section shows 'Allow All Traffic (Default)' selected. The 'Create Policy' button is visible at the bottom right.

CloudVision Portals

The **Create Policy** window lists **CloudVision Portals** connected to DMF and includes the **CloudVision Portal** name, the portal hostname, and the current software version. Select a card to add a CloudVision **Port Mirroring Table**. The card displays similar information and the default Tunnel Endpoint.

Figure 3-78: CloudVision Portals



An empty port mirroring table initializes to add rows to the table for configuring port mirroring sessions.

Use the following guidelines to configure a port mirroring session:

- Each row must contain a **Device** and Source **Interface**. This interface in the CloudVision production network will mirror traffic to DMF.
- Each interface must select a **Monitor Type**: **GRE Tunnel** or **SPAN**.



Note: SPAN requires a physical connection from the CloudVision Portal to DMF. The default value for **Tunnel Endpoint** is the CloudVision Portal's Default Tunnel Endpoint.

- Each device must have the same **Tunnel Endpoint** and **Tunnel Source** values across the policies. Each interface on a device must have an identical destination configuration (**GRE Tunnel**, **GRE Tunnel Source**, and **SPAN Interface**) across the policies.
- The default traffic direction is **Bidirectional** but configurable to **Ingress** or **Egress**.
- After configuring the **Port Mirroring Table**, click **Add Sources** to return to the Main Page of the **Create Policy** configuration page.

Figure 3-79: Edit Policy

Edit Policy

Policy Details ☐ Enable Advanced Options

* Name Description Action

mm Input value Forward

Port Selection Match Traffic

Traffic Sources Create

+ Add Port(s)

cvp689
cvp689
cal433

Version: 2024.1.0
2 Entries

Services

+ Add Service(s)

After configuring **Port Mirroring**, the card appears in the **Traffic Sources** section. To edit the **Port Mirroring Table**, click the **X Entries** link.

3.4.1.3 Services

The **Services** column displays the **Services** and **Managed Services** associated with the policy. The **Add Service(s)** button opens a new page to specify additional services.

Figure 3-80: Services Add Services

Services

+ Add Service(s)

View All Services and **View All Managed Services** open the DMF **Services** and **Managed Services** pages, respectively. The **Add Service** button opens a configuration panel to specify **Service** information. If there are **Services** associated with this policy, they will be listed and available to edit.

Figure 3-81: View All Services / View All Managed Services

The screenshot shows a 'Create Policy' dialog box. At the top, there are navigation icons: a crosshair, a minus sign, a zoom-in icon, and a close 'X' button. Below the title bar, there is a section labeled 'Add Service' with two buttons: 'Cancel' and 'Add Service'. Underneath, there are two links: 'View All Services' and 'View All Managed Services'. The main area of the dialog is light gray and contains a central icon of a folder with a plus sign. Below the icon, the text 'Click here to add services' is displayed, followed by a blue button with a white plus sign and the text 'Add Service'.

For each **Service**, specify:

- **Service Type:** Managed or Unmanaged.
- **Service:** Name of the Service (required).
- **Optional:** Whether the Service is optional.
- **Backup Service:** Name of the backup Service.
- **Del. Service:** If the Managed Service type is selected, whether to use it as a Delivery Service.

Click the **Add Another** button to populate a new row to add another **Service**. The **Add (n) Services** button associates the Services with the policy.

Figure 3-82: Add Another Service

The screenshot shows the 'Create Policy' dialog box with the 'Add Service' section expanded. At the top, the same navigation icons are present. Below the title bar, the 'Add Service' section has 'Cancel' and 'Add 1 Service' buttons. Underneath are the links 'View All Services' and 'View All Managed Services'. The main area contains a table for 'Service 1' with a close 'X' button in the top right corner. The table has five columns: 'Service Type', '* Service', 'Optional', 'Backup Service', and 'Del. Service'. The 'Service Type' column has a dropdown menu with 'Managed Service' selected. The '* Service' column has a dropdown menu with 'MS-APP-ID-ACTIVE' selected. The 'Optional' column has a dropdown menu with 'Yes' selected. The 'Backup Service' column has an empty dropdown menu. The 'Del. Service' column has a dropdown menu with 'No' selected. Below the table, there is a blue button with a white plus sign and the text 'Add Another'.

After adding the services, they appear in the **Services** column. Click the **x** icon on the **Service** tile to disassociate the **Service** from the policy. While remaining on the page, if required, re-associate the Service by clicking **Undo**.

Figure 3-83: Service Added

Create Policy

Policy Details

Enable Advanced Options

Collapse

* Name

Description

Action

Push VLAN

* Priority ?

Active

Input value

Input value

Forward

##

100

Port Selection

Match Traffic

Stat: Bit Rate

Traffic Sources

Create

+ Add Port(s)

Services

+ Add Service(s)

Delete MS-APP-ID-ACTIVE

Undo

MS-SLICE-ACTIVE

Managed Service

MS-SLICE-BACKUP

Destination Tools

Create

+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

Cancel

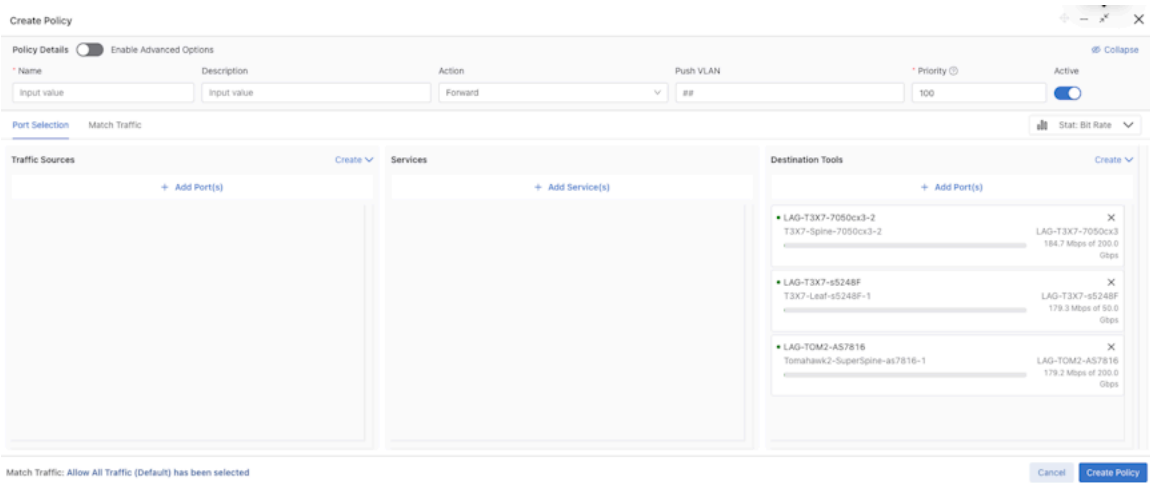
Create Policy

103

3.4.1.4 Destination Tools

The **Destination Tools** column displays the associated Destination Tool ports to a given policy.

Figure 3-84: Destination Tools



Use the **Add Port(s)** button to add more destinations. The configuration page allows adding **Delivery Interfaces/Groups** or **Recorder Node Fabric Interfaces**.


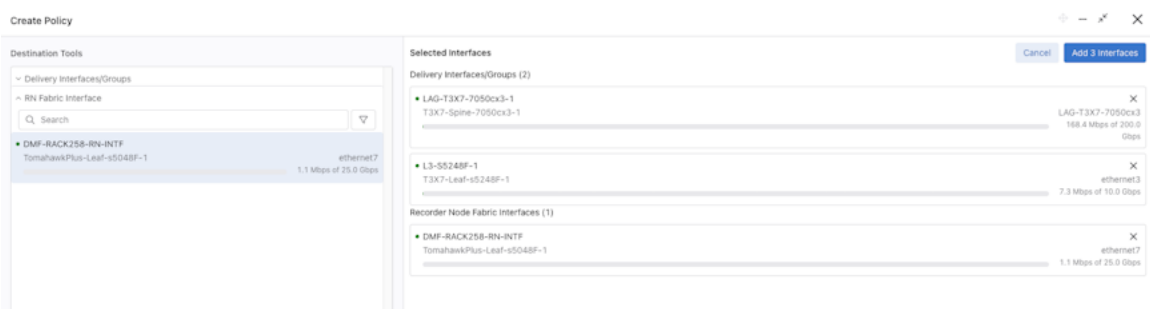

**Note:** The left column has two multiple groups. Select the corresponding type of **Destination Tools** to see the available selections. After making the desired selections, confirm using the **Add (n) Interfaces** button.

Figure 3-85: Add Interfaces



Interfaces can be searched by the available information in the interface tiles using the search bar. Clicking the  icon reveals sorting and filtering options using **Display Data**, which includes:

Sort - By default, DMF sorts the data in descending Bit Rate order. Optionally, sort the data by ascending Bit Rate order or alphabetically.

Bit Rate (default), **Utilization** percentage, or **Packet Rate**

Switch Name

Interface Name(s)**Figure 3-86: Filter Destination Tools**

The screenshot displays the 'Destination Tools' interface. On the left, under 'Delivery Interfaces/Groups', there is a search bar and a list of interfaces. The 'test_del_group' is expanded, showing several LAG interfaces with their respective bit rates and utilization bars. A modal window titled 'Display Data' is open in the foreground, allowing users to filter and sort the data.

Destination Tools

^ Delivery Interfaces/Groups

☐ Select All Delivery Interfaces

Search

LAG-T2-s4048
T2-Leaf-s4048-1 LAG-T2-
165.7 Mbps

test_del_group
Interface Group

LAG-TOM2-7260CX3
Tomahawk2-Spi... LAG-TOM2-
165.1 Mbps

LAG-T3X5-s5212F
T3X5-Leaf-s52... LAG-T3X5-
164.9 Mbps

LAG-T3X7-as7326
T3X7-Leaf-as73... LAG-T3X7-
163.9 Mbps

LAG-TOM-AS7712
Tomahawk-Spin... LAG-TOM-A
161.4 Mbps

Selected Interfaces

Display Data

Bit Rate Util% Packet Rate

Sort

☒ Bit Rate (Descending)
☐ Bit Rate (Ascending)
☐ Alphabetical (A-Z)
☐ Alphabetical (Z-A)

Filter

Bit Rate (Mbps)

0 20000

Switch Name

Select

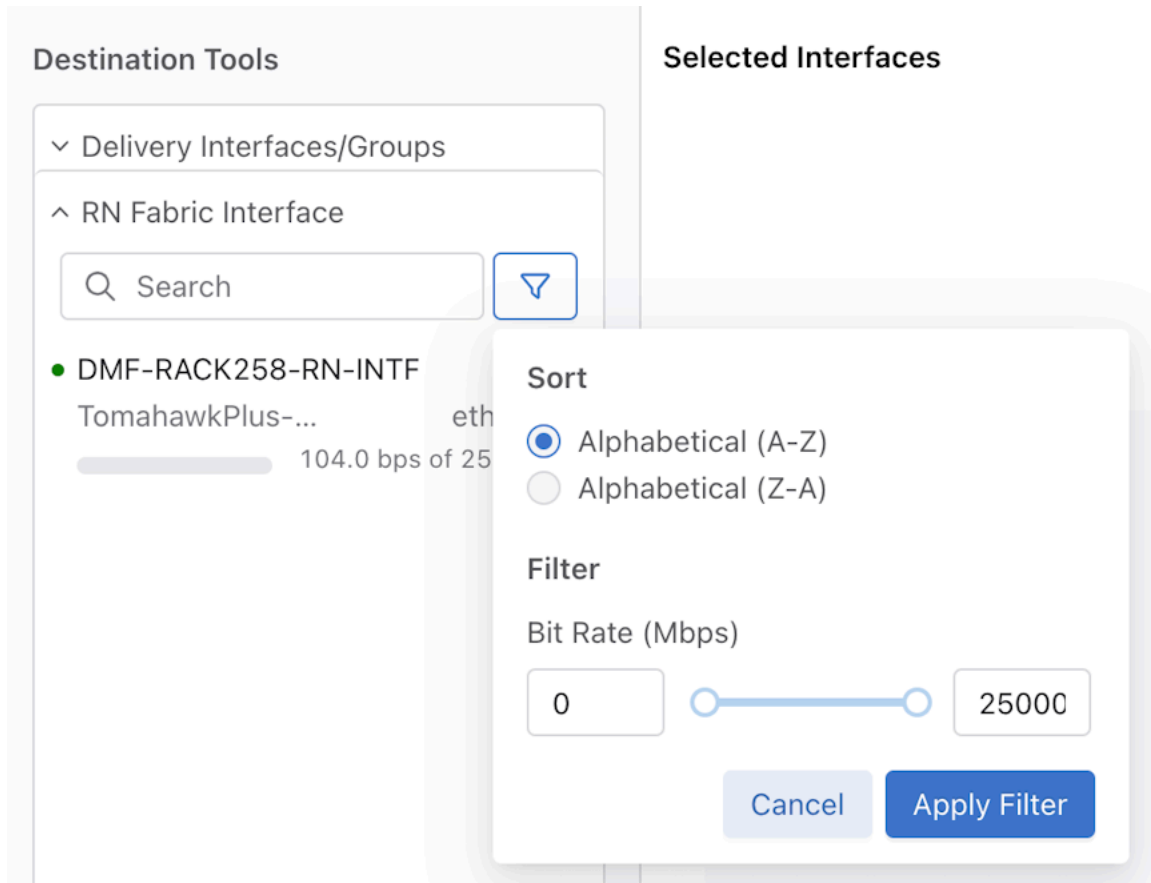
Interface Name(s)

Select

Cancel Apply Filter

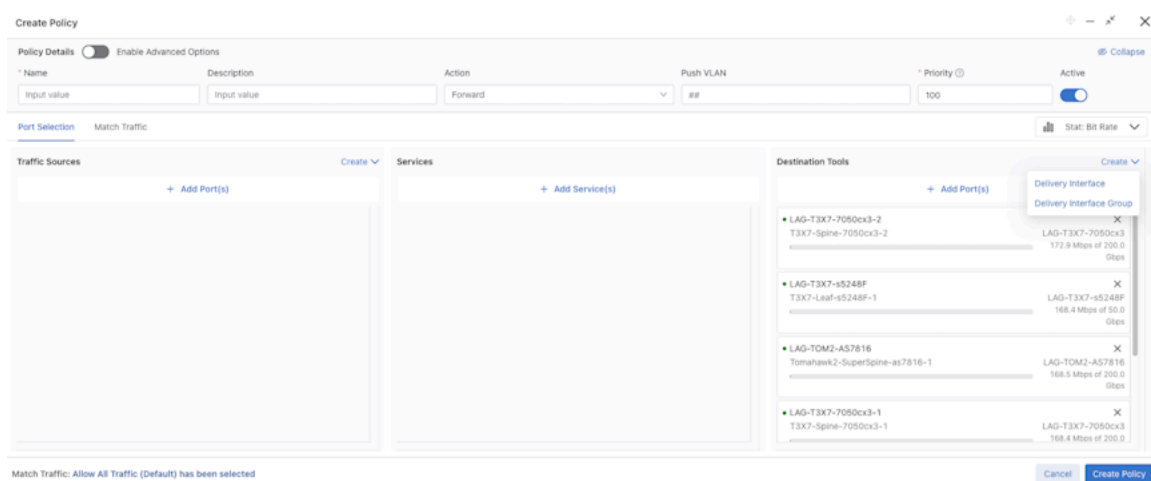
Sort Recorder Node Fabric Interfaces alphabetically (A-Z, by default) and filter by Bit Rate.

Figure 3-87: Sort Destination Tools



Suppose there is still a need to create **Destinations** (Delivery Interfaces). In that case, the **Create** button has two selections: create **Delivery Interfaces** and **Delivery Interface Groups**.

Figure 3-88: Create Delivery Interfaces / Delivery Interface Groups



Clicking the **Delivery Interface** button opens a form to configure a **Delivery Interface**. Enter the required settings to configure the new **Delivery Interface**.

Figure 3-89: Configure Delivery Interface

The screenshot shows the 'Create Policy' dialog box. On the left, under 'Connected Devices', there is a list of devices including 'dmf-rack257-arista-1' and 'dmf-rack237-ex4200-1'. The right panel, titled 'Create New Delivery Interface', contains the following fields:

- Delivery Interface Name**: A text input field.
- Switch Name**: A dropdown menu.
- Interface Name**: A dropdown menu.
- Description**: A text input field.
- IP Address**: A text input field with an 'IPv4 Address' icon.
- Strip VLAN on Egress**: A dropdown menu with 'Default' selected.
- Next Hop IP**: A text input field with an 'IPv4 Address' icon.
- Subnet Mask**: A text input field.

Alternatively, the left column allows the selection of an existing connected device to pre-populate the **Switch Name** and **Interface Name** fields and to configure a **Delivery Interface** based on a connected device. Click the **Create and Select** button to create the **Delivery Interface** and associate it with the current policy.

Figure 3-90: Associate Delivery Interface

The screenshot shows the 'Create Policy' dialog box. On the left, under 'Connected Devices', there is a list of devices including 'T3X5-Leaf-7050x3-1' and 'dmf-rack-258'. The right panel, titled 'Create New Delivery Interface', contains the following fields:

- Connected Device Selected**: A section showing the selected device 'T3X5-Leaf-7050x3-1' with details like 'Arista Networks EOS version 4.8.6 running on an Arista Networks DCS-7050S-64'.
- Delivery Interface Name**: A text input field.
- Switch Name**: A dropdown menu.
- Interface Name**: A dropdown menu.
- Description**: A text input field.
- IP Address**: A text input field with an 'IPv4 Address' icon.
- Strip VLAN on Egress**: A dropdown menu with 'Default' selected.
- Next Hop IP**: A text input field with an 'IPv4 Address' icon.
- Subnet Mask**: A text input field.

To create multiple **Delivery Interface(s)**, click the **Create another** button to create an interface using the current configuration. This action clears the form to allow the creation of an additional **Delivery Interface**.

Figure 3-91: Multiple Delivery Interfaces

The screenshot shows the 'Create Policy' dialog. On the left, under 'Connected Devices', there is a list of devices including T3X5-Leaf-7050x3-1, dmf-rack258, and others. The right pane is titled 'Description' and contains fields for 'Input Description', 'IP Address', 'Strip VLAN on Egress' (set to Default), 'Next Hop IP', 'Subnet Mask', 'ARP Interval' (set to 5 seconds), 'No VLAN Preservation' (checked), 'Preserve User Configured VLANs' (unchecked), and 'VLAN Preserve' (set to No VLAN Configured). At the bottom right, there are two buttons: 'Create another' and 'Create and Select'.

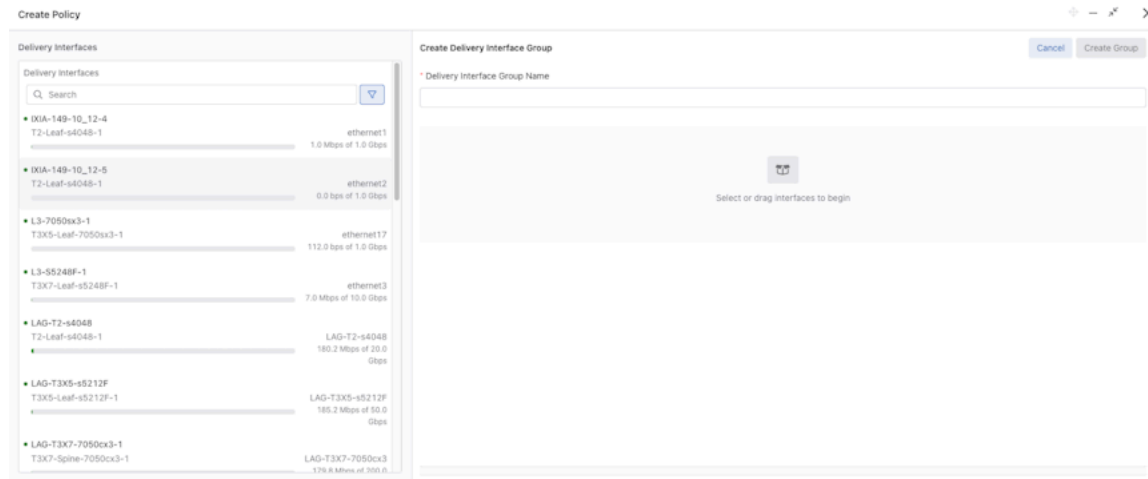
The **Select (n) interface** button associates all created **Delivery Interfaces** to the current policy.

Figure 3-92: Select Number of Interfaces & Associate

The screenshot shows the 'Create Policy' dialog. The left pane is the same as in Figure 3-91. The right pane is titled 'Created DMF Interfaces' and shows a list of interfaces. Below the list is a button that says 'Select a Connected Device from the left-hand panel to get started'. At the bottom right, there is a button labeled 'Select 1 Interface'.

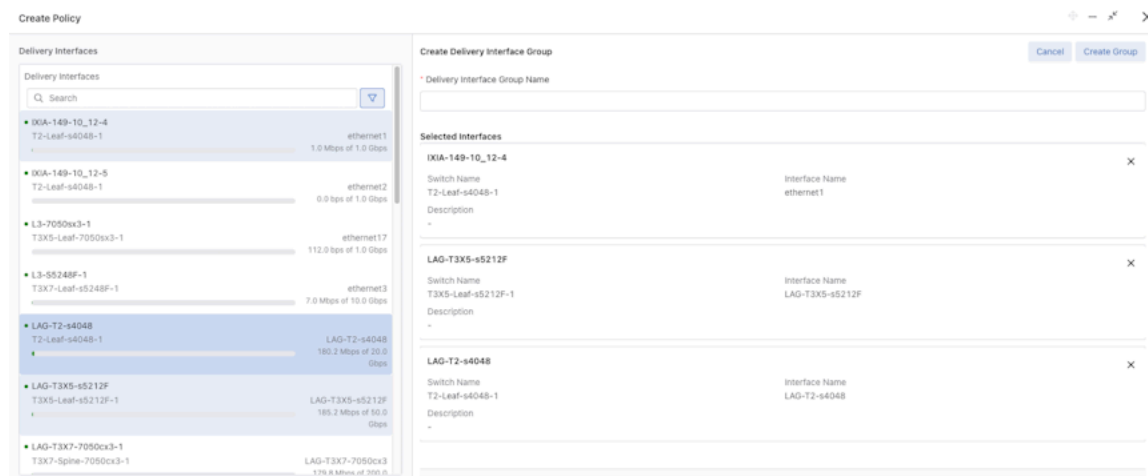
Click the **Create Delivery Interface Group** button to create a group of delivery interfaces.

Figure 3-93: Create Delivery Interface Group



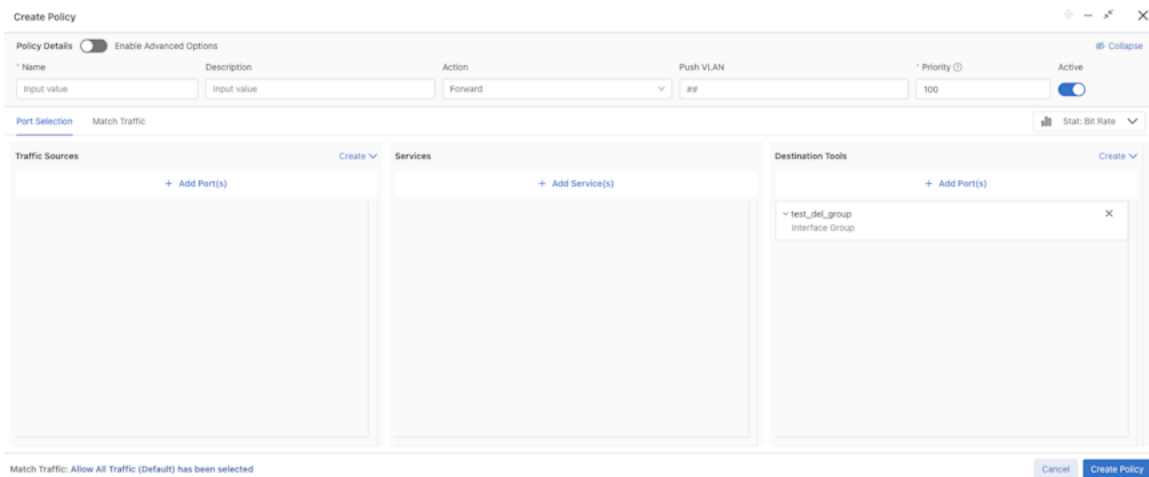
Select one or more delivery interfaces to create a **Delivery Interface Group**.

Figure 3-94: Multiple Delivery Interfaces



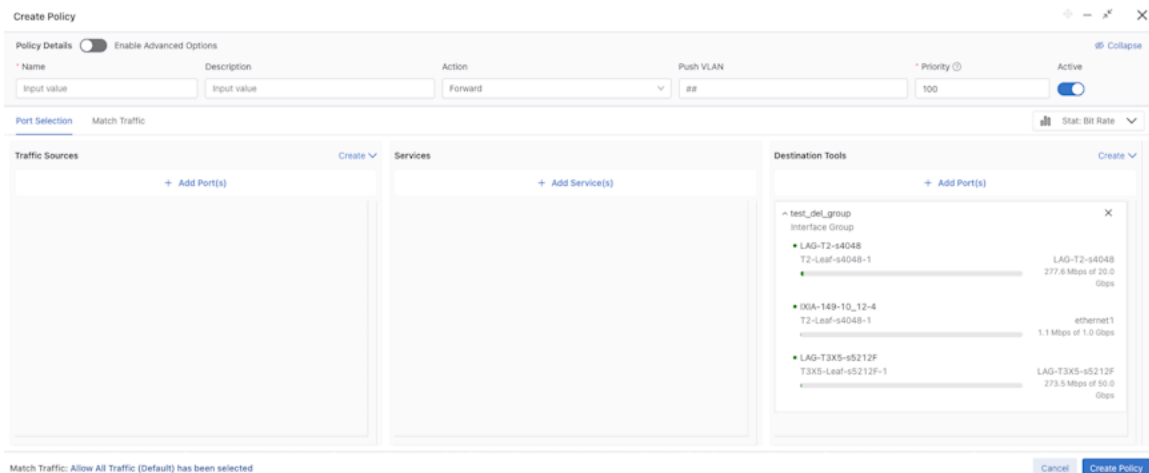
Click the **Create Group** button to create the **Delivery Interface Group** and associate the group with the current policy.

Figure 3-95: Associate Delivery Interface Group



Expand the group tile to view interfaces within an **Interface Group**.


Figure 3-96: Expand Details



3.4.1.5 Stat Picker

Use the **Stat: Packet Rate** drop-down to select view specific data for the associated interfaces.

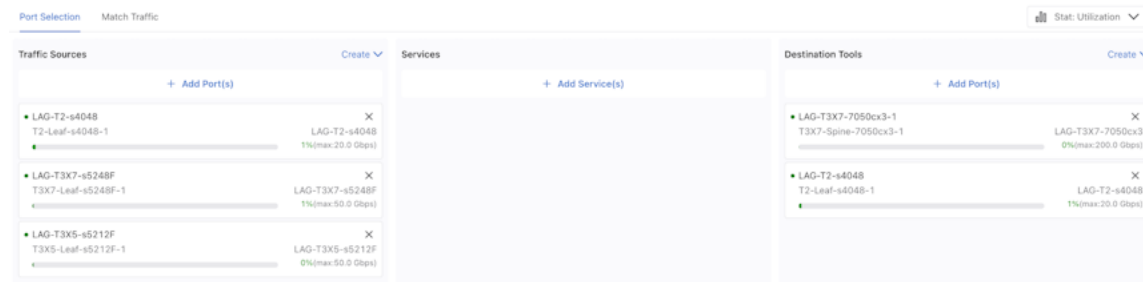
Figure 3-97: None



The data options are:

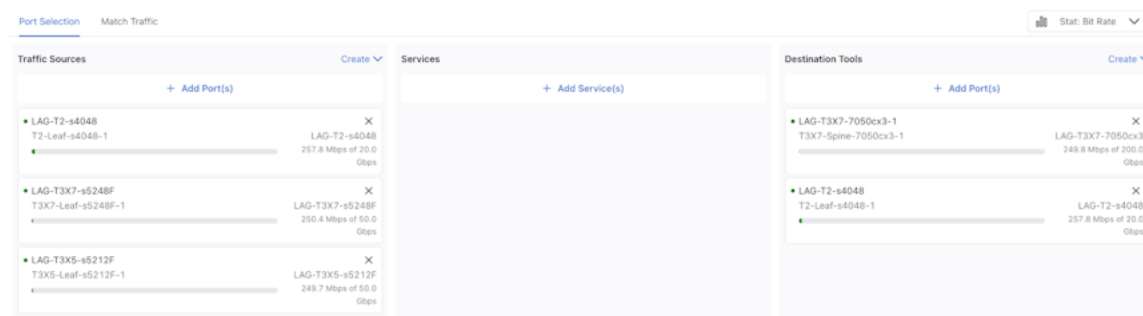
Utilization

Figure 3-98: Utilization



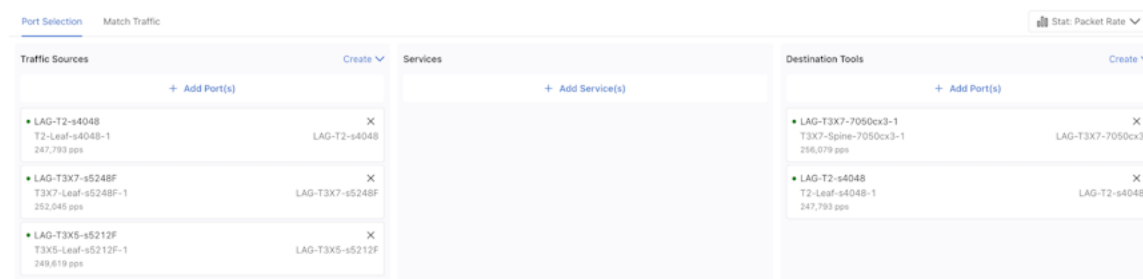
Bit Rate (default)

Figure 3-99: Bit Rate



Packet Rate

Figure 3-100: Packet Rate



3.4.1.6 Match Traffic and Match Traffic Rules

Match Traffic

Use the **Match Traffic** tab to configure rules for the current policy.

Figure 3-101: Match Traffic

The screenshot shows a web interface for configuring a policy. At the top, there are two tabs: "Port Selection" and "Match Traffic", with "Match Traffic" being the active tab. Below the tabs is a large light gray box containing a server icon, the text "Select an option below or configure a custom rule.", and two radio buttons: "Allow All Traffic" (which is selected) and "Deny All Traffic". Below these is the word "or" and a blue link "Configure A Rule". Underneath the gray box is a button with an upward arrow icon and the text "Import Rules". At the bottom of the interface, there is a status message "Match Traffic: No Selection has been added" on the left, and two buttons, "Cancel" and "Create Policy", on the right.

Port Selection **Match Traffic**

Select an option below or configure a custom rule.

☒ Allow All Traffic ☐ Deny All Traffic

or

[Configure A Rule](#)

Import Rules

Match Traffic: No Selection has been added

[Cancel](#) [Create Policy](#)

There are four options to configure traffic rules.

Figure 3-102: Configuration Options

Create Policy

Policy Details

☐ Enable Advanced Options

Collapse

Name	Description	Action	Push VLAN	Priority	Active
Input value	Input value	Forward	##	100	<input checked="" type="checkbox"/>

Port Selection

Match Traffic

Select an option below or configure a custom rule.

☒ Allow All Traffic
☐ Deny All Traffic

or

Configure A Rule

Import Rules

Match Traffic: No Selection has been added

Cancel

Create Policy

Select the **Allow All Traffic** or **Deny All Traffic** radio button to quickly configure a rule for all traffic.

Navigate to the **Rule Details** configuration panel using the **Configure A Rule** button. Refer to the Custom Rule, Match Rule Shortcut, and Match Rule Group sections for more information.

The **Import Rules** button opens the import rule configuration dialog and supports importing `.txt` files using drag and drop or **Browse**.

Example Text File

```
1 match ip
2 match tcp
3 match tcp src-port 80
```

```
4 match tcp dst-port 25
```

Figure 3-103: Import Rules

Import Rule
Drag a txt file into the space below to import rules.

Drop files here
Supported files: .txt

[Browse](#)

✓ rules_2023-11-02T16...

[Cancel](#) [Preview](#)

Click the **Preview** button to verify the import result.

Figure 3-104: Preview

Preview Imported Rule
Review the imported rules before importing

Rule ID	EtherType	IP Protocol	Source Ports	Dest. Ports	VLAN			
1	ARP (2054)						Edit	
2	IPv4 (2048)						Edit	

< 1 >

[Cancel](#) [Import 2 Rules](#)

While using the Preview Imported Rule table, click the **Edit** button to open the **Edit Rule** configuration panel.

Figure 3-105: Edit Rule

Edit Rule

EtherType

IP Protocol ⓘ

TCP Flags

IP DSCP ⓘ

IP ECN ⓘ

IP Fragment

VLAN

Bidirection

ARP (2054) ▾

None ▾

Select Flags ▾

Unspecified ▾

Any ▾

On

Additional Configurations ▾

[Back to List](#) [Confirm](#)

Click the **Confirm** button when finished, and use the **Import x Rules** button to import the rules.

Custom Rule

Click the **Configure a Rule** button to open the **Configure A Traffic Rule** screen.

Figure 3-106: Configure a Traffic rule

Configure A Traffic Rule Cancel Add Rule

Custom Rule ▾

EtherType IP Protocol ⓘ TCP Flags IP DSCP ⓘ

None ▾ None ▾ Select Flags ▾

IP ECN ⓘ IP Fragment VLAN Bidirectional

Unspecified ▾ Any ▾ On Off

[Additional Configurations ▾](#)

By default, the configuration method is **Custom Rule** with several fields disabled by default; hover over the **question mark** icon for more information on enabling an input field.

Figure 3-107: Help Icon

EtherType must be set to IPv4 or IPv6

IP Protocol ⓘ

None ▾

Specific **EtherTypes** will open an **Additional Configurations** panel.

Figure 3-108: Additional Configurations

[Additional Configurations ▾](#)

Source IP Address Ports ⓘ MAC Address ⓘ MAC Mask ⓘ

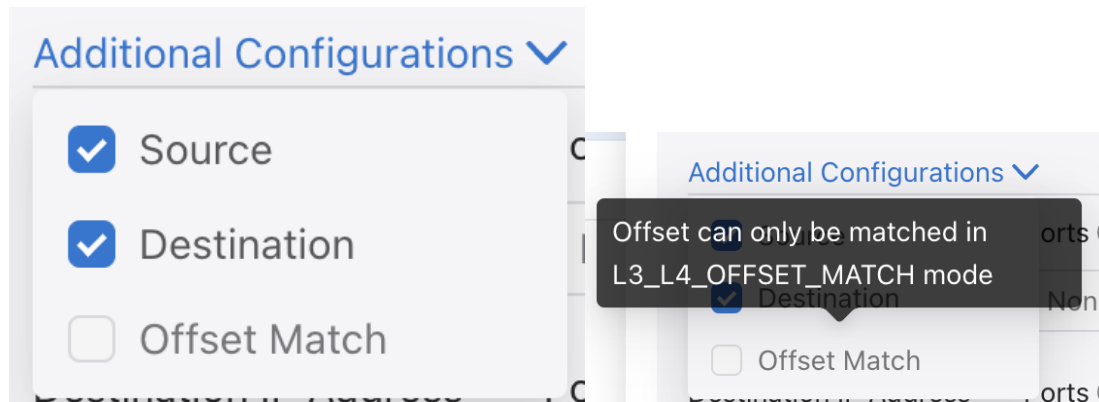
None ▾ None ▾ - Address - - Mask -

Destination IP Address Ports ⓘ MAC Address ⓘ MAC Mask ⓘ

None ▾ None ▾ - Address - - Mask -

Click the drop-down icon to display additional configurations (**Source**, **Destination**, **Offset Match**). Hovering over **Offset Match** allows viewing requirements to enable the **Offset Match**.

Figure 3-109: Offset Match



Match Rule Shortcut

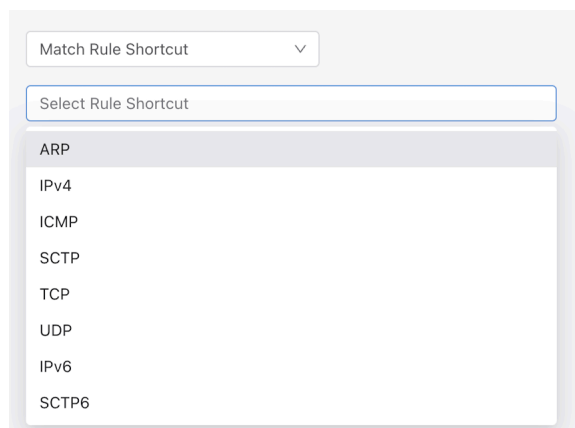
To access the **Match Rule Shortcut**, click the drop-down button and select **Match Rule Shortcut**.

Figure 3-110: Match Rule Shortcut



Click the **Select Rule Shortcut** selector and choose the required shortcut rules (supports multi-selection).

Figure 3-111: Shortcut Rule List



After selecting the rule shortcut:

- All selected rules appear as a card in the selector.
- Delete selected rules using the **x** icons.

- Click the **Customize Shortcut** button to edit a rule shortcut.

Figure 3-112: Edit Shortcut

The screenshot shows the 'Edit Shortcut' interface. At the top, there is a dropdown menu labeled 'Match Rule Shortcut' with a downward arrow. Below it, there are two tabs: 'ARP' and 'IPv4', both with close buttons (X). The main area contains two identical configuration sections. Each section has a 'Rule Shortcut' dropdown (set to 'ARP' and 'IPv4' respectively) and a 'Customize Shortcut' link. To the right of each section are three dropdown menus: 'EtherType' (set to 'ARP (2054)' and 'IPv4 (2048)' respectively) and 'IP Protocol' (set to 'None' for both). A blue 'X' icon is visible to the right of each configuration section.

After editing, click the **Save Edit** button to return to the **Match Rule Shortcut** view.

Figure 3-113: Save Edits

The screenshot shows the 'Customize Shortcut' interface. At the top, there is a 'Customize Shortcut' label and two buttons: 'Back' and 'Save Edit'. The main area contains several configuration options: 'EtherType' (set to 'ARP (2054)'), 'IP Protocol' (set to 'None'), 'TCP Flags' (set to 'Select Flags'), 'IP DSCP' (empty), 'IP ECN' (empty), 'IP Fragment' (set to 'Unspecified'), 'VLAN' (set to 'Any'), and 'Bidirectional' (with 'On' and 'Off' buttons). At the bottom, there is a link for 'Additional Configurations' with a downward arrow.

After configuring the shortcut rules, click the **Add (n) Rules** button to finish the configuration.

The screenshot shows a button labeled 'Add 2 Rules' in blue, next to a 'Cancel' button in light blue. Above the buttons, the text 'Configure A Traffic Rule' is visible.

Match Rule Group

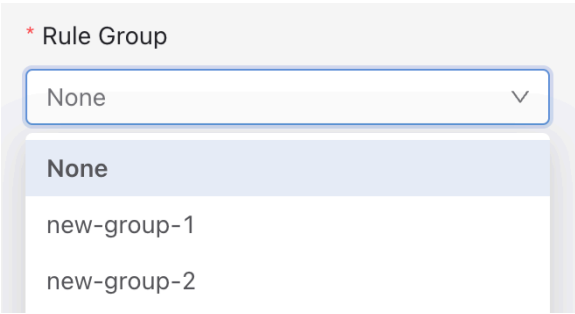
To access the **Match Rule Group**, click the drop-down button and select **Match Rule Group**.

Figure 3-114: Match Rule Group

The screenshot shows a dropdown menu with the following options: 'Match Rule Group' (selected and highlighted in blue), 'Custom Rule', and 'Match Rule Shortcut'.

To select a rule group, click the drop-down button under **Rule Group**. All rule groups appear in the menu. Select one. There is no multi-select available. Repeat the **Match Rule Group** steps to add more than one rule group.

Figure 3-115: Rule Group List



After configuring the rule group, click the **Add Rule** button to finish the configuration.

Figure 3-116: Add Rule



3.4.1.7 Rules Table

All configured rules appear in the **Rules Table**.

Figure 3-117: Rules Table

Import Rules		Export Select Rules		Delete		+ Create New Rule	
<input type="checkbox"/>	Rule ID	EtherType	IP Protocol	Sou			Q
<input type="checkbox"/>	1	IPv4 (2048)			Edit	Delete	
<input type="checkbox"/>	2	IPv4 (2048)	ICMP (1)		Edit	Delete	
<input checked="" type="checkbox"/>	3	new-group-1			Edit	Delete	

• Import Rules

Figure 3-118: Import Rules



- Similar in function to the Import Rules button on the start page. Refer to [Start Page -> Import Rules](#) for more information.
- **Export Select Rules**

Figure 3-119: Export Select Rules



- Disabled by default when no rule is selected.
- Enabled when one or more than one rule is selected.
- Click to export selected rules information as a `.txt` file.
- **Delete**

Figure 3-120: Delete

- Disabled by default when no rule is selected.
- Enabled when one or more than one rule is selected.
- Click to delete the selected rules.
- **Create New Rule** and **Create Rule Group** buttons

Figure 3-121: Create New Rule / Create Rule Group

- The button will appear as **Create New Rule** when no rule is selected. Click to open the **Create New Rule** screen.
- When one or more rules are selected, the button changes to **Create Rule Group**. Click to open the **Create Rule Group** screen.

Figure 3-122: Create Rule Group

Create Rule Group

* Rule Group Name

Description

The **Rule Group** Name is required. Click **Create Group** to confirm the rule group creation.

- **Table Actions**

Figure 3-123: Edit / Delete

- Click the **Edit** button to edit the rule view.
- Click the **Delete** button to delete the rule.

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

2. Configure one or more match rules to identify the aggregated traffic 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 matching rule (**10**) selects IP traffic with a destination address **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 the **config-switch-if** submode.

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



Note: Define the filter interfaces used 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 monitor fabric edge ports connected to destination tools and defined using the **interface** command from the **config-switch-if** submode.

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

Define the delivery interfaces used in the policy 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.
- When inserting a service 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 the same priority with similar match conditions, then incorrect statistics can result 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 reaching its total capacity. When this happens, the detailed status shows a message like **Table ing_flow2 is full <switch_DPID>**.

- **ing_flow1**- 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 deleting a policy, DMF deletes all traffic rules associated with the policy.

3.4.3 Define Out-of-band Match Rules Using the CLI

A policy can contain multiple match rules, each assigned a rule number. However, the rule number does not specify a priority or the sequence in applying the match rule to traffic entering the filter ports included in a policy. Instead, if the traffic matches any 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, DANZ Monitoring Fabric (DMF) calculates the change in policy flows and only sends the difference 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 matched packet. For details, refer to the chapter [Advanced Policy Configuration](#).

3.4.4 Stop, Start, and Schedule a Policy Using the CLI

Enter the **active** or **inactive** command from the **config-policy** submode to enable or disable a policy.

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 action prevents a capture from running indefinitely and utilizes the appliance storage capacity.

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**: start the action immediately.
- **delivery-count**: runs until the specified number of packets are delivered to all delivery interfaces.
- **seconds**: start the action after waiting for the specified number of seconds. For example, **300+** start 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**: DANZ Monitoring Fabric (DMF) assigns **60** seconds by default if no duration is specified. A value of **0** causes the action to run until it is stopped manually. When using the **delivery-count** keyword 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+
```

3.4.5 Clear a Policy Using the CLI

To remove a specific DANZ Monitoring Fabric (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
```

3.4.6 View Policies Using the CLI

To display the policies currently configured in the DANZ Monitoring Fabric (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 policy configuration 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 inserting a service 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 ] ]
```

Use the event history to determine the last time when policy flows were installed or removed. A value of **dynamic** for Type indicates the policy was dynamically created for overlapping policies.

3.4.7 Rename a Policy Using the CLI

Policy Renaming Procedure



Note: A DANZ Monitoring Fabric (DMF) policy must exist to use the renaming feature.

Use the following procedure to rename an existing policy.

1. Use the CLI command **policy *existing-policy-name*** to enter the submode of an existing policy and then enter the **show this** command.

```
dmf-controller-1(config)# policy existing-policy-name
dmf-controller-1(config-policy)# show this
! policy
policy existing-policy-name
```

2. Enter the **rename** command with the new policy name, as shown in the following example.

```
dmf-controller-1(config-policy)# rename new-policy-name
```



Note: Possible traffic loss may occur when renaming a policy.

3. Verify the policy name change using the **show this** command.

```
dmf-controller-1(config-policy)# show this
! policy
policy new-policy-name
dmf-controller-1(config-policy)#
```



Note: A user must have permission to update the policy. The new policy name must follow the requirements for a policy name.

3.5 Using the Packet Capture Action in a Policy

Capture packets into a PCAP file for later processing or analysis. DANZ Monitoring Fabric (DMF) stores the captured packets on the DMF Controller hardware appliance. This feature provides 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, as running the Controller in a virtual machine is impossible. 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 3-128: DMF Controller Hardware Appliance


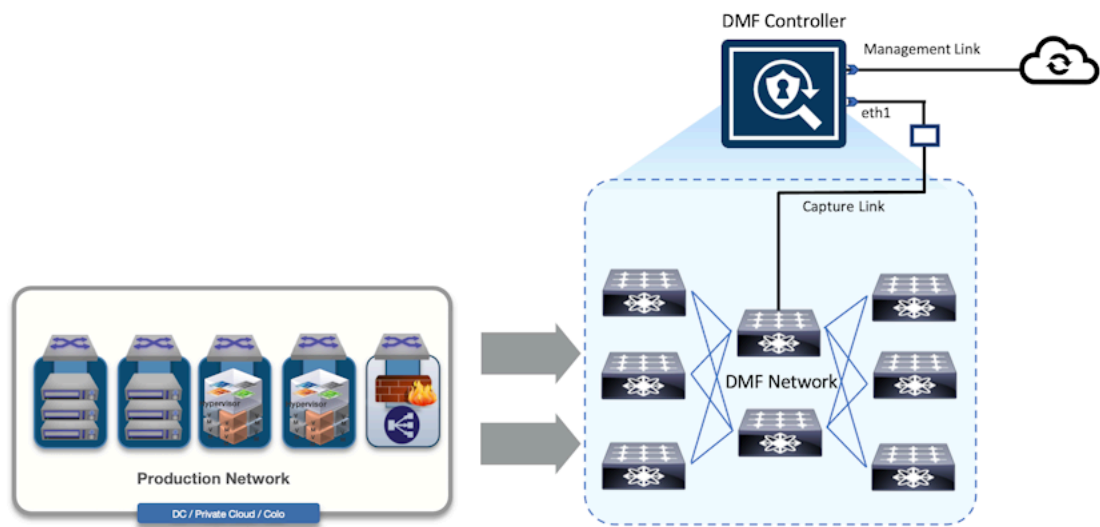


Table 2:

1	1G Management Port	2	10G Management Port
---	--------------------	---	---------------------

Figure 3-129: 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 exists in the policy to modify the packets before capture, but this is optional.

By default, when the policy action is **capture**, the policy is only active after scheduling the policy. Packet captures are always saved on the master (active) Controller. In case of HA failover, previous packet captures remain on the Controller where they were initially saved.

By default, DMF automatically removes PCAP files after seven days. Change the default value using the following CLI command with the command option if preferred:

```
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)#
```

3.6 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. Delete the files affiliated with a specific policy, as shown 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 action 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
# 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/2022-11-01-03-19-cl06efc.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)#
```

Viewing Information about Monitoring Fabric and Production Networks

This chapter describes to view information about the DANZ Monitoring Fabric (DMF) and connected production networks.


4.1 Monitoring DMF Interfaces

Monitoring DANZ Monitoring Fabric (DMF) Interfaces

4.1.1 Using the GUI to Monitor DMF Interfaces

Click the **Menu** control to view statistics for the specific interface and select **Monitor Stats**. The system displays the following dialog box.


Figure 4-1: Monitor Interface Stats



This window displays statistics for up to four selected interfaces and provides a line graph (sparkline) that indicates changes in packet rate or bandwidth utilization. The auto-refresh rate for these statistics is ten seconds—mouse over the sparkline to view the range of values represented. To clear statistics for an interface, click the **Menu** control and select **Clear Stats**.

To view statistics for multiple interfaces, enable the checkbox to the left of the **Menu** control for each interface, click the **Menu** control in table, and select **Monitor Selected Stats**.

Figure 4-2: Monitoring > Interfaces



To view the interfaces assigned a specific role, use the **Monitoring > Interfaces** command and select the **Filter**, **Delivery**, or **Service** sub-option from the menu.

4.1.2 Viewing Oversubscription Statistics

To view peak bit rate statistics used to monitor bandwidth utilization due to oversubscription, click the **Menu** in the **Interfaces** table, select **Show/Hide Columns**, and enable the **Peak Bit Rate** checkbox on the dialog box that appears.

After enabling the **Peak Bit Rate** column, a column appears in the **Interfaces** table that indicates the relative bandwidth utilization of each interface. When using less than **50%** of the bandwidth, the bar appears in green; **50-75%** changes the bar to yellow, and over **75%** switches 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: DANZ Monitoring Fabric (DMF) Controllers generate SNMP traps for link saturation and packet loss. For more information, please refer to the DMF 8.6 Deployment Guide - SNMP Trap Generation for Packet Drops and Link Saturation chapter.

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 DANZ Monitoring Fabric (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#

```

4.2 Viewing Devices Connected to the Monitoring Fabric

4.2.1 Using the GUI to View Fabric-Connected Devices

To view a display of the devices connected to the Controller, select **Fabric > Connected Devices** from the main menu. The system displays the following screen.

Figure 4-3: Connected Devices

The screenshot shows the 'Connected Devices' page in a web application. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. The page title is 'Connected Devices' with a 'Refresh' button. Below the title is a subtitle: 'List of all connected devices in the production network'. The main content area is divided into two sections: 'Unique Device Names' and 'Switch Interfaces'. The 'Unique Device Names' section shows a list of device names, with one example: 'app-accton-tb2-arista-1.qa.bsn.sjc...' (20). The 'Switch Interfaces' section features a search bar with a dropdown arrow and a search icon, and a table with 7 columns: 'DMF Switch', 'DMF Interface', 'Device Name', 'Device Description', 'Chassis ID', 'Port ID', and 'Port Description'. The table contains 10 rows of data, all showing 'app-accton-tb2-arista-1.qa.bsn.sjc...' as the device name and 'Arista Networks EOS version 4.1...' as the device description. The bottom of the page shows a pagination bar: '1-10 of 20 items' with a page number '1' and a '10 / page' dropdown.

DMF Switch	DMF Interface	Device Name	Device Description	Chassis ID	Port ID	Port Description
app-accton-as6812-1	ethernet22/1	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet33	-
app-accton-as6812-1	ethernet22/2	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet34	-
app-accton-as6812-1	ethernet22/3	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet35	-
app-accton-as6812-1	ethernet22/4	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet36	-
app-accton-as7312-1	ethernet10	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet14	-
app-accton-as7312-1	ethernet7	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet11	-
app-accton-as7312-1	ethernet8	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet12	-
app-accton-as7312-1	ethernet9	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet13	-
app-accton-as6712-1	ethernet11/1	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet29	-
app-accton-as6712-1	ethernet11/2	app-accton-tb2-arista-1.qa.bsn.sjc...	Arista Networks EOS version 4.1...	00:1c:73:18:bf:5a	Ethernet30	-

The **Switch Interfaces** table displays the unique devices connected to each out-of-band filter or delivery switch. It lists each interface's MAC address (Chassis ID) on every device connected to the fabric as a separate device.

The **Unique Device Names** table lists all unique device names with a count of interfaces in parentheses. Clicking a row in this list filters the contents of the **Switch Interfaces** table.

To view a display of the devices discovered by the Controller through the Link Aggregation Control Protocol (LACP), select **Fabric > Connected LACP** from the main menu.

The system displays the following screen.

Figure 4-4: Connected LACP

Fabric Monitoring Maintenance Integration Security									
2 warnings DMF-8.0.0 DANZ Monitoring Fabric - Connected LACP									
Connected LACP									
Filter table rows									
Switch	Interface	DMF Name	Device MAC	Port ID	Link Group Key	Partner Device MAC	Partner Port ID	Partner Link Group Key	
leaf-s4112F-1	ethernet11	—	80	4	1	00:00:00:00:00:00	4	1	
leaf-s4112F-1	ethernet9	—	06	13	4	00:00:00:00:00:00	0	0	
leaf-s4112T-1	ethernet10	—	80	7	1	00:00:00:00:00:00	7	1	
leaf-s4112T-1	ethernet9	—	80	6	1	00:00:00:00:00:00	6	1	

Nov 14, 2020, 3:15:44am UTC Show: 10 25 100 All (1 - 4 / 4)

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

4.2.2 Using the CLI to View Switch Configuration

To verify the switch interface configuration, enter the **show topology** command, as shown 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, DANZ Monitoring Fabric (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. View this information on the GUI page under **Monitoring > Host Tracker**. As shown below, the output displays 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 several 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: DNS traffic will not be included in DMF policies when enabling DNS for tracking.

Exclude host tracking for a specific filter interface 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, discovered using LLDP and CDP, can be seen using the **show connected-devices** command. The data helps determine if filter interfaces are connected to the intended production device.

The **show connected-devices** command from login mode displays the devices connected to the DANZ Monitoring Fabric (DMF). This command displays information about devices connected to DMF switch interfaces. DMF extracts the information from link-level protocol packets such as LLDP, CDP, and UDLD and ignores expired link-level data.

```
controller-1# show connected-devices switch FILTER-SWITCH-1
# Switch
# IF Name DMF name SPAN? Device Name Device Description Chassis ID Port ID Port Description Management Address Protocol
-----
1 FILTER-SWITCH-1 ethernet3 TOOL-1 False def-rack257-server-4 Ubuntu 18.04.6 LTS Linux 4.15.0-188-generic #199-Ubuntu SMP Wed :1e d4:ae:52:c6:b7:1f eth3 fe80:: 9 LLDP
2 FILTER-SWITCH-1 ethernet4 False def-rack257-server-4 Ubuntu 18.04.6 LTS Linux 4.15.0-188-generic #199-Ubuntu SMP Wed :1e 90:e2:ba:28:ea:a5 eth5 fe80:: 9 LLDP
3 FILTER-SWITCH-1 ethernet5 False def-rack258-arista-2 Arista Networks EOS version 4.10.6 running on an Arista Network :83 ethernet27 10.24 LLDP
4 FILTER-SWITCH-1 ethernet7 False def-rack258-arista-2 Arista Networks EOS version 4.10.6 running on an Arista Network :83 ethernet28 10.24 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
----- Selected Connected Device -----
Device Name      : DMF-RECORDER-NODE
Device Description : dmf-recorder-node, SN 33R8XQ2
Connection State  : connected
Last Connected Switch : DELIVERY-SWITCH-1
Last Connected Interface : ethernet11
Chassis ID       : 10
Port ID          : f8:f2:1e:00:0b:b0
Port Description  : ens1f0
Management Address : 10.240.180.222
Protocol         : LLDP

----- Connected Device Events -----
# Time of Event      Type of Event      Switch      IF Name      Device Name      Device Description      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 33R8XQ2 10 f8:f2:1e:00:0b:b0 ens1f0 10.2 LLDP
2 2022-11-02 12:49:56.936000 PDT device-removed DELIVERY-SWITCH-1 ethernet11 DMF-RECORDER-NODE dmf-recorder-node, SN 33R8XQ2 10 f8:f2:1e:00:0b:b0 ens1f0 10.2 LLDP
3 2022-11-02 12:45:56.863000 PDT device-added DELIVERY-SWITCH-1 ethernet11 DMF-RECORDER-NODE dmf-recorder-node, SN 33R8XQ2 10 f8:f2:1e:00:0b:b0 ens1f0 10.2 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
----- Selected Connected Device -----
# 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 VShare      ESXi Releasebuild-18426014 disconnected DELIVERY-SWITCH-1 NETWORK-TAP-1 11 Vb38 TO-RELEASE-SETUP 10.240.180.143 LLDP
2 2022-11-02 12:52:24.736000 PDT VC7-ESX-5 VShare      ESXi Releasebuild-18426014 disconnected DELIVERY-SWITCH-1 ethernet11 10 f8:f2:1e:00:0b:b0 0 LLDP
3 2022-11-02 12:46:59.884000 PDT VC7-ESX-6 VShare      ESXi Releasebuild-18426014 disconnected DELIVERY-SWITCH-2 ethernet3 01 f8:f2:1e:0f:08:01 0 LLDP

controller-1#

```

Connecting a DMF switch interface to a SPAN port may result in inaccurate information 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. DMF extracts the information from LACP protocol packets and expired LACP information is ignored as illustrated in the following.

```

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      00:00:00:00:00:00 01 23      10      00:00:00:00:00:00 0 0
2 DMF-SWITCH-1 ethernet14     00:00:00:00:00:00 01 23      10      00:00:00:00:00:00 0 0
3 DMF-SWITCH-1 ethernet16     00:00:00:00:00:00 01 42      3       00:00:00:00:00:00 0 0
4 DMF-SWITCH-1 ethernet2      00:00:00:00:00:00 01 42      3       00:00:00:00:00:00 0 0

```

Using the DMF Service Node Appliance

This chapter describes to configure the managed services provided by the DANZ Monitoring Fabric (DMF) Service Node Appliance.

5.1 Overview

The DANZ Monitoring Fabric (DMF) Service Node has multiple interfaces connected to traffic for processing and analysis. Each interface can be programmed independently to provide any 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 action causes the Controller to forward traffic the policy selects to the service node. The processed traffic is returned to the monitoring fabric using the same interface and sent to the tools (delivery interfaces) defined in the DMF policy.

If the traffic volume the policy selects is too much for a single service node interface, define an LAG on the switch connected to the service node, then use the LAG interface when defining the managed service. All service node interfaces connected to the LAG are configured to perform the same action. The traffic the policy selects is automatically load-balanced among the LAG member interfaces and distributes the return traffic similarly.

5.2 Changing the Service Node Default Configuration

Configuration settings are automatically downloaded to the service node from the DANZ Monitoring Fabric (DMF) Controller to eliminate the need for box-by-box configuration. However, the option exists to override the default configuration for a service node from the **config-service-node** submode for any service node.



Note: 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 to use for the service node. This alias is affiliated 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 the **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 to receive SNMP traps from the service node.

5.3 Using SNMP to Monitor DPDK Service Node Interfaces

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: A three-digit number between **101** and **116** identifies SNI DPDK (traffic) interfaces.

In the following example, interface **sni5** (105) handles 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
```

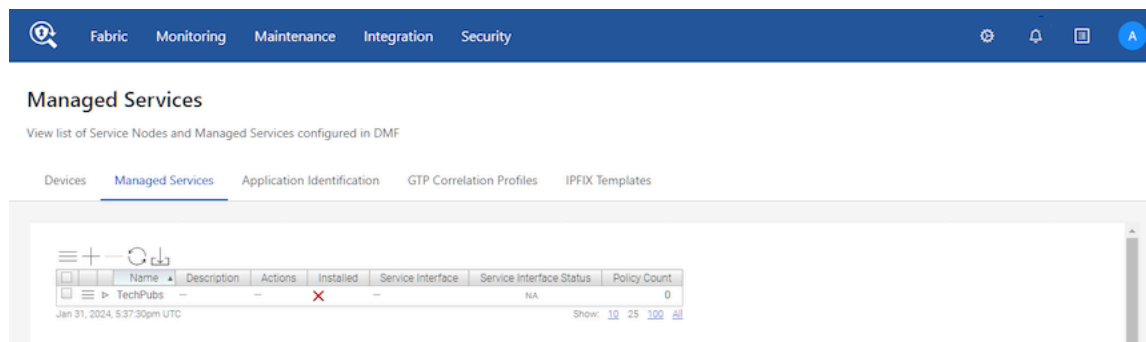
To fetch Link Up and Down status, enter the following command:

```
[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 DANZ Monitoring Fabric (DMF) managed services, select the **Monitoring > Managed Services** option.

Figure 5-1: Managed Services



This page displays the service node appliance devices connected to the DMF Controller and the services configured on the Controller.

5.4.1 Using the GUI to Define a Managed Service

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

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

Figure 5-2: Create Managed Service: Info

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

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

5. Select the action from the Action selection list, which provides the following options.

- **Application ID**
- **Deduplication**: Deduplicate selected traffic, including NATed 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 the 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 forwarding it to specified collectors.
 - **Mask:** Mask sensitive information as specified by the user in packet fields.
 - **NetFlow:** Generate a NetFlow by selecting matching traffic and forwarding it 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 service node receives the packet.
 - **UDP Replication:** Copy UDP messages to multiple IP destinations, such as Syslog or NetFlow messages.
6. (Optional) Identify the starting point for service actions.
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.
 - **I3-header-start:** add the number of bytes specified by the *integer* value to the first byte in the Layer 3 header.
 - **I4-header-start:** add the number of bytes specified by the *integer* value to the first byte in the layer-4 header.
 - **I4-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. Select the backup service from the Backup Service selection list to create a backup service. 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 DANZ Monitoring Fabric (DMF) service node appliance, member links should be of the same speed and can span across multiple service nodes. The maximum number of supported member links per LAG interface is 32, which varies based on the 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 step enters the **config-managed-srv** submode to 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 the 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 for defining 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
```

Use a port channel instead of an interface to increase the bandwidth available to the managed service. 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: Multiple DMF policies can use the same managed service, for example, a packet slicing managed service.

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 MG-LAG-1      MULTIPLE-SERVICES-ACTIVE-1
2 sni2      CORE-SWITCH-1 MG-LAG-2      MULTIPLE-SERVICES-BACKUP-1
3 sni3      CORE-SWITCH-1 MG-LAG-3      MULTIPLE-SERVICES-ACTIVE-2
4 sni4      CORE-SWITCH-1 MG-LAG-4      MULTIPLE-SERVICES-BACKUP-2
5 sni5      CORE-SWITCH-1 MG-LAG-5      MULTIPLE-SERVICES-ACTIVE-3
6 sni6      CORE-SWITCH-1 MG-LAG-6      MULTIPLE-SERVICES-BACKUP-3
7 sni7      CORE-SWITCH-1 MG-LAG-7
8 sni8      CORE-SWITCH-1 MG-LAG-8
9 sni9      CORE-SWITCH-1 MG-LAG-1      MULTIPLE-SERVICES-ACTIVE-1
10 sni10     CORE-SWITCH-1 MG-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 22762691495037 335Mbps 40790968622 22760851128054 335Mbps 38803332951 22533111539451 331Mbps
3 sni3      MULTIPLE-SERVICES-ACTIVE-2 ipfix no 0 0 0 0 0 0 0 0
4 sni3      MULTIPLE-SERVICES-ACTIVE-2 ipfix no 0 0 0 0 0 0 0 0
5 sni5      MULTIPLE-SERVICES-ACTIVE-3 dapper critical 14769229199 1070427935998 29.1Mbps 14769229193 1070427935998 29.1Mbps 1614349 13328879865 367Mbps
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 shows 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 but do not show the bandwidth of the respective data port (SNI) handling traffic, as shown in the following sample snapshot of the **Service Stats** output.

Figure 5-3: Service Node Interface Load Indicator

Service Stats

Show: 10 20 100 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	83,879,701	42,946,406,912	192 Mbit/s	0	0	0 bit/s	83,879,701	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 DANZ Monitoring Fabric (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 enables 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 incoming packets that are identical at the L2/L3/L4 layers.
- Routed packet deduplication: as packets traverse an IP network, the MAC address changes from hop to hop. Routed packet deduplication enables users to match packet contents starting from the L3 header.
- NATed packet deduplication: to perform NATed deduplication, the service node compares packets in the configured window that are identical starting from the L4 payload. To use NATed packet deduplication, perform 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 in which the service looks for duplicate packets is configurable. Select a value among these choices: **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 page. The 'Action' dropdown is set to 'Deduplication'. Under 'Packet Handling', the 'None' radio button is selected. The 'Anchor' dropdown is empty, and the 'Offset' is set to '- Bytes -'. The 'Window Size' is set to '2 ms'. A diagram on the right shows packet layers: Packet Start, L3 Header, L4 Header, and L4 Payload.

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 14-payload-start 0 window 8
service-interface switch CORE-SWITCH-1 ethernet13/3
Controller-1(config)#
```

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

```
Controller-1(config)# show managed-service-device DMF-SN-R740-1 stats
MS-DEDUP dedup
~~~~~ Stats ~~~~~
Interface Name : sn16
Function : dedup
Service Name : MS-DEDUP
Rx packets : 9924950
Rx bytes : 4216466684
Rx Bit Rate : 1.40Gbps
Applied packets : 9923032
Applied bytes : 4216337540
Applied Bit Rate : 1.40Gbps
Tx packets : 9796381
Tx bytes : 4207106113
Tx Bit Rate : 1.39Gbps
Deduped frame count : 126651
Deduped percent : 1.2763336851075358
Load : low
Controller-1(config)#
```

5.6 Header Strip Action

This action removes specific headers from the traffic selected by the associated DANZ Monitoring Fabric (DMF) policy. Alternatively, 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, 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:

- **I3-header-start**
- **I4-header-start**
- **I4-payload-start**
- **packet-start**

Input a positive integer representing the offset from which the strip action begins. When omitting an offset, the header stripping starts from the first byte in the packet.

GUI Configuration

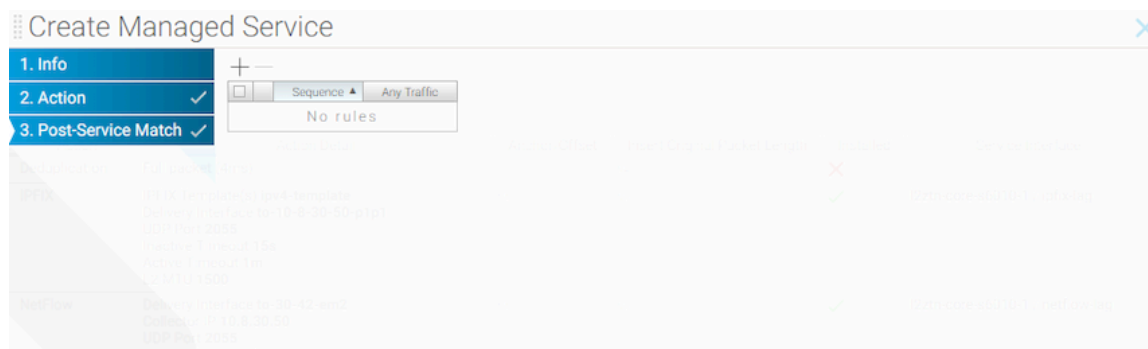
Figure 5-5: Create Managed Service: Header Strip Action



After assigning the required actions to the header stripping service, click **Next** or **Post-Service Match**.

The system displays the **Post Service Match** page, used in conjunction with the header strip service action.

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



CLI Configuration

The header-strip service action strips the header and replaces 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 and add a custom **src-mac**, **dst-mac**, and **ether-type**.

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 the addition of a customized **src-mac**, **dst-mac**, and **ether-type**.

```
! 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 DANZ Monitoring Fabric (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. It is not necessary to change to offset match mode.
- Easier configuration.
- All match conditions are available for the inner packet.
- The 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 example illustrates this configuration.

```
! managed-service
managed-service MS-HEADER-STRIP-4
service-interface switch CORE-SWITCH-1 interface ethernet1
1 decap-l3-mpis
!
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 as a flowset.

Starting with the **DANZ Monitoring Fabric (DMF)-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 except for 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 service node drops the original packets, and the generated flow records, containing metadata about each flow, are forwarded out of the service node interface.

5.7.1 IPFIX Template

The IPFIX 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 that 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 DANZ Monitoring Fabric (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

Use an IPFIX key to specify the characteristics of the traffic 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 flow set having 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 DANZ Monitoring Fabric (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 flow 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:

- The Controller will not allow the key combination of **source-mac-address** and **records-per-dmf-interface** in **push-per-policy** mode.
- The 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 updated for the packets the generator receives that match the specified keys. For example, include fields in the template to record the number of packets, the largest and smallest packet sizes, or the start and end times of the flows. To see a listing of the fields supported in the current release of the DANZ Monitoring Fabric (DMF) Service Node, select the **Monitoring > Managed Service > IPFIX Template** option from the DMF GUI, or type **help** in **config-ipxif-template** submenu. The following are the fields supported:

- **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 flow set 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 that match the flow definition are received, the flow set 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
```

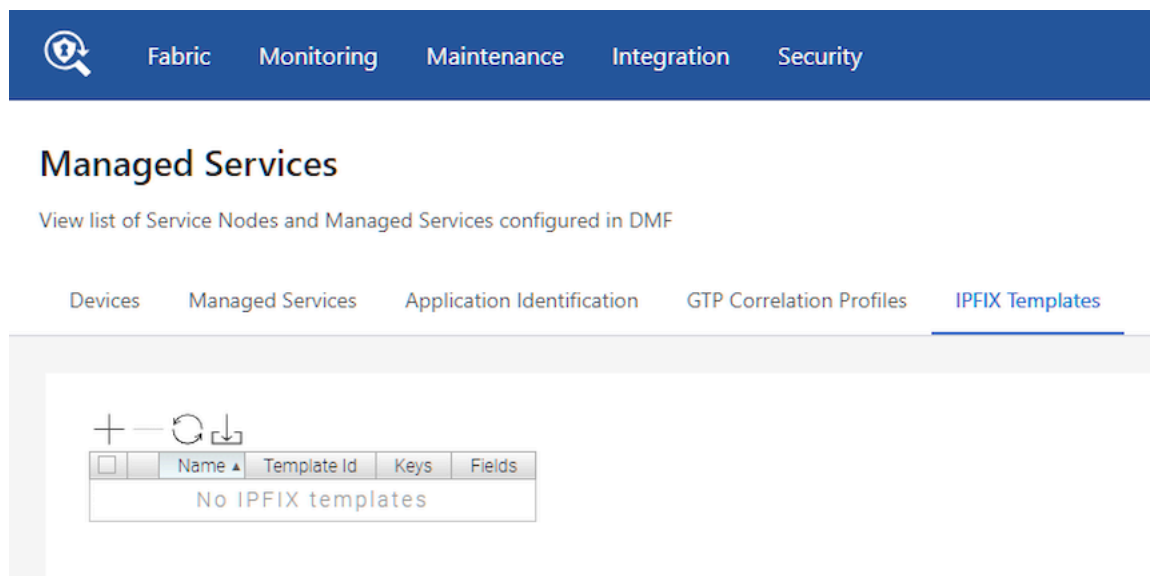
5.7.6 Using the GUI to Define an IPFIX Template

To define an IPFIX template, perform the following steps:

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

The system displays the **IPFIX Templates** section.

Figure 5-8: IPFIX Templates



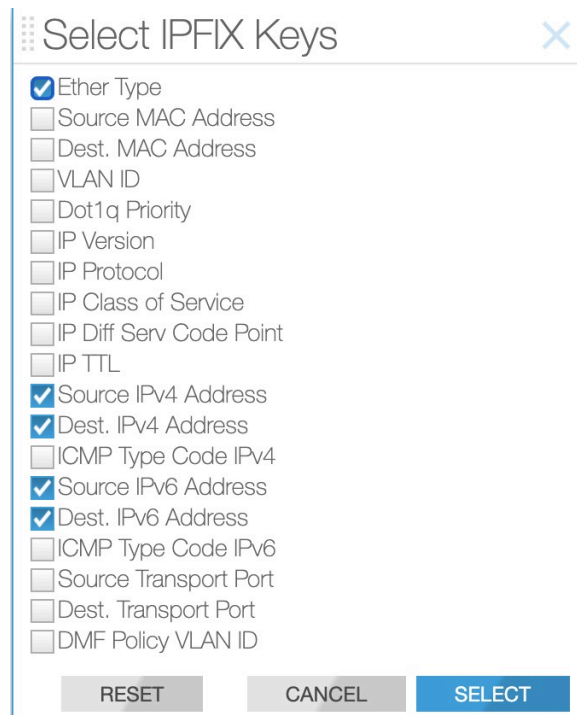
3. To create a new template, click the provision (+) icon in the **IPFIX Templates** section.

Figure 5-9: Create IPFIX Template

The screenshot shows the 'Create IPFIX Template' dialog box. It has a title bar with a close button (X). The dialog contains two input fields: 'Name' with a red asterisk and 'Template ID' with a red asterisk. The 'Name' field contains the text 'Test' and the 'Template ID' field contains the value '256'. Below these fields are two sections: 'Keys' and 'Fields', each with a red asterisk and a dropdown menu. Both dropdown menus are currently set to '- None -'. At the bottom of the dialog 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 dialog box titled "Select IPFIX Keys" contains a list of 20 items, each with a checkbox. The following items are checked: Ether Type, Source IPv4 Address, Dest. IPv4 Address, Source IPv6 Address, and Dest. IPv6 Address. The other items are unchecked. At the bottom are three buttons: RESET, CANCEL, and SELECT.

Key	Selected
Ether Type	Yes
Source MAC Address	No
Dest. MAC Address	No
VLAN ID	No
Dot1q Priority	No
IP Version	No
IP Protocol	No
IP Class of Service	No
IP Diff Serv Code Point	No
IP TTL	No
Source IPv4 Address	Yes
Dest. IPv4 Address	Yes
ICMP Type Code IPv4	No
Source IPv6 Address	Yes
Dest. IPv6 Address	Yes
ICMP Type Code IPv6	No
Source Transport Port	No
Dest. Transport Port	No
DMF Policy VLAN ID	No

5. Enable each checkbox for the keys to add to the template and click **Select**.
6. 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



The dialog box titled "Select IPFIX Fields" contains a list of 17 items, each with a checkbox. The following items are checked: Minimum IP Total Length and Maximum IP Total Length. The other items are unchecked. At the bottom are three buttons: RESET, CANCEL, and SELECT.

Field	Selected
Packet Delta Count	No
Byte Delta Count	No
Minimum IP Total Length	Yes
Maximum IP Total Length	Yes
Flow Start Time (ms)	No
Flow End Time (ms)	No
Flow Start Time (s)	No
Flow End Time (s)	No
Flow End Reason	No
Minimum L2 Total Length	No
Maximum L2 Total Length	No
Minimum TTL	No
Maximum TTL	No
TCP Flags	No

7. Enable the checkbox for each field to add to the template and click **Select**.
8. 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. Use this customized template to apply when defining an IPFIX-managed service.

5.7.7 Using the CLI to Define an IPFIX Template

1. Create an IPFIX template.

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

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

2. 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.

3. 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.

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 'Action' tab selected. The 'Action' dropdown is set to 'IPFIX'. The 'Delivery Interface' is 'DELIVERY-AS-TUNNEL'. Below this, there are input fields for 'Collector IP' and 'Source IP', both currently showing '- IP Address -'. To the right of these are four timeout and port settings: 'Inactive Timeout' (15 seconds), 'Active Timeout' (1 minutes), 'UDP Port' (4739), and 'Template Timeout' (30 seconds). At the bottom, there is a section for 'IPFIX Templates' with a search bar and a list of templates. The first template, 'IPFIX-IP', is selected. At the very bottom of the window are buttons for 'BACK', 'NEXT', 'RESET', 'CANCEL', and 'SAVE'.

Enter the following required configuration details:

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

The following configuration is optional:

- Inactive timeout: the interval of inactivity that marks a flow inactive.
- Active timeout: length of time between each IPFIX flows 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-----
```

5.7.10 Records Per Interface Netflow using DST-MAC Rewrite

Destination MAC rewrite for the records-per-interface NetFlow and IPFIX feature is the default setting and applies to switches running Extensible Operating System (EOS) and SWL and is supported on all platforms.

A configuration option exists for using **src-mac** when overwriting the **dst-mac** isn't preferred.

Configurations using the CLI

Global Configuration

The global configuration is a central place to choose which rewrite option to use for the **records-per-interface**. The following example illustrates using **rewrite-src-mac** or **rewrite-dst-mac** in conjunction with the **filter-mac-rewrite** command.

```
c1(config)# filter-mac-rewrite rewrite-src-mac
c1(config)# filter-mac-rewrite rewrite-dst-mac
```

Netflow Configuration

The following example illustrates a NetFlow configuration.

```
c1(config)# managed-service ms1
c1(config-managed-srv)# 1 netflow
c1(config-managed-srv-netflow)# collector 213.1.1.20 udp-port 2055 mtu 1024
records-per-interface
```

IPFIX Configuration

The following example illustrates an IPFIX configuration.

```
c1(config)# ipfix-template i1
c1(config-ipfix-template)# field maximum-ttl
c1(config-ipfix-template)# key records-per-dmf-interface
c1(config-ipfix-template)# template-id 300

c1(config)# managed-service ms1
c1(config-managed-srv)# 1 ipfix
c1(config-managed-srv-ipfix)# template i1
```

Show Commands

NetFlow Show Commands

Use the **show running-config managed-service** command to view the NetFlow settings.

```
c1(config)# show running-config managed-service
! managed-service
managed-service ms1
!
1 netflow
collector 213.1.1.20 udp-port 2055 mtu 1024 records-per-interface
```

IPFIX Show Commands

Use the **show ipfix-template i1** command to view the IPFIX settings.

```
c1(config)# show ipfix-template i1
~~~~~ Ipfix-templates ~~~~~
# Template Name Keys Fields
-|-----|-----|-----|
1 i1          records-per-dmf-interface maximum-ttl

c1(config)# show running-config managed-service
! managed-service
managed-service ms1
!
1 ipfix
template i1
```

Limitations

- The **filter-mac-rewrite rewrite-src-mac** command cannot be used on the **filter interface** that is part of the policy using **timestamping replace-src-mac**. However, the command works when using a **timestamping add-header-after-l2** configuration.

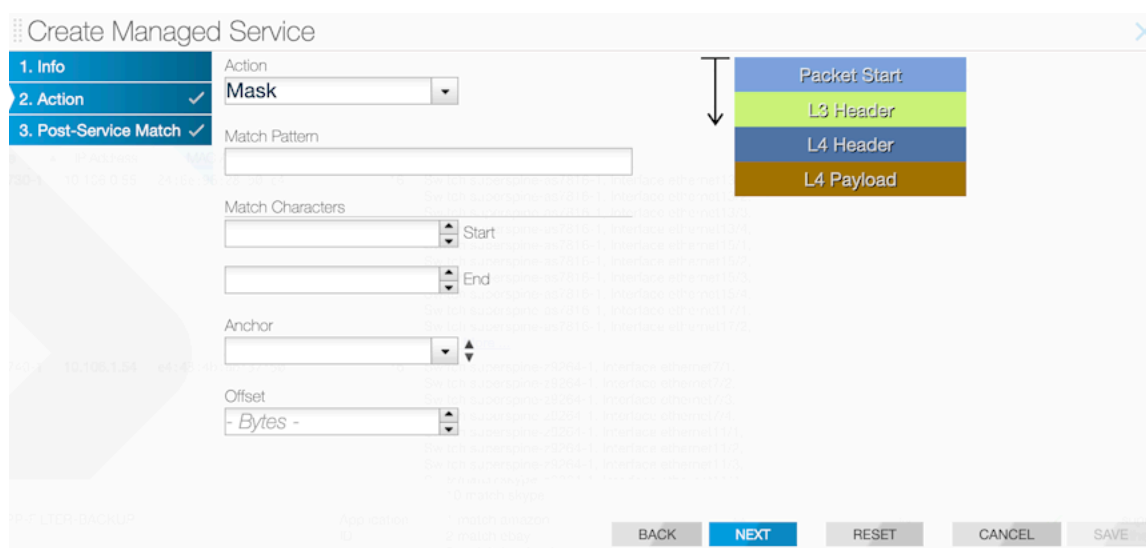
5.8 Packet-masking Action

The packet-masking action can 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



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

Arista Analytics Node capabilities are enhanced to handle NetFlow V5/V9 and IPFIX Packets. All these flow data are represented with the Netflow index.



Note: NetFlow flow record generation is enhanced for selecting VXLAN traffic. For VXLAN traffic, flow processing is based on inner headers, with the VNI 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. Use the DMF Service Node to configure the NetFlow collector IP address and the destination UDP port. 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, apply filter rules 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: The next-hop address must be resolved for the service to be active.

- **Collector IP:** identify the NetFlow collector IP address.
- **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 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 the ARP resolution status.



Note: The DANZ Monitoring Fabric (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.

Enter the **1 netflow** command and identify the configuration name and the submode changes to the **config-managed-srv-netflow** mode for viewing and configuring 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 and packets that are not IPv4 or UDP are passed through. To configure a NetFlow-managed service, perform the following steps:

1. Configure the IP address on the delivery interface.

This IP address is 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.

An option exists to limit the number of flows or change the inactivity timeout using the max-flows or active timeout, or inactive timeout commands.

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

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

The IP address, in IPV4 dotted-decimal notation, is required. The MTU and UDP port are required when changing these parameters from the defaults. Enable the **records-per-interface** option to allow identification of the filter interfaces from which the Netflow originated. Configure the Arista Analytics Node to display this information, as described in the **DMF User Guide**.

The following example illustrates 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 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) to identify HTTPS transactions on non-standard SSL ports. It can filter custom applications and separate control traffic from user data traffic.

Pattern matching is also helpful in 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

The screenshot shows the 'Create Managed Service' configuration window. The 'Action' is set to 'Pattern Drop'. The 'Pattern' field is empty. The 'Anchor' and 'Offset' fields are also empty. A diagram on the right shows the packet structure: Packet Start, L3 Header, L4 Header, and L4 Payload. The 'Offset' is set to '- Bytes -'.

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 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 ports. It can filter custom applications and can separate control traffic from user data traffic.

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

Pattern matching is also helpful in 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 shows the 'Create Managed Service' window with the 'Pattern Match' action selected. The configuration fields include: Action (Pattern Match), Pattern (empty text box), Anchor (empty dropdown), and Offset (- Bytes -). To the right, a diagram illustrates packet slicing with four colored segments: Packet Start (blue), L3 Header (green), L4 Header (dark blue), and L4 Payload (orange). The bottom of the window features buttons for BACK, NEXT, RESET, CANCEL, and SAVE.

CLI Configuration

Use the **pattern-match *pattern*** keyword to enable 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 each packet's vital, relevant portions. Packet slicing can significantly increase the capacity of forensic recording tools. Apply packet slicing 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

Create Managed Service

1. Info

2. Action ✓

3. Post-Service Match ✓

Action

Slice

Insert Original Packet Length

No ☒ Yes

Use to insert original pre-slice packet length into additional header

Anchor

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 -

Packet Structure Diagram:

Packet Start

L3 Header

L4 Header

L4 Payload

This page allows inserting an additional header containing the original header length.

CLI Configuration

Use the **slice** keyword to enable the packet slicing service action and 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 DANZ Monitoring Fabric (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 on 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 and above 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

Configure a **slice-managed** service on a switch using the following steps.

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
```

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
```


Once a managed service for slice action exists, 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. Ensure the offset is within the Min/Max truncate size bounds reported by the **show switch all property** command. If the configured value is beyond the bound, then DMF chooses the closest value of the range.

For example, if a 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, which 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

# 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 - 7820 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: DANZ Monitoring Fabric (DMF) Managed Services

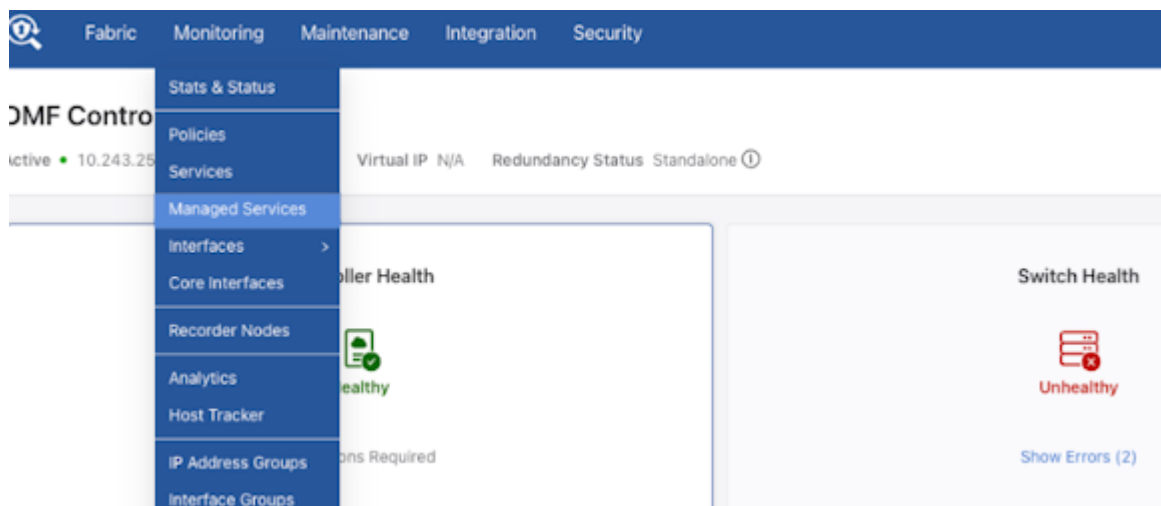
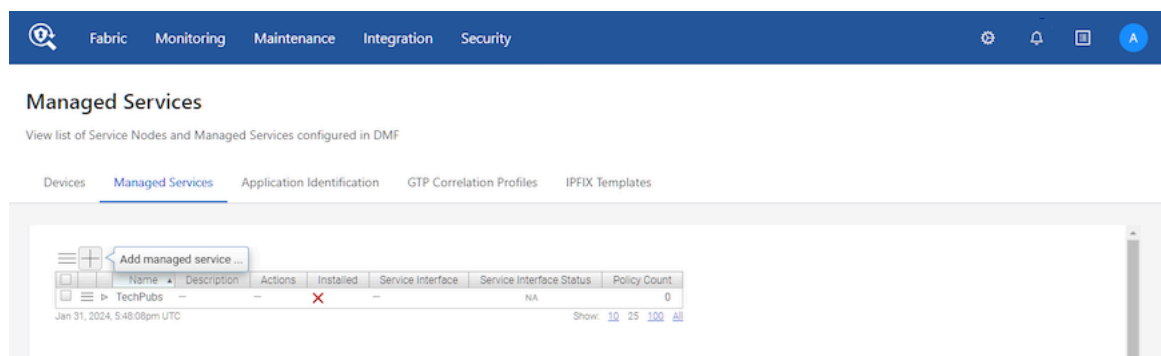


Figure 5-20: DMF Managed Services Add Managed Service



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** to configure 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

5. Click **Append** to continue. The slice action appears on the **Managed Services** page.

Figure 5-24: Slice Action Added

Managed Services

Aug 12, 2023, 3:55:11am UTC

Name	Description	Actions	Installed	Service Interface	Service Interface Status	Policy Count
slice-action-7280-J2-J2C	—	Slice	✓	7280-J2-J2C / Ethernet10/1	✓ Up	0

Show: 10 25 100 All

Interface Loopback Configuration

The managed service interface used for slice action must be in MAC loopback mode.

6. Configure the loopback mode in the **Fabric > Interfaces** page by clicking on the **configuration icon** of the interface.

Figure 5-25: Interfaces

Fabric Monitoring Maintenance Integration Security

Fabric: Healthy

Interfaces

Filter table rows

Switch	Switch DPID	Interface Name	Description	Forward Error Correction	Optics Always Enabled	Disable Transmitting Packets	Management	Span Info	Status	Tunnel	LAG	LAG Parent	Device Name	Speed
apo-accton-as5812T-1	90000000000000000000000000000000	ethernet1	—	—	—	—	—	—	✓ Up	—	—	—	—	1 Gbit/s
apo-accton-as5812X-1	90000000000000000000000000000000	ethernet1	—	—	—	—	—	—	✓ Up	—	—	—	—	1 Gbit/s
apo-accton-as6712-1	90000000000000000000000000000000	ethernet1	—	—	—	—	—	—	✗ Link Down	—	—	—	—	0 bit/s
apo-accton-as6812-1	90000000000000000000000000000000	ethernet1	—	—	—	—	—	—	✗ Link Down	—	—	—	—	0 bit/s
apo-accton-as7312-1	90000000000000000000000000000000	ethernet1	—	—	—	—	—	—	✓ Up	—	—	—	—	10 Gbit/s
apo-accton-as5812T-1	90000000000000000000000000000000	ethernet2	—	—	—	—	—	—	✓ Up	—	—	—	—	1 Gbit/s
apo-accton-as5812X-1	90000000000000000000000000000000	ethernet2	—	—	—	—	—	—	✓ Up	—	—	—	—	1 Gbit/s
apo-accton-as6712-1	90000000000000000000000000000000	ethernet2	—	—	—	—	—	—	✗ Link Down	—	—	—	—	0 bit/s



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.

Policy Configuration

9. Create a new policy from the **DMF Policies** page.

Figure 5-27: DMF Policies Page

DMF Policies

Refresh Clear Stats + Create Policy

Show Quick Filters Table View Interface View

0 Policies Inactive 0 Policies W/ No Traffic 0 Total Policies Errors 0 Interfaces > 90% Util

There seems to be nothing here.
Create Policies to get started

Create Policy

10. Add the previously configured packet slicing managed service.

Figure 5-28: Create Policy

Create Policy

Policy Details

Enable Advanced Options

Collapse

* Name

Input value

Description

Input value

Action

Forward

Push VLAN

##

* Priority

100

Active

Port Selection

Match Traffic

Stat: Bit Rate

Traffic Sources

Create

+ Add Port(s)

Services

+ Add Service(s)

Destination Tools

Create

+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

Cancel

Create Policy

168

11. Select **Add Service** under the **+ Add Service(s)** option shown above.

Figure 5-29: Add Service

The screenshot shows a dialog box titled "Create Policy" with a close button (X) in the top right corner. Below the title bar, there is a section labeled "Add Service" with a "Cancel" button and an "Add Service" button. Underneath, there are two links: "View All Services" and "View All Managed Services". The main area of the dialog is light gray and contains a central icon of a box with an open lid. Below the icon, the text "Click here to add services" is displayed, followed by a blue button with a plus icon and the text "Add Service".

Figure 5-30: Service Type - Service - slice action

The screenshot shows a dialog box titled "Create Policy" with a close button (X) in the top right corner. Below the title bar, there is a section labeled "Add Service" with a "Cancel" button and an "Add 1 Service" button. Underneath, there are two links: "View All Services" and "View All Managed Services". The main area of the dialog is light gray and contains a table with the following structure:

Service Type	* Service	Optional	Backup Service	Del. Service
Managed Service	slice-action	Yes	None	No

Below the table, there is a blue button with a plus icon and the text "Add Another".

12. Click **Add 1 Service** and the slice-managed service (packet-slicing-policy) appears in the **Create Policy** page.

Figure 5-31: Manage Service Added

Create Policy

Policy Details ☐ Enable Advanced Options [Collapse](#)

* Name	Description	Action	Push VLAN ?	* Priority ?	Active
<input type="text" value="Input value"/>	<input type="text" value="Input value"/>	Forward	##	100	<input checked="" type="checkbox"/>

Port Selection Match Traffic Stat: Utilization

Traffic Sources [Create](#)
[+ Add Port\(s\)](#)

Services
[+ Add Service\(s\)](#)
• slice-action
Managed Service
Optional

Destination Tools [Create](#)
[+ Add Port\(s\)](#)

Match Traffic: Allow All Traffic (Default) has been selected

[Cancel](#) [Create Policy](#)

13. Click **Create Policy** and the new policy appears in **DMF Policies**.

Figure 5-32: DMF Policy Configured

Fabric Monitoring Maintenance Integration Security

DMF Policies [Refresh](#) [Clear Stats](#) [+ Create Policy](#)

☒ Show Quick Filters [Table View](#) [Interface View](#)

1 Policy Inactive 1 Policy W/ No Traffic 0 Total Policies Errors 0 Interfaces > 90% Util

Delete Edit Duplicate Showing All 1 Policies [Filter](#)

Policy Name	Filter Interfaces	CloudVision Portal(s)	vCenter(s)	Services	RN Fabric Interfaces	Delivery Interfaces
techpubs-1	-	-	-	1 Service	-	-
				• slice-action		



Note: The images above have been edited for documentation purposes. The actual outputs may 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 DANZ Monitoring Fabric (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. The 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 using a switch interface 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, affecting 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 VXLAN Stripping on the 7280R3 Switch

Virtual Extensible LAN Header Stripping

Virtual Extensible LAN (VXLAN) Header Stripping supports the delivery of decapsulated packets to tools and devices in a DANZ Monitoring Fabric (DMF) fabric. This feature removes the VXLAN header, previously established in a tunnel for reaching the TAP Aggregation switch or inherent to the tapped traffic within the DMF. Within the fabric, DMF supports the installation of the strip VXLAN service on a filter interface or a filter-and-delivery interface of a supported switch.

Platform Compatibility

For DMF deployments, the target platform is DCS-7280R3.

Use the **show switch all property** command to verify which switch in the DMF fabric supports this feature.

The feature is supported if the **Strip Header Supported** property has the value **BSN_STRIP_HEADER_CAPS_VXLAN**.



Note: The following example is displayed differently for documentation purposes than what appears when using the CLI.

```
# show switch all property
#
Switch : 1
Max Phys Port : 1000000
Min Lag Port : 1000001
Max Lag Port : 1000256
Min Tunnel Port : 15000001
Max Tunnel Port : 15001024
Max Lag Comps : 64
Tunnel Supported : BSN_TUNNEL_L2GRE
UDF Supported : BSN_UDF_6X2_BYTES
Enhanced Hash Supported : BSN_ENHANCED_HASH_L2GRE,BSN_ENHANCED_HA
SH_L3,BSN_ENHANCED_HASH_L2,
BSN_ENHANCED_HASH_MPLS,BSN_ENHANCED_HAS
H_SYMMETRIC
Strip Header Supported : BSN_STRIP_HEADER_CAPS_VXLAN
Min Rate Limit : 1Mbps
Max Multicast Replication Groups : 0
Max Multicast Replication Entries : 0
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
Min Truncate Offset : 100
Max Truncate Offset : 9236
```

5.14.1 Using the CLI to Configure VXLAN Header Stripping

Configuration

Use the following steps to configure **strip-vxlan** on a switch:

1. Set the optional field **strip-vxlan-udp-port** at switch configuration, and the default **udp-port** for **strip-vxlan** is 4789.

2. Enable or disable **strip-vxlan** on a **filter** or **both-filter-and-delivery** interface using the **role both-filter-and-delivery interface-name filter-interface strip-vxlan** command.

```
> enable
# config
(config)# switch switch-name
(config-switch)# strip-vxlan-udp-port udp-port-number

(config-switch)# interface interface-name
(config-switch-if)# role both-filter-and-delivery interface-name filter-
interface strip-vxlan
(config-switch-if)# role both-filter-and-delivery interface-name filter-
interface no-strip-vxlan
(config)# show running-config
```

3. After enabling a filter interface with **strip-vxlan**, any policy can use it. From the **config-policy** submode, add the **filter-interface** to the policy:

```
(config)# policy p1
(config-policy)# filter-interface filter-interface
```

Show Commands

Use the **show policy policy name** command to see the runtime state of a policy using a filter interface with **strip-vxlan** configured. It will also show the service interface information and stats.

```
# show policy strip-vxlan
Policy Name                : strip-vxlan
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  : -
Total Delivery Rate        : -
Total Pre Service Rate     : -
Total Post Service Rate    : -
Overlapping Policies       : none
Component Policies         : none
Installed Time              : 2024-05-02 19:54:27 UTC
Installed Duration         : 1 minute, 18 secs
Timestamping enabled       : False
~ 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 f1      lyd598 Ethernet1/1 up    rx  0      0      0      -
  2024-05-02 19:54:27.141000 UTC
~~~~~ Delivery Interface(s) ~~~~~
~~~~~
# DMF IF Switch IF Name      State Dir Packets Bytes Pkt Rate Bit Rate Counter
  Reset Time
-|-----|-----|-----|-----|---|-----|-----|-----|-----|-----|
-----|
1 d1      lyd598 Ethernet2/1 up    tx  0      0      0      -
  2024-05-02 19:54:27.141000 UTC
~ Service Interface(s) ~
None.
~ Core Interface(s) ~
None.
~ Failed Path(s) ~
None.

```

5.14.2 Using the GUI to Configure VXLAN Header Stripping

Filter Interface Configuration

To configure or edit a filter interface, proceed to **Interfaces** from the **Monitoring** menu and select **Interfaces > Filter Interfaces**.

Figure 5-33: Filter Interfaces

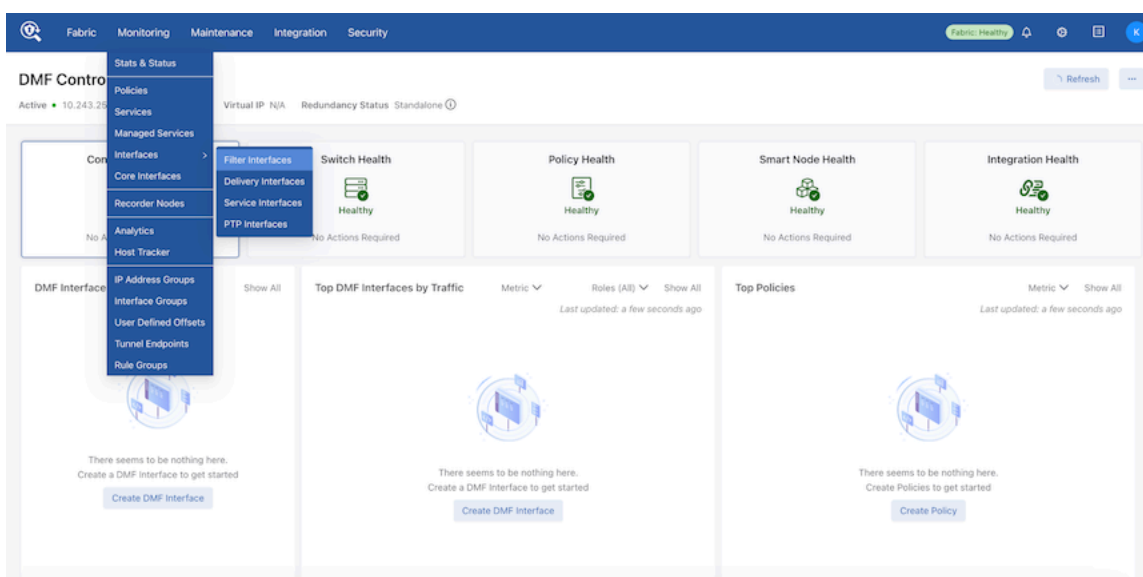



Figure 5-34: DMF Interfaces



Configure a filter interface on a switch that supports **strip-vxlan**.

Figure 5-35: Configure Filter 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

DMF Switch: lyd598 (00:00:fc:59:c0:9d:f6:4b)

Optional: use to filter interface choices below

Interface * Ethernet1/1

Choose from a list of interfaces not already configured for DMF

Show Connected Switches Only

Enable or Disable Strip VXLAN.

Figure 5-36: Enable Strip VXLAN

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

PTP

Interface Name *

f1

Rewrite VLAN

Controller will automatically push VLAN.

Rewrite Dest. MAC Address

- MAC Address -

Analytics

Disabled

Enabled

Strip VxLAN

Disable

Enable ?

sFlow

Disabled

Enabled

sFlow sampling on this interface

IP Address

- IPv4 Address -

For receiving IP datagram traffic

Figure 5-37: DMF Interfaces Updated

Fabric Monitoring Maintenance Integration Security

Fabric: healthy

Interface "Ethernet1/1" saved

DMF Interfaces

js_filter=true

FILTER

Name	Role	Filter	Delivery	PTP	MAC	Switch Name	Interface	State	Speed	Rate Limit	Description	Analytics	Strip VLAN	aflow	VLAN Tag	Rewrite Dest. MAC	Strip VLAN Action	IP Address	Next Hop IP	Next Hop Mask	APP Resolved	APP Internal	VLAN Preservation	Preserve User Configured VLANs	Preserve VLANs	Egress Filter
Ethernet1/1	Filter	✓	—	—	—	spine9	Ethernet1/1	Down	—	—	—	✓	Enable	✓	—	—	—	—	—	—	—	56	—	—	—	

May 2, 2024, 7:40:00pm UTC

Policy Configuration

Create a new policy using **DMF Policies** and add the filter interface with strip VXLAN enabled.

Figure 5-38: Create Policy

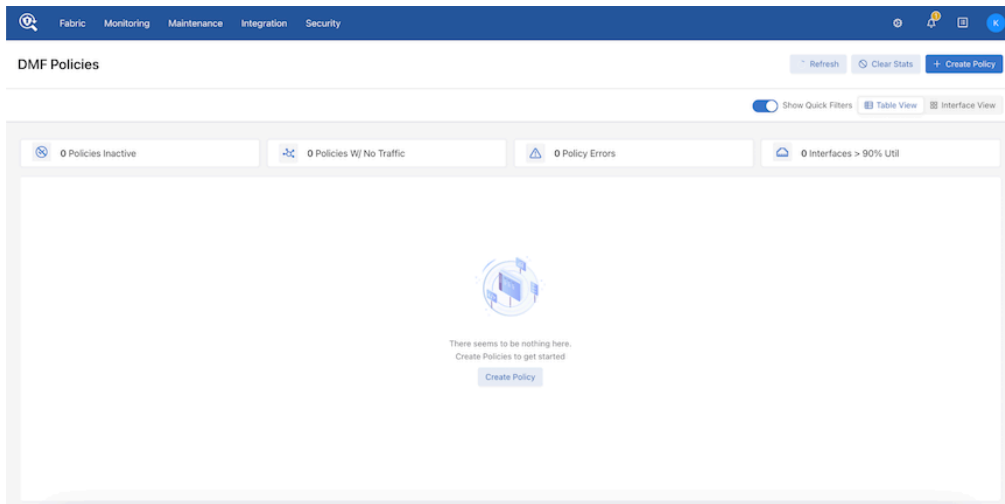
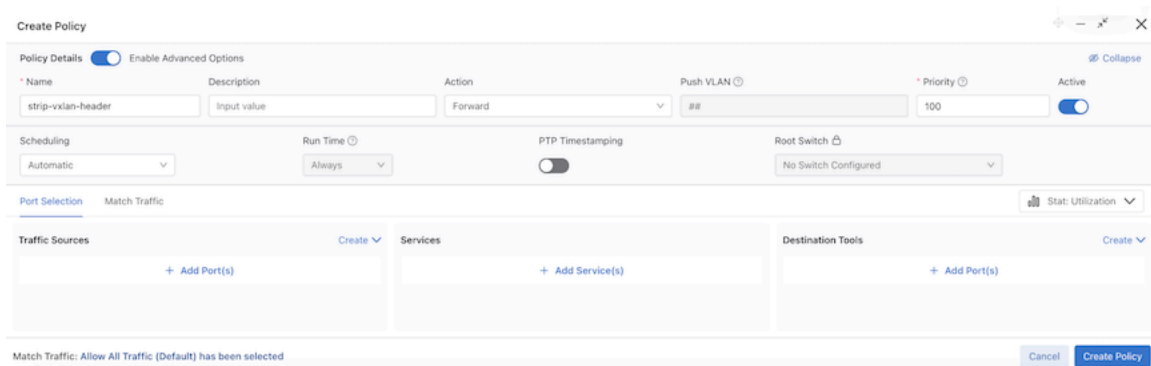
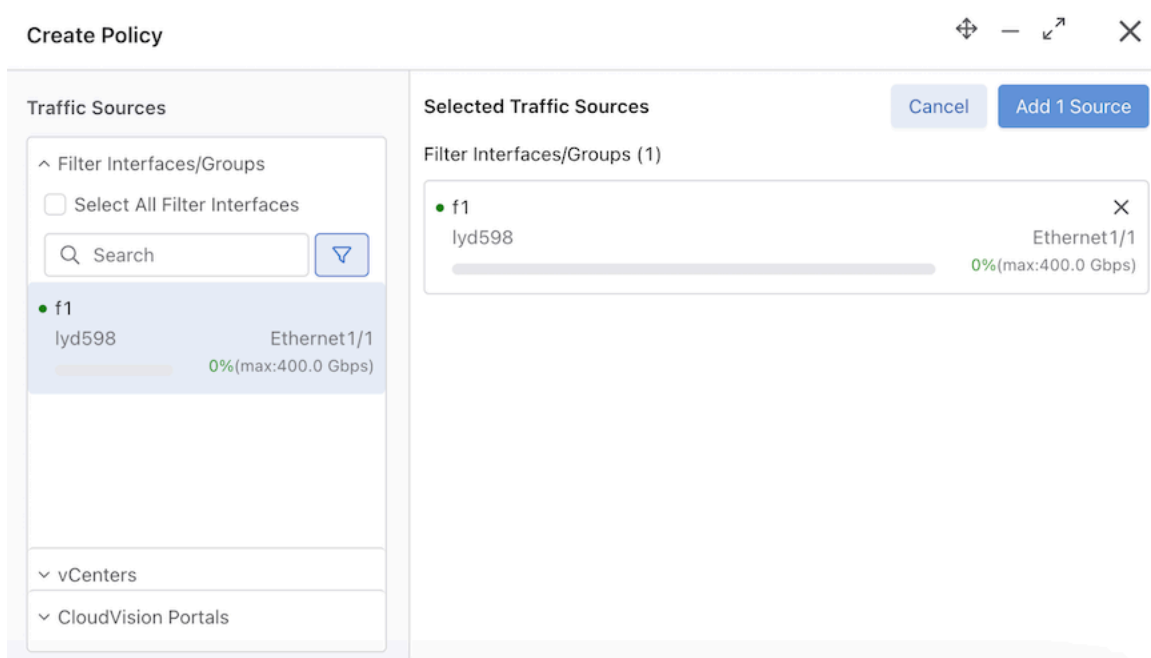


Figure 5-39: Strip VXLAN Header



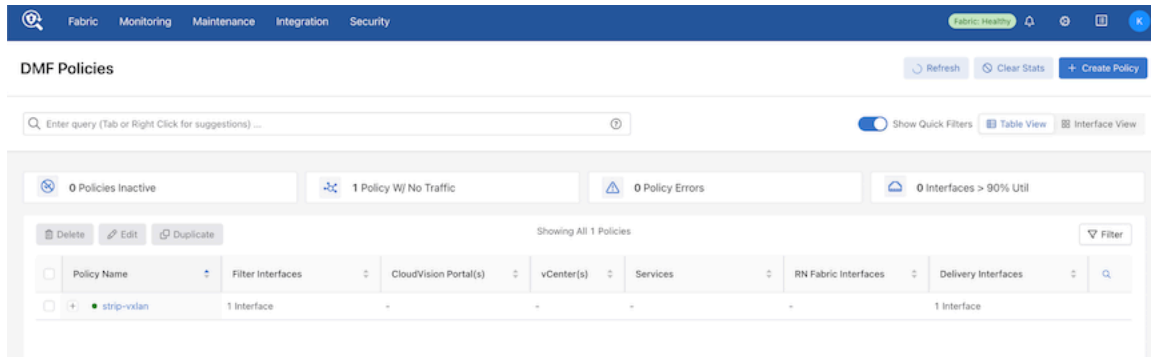
Select **Add port(s)** under the **Traffic Sources** option and add the **Filter Interface**.

Figure 5-40: Selected Traffic Sources



Add another delivery interface and create the policy.

Figure 5-41: Policy Created



5.14.3 Syslog Messages

There are no syslog messages associated with this feature.

5.14.4 Troubleshooting

The `show switch all property` command provides the **Strip Header Supported** property of the switch. If the value `BSN_STRIP_HEADER_CAPS_VXLAN` is present, the feature is supported; otherwise, the switch does not support this feature.

The `show fabric warnings feature-unsupported-on-device` command provides information when DMF fails to enable `strip-vxlan` on an unsupported switch.

The `show switch switch-name table strip-vxlan-header` command provides the gentable details.

The following are examples of several failure cases:

1. The filter interface is down.
2. The interface with `strip-vxlan` is neither a filter interface nor a filter-and-delivery interface.
3. The switch does not support `strip-vxlan`.
4. Tunneling / UDF is enabled simultaneously with `strip-vxlan`.
5. Unsupported pipeline mode with `strip-vxlan` enabled (`strip-vxlan` requires a specific pipeline mode `strip-vxlan-match-push-vlan`).

5.14.5 Limitations

- When configured for a supported switch, the filter interface for `decap-vxlan` action can only be a physical interface or a LAG.
- It is not possible to enable `strip-vxlan` simultaneously with tunneling / UDF.
- When enabling `strip-vxlan` on one or more switch interfaces on the same switch, other filter interfaces on the same switch cannot be matched on the VXLAN header.

5.15 Session Slicing for TCP and UDP Sessions

Session-slice keeps track of TCP and UDP sessions (distinguished by source and destination IP address and port) and counts the number of packets sent in each direction (client-to-server and vice versa). After recognizing the session, the action transmits a user-configured number of packets to the tool node.

For TCP packets, session-slice tracks the number of packets sent in each direction after establishing the TCP handshake. Slicing begins after the packet count in a direction has reached the configured threshold in both directions.

For UDP packets, slicing begins after reaching the configured threshold in either direction.

By default, session-slice will operate on both TCP and UDP sessions but is configurable to operate on only one or the other.



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 dropping packets.

Refer to the *DANZ Monitoring Fabric (DMF) Verified Scale Guide* for session-slicing performance numbers.

Configure `session-slice` in managed services through the Controller as a Service Node action.

- [Using the CLI to Configure Session Slicing](#)
- [Using the GUI to Configure Session Slicing](#)

5.15.1 Using the CLI to Configure Session Slicing

Configure `session-slice` in managed services through the Controller as a Service Node action.

Configuration Steps

1. Create a managed service and enter the service interface.
2. Choose the session-slice service action with the command: `<seq num> session-slice`



Note: The `<seq num> session-slice` command opens the `session-slice` submode, which supports two configuration parameters: `slice-after` and `idle-timeout`.

3. Use `slice-after` to configure the packet threshold, after which the Service Node will stop forwarding packets to tool nodes.
4. Use `idle-timeout` to configure the timeout in milliseconds before an idle connection is removed from the cache. `idle-timeout` is an optional command with a default value of 60000 ms.

```
dmf-controller-1(config)# managed-service managed_service_1
dmf-controller-1(config-managed-srv)# 1 session-slice
dmf-controller-1(config-managed-srv-ssn-slice)# slice-after 1000
dmf-controller-1(config-managed-srv-ssn-slice)# idle-timeout 60000
```

Show Commands

The following `show` commands provide helpful information.

The `show running-config managed-service managed service` command helps verify whether the `session-slice` configuration is complete.

```
dmf-controller-1(config)# show running-config managed-service managed_service_1

! managed-service
managed-service managed_service_1
!
1 session-slice
slice-after 1000
idle-timeout 60000
```

The `show managed-services managed service` command provides status information about the service.

```
dmf-controller-1(config)# show managed-services managed_service_1
# Service Name      Switch      Switch Interface Installed Max Post-Service
  BW Max Pre-Service BW Total  Post-Service BW Total Pre-Service BW
-----|-----|-----|-----|-----|
1 managed_service_1 DCS-7050CX3-32S ethernet2/4      True      25Gbps
                25Gbps                624Kbps                432Mbps
```

5.15.2 Using the GUI to Configure Session Slicing

Perform the following steps to configure session slicing.

1. Navigate to **Monitoring > Managed Services > Managed Services**

Figure 5-42: Managed Services

FabricMonitoringMaintenanceIntegrationSecurityFabric Healthy12

Managed Services

View list of Service Nodes and Managed Services configured in DMF

DevicesManaged ServicesApplication IdentificationGTP Correlation ProfilesIPFIX Templates

	Name	Description	Actions	Installed	Service Interface	Service Interface Status	Policy Count
	app-id	App ID Filter, App ID		✓	DCS-7050CX3-32S / ethernet11/1	✓ Up	1
	Dup	Deduplication		✓	DCS-7050CX3-32S / ethernet29/3	✓ Up	1
	h1	Header Strip		✓	DCS-7050CX3-32S / ethernet7/1	✓ Up	0
	ip-test	IPFIX		✓	DCS-7050CX3-32S / ethernet1/3	✓ Up	0
	L2	Flow Latency and Drops		✓	DCS-7050CX3-32S / ethernet1/4	✓ Up	0
	LD	Flow Latency and Drops		✓	DCS-7280CR3-32P4 / Ethernet31/4	✓ Up	1
	LD1	Flow Latency and Drops		✗	—	NA	0
	Reg-verify	IPFIX		✓	DCS-7050CX3-32S / ethernet4/4	✓ Up	0
	Sam	Sample		✓	DCS-7050CX3-32S / ethernet9/1	✓ Up	0
	t	Session Slice		✗	—	NA	0
	t1	Session Slice, Header Strip		✓	DCS-7280CR3-32P4 / Ethernet29/1	✓ Up	0
	test	Flow Latency and Drops		✗	DCS-7050CX3-32S / ethernet9/4	✓ Up	0

Jul 19, 2024, 6:43:23pm UTCShow: 1025100All(1 - 12 / 12)

2. Select the + icon to create a new managed service.

Figure 5-43: Create Managed Service

Create Managed Service

1. Info

2. Actions ✓

3. Post-Service Match ✓

Name

Description

Switch

Optional: use to filter interface choices below

Interface

- Interface -

Choose from a list of interfaces not already configured for DMF

Show Managed Device Switches Only

Show Connected Switches Only

BACK

NEXT

RESET

CANCEL

SAVE

3. Enter a **Name** for the managed service.

Figure 5-44: Managed Service Name

Create Managed Service DCS-7050CX3-32S / ethernet4/1 Up

1. Info ✓ **Name** DCS-7050CX3-32S / ethernet4/1 Up **Description** 0

2. Actions ✓ **Switch** DCS-7050CX3-32S / ethernet4/1 Up ☒ **Show Managed Device Switches Only**

3. Post-Service Match ✓ **Interface** DCS-7280CR3-32P4 / Ethernet23.1 Up ☐ **Show Connected Switches Only**

IPFIX 0

Sample 0

Session Slice 0

Session Slice, Header Strip 0

Flow Latency and Drops 0

Optional: use to filter interface choices below

Interface - Interface - 0

Choose from a list of interfaces not already configured for DMF

BACK **NEXT** **RESET** **CANCEL** **SAVE**

4. Select a **Switch** from the drop-down list.

Figure 5-45: Manage Service Switch

Create Managed Service DCS-7050CX3-32S / ethernet4/1 Up

1. Info ✓ **Name** DCS-7050CX3-32S / ethernet4/1 Up **Description** 0

2. Actions ✓ **Switch** DCS-7050CX3-32S / ethernet4/1 Up ☒ **Show Managed Device Switches Only**

3. Post-Service Match ✓ **Interface** DCS-7280CR3-32P4 / Ethernet23.1 Up ☐ **Show Connected Switches Only**

IPFIX 0

Sample 0

Session Slice 0

Session Slice, Header Strip 0

Flow Latency and Drops 0

Optional: use to filter interface choices below

Interface - Interface - 0

Choose from a list of interfaces not already configured for DMF

BACK **NEXT** **RESET** **CANCEL** **SAVE**

Figure 5-46: Managed Service Switch Added

Create Managed Service DCS-7050CX3-32S / ethernet4/1 Up

1. Info ✓ **Name** DCS-7050CX3-32S / ethernet4/1 Up **Description** 0

2. Actions ✓ **Switch** DCS-7050CX3-32S / ethernet4/1 Up ☒ **Show Managed Device Switches Only**

3. Post-Service Match ✓ **Interface** DCS-7280CR3-32P4 / Ethernet23.1 Up ☐ **Show Connected Switches Only**

IPFIX 0

Sample 0

Session Slice 0

Session Slice, Header Strip 0

Flow Latency and Drops 0

Optional: use to filter interface choices below

Interface - Interface - 0

Choose from a list of interfaces not already configured for DMF

BACK **NEXT** **RESET** **CANCEL** **SAVE**

5. Select an **Interface** from the drop-down list.

Figure 5-47: Managed Service Interface Added

Create Managed Service

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

Name: TechPubs_ Description:

Switch: DCS-7280CR3-32P4 (00:00:e4:78:76:f2:46) ☒ Show Managed Device Switches Only ☐ Show Connected Switches Only

Optional: use to filter interface choices below

Interface: Ethernet27/3

- ✓ Ethernet25/1
- ✓ Ethernet25/2
- ✓ Ethernet25/3
- ✓ Ethernet25/4
- ✓ Ethernet27/1
- ✓ Ethernet27/2
- ✓ Ethernet27/3
- ✓ Ethernet27/4
- ✓ Ethernet29/1
- ✓ Ethernet29/2
- ✓ Ethernet29/3
- ✓ Ethernet29/4
- ✓ Ethernet31/1
- ✓ Ethernet31/2
- ✓ Ethernet31/3
- ✓ Ethernet31/4

BACK NEXT RESET CANCEL SAVE

6. Select **Actions** or **Next**.

Figure 5-48: Actions Menu

Create Managed Service

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

Click + and - to include or exclude actions.

No actions

Name	Status	Actions
DCS-7280CR3-32P4 / Ethernet25/1	Up	0
DCS-7280CR3-32P4 / Ethernet25/2	Up	0
DCS-7280CR3-32P4 / Ethernet25/3	Up	0
DCS-7280CR3-32P4 / Ethernet25/4	Up	0
DCS-7280CR3-32P4 / Ethernet27/1	Up	0
DCS-7280CR3-32P4 / Ethernet27/2	Up	0
DCS-7280CR3-32P4 / Ethernet27/3	Up	0
DCS-7280CR3-32P4 / Ethernet27/4	Up	0
DCS-7280CR3-32P4 / Ethernet29/1	Up	0
DCS-7280CR3-32P4 / Ethernet29/2	Up	0
DCS-7280CR3-32P4 / Ethernet29/3	Up	0
DCS-7280CR3-32P4 / Ethernet29/4	Up	0
DCS-7280CR3-32P4 / Ethernet31/1	Up	0
DCS-7280CR3-32P4 / Ethernet31/2	Up	0
DCS-7280CR3-32P4 / Ethernet31/3	Up	0
DCS-7280CR3-32P4 / Ethernet31/4	Up	0

BACK NEXT RESET CANCEL SAVE

7. Click the **+** icon to select a managed service action.

Figure 5-49: Configure Managed Service Action List

Configure Managed Service Action

Sequence *
1

Action

None

- None
- Deduplication
- App ID
- App ID Filter
- Flow Latency and Drops
- 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
- Mask
- NetFlow
- Pattern Drop
- Pattern Match
- Sample
- Session Slice
- Slice
- TCP Analysis
- Timestamp
- UDP Replication

BACK NEXT CANCEL APPEND

8. Choose **Session Slice** from the drop-down list. Adjust the **Slice After** and **Idle Timeout** parameters, as required.

Figure 5-50: Configure Managed Service Action Session Slice

Configure Managed Service Action

Sequence *
1

Action
Session Slice

Slice After *
16 packet(s)
After the TCP handshake, forward this many packets and drop all remaining.

Idle Timeout *
60000 milliseconds
Idle timeout before removal from cache

BACK SAVE RESET CANCEL

CANCEL APPEND

9. Select **Append** and then **Save** to add the session slice managed service.

Figure 5-51: Managed Service Session Slice

Managed Services

View list of Service Nodes and Managed Services configured in DMF

Devices **Managed Services** Application Identification GTP Correlation Profiles IPFIX Templates

Name	Description	Actions	Installed	Service Interface	Service Interface Status	Policy Count
app-id	App ID Filter, App ID		✓	DCS-7050CX3-32S / ethernet11/1	✓ Up	1
Dup	Deduplication		✓	DCS-7050CX3-32S / ethernet29/3	✓ Up	1
h1	Header Strip		✓	DCS-7050CX3-32S / ethernet7/1	✓ Up	0
ip-test	IPFIX		✓	DCS-7050CX3-32S / ethernet1/3	✓ Up	0
L2	Flow Latency and Drops		✓	DCS-7050CX3-32S / ethernet1/4	✓ Up	0
LD	Flow Latency and Drops		✓	DCS-7280CR3-32P4 / Ethernet31/4	✓ Up	1
LD1	Flow Latency and Drops		✗	—	NA	0
Reg-verify	IPFIX		✓	DCS-7050CX3-32S / ethernet4/4	✓ Up	0
Sam	Sample		✓	DCS-7050CX3-32S / ethernet9/1	✓ Up	0
t	Session Slice		✗	—	NA	0
t1	Session Slice, Header Strip		✓	DCS-7280CR3-32P4 / Ethernet29/1	✓ Up	0
TechPubs	Session Slice		✓	DCS-7280CR3-32P4 / Ethernet27/3	✓ Up	0
test	Flow Latency and Drops		✗	DCS-7050CX3-32S / ethernet9/4	✓ Up	0

Jul 19, 2024, 6:53:38pm UTC Show: 10 25 100 All (1 - 13 / 13)

5.16 Timestamp Action

The timestamp service action identifies and timestamps every packet it receives with the time the service node receives the packet for matching traffic.

GUI Configuration

Figure 5-52: Create Managed Service: Timestamp Action

Create Managed Service

1. Info ✓ Action

2. Action ✓ Timestamp

3. Post-Service Match ✓

Application ID	Match	Interfaces
amazon	1	ethernet15/3

BACK NEXT RESET CANCEL SAVE

CLI Configuration

```
! managed-service
managed-service MS-TIMESTAMP-1
1 timestamp
service-interface switch CORE-SWITCH-1 ethernet15/3
```

5.17 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.

Configure a rate limit when enabling UDP replication. When upgrading from a version of DANZ Monitoring Fabric (DMF) before release **6.3.1**, the UDP-replication configuration is not applied until a rate limit is applied to the delivery interface.

The following example illustrates 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

Create Managed Service

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

Action: **UDP Replication**

Input Packet Destination IP: **- IPv4 Address -**

Delivery Interface: **eth0**

Output Packet Destination IPs: **+ 0 items**

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 -**

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 function sends 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 5-53: Configure Output Packet Destination IP

Configure Output Packet Dest. IP

IP Address: **- IPv4 Address -**

EtherType: **- Select By Name -**

L2 Header: **- MAC Address -**

CANCEL **ADD**

For the header-strip service action only, configure the policy rules for matching traffic after applying the header-strip service action. After completing pages 1-4, click **Append** and enable the checkbox to apply the policy.

Click **Save** to save the managed service.

CLI Configuration

Enter the **1 udp-replicate** command and identify the configuration name (the submode changes to the **config-managed-srv-udp-replicate** submode) to 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 for sending 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
```

5.18 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 DANZ Monitoring Fabric (DMF) policy. The backup service is activated only when the primary service becomes unavailable. The backup service can be on the same service node or core switch or a different service node and core switch.



Note: Transitioning from active to backup managed service requires reprogramming switches and associated managed appliances. This reprogramming, done seamlessly, will result in a slight traffic loss.

5.18.1 Using the GUI to Configure a Backup Managed Service

To assign a managed service as a backup service in a DANZ Monitoring Fabric (DMF) policy, perform the following steps:

1. Select **Monitoring > Policies** and click the Provision control (+) to create a new policy.
2. Configure the policy as required. From the **Services** section, click the Provision control (+) in the **Managed Services** table.

Figure 5-54: 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**.

5.18.2 Using the CLI to Configure a Backup Managed Service

To implement backup-managed services, perform 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 referring 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
```

5.19 Application Identification

The DANZ Monitoring Fabric (DMF) Application Identification feature allows for the monitoring of applications identified with Deep Packet Inspection (DPI) into packet flows received via filter interfaces and generates IPFIX flow records. These IPFIX flow records are transmitted to a configured collector device via the L3 delivery interface. The feature provides a filtering function by forwarding or dropping packets from specific applications before sending the packet to the analysis tools.



Note: Application identification is supported on R640 Service Nodes (DCA-DM-SC and DCA-DM-SC2) and R740 Service Nodes (DCA-DM-SDL and DCA-DM-SEL).

5.19.1 Using the CLI to Configure Application Identification

Configure the feature through the Controller in managed services.

There are two application identification services to configure:

- [app-id](#)
- [app-id-filter](#)
- [app-id and app-id-filter combined](#)

5.19.2 Using the CLI to Configure app-id

Perform the following steps to configure app-id.

1. Create a managed service and enter the service interface.
2. Choose the app-id managed service using the `<seq num> app-id` command.



Note: The above command should enter the app-id submode, which supports two configuration parameters: `collector` and `l3-delivery-interface`. Both are required.

3. To configure the IPFIX collector IP address, enter the following command: `collector ip-address`. The UDP port and MTU parameters are optional; the default values are **4739** and **1500**, respectively.
4. Enter the command: `l3-delivery-interface delivery interface name` to configure the delivery interface.
5. Add this managed service to a policy. The policy will not have a physical delivery interface.


The following shows an example of an 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
l3-delivery-interface app-to-analytics
```

After configuring the app-id, refer to the analytics node for application reports and visualizations. For instance, a flow is classified internally with the following tuple: **ip**, **tcp**, **http**, **google**, and **google_maps**. Consequently, the analytics node displays the most specific app ID for this flow as **google_maps** under **appName**.

On the Analytics Node, there are AppIDs 0-4 representing applications according to their numerical IDs. 0 is the most specific application identified in that flow, while 4 is the least. In the example above, ID 0 would be the numerical ID for **google_maps**, ID 1 **google**, ID 2 **http**, ID 3 **tcp**, and ID 4 **IP address**. Use the **appName** in place of these since these require an ID to name mapping to interpret.

246,557 hits



Time ↓	appName	appId	qosmosAppID0
> Oct 30, 2023 @ 15:47:37.000	chase_bank	4,311	4,311
> Oct 30, 2023 @ 15:47:37.000	bofa	2,885	2,885
> Oct 30, 2023 @ 15:47:37.000	chase_bank	4,311	4,311
> Oct 30, 2023 @ 15:47:37.000	chase_bank	4,311	4,311
> Oct 30, 2023 @ 15:47:37.000	chase_bank	4,311	4,311

5.19.3 Using the CLI to Configure app-id-filter

Perform the following steps to configure app-id-filter:

1. Create a managed service and enter the service interface.
2. Choose the app-id managed service using the **<seq num> app-id-filter** command.



Note: The above command should enter the **app-id-filter** submode, which supports three configuration parameters: **app**, **app-category**, and **filter-mode**. The category **app** is required, while **app-category** and **filter-mode** are optional. The option **filter-mode** has a default value of **forward**.

3. Enter the command: **app application name** to configure the 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 the **Tab** to see all possible choices to auto-complete the name. The application name provided 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 using the **app application-name** command. An example of a (partial) list of names:

```
dmf-controller-1 (config-managed-srv-app-id-filter)# app ibm
```

```

ibm      ibm_as_central  ibm_as_dtaq  ibm_as_netprt  ibm_as_srvmap
ibm_iseries  ibm_tsm
ibm_app  ibm_as_database  ibm_as_file  ibm_as_rmtcmd  ibm_db2
ibm_tealeaf

```

4. Filter applications by category using the **app-category category name** command. Currently, the applications contained in these categories are not displayed.

```

dmf-controller-1(config-managed-srv-app-id-filter)# app-category
<Category>          <String> : <String>
aaa                  Category selection
adult_content        Category selection
advertising          Category selection
aetls                Category selection
analytics            Category selection
anonymizer           Category selection
audio_chat           Category selection
basic                Category selection
blog                 Category selection
cdn                  Category selection
certif_auth          Category selection
chat                 Category selection
classified_ads        Category selection
cloud_services        Category selection
crowdfunding          Category selection
cryptocurrency        Category selection
db                   Category selection
dea_mail              Category selection
ebook_reader          Category selection
education             Category selection
email                 Category selection
enterprise            Category selection
file_mngt             Category selection
file_transfer         Category selection
forum                 Category selection
gaming                Category selection
healthcare            Category selection
im_mc                 Category selection
iot                   Category selection
map_service           Category selection
mm_streaming          Category selection
mobile                Category selection
networking            Category selection
news_portal           Category selection
p2p                   Category selection
payment_service        Category selection
remote_access          Category selection
scada                 Category selection
social_network         Category selection
speedtest             Category selection
standardized           Category selection
transportation         Category selection
update                Category selection
video_chat            Category selection
voip                  Category selection
vpn_tun               Category selection
web                   Category selection
web_ecom              Category selection
web_search            Category selection
web_sites             Category selection
webmail               Category selection

```

5. The **filter-mode** parameter supports two modes: forward and drop. Enter **filter-mode forward** to allow the packets to be forwarded based on the configured applications. Enter **filter-mode drop** to drop these packets.

An example of an **app-id-filter** configuration that drops all Facebook and IBM Tealeaf packets:

```
managed-service MS
service-interface switch CORE-SWITCH-1 ethernet2
!
1 app-id-filter
  app facebook
  app ibm_tealeaf
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**, or **google_maps**, adding any of these parameters to the filter list permits or drops all the packets matching after determining classification (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 the lower-layer protocols and apps. Also, when forwarding an application, packets will be dropped at the beginning of the session until the application is identified. When dropping, packets at the beginning of the session will be passed until the application is identified.

5.19.4 Using the CLI to Configure app-id and app-id-filter Combined

Follow the configuration steps described in the services earlier 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. The following example illustrates a combined **app-id-filter** and **app-id** configuration.

```
! managed-service
managed-service MS1
service-interface switch CORE-SWITCH-1 ethernet2
!
!
1 app-id-filter
  app facebook
  filter-mode forward
!
2 app-id
  collector 1.1.1.1
  l3-delivery-interface L3-INTF-1
```



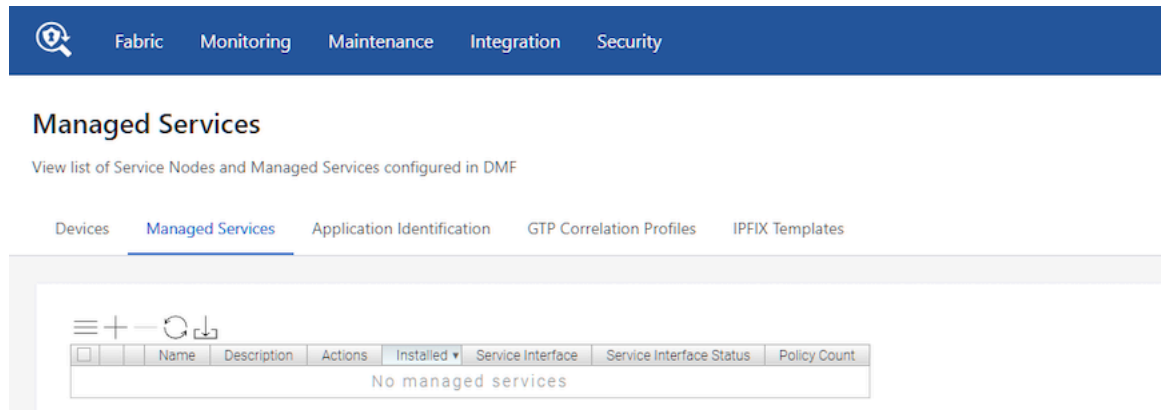
Note: The two drawbacks of this configuration are **app-id** dropping all traffic except *facebook*, and this type of service chaining can cause a performance hit and high memory utilization.

5.19.5 Using the GUI to Configure app-id and app-id-filter

App ID and **App ID Filter** are in the **Managed Service** workflow. Perform the following steps to complete the configuration.

1. Navigate to the **Monitoring > Managed Services** page. Select the table action **+** icon button to add a new managed service.

Figure 5-55: DANZ Monitoring Fabric (DMF) Managed Services



2. Configure the **Name**, **Switch**, and **Interface** inputs in the **Info** step.

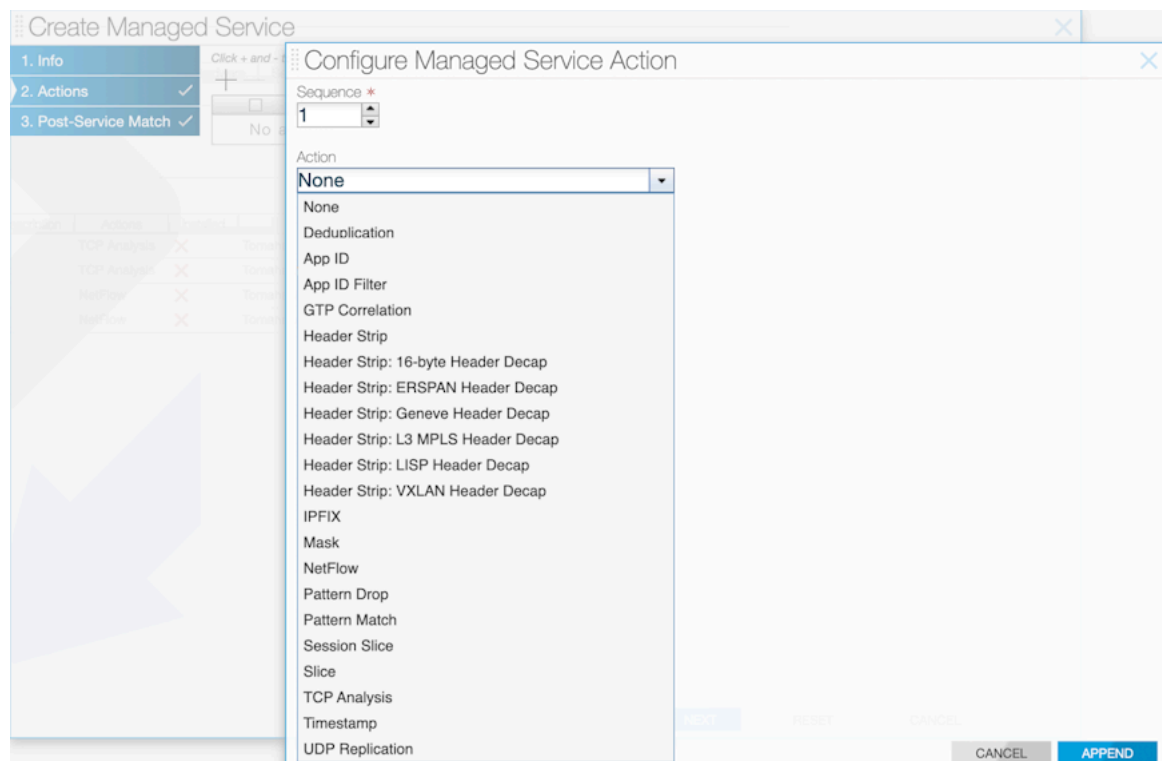
Figure 5-56: Info Step

3. In the **Actions** step, select the **+** icon to add a new managed service action.

Figure 5-57: Add App ID Action

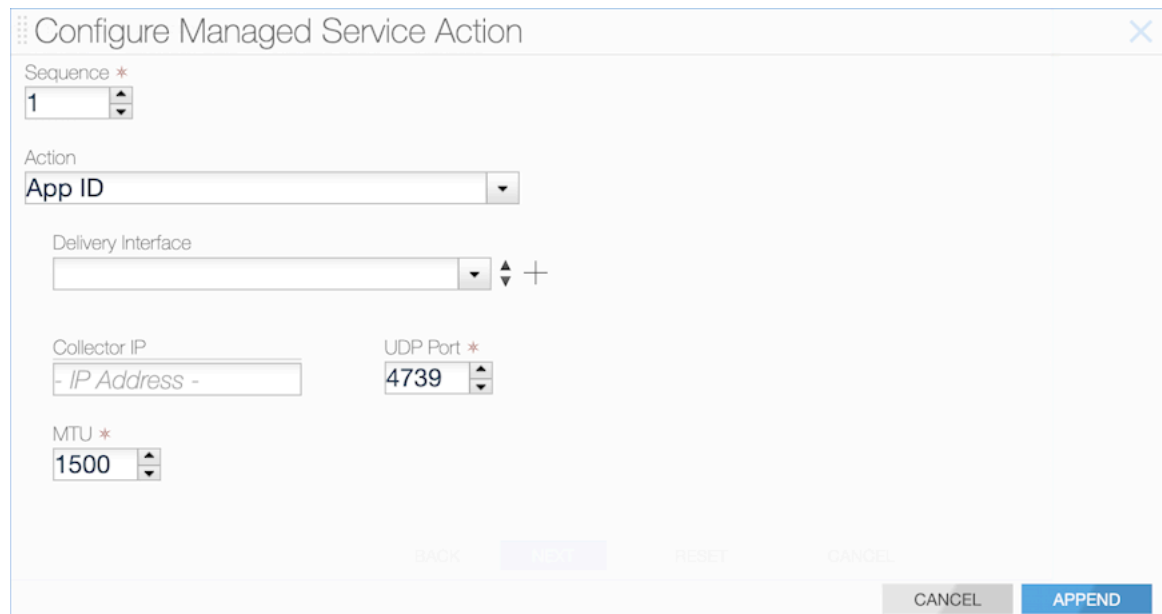
4. To Add the **App ID** Action, select **App ID** from the action selection input:

Figure 5-58: Select App ID



5. Fill in the **Delivery Interface**, **Collector IP**, **UDP Port**, and **MTU** inputs and select **Append** to include the action in the managed service:

Figure 5-59: Delivery Interface



6. To Add the **App ID Filter** Action, select **App ID Filter** from the action selection input:

Figure 5-60: Select App ID Filter

The screenshot shows the 'Configure Managed Service Action' dialog box. On the left, there's a sidebar with '1. Info', '2. Actions' (checked), and '3. Post-Service Match' (checked). The main area shows a table with columns 'Description', 'Actions', 'Installed', and 'Tomahawk'. The 'App ID Filter' action is selected in the 'Action' dropdown menu. The 'Sequence' is set to 1.

7. Select the **Filter** input as **Forward** or **Drop** action:

Figure 5-61: Select Filter Input

The screenshot shows the 'Configure Managed Service Action' dialog box. The 'Action' is set to 'App ID Filter'. The 'Filter' dropdown menu is open, showing 'Forward' and 'Drop' options. The 'Forward' option is selected. Below the filter, there are sections for 'App Categories' and 'App Names', each with a 'Click + and - to include or exclude' instruction. The 'App Categories' section shows 'No categories' and the 'App Names' section shows 'No apps'. At the bottom, there are buttons for 'BACK', 'NEXT', 'CANCEL', and 'APPEND'.

8. Use the **App Names** section to add app names.

- a. Select the **+** button to open a modal pane to add an app name.
- b. The table lists all app names. 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 5-62: Associate App Names

Associate App Names

google FILTER 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

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 5-63: Application Names

Configure Managed Service Action

Sequence *
1

Action
App ID Filter

Filter
Forward

App Categories
Click + and - to include or exclude Category Names.

App Names
Click + and - to include or exclude App Names.

Category
No categories

Name

- google
- google_accounts
- google_ads
- google_analytics
- google_api
- google_appengine
- google_blog
- google_books
- google_cache
- google_calendar

(1 - 10 / 10)

BACK SAVE CANCEL APPEND

10. Select the **Append** button to add the action to the managed service and **Save** to save the managed service.

For existing managed services, add **App ID** or **App ID Filter** using the **Edit** workflow of a managed service.

5.19.6 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 DANZ Monitoring Fabric (DMF) release. Perform such runtime service updates during a maintenance cycle. There can be issues with backward compatibility if attempting to revert to an older bundle. Adopt only supported versions. In the Controller's CLI, perform the following recommended steps:

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: Arista Networks provides the specific update file in the command example below.

2. To pull the signature file onto the Controller node, use the command:

```
dmf-controller-1(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:

```
dmf-controller-1(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 view files currently saved on the Controller node after the fetch operation is successful, use the following command:

```
dmf-controller-1(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 using the following command:

```
dmf-controller-1(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 app-id, service nodes can now 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 the **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: DMF only allows deleting 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.

```
dmf-controller-1(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**: after a successful fetch 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 complete set of applications.



Note: An analytics node cannot display these applications in the current version.

This step is only for informational purposes:

- Verify the bundle version on the service node by entering the **show service-node app-id-bundle-version** command in the service node CLI, as shown below.

Figure 5-64: 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 5-64: 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)
```

5.19.7 CLI Show Commands

Service Node

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 in use. An **app-id** or **app-id-filter** instance must be configured, or an error message is displayed.

```
dmf-servicenode-1# show app-id bundle-version
Name : bundle_version
Data : 1.680.0-22 (build date Sep 26 2023)
dmf-servicenode-1#
```

Controller

To obtain more information about the running version on a Service Node, or when the last push attempt was made and the outcome, use the following Controller CLI commands:

- **show app-id push-results optional SN name**
- **show service-node SN name app-id**

```
dmf-controller-1# show app-id push-results
# Name          IP Address      Current Version Current Push Time
  Previous Version Previous Push Time      Last Attempt Version Last
Attempt Time      Last Attempt Result Last Attempt Failure Reason
```

```

-|-----|-----|-----|-----|
-----|-----|-----|-----|
--|-----|-----|-----|-----|
---|

```

```

1 dmf-servicenode-1 10.240.180.124 1.660.2-33      2023-12-06 11:13:36.6620
00 PST 1.680.0-22      2023-09-29 16:21:11.034000 PDT 1.660.2-33
2023-12-06 11:13:34.085000 PST success

```

```

dmf-controller-1# show service-node dmf-servicenode-1 app-id
# Name IP Address Current Version Current Push Time
Previous Version Previous Push Time Last Attempt Version Last Attempt Time
Last Attempt Result Last Attempt Failure Reason

```

```

-|-----|-----|-----|-----|
-----|-----|-----|-----|
-----|-----|-----|-----|

```

```

1 dmf-servicenode-1 10.240.180.124 1.680.0-22      2023-09-29 16:21:11.034000
PDT

```

The **show app-id signature-files** command displays the validated files that are available to push to Service Nodes.

```

dmf-controller-1# show app-id signature-files
# Signature-file Checksum Fetch time
-|-----|-----|-----|
1 file1.tar.gz abcdefgh12345 2023-08-02 22:20:49.422000 UTC
2 file2.tar.gz ijklmnop67890 2023-08-03 07:10:22.123000 UTC

```

The **show analytics app-info filter-interface-name** command displays aggregated information over the last 5 minutes about the applications seen on a given filter interface, sorted by unique flow count. This command also has an optional **size** option to limit the number of results, default is all.



Note: This command only works in **push-per-filter** mode.

```

dmf-controller-1# show analytics app-info filter-interface f1 size 3
# App name Flow count
-|-----|-----|
1 app1 1000
2 app2 900
3 app3 800

```

5.19.8 Syslog Messages

Syslog messages for configuring the **app-id** and **app-id-filter** services appear in a service node's syslog through **journalctl**.

A Service Node syslog registers events for the **app-id** **add**, **modify**, and **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.
```

The addition, modification, or 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 register in **floodlight.log**.

5.19.9 Scale

- Max concurrent sessions are currently set to permit less than 200,000 active flows per core. Performance may drop the more concurrent flows there are. This value is a maximum value to prevent the service from overloading. Surpassing this threshold may cause some flows not to be processed, and the new flows will not be identified or filtered. Entries for inactive flows will time out after a few minutes for ongoing sessions and a few seconds after the session ends.
- If there are many inactive sessions, DMF holds the flow contexts, reducing the number of available flows used for DPI. The timeouts are approximately 7 minutes for TCP sessions and 1 minute for UDP.
- Heavy application traffic load degrades performance.

5.19.10 Troubleshooting

- If IPFIX reports do not appear on an Analytics Node (AN) or Collector, ensure the UDP port is configured correctly and verify the AN receives traffic.
- If the **app-id-filter** app list does not appear, ensure a Service Node (SN) is connected using the **show service-node** command on the Controller.
- For **app-id-filter**, enter at least one valid application from the list that appears using <Tab>. If not, the policy will fail to install with an error message `app-id-filter specified without at least one name TLV identifying application`.
- 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 or dropped.
- An IPFIX element, BSN type 14, that existed in DMF version 8.4 was removed in 8.6.
- During a dynamic signature update, if a SN reboot occurs, it will likely boot up with the correct version. To avoid issues of traffic loss, perform the update during a maintenance window. Also, during an update, the SN will temporarily not send LLDP packets to the Controller and disconnect for a short while.
- After a dynamic signature update, do not change configurations or push another signature file for several minutes. The update will take some time to process. If there are any VFT changes, it may lead to warning messages in floodlight, such as:

```
Sync job 2853: still waiting after 50002 ms
Stuck switch update: R740-25G[00:00:e4:43:4b:bb:38:ca], duration=50002ms,
stage=COMMIT
```

These messages may also appear when configuring DPI on a large number of ports.

5.19.11 Limitations

- When using a drop filter, a few packets may slip through the filter before determining an application ID for a flow, and when using a forward filter, a few packets may not be forwarded. Such a small amount is estimated to be between 1 and 6 packets at the beginning of a flow.

- When using a drop filter, add the **unknown** app ID to the filter list to drop any unidentified traffic if these packets are unwanted.
- The Controller must be connected to a Service Node for the **app-id-filter** app list to appear. If the list does not appear and the application names are unknown, use the **app-id** to send reports to the analytics node. Use the application names seen there to configure an **app-id-filter**. The name must match exactly.
- Since **app-category** does not currently show the applications included in that category, do not use it when targeting specific apps. Categories like basic, which include all basic networking protocols like TCP and UDP, may affect all flows.
- For **app-id**, a report is only generated for a fully classified flow after that flow has been fully classified. Therefore, the number of reported applications may not match the total number of flows. These reports are sent after enough applications are identified on the Service Node. If many applications are identified, DMF sends the reports quickly. However, DMF sends these reports every 10 seconds when identifying only a few applications.
- DMF treats a bidirectional flow as part of the same n-tuple. As such, generated reports contain the client's source IP address and the server's destination IP address.
- While configuring many ports with the **app-id**, there may occasionally be a few Rx drops on the 16 port machines at a high traffic rate in the first couple of seconds.
- The feature uses a cache that maps dest ip and port to the application. Caching may vary the performance depending on the traffic profile.
- 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.
- At scale, such as configuring 16 ports on the R740 DCA-DM-SEL, **app-id** may take a few minutes to set up on all these ports, and this is also true when doing a dynamic signature update.
- The **show analytics app-info** command only works in **push-per-filter** VLAN mode.

5.20 Redundancy of Managed Services Using Two DMF Policies

In this method, users can employ 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 DANZ Monitoring Fabric (DMF) policy. In this case, the backup policy (and, by extension, the backup service) will always be active but only receive relevant traffic once the primary policy goes down. This method provides true redundancy at the policy, service-node, and core switch levels but uses additional network and node resources.

Example

```
! 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

```

5.21 Cloud Services Filtering

The DANZ Monitoring Fabric (DMF) supports traffic filtering to specific services hosted in the public cloud and redirecting filtered traffic to customer tools. DMF achieves this functionality by reading the source and destination IP addresses of specific flows, identifying the Autonomous System number they belong to, tagging the flows with their respective AS numbers, and redirecting them 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 configuring the last rule as **match default**. It instructs the DMF Service Node to drop packets when either of the following conditions occurs:

- The stream's source IP address or destination IP address doesn't belong to any AS number.
- The stream's source IP address or destination IP address is affiliated with an AS number but has no specific action set.

5.21.1 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
```

5.22 Multiple Services Per Service Node Interface

The service-node capability is augmented to support more than one service action per service-node interface. Though this feature is economical regarding per-interface cost, it could cause packet drops in high-volume traffic environments. Arista Networks recommends using this feature judiciously.

Example

```
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 14-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: After upgrading to **DANZ Monitoring Fabric (DMF) release 8.1.0** and later, the service-action CLI is 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	-

When using the DMF GUI, view the above information in **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: The following limitations apply to this mode of configuration:

- The NetFlow/IPFIX-action configuration should not be followed by the timestamp service action.
- Ensure the UDP-replication action configuration is 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.
- When configuring a header strip and slice action, the header strip action must precede the slice action.

5.23 Sample Service

The Service Node forwards packets based on the **max-tokens** and **tokens-per-refresh** parameters using the DANZ Monitoring Fabric (DMF) Sample Service feature. The sample service uses one token to forward one packet.

After consuming all the initial tokens from the **max-tokens** bucket, the system drops subsequent packets until the **max-tokens** bucket refills using the **tokens-per-refresh** counter at a recurring predefined time interval of 10ms. Packet sizes do not affect this service.

Arista Networks recommends keeping the **tokens-per-refresh** value at or below **max-tokens**. For example, **max-tokens = 1000** and **tokens-per-refresh = 500**.

Setting the **max-tokens** value to 1000 means that the initial number of tokens is 1000, and the maximum number of tokens stored at any time is 1000.

The **max-tokens** bucket will be zero when the Service Node has forwarded 1000 packets before the first 10 ms period ends, leading to a situation where the Service Node is no longer forwarding packets. After every 10ms time interval, if the **tokens-per-refresh** value is set to 500, the **max-tokens** bucket is refilled using the **tokens-per-refresh** configured value, 500 tokens in this case, to pass packets the service tries to use immediately.

Suppose the traffic rate is higher than the refresh amount added. In that case, available tokens will eventually drop back to 0, and every 10ms, only 500 packets will be forwarded, with subsequent packets being dropped.

If the traffic rate is lower than the refresh amount added, a surplus of tokens will result in all packets passing. Since the system only consumes some of the tokens before the next refresh interval, available tokens will accumulate until they reach the **max-tokens** value of 1000. After 1000, the system does not store any surplus tokens above the **max-tokens** value.

To estimate the maximum possible packets passed per second (pps), use the calculation $(1000\text{ms}/10\text{ms}) * \text{tokens-per-refresh}$ and assume the **max-tokens** value is larger than **tokens-per-refresh**. For example, if the **tokens-per-refresh** value is 5000, then 500000 pps are passed.

The Sample Service feature can be used as a standalone Managed Service or chained with other Managed Services.

Use Cases and Compatibility

- Applies to Service Nodes
- Limit traffic to tools that cannot handle a large amount of traffic.
- Use the Sample Service before another managed service to decrease the load on that service.
- The Sample Service is applicable when needing only a portion of the total packets without specifically choosing which packets to forward.

5.23.1 Sample Service CLI Configuration

1. Create a managed service and enter the service interface.
2. Choose the sample managed service with the **seq num sample** command.
 - a. There are two required configuration values: **max-tokens** and **tokens-per-refresh**. There are no default values, and the service requires both values.
 - b. The **max-tokens** value is the maximum size of tokens in the token bucket. The service will start with the number of tokens specified when first configured. Each packet passed consumes one token. If no tokens remain, packet forwarding stops. Configure the **max-tokens** value from a range of **1** to the maximum uint64 (unsigned integer) value of **9,223,372,036,854,775,807**.
 - c. DMF refreshes the token bucket every 10 ms. The **tokens-per-refresh** value is the number of tokens added to the token bucket on each refresh. Each packet passed consumes one token, and when the number of tokens drops to zero, the system drops all subsequent packets until the next refresh. The number of tokens in the bucket cannot exceed the value of **max-tokens**. Configure the **tokens-per-refresh** value from a range of **1** to the maximum uint64 (unsigned integer) value of **9,223,372,036,854,775,807**.

The following example illustrates a typical Sample Service configuration

```
dmf-controller-1(config-managed-srv-sample)# show this
! managed-service
managed-service MS
!
3 sample
max-tokens 50000
tokens-per-refresh 20000
```

3. Add the managed service to the policy.

Show Commands

Use the **show running-config managed-service sample_service_name** command to view pertinent details. In this example, the **sample_service_name** is **techpubs**.

```
DMF-SCALE-R450> show running-config managed-service techpubs

! managed-service
managed-service techpubs
```

```

!
1 sample
  max-tokens 1000
  tokens-per-refresh 500
DMF-SCALE-R450>

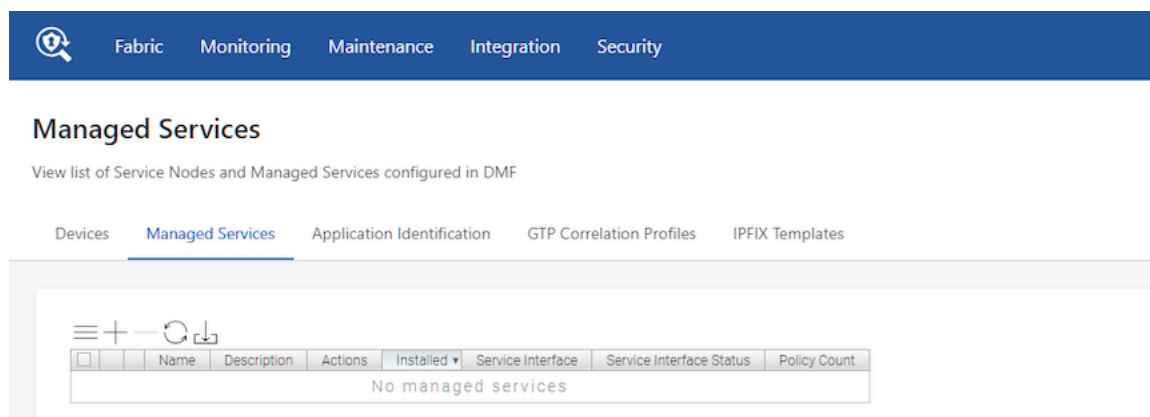
```

5.23.2 Sample Service GUI Configuration

Use the following steps to add a Sample Service.

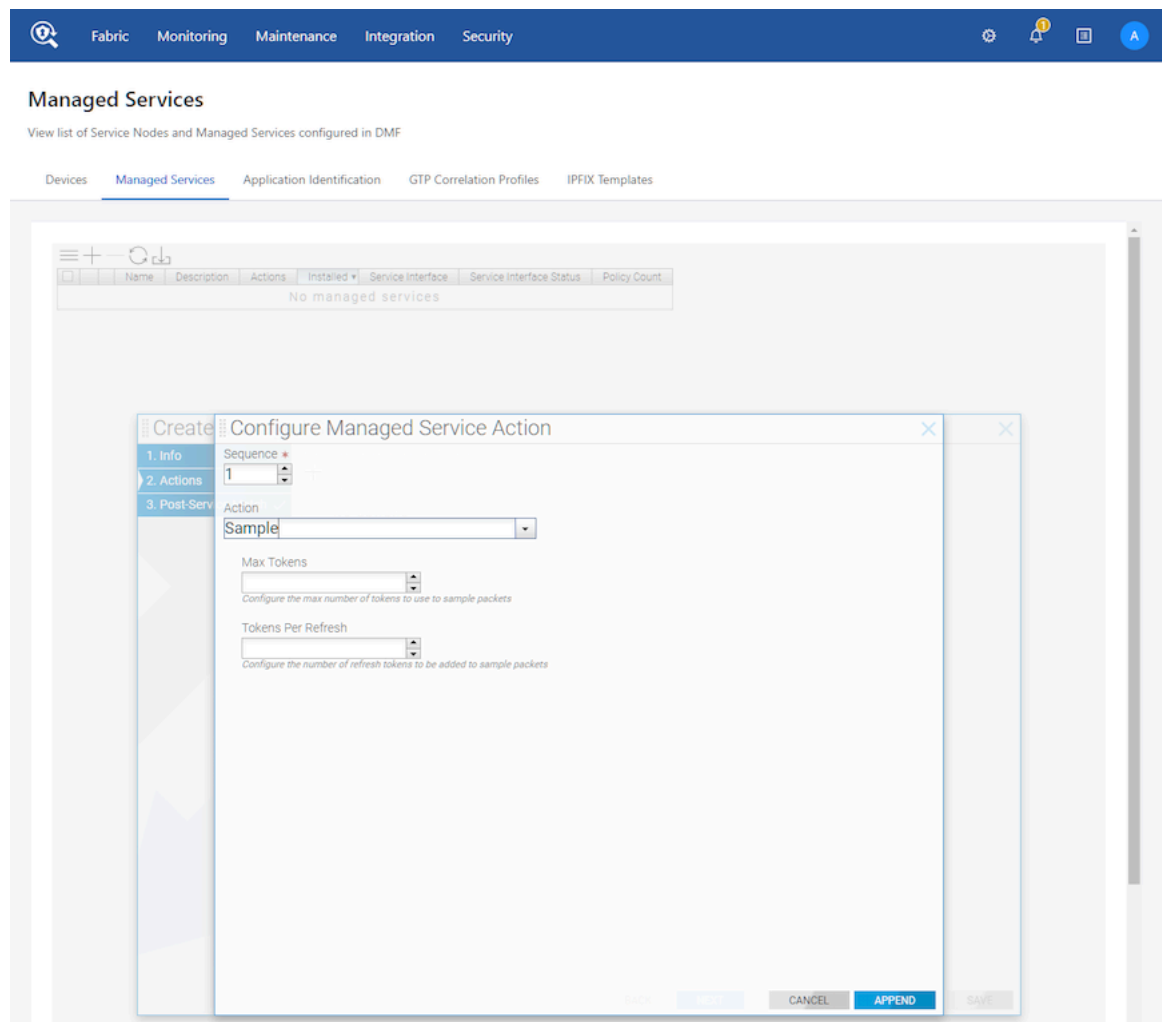
1. Navigate to the **Monitoring > Managed Services** page.

Figure 5-66: DMF Managed Services



2. Under the **Managed Services** section, click the **+** icon to create a new managed service. Go to the **Actions**, and select the **Sample** option in the **Action** drop-down. Enter values for **Max tokens** and **Tokens per refresh**.

Figure 5-67: Configure Managed Service Action



3. Click **Append** and then **Save**.

5.23.3 Troubleshooting Sample Service

Troubleshooting

- If the number of packets forwarded by the Service Node interfaces is few, the **max-tokens** and **tokens-per-refresh** values likely need to be higher.
- If fewer packets than the **tokens-per-refresh** value forward, ensure the **max-tokens** value is larger than the **tokens-per-refresh** value. The system discards any surplus refresh tokens above the **max-tokens** value.
- When all traffic forwards, the initial **max-tokens** value is too large, or the tokens refreshed by **tokens-per-refresh** are higher than the packet rate.
- When experiencing packet drops after the first 10ms post commencement of traffic, it may be due to a low **tokens-per-refresh** value. For example, calculate the minimum value of **max-tokens** and **tokens-per-refresh** that would lead to forwarding all packets.

Calculation Example

```
Traffic Rate : 400 Mbps
Packet Size - 64 bytes
400 Mbps = 400000000 bps
400000000 bps = 500000000 Bps
500000000 Bps = 595238 pps (Includes 20 bytes of inter packet gap in addition to
the 64 bytes)
1000 ms = 595238 pps
1 ms = 595.238 pps
10 ms = 5952 pps
max-tokens : 5952 (the minimum value)
tokens-per-refresh : 5952 ( the minimum value)
```

5.23.4 Limitations

- In the current implementation, the Service Sample action is bursty. The token consumption rate is not configured to withhold tokens over time, so a large burst of incoming packets can immediately consume all the tokens in the bucket. There is currently no way to select what traffic is forwarded or dropped; it only depends on when the packets arrive concerning the refresh interval.
- Setting the **max-tokens** and **tokens-per-refresh** values too high will forward all packets. The maximum value is 9,223,372,036,854,775,807, but Arista Networks recommends staying within the maximum values stated under the description section.

5.24 Flow Diff Latency and Drop Analysis

Latency and drop information help determine if there is a loss in a particular flow and where the loss occurred. A Service Node action configured as a DANZ Monitoring Fabric (DMF) managed service has two separate taps or spans in the production network and can measure the latency of a flow traversing through these two points. It can also detect packet drops between two points in the network if the packet only appears on one point within a specified time frame, currently set to 100ms.

Latency and drop analysis require PTP time-stamped packets. The DMF PTP timestamping feature can do this as the packets enter the monitoring fabric, or the production network switches can also timestamp the packet.

The Service Node accumulates latency values by flow and sends IPFIX data records with each flow's 5-tuple and ingress and egress identifiers. It sends IPFIX data records to the Analytics Node after collecting a specified number of values for a flow or when a timeout occurs for the flow entry. The threshold count is 10,000, and the flow timeout is 4 seconds.



Note: Only basic statistics are available: **min**, **max**, and **mean**. Use the Analytics Node to build custom dashboards to view and check the data.

Use the DMF Analytics Node to build custom dashboards to view and check the data.



Attention: The flow diff latency and drop analysis feature is switch dependent and requires PTP timestamping. It is supported on 7280R3 switches.

5.24.1 Configure Flow Diff Latency and Drop Analysis Using the CLI

Configure this feature through the Controller as a Service Node action in managed services using the managed service action **flow-diff**.

Latency configuration configures multiple tap point pairs to analyze latency or drops between which analysis of latency or drops occurs. A tap point pair comprises a source and a destination tap point, identified by the filter interface, policy name, or filter interface group. Based on the latency configuration, configuring the Service Node with traffic metadata tells the Service Node where to look for tap point information, timestamps, and the IPFIX collector.

Configure appropriate DMF Policies such that traffic tapped from tap point pairs in the network is delivered to the configured Service Node interface for analysis.

Configuration Steps for `flow-diff`.

1. Create a managed service and enter the service interface.
2. Choose the `flow-diff` service action with the command: `seq num flow-diff`



Note: The command should enter the `flow-diff` submode, which supports three configuration parameters: **collector**, **l3-delivery-interface**, and **tap-point-pair**. These all are required parameters.

3. Configure the IPFIX collector IP address by entering the following command: `collector ip-address` (the UDP port and MTU parameters are optional; the default values are 4739 and 1500, respectively).
4. Configure the delivery interface by entering the command `l3-delivery-interface delivery interface name`.
5. Configure the points for flow-diff and drop analysis using `tap-point-pair` parameters as specified in the following section. Multiple options to identify the tap-point include `filter-interface`, `filter-interface-group`, and `policy-name`. This command will require a source and a destination tap point.
6. Optional parameters are `latency-table-size`, `sample-count-threshold`, `packet-timeout`, and `flow-timeout`. The default values are large, 10000, 100 ms, and 4000 ms, respectively.
7. `latency-table-size` determines the memory footprint of flow-diff action on the Service Node.
8. `sample-count-threshold` specifies the number of samples needed to generate a latency report. Every time a packet times out, it generates a sample for that flow. DMF generates a report if the flow reaches the sample threshold and resets the flow stats.
9. `packet-timeout` is the time interval in which timestamps are collected for a packet. It must be larger than the time it takes the same packet to appear at all tap points. Every timeout generates a sample for the flow associated with the packet.
10. `flow-timeout` is the time after which, when the flow no longer receives any packets, the flow will be evicted, and a report for the flow is generated. The timeout for a flow refreshes each time a new packet is received. If packets are continuously received for a flow below the flow timeout value, then the flow will never be evicted.

The following example illustrates configuring `flow-diff` using the steps mentioned earlier:

```
dmf-controller-1(config)# managed-service managed_service_1
dmf-controller-1(config-managed-srv)# service-interface switch delivery1
ethernet1
dmf-controller-1(config-managed-srv)# 1 flow-diff
dmf-controller-1(config-managed-srv-flow-diff)# collector 192.168.1.1
dmf-controller-1(config-managed-srv-flow-diff)# l3-delivery-interface l3-
iface-1
dmf-controller-1(config-managed-srv-flow-diff)# tap-point-pair source filter-
interface f1 destination filter-interface f2
dmf-controller-1(config-managed-srv-flow-diff)# latency-table-
size small|medium|large
dmf-controller-1(config-managed-srv-flow-diff)# packet-timeout 100
dmf-controller-1(config-managed-srv-flow-diff)# flow-timeout 4000
dmf-controller-1(config-managed-srv-flow-diff)# sample-count-threshold 10000
```


5.24.2 Configuring Tap Points

Configure tap points using **tap-point-pair** parameters in the **flow-diff** submode specifying three identifiers: **filter interface name**, **policy name**, and **filter-interface-group**.

```
dmf-controller-1(config-managed-srv-flow-diff)# tap-point-pair source <Tab>
filter-interface          policy-name          filter-interface-group
```

Both **source** and **destination** tap points must use identifiers compatible; **policy-name** cannot be used with **filter-interface** or **filter-interface-group**.

The **filter-interface-group** option takes in any configured filter interface group used to represent a collection of tap points in **push-per-filter** mode. This is an optional command to use when a group of tap points exists, all expecting traffic from the same source or group of source tap points for ease of configuration. For example:

1. Instead of having two separate **tap-point-pairs** to represent $A \rightarrow B$, $A \rightarrow C$, use a **filter-interface-group** $G = [B, C]$, and only one tap-point-pair $A \rightarrow G$.

```
dmf-controller-1(config-managed-srv-flow-diff# tap-point-pair source type A
destination filter-interface-group G
```

2. With a topology like $A \rightarrow C$ and $B \rightarrow C$, configure a **filter-interface-group** $G = [A, B]$, and **tap-point-pair** $G \rightarrow C$.

```
dmf-controller-1(config-managed-srv-flow-diff # tap-point-pair source
filter-interface-group G destination type C
```

There are some restrictions to keep in mind while configuring **tap-point-pairs**:

- **source** and **destination** must exist and cannot refer to the same tap point
- You can only configure a maximum of 1024 tap points, therefore a maximum of 512 tap-point-pairs. These limits are not for a single flow-diff managed service but across all managed services with flow-diff action.
- **filter-interface-group** must not overlap with other groups within the same managed service and cannot have more than 128 members.
- When configuring multiple **tap-point-pairs** using **filter-interface** and **filter-interface-group**, an interface part of a filter interface group cannot be used simultaneously as an individual source and destination within the same managed service.

There are several caveats on what is accepted based on the VLAN mode:

Push per Filter

- Use the **filter-interface** option to provide the filter interface name.
- Use the **filter-interface-group** option to provide the filter interface group name.
- The **policy-name** identifier is invalid in this mode.

Push per Policy

- Accepts a **policy-name** identifier as a tap point identifier.
- The **filter-interface** and **filter-interface-group** identifiers are invalid in this mode.
- Policies configured as **source** and **destination** tap points within a **tap-point-pair** must not overlap.

5.24.3 Configuring Policy

Irrespective of the VLAN mode, configure a policy or policies so that the same packet can be tapped from two independent points in the network and then sent to the Service Node.

After creating a policy, add the managed service with flow-diff action as shown below:

```
dmf-controller-1 (config-policy)# use-managed-service service name sequence 1
```

There are several things to consider while configuring policies depending on the VLAN mode:

Push per Filter

- Only one policy can contain the **flow-diff** service action.
- A policy should have all **filter-interfaces** and **filter-interface-groups** configured as tap points in the flow-diff configuration. Missing filter interfaces and groups from policy may result in drops being reported when in reality we simply won't be forwarding the packets from one end of the **tap-point-pair** to the Service Node.
- It's also advisable to not add any filter interfaces (groups) that are not in the tap-point-pairs as their latency and drop analysis will not be done, so it will cause unnecessary packets forwarded to the Service Node, which will be reported as **unexpected**.

Push per Policy

- Add the flow-diff service to two policies when using **policy-name** as source and destination identifiers. In this case, there are no restrictions on how many filter interfaces a policy can have.
- A policy configured as one of the tap points will fail if it overlaps with the other policy in a tap-point-pair, or if the other policy does not exist.

In both VLAN modes, policies must have PTP timestamping enabled. To do so, use the following command:

```
dmf-controller-1 (config-policy)# use-timestamping
```

Configuring PTP Timestamping

This feature depends on configuring PTP timestamping for the packet stream going through the tap points. Refer to the [Resources](#) section for more information on setting up PTP timestamping functionality.

5.24.4 Show Commands

The following **show** commands provide helpful information.

The **show running-config managed-service managed service** command helps check whether the **flow-diff** configuration is complete.

```
dmf-controller-1(config)# show running-config managed-service flow-diff
! managed-service
managed-service flow-diff
service-interface switch DCS-7050CX3-32S ethernet2/4
!
1 flow-diff
collector 192.168.1.1
l3-delivery-interface AN-Data
tap-point-pair source filter-interface ingress destination filter-interface
egress
```

The **show managed-services managed service** command provides status information about the service.

```
dmf-controller-1(config)# show managed-services flow-diff
# Service Name Switch      Switch Interface Installed Max Post-Service BW Max Pre-Service BW Total Post-Service BW Total Pre-
Service BW
-----|-----|-----|-----|-----|-----|-----|-----|
1 flow-diff  DCS-7050CX3-32S ethernet2/4      True      25Gbps      25Gbps      624Kbps      432Mbps
```

The **show running-config policy *policy*** command checks whether the policy flow-diff service exists, whether **use-timestamping** is enabled, and the use of the correct filter interfaces.

The **show policy *policy*** command provides detailed information about a policy and whether any errors are related to the flow-diff service. The Service Interfaces tab section shows the packets transmitted to the Service Node and IPFIX packets received from the Service Node.



Note: Regarding two policies in a tap-point-pair (one source and one destination), only one policy will show stats about packets received from the Service Node. This output is by design, as only a single VLAN exists for the IPFIX packets transmitted from the Service Node, so there can only be one policy.

```
dmf-controller-1 (config)# show policy flow-diff-1
Policy Name          : flow-diff-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 : 1
# of filter interfaces : 1
# of delivery interfaces : 1
# of core interfaces : 4
# of services : 1
# of pre service interfaces : 1
# of post service interfaces : 1
Push VLAN           : 2
Post Match Filter Traffic : 215Mbps
Total Delivery Rate  : -
Total Pre Service Rate : 217Mbps
Total Post Service Rate : -
Overlapping Policies : none
Component Policies   : none
Runtime Service Names : flow-diff
Installed Time        : 2023-11-16 18:15:27 PST
Installed Duration    : 19 minutes, 45 secs
~ Match Rules ~
# Rule
-|-----|
1 1 match any
~~~~~
# DMF IF Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 BP1 7280SR3E Ethernet25 up rx 24319476 27484991953 23313 215Mbps 2023-11-16 18:18:18.837000 PST
~~~~~
# DMF IF Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 AN-Data 7050SX3-1 ethernet41 up tx 81 117222 0 - 2023-11-16 18:18:18.837000 PST
~~~~~
# Service name Role Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
-----|
1 flow-diff pre-service DCS-7050CX3-32S ethernet2/4 up tx 23950846 27175761734 23418 217Mbps 2023-11-16 18:18:18.837000 PST
2 flow-diff post-service DCS-7050CX3-32S ethernet2/4 up rx 81 117546 0 - 2023-11-16 18:18:18.837000 PST
~~~~~
# Switch IF Name State Dir Packets Bytes Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 7050SX3-1 ethernet7 up rx 23950773 27175675524 23415 217Mbps 2023-11-16 18:18:18.837000 PST
2 7050SX3-1 ethernet56 up rx 81 117222 0 - 2023-11-16 18:18:18.837000 PST
3 7050SX3-1 ethernet56 up tx 23950773 27175675524 23415 217Mbps 2023-11-16 18:18:18.837000 PST
4 7280SR3E Ethernet7 up tx 24319476 27484991953 23313 215Mbps 2023-11-16 18:18:18.837000 PST
5 DCS-7050CX3-32S ethernet28 up tx 81 117546 0 - 2023-11-16 18:18:18.837000 PST
6 DCS-7050CX3-32S ethernet28 up rx 23950846 27175761734 23418 217Mbps 2023-11-16 18:18:18.837000 PST
~ Failed Path(s) ~
None.
```

5.24.5 Syslog Messages

The Flow Diff Latency and Drop Analysis feature does not create Syslog messages.

5.24.6 Troubleshooting

Controller

Policies dictate how and what packets are directed to the Service Node. Policies must be able to stream packets from two distinct tap points so that the same packet gets delivered to the Service Node for flow-diff and drop analysis.

Possible reasons for flow-diff and drop analysis not working are:

- flow-diff action is not added to both policies in a **tap-point-pair** in **push-per-policy** mode.
- flow-diff action exists in a policy that does not have all filter interfaces or groups configured as **tap-point-pairs** in **push-per-filter** mode.

A policy programmed to use flow-diff service action can fail for multiple reasons:

- An incomplete flow-diff configuration exists, there are missing **source** or **destination** tap points, or the use of a **policy-name** identifier in **push-per-filter** or **filter-interface** or **filter-interface-group** in **push-per-policy** mode.
- In the **push-per-policy** mode, the policy doesn't match the **source** or **destination** tap point for any configured **tap-point-pair**.
- In the **push-per-policy** mode, policy names configured as tap points within a **tap-point-pair** overlap or do not exist.
- There are more than 512 **tap-point-pairs** configured or more than 1024 distinct tap points. These limits are global across all flow-diff managed services.
- **filter-interface-groups** overlap with each other within the same managed service or have more than 128 group members
- **filter-interface** is being used individually as a tap point and as a part of some **filter-interface-group** within the same managed service

Reasons for failure are available in the runtime state of the policy and viewed using the **show policy *policy name*** command.

After verifying the correct configuration of the policy and flow-diff action, enabling the log debug mode and reviewing the floodlight logs (use the **fl-log** command in bash mode) should show the ipfix-collector, traffic-metadata, and flow-diff gentable entries sent to the Service Node.

If no computed latency reports are generated, this can mean two things. First, not reaching the **sample-count-threshold** value. Either lower the **sample-count-threshold** until reports are generated or increase the number of unique packets per flow. Second, there are no evicted flows. The time specified in **flow-timeout** will refresh every time a new packet is received on that flow before the timeout occurs. If a flow continuously receives packets, it will never time out. Lower the **flow-timeout** value if the default of 4 seconds causes flows not to expire. After a flow expires, the Controller generates a report for that flow.

For **packet-timeout**, this value must be larger than the time expected to receive the same packet on every tap-point.

For A->B, if it takes 20ms for the packet to appear at B, **packet-timeout** must be larger than this time frame to collect timestamps for both these tap points and compute a latency in one sample.

If the Controller generates many unexpected reports, ensure taps are in the correct order in the Controller configuration. If there are many drop reports, ensure the same packet is received at all relevant taps within the **packet-timeout** window.

5.24.7 Limitations

- Only 512 **tap-point-pairs** and 1024 distinct tap points are allowed.
- **filter-interface-group** used as a tap point must not overlap with any other group within the same managed service and must not have more than 128 members.

- A filter interface cannot be used individually as a tap point and as a part of a **filter-interface-group** simultaneously within the same managed service.
- There is no chaining if a packet flows through three or more tap points in an A->B->C->...->Z topology. The only computed latency reports will be for `tap-point-pairs` A->B, A->C, ..., and A->Z if these links are specified, but B->C, C->D, etc., will not be computed.
- Hardware RSS firmware in the Service Node currently cannot parse L2 header timestamps, so packets for all L2 header timestamps are sent to the same lcore; however, RSS does distribute packets correctly to multiple lcores when using `src-mac` timestamping.
- PTP timestamping doesn't allow **rewrite-dst-mac**, so filter interfaces cannot be used as tap points in **push-per-policy** mode.
- In **push-per-policy**, a policy can only have one filter interface.
- Each packet from the L3 header and onwards gets hashed to a 64-bit value; if two packets hash to the same value, we assume the underlying packets are the same.
- Currently, on the flow-diff action in the Service Node, if packets are duplicated so that N copies of the same packet are received:
 - N-1 latencies are computed.
 - The ingress identifier is the earliest timestamp.
- The system reports timestamps as unsigned 32-bit values, with the maximum timestamp being $2^{32}-1$, corresponding to approximately 4.29 seconds.
- Only min/mean/max latencies are currently reported.
- Packets are hashed from the L3 header onwards, meaning if there is any corrupted data past the L3 header, it will lead to drop reports. The same packet must appear at two tap points to generate a computed latency report.
- In A->B, if B packets appear before A, an **unexpected** type report is generated.
- At whichever tap point the packet first appears with the earliest timestamp, it is considered the source.
- While switching the latency configuration, the system may generate a couple of unexpected or drop reports at the beginning.
- Users must have good knowledge of network topology when setting up tap points and configuring timeout or sample threshold values. Improper configuration may lead to drop or unexpected reports.

5.24.8 Resources

- **PTP Timestamping** (<https://www.arista.com/en/support/toi/dmf-8-4-0/18066-ptp-timestamping-on-7280-switches>)

5.25 Service Node Management Migration L3ZTN

After the first boot (initial configuration) is completed, in the case of Layer-3 topology mode an administrator can move a Service Node (SN) from an old DANZ Monitoring Fabric (DMF) Controller to a new one via the CLI.



Note: For appliances to connect to the Controller in Layer-3 Zero Touch Network (L3ZTN) mode, you must configure the Controller's deployment mode as `pre-configure`.

To migrate a Service Node's management to a new Controller, follow the steps outlined below:

1. Remove the Service Node from the old Controller using the following command:

```
controller-1(config)# no service-node service-node-1
```

2. Connect the data NICs' `sni` interfaces to the new core fabric switch ports.

- SSH to the Service Node and configure the new Controller's IP address using the **zerotouch l3ztn controller-ip** command:

```
service-node-1(config)# zerotouch l3ztn controller-ip 10.2.0.151
```

- Get the management MAC address (of interface **bond0**) of the Service Node using the following command:

```
service-node-1(config)# show local-node interfaces
```

Interfaces						
Interface	Master	Hardware address	Permanent hardware address	Operstate	Carrier	Bond mode
bond0		78:ac:44:94:2b:b6 (Dell)		up	up	active-backup

- Add the Service Node and its **bond0** interface's MAC address (obtained in the step above) to the new Controller:

```
controller-2(config)# service-node service-node-1
controller-2(config-service-node)# mac 78:ac:44:94:2b:b6
```

- After associating the Service Node with the new Controller, reboot the Service Node.
- Once the Service Node is back online, the Controller should receive a ZTN request. If the Service Node's image differs from the Service Node image file on the new Controller, the mismatch triggers the Service Node to perform an auto-upgrade of the image and to reboot twice.

```
controller-2# show zerotouch request
```

41	78:ac:44:94:2b:b6 (Dell)	10.240.156.10	get-manifest	2024-06-12 23:14:42.284000	UTC ok	The request has succeeded
56	24:6e:96:78:58:b4 (Dell)	10.240.156.91	get-manifest	2024-06-12 23:13:38.633000	UTC ok	The request has succeeded

- Then the Service Node should appear as a member of the new DMF fabric, which you can verify by using the following command:

```
controller-2# show service-node service-node-1 details
```

Using the DMF Recorder Node

This chapter describes configuring the DANZ Monitoring Fabric (DMF) Recorder Node (RN) 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 DANZ Monitoring Fabric (DMF) Recorder Node (RN) integrates with the DMF for single-pane-of-glass monitoring. A single DMF Controller can manage multiple RNs, delivering packets for recording through Out-of-Band policies. The DMF Controller also provides central APIs for packet queries across one or multiple RNs and for viewing errors, warnings, statistics, and the status of connected RNs.

A DMF out-of-band policy directs matching packets for recording to one or more RNs. An RN interface identifies the switch and port used to attach the RN 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 RN interfaces.

6.2 Configuration Summary

At a high level, follow the below three steps to use the Recorder Node (RN).

Step 1: Define the RN.

Step 2: Define a DANZ Monitoring Fabric (DMF) policy to select the traffic to forward to the RN.

Step 3: View and analyze the recorded traffic.

The RN configuration on the DMF Controller includes the following:

- **Name:** Each RN requires a unique name among recorder nodes in the connected fabric. Removing the name removes the entire configuration for the given RN.
- **Management MAC address:** Each RN must have a unique management interface MAC address in the connected fabric.
- **Packet removal policy:** Defines the behavior when the RN disks reach capacity. The default policy causes the overwriting of the earliest recorded packets by the most recent packets. The other option is to stop recording and wait until space is available.
- **Record enable, or Record disable:** DMF enables packet recording by default but can be enabled or disabled for a specific RN.
- **Static auth tokens:** Static auth tokens are pushed to each RN 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 RN treats the controller as an ordinary client and requires it to present valid credentials as an authentication token. The DMF Controller authentication token is automatically generated and resettable upon request.
- **Pre-buffer:** This buffer, defined in minutes, is used for proactive network monitoring without recording and retaining unnecessary packets. Once the buffer is full, DMF deletes the oldest packets.
- **Maximum disk utilization:** This defines the maximum disk utilization as a percentage between 5% and 95%. After reaching the configured utilization, DMF enforces the packet removal policy. The default maximum disk utilization is 95%.
- **Maximum packet age:** This defines the maximum age in minutes of any packet in the RN. Use it with the packet removal policy to control when packets are deleted based on age rather than disk utilization alone. When not set, DMF does not enforce the maximum packet age and keeps packets until the maximum disk utilization is reached.

6.3 Indexing Configuration

The Recorder Node (RN) indexing configuration defines the fields used to query packets on the RN. By default, DMF enables all indexing fields in the indexing configuration. Selectively disable the specific indexing fields not required in RN 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. Arista recommends disabling unnecessary indexing fields.

The RN 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: Enable the Outer VLAN ID indexing field to query the RN using a DANZ Monitoring Fabric (DMF) policy name or a DMF filter interface name.

To understand leveraging an indexing configuration, 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, saving only transport ports as metadata for each packet recorded. This technique greatly reduces per-packet index space consumption and increases RN query speed.

However, this will impact an effective query on any other indexing field because that metadata was not saved when the packets were recorded.

Example 2: The RN supports community ID indexing, a hash of IP addresses, IP protocol, and transport ports that identify a flow of interest. Suppose the RN use case is to query based on community ID. In that case, indexing on IPv4 source and destination addresses, IPv6 source and destination addresses, IP protocol, and transport port source and destination addresses might be redundant.

6.4 Pre-buffer Configuration and Events

The Recorder Node (RN) pre-buffer is a circular buffer recording received packets. When enabled, the pre-buffer feature allows for the retention of the packets received by the RN for a specified length of time prior to an event that triggers the recording of buffered and future packets to disk. Without an event, the RN will record into this buffer, deleting the oldest packets when the buffer reaches capacity.

When an RN event is triggered, DMF saves packets in the pre-buffer to disk. 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. However, the received packets are also retained in the pre-buffer until the next event is triggered. By default, the pre-buffer feature is disabled, indicated by a value of zero minutes.

For example, when configuring the pre-buffer to thirty minutes, the buffer will receive up to thirty minutes of packets. When triggering an event, DMF records the packets currently in the buffer to disk, and packets newly received by the RN bypass the buffer and are written directly to disk until the termination of the event. When terminating the event, the pre-buffer resets, accumulating received packets for up to the defined thirty-minute pre-buffer size.

The packets affiliated with an event can be queried, replayed, or analyzed using any RN query. Each triggered event is identified by a unique, user-supplied name, used in the query to reference packets recorded in the pre-buffer before and during the event.

6.5 Using an Authentication Token

When using a DANZ Monitoring Fabric (DMF) Controller authentication token, the Recorder Node (RN) treats the DMF Controller as an ordinary client, requiring 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 RN 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.

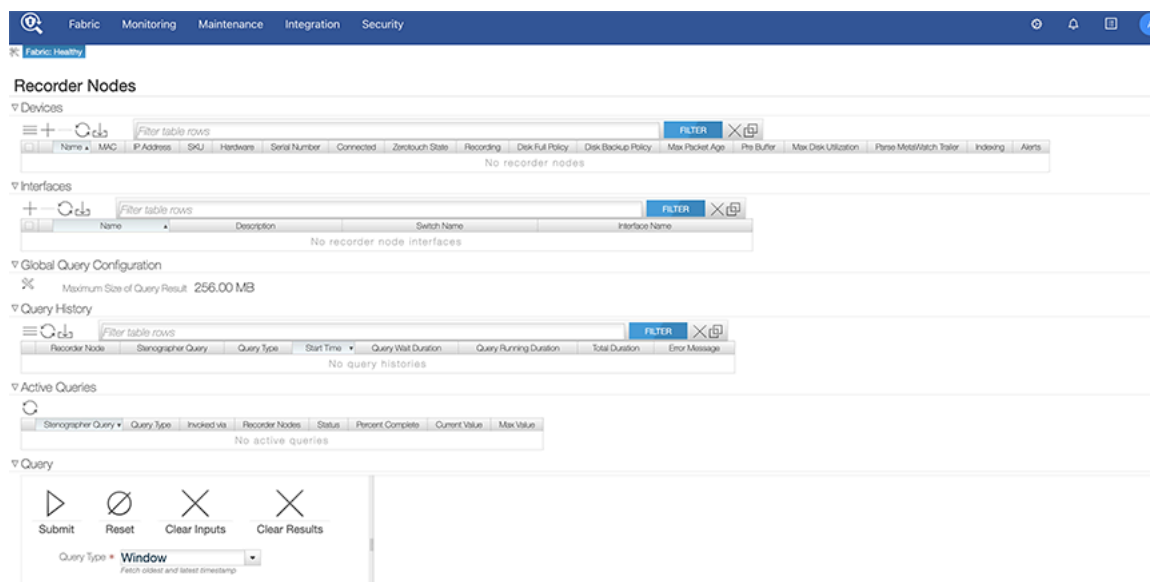
6.6 Using the GUI to Add a Recorder Device

To configure a Recorder Node (RN) or update the configuration of an existing RN, follow the steps below:

1. Select **Monitoring > Recorder Nodes** from the main menu bar of the DANZ Monitoring Fabric (DMF) GUI.

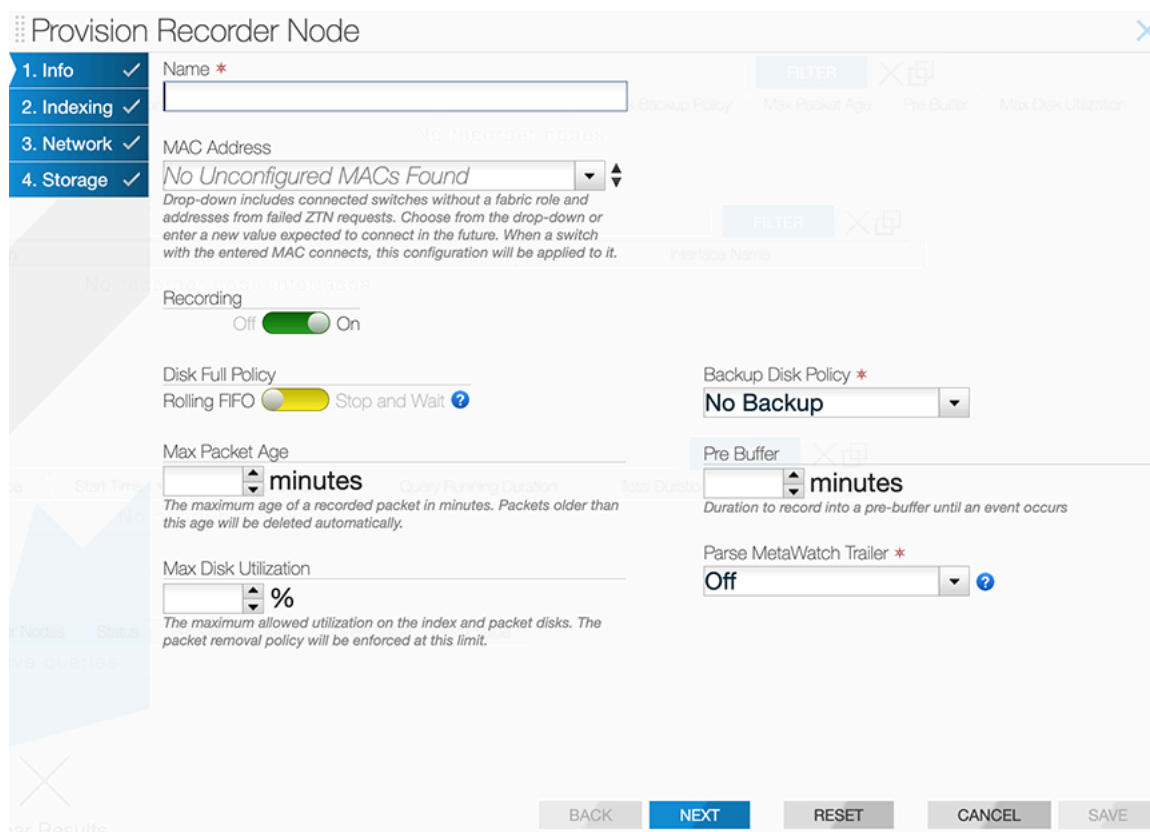
The system displays the page shown below.

Figure 6-1: Recorder Nodes




2. To add a new RN, click the provision control (+) in the **Recorder Nodes Devices** table.

Figure 6-2: Provision Recorder Node



3. Enter the following information in the required fields:
 - Assign a name to the RN.

- Set the MAC address of the RN. Obtain the MAC address from the chassis ID of the connected device, using the **Fabric > Connected Devices** option.
4. Configure the following options as needed:
- **Recording:** Recording is enabled by default. To disable recording on the RN, move the **Recording** toggle switch to **Off**. When recording is enabled, the RN 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 RN disks are full. The alternative removal policy is **Stop and Wait**, which causes the RN to stop recording when the disks are full and wait until disk space becomes available. Disk space can be made available by leveraging the RN delete operation to remove all or selected time ranges of recorded packets.
 - **Backup Disk Policy:** Specify the disk backup policy to as desired. This is a mandatory field while creating a new recorder. Select from one of the three following options:
 - **No Backup:** This is the default option and is also the recommended option when no extra disk is available. It is also a continuation of the behavior supported in previous releases.
 - **Remote Extend:** In this option, recording is performed on the local disks. When full, the recording continues on a remote Isilon cluster mounted over NFS. In this mode, the remote disks are called **backup disks**. With regard to the *Disk Full Policy*, if set to:
 - **Stop and Wait:** Recording stops when both local and remote disks become full.
 - **Rolling FIFO:** When the configured threshold is reached, the oldest files from both disks are removed until the disk usage returns below the threshold number.
 - **Local Fallback:** In this option, recording is performed on a remote Isilon cluster mounted over NFS. If the connection between the Recorder Node and the remote cluster fails, the recording is performed on the local disks until the failure is resolved. In this mode, the local disks are called **backup disks**. With regard to the *Disk Full Policy*, if set to:
 - **Stop and Wait:** Recording stops when the remote disks become full.
 - **Rolling FIFO:** When the configured threshold is reached, the oldest files from both disks are removed until the disk usage returns below the threshold number.
-  **Note:** A connection failure should not occur due to a misconfiguration of the NFS server on the DMF Controller. In such cases, the recording stops until the Controller's configuration is fixed.
- **Max Packet Age:** Change the **Max Packet Age** to set the maximum number of minutes that recorded packets will be kept on the RN. Packets recorded are discarded after the specified number of minutes. This defines the maximum age in minutes of any packet in the RN. 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 RN pre-buffer allows for windowed retention of packets received by the RN for a specified length of time. By default, the **Pre-Buffer** is set to zero minutes (disabled). With a nonzero **Pre-Buffer** setting and triggering 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 terminating 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 the parsing of the MetaWatch trailer.
 - **Off:** When set to **Off**, the RN will not parse the MetaWatch trailer, even if it is present in incoming packets.
 - **Auto:** When set to **Auto**, the RN 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 RN.

- **Force:** When set to **Force**, RN 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. Click **Save** to save and close the configuration page or click **NEXT** to continue with the configuration. Displays the **Indexing** tab of the **Provision Recorder Node** page.

Figure 6-3: Provision Recorder Node-Indexing

Provision Recorder Node

1. Info ✓
2. Indexing ✓
3. Network ✓
4. Storage ✓

Indexing
Defines indexing behavior when processing received frames

- ☒ 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

BACK NEXT RESET CANCEL SAVE

6. All the indexing options are enabled by default. To disable any of the indexing behaviors, move the toggle switch of the respective item to the left. For more details, see the [Indexing Configuration](#) section.

- Click **Save** to save and close the configuration page or click **NEXT** to continue with the configuration. Displays the **Network** tab of the **Provision Recorder Node** page.

Figure 6-4: Provision Recorder Node

The screenshot shows the 'Provision Recorder Node' configuration window. The left sidebar contains four tabs: '1. Info', '2. Indexing', '3. Network', and '4. Storage'. The '3. Network' tab is currently selected. The main area of the window displays the 'Auxiliary NIC Configuration' section, which includes a toggle switch set to 'No' (with 'Yes' as an alternative option). Below this, there are several other configuration sections, each with a 'Save' button and a close icon: 'Download Client', 'Recording', 'Disk Full Policy', 'Disk Backup Policy', 'Max Packet Age', 'Pre Buffer', 'Max Disk Utilization', and 'Parse Method'. At the bottom of the window, there are five buttons: 'BACK', 'NEXT', 'RESET', 'CANCEL', and 'SAVE'.

- By default, the Auxiliary NIC Configuration is set to **No**. This option is selected while configuring a node for local storage. For details, see the [Configuring a Node to Use Local Storage](#) section. Move the toggle switch to **Yes** when configuring a node for external storage. For details, see the [Configuring a Node to Use External Storage](#) section.

6.6.1 Configuring a Node to Use Local Storage

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 6-5: Network Provisioning

The screenshot shows the 'Provision Recorder Node' dialog with the 'Network' tab selected. The left sidebar shows a progress list: 1. Info (checked), 2. Indexing (checked), 3. Network (active), and 4. Storage (checked). The main content area is titled 'Configure the unused fat NIC of the Recorder Node for use on an L3 network'. It contains a section for 'Auxiliary NIC Configuration' with a toggle switch set to 'No'. Below this is a table with columns 'Description', 'Switch Name', and 'Interface Name'. The table is currently empty. At the bottom right, there is a 'Next' button. At the bottom of the dialog, there are 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 6-6: Configure to Use Local Storage


The screenshot shows the 'Provision Recorder Node' dialog with the 'Storage' tab selected. The left sidebar shows a progress list: 1. Info (checked), 2. Indexing (checked), 3. Network (checked), and 4. Storage (active). The main content area is titled 'Configure how and where we store packets/metadata'. It contains two sections: 'Index Disk' with 'Index Disk Configuration' set to 'No', and 'Packet Disk' with 'Packet Disk Configuration' set to 'No'. Both sections have toggle switches. Below these is a table with columns 'Description', 'Switch Name', and 'Interface Name'. The table is currently empty. At the bottom right, there is a 'Next' button. At the bottom of the dialog, there are 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

In order to store packets on external storage using an NFS mount, connect the Recorder Node's (RN) auxiliary interface to the same network and subnet where the NFS storage resides, as displayed in the figure below.

Figure 6-7: Topology to Use External Storage



Note: Create the volume for the index and packet on the NFS storage first. Refer to the vendor-specific NFS storage documentation about creating the volume (or path).

To configure an RN for external NFS storage, update the configuration of an existing recorder node or add a new node with the following steps:



Note: DMF *release 7.2* only supports Isilon NFS storage.

1. **Network:** For external NFS storage, such as Isilon, connect the auxiliary interface of the RN to a network and subnet that is reachable to Isilon NFS storage. Set the **Auxiliary NIC Configuration** toggle switch

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 RN management IP address.

Figure 6-8: Provision External Storage

Provision Recorder Node

1. Info ✓

2. Indexing ✓

3. Network ✓

4. Storage ✓

Auxiliary NIC Configuration

No ☐ Yes ☐

IP Address *

10.1 / 24

Interface Name

eth0

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 (toggle switch set to **No**). Set the toggle switch for both **Index Disk Configuration** and **Packet Disk Configuration** to **Yes**.
- **NFS Server [Index Disk Configuration and Packet Disk Configuration]:** assign the IP address or hostname for the NFS Server (e.g., Isilon Smart Connect hostname).
- **Transport Port of NFS Service [Index Disk Configuration and Packet Disk Configuration]:** DMF uses the default value if no value is specified (2049). Specify a value if the NFS storage has been configured to use something other than the default value.
- **Transport Port of Mounted Service [Index Disk Configuration and Packet Disk Configuration]:** if no value is specified, DMF uses the default value. Specify a value for this if the NFS storage-mounted service has been configured to use something other than the default value.

- **Volume:** [Index Disk Configuration and Packet Disk Configuration] - Specify the storage location or path on the NFS server for the **index** and **packets**.

Figure 6-9: Provision External Storage

3. Click **Save** to add the RN configuration to the Controller.



Note: When editing the configuration of a previously added RN to use external storage versus local storage or vice versa, **reboot** the RN.

6.7 Configuring a Recorder Node Interface

To record packets to a recorder node using a DANZ Monitoring Fabric (DMF) policy, configure a DMF Recorder Node (RN) interface that defines the switch and interface in the monitoring fabric where the RN is connected. The DMF RN interface is referenced by name in the DMF policy as the destination for traffic matched by the policy. To configure a DMF RN interface, perform the following steps:

1. Click the provision control (+) at the top of the **Recorder Node Interfaces** table. The system displays the following page:

Figure 6-10: 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

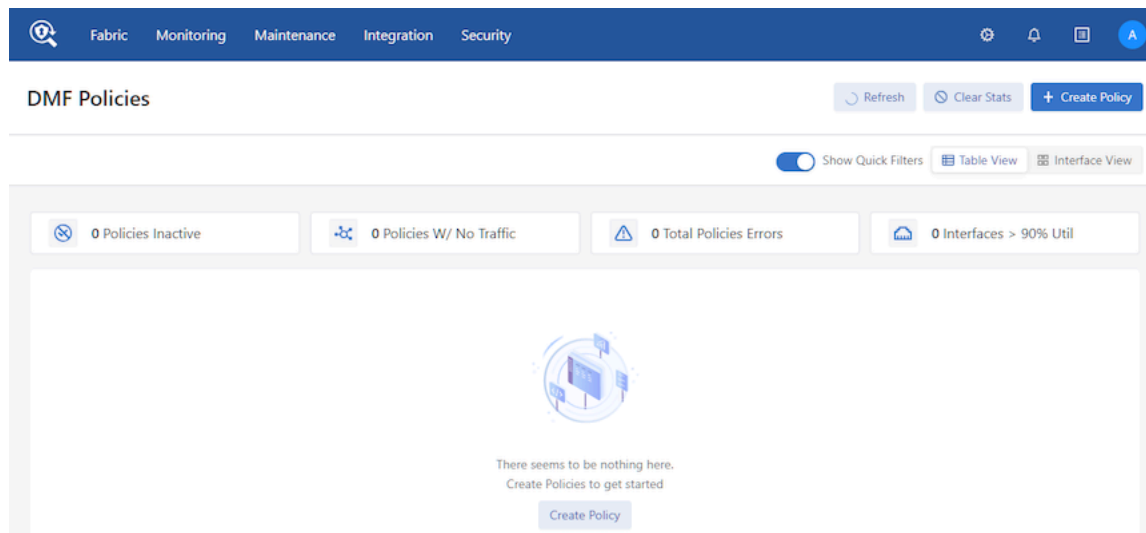
2. Assign a name for the DMF RN interface in the **Name** field.
3. Select the switch containing the interface that connects the RN to the monitoring fabric.
4. Select the interface that connects the RN to the monitoring fabric.
5. (Optional) Type information about the interface in the **Description** 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 (RN), include one or more RN interfaces as a delivery interface in a DANZ Monitoring Fabric (DMF) policy.

When creating a new policy or editing an existing policy, select the RN interfaces from the **Monitoring > Policies** dialog, as shown in the following screen.

Figure 6-11: DMF Policies



Note: To create an RN interface, proceed to the **Monitoring > Recorder Nodes** page and click the **+** in the Interface section.

To create a policy, select **Destination Tools > Add Ports(s)** and use the **RN Fabric Interface** to select a previously configured RN interface. Select or drag the **Interfaces** or **Recorder Nodes** to add **Destination Tools**.

Figure 6-12: Recorder Node - Create Policy

Create Policy

Policy Details ☐ Enable Advanced Options [Collapse](#)

Name	Description	Action	Push VLAN ?	Priority ?	Active
<input type="text" value="Input value"/>	<input type="text" value="Input value"/>	Forward	##	100	<input checked="" type="checkbox"/>

Port Selection Match Traffic Stat: Bit Rate

Traffic Sources [Create](#)
[+ Add Port\(s\)](#)

Services
[+ Add Service\(s\)](#)

Destination Tools [Create](#)
[+ Add Port\(s\)](#)

Match Traffic: Allow All Traffic (Default) has been selected

[Cancel](#) [Create Policy](#)

Figure 6-13: Add Recorder Nodes

Create Policy

Destination Tools
▼ Delivery Interfaces/Groups
^ RN Fabric Interface

You currently have no RN Fabric Interface configured in DMF.

Selected Interfaces [Cancel](#) [Add Interface](#)

Select or drag interface(s)/Recorder Node(s) to add Destination Tools



Note: The RN 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 (RN) records all the packets received on a filter interface that match the criteria defined in a DANZ Monitoring Fabric (DMF) policy. Recorded packets can be recalled from or analyzed on the RN using a variety of queries. Use the options in the RN **Query** section to create a query and submit it to the RN 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 metadata.
- **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 to further filter the retrieved results, as shown below:

Figure 6-14: 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 restrict the query to those interfaces. 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 restrict the query to those policies. 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, use **Global Query Configuration** to set the byte limit on packet query results.

Figure 6-15: Global Query Configuration



6.10 Viewing Query History

View Recorder Node (RN) submitted queries using the GUI or CLI.

To view the query history using the DANZ Monitoring Fabric (DMF) GUI, select **Monitoring > Recorder Nodes** and scroll down to the **Query History** section.

Figure 6-16: 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 RN and the query status.

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
3	HW-PR-2	after 10m ago	abort	2019-03-20 09:52:40.439000 PDT 711
4	HW-PR-1	after 10m ago	abort	2019-03-20 09:52:40.439000 PDT 711

--output truncated--

6.11 Using the CLI to Manage the DMF Recorder Node

6.11.1 Basic Configuration

To perform basic Recorder Node (RN) configuration, perform the following steps:

1. Assign a name to the RN device.

```
controller-1(config)# recorder-node device rn-alias
```

2. Set the MAC address of the RN.

```
controller-1(config-recorder-node)# mac 18:66:da:fb:6d:b4
```

Determine from the chassis ID of connected devices if the management MAC is unknown.

```
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 HLEYH2	18:66:da:fb:6d:b4	3c:fd:fe:1f:0f:02	esp130s0f1	10.4.100.200	LLDP

3. Define the RN interface name.

```
controller-1(config)# recorder-fabric interface Intf-alias
controller-1(config-pkt-rec-intf)#
```

Assign any alphanumeric identifier for the recorder node interface name, which changes the submode to **config-pkt-rec-intf**, to provide an optional description. This submode allows specifying the switch and interface where the RN is connected.

4. Provide an optional description and identify the switch interface connected to the RN.

```
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**: 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**, deleting the oldest packets to make room for newer packets when RN disks are full. This configuration can be changed to **stop-and-wait**, allowing the RN 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. **Backup Disk Policy**: Define the backup disk policy to select the secondary volume and select one of the following three options:

```
controller-1(config-recorder-node)# backup-volume
local-fallback Set local disk as backup when remote disk is unreachable
```



```
no-backup      Do not use any backup volume (default selection)
remote-extend  Set remote volume to extend local main disk
```

The **no-backup** mode is the default mode. The other two modes require that the Recorder Node have a set of recording disks and a connection to an Isilon cluster mounted via NFS. Configure this remote storage from the DMF Controller.

8. (Optional) **Max Packet Age**: This defines the maximum age in minutes of any packet in the RN. By default, **Max Packet Age** is unset, which means no limit is enforced. When setting a **Max Packet Age**, packets recorded on the RN are discarded after the minutes specified. To set the maximum number of minutes that recorded packets will be kept on the RN, 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: Combine **Max Packet Age** with the packet removal policy to control when packets are deleted based on age rather than disk utilization alone.

9. (Optional) **Max Disk Utilization**: This defines the maximum disk utilization as a percentage between 5% and 95%. The **Disk Full Policy** (**rolling-fifo** or **stop-and-wait**) is enforced when reaching this value. If unset, the default maximum disk utilization is **95%**; however, configure it, as required, using the following commands:

```
controller-1(config)# recorder-node device rn-alias
controller-1(config-recorder-node)# max-disk-utilization 80
```

10. (Optional) Disable unused or unneeded indexing configuration fields in subsequent recorder node queries. DMF enables all indexing fields by default. To disable a specific indexing option, enter the following commands from the **config-recorder-node-indexing** submodule. To re-enable a disabled option, enter the command without the **no** prefix.

Use the following command to enter the RN indexing submodule:

```
controller-1(config-recorder-node)# indexing
controller-1(config-recorder-node-indexing)#
```

Use the following commands to disable any unused fields 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-src
```

11. Identify the RN 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)#
```

12. Configure the DANZ Monitoring Fabric (DMF) policy to identify the traffic to send to the RN.

```
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 RN 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 (RN) as an alternative form of authentication in headless mode when the DANZ Monitoring Fabric (DMF) Controller is unreachable or by third-party applications that do not have or do not need DMF controller credentials to query the RN.

To configure the RN 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$12$cwt4PvsPySXrmMLYA.Mnyus9DpQ/bydGWD4LEhNL6xhPpkKNLzqWS <---hashed token shows in running
config
```

The DMF Controller uses its hidden authentication token to query the RN. 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 recorder node name. 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 (RN), enter the following command from enable mode:

```
controller-1# trigger recorder-node <name> event <event-name>
```

Replace **name** with the RN name 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 (RN) event, use the following command:

```
controller-1# terminate recorder-node <name> event <event-name>
```

Replace **name** with the RN name and replace **event-name** with the RN event name 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 DANZ Monitoring Fabric (DMF) Controller prompt is displayed immediately after entering a query or replay request, but the query continues in the background. Attempting to enter another replay or query command before the previous command is completed, an error message is displayed.

6.12.1 Packet Replay

Enter the **replay recorder-node** command from enable mode to replay the packets recorded by a Recorder Node (RN).

```
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 RN from which to replay the recorded packets.
- **interface:** The DMF delivery interface name receiving the packets.
- **stenographer-query:** The filter used to look up desired packets.
- (Optional) **real-time:** Replay the packets at the original rate recorded by the specified RN. The absence of this parameter will result in a replay up to the line rate of the RN interface.

- (Optional) **replay-rate bps**: Specify the number of bits per second used for replaying the packets recorded by the specified RN. The absence of this parameter will result in a replay up to the line rate of the RN 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

Use a packet query to search the packets recorded by a specific Recorder Node (RN). The operation uses a Stenographer query string to filter only the interesting traffic. The query returns a URL 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 RN.
- **packet-data filter stenographer-query**: Look up only the packets that match the specified Stenographer query.

The following example illustrates 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-1-2018-10-19-08-19-59-40b3d58a.pcap
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-2018-09-20-13-13-33-5f70c9ea.pcap 642KB 2018-09-20 13:13:33 PDT https://10.100.6.15/pcap/___packet_recorder___/coalesced-all-2018-09-20-13-13-33-5f70c9ea.pcap
2 coalesced-all-2018-09-20-16-29-16-0f790c5.pcap 106KB 2018-09-20 16:29:16 PDT https://10.100.6.15/pcap/___packet_recorder___/coalesced-all-2018-09-20-16-29-16-0f790c5.pcap
3 coalesced-all-2018-09-21-10-13-48-524682b1.pcap 2.34MB 2018-09-21 10:13:48 PDT https://10.100.6.15/pcap/___packet_recorder___/coalesced-all-2018-09-21-10-13-48-524682b1.pcap
4 coalesced-all-2018-09-21-10-13-08-23d8266c.pcap 3.30MB 2018-09-21 10:13:08 PDT https://10.100.6.15/pcap/___packet_recorder___/coalesced-all-2018-09-21-10-13-08-23d8266c.pcap
5 coalesced-all-2018-09-21-10-45-53-b94f75b.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-b94f75b.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-2018-10-03-08-17-34-34a259c6.pcap 11.4MB 2018-10-03 08:17:35 PDT https://10.100.6.15/pcap/___packet_recorder___/coalesced-all-2018-10-03-08-17-34-34a259c6.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 illustrates 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

Use a size query to analyze the number of packets and the total size recorded by a specific Recorder Node (RN). The operation uses a Stenographer query string to filter only the interesting traffic.

Enter the **query recorder-node** command from enable mode to run a size query.

```
# query recorder-node <name> size filter <stenographer_query>
```

The following is the meaning of each parameter:

- **name:** Identify the RN.
- **size filter stenographer-query:** Analyze only the packets that match the specified Stenographer query.

The following example illustrates 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

Use a window query to analyze the oldest and most recent available packets recorded by a specific Recorder Node (RN).

Enter the **query recorder-node** command from enable mode to run a window query.

```
switch# query recorder-node <name> window
```

The following is the meaning of each parameter:

- **name:** Identify the RN.

The following example illustrates 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

Use the abort **recorder-node** command to stop the query running on the specified Recorder Node (RN). From enable mode, enter the following command:

```
controller-1# abort recorder-node <name> filter <string>
```

Replace **name** with the RN name, and use the **filter** keyword to identify the specific filter used to submit the query. If the specific running query is unknown, use an empty-string filter of "" 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

Use Role-Based Access Control (RBAC) to manage access to the DANZ Monitoring Fabric (DMF) Recorder Node (RN) by associating the RN with an RBAC group.

To restrict access for a specific RN 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 associate the Recorder Node (RN).

Enter the following command from config mode on the active DANZ Monitoring Fabric (DMF) controller:

```
controller-1(config)# group test
controller-1(config-group)#
```

2. Associate one or more RNs with the group.

Enter the following CLI command from the **config-group** submode:

```
controller-1(config-group)# associate recorder-node <device-name>
```

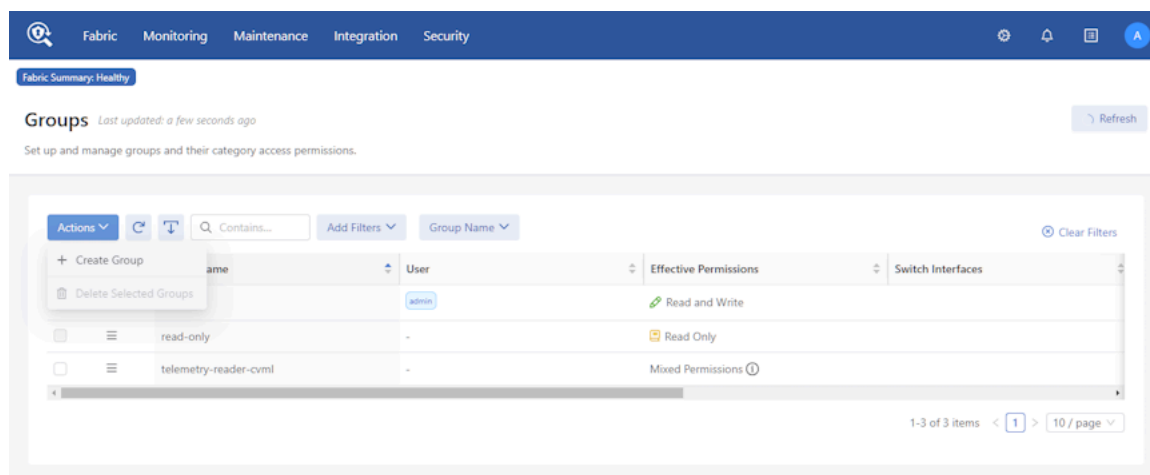
Replace **device-name** name with the RN name, as in the following example:

```
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 6-17: Create Security Group



2. Enter a **Group Name**.

Figure 6-18: Create Group

Create Group X

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

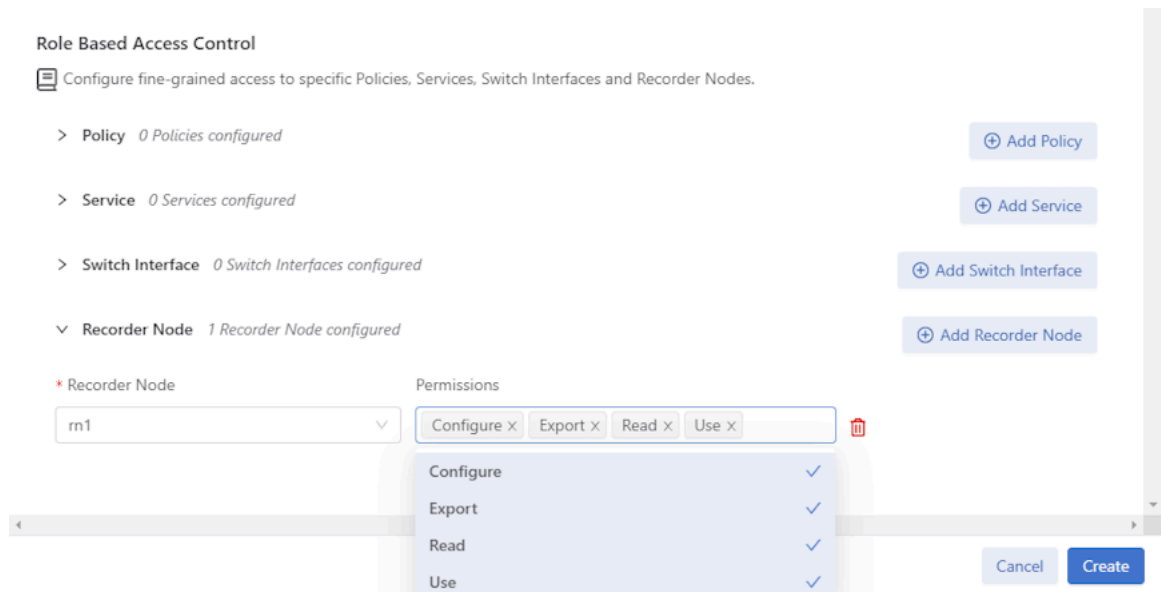
3. Under the **Role Based Access Control** section select **Add Recorder Node**.

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 6-19: Associate Recorder Node



5. Click **Create**.

6.14 Using the CLI to View Information About a Recorder Node

This section describes monitoring and troubleshooting the Recorder Node (RN) status and operation. The RN 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 RN 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 Bytes Rx 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 (RN) interface used as a delivery port in a DANZ Monitoring Fabric (DMF) out-of-band policy appears in a list. It lists RN interfaces 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 IF      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 IF      Switch      IF Name      State Dir Packets      Bytes      Pkt Rate Bit Rate Counter Reset Time
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 PR-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 a recorder node may display. In the CLI, display these errors and warnings 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 DANZ Monitoring Fabric (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

View Recorder Node (RN) statistics by clicking the RN alias from the **Monitoring > Recorder Nodes** page.

Figure 6-20: DANZ Monitoring Fabric (DMF) List of Connected Recorder Nodes

Name	MAC	IP Address	SKU	Hardware	Serial Number	Connected	Zerotouch State	Recording	Disk Full Policy
m1	...	10.2...	-	-	Not Specified	✓	Ok	Off	-

Click a **Recorder Node** to display the available recorder node statistics. All statistics are disabled or hidden by default.

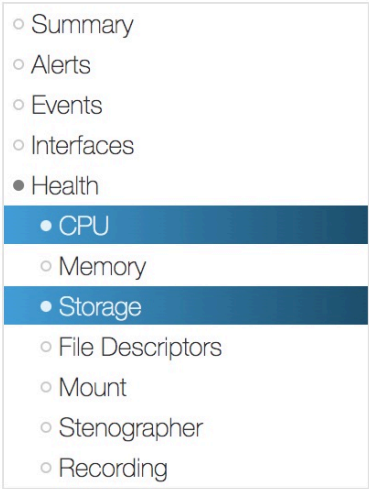
Figure 6-21: Available Recorder Node Statistics

The screenshot displays the Recorder Node Statistics page for node **rn1**. The page is divided into several sections:

- Navigation Bar:** Includes tabs for Fabric, Monitoring, Maintenance, Integration, and Security. A status indicator shows "Fabric: Healthy".
- Left Sidebar:** Contains a list of navigation items: Alerts, Events, Interfaces, Health (selected), CPU, Memory, Storage, File Descriptors, Mount, Stenographer, and Recording. Below this is a section for "Options & Shortcuts" with "Selected Content Action" (Replace/Append), "Selected Content Placement" (Top of content area, Bottom of content area, In order listed above), and "Select All", "Unselect All", "Restore Defaults".
- Main Content Area:**
 - Summary:**
 - Configuration:**
 - MAC: 02:15:58:00:00:00 f8
 - Recording: Off (toggle to On)
 - Disk Full Policy: Rolling FIFO (toggle to Stop and Wait)
 - Max Packet Age: —
 - Pre Buffer: 0 minutes
 - Max Disk Utilization: —
 - Parse MetaWatch Trailer: —
 - Indexing:**
 - MAC Source: X
 - MAC Destination: X
 - VLAN 1: X
 - VLAN 2: X
 - VLAN 3: X
 - IPv4 Source: X
 - IPv4 Destination: X
 - IPv6 Source: X
 - IPv6 Destination: X
 - IP Protocol: X
 - Port Source: X
 - Port Destination: X
 - MPLS: X
 - Community ID: X
 - MetaWatch Device ID: X
 - MetaWatch Port ID: X
 - Fabric Connection:** Connected ✓
 - Network:** Auxiliary NIC ✗
 - Storage: Index Disk:** Configured ✗
 - Storage: Packet Disk:** Configured ✗

Enable or view statistics by clicking on them. Selected statistics appear as highlighted in blue.

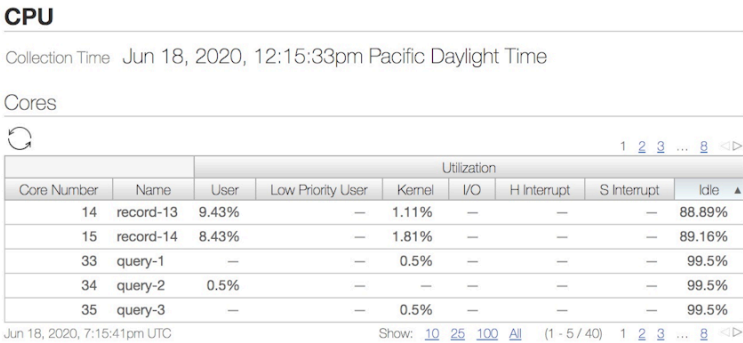
Figure 6-22: Selected Recorder Node Statistics



The RN shows health statistics for the following:

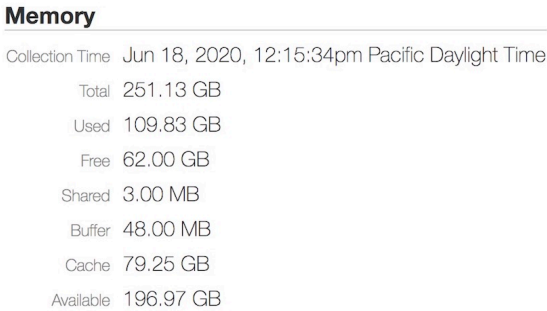
CPU: CPU health displays the compute resource utilization of the recorder node.

Figure 6-23: Recorder Node CPU Health Statistics




Memory: DMF displays memory-related stats such as total memory, used, free, available, etc.

Figure 6-24: Recorder Node Memory Statistics



Storage: Storage health displays the storage utilization percentage and total and available capacity of Index and Packet virtual disks.

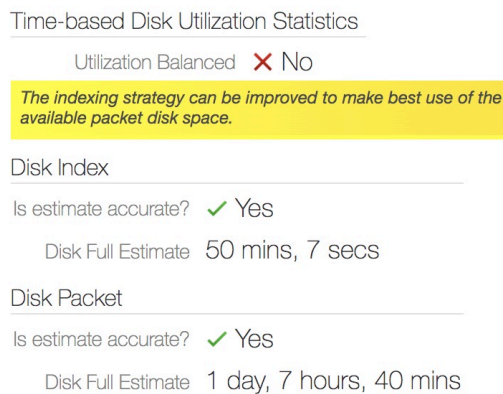
Figure 6-25: Recorder Node Storage Statistics



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 the RN for a certain duration.

Note: Inaccurate data may appear if the collected data points are insufficient to calculate the disk-full estimate. However, once sufficient data points are collected, the estimate recalculates and updates automatically.

Figure 6-26: Time-based Disk Utilization Statistics



Virtual Disks: Virtual Disk's health stats display the Index and Packet virtual disks' size, state, health, and RAID level configuration.

Figure 6-27: 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 6-28: 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

Slot Number	Type	Size	State	Temperature (°C)	Predictive Failure Count	SMART	Media Error Count	Other Error Count
12	SSD	3.49 TB	Online, Spun Up	44	0		0	0
13	SSD	3.49 TB	Online, Spun Up	45	0		0	0

Jun 18, 2020, 7:40:21pm UTC

Show: 10 25 100 All (1 - 2 / 2)

packet	142.55 TB	Optimal	Good	RAID 0	16
--------	-----------	---------	------	--------	----

Physical Disks

Slot Number	Type	Size	State	Temperature (°C)	Predictive Failure Count	SMART	Media Error Count	Other Error Count
0	HDD	8.91 TB	Online, Spun Up	26	0		0	0
1	HDD	8.91 TB	Online, Spun Up	26	0		0	0
2	HDD	8.91 TB	Online, Spun Up	26	0		0	0
3	HDD	8.91 TB	Online, Spun Up	26	0		0	0
4	HDD	8.91 TB	Online, Spun Up	26	0		0	0
5	HDD	8.91 TB	Online, Spun Up	26	0		0	0
6	HDD	8.91 TB	Online, Spun Up	26	0		0	0
7	HDD	8.91 TB	Online, Spun Up	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 6-29: 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: The Mount section displays the Index and Packet disk mount information, such as volume name, mount point, file system type, and mount health.

Figure 6-30: 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 6-31: 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 the application was successfully initialized. When the Stenographer application starts, a red x mark appears and disallows recording and querying during this time, which is expected behavior.
- **Tracked Files:** Tracked files are the total number of files stored under each CPU instance thread.
- **Cached Files:** Cached files are the number of open files with a file descriptor.
- **Max Cached Files:** The maximum cached files are the total openable files allowed.

These numbers are further divided and displayed for each recording thread and viewable in the **Recording Threads** table:

Figure 6-32: 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 6-33: 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 6-34: 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 6-35: 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 (RN) from the DANZ Monitoring Fabric (DMF) Controller, eliminating the need for box-by-box configuration. However, the option exists to override the default configuration for the RN from the **config-recorder-node** submode for any RN.



Note: These options are available only from the CLI, not 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 to use for the RN. This alias is affiliated with the MAC hardware address using the **mac** command.

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

- **banner**: Set the RN pre-login banner message
- **mac**: Configure the MAC address for the RN

Additionally, the option exists to override the configurations shown below to use values specific to the RN or used in a merge-mode along with the configuration inherited from the DMF controller:

- **ntp**: Configure RN to override default timezone and NTP parameters.
- **snmp-server**: Configure RN SNMP parameters and traps.
- **logging**: Enable RN logging to Controller.
- **tacacs**: Set TACACS defaults, server IP address(es), timeouts and keys.

Use the following commands from the **config-recorder-node** submode to change the default configuration on the RN:

- **ntp override-global**: Override global time configuration with RN time configuration.
- **snmp-server override-global**: Override global SNMP configuration with RN SNMP configuration.
- **snmp-server trap override-global**: Override global SNMP trap configuration with RN SNMP trap configuration.
- **logging override-global**: Override global logging configuration with packet recorder logging configuration.
- **tacacs override-global**: Override global TACACS configuration with RN TACACS configuration.

To configure the RN 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 RN time configuration.
- **snmp-server merge-global**: Merge global SNMP configuration with RN SNMP configuration.
- **snmp-server trap merge-global**: Merge global SNMP trap configuration with RN SNMP trap configuration.
- **logging merge-global**: Merge global logging configuration with RN logging configuration.

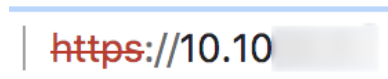
TACACS configuration does not have a merge option. It can either be inherited from the DMF Controller or overridden to use only the RN-specific configuration.

6.17 Large PCAP Queries

Access the RN via a web browser to run large PCAP queries to the Recorder Node (RN). This allows running packet queries directly to the RN without specifying the maximum byte or packet limit for the PCAP file (which is required when executing the query from the DANZ Monitoring Fabric (DMF) Controller).

To access the RN directly, use the URL **https://RecorderNodeIP** in a web browser, as shown below:

Figure 6-36: URL to Recorder Node



| **https://10.10** [blurred IP address]

The following page will be displayed:

Figure 6-37: 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 target RN IP address.
- **DMF Controller Username:** Provide the DMF Controller username.
- **DMF Controller Password:** Provide the password for authentication.
- **Stenographer Query Filter:** Use the query filter 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 for this PCAP query result.
- **Submit Request:** Sends a query to the specified RN and saves the PCAP file with the provided file name to the default download location for the browser.

6.18 Recorder Node Management Migration L3ZTN

After completing the first boot (initial configuration), remove the Recorder Node (RN) from the old Controller and point it to a new Controller via the CLI in the case of a Layer-3 topology mode.



Note: For appliances to connect to the DANZ Monitoring Fabric (DMF) Controller in Layer-3 Zero Touch Network (L3ZTN) mode, configure the DMF Controller deployment mode as **pre-configure**.

To migrate management to a new Controller, follow the steps below:

1. Remove the RN 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.

- SSH to the RN and configure the new Controller IP using the **zerotouch 13ztn controller-ip** command:

```
controller-1(config)# zerotouch 13ztn controller-ip 10.2.0.151
```

- After pointing the RN to use the new Controller, reboot the RN.
- Once the RN is back online, the DMF Controller receives the ZTN request.

```
controller-1(config)# show zerotouch request
# Request-history
-----|-----|-----|-----|-----|
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
```

- After the DMF Controller has received a ZTN request from the RN, add it 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
```

- Verify the addition of the RN to the new DMF Controller using the command below:

```
controller-1(config)# show zerotouch request
# Request-history
-----|-----|-----|-----|-----|
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 (RN):

Use the **show version** command to view the version and image information that RN is running on.

```
RecNode(config)# show version
Controller Version : DMF Recorder Node 8.1.0 (bigswitch/enable/dmf-8.1.x #5)
RecNode(config)#
```

Use the **show controllers** command to view the connected DANZ Monitoring Fabric (DMF) controllers to the recorder node.



Note: All cluster nodes appear in the command output if the RN is connected to a DMF Controller cluster.

```
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

Use the GUI to run multiple Recorder Node (RN) queries.

To run queries on recorded packets by the RN, navigate to the **Monitoring > Recorder Nodes** page.

Under the **Query** section, click on the **Query Type** drop-down to select the type of analysis to run on the recorded packets, as shown below:

Figure 6-38: Query Type

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 * Absolute Time

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 [dropdown]

MetaWatch Port ID [dropdown]





Deduplication No ☐ Yes ☒

Dedup. Time Window [dropdown] ms

After selecting the query type, 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.

Figure 6-39: Query - IP Address and Time

▽ Query



SubmitResetClear InputsClear Results

Query Type * ▼
Fetch TCP conversation analysis

☒ Relative Time☐ Absolute Time


Last * ▼ ▼

☒ Any IP☐ Directional IP

IP / ▼

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:

Figure 6-40: Query Progress



While a query is in progress, initiate another query from a new DANZ Monitoring Fabric (DMF) Controller web session. View the query progress under the **Active Queries** section:

Figure 6-41: Active Queries

▽ 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 (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 DANZ Monitoring Fabric (DMF) Controller CLI, packet deduplication is available for the packet data, packet object, size, and replay query types. Deduplication is **off** by default for these queries. Add the **dedup** option to the end of the query command after all optional values (if any) have been selected to enable deduplication.

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
```

Specify a time window (in milliseconds) 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, use the **dedup-window** option followed by an integer value for the time window after the **dedup** option.

```
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 DANZ Monitoring Fabric (DMF) Controller GUI, packet deduplication is available for the packet data, packet object, size, replay, application, and analysis query types. Deduplication is off 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 the time window value is unset.
3. Click **Submit** to continue.







Note: If a time window value is specified but deduplication is off, packet deduplication will not occur.

The following example illustrates enabling deduplication for a size query specifying a time window value.

Figure 6-42: 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 packet deduplication disabled. Hence, packet deduplication, by default, is off.

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.

6.22 Enabling Egress sFlow[®] on Recorder Node Interfaces

Enable egress sFlow[®]* to sample traffic sent to any DANZ Monitoring Fabric (DMF) Recorder Node (RN) attached to the fabric. Examining these sampled packets on a configured sFlow collector allows the identification of post-match-rule flows recorded by the RNs without performing a query against the RNs. While not explicitly required, Arista Networks highly recommends using the DMF Analytics Node (AN) as the configured sFlow collector, as it can automatically identify packets sampled utilizing this feature.

Platform Compatibility

All platforms apart from the following series:

- DCS-7280R
- DCS-7280R2
- DCS-7500R
- DCS-7020
- DCS-7050X4

Configuration

Use the DMF [CLI](#) or the [GUI](#) to turn the feature on or off.

6.22.1 Using the CLI to Enable Egress sFlow

The egress sFlow feature requires a configured sFlow collector. After configuring the sFlow collector, enter the following command from the config mode to enable the feature:

```
Controller-1(config)# recorder-node sflow
```

To disable the feature, enter the command:

```
Controller-1(config)# no recorder-node sflow
```

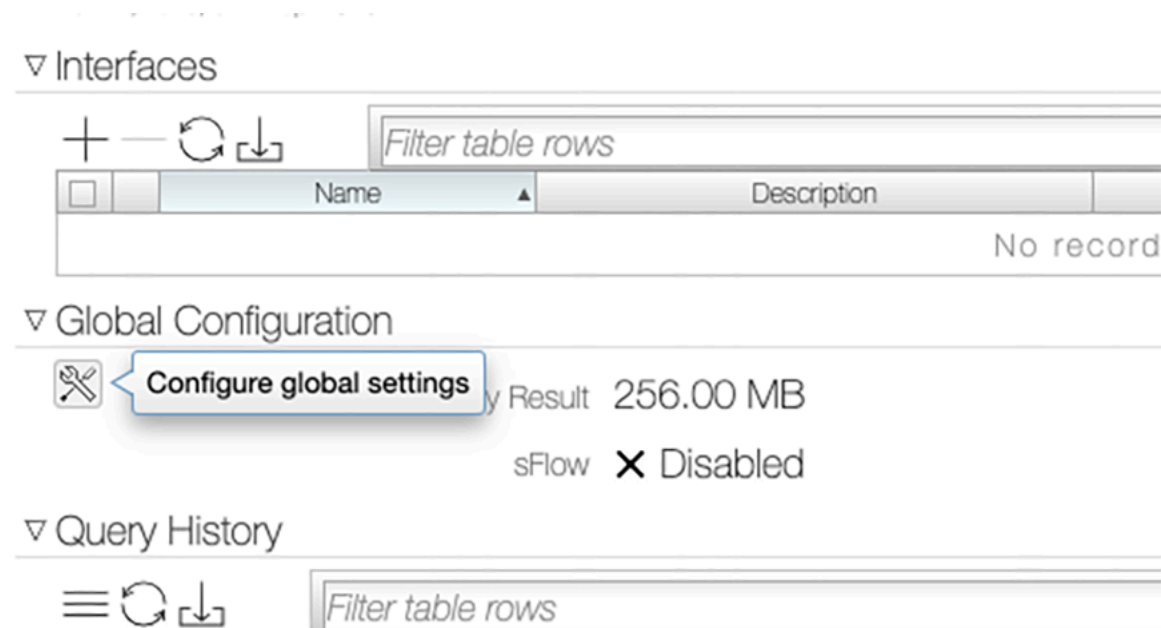
* sFlow[®] is a registered trademark of Inmon Corp.

6.22.2 Using the GUI to Enable Egress sFlow

Using the GUI

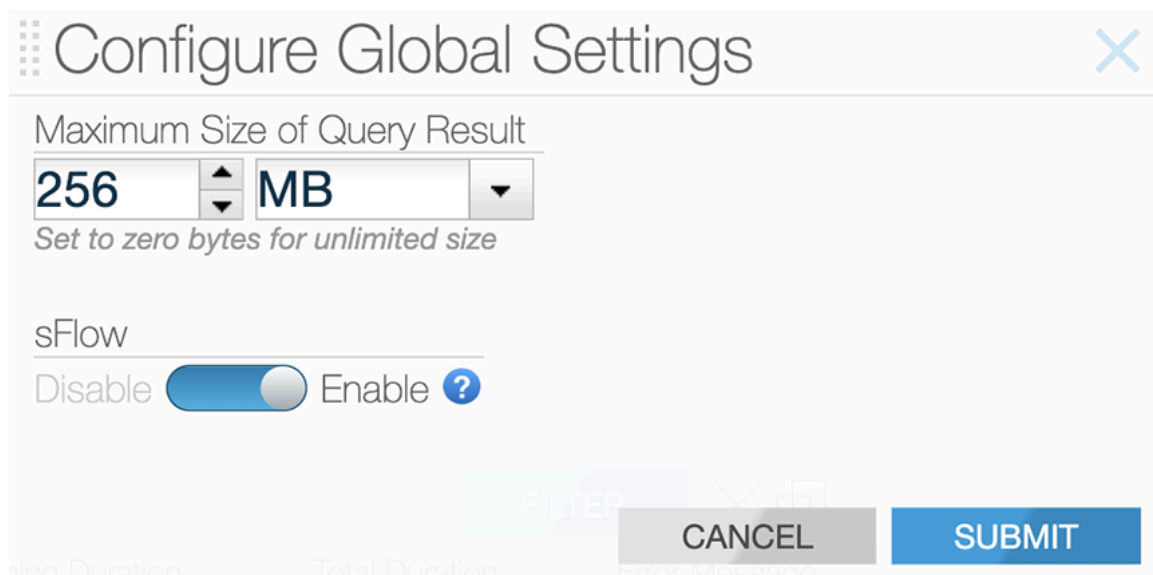
After configuring the fabric for sFlow and setting up the sFlow collector, navigate to the **Monitoring > Recorder Node** page. Under the **Global Configuration** section, click the **Configure global settings** button.

Figure 6-43: Configure Global Settings



In the Configure Global Settings pop-up window, enable the **sFlow** setting and click **Submit**.

Figure 6-44: Enable sFlow



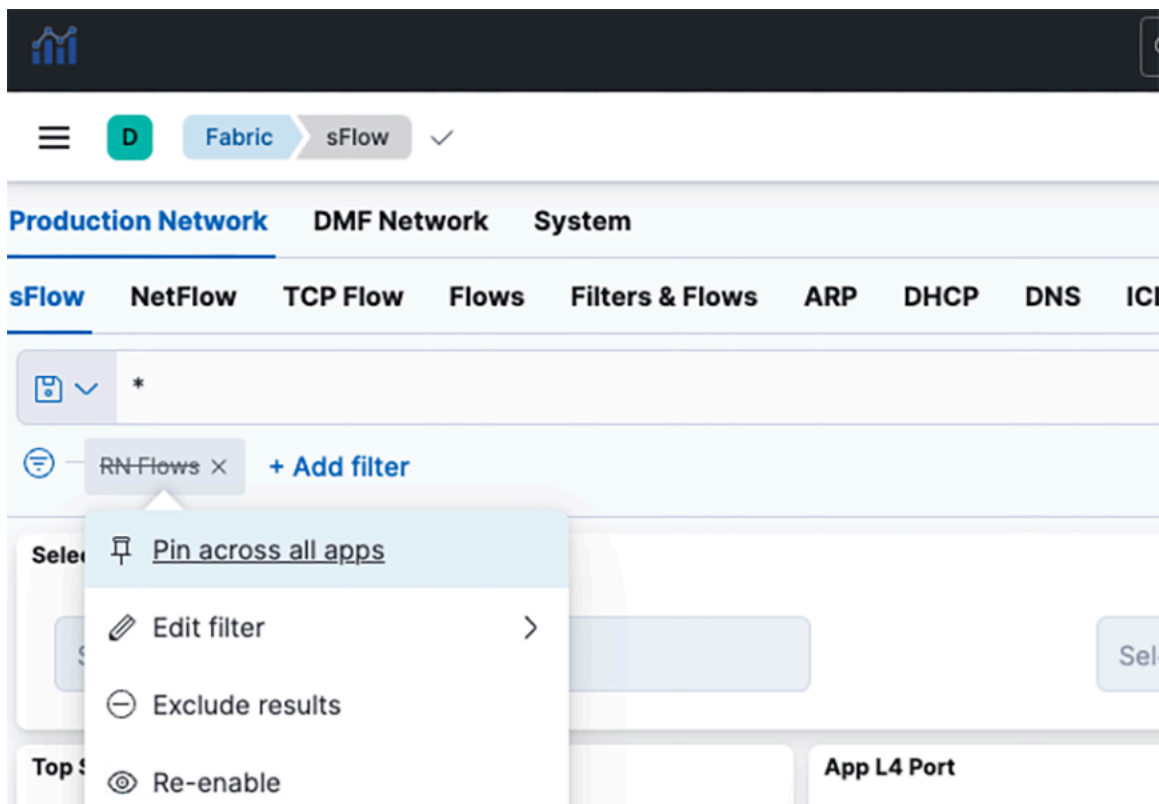
Analytics Node

When using a DMF Analytics Node as the sFlow collector, it has a dashboard to display the results from this feature. To access the results:

1. Navigate to the **sFlow** dashboard from the **Fabric** dashboard.

2. Select the disabled **RN Flows** filter.
3. Select the option to **Re-enable** the filter, as shown below.

Figure 6-45: Re-enable sFlow



6.22.3 Troubleshooting Egress sFlow Configurations

Switches not affiliated with a sFlow collector (either a global sFlow collector or a switch-specific sFlow collector) do not have an active feature even if the feature is enabled. Ensure the fabric is set up for sFlow and a configured sFlow collector exists. To verify that a configured global sFlow collector exists, use the command:

```
Controller-1# show sflow default
```

A configured collector appears as an entry in the table under the column labeled **collector**. Alternatively, to verify a configured collector exists for a given switch, use the command:

```
Controller-1# show switch switch-name table sflow-collector
```

This command displays a table with one entry per configured collector.

A **feature-unsupported-on-device** warning appears when connecting an unsupported switch to an RN. The feature does not sample packets passing to an RN from an unsupported switch. View any such warnings using the GUI or using the following CLI command:

```
Controller-1# show fabric warnings feature-unsupported-on-device
```

To verify the feature is active on a given switch, use the command:

```
Controller-1# show switch switch-name table sflow-sample
```

If the feature is enabled, the entry values associated with the ports connected to an RN would include an `EgressSamplingRate(number)` with a number greater than 0. The following example illustrates Port (1) on <switch-name> connecting to an RN.

```
Controller-1# show switch <switch-name> table sflow-sample
# Sflow-sample Device name      Entry key Entry value
-----|-----|-----|-----|
53  52          <switch-name> Port (1)  SamplingRate(0), EgressSamplingRate(10000), HeaderSize(128), Interval(10000)
```

6.22.4 Guidelines and Limitations for Enabling Egress sFlow

Consider the following guidelines and limitations while enabling Egress sFlow:

- The Egress sFlow support for the Recorder Nodes (RN) feature requires a configured sFlow collector in a fabric configured to allow sFlows.
- If a packet enters a switch through a filter interface with sFlow enabled and exits through a port connected to an RN while the feature is enabled, only one sFlow packet (i.e., the ingress sFlow packet) is sent to the collector.
- The Egress sFlow feature does not identify which RN has recorded a given packet in a fabric when there are multiple RNs. This is fine in a normal case as the queries are issued to the RNs in aggregate rather than to individual RNs, and hence, the information that any RN has received a packet is sufficient. In some cases, it may be possible to make that determination from the outport of the sFlow packet, but that information may not be available in all cases. This is an inherent limitation of egress sFlow.
- An **enabled** egress sFlow feature captures the packets sent to any RN with recording enabled, regardless of whether the RN is actively recording or not.

6.23 Recorder Node Recording State API

The Recorder Node (RN) recording statistics API on the DANZ Monitoring Fabric (DMF) Controller includes information about the operational state of the ongoing recording.

The information is available in `device/state/packet-recorder/recording/recording-state` where **recording-state** is the container holding the opstate information.



Note: This is a simplified version of the path; the path requires the **name** and **type** to be provided at the **device** to select a specific RN device to display state information.

Within the container, there are the following two values:

- **state:** an enumeration describing the recording opstate that has one of the following values: **initializing**, **ready**, **active**, and **stopped**.
- **description:** a string describing the state in more detail.

The following table lists the possible recording state values and describes how to interpret each scenario.

state	description	Interpretation
initializing	Recording application syncing files, {}% complete.	The application is synchronizing previously recorded packets, so it cannot record yet. The formatted value is the percentage of synchronization completed so far.
initializing	Recording application is starting up.	The application has finished synchronization and is now completing the setup for recording.
ready	Ready to record, no recordable traffic received.	No traffic has been received after the application completed initializing.
active	Recording traffic.	Traffic is being received and recorded to disk.
stopped	Recording threads not running.	Recording threads were interrupted or terminated, so recording has stopped.
stopped	Recording not enabled.	Recording is not enabled, and packets are not recorded.
stopped	Disk space exhausted with stop-and-wait mode configured. Awaiting user input.	Recording has stopped as disk space has been exhausted, and the removal policy has been configured as stop-and-wait. Delete packets to resume recording.
stopped	Recording state not found.	The recording application isn't running, so there is no recording state information.

When the state is not active, an RN warning is generated on the DMF Controller to indicate that a nonactive RN is connected to the fabric.



Note: Any other state than `active` is considered to be inactive. This is not necessarily indicative of a problem. However, you should take action accordingly using the provided interpretations.

Additionally, the type of logical disk volume used by each recording thread is part of the RN recording statistics API on the DMF Controller.

This information is available in `device/state/packet-recorder/recording/recording-thread/disk` where a `disk` is an enumeration having one of two values, `primary` or `backup`, if the thread is recording to the primary or backup configured logical disk volume, respectively. This is a simplified version of the path. The disk type does not imply whether the disk volume is local or remote, as the storage configuration determines this.

Show Commands

The `show recorder-node device <rn name> recording-health` command includes the state from the recording state and the disk type for each recording thread. Use the **details** option of this command to view the description of the recording state in addition to the previously stated information.

The following examples highlight the additional information in **bold** for illustrative purposes.

```
C1> show recorder-node device <rn name> recording-health
~~~~~ Recording Application ~~~~~
Recording Collection Time      : 2024-04-23 15:12:30 UTC
Recording CRC Errors          : 0
Recording Frame Length Errors : 0
Recording Back Pressure Errors : 0
Recording Undersized Frames   : 0
Recording Frames 64B          : 0
Recording Frames 65-127B      : 0
Recording Frames 128-255B     : 0
Recording Frames 256-511B     : 0
Recording Frames 512-1023B    : 0
Recording Frames 1024-1522B   : 0
Recording Frames 1523-9522B   : 0
Recording Oversized Frames    : 0
Recording state                : Active

~~~~~ Recording Instance ~~~~~
Core Disk   Dropped Packets Total Packets Collection Start Time
----|-----|-----|-----|-----|
1    primary 0                204                2024-04-23 15:00:34 UTC
```

```
C1> show recorder-node device <rn name> recording-health details
~~~~~ Recording Application ~~~~~
Recording Collection Time      : 2024-04-23 15:12:30 UTC
Recording CRC Errors          : 0
Recording Frame Length Errors : 0
Recording Back Pressure Errors : 0
Recording Undersized Frames   : 0
Recording Frames 64B          : 0
Recording Frames 65-127B      : 0
Recording Frames 128-255B     : 0
Recording Frames 256-511B     : 0
Recording Frames 512-1023B    : 0
Recording Frames 1024-1522B   : 0
Recording Frames 1523-9522B   : 0
Recording Oversized Frames    : 0
Recording state                : Active
Recording state Details        : Recording traffic.

~~~~~ Recording Instance ~~~~~
Core Disk   Dropped Packets Total Packets Collection Start Time
----|-----|-----|-----|-----|
1    primary 0                204                2024-04-23 15:00:34 UTC
```

```
C1> show recorder-node device <rn name> recording-health
~~~~~ Recording Application ~~~~~
Recording Collection Time      : 2024-04-23 15:12:30 UTC
Recording CRC Errors          : 0
Recording Frame Length Errors : 0
Recording Back Pressure Errors : 0
Recording Undersized Frames   : 0
Recording Frames 64B          : 0
Recording Frames 65-127B      : 0
Recording Frames 128-255B     : 0
Recording Frames 256-511B     : 0
Recording Frames 512-1023B    : 0
Recording Frames 1024-1522B   : 0
Recording Frames 1523-9522B   : 0
Recording Oversized Frames    : 0
```

```

Recording state           : Active

~~~~~ Recording Instance ~~~~~
Core Disk   Dropped Packets Total Packets Collection Start Time
----|-----|-----|-----|-----|
1    backup    0                204          2024-04-23 15:00:34 UTC

```

```

C1> show recorder-node device <rn name> recording-health details
~~~~~ Recording Application ~~~~~
Recording Collection Time      : 2024-04-23 15:12:30 UTC
Recording CRC Errors          : 0
Recording Frame Length Errors : 0
Recording Back Pressure Errors : 0
Recording Undersized Frames   : 0
Recording Frames 64B          : 0
Recording Frames 65-127B      : 0
Recording Frames 128-255B     : 0
Recording Frames 256-511B     : 0
Recording Frames 512-1023B    : 0
Recording Frames 1024-1522B   : 0
Recording Frames 1523-9522B   : 0
Recording Oversized Frames    : 0
Recording state               : Active
Recording state Details       : Recording traffic.

~~~~~ Recording Instance ~~~~~
Core Disk   Dropped Packets Total Packets Collection Start Time
----|-----|-----|-----|-----|
1    backup    0                204          2024-04-23 15:00:34 UTC

```

The **show recorder-node warnings** command includes nonactive RN warnings if any exist.

```

C1> show recorder-node warnings

~ RAID Health Warning(s) ~
None.

~ Low Disk Space Warning(s) ~
None.

~ Full Disk Warning(s) ~
None.

~~~~~ Non Active Recorder Node Warning(s) ~~~~~
# Recorder Node Message
-|-----|-----|-----|-----|-----|
-----|
1 rn1          Inactive Recording State: Ready - Ready to record, no
recordable traffic received.

```

Limitations

The **ready** state only occurs after the recording application has finished initializing if no recordable traffic has been received yet. The recording application must undergo its initialization process whenever the RN is rebooted, restarted, or after restarting the RN application from the DMF Controller. If the RN is in the **active** state and stops receiving packets, it will not regress into the **ready** state; it will remain in the **active** state.

Link Aggregation

This chapter describes configuring link aggregation groups between switches, switches, and tools or between switches and taps.

7.1 Configuring Link Aggregation

Link aggregation combines multiple LAN links and 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 LAG members. The enhanced hashing algorithm automatically assigns the best hashing type for the switch and traffic. This setting allows the manual selection of the packet types and fields used for load distribution among the members of a port-channel interface. Enhanced mode and symmetric hashing are enabled by default for the supported switch platforms. 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 apply, DMF uses Layer-2 header fields (source MAC address, destination MAC address, VLAN-ID, and ethertype) 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: DMF treats VN-tagged and QinQ packets as L2 packets and uses Layer-2 headers to distribute traffic among LAG member interfaces for these packets.

7.1.1 Using the CLI to Configure Link Aggregation Groups

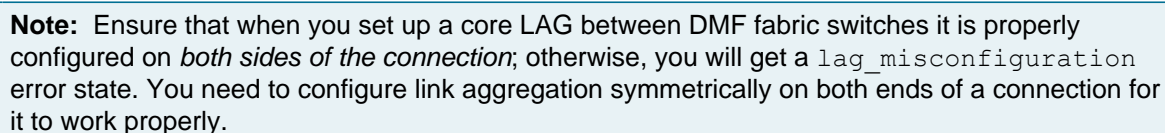
1. Use the `lag-interface` command to enter the config-switch-lag-if submode to define the LAG member interfaces and specify the type of load distribution (hashing) to use for the LAG.

- For example, the following commands add two interfaces to a LAG named *my-lag*.

- a. Assign a name to the LAG and enter the **config-switch-lag-if** submode.

b. Assign members to the LAG.

```
controller-1(config)# switch SWITCH-2
controller-1(config-switch)# lag-interface core-link
controller-1(config-switch-lag-if)# member ethernet49
controller-1(config-switch-lag-if)# member ethernet51
```



- ```

controller-1> show lag
~~~~~
# Switch  LAG      State Rx Rate Pkt Rate Peak Rate Peak Pkt Rate TX Rate Pkt Rate Peak Rate Peak Pkt Rate
- - - - -
1 SWITCH-1 core-link up    720bps  0      2.59Gbps 608181      184bps  0      53.2Mbps 37150
2 SWITCH-2 core-link up    184bps  0      424bps   0      720bps  0      1.08Kbps  1
~~~~~

Switch LAG Member State Fault Rx Rate Pkt Rate Peak Rate Peak Pkt Rate Tx Rate Pkt Rate Peak Rate Peak Pkt Rate
- - - - -
1 SWITCH-1 core-link ethernet1 up 624bps 0 800bps 1 88bps 0 152bps 0
2 SWITCH-1 core-link ethernet3 up 96bps 0 18.5Mbps 1735 88bps 0 136bps 0
3 SWITCH-2 core-link ethernet49 up 88bps 0 152bps 0 624bps 0 800bps 1
4 SWITCH-2 core-link ethernet51 up 88bps 0 280bps 0 96bps 0 18.5Mbps 1735
~~~~~

```

Navigate to **Fabric > LAGs**.

[Fabric](#)
[Monitoring](#)
[Maintenance](#)
[Integration](#)
[Security](#)

Fabric: Unhealthy
1
⚙️
📄
A

## Link Aggregation Groups (LAGs)

View and configure LAGs

LAG Information

☐ LAG Name
 ☐ Switch Name
 ☐ Status

|                      |    |                |
|----------------------|----|----------------|
| <a href="#">LAG1</a> | F1 | 2/2 Members Up |
|----------------------|----|----------------|

Edit

Delete

LAG Alerts

View All Alerts

Switch

C1

1

>

Switch LAG Enhanced Hash Settings

Active Setting

Inactive Setting

|    | L2 | L2 GRE Inner L2 | L2 GRE Inner L3 | IPv4 | IPv6 | VxLAN Inner L2 | VxLAN Inner L3 | MPLS | GTP        | Symmetric | Seed 1 | Seed 2 |
|----|----|-----------------|-----------------|------|------|----------------|----------------|------|------------|-----------|--------|--------|
| C1 |    |                 | 1               | 1    | 1    |                | 1              |      | Configured |           |        |        |
| C2 | 1  |                 |                 |      |      |                |                |      | Configured |           |        |        |
| D1 | 4  |                 | 2               | 2    | 2    |                |                |      |            |           |        |        |
| D2 | 1  |                 |                 |      |      |                |                |      |            |           |        |        |
| F1 | 4  |                 |                 |      |      |                |                |      | Configured |           |        |        |
| F2 | 2  |                 |                 |      |      |                |                |      | Configured |           |        |        |

LAG Members Utilization

| LAG Name             | Switch Name | RX Members Utilization Range | RX Avg. Utilization | TX Members Utilization Range | TX Avg. Ut |
|----------------------|-------------|------------------------------|---------------------|------------------------------|------------|
| <a href="#">LAG1</a> | F1          | <div></div>                  | 0%                  | <div></div>                  | 0%         |

## Main Page

The main page contains four components:

- **LAG Information**
- **LAG Alerts**
- **Switch LAG Enhanced Hash Settings**
- **LAG Members Utilization**

## LAG Information

LAG Information

☐ LAG Name
 ☐ Switch Name
 ☐ Status

|                      |    |                |
|----------------------|----|----------------|
| <a href="#">LAG1</a> | F1 |                |
| <a href="#">LAG2</a> | C2 | 1/2 Members Up |

Edit

Delete

LAG2 Member Status

ethernet1

Up

ethernet5

Down

The LAG Information table displays all the LAGs configured in DMF and each member's status. A row contains the **LAG Name**, **Switch Name**, and **Member Status**.

Hover over a member under **Status** to view the status of each member interface.

Use the checkbox to **Edit** or **Delete** a LAG entry. Select the **LAG Name** link to open the properties tab of the LAG details panel.

## Create or Edit LAG

Create LAG

×

---

\* LAG Name

LAG3

\* Switch Name

● D2

▼

Interfaces

● ethernet2 ×

● ethernet3 ×

Minimum Links ?

2

Cancel

Submit

Select **Create LAG** to open the configuration panel.

Enter a **LAG Name** and the **Switch Name** from the drop-down list, and choose the member **Interfaces** from that switch.

As an option, set the **Minimum Links** configuration required for the LAG to stay up. Refer to [Minimum Link and Activation Configuration for Lag Interfaces](#) for more information.

Select **Submit** to create the LAG.



**Tip:** If no configured LAGs exist, select **Create LAG** on the main page.

To edit a LAG, choose it in the Information table and select **Edit**.

While the LAG Name and switch name are not editable, modifying member **Interfaces** and **Minimum Links** values is.

Only one LAG can be edited at a time; however, the system supports deleting multiple LAGs simultaneously.

Select **Submit** to commit the changes to the LAG.

LAG Alerts

LAG Alerts

[View All Alerts](#)

SwitchC1

1

Unsupported LAG Config

LAGLAG2

1

Down LAG Member Link (1)

**LAG Alerts** displays all LAG Alerts and LAG-related Switch Alerts. The **Switch** or **LAG** tag indicates the type of alert. **Expand** each item to view the alert description and the number of occurrences.

LAG Enhanced Hash Grid

Switch LAG Enhanced Hash Settings

Active SettingInactive Setting

|    | L2 | L2 GRE<br>Inner L2 | L2 GRE<br>Inner L3 | IPv4 | IPv6 | VxLAN<br>Inner L2 | VxLAN<br>Inner L3 | MPLS | GTP        | Symmetric | Seed 1 | Seed 2 |
|----|----|--------------------|--------------------|------|------|-------------------|-------------------|------|------------|-----------|--------|--------|
| C1 |    |                    | 1                  | 1    | 1    |                   | 1                 |      | Configured |           |        |        |
| C2 | 1  |                    |                    |      |      |                   |                   |      | Configured |           |        |        |
| D1 | 4  |                    | 2                  | 2    | 2    |                   |                   |      |            |           |        |        |
| D2 | 1  |                    |                    |      |      |                   |                   |      |            |           |        |        |
| F1 | 4  |                    |                    |      |      |                   |                   |      | Configured |           |        |        |
| F2 | 2  |                    |                    |      |      |                   |                   |      | Configured |           |        |        |

The **Switch Lag Enhanced Hash Settings** grid displays a switch's current LAG-enhanced hash settings across various hash fields and categories. If a cell is gray, there are no currently active sub-settings. Hover over the blue cell to view active fields. Select a **switch name** link to open the Enhanced Hash Configuration tab of the LAG Details pane.

LAG Members Utilization

LAG Members Utilization

| LAG Name | Switch Name | RX Members Utilization Range | RX Avg. Utilization | TX Avg. Utilization |
|----------|-------------|------------------------------|---------------------|---------------------|
| LAG1     | F1          | <div></div>                  | 0%                  | <div></div>         |
| LAG2     | C2          | <div></div>                  | 0%                  | <div></div>         |

LAG1 Utilization Range (TX)

Speed10.0 Gbps

Min. Bit Rate72.0 bps

Max. Bit Rate72.0 bps

The **LAG Members Utilization** table summarizes the utilization for each LAG.

In the Utilization Range column, a colored bar indicates the start of the minimum utilization and the end of the maximum utilization of its members. Hover over the bar to view the Speed, minimum bit Rate, and maximum bit Rate of the LAG for each direction. The bar changes color based on LAG alerts:

- Green – indicates no alerts.
- Yellow – only warnings.
- Red – there is more than one alert.

Select the **LAG Name** link to open the Utilization tab of the LAG Details pane.

## Details Pane

Select a **LAG** or **switch link** to open the details pane. Choose a **LAG** to view its details. When choosing a switch link in the Enhanced Hash grid, the first LAG on that switch is selected. If there isn't a LAG on the switch, the system displays an empty string, and the Enhanced Hash opens for that switch. The Properties and Utilization tab will be disabled. Select a different LAG to access the other tabs.

## Properties

LAG1 (F1) ▾

Properties

Utilization

LAG Enhanced Hash

| LAG Name | Switch Name | Members                 | LAG Status | Minimum Links |
|----------|-------------|-------------------------|------------|---------------|
| LAG1     | F1          | • ethernet1 • ethernet2 | • Up       |               |

LAG Members Configuration [Show/Hide Columns](#)

| Interface Name | Admin Status | Forward Error Correction | MAC Loopback Mode | Auto-Negotiation     | Enable Optics |
|----------------|--------------|--------------------------|-------------------|----------------------|---------------|
| • ethernet1    | • Up         | • Use Switch Default     | • No              | • Use Switch Default |               |
| • ethernet2    | • Up         | • Use Switch Default     | • No              | • Use Switch Default |               |

The LAG Name, Switch Name, Members tags, and Minimum Link values appear in the descriptions section. The Members Configuration table shows the configuration of each member interface. Use **Show/Hide Columns** to determine which columns to display.

## Utilization

LAG1 (F1) ▾

Properties

Utilization

LAG Enhanced Hash

LAG Stats [Clear/Reset LAG Stats](#) Bytes Packets

| RX Stats    |           |               |                      | TX Stats    |           |               |                      |
|-------------|-----------|---------------|----------------------|-------------|-----------|---------------|----------------------|
| Bytes Count | Bit Rate  | Peak Bit Rate | Peak Byte Time       | Bytes Count | Bit Rate  | Peak Bit Rate | Peak Byte Time       |
| 52.3 kB     | 152.0 bps | 152.0 bps     | May 09, 2024 12:43PM | 52.4 kB     | 152.0 bps | 192.0 bps     | May 09, 2024 12:50PM |

LAG Members RX Stats [All ▾](#) LAG Members TX Stats [All ▾](#)

80.0 bps

60.0 bps

40.0 bps

20.0 bps

0.0 bps

13:27:35 13:27:40

• ethernet1 • ethernet2

| Interface   | Byte Count | Bit Rate | Peak Bit Rate |
|-------------|------------|----------|---------------|
| • ethernet1 | 26.1 kB    | 72.0 bps | 112.0 bps     |
| • ethernet2 | 26.1 kB    | 72.0 bps | 112.0 bps     |

**Utilization** displays the utilization statistics for the LAG.

There are two time-series charts for RX and TX stats. These charts are updated every 10 seconds and display each LAG member's bit-rate / packet rate. For LAGs with over five members, select **Top 5** or **Bottom 5** from the **All** drop-down in the chart to display only the top or bottom five interfaces.

Alternate between **Bytes** and **Packets** to change the statistics displayed. Alternating between bytes and packets will not reset the time-series chart since both bytes and packet data are polled every 10 seconds.

Select **Clear/Reset LAG Stats** to re-initialize the chart. Selecting a new LAG re-initializes the charts, and polling begins for new LAG utilization data.

## Enhanced Hash Configuration

LAG1 (F1) ×

Properties **Configuration - F1** Save Changes Cancel

Utilization Symmetric Disabled Seed 1  Seed 2

LAG Enhanced Hash

| L2                                          | L2 GRE Inner L2                             | VxLAN Inner L2                   | IPv4                                 | IPv6                                 | L2 GRE Inner L3                      | VxLAN Inner L3                       | MPLS                                    |
|---------------------------------------------|---------------------------------------------|----------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|-----------------------------------------|
| <input checked="" type="radio"/> Source MAC | <input checked="" type="radio"/> Source MAC | <input type="radio"/> Source MAC | <input type="radio"/> Source IP      | <input type="radio"/> Source IP      | <input type="radio"/> Source IP      | <input type="radio"/> Source IP      | <input type="radio"/> Label 1           |
| <input checked="" type="radio"/> Dest. MAC  | <input type="radio"/> Dest. MAC             | <input type="radio"/> Dest. MAC  | <input type="radio"/> Dst. IP        | <input type="radio"/> Dst. IP        | <input type="radio"/> Dst. IP        | <input type="radio"/> Dst. IP        | <input type="radio"/> Label 2           |
| <input type="radio"/> Eth Type              | <input type="radio"/> Eth Type              | <input type="radio"/> Eth Type   | <input type="radio"/> L4 Source Port | <input type="radio"/> L4 Source Port | <input type="radio"/> L4 Source Port | <input type="radio"/> L4 Source Port | <input type="radio"/> Label 3           |
| <input checked="" type="radio"/> VLAN ID    | <input type="radio"/> VLAN ID               | <input type="radio"/> VLAN ID    | <input type="radio"/> L4 Dest. Port  | <input type="radio"/> L4 Dest. Port  | <input type="radio"/> L4 Dest. Port  | <input type="radio"/> L4 Dest. Port  | <input type="radio"/> Label High Bits   |
|                                             |                                             |                                  | <input type="radio"/> Protocol       | <input type="radio"/> Protocol       | <input type="radio"/> Protocol       | <input type="radio"/> Protocol       | <input type="radio"/> Payload Dest. IP  |
|                                             |                                             |                                  | <input type="radio"/> VLAN ID        | <input type="radio"/> VLAN ID        | <input type="radio"/> VLAN ID        | <input type="radio"/> VLAN ID        | <input type="radio"/> Payload Source IP |
| <input type="radio"/> L2 GRE Key            |                                             |                                  |                                      |                                      |                                      |                                      |                                         |

GTP Match Fields

| First Byte | First Byte Mask | Sequence | Port Match Entries | Port Combo | Src | Dst | Src Port | Dst Port |
|------------|-----------------|----------|--------------------|------------|-----|-----|----------|----------|
| 0x         |                 | 0        |                    |            | Src |     | 12       |          |
|            |                 | 1        |                    |            | Dst |     | 33       |          |
|            |                 | 2        |                    |            | Src |     | 44       |          |
|            |                 | 3        |                    |            | Src |     | 332      |          |

⊕ Add Another

Before the DMF 8.6 release, the LAG Enhanced Hash configuration occurred on the Switches page. Now, it is done in the **LAGs Detail** pane.

Select **Unlock to Edit** to enable changing configuration settings and **Save Changes** to submit the changes.

The LAGs detail pane closes, and the updated settings appear in the **Switch Lag Enhanced Hash Settings** grid table.

Each column lists sub-settings for **L2**, **L2 GRE Inner L2**, and other configurations.

Select a link in **Switch Lag Enhanced Hash Settings** to turn these on and off, depending on the data required. There are validations to ensure the correct combination of settings (such as **L2 GRE Inner L2** and **L2 GRE Inner L3** cannot be set at the same time) that will appear in an `AlertMessageList` in the pane. The configuration must pass all validations to submit the changes successfully.

The Symmetric field has three options: **Default**, **Enabled**, and **Disabled**.

- Disable the **Symmetric** field for the GTP Match settings.
- Enter **First Byte** and **First Byte Mask** as hex values 0–255 with up to four **Port Match Entries**.
- Each port match entry must have a port combo and number values for the corresponding ports.

When setting the Port Combo to **AND** or **OR**, both **Src Port** and **Dst Port** must have non-zero values.

### 7.1.3 Minimum Link and Activation Configuration for LAG Interfaces

When some LAG member links go down, it may be preferable to isolate the filter switch by bringing down the entire LAG interface rather than delivering unreliable data to tools and devices.

Two additional commands are part of the DANZ Monitoring Fabric (DMF) **lag-interface** configuration to aid in managing the LAG interface when a specified number of links go down. These commands are:

- **minimum-link**
- **activate**

#### minimum-link

The **minimum-link** command configures the minimum number of links that must be up for a functional LAG. If the total number of active links in the **interface-group** is less than the configured minimum threshold, DMF brings down all active links belonging to the interface group.

The **minimum-link** command is optional, with a default value of 0. You should ensure this value is reasonable. If the value exceeds the total member links, the LAG will always be down.

#### activate

You must use the **activate** command to manually try to bring up the LAG if it was shut down because the number of active links was below the minimum link.



**Note:** If you make any changes to fix the downed LAG members and subsequently use the **activate** command to re-enable the LAG, it will go back down if the active members are still less than the minimum links value.

### Configuration

The **minimum-link** and **activate** commands are available within a switch's lag-interface submode. The following example illustrates configuring these options:

Use the CLI to enter the switch's submode.

1. Enter the lag interface submode by using the **lag-interface** command.
2. Use the **minimum-link** command to configure the minimum threshold value needed before bringing down a LAG.
3. Use the **activate** command to bring back up the previously shutdown LAG (due to having fewer active member links than the minimum threshold).

```
dmf-controller-1(config)# switch s1
dmf-controller-1(config-switch)# lag-interface lag1
dmf-controller-1(config-switch-lag-if)# minimum-link 2
dmf-controller-1(config-switch-lag-if)# activate
```

### Show Commands

While in the **switch** and **lag-interface** submodes, the **show running-config** and **show this** commands list the configured **minimum-link** threshold value. However, since **activate** is an action command rather than a configuration value, it doesn't appear when running the **show** commands.

```
dmf-controller-1# show running-config switch core1 lag lag1
! switch
switch core1
!
lag-interface lag1
member a
member b
```

```
member c
member d
minimum-link 2
```

### Troubleshooting

- If a LAG always shows down, ensure the **minimum-link** value is not greater than the total number of member links in that LAG.

### Limitations

- The system doesn't provide any warning when the **minimum-link** value exceeds the total number of member links in that LAG.

## 7.1.4 Tunnel Endpoint for vCenter Integration

This feature supports using Link Aggregation Group (LAG) in the tunnel endpoint configuration and runs on DANZ Monitoring Fabric (DMF) compatible switches supporting LAGs. Refer to the DMF 8.6 Hardware Compatibility Guide to view the list of compatible switches.

### Configuration

Add the configured LAG as an interface in the tunnel endpoint configuration, as shown in the following example.

```
dmf-controller-1(conf)# tunnel-endpoint tep1 switch sw1-leaf-RU35 lag1 ip-  
address 192.168.199.254 mask 255.255.255.0 gateway 192.168.199.1
```

### Show Commands

Review the tunnel endpoint configuration using the following command:

```
dmf-controller-1# show running-config tunnel-endpoint  
  
! tunnel-endpoint  
tunnel-endpoint tep1 switch sw1-leaf-RU35 lag1 ip-address 192.168.199.254 mask  
255.255.255.0 gateway 192.168.199.1
```

### Limitations

- Do not use LAG members as a Tunnel endpoint interface.
- Do not configure multiple gateways for the same interface used in Tunnel endpoints.
- Do not use a LAG configured as management or DMF interface (filter or delivery) in the Tunnel endpoint configuration.

## 7.1.5 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, use 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, use 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, use 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, use 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, use 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, use 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 or disable symmetric hashing:

```
controller1(config-switch-hash)# hash symmetric
<cr>
disable Disable symmetric hashing
enable Enable symmetric hashing
```

### 7.1.6 L2 GRE Key Hashing

The L2 GRE Key-based hashing feature 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.

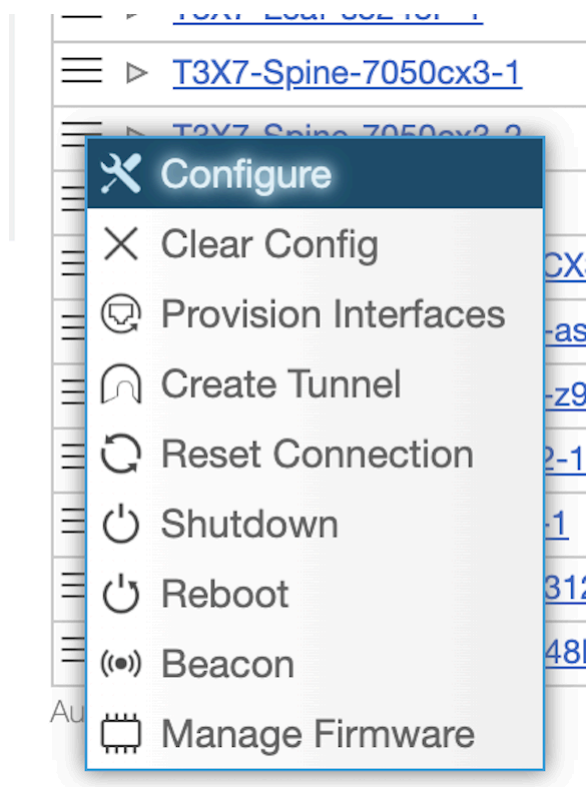
Enable the L2 GRE Key hashing by setting the **l2-gre-key** parameter, as shown in the following example.

```
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
```

#### GUI 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 7-1: 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 7-2: Configure Switch L2 GRE Key**

**Configure Switch**

1. Info ✓ 2. PTP ✓ 3. Clock ✓ 4. SNMP ✓ 5. SNMP Traps ✓ 6. Logging ✓ 7. TACACS ✓ 8. sFlow ✓ 9. LAG Enhanced Hash ✓

☒ Clear All Inputs

**Options**

Symmetric ☐ Seed 1: 0 - 4,294,967,295 Seed 2:

**L2 Fields**

| Field             | Enabled                             | Src. MAC                            | Dst. MAC                            | Ether-Type                          | VLAN ID                             |
|-------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| L2                | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| L2 GRE Inner L2   | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| L2 VXLAN Inner L2 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

**L3 Fields**

| Field           | Enabled                             | Src. IP                             | Dst. IP                             | L4 Src. Port                        | L4 Dst. Port                        | Protocol                            | VLAN ID                             | L2 GRE Key                          |
|-----------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| IPv4            | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| IPv6            | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| L2 GRE Inner L3 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| VXLAN Inner L3  | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

**MPLS Fields**

| Field | Enabled                             | Label 1                             | Label 2                             | Label 3                             | Label High Bits                     | Payload Src. IP                     | Payload Dst. IP                     |
|-------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| MPLS  | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

**GTP Match Fields**

First Byte: Decimal  Hex   
 First Byte Mask: Decimal  Hex

+ -

☐ Sequence  Src. Port  Port Combo  Dst. Port

No port matches

BACK NEXT RESET CANCEL SUBMIT

3. Click the **Submit** button to save the configuration.

### CLI Commands

Use the following CLI commands to verify settings and 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.

## 7.1.7 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:

```

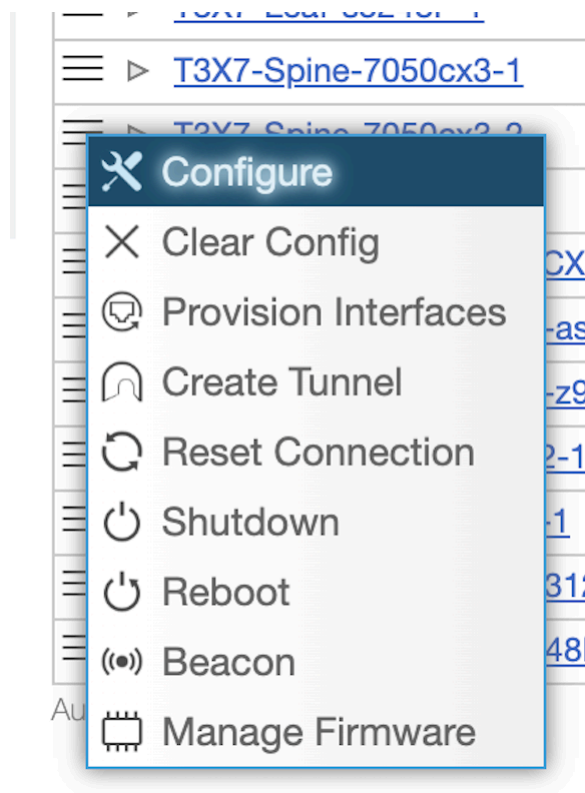
# lag-enhanced-hash
hash vxlan inner-l2 dst-mac
hash vxlan inner-l3 dst-ip

```

### 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 7-3: 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 7-4: Configure Switch LAG Enhanced Hash**

**Configure Switch**

1. Info ✓ 2. PTP Timestamping ✓ 3. Clock ✓ 4. SNMP ✓ 5. SNMP Traps ✓ 6. Logging ✓ 7. TACACS ✓ 8. sFlow ✓ 9. LAG Enhanced Hash ✓

Clear All Inputs

**Options**

Symmetric ☐ Seed 1: 0 - 4,294,967,295 Seed 2: 0 - 4,294,967,295

**L2 Fields**

|                   | Src. MAC                 | Dst. MAC                 | Ether-Type               | VLAN ID                  |
|-------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| L2                | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| L2 GRE Inner L2   | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| L2 VXLAN Inner L2 | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

**L3 Fields**

|                 | Src. IP                  | Dst. IP                  | L4 Src. Port             | L4 Dst. Port             | Protocol                 | VLAN ID                  | L2 GRE Key               |
|-----------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| IPv4            | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| IPv6            | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| L2 GRE Inner L3 | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| VXLAN Inner L3  | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

**MPLS Fields**

|      | Label 1                  | Label 2                  | Label 3                  | Label High Bits          | Payload Src. IP          | Payload Dst. IP          |
|------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| MPLS | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

**GTP Match Fields**

First Byte: Decimal: -Decimal- Hex: -Hex-  
 First Byte Mask: Decimal: -Decimal- Hex: -Hex-  
 Sequence: Src. Port Port Combo Dst. Port  
 No port matches

BACK NEXT RESET CANCEL SUBMIT

- Click the **Submit** button to save the configuration.

### CLI Commands

Use the following CLI commands to verify settings and 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>

```

### 7.1.8 VXLAN LAG Hashing on DCS7280 Platforms

The following information addresses Virtual Extensible LAN (VXLAN) hashing capabilities and behavior, specifically on the DCS-7280 platforms.

#### Platform Compatibility

There are two use cases:

1. **Decapsulation** – Strips the VXLAN header.
2. **Transit** – Retains the VXLAN header.

Since VXLAN Header Stripping is only supported on DCS-7280R3 platforms, the **Decapsulation** use case is supported only on these platforms.

The **Transit** use case is supported on all DCS-7280 platforms.

#### Configuration

There are two assumptions:

1. The VXLAN packet is IPv4 (outer).
2. The VXLAN packet has a UDP destination port value matching the configured value (default – 4789).

For the decapsulation use case, the packet context is advanced to the start of the payload upon VXLAN parsing on ingress. Hence, you can configure normal (outer) hash fields to hash against the VXLAN payload (inner) as if the outer encapsulation has already been discarded.

Traffic is hashed based on the outer header for the transit use case, just like any non-VXLAN UDP packet. Arista advises configuring L4-src-port hashing to hash against the UDP source port of the VXLAN packet, typically used as an entropy of its payload.

```

> enable
# config
(config)# switch switch-name
(config-switch)# lag-enhanced-hash

```

#### Limitations

- Only IPv4 VXLAN tunnel header stripping is supported.



- Currently, it is not possible to limit VXLAN packet parsing to selected interfaces on a given switch. Configuring **strip-vxlan** on any switch interface will trigger VXLAN packet parsing to be active globally on the switch. In this situation, the packet context of any ingressed VXLAN packet would be advanced to the start of its payload. As a result, policy matching and hashing behaviors will be generally affected for VXLAN packets on this switch, even for packets not subject to **strip-vxlan**.

### 7.1.9 MLAG RTAG7 Hash Computation for Unbalanced Load Balancing

In DANZ Monitoring Fabric (DMF), all SWL OS switches with the same underlying ASIC have the same RTAG7 hash parameters configured by default. After configuring MLAG to load balance traffic from filter to delivery switches, and when the number of interfaces in the MLAG and LAG in the delivery switch is the same, the traffic will not be load balanced in the delivery switch because of hash polarization.

Use the feature to avoid LAG hash polarization by choosing different hash algorithms in separate switches. The other option is using a different packet field set in each switch.

The feature supports using different hash algorithms by choosing different hash seed values.



**Note:** MLAG RTAG7 hash computation for unbalanced load balancing only applies to switches running SWL OS, does **not** apply to switches running EOS and is only supported on select Broadcom® switch ASICs that run SWL OS.

#### CLI Configuration

Configure the feature on each switch using the following CLI commands:

```
(config)# switch S4048T
(config-switch)# lag-enhanced-hash
(config-switch)# hash symmetric enable
(config-switch-hash)# hash seeds 1234 3456
(config-switch-hash)# hash 12 dst-mac
(config-switch-hash)# hash 12 src-mac
(config-switch-hash)# hash 12 eth-type
(config-switch-hash)# hash 12 vlan-id
```

Use the following **show** command to view the configured hash seed.

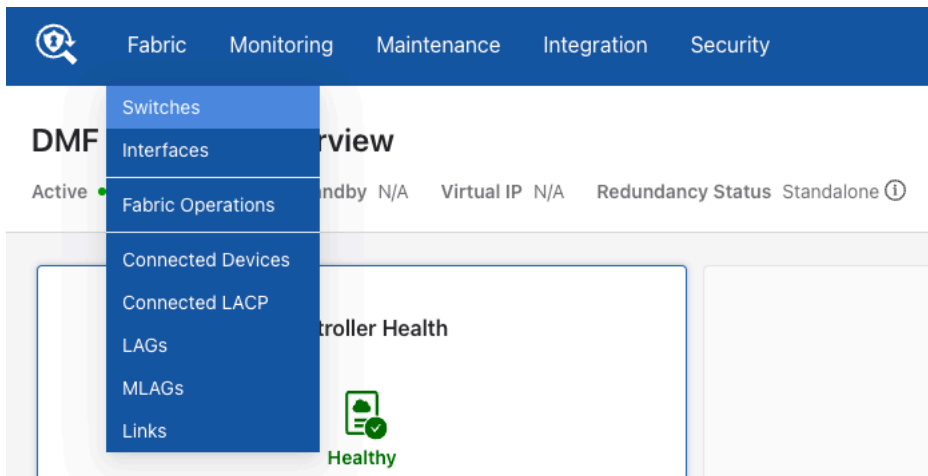
```
R450-C1# show lag-enhanced-hash switch S4048T
~~~~~ L2 Enhanced Hash ~~~~~
Switch DPID Dst mac Eth type Src mac Vlan id Seed1 Seed2 Symmetric
-|-----|-----|-----|-----|-----|-----|-----|
1 S4048T True True True True 1234 3456 True
```



**Note:** The CLI example above illustrates an **L2** configuration.

#### GUI Configuration

From the home page, select the switches using **Fabric > Switches**.



A list of available switches displays.

| Name            | Connected | Guaranteed Reason | Admin Status | Zebrosch State | Management Interface | Genstat | Priority | Source Pkts Address | Source Pkts Address | Connected Since            | Connection Time | Management IP Addressing | Allocated IP Address | Pkts  | Interface Count | Flow Count |
|-----------------|-----------|-------------------|--------------|----------------|----------------------|---------|----------|---------------------|---------------------|----------------------------|-----------------|--------------------------|----------------------|-------|-----------------|------------|
| A-7050SX3-48YC8 | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:31:18pm UTC | 2d 23h 2m       |                          |                      | 56970 | 58              | na         |
| A-7050SX3-48YC8 | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:31:07pm UTC | 2d 23h 2m       |                          |                      | 37326 | 68              | na         |
| A-7050SX3-48YC8 | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:30:30pm UTC | 2d 23h 2m       |                          |                      | 49166 | 68              | na         |
| S4048E          | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:30:17pm UTC | 2d 23h 2m       |                          |                      | 38886 | 56              | na         |
| S4048E          | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:30:06pm UTC | 2d 23h 2m       |                          |                      | 36942 | 56              | na         |
| S4132T          | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:34:17pm UTC | 2d 23h 2m       |                          |                      | 59556 | 17              | na         |
| S4132T          | ✓         |                   | Up           | OK             |                      |         |          |                     |                     | Dec 8, 2023, 9:36:18pm UTC | 2d 22h 48m      |                          |                      | 57834 | 17              | na         |

Choose the required switch for configuring the hash seeds. Hash seeds are integers used to generate random numbers in the RTAG7 hash algorithm.




Select **Actions** from the menu.

Fabric Monitoring Maintenance Integration Security

Fabric: Healthy

DMF SWITCH A-7050SX3-48YC8 [ARISTA 7050SX3-48YC8]



Properties Environment Interfaces **Actions** LAG Enhanced Hash

- Configure Switch
- Clear Switch Config
- Create Tunnel
- Manage Firmware
- Provision Interfaces
- Reset Connection
- Shutdown
- Reboot
- Beacon

On the **Configure Switch** page, select **LAG Enhanced Hash**.

**Configure Switch**

1. Info ✓
2. PTP ✓
3. Clock ✓
4. SNMP ✓
5. SNMP Traps ✓
6. Logging ✓
7. TACACS ✓
8. sFlow ✓
- 9. LAG Enhanced Hash ✓**

☐ Clear All Inputs

**Options**

Symmetric ☒ Seed 1: 1234 (0 - 4,294,967,295) Seed 2: 3456 (0 - 4,294,967,295)

**L2 Fields**

|                 | Src. MAC                            | Dst. MAC                            | Ether-Type                          | VLAN ID                             |                          |
|-----------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------|
| L2              | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| L2 GRE Inner L2 | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/> |
| VxLAN Inner L2  | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/>            | <input type="checkbox"/> |

**L3 Fields**

|                 | Src. IP                  | Dst. IP                  | L4 Src. Port             | L4 Dst. Port             | Protocol                 | VLAN ID                  | L2 GRE Key               |                          |
|-----------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| IPv4            | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| IPv6            | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| L2 GRE Inner L3 | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| VxLAN Inner L3  | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

**MPLS Fields**

|      | Label 1                  | Label 2                  | Label 3                  | Label High Bits          | Payload Src. IP          | Payload Dst. IP          |                          |
|------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| MPLS | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

**GTP Match Fields**

First Byte: Decimal: -Decimal- Hex: -Hex-


First Byte Mask: Decimal: -Decimal- Hex: -Hex-

+ -

☐ Sequence ☐ Src. Port ☐ Port Combo ☐ Dst. Port

No port matches

BACK NEXT RESET CANCEL **SUBMIT**



Toggle **on** the following:

- Options

- Symmetric
- L2 Fields
  - Src. MAC
  - Dst. MAC
  - Ether-Type
  - VLAN ID

Enter the hash **Seed 1** and **Seed 2** values in the LAG Enhanced Hash window.

Click **Submit**.



**Note:** The GUI example above illustrates an **L2** configuration.

**Optional:** View the configured hash seed using the CLI and the following **show** command.

```
R450-C1# show lag-enhanced-hash
~~~~~ L2 Enhanced Hash ~~~~~
# Switch DPID      Dst mac Eth type Src mac Vlan id Seed1 Seed2 Symmetric
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 A-7050SX3-48YC8 True      True      True      True      1234  3456  True
```

### 7.1.9.1 Troubleshooting a Configured Hash Seed

Use the following procedure to ensure the configured hash seed is correctly applied on the switch.

1. Log in to the switch using the **connect switch switch\_name** command.
2. Use the **ofad-ctl gt PORT\_CHANNEL\_ENHANCED\_HASH\_SEED** command to print the seed values and the hash algorithm.
3. Ensure the hash algorithms differ in the switches where hash polarization occurs.
4. Configure the hash seeds such that the hash algorithms are different.

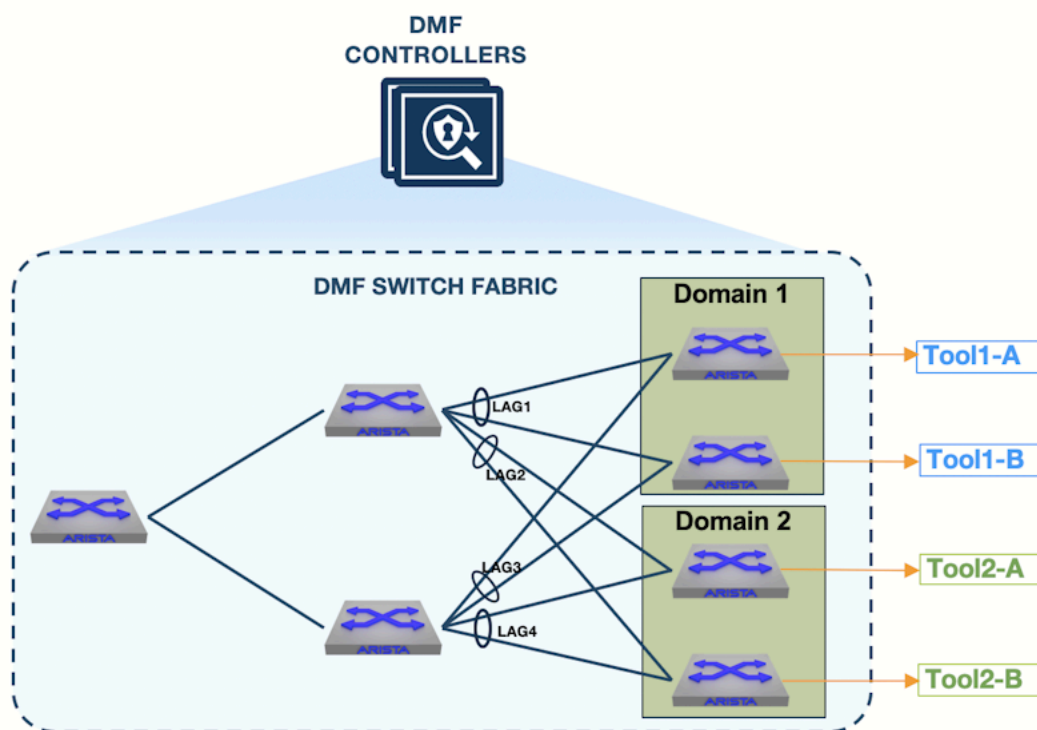
#### Example

```
R450-C1(config)# connect switch A-7050SX3-48YC8
Switch Light OS SWL-OS-DMF-8.6.x(0), 2023-12-07.09:17-42d8658
Linux A-7050SX3-48YC8 4.19.296-OpenNetworkLinux #1 SMP Thu Dec 7 09:28:30 UTC
2023 x86_64
Switch Light ZTN Manual Configuration. Type help or ? to list commands.
(ztn-config) debug bash
***** WARNING *****
Any/All activities within bash mode are UNSUPPORTED
This is intended ONLY for additional debugging ONLY by Arista TAC.
Please type "exit" or Ctrl-D to return to the CLI
***** WARNING *****
root@A-7050SX3-48YC8:~# ofad-ctl gt PORT_CHANNEL_ENHANCED_HASH_SEED
GENTABLE : port_channel_enhanced_hash_seed
GENTABLE ID : 0x0006
hash_seed_valid: true
HASH_SEED1=0x000004d2
HASH_SEED2=0x000004d2
computed hash algorithm from seed: CRC32LO
root@A-7050SX3-48YC8:~#
```

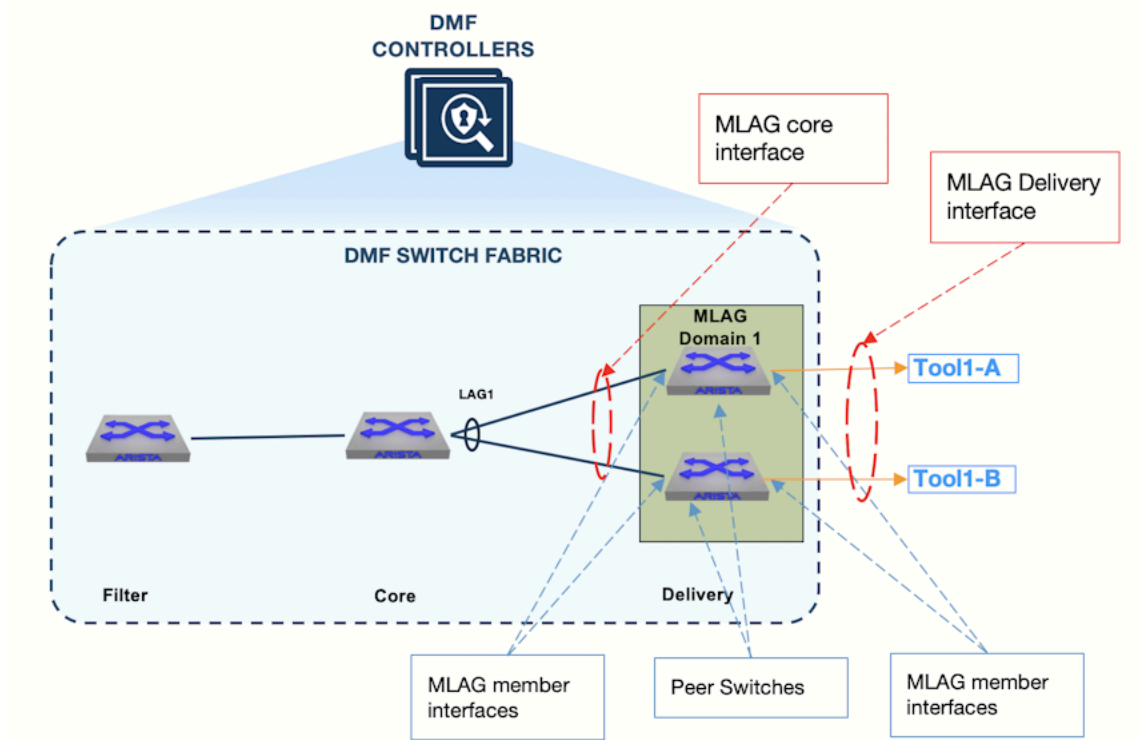
## 7.2 Pseudo Multi-Chassis Link Aggregation

DMF supports 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 connect to the same DMF delivery switch, typically when co-locating customer tools in the same physical location. However, in cases where tools reside 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 redundancy for each delivery switch 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. If one of the switches fails, the traffic will fail 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
```

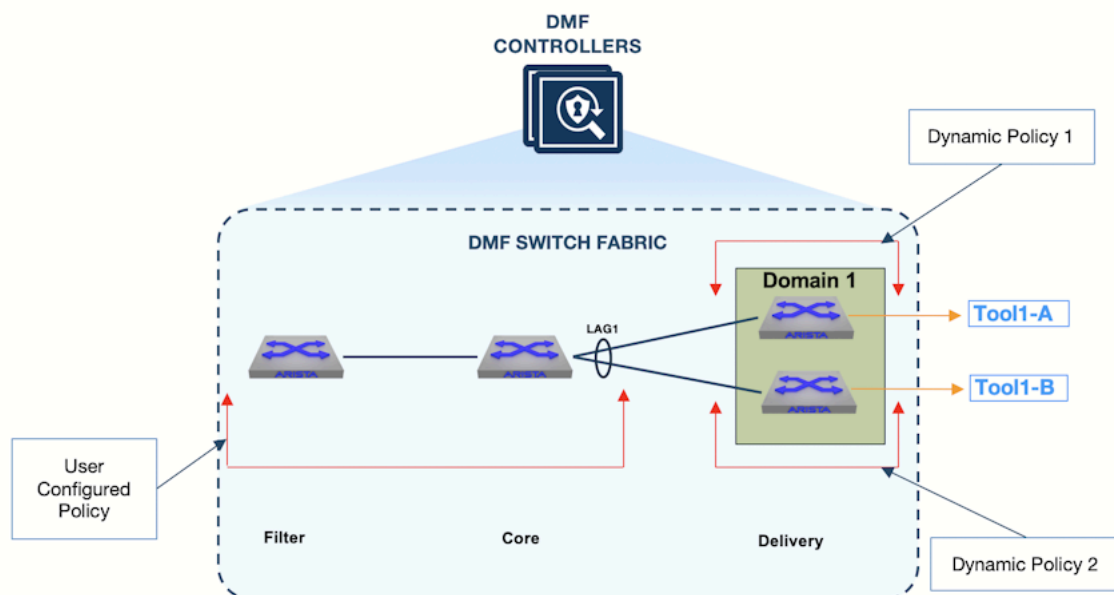
The above MLAG interface configuration selects one physical interface from each peer switch added into the MLAG domain. The members of this MLAG interface, **ethernet1** of **DeliverySwitch-1** and **ethernet1** of **DeliverySwitch-2**, are connected to multi-homed tools. Note that unlike the core MLAG 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 7-5: 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     | 61296bps    | 61296bps                  |                  |          |
| 3 | _mlag_Policy-1_DeliverySwitch-2 | forward | installed      | Dynamic    | 100      | 1       | 1             | 40Gbps | 1Gbps     | 61386bps    | 61386bps                  |                  |          |

```
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

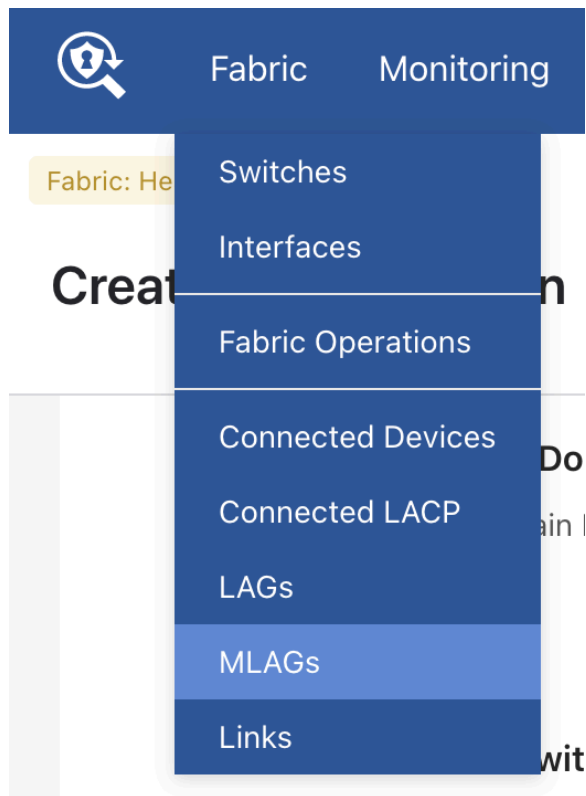
Link Layer Discovery Protocol (LLDP) is used to discover MLAG links. When the DMF Controller receives an LLDP message, it looks for the switch and interface names. If the switch is a part of an MLAG domain, and the reported interface corresponds to the MLAG interface, then it is classified as an MLAG link.

```
Controller-1(config)# show link all link-type mlag-member
----- Links -----
# 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 7-6: Fabric → MLAGs page**



**Figure 7-7: 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 **Create** to save the configuration.

**Figure 7-8: MLAG Domain State**

Multi-Chassis LAGs (MLAGs) Refresh

View and configure MLAG domains and interfaces

Download

Add Filters
Domain Name
Clear current filters
Delete MLAG Domain(s)
Edit MLAG Domain
Create MLAG Domain

| <input type="checkbox"/>            | Domain Name | Peer Switches     | MLAG Interfaces | MLAG Delivery Interfaces | Status |
|-------------------------------------|-------------|-------------------|-----------------|--------------------------|--------|
| <input checked="" type="checkbox"/> | Domain1     | DMF-D1 and DMF-D2 | 2               | 1                        | Active |

1-1 of 1 items 1 / page

**Figure 7-9: MLAG Domain Expanded View**

Multi-Chassis LAGs (MLAGs) Refresh

View and configure MLAG domains and interfaces

Download

Add Filters
Domain Name
Clear current filters
Delete MLAG Domain(s)
Edit MLAG Domain
Create MLAG Domain

| <input type="checkbox"/>            | Domain Name | Peer Switches     | MLAG Interfaces | MLAG Delivery Interfaces | Status |
|-------------------------------------|-------------|-------------------|-----------------|--------------------------|--------|
| <input checked="" type="checkbox"/> | Domain1     | DMF-D1 and DMF-D2 | 2               | 1                        | Active |

**MLAG Interfaces**

| Name          | Member Interfaces                          | Status |
|---------------|--------------------------------------------|--------|
| Intf-core-del | DMF-D1 / ethernet51<br>DMF-D2 / ethernet51 | Active |

**MLAG Delivery Interfaces**

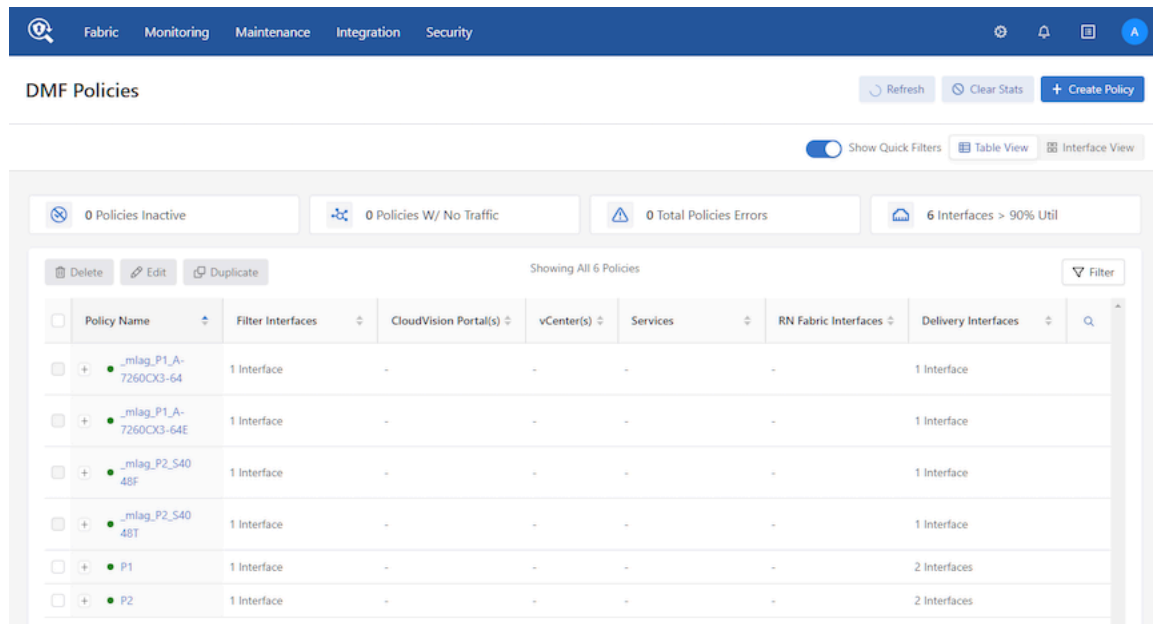
| Name          | Dmf Name    | Role     | Member Interfaces                          | Strip Vlan on Egress | Status |
|---------------|-------------|----------|--------------------------------------------|----------------------|--------|
| Intf-del-tool | MLAG-Tool-1 | Delivery | DMF-D1 / ethernet23<br>DMF-D2 / ethernet11 | -                    | Active |

The above screenshot displays the MLAG domain status. Click **+** to expand each MLAG interface's domain configuration and status of each MLAG interface.

## 7.8 Create MLAG Policy from GUI

To configure an MLAG policy, go to the **Monitoring > Policies** page.

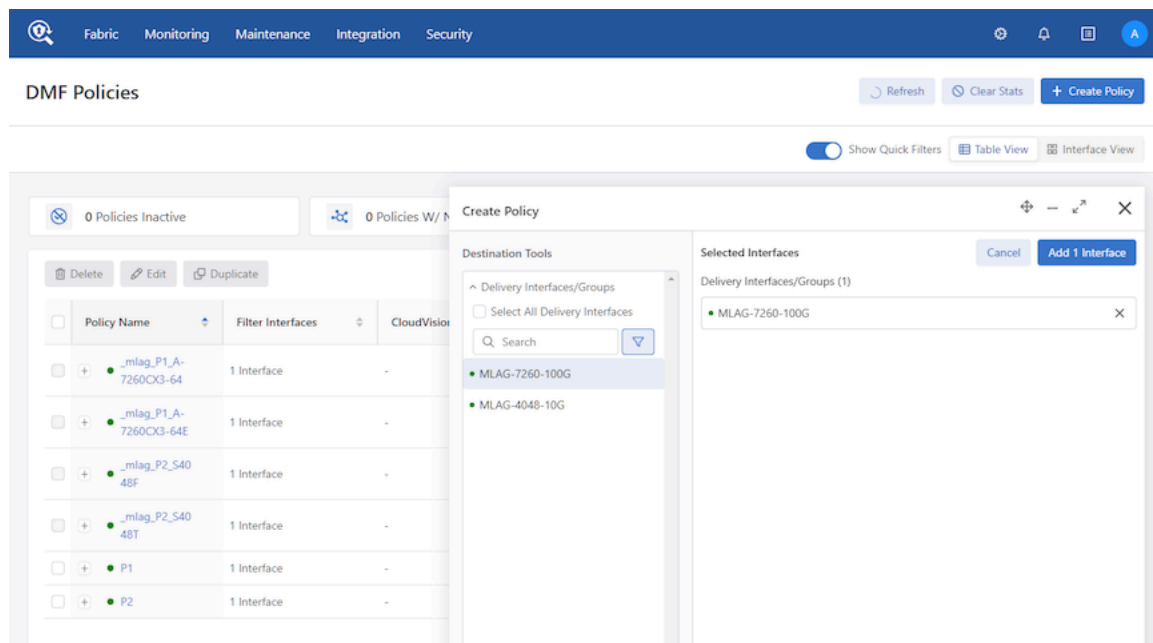
**Figure 7-10: Configure MLAG Policy**



Select **+ Create Policy** to add a new policy.

Click **+ Add Ports** in **Destination Tools** and select or drag the MLAG interface to associate with the policy.

**Figure 7-11: Add MLAG Interface**



Click **Add 1 Interface** and enter the following to configure the policy association.

- **Name:** Assign a unique name to the policy.
- **Description:** A description for the policy.
- **Action:** **Forward** (default), **None**, **Capture**, or **Drop**.

- **Priority: 100** (default), or enter a value.
- **Scheduling: Automatic** (default), **Now**, **Set Time**, or **Set Delay**.
- **Port Selection > Traffic Sources > + Add Port(s)**: Select the filter interface (traffic source) for the policy.
- **Match Traffic > Allow All Traffic / Deny All Traffic or Configure A Rule**: Specify the traffic rule for the policy.

Click **Create Policy**.

## 7.9 Viewing Policy Statistics in the GUI

After configuring the MLAG policy, view it under **Monitoring > Policies** along with the dynamic policies created as part of the MLAG policy.

**Figure 7-12: MLAG Policy**

| Policy Name            | Filter Interfaces | CloudVision Portal(s) | vCenter(s) | Services | RN Fabric Interfaces | Delivery Interfaces |
|------------------------|-------------------|-----------------------|------------|----------|----------------------|---------------------|
| _mlag_P1_A-7260CX3-64  | 1 Interface       | -                     | -          | -        | -                    | 1 Interface         |
| _mlag_P1_A-7260CX3-64E | 1 Interface       | -                     | -          | -        | -                    | 1 Interface         |
| _mlag_P2_S404_8F       | 1 Interface       | -                     | -          | -        | -                    | 1 Interface         |
| _mlag_P2_S404_8T       | 1 Interface       | -                     | -          | -        | -                    | 1 Interface         |
| P1                     | 1 Interface       | -                     | -          | -        | -                    | 2 Interfaces        |
| P2                     | 1 Interface       | -                     | -          | -        | -                    | 2 Interfaces        |

To view the policy statistics, roll over a **Policy Name** or click on the **MLAG Policy Name**. The following screens appear.

Figure 7-13: MLAG Policy - Interface Statistics

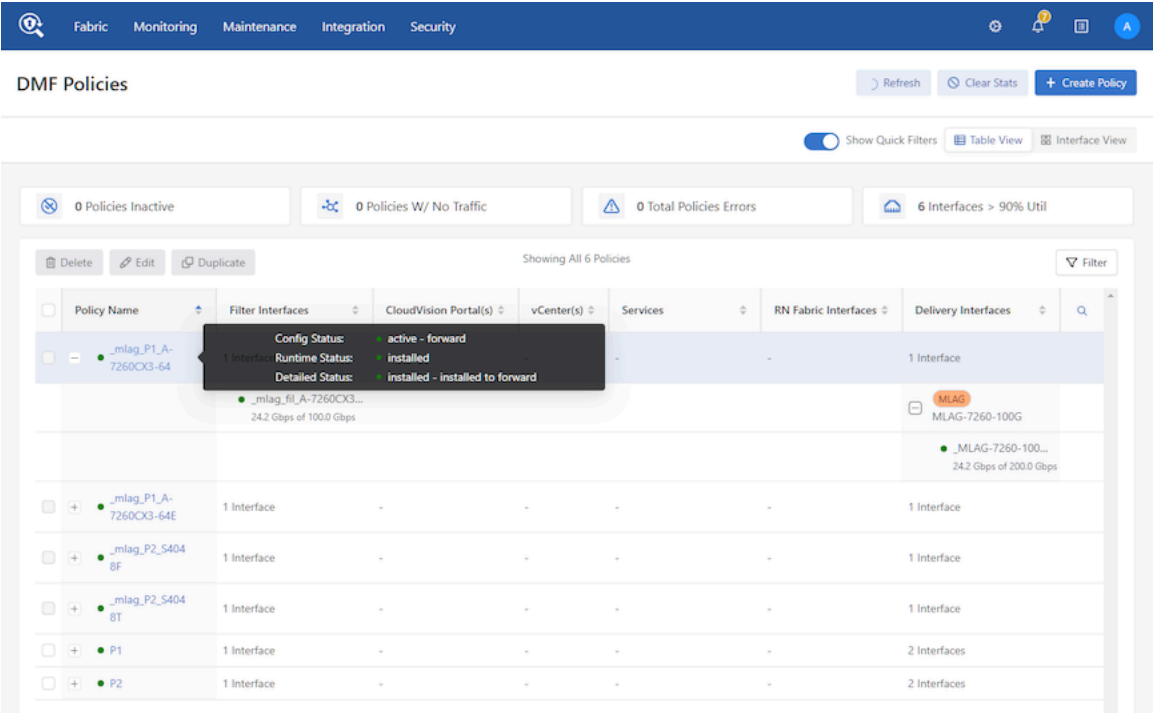
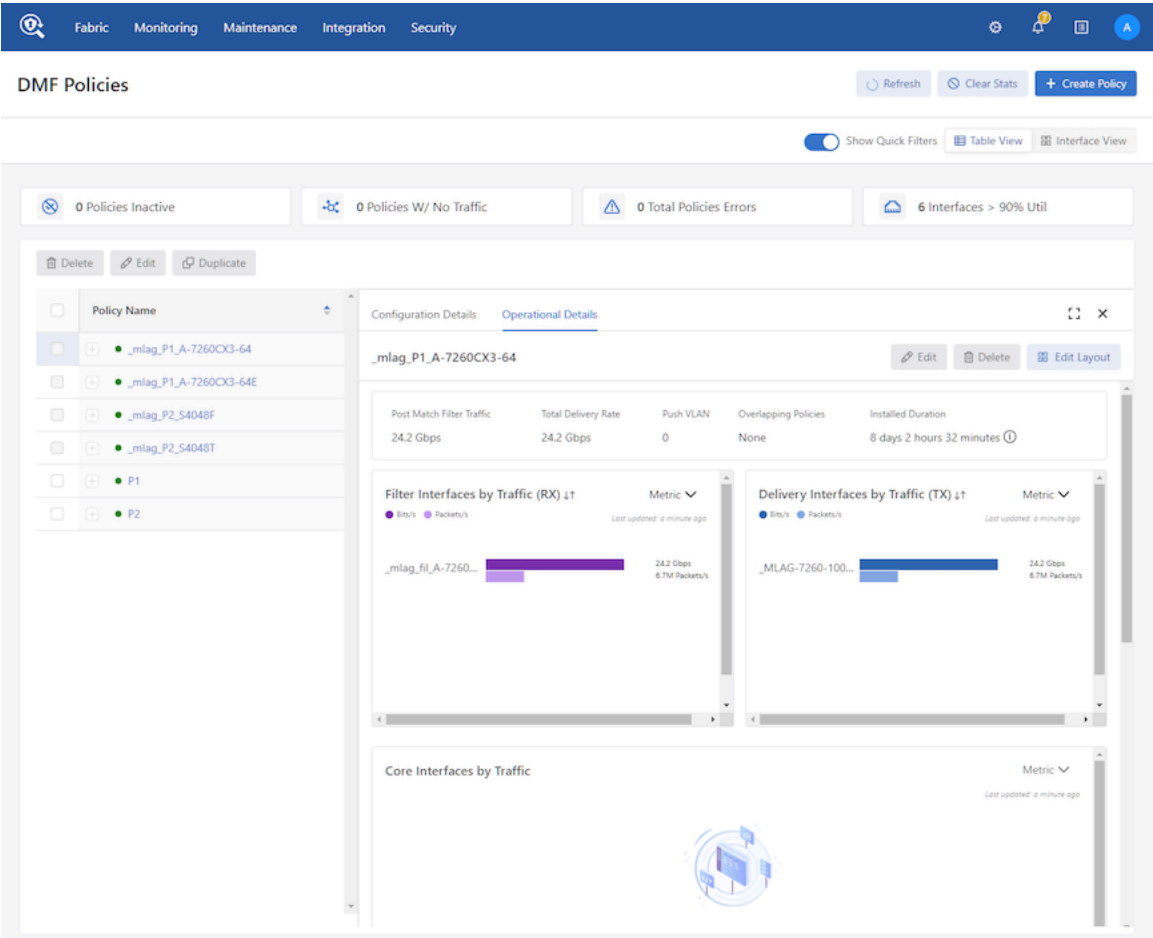


Figure 7-14: MLAG Policy - Operational Details



## 7.10 Viewing MLAG Links in the GUI

To view the MLAG links, go to **Fabric > Links > MLAG Member Links** tab.

**Figure 7-15: MLAG Member Links**

Links *Last updated: a few seconds ago*

Switch Links   LAG Links   **MLAG Member Links**   History

| 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 LLDP message exchange discovers the links.

## 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 differs 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 are configured to use an MLAG interface as a delivery interface overlap, the following policies are created:

1. **Policy-1** uses the filter interface **Filter-1** and the delivery interface **MLAG-Tool-1**.
2. **Policy-2** uses 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**.

After calculating the overlapping policy for the two user-configured policies, DMF configures two dynamic policies: one for each peer switch in the MLAG domain and one for each of the three policies listed above.

**Table 4: Dynamic Policies**

| MLAG Dynamic Policy                               | Parent Policy               | Delivery Switch/Peer switch |
|---------------------------------------------------|-----------------------------|-----------------------------|
| <b>_mlag_Policy-1_DeliverySwitch-1</b>            | <b>Policy-1</b>             | <b>DeliverySwitch-1</b>     |
| <b>_mlag_Policy-1_DeliverySwitch-2</b>            | <b>Policy-1</b>             | <b>DeliverySwitch-2</b>     |
| <b>_mlag_Policy-2_DeliverySwitch-1</b>            | <b>Policy-2</b>             | <b>DeliverySwitch-1</b>     |
| <b>_mlag_Policy-2_DeliverySwitch-2</b>            | <b>Policy-2</b>             | <b>DeliverySwitch-2</b>     |
| <b>_mlag_Policy-1_o_Policy-2_DeliverySwitch-1</b> | <b>_Policy-1_o_Policy-2</b> | <b>DeliverySwitch-1</b>     |
| <b>_mlag_Policy-1_o_Policy-2_DeliverySwitch-2</b> | <b>_Policy-1_o_Policy-2</b> | <b>DeliverySwitch-2</b>     |

The following policies, **Policy-1** and **Policy-2**, share the same filter interface, **Filter-1**, but are configured to use different delivery interfaces, **MLAG-Tool-1** and **MLAG-Tool-2**. No priority is configured; these policies use 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 Fil Traff Del Traffic Services
-----
1 Policy-1         forward installed      Configured      100          0          1      25Gbps      80Gbps
2 Policy-2         forward installed      Configured      100          0          3      25Gbps      80Gbps
3 _Policy-1_o_Policy-2 forward installed      Dynamic         100          1          5      25Gbps      80Gbps

```

- **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 forwards to both **MLAG-Tool-1** and **MLAG-Tool-2**.

The 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 Delivery 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 than the rest to prioritize the overlapping traffic forwarded to **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      157Mbps      315Mbps
9 _mlag_Policy-1_o_Policy-2_DeliverySwitch-2 forward installed      Dynamic         100          3          8      40Gbps      11Gbps      157Mbps      315Mbps

```

- **\_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 establishing 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. Use this feature to extend a DMF deployment across multiple data centers or branch offices over networks connected by Layer-3 networks. This feature supports the centralization or distribution of tools and taps across multiple locations when they cannot be cabled directly.



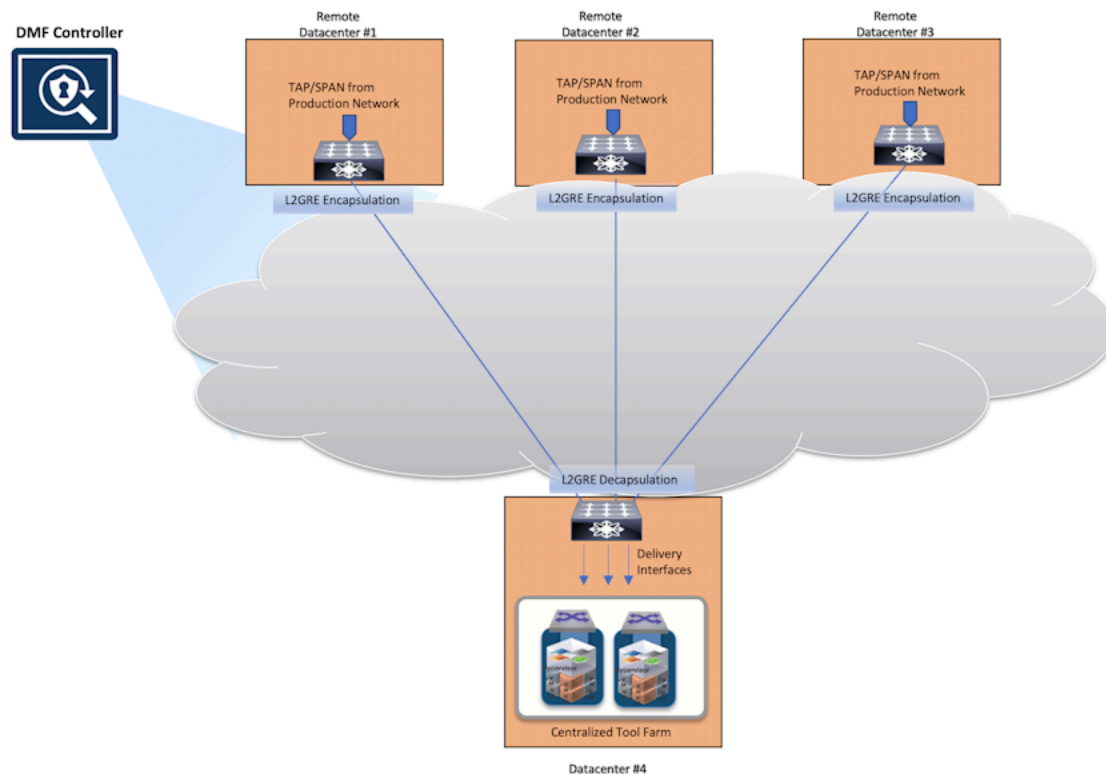
**Note:** Refer to the *DANZ Monitoring Fabric 8.6 Hardware Compatibility List* for a list of the switches that support tunneling. The *DANZ Monitoring Fabric 8.6 Verified Scale Guide* indicates the number of tunnels supported by each supported switch (Verified Scalability Values/Encap Tunnels/Decap 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 ensure 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 is disabled, re-enable it before configuring a tunnel.
- In the case of GRE tunnels, the optional **gre-key-decap** value on the receiving end must match the GRE key value of the sender. The option exists to set multiple values on the same tunnel to decapsulate traffic with different keys.
- A single switch can initiate multiple tunnels. Configure a separate encap-loopback-interface for each tunnel (transmit-only or bidirectional).

- Set the **loopback-mode** 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 8.5: 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.
- 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 work around this limitation.
- VXLAN tunnels are currently NOT supported on 7280 switches.

## 8.2 Encapsulation Type

DANZ Monitoring Fabric (DMF) supports VXLAN tunnel type and Level-2 Generic Routing Encapsulation (L2GRE). The tunnel type is a per-switch configuration, setting the switch pipeline 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:

| Name     | Quarantine Reason | Admin Status | Management Interface                  | ASIO | Connected Since            | Connection Time | MAC                         |
|----------|-------------------|--------------|---------------------------------------|------|----------------------------|-----------------|-----------------------------|
| Switch-1 |                   | Up           | Prefer Dedicated Management Interface |      | Mar 4, 2022, 4:45:17pm UTC | 5d 5h 2m        | fe80::d6af:f7ff:fe99:e2b0%9 |

The encapsulation type can be edited for an existing switch from the **Fabric > Switches > Configure Switch** page, as shown in the figure below:

| Name     | Connected | Quarantine Reason | Admin Status |
|----------|-----------|-------------------|--------------|
| Switch-1 | ✓         | —                 | ✓ Up         |
| Switch-2 | ✗         | —                 | ✓ Up         |

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  l3-14-offset-match-push-vlan-vxlan
2 Switch-2         fe80::e6f0:4ff:fe69:6aee%9   connected  l3-14-offset-match-push-vlan
3 Switch-3         fe80::e6f0:4ff:fe78:1ffe%9   connected  l3-14-offset-match-push-vlan-vxlan
Ctrl-1(config)#
```

In the above CLI output, **Switch-1** and **Switch-3** use 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, filter interface, or delivery interface. The most common use case is linking multiple sites, using the tunnel as a core link. If used as a core link, DMF automatically discovers the link as if it were a physical link and similarly determines connectivity (link-state). 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. Use the tunnel endpoint as a delivery interface to encapsulate filtered monitoring traffic to send to analysis tools or another DANZ Monitoring Fabric managed by a different DMF Controller.



**Note:** By default, sFlow<sup>®\*</sup> and other Arista Analytics metadata cannot be generated for decapsulated L2 GRE/VXLAN tunneled production traffic on a tunnel interface configured as a filter interface.

To generate this metadata, 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 metadata will now be generated for the decapsulated tunnel traffic.

## 8.4 Using the GUI to Configure a GRE Tunnel

To configure a VXLAN tunnel, perform the following steps:

1. Select **Fabric > Switch**.
2. On the **Switches** page, click the **Menu** control next to the switch or interface to include in the tunnel and select **Create Tunnel**.

\* sFlow<sup>®</sup> is a registered trademark of Inmon Corp.

Alternatively, configure tunnels from the **Fabric > Interfaces** page by clicking on the **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, set the tunnel direction to 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. Use this physical interface for tunnel encapsulation and not 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:** Configure this procedure on both switches at each end of the tunnel. Set the **Auto VLAN** mode 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, perform the following steps:

1. Connect switch ports (on remote and main datacenter) to their respective WAN routers and ensure 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 encap-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 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
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 encap-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**. When attempting to set a value below this limit, the switch will still set the rate limit value to **25 kbps**.

```
switch DMF-CORE-SWITCH-1
gre-tunnel-interface tunnel1
direction bidirectional encap-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 apply to GRE tunnel interfaces.

Use the **show running-config** command to view the configuration of tunnel interfaces.

Enter the **show tunnel** command to see a tunnel interface's configuration parameters and runtime state.

```
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, perform the following steps:

1. Select **Fabric > Switch**.
2. On the **Switches** page, click the **Menu** control next to the switch or interface to include in the tunnel and select **Create Tunnel**.



Alternatively, configure tunnels from the **Fabric > Interfaces** page by clicking on the **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 **tunnel**.
- **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 encaps side.
- **Direction:** bidirectional, transmit-only, or receive-only. For bidirectional tunnels, set tunnel-direction to 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. Use this physical interface for tunnel encapsulation and not 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:** Configure this procedure on both switches at each end of the tunnel. Set the **Auto VLAN** mode 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, perform 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 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 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**. When attempting to set a value below this limit, the switch will still set the rate limit value to **25 kbps**.

```
switch DMF-CORE-SWITCH-1
vxlan-tunnel-interface tunnell
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 apply to tunnel interfaces.

Use the **show running-config** command to display the configuration of tunnel interfaces.

Enter the **show tunnel** command to see the configuration parameters and 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 tunnell ESTABLISHED_TUNNEL bidirectional 198.82.215.1 216.47.143.1 ethernet5:1 ethernet6
2 DMF-CORE-SWITCH-2 tunnell 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 tunnell 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 tunnell
# Switch DPID Tunnel Name Tunnel Status Direction Src IP Dst IP Parent Name Loopback Name
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 DMF-CORE-SWITCH-2 tunnell 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. DMF displays the tunnel configuration as illustrated in the following figure.



**Tip:** When multiple interfaces are present, use the **Filter** feature to locate tunnel interfaces after typing the first few letters of the word *tunnel*.

**Figure 8-4: Tunnel Interfaces**

| Switch            | Switch DPID       | Interface Name | Description | Forward Error Correction | Optics Always Enabled | Disable Transmitting Packets | Management | Span Info | Status | Tunnel |
|-------------------|-------------------|----------------|-------------|--------------------------|-----------------------|------------------------------|------------|-----------|--------|--------|
| acc-5712-254-ru20 | acc-5712-254-ru20 | ethernet6      |             |                          |                       |                              |            |           | Up     |        |
| acc-5712-254-ru20 | acc-5712-254-ru20 | ethernet7      |             |                          |                       |                              |            |           | Up     |        |
| acc-5712-254-ru20 | acc-5712-254-ru20 | ethernet8      |             |                          |                       |                              |            |           | Up     |        |
| acc-5712-254-ru20 | acc-5712-254-ru20 | ethernet9      |             |                          |                       |                              |            |           | Up     |        |
| ari-7050-254-ru23 | ari-7050-254-ru23 | Arista-Tunnel  |             |                          |                       |                              |            |           | Up     | ✓      |

| Tunnel Configuration |               |               |                    |
|----------------------|---------------|---------------|--------------------|
| Encapsulation Type   | GRE           | Local IP      | 192.37.1.2/24      |
| Interface Name       | Arista-Tunnel | Remote IP     | 192.33.1.2         |
| Loopback Interface   | ethernet10    | Gateway IP    | 192.37.1.1         |
| Parent Interface     | ethernet39    | GRE Key Decap |                    |
| Direction            | transmit-only | DSCP          |                    |
|                      |               | Rate Limit    |                    |
| Status               | Up            | Admin Status  | Up                 |
| Link Status          | Up            | Tunnel Status | ESTABLISHED_TUNNEL |

| Core Links                                  |                         |              |                                             |                       |            |                |           |
|---------------------------------------------|-------------------------|--------------|---------------------------------------------|-----------------------|------------|----------------|-----------|
| Source Switch/Interface                     | Source Interface Number | Source Speed | Dest Switch/Interface                       | Dest Interface Number | Dest Speed | Link Type      | Last Seen |
| ari-7050-254-ru23 / Arista-Tunnel           | 321                     | 0 bit/s      | dell-5232-254-ru18-core2 / Arista-rx-tunnel | 321                   | 0 bit/s    | tunnel-allowed | 10s       |
| dell-5232-254-ru18-core2 / Arista-rx-tunnel | 321                     | 0 bit/s      | ari-7050-254-ru23 / Arista-Tunnel           | 321                   | 0 bit/s    | tunnel-blocked | 10s       |

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 tunnel header throws off the offset calculation, and the selected traffic may be dropped. This behavior is due to how switch hardware calculates the anchor and offset concerning incoming packets. When the core link is a tunnel, the anchor and offset calculation differs when encapsulating packets compared to when decapsulating.

There are two ways to work around this issue:

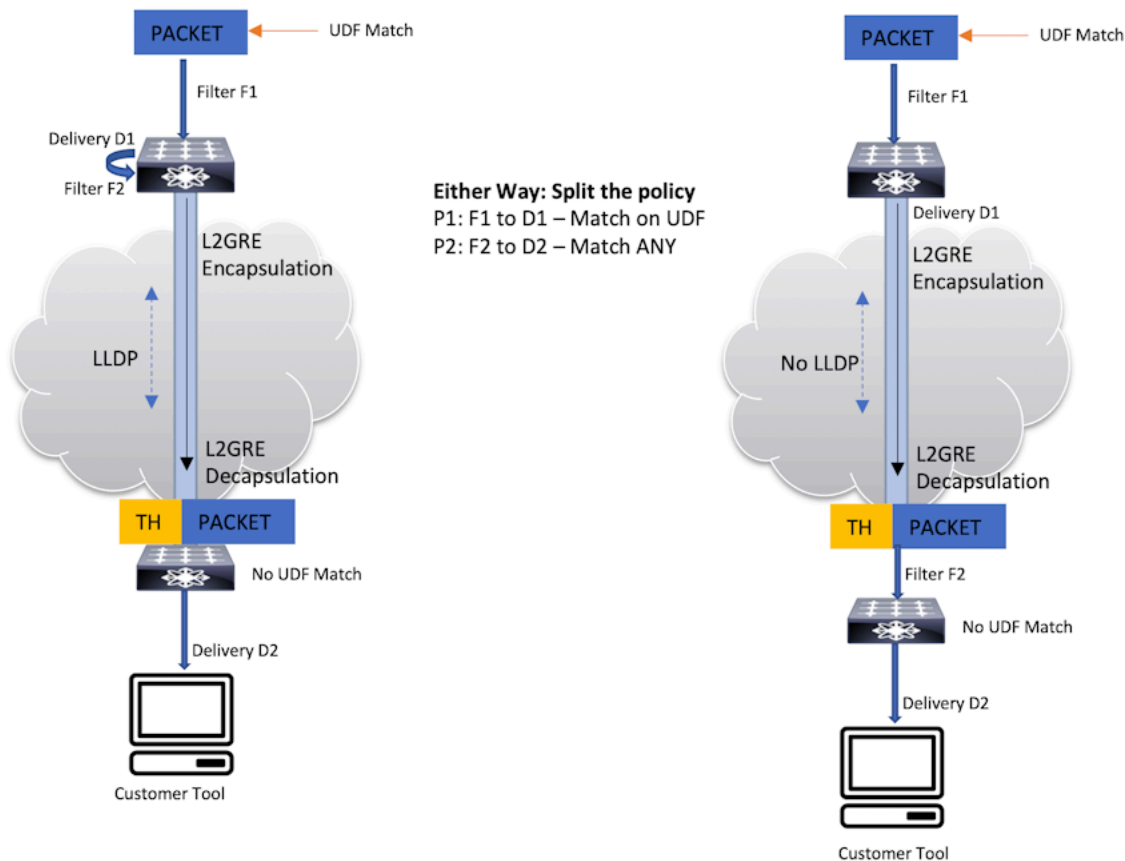
- Avoid matching on user-defined offsets on tunnel interfaces
- Avoid using a tunnel as a core link when matching on a user-defined offset

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. The 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.

## 8.10 Wildcard Tunnels on SAND Platforms

The Wildcard tunneling feature allows the DANZ Monitoring Fabric (DMF) to decapsulate L2GRE-based tunneled traffic from any remote source. This feature, supported on Switch Light OS (SWL) based DMF switches in prior releases, now allows wildcard tunnels on Arista EOS-based DMF switches.

### Platform Compatibility

EOS switches with Jericho2 or higher ASICs compatible with DMF 8.5.0 that support L2GRE tunneling. L2GRE tunneling on EOS SAND platforms is only supported on Arista 7280R3 switches.

### 8.10.1 Configuring Wildcard Tunnels Using the CLI

Perform the following steps to configure the tunnels.

1. Enable the tunneling feature before configuring tunnels on a switch using the **tunneling** command. Enter **yes** when prompted to continue.

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

2. Configure a tunnel by using the **remote-ip** as 0.0.0.0.

```
dmf-controller-1(config)# switch main-dc-delivery-sw
dmf-controller-1(config-switch)# gre-tunnel-interface tunnel1
dmf-controller-1(config-switch)# remote-ip 0.0.0.0 == this is to enable wildcard tunnel
dmf-controller-1(config-switch)# gre-key-decap 4097
dmf-controller-1(config-switch)# parent-interface ethernet5
dmf-controller-1(config-switch)# local-ip 192.168.200.50 mask 255.255.255.0 gateway-ip 192.168.200.1
dmf-controller-1(config-switch)# direction receive-only
```

Please refer to the [Tunneling Between Data Centers](#) chapter for more information on using L2GRE tunnels in DANZ Monitoring Fabric (DMF).

### Show Commands

All CLI show commands for regular interfaces apply to GRE tunnel interfaces. Use the **show running-config** command to view the configuration of tunnel interfaces.

Enter the **show tunnel** command to view a tunnel interface's configuration parameters and runtime state.

### Example

```
dmf-controller-1# show tunnel
# Switch DPID      Tunnel Name Tunnel Status      Direction      Src IP      Dst IP      Parent Name Loopback
-----|-----|-----|-----|-----|-----|-----|-----|
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
```

## 8.10.2 Configuring Wildcard Tunnels Using the GUI

In the DANZ Monitoring Fabric (DMF) UI, enable **Tunneling** by navigating to the **DMF Features** page and clicking on the **gear icon** in the navigation bar.

**Figure 8-6: DMF Navigation Menu**

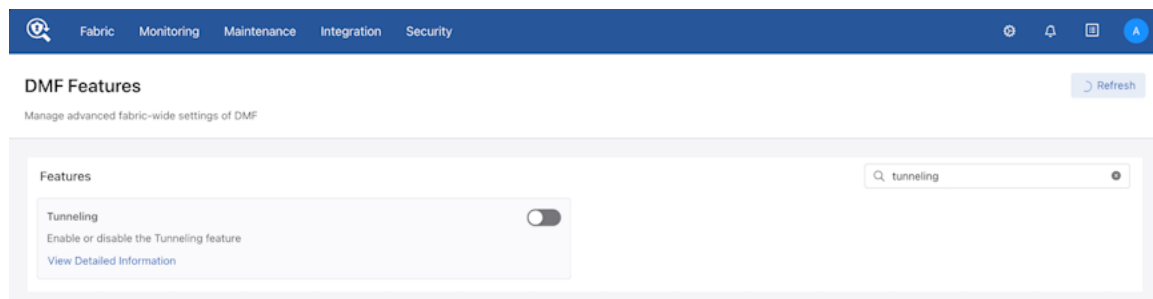


Locate the **Tunneling** card feature on the page.



**Tip:** Use the search bar to quickly locate the feature.

**Figure 8-7: DMF Features**

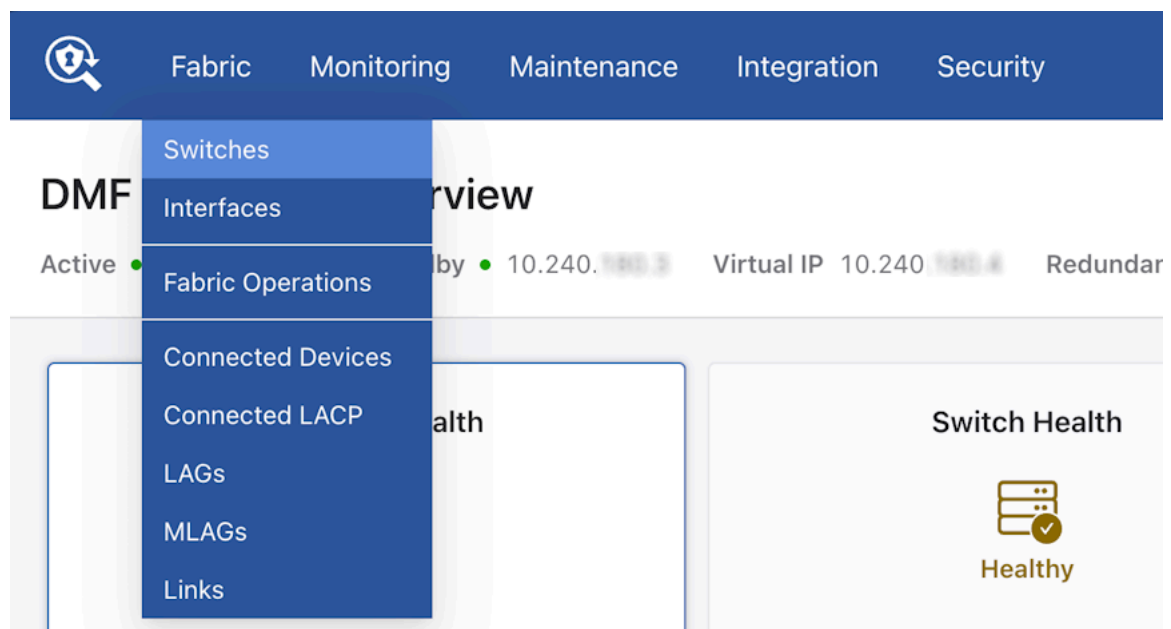


Turn on the **toggle** switch to enable the Tunneling feature on DMF.

### Steps to Enable Wildcard Tunnels

1. Locate the **Switches** tab on the **Fabric > Switches** page.

**Figure 8-8: Fabric > Switches**

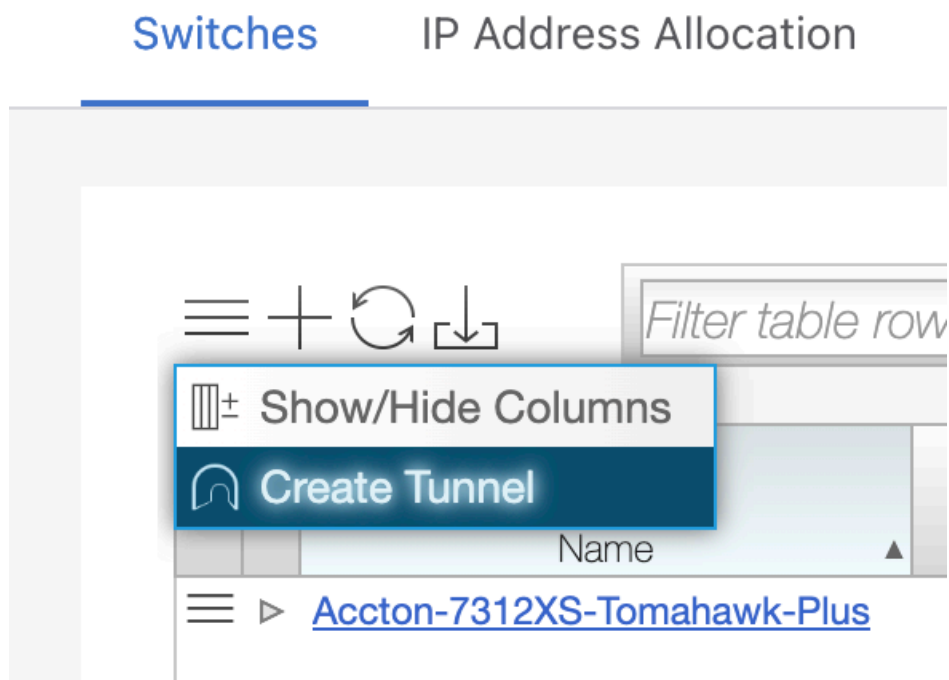


2. Click the **Create Tunnel** option from the table menu.

**Figure 8-9: Create Tunnel**

# Switches

Manage switches connected to DMF Contr





3. Create a **Tunnel** (select **Encapsulation Type** as **GRE**) by entering the required fields. To enable wildcard support, enter the **Remote IP** input with data as shown below.

**Figure 8-10: Configure Tunnel**

The screenshot shows the 'Configure Tunnel' window with the following fields and values:

- Switch \***: - DMF Switch -
- Encapsulation Type \***: GRE
- Name \***: (empty field, note: Tunnel name must begin with tunnel, 16 characters max.)
- Rate Limit**: ☐ 0 kbps
- Direction \***: Bidirectional (Receive and Transmit toggles are both checked)
- Local IP \***: - IPv4 Address - / - Prefix -
- Subnet Mask**: - Subnet Mask -
- Gateway IP \***: - IPv4 Address -
- ARP Interval**: 5 seconds
- Remote IP \***: ☒ Use Wildcard
- Remote IP Value**: 0.0.0.0
- Parent Interface \***: (empty field)
- Loopback Interface \***: (empty field)
- GRE Decap Keys \***: - None - (Note: A list of 32-bit integers)

At the bottom right, there is a 3D illustration of a blue tunnel pipe with light rays passing through it. Below the form are three buttons: RESET, CANCEL, and SUBMIT.

4. Click **Submit** to save the tunnel configuration.

## Integrating vCenter with DMF

This chapter describes integrating VMware vCenter with the DANZ Monitoring Fabric (DMF) and monitoring Virtual Machines (VM) in the vCenter.

### 9.1 Overview

The DANZ Monitoring Fabric (DMF) allows the integration and monitoring of VMs in a VMware vCenter cluster. After integrating a vCenter with the DMF fabric, 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, DMF supports the following versions of VMware vCenter 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:** This method monitors 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 to monitor traffic.

The following figure illustrates the topology used for local SPAN configuration:

**Figure 9-1: Mirroring on a Separate SPAN Physical NIC (SPAN)**



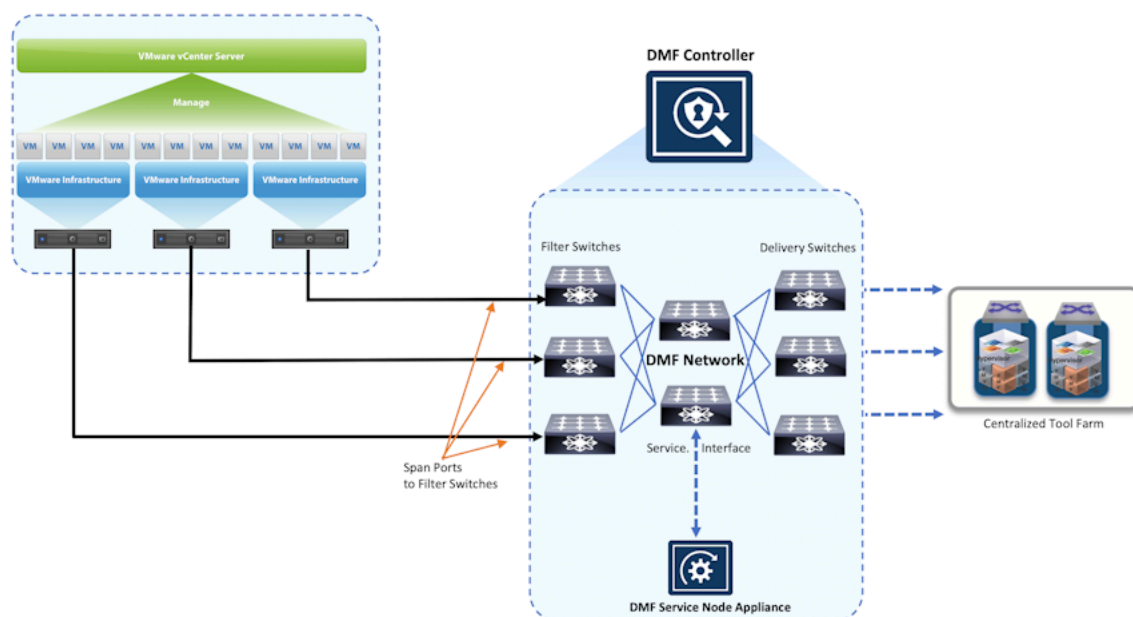
- **Monitoring using ERSPAN/L2GRE tunnels:** Use Remote SPAN (ERSPAN) to monitor VMs running on the ESX hosts within a vCenter instance integrated with DMF. ERSPAN monitors traffic to and from VMs anywhere in the network and does not require a dedicated physical interface card on the ESX host. However, ERSPAN can affect network performance, especially when 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 DANZ Monitoring Fabric (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 to integrate a vCenter instance with the monitoring fabric and deliver 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 connected to the vCenter instance.

## 9.3 Using ERSPAN to Monitor VMs

Use Remote SPAN (ERSPAN) to monitor VMs running on the ESX hosts within a VMware vCenter instance integrated with the DANZ Monitoring Fabric (DMF). ERSPAN monitors 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 affect network performance, especially when 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 DANZ Monitoring Fabric (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, the option exists to 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 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 and cannot be specified manually). All managed-service actions are supported.

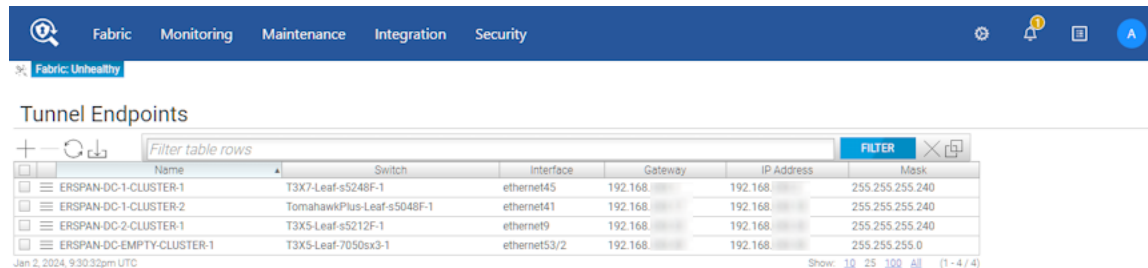
## 9.4 Defining a Tunnel Endpoint

Predefine the tunnel endpoints 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**



The screenshot shows the 'Tunnel Endpoints' page in the GUI. It features a navigation bar with 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. Below the navigation bar, there's a 'Tunnel Endpoints' section with a 'Filter table rows' input and a 'FILTER' button. The table lists four endpoints:

| Name                      | Switch                     | Interface    | Gateway  | IP Address | Mask            |
|---------------------------|----------------------------|--------------|----------|------------|-----------------|
| ERSPAN-DC-1-CLUSTER-1     | T3X7-Leaf-s5248F-1         | ethernet45   | 192.168. | 192.168.   | 255.255.255.240 |
| ERSPAN-DC-1-CLUSTER-2     | TomahawkPlus-Leaf-s5048F-1 | ethernet41   | 192.168. | 192.168.   | 255.255.255.240 |
| ERSPAN-DC-2-CLUSTER-1     | T3X5-Leaf-s5212F-1         | ethernet9    | 192.168. | 192.168.   | 255.255.255.240 |
| ERSPAN-DC-EMPTY-CLUSTER-1 | T3X5-Leaf-7050sx3-1        | ethernet53/2 | 192.168. | 192.168.   | 255.255.255.0   |

At the bottom right, it says 'Show: 10 25 100 All (1 - 4 / 4)'.

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 in the **Tunnel Endpoints** table.

**Figure 9-5: Create Tunnel Endpoint**



The screenshot shows the 'Create Tunnel Endpoint' form. It has the following fields:

- Name \***: A text input field.
- Switch**: A dropdown menu.
- Interface \***: A dropdown menu with a note: 'Optional: use to filter interface choices below' and 'Choose from a list of interfaces not already configured for DMF'.
- Gateway \***: A text input field with a placeholder '- IP Address -'.
- IP Address \***: A text input field with a placeholder '- IP Address -'.
- Mask \***: A text input field with a placeholder '- Subnet Mask -'.

At the bottom, there are three buttons: 'RESET', 'CANCEL', and 'SAVE'.

To create the tunnel endpoint, enter the following information and click **Save**:

- **Name**: Type a descriptive name for the endpoint.
- **Switch**: Select the DMF switch from the selection list for the configured endpoint interface.
- **Interface**: Select the interface from the selection list for the endpoint.
- **Gateway**: Type the address of the default gateway.
- **IP Address**: Type the endpoint IP address.
- **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**:

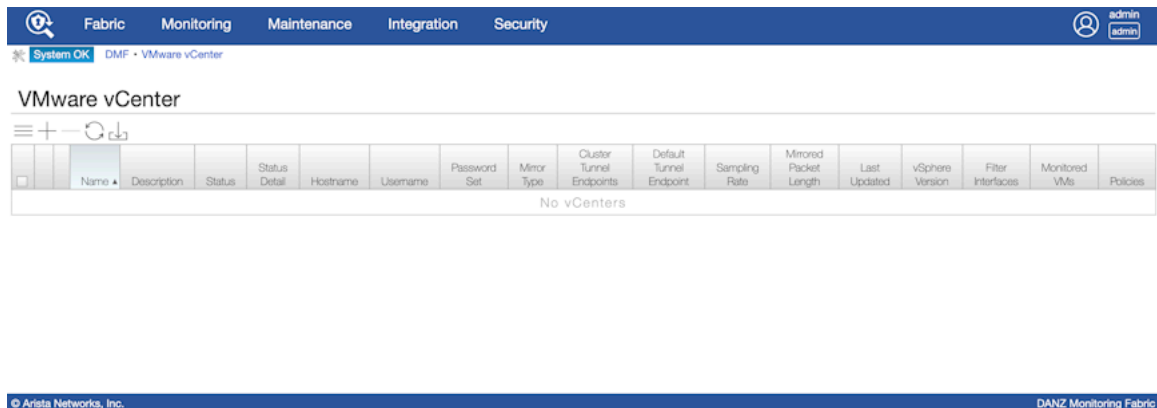
```
controller-1(config)# tunnel-endpoint ERSPAN switch CORE-SWITCH ethernet7 ip-  
address 172.27.1.1  
mask 255.255.255.0 gateway 172.27.1.2
```

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 Using the GUI to Integrate a vCenter Instance

To integrate a vCenter instance with DANZ Monitoring Fabric (DMF) to begin monitoring VMs, select **Integration > vCenter** from the **DMF** menu bar.

**Figure 9-6: Integration > vCenter**



This page displays information about the vCenter instances integrated with DMF. To add a vCenter instance for integration with DMF, perform the following steps:

1. Click the provision control (+) in the table.

**Figure 9-7: Create vCenter: Info**

**Create vCenter**

**Info** ✓

**Options** ✓

**VMs (0)** ✓

Name: my-vcenter

Description:

Host Name:   
Up to 255 characters in length

Username:   
Up to 255 characters in length

Password: - Password -   
☐ Show Password   
Up to 255 characters in length

BACK NEXT CANCEL SAVE

2. Type an alphanumeric identifier for the vCenter instance, and (optionally) add a description in the fields provided.
3. Identify the vCenter hostname 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 9-8: Create vCenter: Options (page 2)**

**Create vCenter**

**Info** ✓

**Options** ✓

**VMs (0)** ✓

Mirror Type: Encapsulated Remote

Cluster Tunnel Endpoints:   
+ -   
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:

Create Wildcard Tunnels:   
No Yes

BACK NEXT CANCEL SAVE

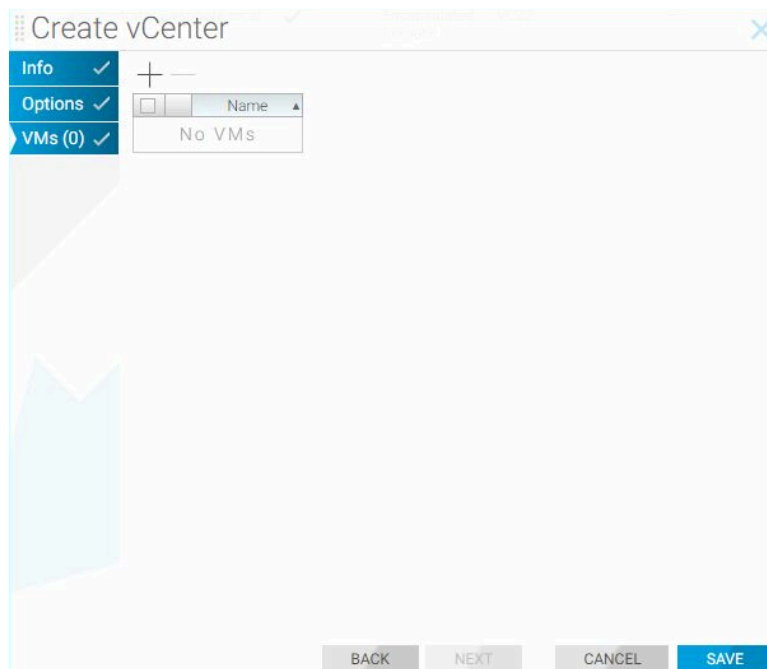
This page defines the mirror type as SPAN or ERSPAN. When selecting ERSPAN, the following additional fields complete the ERSPAN configuration:

- **Cluster Tunnel Endpoints** (optional)
- **Default Tunnel Endpoint** (required)
- **Sampling Rate** (optional)
- **Mirrored Packet Length** (optional)
- **Create Wildcard Tunnels** (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 configuring both cluster and default tunnel endpoints, 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.

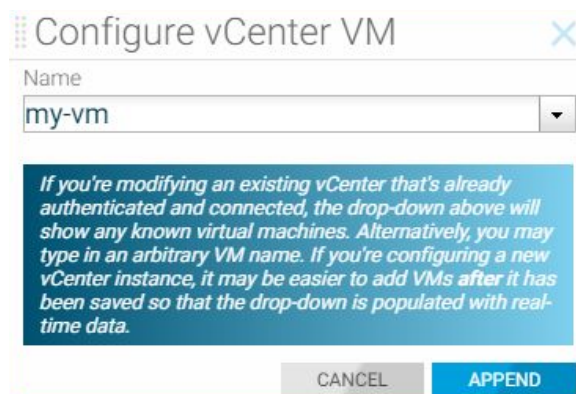
6. Click **Next**.

**Figure 9-9: Create vCenter/VMs**



7. To add a VM for monitoring, click the provision control (+).

**Figure 9-10: Configure vCenter VM**



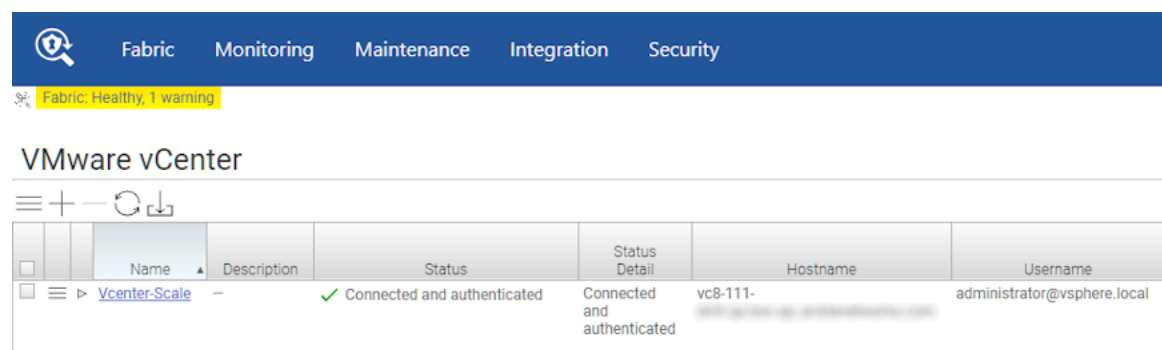


- Select VMs from the selection list after integrating vCenter and discovering the VMs, or manually add the VM hostname.
- After identifying the VM to monitor, click **Append**.
  - On the **VMs** of the **Create vCenter** dialog, click **Save**.

9.5.1 Using a vCenter Instance as the Traffic Source in a DMF Policy

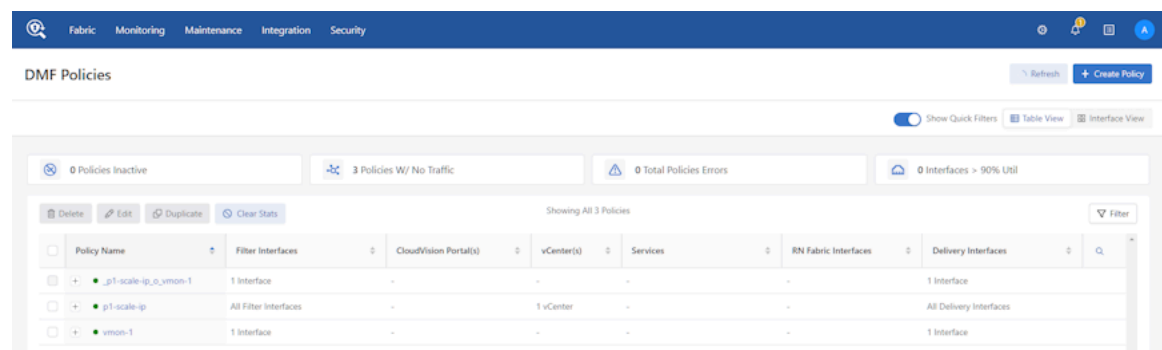
To identify a vCenter instance integrated with the DANZ Monitoring Fabric (DMF) Controller as the traffic source for a DMF policy, click the **VMware vCenter** tab on the **Integration** page. Locate the vCenter instance name.

Figure 9-11: VMware vCenter Name



Proceed to the **Monitoring > Policies** page.

Figure 9-12: DMF Policies



Click the **+ Create Policy** button to add a policy.

**Figure 9-13: Create Policy**

Create Policy

Policy Details ☐ Enable Advanced Options [Collapse](#)

| * Name                                   | Description                              | Action                               | Push VLAN ?                     | * Priority ?                     | Active                              |
|------------------------------------------|------------------------------------------|--------------------------------------|---------------------------------|----------------------------------|-------------------------------------|
| <input type="text" value="Input value"/> | <input type="text" value="Input value"/> | <input type="text" value="Forward"/> | <input type="text" value="##"/> | <input type="text" value="100"/> | <input checked="" type="checkbox"/> |

Port Selection Match Traffic Stat: Bit Rate

Traffic Sources [Create](#)  

+ Add Port(s)

Services  

+ Add Service(s)

Destination Tools [Create](#)  

+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

[Cancel](#) [Create Policy](#)

Enter a **Name** and **Description** for the vCenter policy. From the **Traffic Sources** column, select **+ Add Ports(s)**.

**Figure 9-14: Traffic Sources - Add Ports**

Create Policy

Policy Details ☐ Enable Advanced Options [Collapse](#)

| * Name                                  | Description                                   | Action                               | Push VLAN ?                     | * Priority ?                     | Active                              |
|-----------------------------------------|-----------------------------------------------|--------------------------------------|---------------------------------|----------------------------------|-------------------------------------|
| <input type="text" value="techpubs-1"/> | <input type="text" value="TechPubs vCenter"/> | <input type="text" value="Forward"/> | <input type="text" value="##"/> | <input type="text" value="100"/> | <input checked="" type="checkbox"/> |

Port Selection Match Traffic Stat: Bit Rate

Traffic Sources [Create](#)  

+ Add Port(s)

Services  

+ Add Service(s)

Destination Tools [Create](#)  

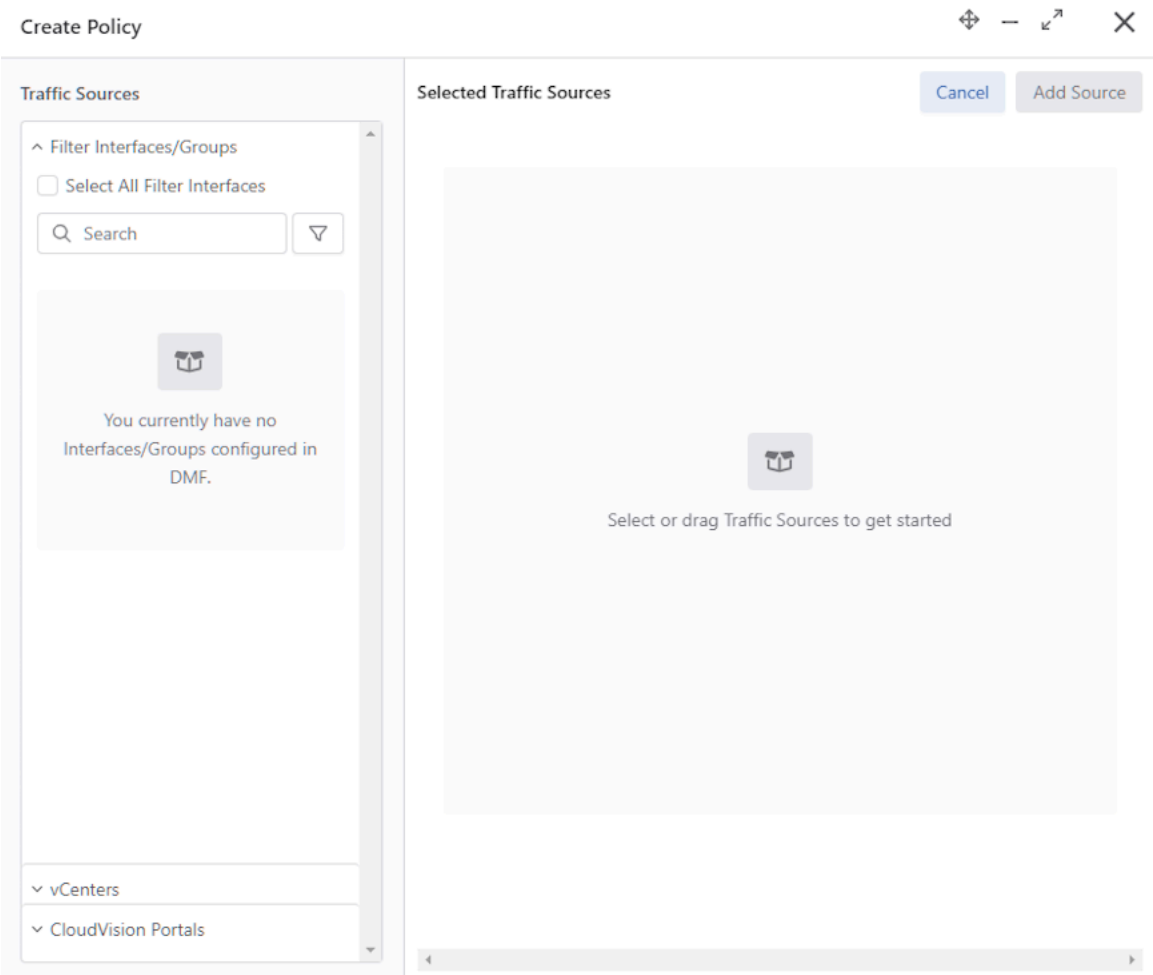
+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

[Cancel](#) [Create Policy](#)

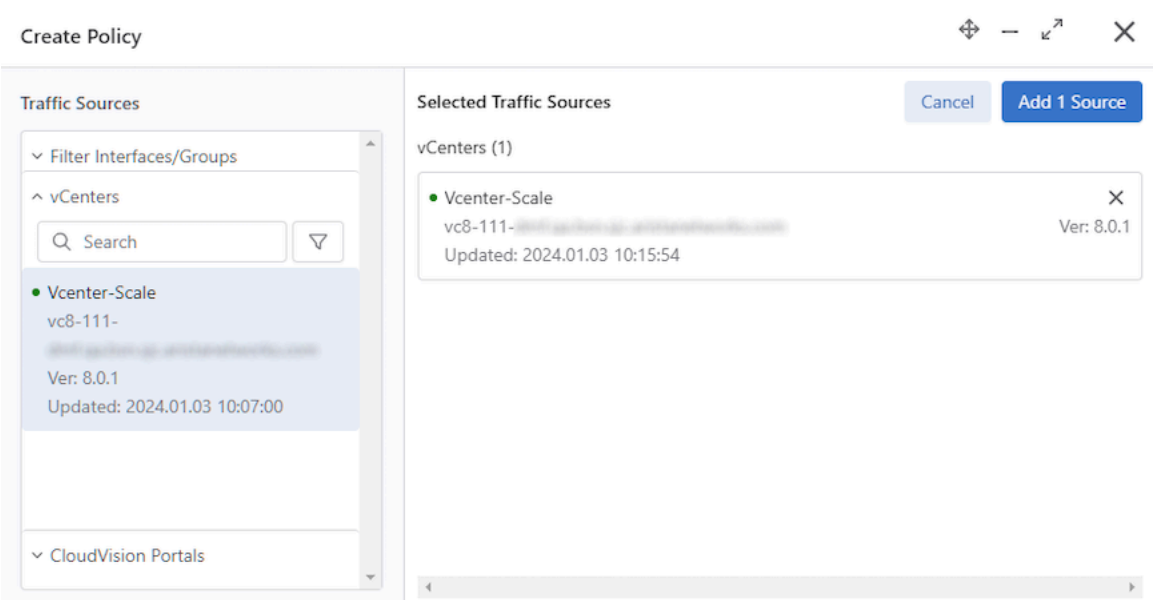
Click **vCenters**.

**Figure 9-15: vCenters**



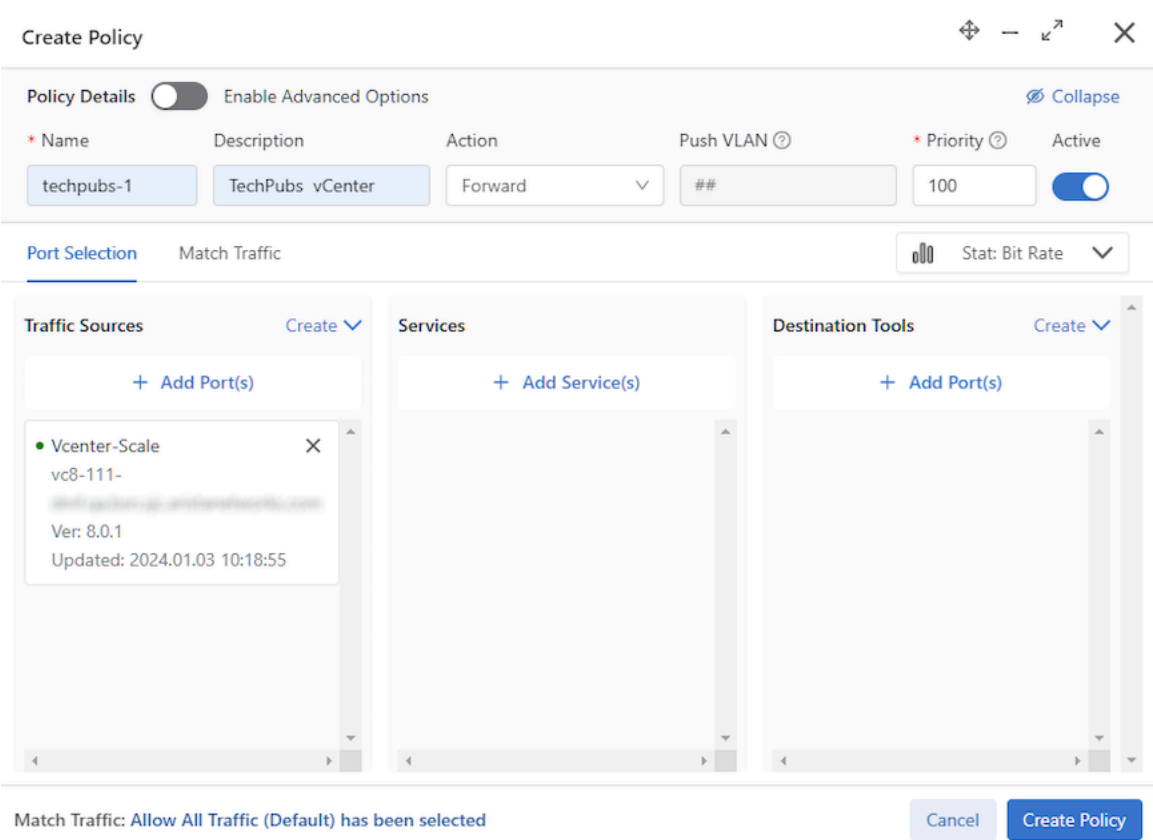
Available vCenter instances display. Select the required vCenter instance which then appears in the Selected traffic Sources panel.

Figure 9-16: vCenter Instance



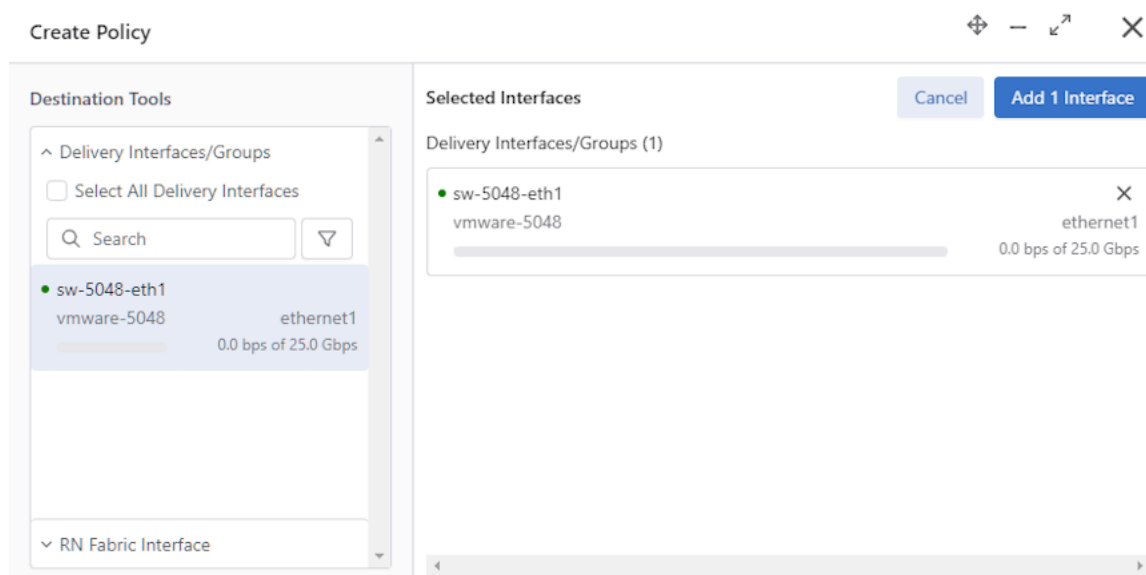
Click **Add 1 Source**. The vCenter instance appears in the **Traffic Sources** column.

Figure 9-17: vCenter Traffic Sources



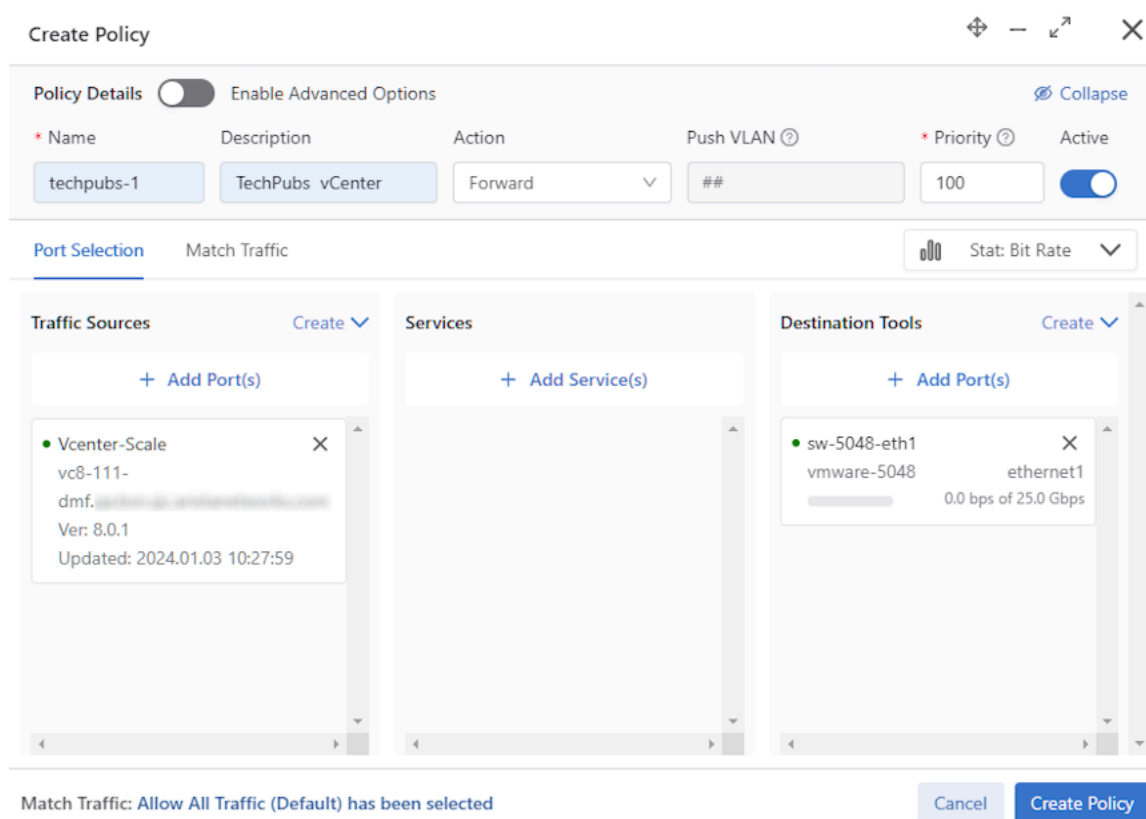
From the **Destination Tools** column, select **+ Add Ports(s)**. Select the interface under **Destination Tools**.

**Figure 9-18: Destination Tools - Add Ports**



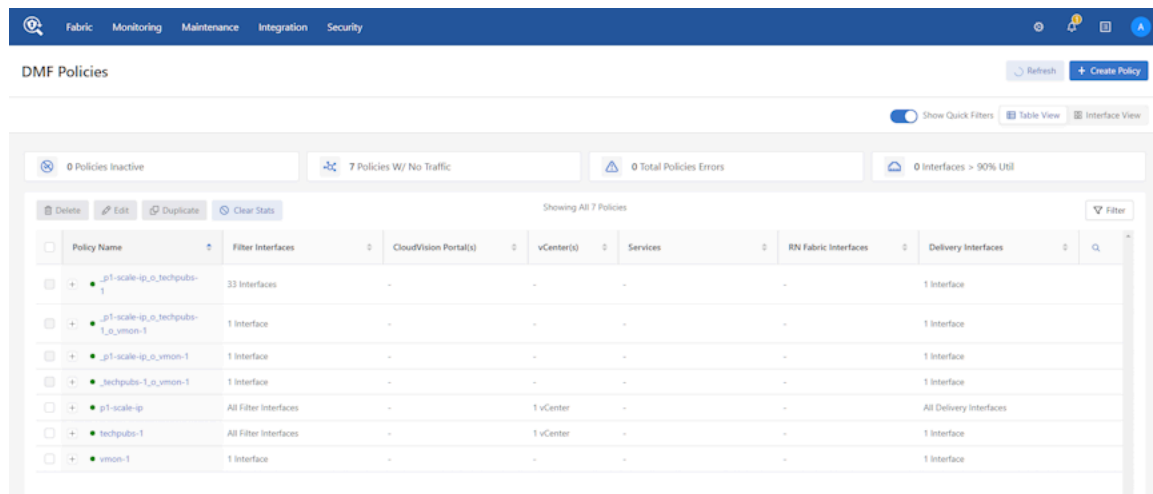
Click the **Add 1 Interface** button. The interface appears under the **Destination Tools** column.

**Figure 9-19: Add Interface**



Click **Create Policy**. The new vCenter policy appears in the DMF Policies dashboard.

**Figure 9-20: Create vCenter Policy**



## 9.6 Using the CLI to Integrate a vCenter Instance

Refer to the following topics to monitor VMs using Encapsulated Remote SPAN (ERSPAN) or Switch Port Analyzer (SPAN) on a locally connected vCenter instance and VMs on a second locally connected vCenter instance.

- [Monitoring VMs using ERSPAN on a Locally Connected vCenter Instance](#)
- [Monitoring VMs using SPAN on a Locally Connected vCenter Instance](#)
- [Monitoring VMs on a Second Locally Connected vCenter Instance](#)

### 9.6.1 VMs using ERSPAN on a Locally Connected vCenter Instance

To configure the DANZ Monitoring Fabric Controller for monitoring VMs using ERSPAN on a locally connected vCenter instance, perform 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 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. ERSPAN mirroring requires a tunnel endpoint configuration. 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 using both the **cluster** and **default-tunnel-endpoint** commands, all hosts in clusters form tunnels

using the cluster-specific configuration, and all the other hosts 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 a static route to the default or cluster tunnel-endpoint in each ESXI host.

```
esxcli network ip route ipv4 add -n <network> -g <gateway>
Example: esxcli network ip route ipv4 add -n 192.168.200.0/24 -g
192.168.150.1
```

5. Add the VMs 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
```

6. Receive-only GRE tunnel-interfaces will be auto-configured under **switch** for all the hosts belonging to **vc-1** that have a route to the default or cluster tunnel-endpoint.

```
! switch
switch DMF-RU34
mac 94:8e:d3:fd:6b:96
!
gre-tunnel-interface vcenter-abd08a18
direction receive-only
local-ip 192.168.200.254 mask 255.255.255.0 gateway-ip 192.168.200.1
origination vc8--interface
parent-interface ethernet55
remote-ip 192.168.150.27
gre-key-decap 33554432
!
gre-tunnel-interface vcenter-abd08a37
direction receive-only
local-ip 192.168.200.254 mask 255.255.255.0 gateway-ip 192.168.200.1
origination vc8--interface
parent-interface ethernet55
remote-ip 192.168.150.28
gre-key-decap 33554432
!
gre-tunnel-interface vcenter-abd08a56
direction receive-only
local-ip 192.168.200.254 mask 255.255.255.0 gateway-ip 192.168.200.1
origination vc8--interface
parent-interface ethernet55
remote-ip 192.168.50.29
gre-key-decap 33554432
```

7. Enter the **show running-config vcenter** command to view the vCenter configuration.

```
controller-1# show running-config vcenter
! vcenter
vcenter vc-1
hashed-password 752a3a3211040e0200090409090611
host-name 10.8.23.70
mirror-type erspan
mirrored-packet-length 60
sampling-rate 60
```

```
user-name administrator@vsphere.local
!  
vm-monitoring  
vm vm-2001  
vm vm-2002
```

8. 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)# filter-vcenter vc-1  
controller-1(config-policy)# 1 match any  
controller-1(config-policy)# delivery-interface TOOL-PORT-03
```

9. Enter the **show running-config policy** command to view the automatically assigned filter interfaces.

```
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 DMF-RU34-filter-vcenter-abd08a18 vc-1--interface  
filter-interface DMF-RU34-filter-vcenter-abd08a37 vc-1--interface  
filter-interface DMF-RU34-filter-vcenter-abd08a56 vc-1--interface  
filter-vcenter vc-1  
1 match any
```

All the host tunnels belonging to **vc-1** will become the filter interfaces. If new hosts are added, deleted, or modified, policies will be recomputed with the new interfaces.

## 9.6.2 VMs using SPAN on a Locally Connected vCenter Instance

To configure the DANZ Monitoring Fabric Controller for monitoring VMs using SPAN on a locally connected vCenter instance, perform 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  
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. Add the VMs 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
```

4. 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
```

5. 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)# filter-vcenter vc-1
controller-1(config-policy)# 1 match any
controller-1(config-policy)# delivery-interface TOOL-PORT-03
```



**Note:** LLDP automatically learns the filter interfaces. All the hosts belonging to **vc-1** that have physical connections to DMF switches become the filter interfaces. If new connections are made later (or existing connections are changed), policies will be recomputed with the new interfaces.

6. To view the automatically assigned filter interfaces, enter the **show running-config policy** command.

```
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
```

### 9.6.3 VMs on a Second Locally Connected vCenter Instance

To configure the DMF Controller for monitoring VMs on a second locally connected vCenter instance, perform the following steps:

1. Add the VMs 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 | erspan
(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
```

2. Configure the policy for the second vCenter instance.

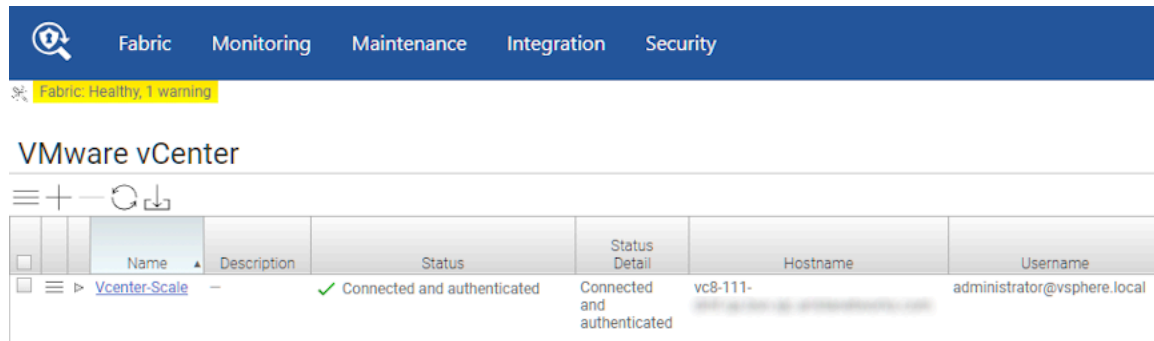
```
(config)# policy dmf-policy-with-vcenter-2
(config-policy)# vcenter vc-2
(config-policy)# 1 match any
```

```
(config-policy)# delivery-interface TOOL-PORT-02
```

## 9.7 Using the GUI to View vCenter Configuration

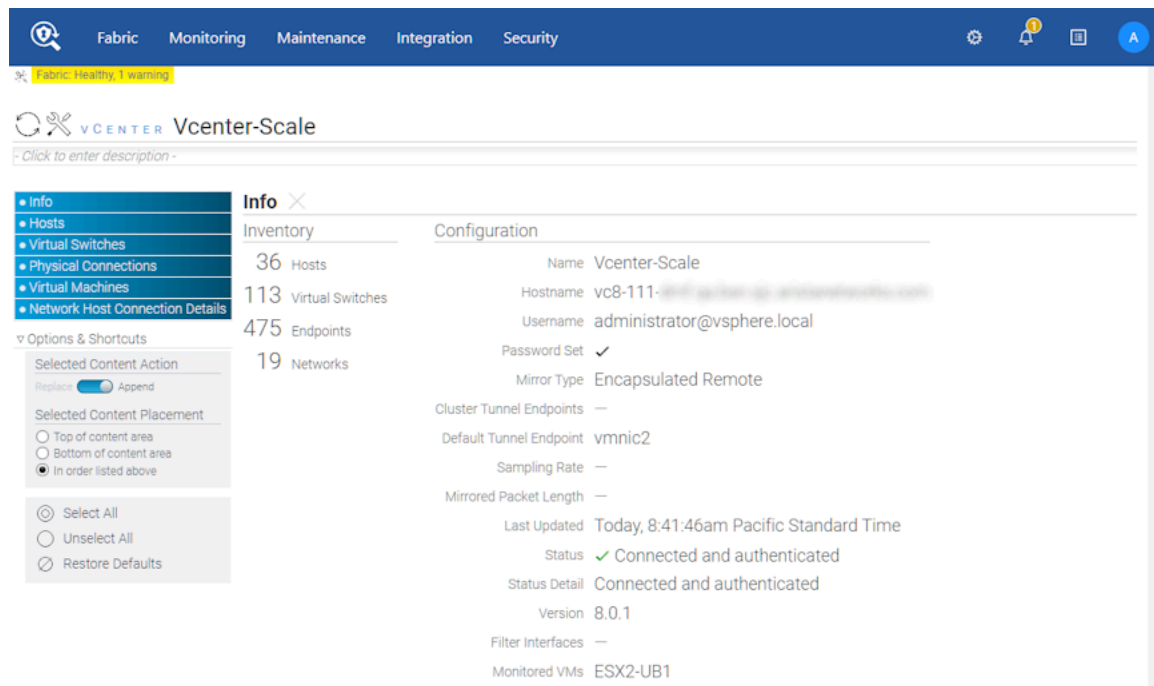
After integrating a vCenter instance, click the link in the **Name** column in the **vCenter** table to view vCenter activity.

**Figure 9-21: VMware vCenter Instance Name**



DANZ Monitoring Fabric (DMF) displays the **vCenter Info** page.

**Figure 9-22: VMware vCenter Configuration**



The **Info** 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 9-23: Hosts, Virtual Switches, and Physical Connections

**Hosts**

| Inventory Host Name | Host DNS Name | Product Name                    | Hardware Model | CPU usage (%) | Memory usage (%) | Virtual Switches | State Summary        | State Description            |
|---------------------|---------------|---------------------------------|----------------|---------------|------------------|------------------|----------------------|------------------------------|
| 10.9.30.41          | localhost     | VMware ESXi 6.7.0 build-6169022 | PowerEdge R320 | 75            | 90               | 2                | [duplicate-dns-name] | [WARNING] duplicate-dns-name |
| 10.9.30.60          | app-anim      | VMware ESXi 6.5.0 build-6294253 | PowerEdge R210 | 0             | 33               | 2                | ---                  | ---                          |
| 10.9.30.61          | localhost     | VMware ESXi 6.7.0 build-6214904 | ---            | 4             | 13               | 2                | [duplicate-dns-name] | [WARNING] duplicate-dns-name |

Oct 24, 2018, 3:11:20am UTC

**Virtual Switches**

| Name               | Inventory Host Name | Host DNS Name | Status                                                                                     | NICs | Link Discovery Protocol | Link Discovery Direction | Virtual Switch Type |
|--------------------|---------------------|---------------|--------------------------------------------------------------------------------------------|------|-------------------------|--------------------------|---------------------|
| VC22-V55-Dswitch-1 | 10.9.30.60          | app-anim      | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected                     | 2    | LLDP                    | both                     | dvb                 |
| VC22-V57-Dswitch-1 | 10.9.30.61          | localhost     | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected                     | 2    | LLDP                    | both                     | dvb                 |
| VC22-V57-Dswitch-1 | 10.9.30.41          | localhost     | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected                     | 2    | LLDP                    | both                     | dvb                 |
| vswitch0           | 10.9.30.60          | app-anim      | [WARNING] vmnic0 is not connected<br>[WARNING] Link Discovery advertisement not configured | 1    | CDP                     | listen                   | vss                 |
| vswitch0           | 10.9.30.61          | localhost     | [WARNING] vmnic0 is not connected<br>[WARNING] Link Discovery advertisement not configured | 1    | CDP                     | listen                   | vss                 |
| vswitch0           | 10.9.30.41          | localhost     | [WARNING] vmnic0 is not connected<br>[WARNING] Link Discovery advertisement not configured | 1    | CDP                     | listen                   | vss                 |

Oct 24, 2018, 3:11:20am UTC

**Physical Connections**

| Inventory Host Name | Host DNS Name      | Virtual Switch     | NIC Team Group | NIC Name | Connected Interface | Status      |
|---------------------|--------------------|--------------------|----------------|----------|---------------------|-------------|
| 10.9.30.41          | localhost          | VC22-V57-Dswitch-1 | ---            | vmnic3   | ---                 | ✓ Connected |
| 10.9.30.41          | localhost          | VC22-V57-Dswitch-1 | ---            | vmnic2   | ---                 | ✓ Connected |
| 10.9.30.41          | localhost          | vswitch0           | ---            | vmnic0   | ---                 | ✓ Connected |
| 10.9.30.60          | app-animesh-esxi-4 | vswitch0           | ---            | vmnic0   | ---                 | ✓ Connected |

Figure 9-24: Virtual Machines and Network Host Connection Details

**Virtual Machines**

| VM Name           | VM Interface Name | Endpoint Name | Power State  | Monitored in Big Top | Tunnel Endpoint | Status                                                                                                                                                                                 | UPLN     | Inventory Host Name | Physical Connection        | Virtual Switch Type | Endpoint Type   | MAC Address | IP Address | Default Gateway | NIC Used (Balanced Policy) |
|-------------------|-------------------|---------------|--------------|----------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|---------------------|----------------------------|---------------------|-----------------|-------------|------------|-----------------|----------------------------|
| BM-50-Analysis    | Network adapter 1 | localhost     | ✓ Powered On | ---                  | ---             | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] VM Adapter not connected<br>[WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected | 2000     | 10.9.30.41          | vmnic3<br>vmnic2           | dvb                 | virtual-machine | ---         | ---        | ---             | IP4only                    |
| 18-vc22-v57-404-1 | Network adapter 2 | app-anim      | ✓ Powered On | ---                  | ---             | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] Link Discovery advertisement not configured                                                        | Untagged | 10.9.30.60          | vmnic0<br>vmnic3<br>vmnic2 | dvb                 | virtual-machine | ---         | ---        | ---             | Originating Port ID        |
| 18-vc22-v57-404-1 | Network adapter 2 | localhost     | ✓ Powered On | ✓                    | VC22            | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] Link Discovery advertisement not configured                                                        | Untagged | 10.9.30.41          | vmnic3<br>vmnic2           | dvb                 | virtual-machine | ---         | ---        | ---             | Originating Port ID        |
| 18-vc22-v57-404-2 | Network adapter 2 | localhost     | ✓ Powered On | ✓                    | VC22            | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] Link Discovery advertisement not configured                                                        | Untagged | 10.9.30.61          | vmnic0<br>vmnic3<br>vmnic2 | dvb                 | virtual-machine | ---         | ---        | ---             | Originating Port ID        |
| 18-vc22-v57-404-1 | Network adapter 2 | localhost     | ✓ Powered On | ✓                    | VC22            | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] Link Discovery advertisement not configured                                                        | Untagged | 10.9.30.41          | vmnic3<br>vmnic2           | dvb                 | virtual-machine | ---         | ---        | ---             | IP4only                    |
| 18-vc22-v57-404-2 | Network adapter 2 | localhost     | ✓ Powered On | ✓                    | VC22            | [WARNING] vmnic3 is not connected<br>[WARNING] vmnic2 is not connected<br>[WARNING] Link Discovery advertisement not configured                                                        | Untagged | 10.9.30.61          | vmnic3<br>vmnic2           | dvb                 | virtual-machine | ---         | ---        | ---             | IP4only                    |

Oct 24, 2018, 3:11:20am UTC

**Network Host Connection Details**

| UPLN     | Inventory Host Name                | Virtual Switch     | VM Guest | NIC    | Connection Status | Switch | Endpoint |
|----------|------------------------------------|--------------------|----------|--------|-------------------|--------|----------|
| Untagged | 10.9.30.41, 10.9.30.41             | VC22-V57-Dswitch-1 | 2        | vmnic3 | ---               | ---    | ---      |
| Untagged | 10.9.30.61, 10.9.30.41             | VC22-V57-Dswitch-1 | 2        | vmnic2 | ---               | ---    | ---      |
| Untagged | 10.9.30.60                         | VC22-V55-Dswitch-1 | 1        | vmnic3 | ---               | ---    | ---      |
| Untagged | 10.9.30.60                         | VC22-V55-Dswitch-1 | 1        | vmnic2 | ---               | ---    | ---      |
| Untagged | 10.9.30.41, 10.9.30.60, 10.9.30.61 | vswitch0           | 3        | vmnic0 | ---               | ---    | ---      |
| Untagged | 10.9.30.41, 10.9.30.61             | VC22-V57-Dswitch-1 | 0        | vmnic3 | ---               | ---    | ---      |
| Untagged | 10.9.30.41, 10.9.30.61             | VC22-V57-Dswitch-1 | 0        | vmnic2 | ---               | ---    | ---      |

Oct 24, 2018, 3:11:20am UTC

## 9.8 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#

```

## 9.9 Integrating vCenter with DMF using Mirror Stack

DANZ Monitoring Fabric (DMF) vCenter integration supports mirroring from vCenter hosts using the default TCP/IP stack. However, this can result in traffic drops and affect production traffic since mirror traffic can conflict with production traffic. DMF vCenter integration with Mirror Stack provides the functionality to use the mirror TCP/IP stack for mirror sessions. Mirror stack in the ESXi host allows decoupling the traffic and keeps the production traffic unaffected.

vCenter configurations in DMF will use a mirror stack by default; however, if upgrading from previous DMF versions, the already configured vCenter will be set to use the default TCP/IP stack.

### Platform Compatibility

vCenter integration with Mirror Stack requires an **extra** NIC on the ESXi host with following versions:

- vCenter Server 7.0.x
- vCenter Server 8.0.x

### 9.9.1 vCenter Configuration

DMF vCenter integration with Mirror Stack requires a mirror stack configuration on the ESXi host and vCenter. Perform the following steps to configure the mirror stack on vCenter.

Repeat the steps for each ESXi host containing VMs to be monitored.

1. Enable the mirror stack in the ESXi host if not already enabled.
  - a. Use the **esxcli network ip netstack list** command to review the current network stacks.

```

[root@ESX33:~] esxcli network ip netstack list
defaultTcpipStack
  Key: defaultTcpipStack
  Name: defaultTcpipStack
  State: 4660

mirror
  Key: mirror

```

Name: mirror  
State: 4660

To view the TCP/IP configuration from vCenter UI, navigate to **Host > Configure > TCP/IP**.

**Figure 9-25: TCP/IP Configuration**

| TCP/IP Configuration |              |                |                 |                      |                      |                      |  |
|----------------------|--------------|----------------|-----------------|----------------------|----------------------|----------------------|--|
| Edit...              |              |                |                 |                      |                      |                      |  |
| TCP/IP Stack         | Type         | VMkernel Ad... | IPv4 Gateway... | IPv6 Gateway Address | Preferred DNS server | Alternate DNS server |  |
| Default              | System stack | 1              | 10.240.166.1    | --                   | 10.240.48.6          | 10.242.35.6          |  |
| Provisioning         | System stack | 0              | --              | --                   | --                   | --                   |  |
| vMotion              | System stack | 1              | --              | --                   | --                   | --                   |  |
| mirror               | Custom stack | 1              | --              | --                   | --                   | --                   |  |

- b. If the mirror stack is not configured, use the `esxcli network ip netstack set -N mirror` command to enable it.



**Note:** The `mirror` setting is required to enable the Mirror TCP/IP stack and DMF integration.

2. From vCenter create a VMkernel adapter with the mirror stack.

**Figure 9-26: VMkernel Network Adapter**

10.240.166.30 - Add Networking ×

1 Select connection type

2 Select target device

3 Port properties

4 IPv4 settings

5 Ready to complete

Select connection type

Select a connection type to create.

☒ VMkernel Network Adapter

The VMkernel TCP/IP stack handles traffic for ESXi services such as vSphere vMotion, iSCSI, NFS, FCoE, Fault Tolerance, vSAN, host management and etc.

☐ Virtual Machine Port Group for a Standard Switch

A port group handles the virtual machine traffic on standard switch.

☐ Physical Network Adapter

A physical network adapter handles the network traffic to other hosts on the network.

CANCEL

BACK

NEXT

Select the appropriate network using the **Browse** option.

**Figure 9-27: Browse**

10.240.166.30 - Add Networking ×

✓ 1 Select connection type

**2 Select target device**

3 Port properties

4 IPv4 settings

5 Ready to complete

Select target device

Select a target device for the new connection.

☒ Select an existing network

☐ Select an existing standard switch

☐ New standard switch

MTU (Bytes)

1500

BROWSE ...

BROWSE ...

CANCEL

BACK

NEXT

Click **Next** and select **Port properties** and choose **mirror**.

Figure 9-28: Port Properties - Mirror

10.240.166.30 - Add Networking ×

✓ 1 Select connection type

✓ 2 Select target device

**3 Port properties**

4 IPv4 settings

5 Ready to complete

Port properties

Specify VMkernel port settings.

VMkernel port settings

Network label4004-mirror-stack (DMF-70-Mirror-stack)

IP settings

MTU

TCP/IP stack

Available services

Enabled services

Default

Provisioning

vMotion

**✓ mirror**

switch

9000

☐ vMotion

☐ Provisioning

☐ Fault Tolerance logging

☐ Management

☐ vSphere Replication

☐ vSphere Replication NFC

☐ vSAN

☐ vSphere Backup NFC

CANCEL

BACK

NEXT

Figure 9-29: Mirror Stack Added

10.240.166.30 - Add Networking ×

✓ 1 Select connection type

✓ 2 Select target device

**✓ 3 Port properties**

✓ 4 IPv4 settings

5 Ready to complete

Port properties

Specify VMkernel port settings.

VMkernel port settings

Network label4004-mirror-stack (DMF-70-Mirror-stack)

IP settingsIPv4

MTUGet MTU from switch9000

TCP/IP stackmirror

Available services

Enabled services

☐ vMotion

☐ Provisioning

☐ Fault Tolerance logging

☐ Management

☐ vSphere Replication

☐ vSphere Replication NFC

☐ vSAN

☐ vSphere Backup NFC

Add the **IPv4 address** and the **Default gateway** address according to your local network requirements.

**Figure 9-30: IP Address and Gateway Address**

10.240.166.30 - Add Networking ×

✓ 1 Select connection type

✓ 2 Select target device

✓ 3 Port properties

**4 IPv4 settings**

5 Ready to complete

IPv4 settings

Specify VMkernel IPv4 settings.

☐ Obtain IPv4 settings automatically
   
☒ Use static IPv4 settings

IPv4 address

192.168.150.30

Subnet mask

255.255.255.0

Default gateway

☒ Override default gateway for this adapter
   
 192.168.150.1

CANCEL

BACK

NEXT

Click **Next**.

**Figure 9-31: VMkernel Adapters**

VMkernel adapters

[ADD NETWORKING...](#) [REFRESH](#)

|      | Device | Network Label     | Switch              | IP Address     | TCP/IP Stack | Enabled Services |
|------|--------|-------------------|---------------------|----------------|--------------|------------------|
| ⋮ >> | vmk0   | untagged-mgmt     | DVS-MGMT            | 10.240.166.30  | Default      | --               |
| ⋮ >> | vmk1   | untagged-mgmt     | DMF-70-Mirror-stack | 10.240.166.211 | Default      | Management       |
| ⋮ >> | vmk2   | 4004-mirror-stack | DMF-70-Mirror-stack | 192.168.150.30 | mirror       | --               |

- Based on the networking requirements, configure the default gateway of the mirror stack in the host's TCP/IP configuration or a static route entry in the ESXi host to the DMF tunnel endpoint. The following example illustrates adding a static route entry to the DMF tunnel endpoint.

```
[root@ESX33:~] esxcli network ip route ipv4 add -n 192.168.200.0/24 -g 192.168.150.1 -N mirror
```

```
[root@ESX31:~] esxcli network ip route ipv4 list -N mirror
```

| Network       | Netmask       | Gateway       | Interface | Source |
|---------------|---------------|---------------|-----------|--------|
| 192.168.150.0 | 255.255.255.0 | 0.0.0.0       | vmk2      | MANUAL |
| 192.168.200.0 | 255.255.255.0 | 192.168.150.1 | vmk2      | MANUAL |



4. Navigate to **Configure > TCP/IP Configuration > Select mirror stack > IPv4 Routing Table** to view the routes.

**Figure 9-32: TCP/IP Configuration & IPv4 Routing Table**

TCP/IP Configuration Edit...

| TCP/IP Stack  | Type                | VMkernel A... | IPv4 Gatew... | IPv6 Gateway Address | Preferred DNS server | Alterna   |
|---------------|---------------------|---------------|---------------|----------------------|----------------------|-----------|
| Monitoring    | System stack        | 0             | --            | --                   | --                   | --        |
| <b>mirror</b> | <b>Custom stack</b> | <b>1</b>      | <b>--</b>     | <b>--</b>            | <b>--</b>            | <b>--</b> |

TCP/IP Stack: mirror

Routing **IPv4 Routing Table** IPv6 Routing Table Advanced

| Network Address | Subnet Mask   | Gateway       | Device |
|-----------------|---------------|---------------|--------|
| 192.168.150.0   | 255.255.255.0 | Local Subnet  | vmk2   |
| 192.168.200.0   | 255.255.255.0 | 192.168.150.1 | vmk2   |

**Figure 9-33: Virtual Switch**

Virtual switches ADD NETWORKING... RI

▼ Distributed Switch: DMF-70-Mirror-stack | [MANAGE PHYSICAL ADAPTERS](#) ...

4004-mirror-stack ...

VLAN ID: 4004

▼ VMkernel Ports (1)

vmk2 : 192.168.150.30 ...

Virtual Machines (0)

untagged-mgmt ...

VLAN ID: --

> VMkernel Ports (1)

Virtual Machines (0)

VLAN31 ...

VLAN ID: 31

> Virtual Machines (1)

DMF-70-Mirror-st-DVUplinks-2... ...

▼ monitoring-tunnel-traffic (1 NIC A...

vmnic3 10.240.166.30 ...

> vm-production-traffic (1 NIC Ada...

## 9.9.2 Configuring DMF

### Using the CLI

From the DMF Controller configure the TCP/IP stack using the `tcp-ip-stack` option in the vCenter config. The default and recommended value is `mirror-stack`.

```
dmf-controller-1(conf)# vcenter vc8
dmf-controller-1(config-vcenter)# tcp-ip-stack
default-stack mirror-stack
dmf-controller-1(config-vcenter)# tcp-ip-stack mirror-stack
```

## Using the GUI

To configure TCP/IP Stack, navigate to **Integration > VMware vCenter**. While adding or editing a vCenter configuration, select the appropriate choice using **TCP/IP Stack**. **Default Stack** and **Mirror Stack** are the options.

**Figure 9-34: Create vCenter TCP/IP Stack**

The screenshot shows the 'Create vCenter' configuration window. On the left, there are tabs for 'Info', 'Options' (checked), and 'VMs (0)' (checked). The main area is titled 'TCP/IP Stack' and contains several configuration fields: a dropdown menu for 'TCP/IP Stack' (highlighted with an orange box), a dropdown for 'Mirror Type' (set to '- None -'), a 'Sampling Rate' field (set to 'packet/s'), a 'Mirrored Packet Length' field, and a 'Create Wildcard Tunnels' toggle switch (set to 'Yes'). At the bottom, there are four buttons: 'BACK', 'NEXT' (highlighted in blue), 'CANCEL', and 'SAVE'.



**Attention:** Encapsulated Remote mirroring with Default Stack is **not** recommended. Use **Mirror Stack** for optimal performance.

## Show Commands

Use the **show running-config** command to view the **tcp-ip-stack** configuration.



**Note:** If **mirror-stack** is configured, it will only show when using the **details** token.

```
dmf-controller-1(config-vcenter)# show running-config vcenter vc8 details

! vcenter
vcenter vc8
  default-tunnel-endpoint r34-lag-leaf1b
  hashed-password <hashed-password>
  host-name <ip-address>
  mirror-type encapsulated-remote
  tcp-ip-stack mirror-stack
  user-name administrator@vsphere.local
```

View the existing mirror stack NICs and IPs of the host using the **show vcenter vCenter name inventory** command.



**Note:** **vc8** is an example vCenter name.

```
dmf-controller-1# show vcenter vc8 inventory
# vCenter ESXi Host      Host DNS Name      Cluster
Product Name      Hardware Model CPU Usage (%) Memory Usage (%)
Virtual switches Mirror Stack VMkernel Adapter VMkernel Adapter IP Address
-|-----|-----|-----|-----|-----|-----|
---|-----|-----|-----|-----|-----|
1 vc8      10.240.166.27 ESX27.qa.bsn.sjc.aristanetworks.com BSN-NSX-1
  VMware ESXi 8.0.2 build-22380479 PowerEdge R430 2      15
  3          vmk1          192.168.60.27
2 vc8      10.240.166.28 ESX28.qa.bsn.sjc.aristanetworks.com BSN-NSX-2
  VMware ESXi 8.0.2 build-22380479          0      4
  4
3 vc8      10.240.166.29 ESX29.qa.bsn.sjc.aristanetworks.com Edge
  VMware ESXi 8.0.0 build-20513097 PowerEdge R430 4      23
  3
4 vc8      10.240.166.33 ESX33.qa.bsn.sjc.aristanetworks.com vc8-mixed-stack
  VMware ESXi 8.0.2 build-22380479          0      6
  3          vmk1          192.168.60.33
5 vc8      10.240.166.35 ESX35.qa.bsn.sjc.aristanetworks.com MGMT
  VMware ESXi 7.0.2 build-17867351 PowerEdge R430 26     23
  2
6 vc8      10.240.166.38 ESX38.qa.bsn.sjc.aristanetworks.com vc8-mixed-stack
  VMware ESXi 8.0.2 build-22380479          1      23
  3          vmk1          192.168.60.38
dmf-rack#
```

### 9.9.3 Troubleshooting

Use the **show fabric errors** and **show fabric warnings** commands to troubleshoot and verify that everything is functioning as expected.

In the following example, the error message indicates that DMF could not find a route from the ESXi host to the DMF tunnel endpoint.

```
dmf-controller-1# show fabric errors
~~~~~
vCenter related error ~~~~~
~~~~~
# vCenter Name Error
--|-----|-----|-----|-----|
-----|
1 vc701      Unable to locate a matching route for Mirror TCP/IP stack in
  host ESX37.qa.bsn.sjc.aristanetworks.com for DMF endpoint 192.168.200.254
```

### 9.9.4 Limitations

- A port mirroring session remains on the original distributed virtual switch (DVS) when a VM migrates between DVSSs.
- Port mirroring sessions will persist on the DVS if a VM is renamed in vCenter while being monitored by DMF.

- 
- DMF cannot create a port mirroring session in the DVS if a conflicting session with the same VM exists in the DVS. This is not a limitation in vCenter 7.
  - When using mirror stack configuration in DMF, mirror sessions may still be created on the DVS for the ESXi host that doesn't have a mirror stack configuration. This will result in no traffic being mirrored from the VM.
  - Auto-generated filter interfaces by vCenter integration should not be deleted from the policy. If they are deleted manually from the policy, they will not be automatically re-added.
  - DMF cannot monitor VMkernel adapters.

### 9.9.5 Resources

- Create a Port Mirroring Session - <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-networking/GUID-68B5DD45-DD3F-4E9B-A6CD-BE97026A846A.html>
- Encapsulated Remote Mirroring (L3) Source - <https://kb.vmware.com/s/article/67973>

## 9.10 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.

### 9.10.1 Configuration

Configure wildcard tunnels using the CLI or the GUI.

### 9.10.2 Using the CLI to Create Wildcard Tunnels

The CLI construct `wildcard-tunnels` is available as a configuration option when configuring a VMware vCenter in DANZ Monitoring Fabric (DMF), as shown below:

**Table 5: Commands**

|                                |                                                 |
|--------------------------------|-------------------------------------------------|
| <b>cluster</b>                 | Configure tunnel-endpoint for cluster           |
| <b>default-tunnel-endpoint</b> | Configure tunnel endpoints                      |
| <b>description</b>             | Describe this vCenter                           |
| <b>hashed-password</b>         | Set the vCenter password (to log into vCenter)  |
| <b>host-name</b>               | Set the vCenter hostname                        |
| <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.

```
dmf-controller-1(config)# vcenter VC1
dmf-controller-1(config-vcenter)# wildcard-tunnels
dmf-controller-1(config-vcenter)# show this
! vcenter
vcenter VC1
wildcard-tunnels
dmf-controller-1(config-vcenter)#
```

Similarly, disable wildcard tunnels by issuing the **no** command as shown below:

```
dmf-controller-1(config-vcenter)# show this
! vcenter
vcenter VC1
wildcard-tunnels
dmf-controller-1(config-vcenter)# no wildcard-tunnels
dmf-controller-1(config-vcenter)# show this
! vcenter
vcenter VC1
dmf-controller-1(config-vcenter)#
```

### Show Commands

There is no specific show command for wildcard tunnels; however, check them in the vCenter running config. In addition, the **show tunnels** command shows the tunnels created for the selected vCenter configuration with a wildcard remote IP address.

### Troubleshooting

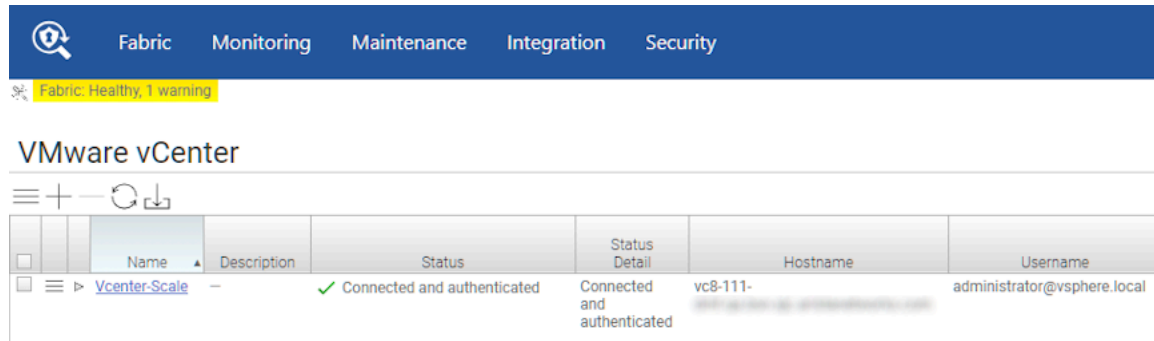
Verify errors and warnings are clear using the **show fabric errors** and **show fabric warnings** commands. The **show tunnels** command displays tunnels created based on the vCenter configuration on the Controller with a wildcard remote IP address. Use the **show switch <name> table gre-tunnel** command to display tunnels programmed on the switch.

### 9.10.3 Using the GUI to Create Wildcard Tunnels

Use the DANZ Monitoring Fabric (DMF) GUI to create wildcard tunnels as outlined below.

Navigate to the **Integration > VMware vCenter** page.

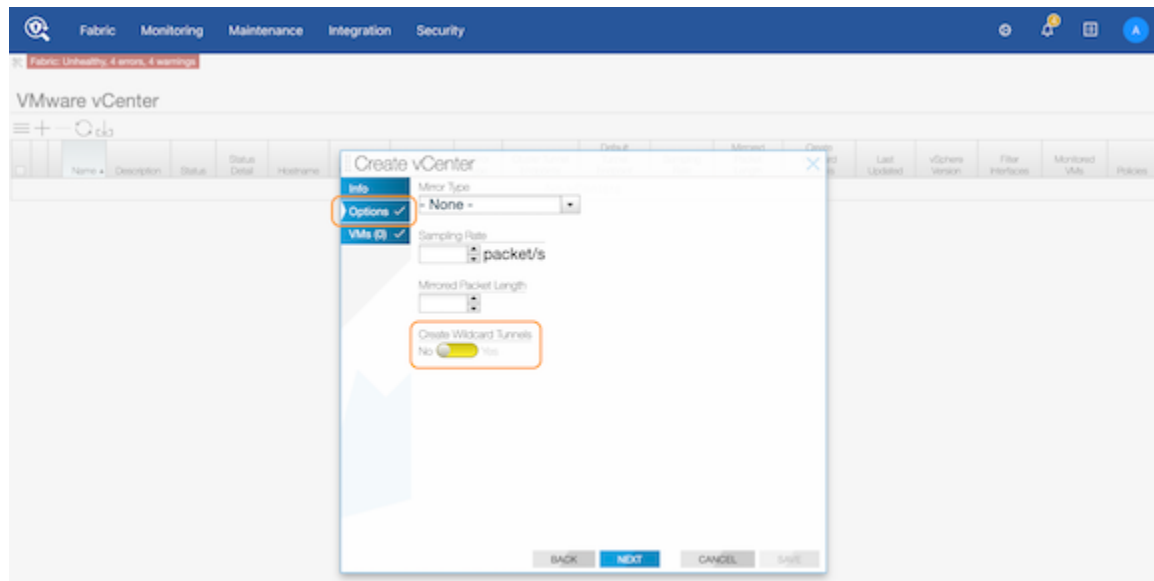
**Figure 9-35: VMware vCenter Add/Edit**



Click the **Menu** icon.

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.

**Figure 9-36: VMware vCenter Create vCenter Options**



### 9.10.4 Limitations

Select Broadcom® switch ASICs support wildcard tunnels; ensure your switch model supports this feature before configuring it for vCenter.

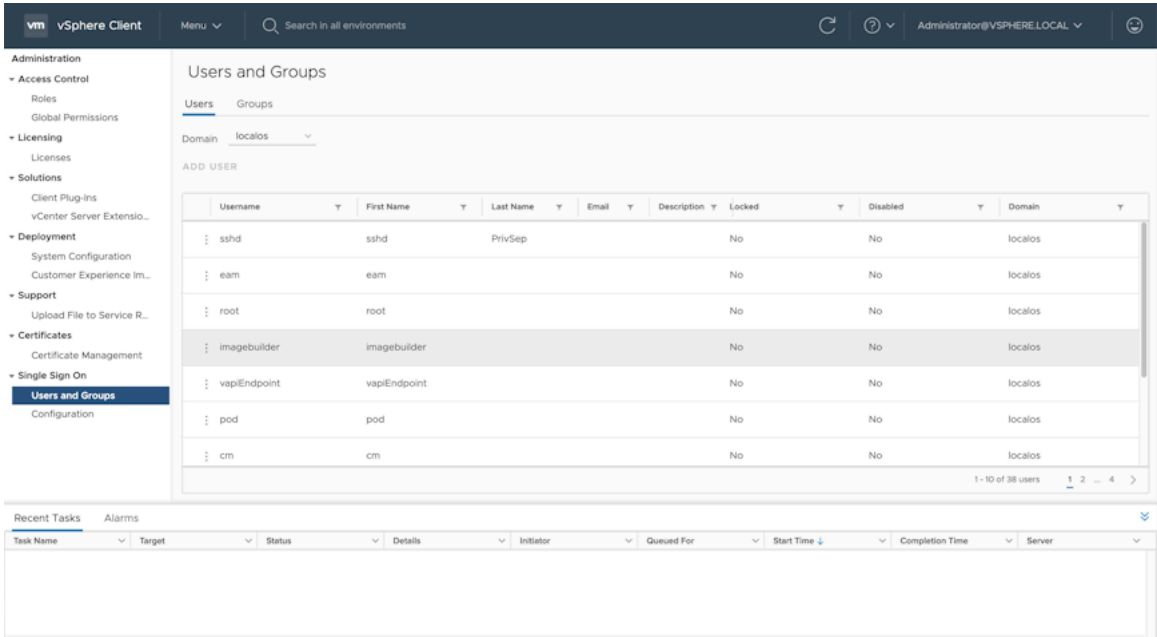
Please refer to the Platform Compatibility section for more information.

# 9.11 Minimum Permissions for Non-admin Users

For a non-admin user to add, remove, edit, or monitor a vCenter via the DANZ Monitoring Fabric (DMF), the privilege level assigned to the non-admin user is VSPAN operation. To assign VSPAN operation privileges to a user, perform the following steps:

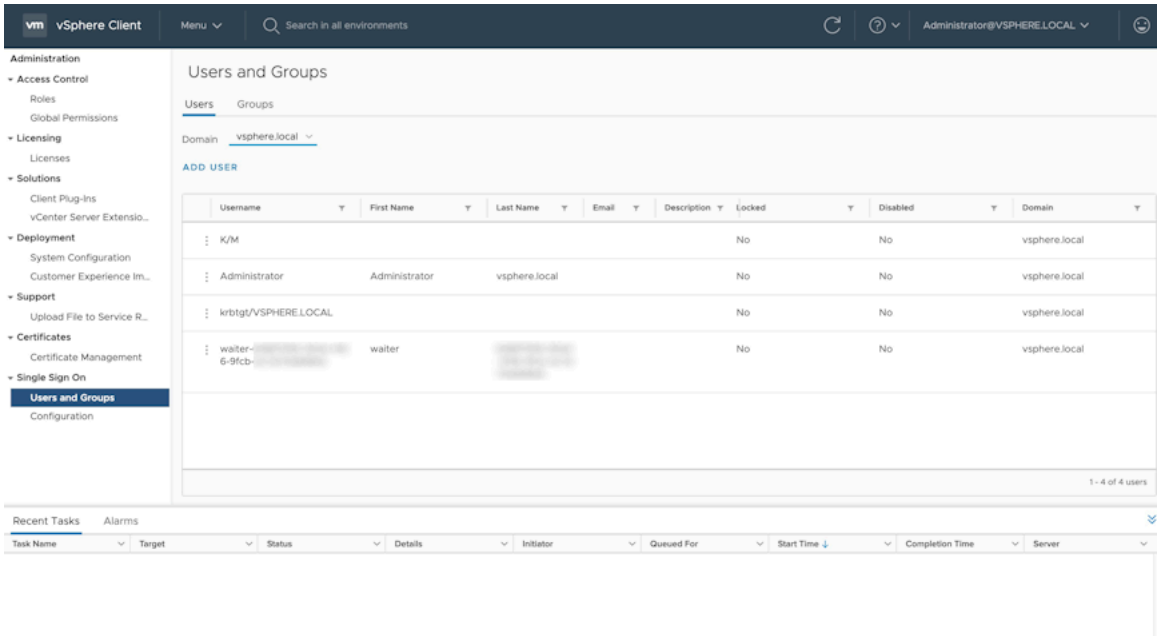
- 1. From the vCenter GUI, navigate to **Menu > Administration**.
- 2. Once on the page, click on the **Users and Groups** link in the navigation bar on the left.

Figure 9-37: Users and Groups



- 3. Click on the **Users** tab and ensure the appropriate domain is selected (in this case, the domain is `vsphere.local`).

Figure 9-38: Domain Selection



4. Next, click on the **ADD USER** link and create the desired user. (In the example below, a user called **dmf-*alice*** is created.)

**Figure 9-39: Add a New User**

Add User

Username \*

dmf-*alice*

Password \*

.....

i

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 9-40: Verify User Created**

vm vSphere Client

Menu

Search in all environments

Administrator@VSPHERE.LOCAL

Administration

Access Control

Roles

Global Permissions

Licensing

Solutions

Client Plug-Ins

vCenter Server Extension...

Deployment

System Configuration

Customer Experience Im...

Support

Upload File to Service R...

Certificates

Certificate Management

Single Sign On

**Users and Groups**

Configuration

Users and Groups

Users Groups

Domain vsphere.local

ADD USER

|  | Username             | First Name    | Last Name | Email                    | Description | Locked | Disabled | Domain        |
|--|----------------------|---------------|-----------|--------------------------|-------------|--------|----------|---------------|
|  | K/M                  |               |           |                          |             | No     | No       | vsphere.local |
|  | Administrator        | Administrator |           |                          |             | No     | No       | vsphere.local |
|  | krbtgt/VSPHERE.LOCAL |               |           |                          |             | No     | No       | vsphere.local |
|  | walter-6-9fcb        | walter        |           |                          |             | No     | No       | vsphere.local |
|  | dmf- <i>alice</i>    | Alice         | D         | <i>alice</i> @arista.com |             | No     | No       | vsphere.local |

1 - 5 of 5 users

Recent Tasks

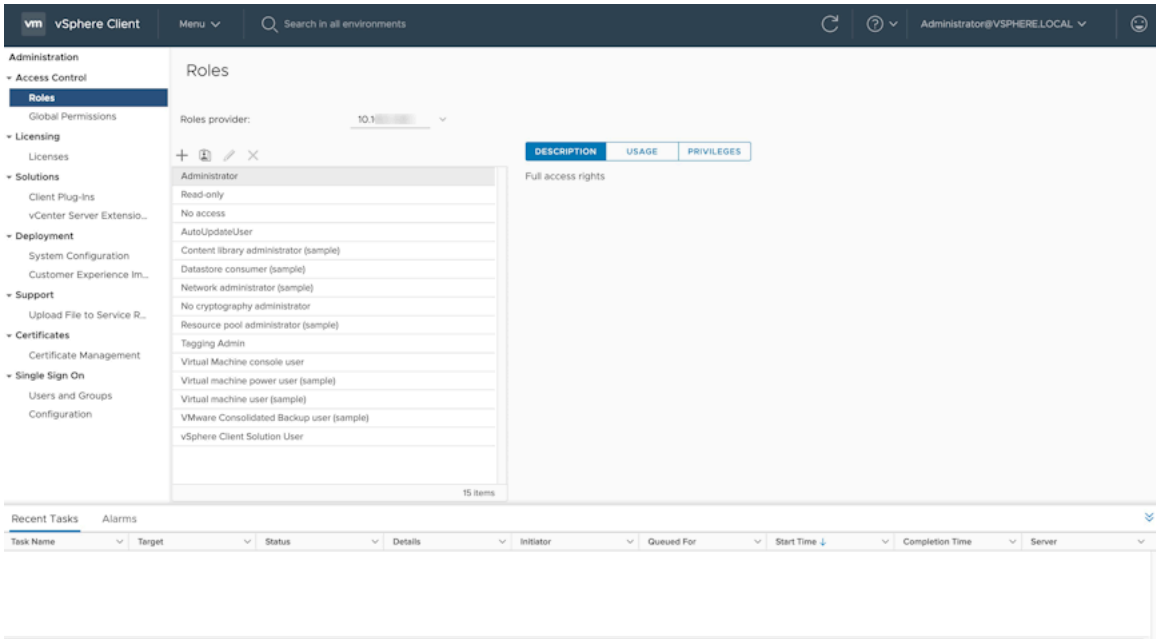
Alarms

| Task Name | Target | Status | Details | Initiator | Queued For | Start Time | Completion Time | Server |
|-----------|--------|--------|---------|-----------|------------|------------|-----------------|--------|
|-----------|--------|--------|---------|-----------|------------|------------|-----------------|--------|



6. After creating the desired user, 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 9-41: Add a New Role



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 the new role *monitor-dmf* is created.) Click **Finish** to create the new role.

**Figure 9-42: 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 9-43: Save New Role**

New Role

Role name

monitor-dmf

Description

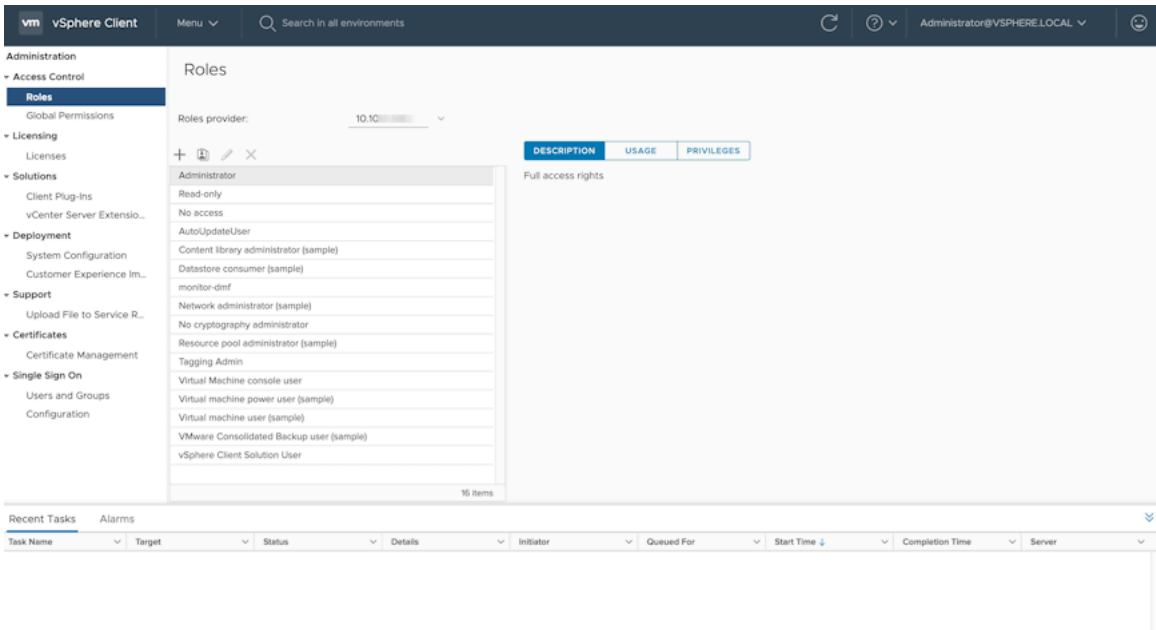
CANCEL

BACK

FINISH

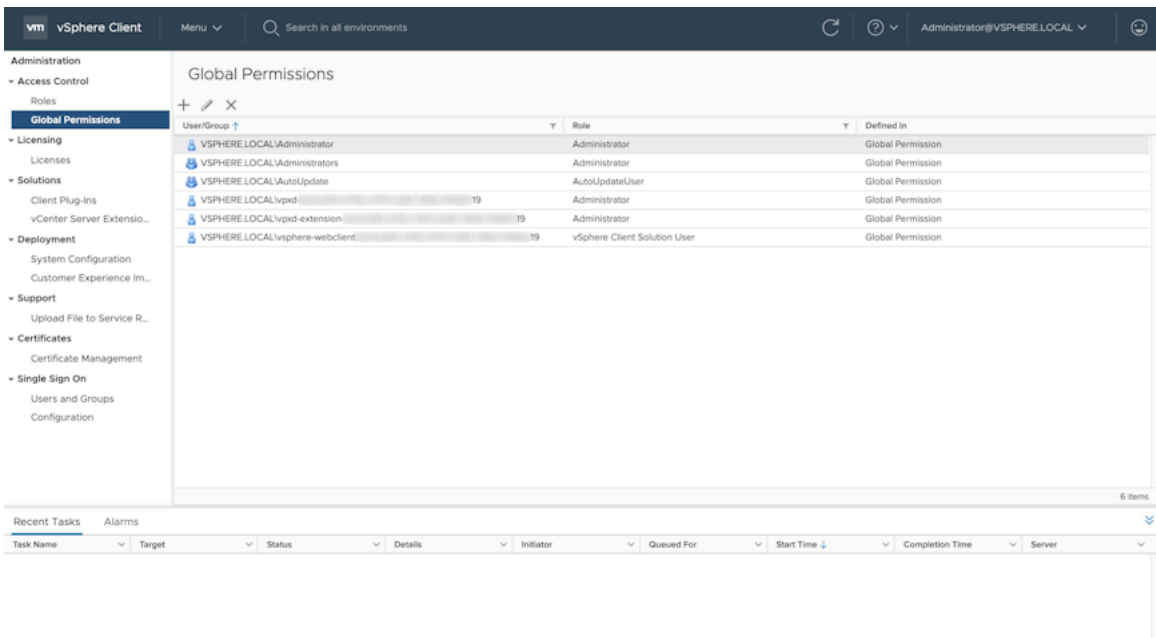
8. Verify the creation of the new role on the **Roles** page.

**Figure 9-44: Verify New Role 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 9-45: 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 mark the **Propagate to children** checkbox.

**Figure 9-46: Assign Role to User**

Add Permission

Global Permission Root

×

User

vsphere.local

▼

Q dmf-

Role

monitor-dmf

▼

☒ Propagate to children

CANCEL

OK

11. Verify assigning the newly created role to the newly created user.

**Figure 9-47: Verify Role Assignment to User**

| 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\vpd-extension-19  | Administrator                | Global Permission |
| VSPHERE.LOCAL\vpd-extension-19  | Administrator                | Global Permission |
| VSPHERE.LOCAL\vsphere-webclient | vsphere Client Solution User | Global Permission |

## 9.12 Monitor VMware vCenter Traffic by VM Names

Match VMware vCenter-specific information in the policy. Specifically, this feature matches traffic using VMware vCenter Virtual Machine (VM) names and requires DANZ Monitoring Fabric (DMF) vCenter integration.

## 9.12.1 Using the CLI to Monitor vCenter Traffic by VM Names

### Configuration

This feature works with vCenter integration; therefore, configure vCenter Integration in DANZ Monitoring Fabric (DMF). Configure vCenter mapping in the policy, then define a policy match using VM names in the vCenter as illustrated in the following configuration example:

```
dmf-controller-1(config)# policy v1
dmf-controller-1(config-policy)# action forward
dmf-controller-1(config-policy)# filter-interface filter-interface
dmf-controller-1(config-policy)# delivery-interface delivery-interface
dmf-controller-1(config-policy)# filter-vcenter vcenter-name
dmf-controller-1(config-policy)# 1 match ip src-vm-name vm-name dst-vm-name vm-
name
dmf-controller-1(config-policy)# 2 match ip6 src-vm-name vm-name
```

### Show Commands

Enter the **show running-config policy policy name** command to display the configuration.

```
dmf-controller-1# show running-config policy v1

! policy
policy v1
action forward
delivery-interface delivery-interface
filter-interface filter-interface
filter-vcenter vcenter-name
1 match ip src-vm-name vm-name dst-vm-name vm-name
2 match ip6 src-vm-name vm-name
```

The **show policy policy name** command displays the policy information, including stats.

```
dmf-controller-1# show policy v2
Policy Name           : v2
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             : 5
Post Match Filter Traffic : -
Total Delivery Rate: -
Total Pre Service Rate : -
Total Post Service Rate : -
Overlapping Policies   : none
Component Policies    : none
Installed Time         : 2023-12-21 19:00:39 UTC
Installed Duration     : 50 minutes, 11 secs
```

```

~~~~~ Match Rules ~~~~~
Rule
-|-----|
1 1 match ip src-vm-name DMF-RADIUS-SERVER-1 dst-vm-name DMF-TACACS-SERVER-1

~~~~~ Filter Interface(s) ~~~~~
# DMF IF          Switch      IF Name      State Dir Packets Bytes Pkt Rate Bit
  Rate Counter Reset Time
-|-----|-----|-----|-----|---|---|-----|-----|-----|
|-----|-----|-----|-----|
1 span_from_arista DELL-S4048 ethernet20 up    rx    0        0        0        -
  2023-12-21 19:00:39.941000 UTC

~~~~~ Delivery Interface(s) ~~~~~
DMF IF Switch IF Name State Dir Packets Bytes Pkt Rate Bit
 Rate Counter Reset Time
-|-----|-----|-----|-----|---|---|-----|-----|-----|
|-----|-----|-----|-----|
1 ubuntu-tools DELL-S4048 ethernet50/2 up tx 0 0 0 -
 2023-12-21 19:00:39.941000 UTC
~ Service Interface(s) ~
None.
~ Core Interface(s) ~
None.
~ Failed Path(s) ~
None.

```

The **show vcenter vcenter name endpoint** command displays the vCenter VM information, including networks.

```

dmf-controller-1# show vcenter vcenter1 endpoint
vCenter Name VM Name ESXi Host Name Network Interface Name MAC Address
 IP Address Virtual Switch Portgroup
 Power State
--|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|
1 vcenter1 ub-11-216 10.240.155.216 Network adapter 1 00:50:56:8b:4
d:03 (VMware) 1.1.11.216/24, fe80::250:56ff:fe8b:4d03/64 DVS-DMF vlan11
 powered-on
2 vcenter1 ub-12-216 10.240.155.216 Network adapter 1 00:50:56:8b:7
2:a0 (VMware) 1.1.12.216/24, fe80::250:56ff:fe8b:72a0/64 DVS-DMF vlan12
 powered-on
3 vcenter1 ub-13-216 10.240.155.216 Network adapter 1 00:50:56:8b:c
0:06 (VMware) 1.1.13.216/24, fe80::250:56ff:fe8b:c006/64 DVS-DMF vlan-10
 powered-on
4 vcenter1 ub-14-216 10.240.155.216 Network adapter 1 00:50:56:8b:d
1:d9 (VMware) 1.1.14.216/24, fe80::250:56ff:fe8b:d1d9/64 DVS-DMF vlan-10
 powered-on


```

## 9.12.2 Using the GUI to Monitor vCenter Traffic by VM Names

Configure vCenter VM name matches under the DANZ Monitoring Fabric (DMF) policies match rules section. For example:

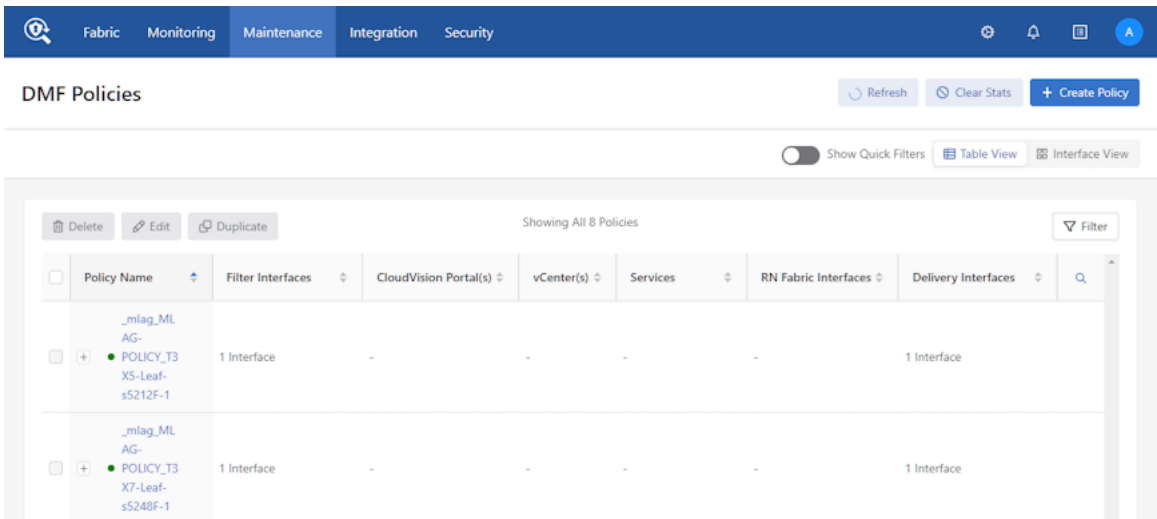
- 1. In the DMF GUI, navigate to the **Monitoring > Policies** page.

Figure 9-48: DMF Policies



- 2. Click **Create Policy** to create a new policy or edit an existing one by selecting a row from the **Policies Table** and clicking **Edit**.

Figure 9-49: Create / Edit Policy



3. Navigate to the **Match Traffic** tab.

**Figure 9-50: Match Traffic**

Create Policy

Policy Details

Enable Advanced Options

Collapse

\* Name

Input value

Description

Input value

Action

Forward

Push VLAN

##

\* Priority

100

Active

Port Selection

Match Traffic

Stat: Bit Rate

Traffic Sources

Create

+ Add Port(s)

Services

+ Add Service(s)

Destination Tools

Create

+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

Cancel

Create Policy

354



4. Click **Configure a Rule** to configure a custom match rule.

**Figure 9-51: Configure a Rule**

Create Policy

Policy Details

Enable Advanced Options

Collapse

| * Name                 | Description            | Action                         | Push VLAN     | * Priority ?   | Active      |
|------------------------|------------------------|--------------------------------|---------------|----------------|-------------|
| <div>Input value</div> | <div>Input value</div> | <div>Forward</div> <div></div> | <div>##</div> | <div>100</div> | <div></div> |

Port Selection

Match Traffic

Select an option below or configure a custom rule.

Allow All Traffic

Deny All Traffic

or

Configure A Rule

Import Rules

Match Traffic: No Selection has been added

Cancel

Create Policy

5. Set the **EtherType** to **IPv4** or **IPv6**.

6. Add the **Source IP Address** as the vCenter VM name. Select the **Virtual Machine** option from the **Source IP Address** drop-down and select a virtual machine from the **VM Name** drop-down.

**Figure 9-52: Source IP Address VM Name**

Configure A Traffic Rule Cancel Add Rule

Custom Rule

EtherType: IPv4 (2048) | IP Protocol: None | TCP Flags: Select Flags | IP DSCP: | IP ECN: | IP Fragment: Unspecified

VLAN: Any | Bidirectional: On | Off

Additional Configurations

Source IP Address: Virtual Machi... | VM name: Select | Ports: None | MAC Address: - Address - | MAC Mask: - Mask -

Destination IP Address: None | MAC Mask: - Mask -

VM name dropdown options:

- ACCTON-TB2-1
- ACCTON-TB2-2
- ANIMESH-DMF-EOS-TB-1
- ANIMESH-DMF-EOS-TB-2
- BSN-AN-Active-Directory
- BSN-AN-WIN-AD-1
- BSN-RDP-001
- BSN-RDP-002

7. Add the Destination IP address as the vCenter VM name. Select the **Virtual Machine** option from the **Destination IP Address** drop-down and select a virtual machine from the **VM Name** drop-down.

**Figure 9-53: Destination IP VM**

Create Policy ✕

Policy Details ☐ Enable Advanced Options Collapse

| * Name      | Description | Action  | Push VLAN | * Priority ? | Active                              |
|-------------|-------------|---------|-----------|--------------|-------------------------------------|
| Input value | Input value | Forward | ##        | 100          | <input checked="" type="checkbox"/> |

Configure A Traffic Rule Cancel Add Rule

Custom Rule

|             |             |              |               |
|-------------|-------------|--------------|---------------|
| EtherType   | IP Protocol | TCP Flags    | IP DSCP       |
| IPv4 (2048) | None        | Select Flags |               |
| IP ECN      | IP Fragment | VLAN         | Bidirectional |
|             | Unspecified | Any          | On Off        |

Additional Configurations

|                        |           |         |             |
|------------------------|-----------|---------|-------------|
| Source IP Address      | VM name ? | Ports ? | MAC Address |
| Virtual Machine        | Select    | None    | - Address - |
| MAC Mask               |           |         |             |
| - Mask -               |           |         |             |
| Destination IP Address | VM name ? | Ports ? | MAC Address |
| Virtual Machine        | Select    | None    | - Address - |
| MAC Mask               |           |         |             |
| - Mask -               |           |         |             |



**Note:** If the **VM Name** drop-down shows **No Data**, ensure only one vCenter is affiliated with the policy (under **Traffic Sources**).

8. Click **Add Rule** to add the match rule to the policy.
9. After entering other inputs as required, click **Create Policy** (or **Save Policy**) to save the configuration.

### 9.12.3 Troubleshooting

Fabric errors and warnings are very useful for troubleshooting this feature.

When using the `show fabric warnings` command, the following validation message displays when the vCenter integration cannot resolve the IP address for the VM name used in the policy.

```
dmf-controller-1# show fabric warnings
~~~~~ Policy related warning
~~~~~
Policy Name Warning
```

```

-|-----|-----
-----|
1 v1 No IP found for VMs [ub-15-216, ub-216-multinic, ub-217-vlan10,
ub-14-216, ub-11-216] associated with policy

```

When VM names used in a policy are matched, the following validation message content appears when a vCenter instance is not associated with the policy.

```

dmf-controller-1# show fabric warnings
~~~~~ Policy related warning ~~~~~
# Policy Name Warning
-|-----|-----
1 v1          No vCenter associated to policy with VM matches

```

## 9.12.4 Limitations

- This feature only works with vCenter integration and a direct Switch Port Analyzer (SPAN) from a switch with ESXi traffic.
- VM interface IP addresses connected to **dvs** will only be added to policy matches.
- The system may use extra TCAM entries if the management network uses **dvs**.
- **Vmkernal** names cannot be matched in the policy.
- When a VM name with multiple vNICs (multiple IP addresses) matches the policy, a TCAM entry is added for all the IP addresses.
- VM Names cannot be matched with the MAC option in the policy.
- If the vCenter becomes disconnected, policies associated with the VM names may not get correct matches or traffic.

## CloudVision DMF Integration

This chapter describes integrating CloudVision with the DANZ Monitoring Fabric (DMF).

### 10.1 Overview

In a typical CloudVision-DMF integration deployment, CloudVision Portal (CVP) deploys alongside the DANZ Monitoring Fabric (DMF). The DMF Controller communicates with CVP to retrieve its managed device inventory and configures port mirroring sessions on any CVP-managed production devices that are Arista Extensible Operating System (EOS) switches.

Configuration on the DMF Controller provides the information necessary to communicate with CVP: the CVP hostname or IP address and user credentials.

Policy configuration on the DMF Controller specifies what to monitor in the production network managed by CVP, such as the production switches, the switch interfaces to monitor traffic from, and the direction of mirrored traffic (bidirectional, ingress, or egress). In addition, the configuration on the DMF Controller can define whether to use a Switch Port Analyzer (SPAN) session or a Layer-2 Generic Routing Encapsulation (L2GRE) tunnel session on a CVP-managed device. When using SPAN, the DMF configuration includes the switch interface to monitor traffic. When using L2GRE, the DMF configuration includes monitoring traffic to the Tunnel End Point (TEP).

**Figure 10-1: Simple CloudVision-DMF Integration Deployment**




Figure 1. Example of a simple CloudVision-DMF integration deployment

The preceding figure illustrates a simple CloudVision-DMF integration configuration where CloudVision Portal and DMF can communicate with each other. DMF monitors traffic from CVP directly to one of its fabric switches (a filter switch), as indicated by the red arrow labeled “SPAN.” DMF also monitors traffic from CVP to a TEP configured on the fabric using an L2GRE tunnel, as indicated by the green arrow labeled “L2GRE Tunnel”. Since DMF initiates monitoring using policy configuration, the policies monitoring CVP will handle the traffic according to their configuration, for example, forwarding it to a delivery interface. This feature enables the automation of the creation, modification, and deletion of filter interfaces and tunnel interfaces in DMF and mirroring sessions on CVP-managed devices.

---

## 10.2 Compatibility Requirements

### EOS Platform Compatibility

- Any platform supporting port mirroring or mirroring to a GRE tunnel specifying the GRE key is supported; refer to [EOS Port Mirroring TOI](#) and [EOS Mirroring to GRE Tunnel TOI](#) for more details.

### CloudVision Compatibility

- On-premise 2024.2.0 and newer is recommended.

### CloudVision Requirements

- The user configured in DMF for CVP integration must have sufficient permissions in CVP. The minimum permissions required are:
  - Devices: Read access to inventory management.
  - gNMI: Read and write access to the gNMI service.
- Register the devices that DMF will monitor for use in Studios using the [Inventory and Topology](#) Studio.

## 10.3 CloudVision DMF Integration using the CLI

To integrate with CloudVision Portal, configure a CVP instance in the DMF Controller, enabling communication between CVP and DMF. The CVP hostname or IP address must be reachable from the DMF Controller. If CVP is a multi-node system, using a hostname that will resolve to the primary node is recommended to maintain the connection in case of a primary node failure. The user in the CVP integration configuration must have at least the permissions in CloudVision as outlined in [CloudVision Requirements](#).

### Configure using the CLI

Configure a CVP instance on the DMF Controller using the following series of commands:

```
dmf-controller(config)# cvp cvp_instance_name
dmf-controller(config-cvp)# host-name cvp_hostname_or_ip
dmf-controller(config-cvp)# username username
dmf-controller(config-cvp)# password password
```

Add a description to the CVP instance using the **description** command, as required.

```
dmf-controller(config-cvp)# description description_of_cvp_instance
```

Refresh the connection between DMF and CVP with the **sync** command, which sends a request to CVP to re-authenticate the connection and to re-fetch the inventory:

```
dmf-controller(config)# sync cvp cvp_instance_name
```

To use L2GRE tunnels in the integration, enable tunneling in the DMF Controller and set the match mode to one of the following that is compatible with tunneling: **full-match** or **13-14-offset-match**. Configure tunnel endpoints to allow monitoring from CVP to DMF using an L2GRE tunnel; add a tunnel endpoint to a policy configuration or to a CVP integration instance's configuration to **optionally** define a default tunnel endpoint for this instance.

Before the DMF 8.5 release, the following command string was mandatory:

```
dmf-controller(config)# tunnel-endpoint tep_name switch fabric_switch fabric_switch_interface ip-  
address tep_ip mask subnet_mask gateway gateway_ip
```

However, starting with DMF 8.5.0, the mask and gateway parameters in the tunnel-endpoint command are now optional. Thus, configure a tunnel endpoint using the following command:

```
dmf-controller(config)# tunnel-endpoint tep_name switch fabric_switch fabric_switch_interface ip-  
address tep_ip
```

To set a default tunnel endpoint for a CVP integration instance, use the following commands:

```
dmf-controller(config)# cvp cvp_instance_name  
dmf-controller(config-cvp)# default-tunnel-endpoint tep_name
```

To remove a default tunnel endpoint for a CVP integration instance, use the following commands:

```
dmf-controller(config)# cvp cvp_instance_name  
dmf-controller(config-cvp)# no default-tunnel-endpoint tep_name
```

Starting with the DMF 8.6.0 release, a configuration flag called **preserve-mirror-sessions** per CVP instance indicates whether mirroring sessions will be preserved for the CVP instance when uninstalling DMF policies configured with it. By default, the flag is **false**, meaning existing mirroring sessions are automatically removed if the relevant DMF policies are uninstalled.

Enable preserving mirroring sessions using the **preserve-mirror-sessions** command.

```
dmf-controller(config-cvp)# preserve-mirror-sessions
```

Conversely, disable preserving mirroring sessions (default behavior) using the **no preserve-mirror-sessions** command.

```
dmf-controller(config-cvp)# no preserve-mirror-sessions
```

## Monitoring Configuration in Policies

DMF uses policies to create, update, or remove the monitoring of CVP-managed devices. DMF supports monitoring multiple CVP instances, switches, and interfaces as mirroring sources in a single policy or across policies. Configure the mirrored traffic direction to one of the following settings:

- bidirectional (default)
- ingress
- egress

After enabling CVP integration in a DMF policy (i.e., adding a CVP instance as a traffic source), the DMF Controller will automatically create filter interfaces and tunnel interfaces, with origination "auto-generated." A mirroring session is automatically created on the CVP-managed switch; DMF does this via the mirroring Studio and the change control process on CVP.



**Note:** If a DMF-managed mirroring session exists on a switch for one DMF policy with identical sources and the same destination as needed for another DMF policy, both policies use the same mirroring session.

Add a CVP instance as a traffic source in a DMF policy using the following series of commands:

```
dmf-controller(config)# policy policy_name  
dmf-controller(config-policy)# filter-cvp cvp_instance_name  
dmf-controller(config-policy-cvp)#
```

To monitor traffic using SPAN, configure a SPAN interface (on the CVP-managed device) as the destination in a DMF policy, along with the source interfaces (on the CVP-managed device) and optionally the direction for each source interface.

Select the switch interfaces on CVP-managed devices individually as source interfaces to a SPAN interface on that switch using the following series of commands where including the direction is optional :

```
dmf-controller(config-policy-cvp)# device device_hostname
dmf-controller(config-policy-cvp-device)# src-interface source_interface
span-interface span_interface direction ingress | egress | bidirectional
```

Select the switch interfaces on CVP-managed devices as source interfaces using an interface range to a SPAN interface on that switch using the following series of commands where including the direction is optional:

```
dmf-controller(config-policy-cvp)# device device_hostname
dmf-controller(config-policy-cvp-device)# src-interface-range
start start_of_range
end end_of_range span-interface span_interface
direction ingress | egress | bidirectional
```

To monitor traffic using L2GRE tunneling, choose from two options: (1) configure a tunnel endpoint (in DMF) as the destination in a DMF policy along with the source interfaces (on the CVP-managed device) and optionally the direction for each source interface, or (2) omit the destination in a DMF policy along with configuring the source interfaces (on the CVP-managed device) and optionally the direction for each source interface.

A GRE tunnel source IP can be optionally configured on DMF as the tunnel source IP on the CVP-managed device to overcome reachability issues due to possible Reverse Path Forwarding (RPF) checks between the CVP and DMF deployment. By default, the tunnel source IP is the switch's management IP.

Select the switch interfaces on CVP-managed devices individually as source interfaces to a tunnel endpoint configured in DMF using the following series of commands where the direction is optional:

```
dmf-controller(config-policy-cvp)# device device_hostname
dmf-controller(config-policy-cvp-device)# src-interface source_interface gre-
tunnel-src src_ip
gre-tunnel-endpoint tep_name direction ingress | egress | bidirectional
```

The **src-interface-range** command is also supported for GRE tunnel configuration in a policy.

The following example illustrates two DMF policies' configuration, where **testPolicy1** is monitoring traffic from **Ethernet1** on the CVP-managed device (production switch) called **dev1** in the CVP instance, **test**, to **Ethernet2** on the same device, using SPAN, and forwarding the traffic to the delivery interface called **tool1**; **testPolicy2** is monitoring traffic from **Ethernet5** on the CVP-managed device called **dev2** in the same CVP instance, **test**, to the default tunnel endpoint called **TEP1** defined in the CVP integration instance configuration, using L2GRE tunneling, and forwarding the traffic to the delivery interface called **tool2**.

```
! cvp
cvp test
  default-tunnel-endpoint TEP1
  hashed-password abc123
  host-name test.arista.com
  user-name cvpadmin

! policy
policy testPolicy1
  action forward
  delivery-interface tool1
  1 match any
  filter-cvp test
```



```

!
device dev1
src-interface Ethernet1 span-interface Ethernet2

policy testPolicy2
action forward
delivery-interface tool2
1 match any
filter-cvp test
!
device dev2
src-interface Ethernet5

```

Suppose you remove the configuration to monitor CVP-managed devices from the DMF Controller. In that case, the system removes the corresponding auto-generated filter interfaces and tunnel interfaces from DMF and deletes the auto-created mirroring sessions on the switch.

To stop monitoring a source interface or a range of source interfaces on a CVP-managed device, remove its configuration from a DMF policy using the following series of commands:

```

dmf-controller(config-policy-cvp-device)# no src-interface source_interface
dmf-controller(config-policy-cvp-device)# no src-interface-range
start start_of_range end end_of_range

```

Stop monitoring a device in a CVP instance in a DMF policy.

```

dmf-controller(config-policy-cvp)# no device device_hostname

```

Stop monitoring a CVP instance in a DMF policy and all its devices.

```

dmf-controller(config)# policy policy_name
dmf-controller(config-policy)# no filter-cvp cvp_instance_name

```

To disable integration with a CVP instance, remove its CVP integration instance configuration and remove it from all DMF policies using the following series of commands:

```

dmf-controller(config)# no cvp cvp_instance_name
dmf-controller(config)# policy policy_name
dmf-controller(config-policy)# no filter-cvp cvp_instance_name

```

## Show Commands

After configuring a CVP instance, the **show cvp *cvp\_instance\_name*** command displays the configuration and connection status information.

```

dmf-controller(config)# show cvp test
# CVP  Hostname      State      Last Update Time      Detail State
Version
-|----|-----|-----|-----|-----
-----|-----|
1 test test.arista.com connected 2023-12-13 05:17:28.512000 UTC connected
2024.1.0

```

The **show cvp *cvp\_instance\_name* detail** command displays detailed status information about the integration.

```

dmf-controller(config)# show cvp test detail
CVP           : test
Hostname      : test.arista.com
State         : connected

```

```
Last Update Time : 2023-12-13 05:17:54.072000 UTC
Detail State      : connected
Version          : 2024.1.0
```

The **show cvp cvp\_instance\_name alert** and **show cvp cvp\_instance\_name error** commands display runtime warnings and alerts, and errors, if any.

```
dmf-controller(config)# show cvp cvp_instance_name alert
dmf-controller(config)# show cvp cvp_instance_name error
```



**Note:** It is possible to specify **all** in the above **show cvp** commands to see the information for all CVP integration instances on the DMF Controller; for example, **show cvp all alert**.

The **show cvp cvp\_instance\_name device device\_hostname** command displays the device inventory in the CVP deployment; only EOS devices are supported. Using **all** is possible for the CVP instance name and the device hostname in the **show cvp cvp\_instance\_name device device\_hostname** command.

```
dmf-controller(config)# show cvp test device all
~~~~~ Device Inventory ~~~~~
CVP Device FQDN Streaming Model Software IP Address MAC Address Device ID
-|---|-----|-----|-----|-----|-----|-----|-----|
1 test dev123 dev123.arista.com active ABC-123 4.31.2F 10.10.10.10 aa:bb:cc:dd:ee:ff (Arista) DEV123
```

The **show cvp cvp\_instance\_name device device\_hostname interface** command includes a list of all the device interfaces.

```
dmf-controller(config)# show cvp test device all
~~~~~ Device Inventory ~~~~~
# CVP Device FQDN Streaming Model Software IP Address MAC Address Device ID
-|---|-----|-----|-----|-----|-----|-----|-----|
1 test dev123 dev123.arista.com active ABC-123 4.31.2F 10.10.10.10 aa:bb:cc:dd:ee:ff (Arista) DEV123

~~~~~ Device Interfaces ~~~~~
CVP Device Interface
-|---|-----|-----|
1 test dev123 Ethernet1
2 test dev123 Ethernet2
3 test dev123 Ethernet3
```

After configuring a DMF policy to take a CVP instance as a traffic source, the **show fabric errors** command displays any errors with the integration relating to monitoring, if any.

```
dmf-controller(config)# show fabric errors
```

In addition, the DMF Controller can show the current mirroring sessions configured on a CVP-managed device used to confirm the current state of a mirroring session created by DMF (thus, managed by DMF) or otherwise (non-DMF-managed sessions are only displayed in the **detail** command). There are three commands to display the mirroring state on a CVP-managed device in varying levels of detail, as follows:

1) The **show cvp cvp\_instance\_name device device\_hostname session** command displays only the mirroring sessions managed by DMF.

```
dmf-controller(config)# show cvp test device dev123 session
~~~~~ L2-GRE Mirroring Sessions ~~~~~
# CVP Hostname Tunnel Endpoint Programmed in hardware Tunnel Src Tunnel Dst Src Interface Src Link Status Src Direction
-|---|-----|-----|-----|-----|-----|-----|-----|
1 test dev123 unknown 3.3.3.3 4.4.4.4 Ethernet2 unspecified bidirectional

~~~~~ SPAN Mirroring Sessions ~~~~~
CVP Hostname SPAN Interface SPAN Status Programmed Src Interface Src Link Status Src Direction
-|---|-----|-----|-----|-----|-----|-----|-----|
1 test dev123 Ethernet5 up true Ethernet4 up bidirectional
```

2) The `show cvp cvp_instance_name device device_hostname session brief` command displays a summary of the state of the mirroring sessions managed by DMF.

```
dmf-controller(config)# show cvp cvp_instance_name device device_hostname
session brief
```

3) The `show cvp cvp_instance_name device device_hostname session detail` command displays all the mirroring sessions on the device, both managed by DMF and otherwise, as well as the name of each session.

```
dmf-controller(config)# show cvp cvp_instance_name device device_hostname
session detail
```

## 10.4 CloudVision DMF Integration using the GUI

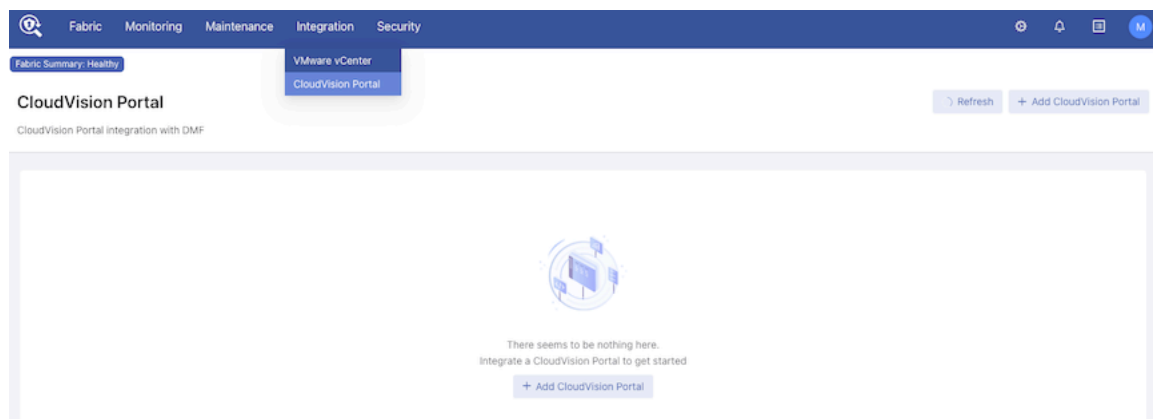
To integrate with CloudVision Portal, configure a CVP instance in the DMF Controller, enabling communication between CVP and DMF. The CVP hostname or IP address must be reachable from the DMF Controller. If CVP is a multi-node system, using a hostname that will resolve to the primary node is recommended to maintain the connection in case of a primary node failure. The user in the CVP integration configuration must have at least the permissions in CloudVision as outlined in CloudVision Requirements.

### Using the GUI

Navigate to **Integration > CloudVision Portal**.

Select **Add CloudVision Portal** to create a new CVP integration instance.

**Figure 10-2: DMF Integration CloudVision Portal**



Enter the CVP integration instance configuration details and click **Submit**.

**Figure 10-3: Add CloudVision Portal**

### Add CloudVision Portal

X

\* Name

Description

\* Hostname/IP Address

\* Username

\* Password

Default Tunnel Endpoint

None

▼

Cancel

Submit

If there are any warnings or errors with the CVP integration instance, click the **alarm bell** icon to view more details.

**Figure 10-4: CloudVision Portal**

## CloudVision Portal

CloudVision Portal integration with DMF

| <div> <div>Edit</div> <div>Delete</div> <div>Sync</div> </div> |      |             |              |  |
|----------------------------------------------------------------|------|-------------|--------------|--|
| <input type="checkbox"/>                                       | Name | Description | Status       |  |
| <input type="checkbox"/>                                       | test |             | Disconnectec |  |

Select the **CVP instance name** to view its details, including the device inventory, port mirroring status, and port mirroring entries in DMF policies.



**Note:** The **Port Mirroring Entries** table will not be visible if no policies using this CVP instance as a traffic source have been configured.

**Figure 10-5: CloudVision Portal Dashboard**

CloudVision Portal

Refresh

+ Add CloudVision Portal

CloudVision Portal integration with DMF

EditDeleteSync

☐

Name

☐

test

test

EditDeleteX

Devices (1)

| Device ID   | FQDN               | Host Name | Total Interfaces | IP Address   | MAC Address       | Software | Streaming Status |
|-------------|--------------------|-----------|------------------|--------------|-------------------|----------|------------------|
| JPE20522931 | Id603.sjc.arist... | Id603     | 131 Interfaces   | 10.240.22... | 94-8e:d3-05:ba... | 4.31.2F  | Active           |

Port Mirroring Status (0)

| Hostname | Session Name | Source Interface | Protocol | GRE Tunnel Source | GRE Tunnel Endpoint | SPAN Interface | Managed by DMF |
|----------|--------------|------------------|----------|-------------------|---------------------|----------------|----------------|
| No data  |              |                  |          |                   |                     |                |                |

Port Mirroring Entries (1)

| Hostname | Source Interface | Tunnel Source | Tunnel Endpoint | SPAN Interface | Direction     | Policies  |
|----------|------------------|---------------|-----------------|----------------|---------------|-----------|
| Id603    | Ethernet2/1      |               |                 | Ethernet1/1    | Bidirectional | fooPolicy |

To start monitoring traffic from the production network:

1. Navigate to **Monitoring > Policies**.
2. Select **Create Policy** and add CVP instances as traffic sources.

3. Select **Add Row** to configure monitoring, such as the device interfaces to monitor, the monitor type (e.g., SPAN or L2GRE), the mirrored traffic direction, and the destination.

Figure 10-6: Create DMF Policy

Create Policy

Traffic Sources

Filter Interfaces/Groups

vCenters

CloudVision Portals

Q Search

test

cvp621.sjc.aristanetworks.com

Version: 2024.1.0

Selected Traffic Sources

CancelAdd 1 Source

CloudVision Portals (1)

test

cvp621.sjc.aristanetworks.com

Default Tunnel Endpoint: Not Configured ⓘ

Version: 2024.1.0

Port Mirroring Entries

| ID      | Device | Source Interface | Monitor Type | Tunnel Source | Tunnel Endpoint | SPAN |
|---------|--------|------------------|--------------|---------------|-----------------|------|
| No data |        |                  |              |               |                 |      |

Add Row

Policy configuration details related to CVP integration appear in the policy’s **Configuration Details** page.

Figure 10-7: Configuration Details

Configuration DetailsOperational Details

fooPolicy

EditDelete

|                       |            |                  |             |        |
|-----------------------|------------|------------------|-------------|--------|
| Description           | Action     | Push VLAN        | Priority    | Active |
| -                     | Forward    | Auto ⓘ           | 100         | Yes    |
| Scheduling Start Time | Run Policy | PTP Timestamping | Root Switch |        |
| Automatic             | Always     | -                | -           |        |

Rules (1)

Allow All Traffic has been selected.

CloudVision (1)

test

Id603: 1 Entry

test Port Mirroring Session Entries (1)

| ID | Device | Source Interface | Monitor Type | Tunnel Source | Tunnel Endpoint | SPAN Interface | Direction     |
|----|--------|------------------|--------------|---------------|-----------------|----------------|---------------|
| 1  | Id603  | Ethernet2/1      | SPAN         |               |                 | Ethernet1/1    | Bidirectional |

To edit the CVP monitoring configuration in a policy, click **Edit** and select **1 Entry** on the **Edit Policy** page, make the required changes, and click **Save Policy**.

**Figure 10-8: Edit Policy**

**Edit Policy**

Policy Details ☐ Enable Advanced Options [Collapse](#)

| Name      | Description | Action  | Push VLAN | Priority | Active                              |
|-----------|-------------|---------|-----------|----------|-------------------------------------|
| fooPolicy | Input value | Forward | ##        | 100      | <input checked="" type="checkbox"/> |

[Port Selection](#) Match Traffic Stat: Utilization

**Traffic Sources** [Create](#)

+ Add Port(s)

test

cvp621.sjc.aristanetworks.com

Version: 2024.1.0

Id603

1 Entry

**Services**

+ Add Service(s)

**Destination Tools** [Create](#)

+ Add Port(s)

Match Traffic: Allow All Traffic (Default) has been selected

[Cancel](#) [Save Policy](#)

## 10.5 Limitations

The following limitations apply to the DANZ Monitoring Fabric (DMF) Controller and CloudVision integration.

- Modifying or deleting auto-generated filter interfaces for CVP integration or adding them manually to non-CloudVision policies will result in unexpected behavior.
- If two or more DMF policies have overlapping CVP monitoring configurations with the same destination, the configuration should be the same in these policies if unexpected traffic is undesired. For example, if a policy has source interfaces A and B on a device in a CVP instance with a SPAN interface as the destination, and if another policy has one source interface, A, on the same device in the same CVP instance with the same SPAN interface as the former policy, then the latter policy will unexpectedly receive

traffic from B as well as the expected source A. This is because the mirroring session in the production switch is reused for both policies.

- An error state in the CVP deployment may prevent DMF from configuring mirroring sessions on the production switches. If DMF encounters such a failure, a fabric error will be displayed, and user action in CVP is required. Examples of such a state include the production switch not being in compliance in CVP, in which case it needs to be brought into compliance, or if there are any pending change controls in CVP, they must be addressed. After taking the appropriate corrective action in CVP, deactivate and reactivate the DMF policy with the error (or delete and recreate it) to retry configuring mirroring.

## 10.6 Troubleshooting

- If there are any fabric errors in creating a mirroring session on a CVP-managed device that specify that user action in CVP is required, take the appropriate corrective action in CVP. Next, delete the CVP monitoring config from the policy and reconfigure it, or delete the policy and then recreate it.
- If there is a message in the `show cvp alert` command output stating a session failed to be updated and an authentication error message in `/var/log/vsphere-extension/vsphere-extension.log` containing `GNMI7001...Status unauthenticated`, run the `sync cvp` command, and then deactivate and reactivate the relevant DMF policy.
- If there is a fabric error containing `No change controls to approve and execute`, a possible cause is that the device had not been registered for use in Studios in CVP using the Inventory and Topology studio. If so, register the device in CVP and delete and recreate the DMF policy.

## 10.7 Resources

Please refer to [tunneling and policies](#) in this guide for further information.



## Advanced Fabric Settings

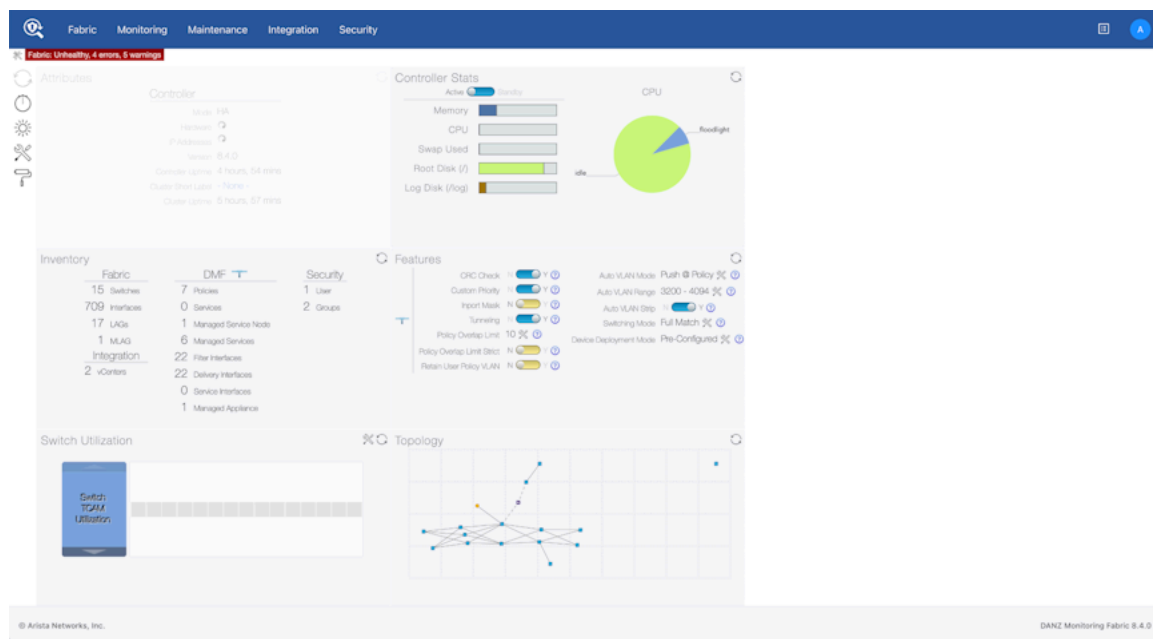
This chapter describes fabric-wide configuration options required in advanced use cases for deploying DMF policies.

### 11.1 Configuring Advanced Fabric Settings

#### Overview

Before 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 11-1: DMF Legacy Page (pre 8.4)**



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

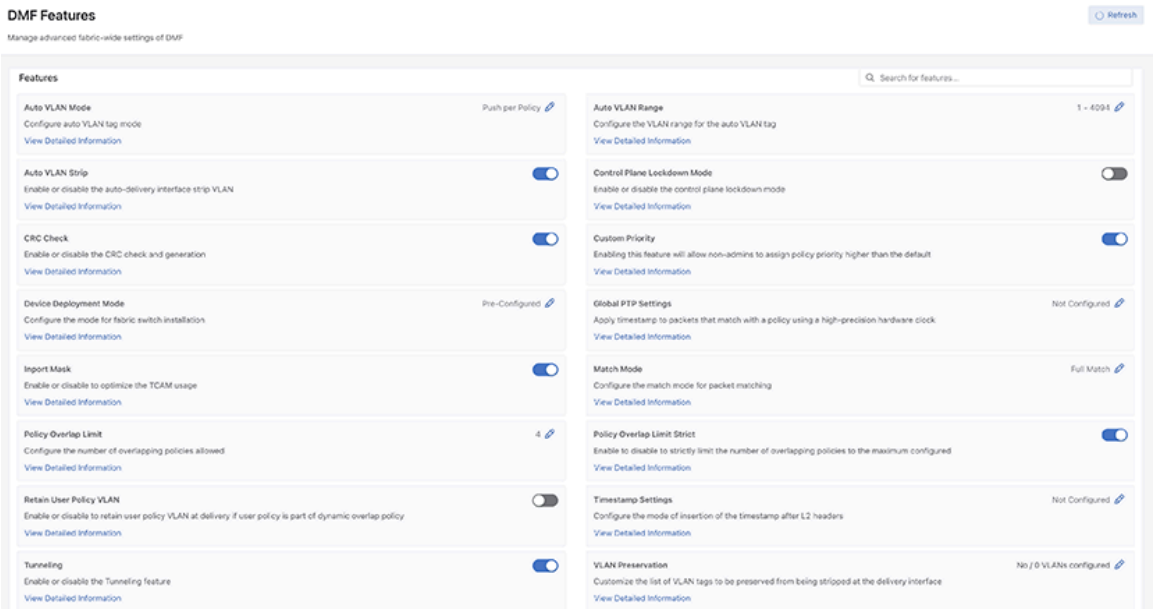
**Figure 11-2: Gear Icon**



Page Layout

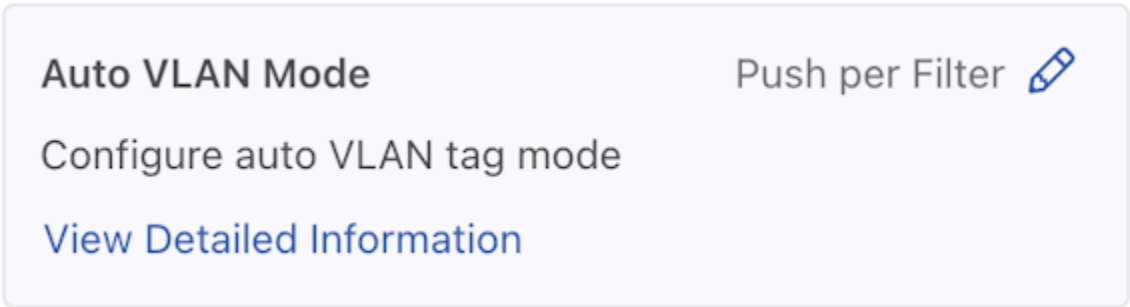
All fabric-wide configuration settings required in advanced use cases for deploying DMF policies appear in the new DMF Features Page.

Figure 11-3: DMF Features Page



Each card on the page corresponds to a feature set.

Figure 11-4: 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 6: 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           |

## 11.2 Managing VLAN Tags in the Monitoring Fabric

Analysis tools often use VLAN tags to identify the filter interface receiving traffic. 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, from **1** to **4094** by default. 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. Manually assign any VLAN ID not in the auto-VLAN range to a filter interface.

The VLAN ID assigned to policies or filter interfaces remains unchanged 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 the outer VLAN tag is rewritten with the assigned VLAN ID if the packet already has two VLAN tags.

The following table summarizes how VLAN tagging occurs with the different auto-VLAN modes:

**Table 7: 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 outer (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 enabling push-per-policy, the **auto-delivery-interface-vlan-strip** feature is enabled (if disabled) before enabling push-per-policy. When enabling push-per-filter, the global delivery strip option is not enabled if 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 8: 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 choosing 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
```

## 11.3 Auto VLAN Mode

Analysis tools often use VLAN tags to identify the filter interface receiving traffic. 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 from **1** to **4094** by default. 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 new vlan range can be specified using the auto-vlan-range command. Manually assign any VLAN ID not in the auto-VLAN range to a filter interface.

The following table summarizes how VLAN tagging occurs with the different Auto VLAN modes.

| Traffic with VLAN tag type | push-per-policy Mode<br>(Applies to all supported switches)                                                                                          | push-per-filter Mode<br>(Applies to all supported switches)                             |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Untagged                   | Pushes a single tag                                                                                                                                  | Pushes a single tag                                                                     |
| Single tag                 | Pushes an outer (second) tag                                                                                                                         | Pushes an outer (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 enabling push-per-policy, the **auto-delivery-interface-vlan-strip** feature is enabled (if disabled) before enabling push-per-policy. When enabling push-per-filter, the global delivery strip option is not enabled if previously disabled.

The following table summarizes how different Auto VLAN modes affect supported applications and services.



**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 only enabled when choosing 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](#)

### 11.3.1 Configuring Auto VLAN Mode using the CLI

To set the auto VLAN mode, perform the following steps:

1. When setting the auto VLAN mode to push-per-filter, define the range of automatically assigned VLAN IDs 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. 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 }
```

Find details of the impact of these options 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.

## 11.3.2 Configuring Auto VLAN Mode using the GUI

### Auto VLAN Mode

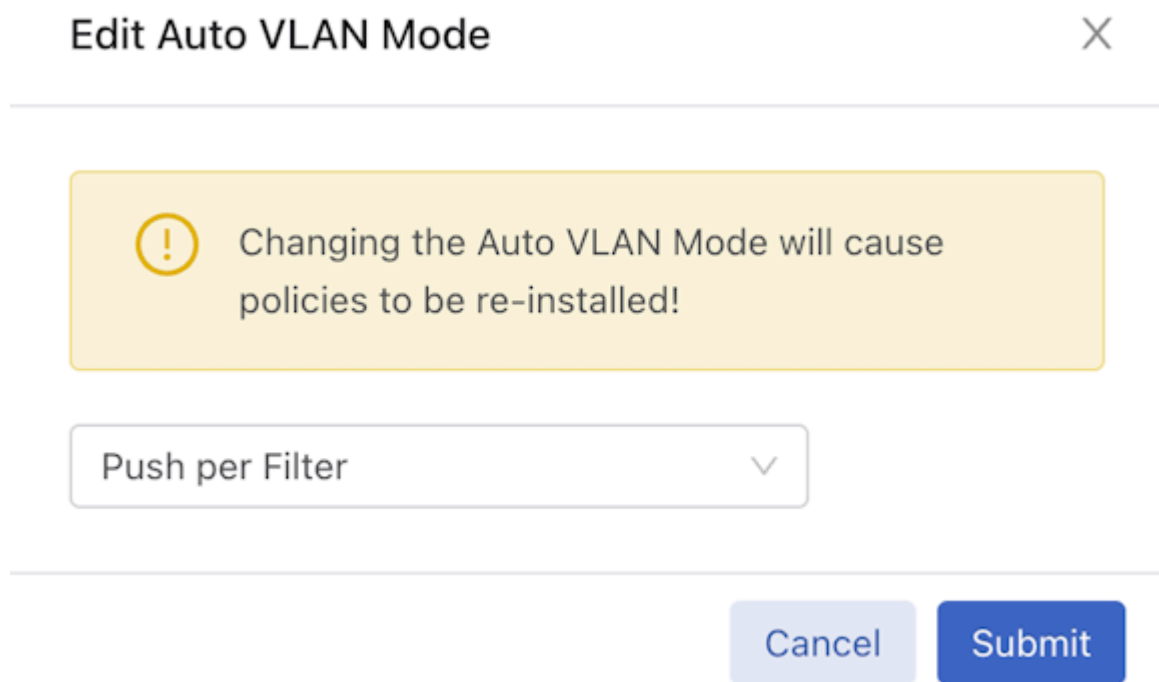
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 11-5: Auto VLAN Mode Config**



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

**Figure 11-6: 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 11-7: 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 11-8: Push Per Policy**



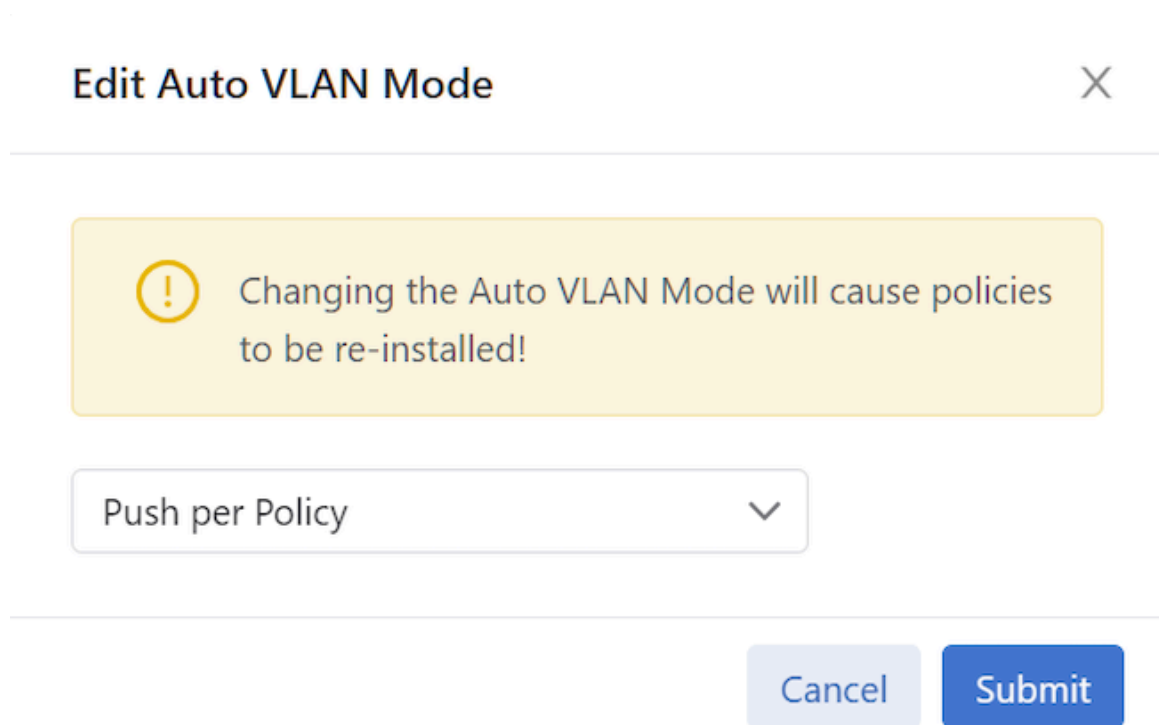
5. Alternatively, enter the desired mode name in the input area.

**Figure 11-9: Push Per Policy**



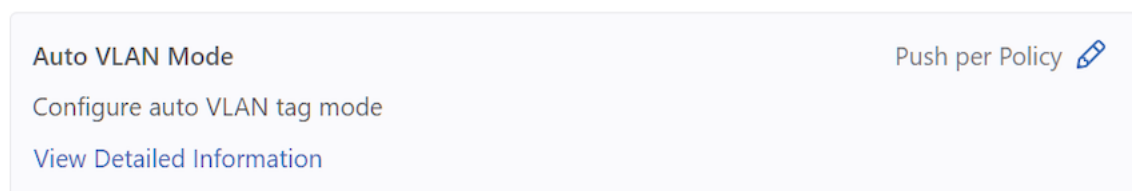
- Click the **Submit** button to confirm the configuration changes or the **Cancel** button to discard the changes.

**Figure 11-10: Submit Button**



- After successfully setting the configuration, the current configuration status displays next to the edit button.

**Figure 11-11: 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**

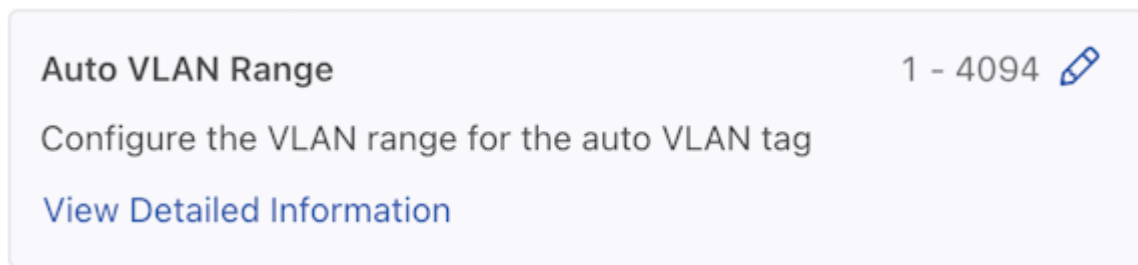
## 11.4 Auto VLAN Range

### 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 specified.

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

**Figure 11-12: 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 11-13: Edit Auto VLAN Range**

Edit Auto VLAN Range

×

### Configure Auto VLAN Range

!

Existing policies will be re-computed!

Default

Custom

Default Range is 1 - 4094

Cancel

Submit

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

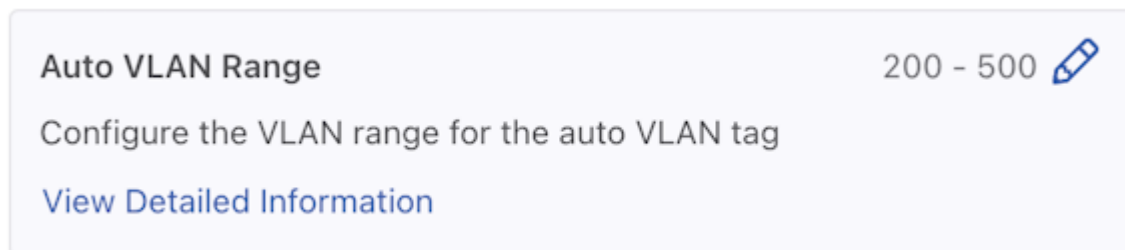
**Figure 11-14: Custom Button**



The image shows a user interface for configuring a custom range. At the top, there are two buttons: 'Default' and 'Custom'. The 'Custom' button is highlighted in blue. Below the buttons, there is a slider control. On the left side of the slider is a text input field containing the number '1'. On the right side is a text input field containing the number '4094'. A blue horizontal line with circular knobs at both ends represents the range between these two values.

- 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 11-15: Configuration Change Success**



The image shows a light blue notification box with rounded corners. At the top left, it says 'Auto VLAN Range'. At the top right, it displays '200 - 500' followed by a small blue pencil icon. Below this, the text reads 'Configure the VLAN range for the auto VLAN tag'. At the bottom, there is a blue link that says 'View Detailed Information'.

### Configuring Auto VLAN Range using the CLI

To set the Auto VLAN Range, use the following command:

```
auto-vlan-range vlan-min start vlan-max end
```

To set the Auto VLAN Range, 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
```

## 11.5 Auto VLAN Strip

The strip VLAN option removes the outer VLAN tag before forwarding the packet to a delivery interface. Only the outer tag is removed if the packet has two VLAN tags. 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 the 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. Preserves the 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 the DANZ Monitoring Fabric added using the strip-one-vlan option during delivery interface configuration.
  - Remove only the second (inner) tag. Preserves the VLAN (outer) tag added by DMF Fabric and removes 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 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 DMF adds it to enable 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 depending on whether the switch is a Trident 3-based.

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](#)

### 11.5.1 Auto VLAN Strip using the CLI

The strip VLAN option removes the outer VLAN tag before forwarding the packet to a delivery interface. Only the outer tag is removed if the packet has two VLAN tags. 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 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. Preserves the 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 the DANZ Monitoring Fabric added using the **strip-one-vlan** option during delivery interface configuration.
  - Remove only the second (inner) tag. Preserves the VLAN (outer) tag added by DMF and removes 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 DMF adds it as a result of enabling the following options:

- push-per-policy
- push-per-filter
- rewrite vlan under filter-interfaces

To view 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 view 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 : 13-14-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 to override the default strip option on a delivery interface use the below options:

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. 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 strip only the inner VLAN ID and preserve the outer VLAN ID that DMF added, 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 ID stripping is supported at both the global and delivery interface levels. The **rewrite-per-policy** and **rewrite-per-filter** options have been removed in **DMF Release 6.0** because the **push-per-policy** and **push-per-filter** options now support the related use cases.

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]
```

Use the option to leave all VLAN tags intact, remove the outermost tag, remove the second (inner) tag, or remove the outermost two tags, as required.

By default, VLAN stripping is enabled and the outer VLAN added by DMF is removed.

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 the 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 9: 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 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-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 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-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 4**

- VLAN Mode - ***Push per Filter***
- Filter interface on any 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 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 5

- VLAN Mode - **Push per Filter**
- Filter interface on any switch
- Delivery interface on any switch
- Global VLAN strip is **disabled**

**Table 13: Behavior of Traffic as It Egresses with Different Strip Options on a Delivery Interface**

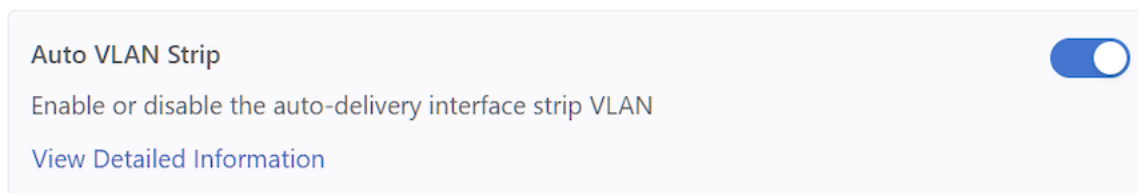
| VLAN tag type | No Configuration                                                                                                            | strip-no-VLAN                                                                                            | strip-one-VLAN                                                                            | 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 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.                                       |

## 11.5.2 Auto VLAN Strip using the GUI

### Auto VLAN Strip

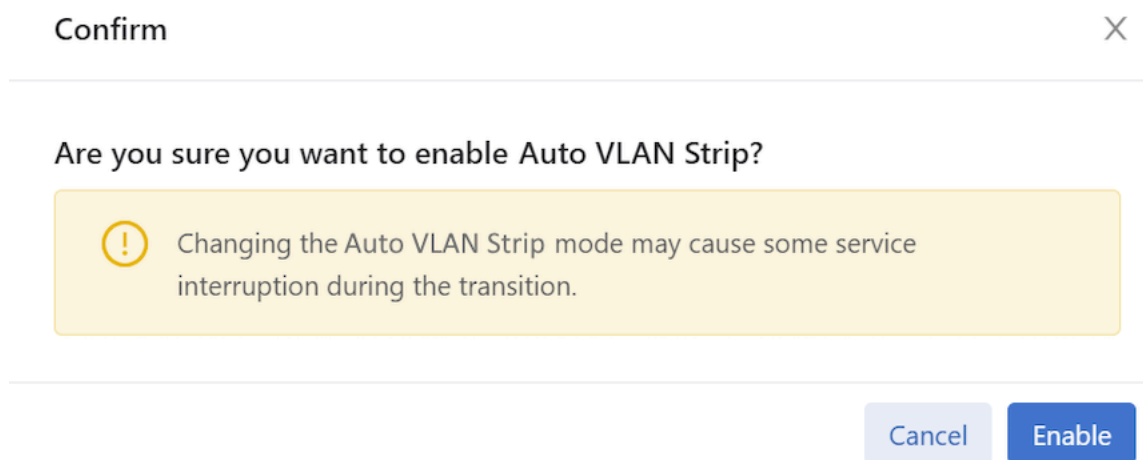
1. A **toggle button** controls the configuration of this feature. Locate the corresponding card and click the **toggle** switch.

**Figure 11-16: Toggle Switch**



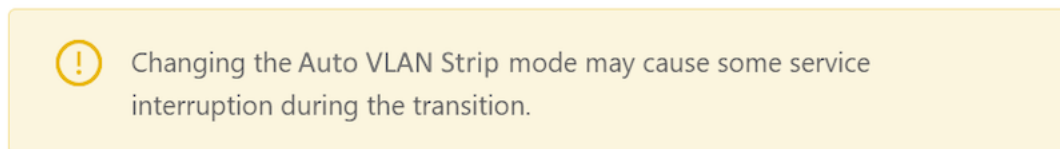
2. A confirm window pops up, displaying the corresponding prompt message. Click the **Enable** button to confirm the configuration changes or the **Cancel** button to cancel the configuration. Conversely, to disable the configuration, click **Disable**.

**Figure 11-17: Confirm / Enable**



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

**Figure 11-18: 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**

## 11.6 CRC Check

If the Switch CRC option is enabled, which is the default, each DMF switch drops incoming packets that enter the fabric with a CRC error. The switch generates a new CRC if the incoming packet was modified using an option that modifies the original CRC checksum, which includes the push VLAN, rewrite VLAN, strip VLAN, and L2 GRE tunnel options.



**Note:** Enable the Switch CRC option 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 helpful if packets with CRC errors need to be delivered to a destination tool unmodified for analysis. When disabling the Switch CRC option, ensure 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](#)

### 11.6.1 CRC Check using the CLI

If the Switch CRC option is enabled, which is the default, each DMF switch drops incoming packets that enter the fabric with a CRC error. The switch generates a new CRC if the incoming packet was modified using one option that modifies the original CRC checksum, which includes the push VLAN, rewrite VLAN, strip VLAN, and L2 GRE tunnel options.



**Note:** Enable the Switch CRC option 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 helpful if packets with CRC errors need to be delivered to a destination tool unmodified for analysis. When disabling the Switch CRC option, ensure 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
```

In the event the Switch CRC option is disabled, re-enable the Switch CRC option using 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 or 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.

## 11.6.2 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.

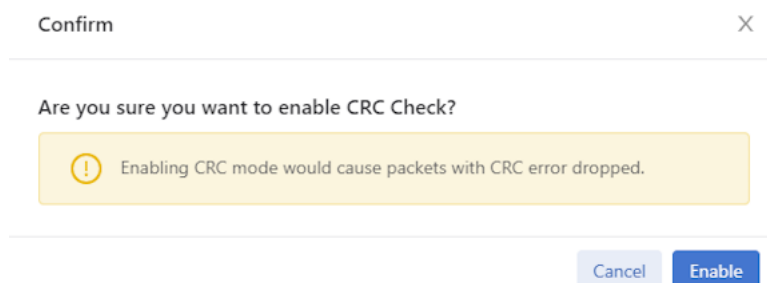
1. Select the **CRC Check** card.

**Figure 11-19: CRC Check Disabled**




2. Toggle the CRC Check to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 11-20: Enable CRC Check**



4. CRC Check is running.


**Figure 11-21: CRC Check Enabled**






5. To disable the feature, toggle the CRC Check to **Off**. Click **Disable** and confirm.

**Figure 11-22: Disable CRC Check**



The feature card updates with the status.

**Figure 11-23: CRC Check Disabled**



## 11.7 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 of 100 will be automatically assigned to non-admin users' policies.

Enable and disable Custom Priority using the steps described in the following topics.

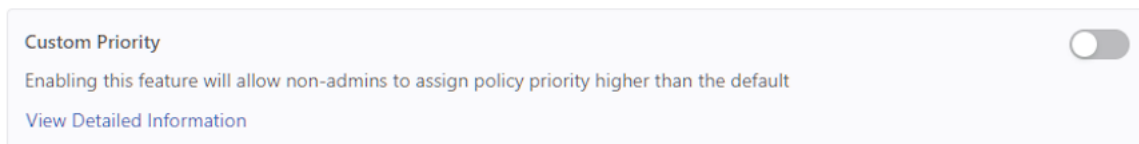
- [Custom Priority using the CLI](#)
- [Custom Priority using the GUI](#)

### 11.7.1 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.

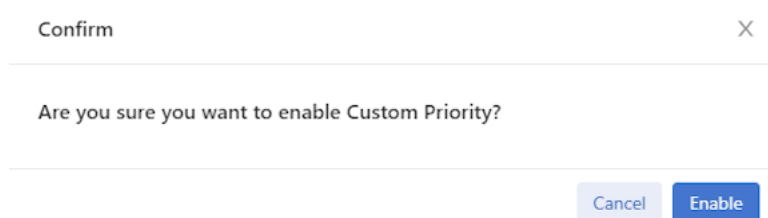
**Figure 11-24: Custom Priority Disabled**



2. Toggle the Custom Priority to **On**.

3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

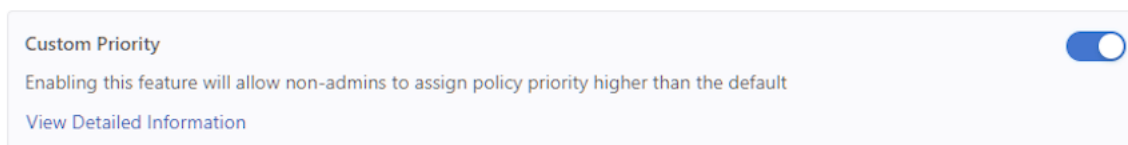
**Figure 11-25: Enable Custom Priority**



A confirmation dialog box with a title bar containing the word "Confirm" and a close button (X). The main text asks, "Are you sure you want to enable Custom Priority?". At the bottom, there are two buttons: "Cancel" and "Enable".

4. Custom Priority is running.

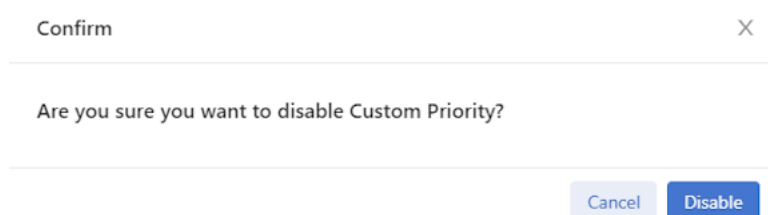
**Figure 11-26: Custom Priority Enabled**



A feature card for "Custom Priority". It has a toggle switch on the right that is turned on (blue). Below the title, it says "Enabling this feature will allow non-admins to assign policy priority higher than the default". There is a link "View Detailed Information" below that.

5. To disable the feature, toggle the **Custom Priority** to **Off**. Click **Disable** and confirm.

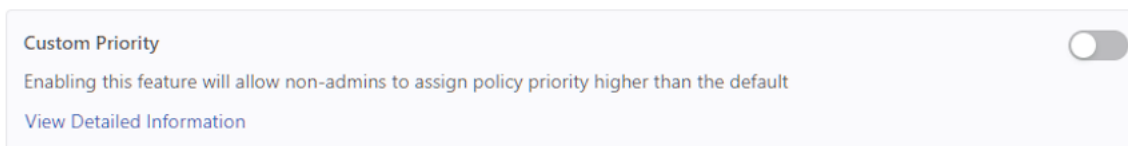
**Figure 11-27: Disable Custom Priority**



A confirmation dialog box with a title bar containing the word "Confirm" and a close button (X). The main text asks, "Are you sure you want to disable Custom Priority?". At the bottom, there are two buttons: "Cancel" and "Disable".

The feature card updates with the status.

**Figure 11-28: Custom Priority Disabled**



A feature card for "Custom Priority". It has a toggle switch on the right that is turned off (grey). Below the title, it says "Enabling this feature will allow non-admins to assign policy priority higher than the default". There is a link "View Detailed Information" below that.

## 11.7.2 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
```

## 11.8 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 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.

### 11.8.1 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 11-29: Device Deployment Mode - Auto Discovery**



- 2. Enter the edit mode using the **pencil icon**.

**Figure 11-30: Configure Device Deployment Mode**

Edit Device Deployment Mode

X

!

Changing device deployment mode requires modifying switch configuration.

Auto Discovery

▼

Cancel

Submit

- 3. Change the switching mode as required using the drop-down menu. The default mode is **Auto Discovery**.

**Figure 11-31: Device Deployment Mode Options**

Edit Device Deployment Mode

X

!

Changing device deployment mode requires modifying switch configuration.

Auto Discovery

🔍

Auto Discovery

Pre-Configured

Cancel

Submit

**Figure 11-31: Device Deployment Mode - Pre-Configured Option**

Edit Device Deployment Mode

X

!

Changing device deployment mode requires modifying switch configuration.

Pre-Configured

🔍

Auto Discovery

Pre-Configured

Cancel

Submit

- 4. Click **Submit** and confirm the operation when prompted.
- 5. The Device Deployment Mode status updates.

**Figure 11-33: Device Deployment Mode - Status Update**

Device Deployment Mode

Pre-Configured

Configure the mode for fabric switch installation

[View Detailed Information](#)

## 11.8.2 Configuring Device Deployment Mode using the CLI

Device Deployment Mode has two options: select the 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
```

## 11.9 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](#)

### 11.9.1 InPort Mask using the CLI

DANZ Monitoring Fabric implements multiple flow optimizations to reduce the number of flows 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 was 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 11-34: 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 11-35: 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 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. The option exists to 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
```

## 11.9.2 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.

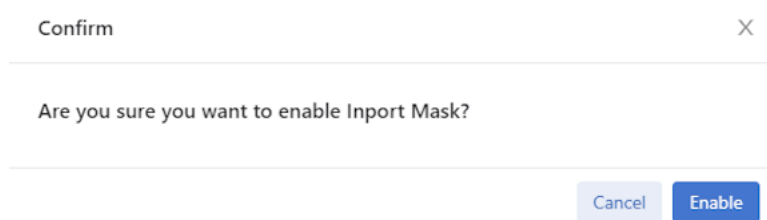
**Figure 11-36: Inport Mask Disabled**



2. Toggle the Inport Mask to **On**.


- Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 11-37: Enable Inport Mask**



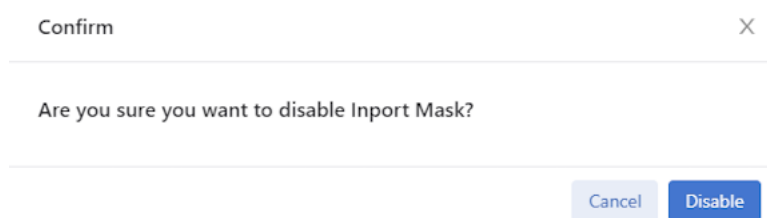
- Inport Mask is running.

**Figure 11-38: Inport Mask Enabled**




- To disable the feature, toggle the Inport Mask to **Off**. Click **Disable** and confirm.

**Figure 11-39: Disable Inport Mask**



The feature card updates with the status.

**Figure 11-40: Inport Mask Disabled**



## 11.10 Match Mode

Switches have finite hardware resources available for packet matching on aggregated traffic streams. This resource allocation is relatively static and configured in advance. The DANZ Monitoring Fabric supports three allocation schemes, 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 no policies use **src-mac** or **dst-mac**, the L3-L4 mode allows more 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. Switching between modes is possible, but any match rules incompatible with the new mode will fail.

### 11.10.1 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
```

### 11.10.2 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 11-41: L3-L4 Match Mode**



2. Enter the edit mode using the **pencil icon**.

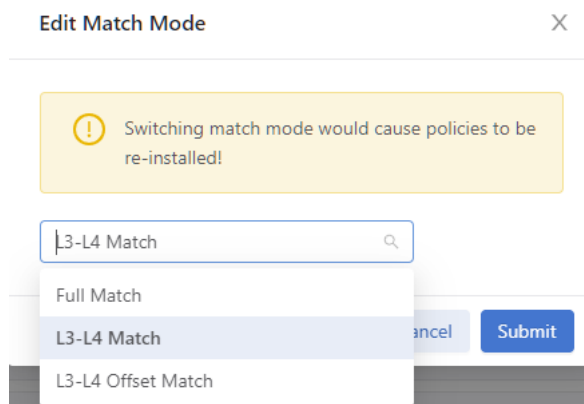
**Figure 11-42: Configure Switching Mode**





3. Change the switching mode as required using the drop-down menu. The default mode is **L3-L4 Match**.

**Figure 11-43: L3-L4 Match Options**



4. Click **Submit** and confirm the operation when prompted.



**Note:** An error message is displayed if the existing configuration of the monitoring fabric is incompatible with the specified switching mode.

## 11.11 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](#)

### 11.11.1 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:

- An 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 from 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
Inport 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
```

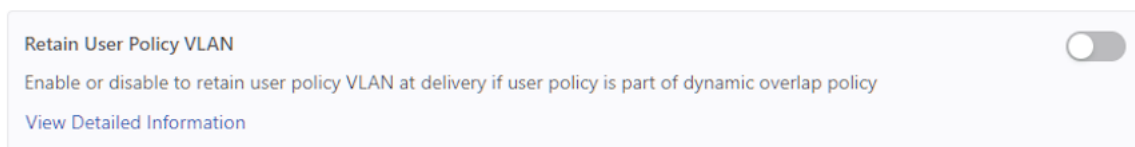
Use this feature with the **strip-second-vlan** option during delivery interface configuration to preserve the outer DMF fabric policy VLAN, strip the inner VLAN of traffic forwarded to a tool, or the **strip-no-vlan** option during delivery interface configuration.

### 11.11.2 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 11-44: Retain User Policy VLAN Disabled**



2. Toggle the Retain User Policy VLAN to **On**.

3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

**Figure 11-45: Enable Retain User Policy VLAN**

Confirm X

Are you sure you want to enable Retain User Policy VLAN?

⚠ Non-IP packets will be dropped at delivery.

Cancel Enable

4. Retain User Policy VLAN is running.

**Figure 11-46: Retain User Policy VLAN Enabled**

Retain User Policy VLAN ☒

Enable or disable to retain user policy VLAN at delivery if user policy is part of dynamic overlap policy

[View Detailed Information](#)

5. To disable the feature, toggle the Retain User Policy VLAN to **Off**. Click **Disable** and confirm.

**Figure 11-47: Disable Retain User Policy VLAN**

Confirm X

Are you sure you want to disable Retain User Policy VLAN?

⚠ Overlap policy VLAN will not be re-written with original policy VLAN.

Cancel Disable

The feature card updates with the status.

**Figure 11-48: Retain User Policy VLAN Disabled**

Retain User Policy VLAN ☐

Enable or disable to retain user policy VLAN at delivery if user policy is part of dynamic overlap policy

[View Detailed Information](#)

## 11.12 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](#)

### 11.12.1 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.

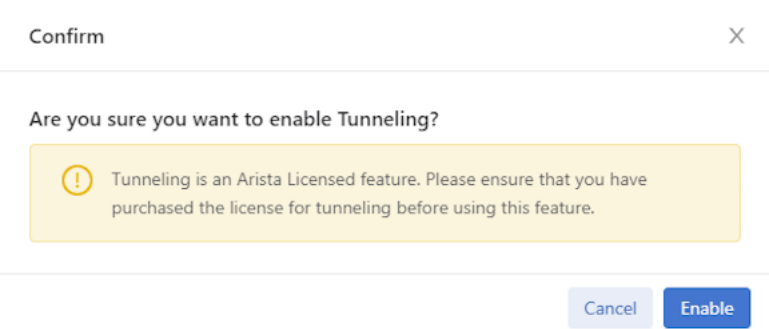
1. Select the **Tunneling** card.

Figure 11-49: Tunneling Disabled



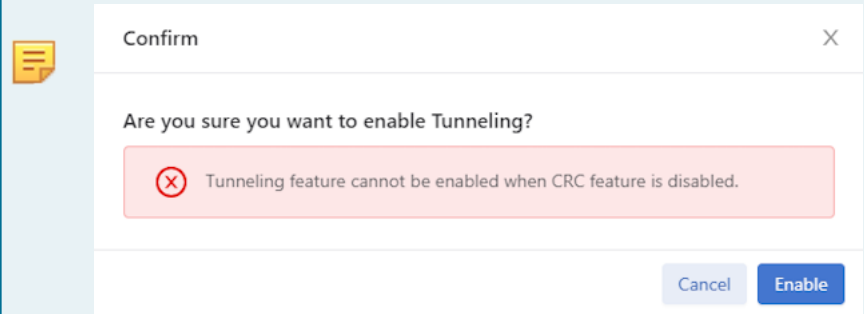
2. Toggle Tunneling to **On**.
3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

Figure 11-50: 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 11-51: CRC Check Warning Message



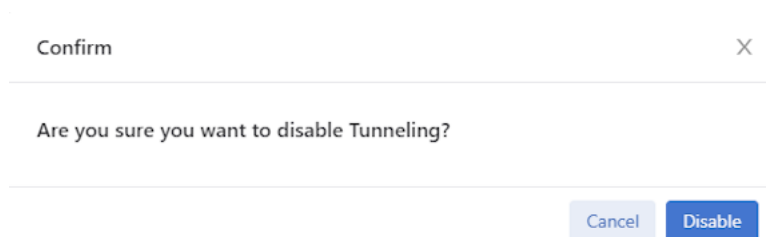
4. Tunneling VLAN is running.

Figure 11-52: Tunneling Enabled



- To disable the feature, toggle Tunneling to **Off**. Click **Disable** and confirm.

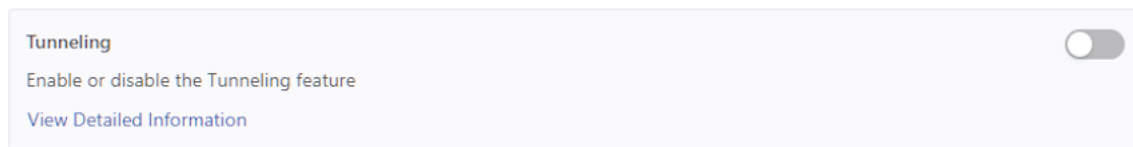
**Figure 11-53: Disable Tunneling**



A confirmation dialog box with a title bar containing the word "Confirm" and a close button (X). The main text asks, "Are you sure you want to disable Tunneling?". At the bottom, there are two buttons: "Cancel" and "Disable".

The feature card updates with the status.

**Figure 11-54: Tunneling VLAN Disabled**



A feature card for "Tunneling". It has a title "Tunneling" and a description "Enable or disable the Tunneling feature". On the right side, there is a toggle switch that is currently turned off. Below the description, there is a link that says "View Detailed Information".

### 11.12.2 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)#
```

## 11.13 VLAN Preservation

In DANZ Monitoring Fabric (DMF), metadata is appended to the packets forwarded by the fabric to a tool attached to a delivery interface. This metadata 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 IDs 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.

### 11.13.1 Using the CLI to Configure VLAN Preservation

Configure VLAN preservation 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 device name 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)
```

## 11.13.2 Using the GUI to Configure VLAN Preservation

VLAN preservation can be configured at global and local levels. 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 an MLAG domain (step 8).

### Global Configuration

1. To view or edit the global configuration, navigate to the DANZ Monitoring Fabric (DMF) Features page by clicking the gear icon in the top right of the navigation bar.

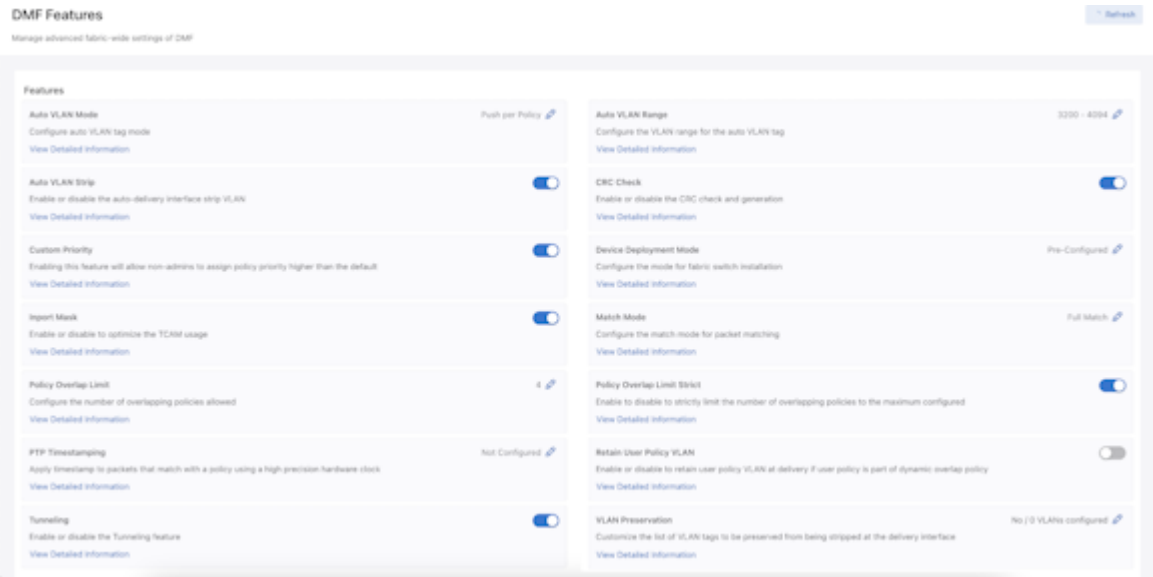
**Figure 11-55: DMF Menu Bar**



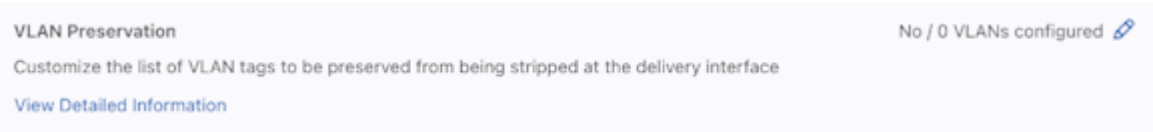
The DMF Feature allows for managing fabric-wide settings for DMF.

2. Scroll to the **VLAN Preservation** card.

**Figure 11-56: DMF Features Page**



**Figure 11-57: VLAN Preservation Card**





- Click the **Edit** button (pencil icon) to configure or modify the global VLAN Preservation feature settings.

**Figure 11-58: 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 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.
- Click the **Submit** button to save the configuration.

### 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 11-59: 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 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. VLAN Preservation 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 the **Fabric > MLAGs** page while configuring an MLAG Domain toggling the **VLAN Preservation** and **Preserve User Configured VLANs** switches to **on** (as required).

**Figure 11-60: Create MLAG Domain**

**Figure 11-61: MLAG VLAN Preservation & Preserve User Configured VLANs (expanded view)**

### 11.13.3 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.

- A partial policy installation may occur if any delivery interface fails to preserve the VLAN tag. This can happen when exceeding the 2000 VLAN ID/Port combination limit. Use the **show policy policy-name** command to obtain a detailed status, as shown in the following example:

```
(config)# show policy vlan-999
Policy Name          : vlan-999
Config Status        : active - forward
Runtime Status       : installed but partial failure
Detailed Status      : installed but partial failure -
                        Failed to preserve VLAN's on some/
all                  delivery interfaces, see warnings
for details
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 : 2
# of services : 0
# of pre service interfaces : 0
# of post service interfaces : 0
Push VLAN : 999
Post Match Filter Traffic : -
Total Delivery Rate : -
Total Pre Service Rate : -
Total Post Service Rate : -
Overlapping Policies : none
Component Policies : none
Installed Time : 2023-11-06 21:01:11 UTC
Installed Duration : 1 week
```

2. 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
```

3. 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)
```

4. 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
```

5. Verify if a switch has any associated preserve VLAN warnings among the fabric warnings:

```
(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.
```

6. The **show fabric warnings feature-unsupported-on-device** command provides information on whether VLAN preservation is configured on any unsupported devices:

```
(config-switch)# show fabric warnings feature-unsupported-on-device
# Name Warning
-|-----|
1 dell VLAN preservation feature is not supported on EOS switch eos
```

In the event of any preserve VLAN fabric warnings, please contact the [Arista Support Team](#) for assistance.

## 11.14 Reuse of Policy VLANs

Policies can reuse VLANs for policies in different switch islands. A switch island is an isolated fabric managed by a single pair of controllers; there is no data plane connection between fabrics in different switch islands. For example, with a single Controller pair managing six switches (**switch1**, **switch2**, **switch3**, **switch4**, **switch5**, and **switch6**) the option exists to 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 data plane connection between switches in the different switch islands.

There is no command needed to enable this feature. If the above condition is met, creating policies in each switch island with the same policy VLAN tag is supported.

In the condition mentioned above, 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 a data plane link connects two switch islands, it becomes one switch island. In that case, two policies cannot use the same policy vlan tag, so one of the policies (**P1** or **P2**) will become inactive.

## 11.15 Rewriting the VLAN ID for a Filter Interface

When sharing a destination tool with multiple filter interfaces, 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, assign a unique VLAN identifier to each filter interface. Ensure this VLAN ID is outside of the auto-VLAN range.



**Note:** In push-per-policy mode, enabling the rewrite VLAN feature on filter interfaces is impossible. When doing so, 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, enabled by default, rewrites the CRC after the frame has been modified so that a CRC error does not occur.



**Note:** Starting with **DMF Release 7.1.0**, simultaneously rewriting the VLAN ID and MAC address is supported and uses VLAN rewriting to isolate traffic while using MAC rewriting to forward traffic to specific VMs.

## 11.16 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 affiliated 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. When trying to reuse the same VLAN ID within an island, DMF generates a fabric error, and only the first Filter Interface (as sorted alphanumerically by DMF name) remains in use.

## Configuration

This feature requires no 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-vland
----- 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.
```

The following is a vertical representation of the CLI output above for illustrative purposes only.

```
>show fabric errors filter-interface-invalid-vlan
~~ Invalid Filter Interface VLAN(s) ~~
# 1
DMF Name filter1-f1
IF Name ethernet2
Switch filter1 (00:00:52:54:00:4b:c9:bc)
Rewrite VLAN 1
Details 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.

## 11.17 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 **1** to **4094**. To change the range, use the **auto-vlan-range** command.

The option exists to manually assign a VLAN not included in the 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
```

## 11.18 Tag-based Forwarding

The DANZ Monitoring Fabric (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 to 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. DMF uses the assigned VLAN tags to forward traffic to the correct delivery interface, saving TCAM space. This feature is handy when using switches based on the Tomahawk chipset because these switches have higher throughput but reduced TCAM space.

## 11.19 Policy Rule Optimization

### 11.19.1 Prefix Optimization

A policy can match with a large number of IPv4 or IPv6 addresses. These matches can be configured explicitly on each match rule, or the match rules can use an address group. With prefix optimization based on IPv4, IPv6, and TCP ports, DANZ Monitoring Fabric (DMF) uses efficient masking algorithms to minimize the number of flow entries in hardware.

**Example 1:** Optimize the 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:** In this case, if a generic prefix exists, 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 a generic prefix exists, 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

```



## 11.19.2 Transport Port Range and VLAN Range Optimization

The DANZ Monitoring Fabric (DMF) optimizes transport port ranges and VLAN ranges within a single match rule. Improvements in DMF now support cross-match rule optimization.

### Show Commands

To view the optimized match rule, use the show command:

```
show policy policy-name optimized-match
```

To view the configured match rules, use the following command:

```
show running-config policy policy-name
```

Consider the following DMF policy configuration.

```
show running-config policy p1
! policy
policy p1
action forward
delivery-interface d1
filter-interface f1
1 match ip vlan-id-range 1 4
2 match ip vlan-id-range 5 8
3 match ip vlan-id-range 7 16
4 match ip vlan-id-range 10 12
```

With the above policy configuration and before the DMF 8.5.0 release, the four match conditions would be optimized into the following TCAM rules:

```
show policy p1 optimized-match
Optimized Matches :
1 ether-type 2048 vlan 0 vlan-mask 4092
1 ether-type 2048 vlan 4 vlan-mask 4095
2 ether-type 2048 vlan 5 vlan-mask 4095
2 ether-type 2048 vlan 6 vlan-mask 4094
3 ether-type 2048 vlan 16 vlan-mask 4095
3 ether-type 2048 vlan 8 vlan-mask 4088
```

However, with the cross-match rule optimizations introduced in this release, the rules installed in the switch would further optimize TCAM usage, resulting in:

```
show policy p1 optimized-match
Optimized Matches :
1 ether-type 2048 vlan 0 vlan-mask 4080
1 ether-type 2048 vlan 16 vlan-mask 4095
```

A similar optimization technique applies to L4 ports in match conditions:

```
show running-config policy p1
! policy
policy p1
action forward
delivery-interface d1
filter-interface f1
1 match tcp range-src-port 1 4
2 match tcp range-src-port 5 8
3 match tcp range-src-port 7 16
```

```
4 match tcp range-src-port 9 14

show policy p1 optimized-match
Optimized Matches :
1 ether-type 2048 ip-proto 6 src-port 0 -16
1 ether-type 2048 ip-proto 6 src-port 16 -1
```

## 11.20 Switch Dual Management Port

### 11.20.1 Overview

When a DANZ Monitoring Fabric (DMF) switch disconnects from the Controller, the switch is taken out of the fabric, causing service interruptions. The dual management feature solves the problem by providing physical redundancy of the switch-to-controller management connection. DMF achieves this by allocating a switch data path port to be bonded with its existing management interface, 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 explicitly disabling the management configuration on the Controller.

Configure an interface for dual management using 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.

### 11.20.2 Configuring Dual Management Using the CLI

1. From config mode, specify the switch 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 submode, to configure the specified switch.

2. From **config-switch** mode, enter the **interface** command to specify the interface 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 submode, to 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)#
```

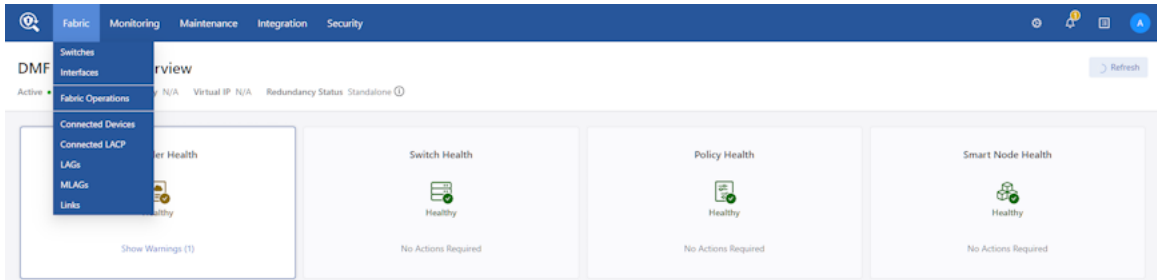


**Note:** When assigning an interface to a management role, no other interface-specific commands are honored for that interface (e.g., **shut-down**, **role**, **speed**, etc.).

## 11.20.3 Configuring Dual Management Using the GUI

1. Select **Fabric > Switches** from the main menu.

**Figure 11-62: Controller GUI Showing Fabric Menu List**



2. Click on the switch name to be configured with dual management.

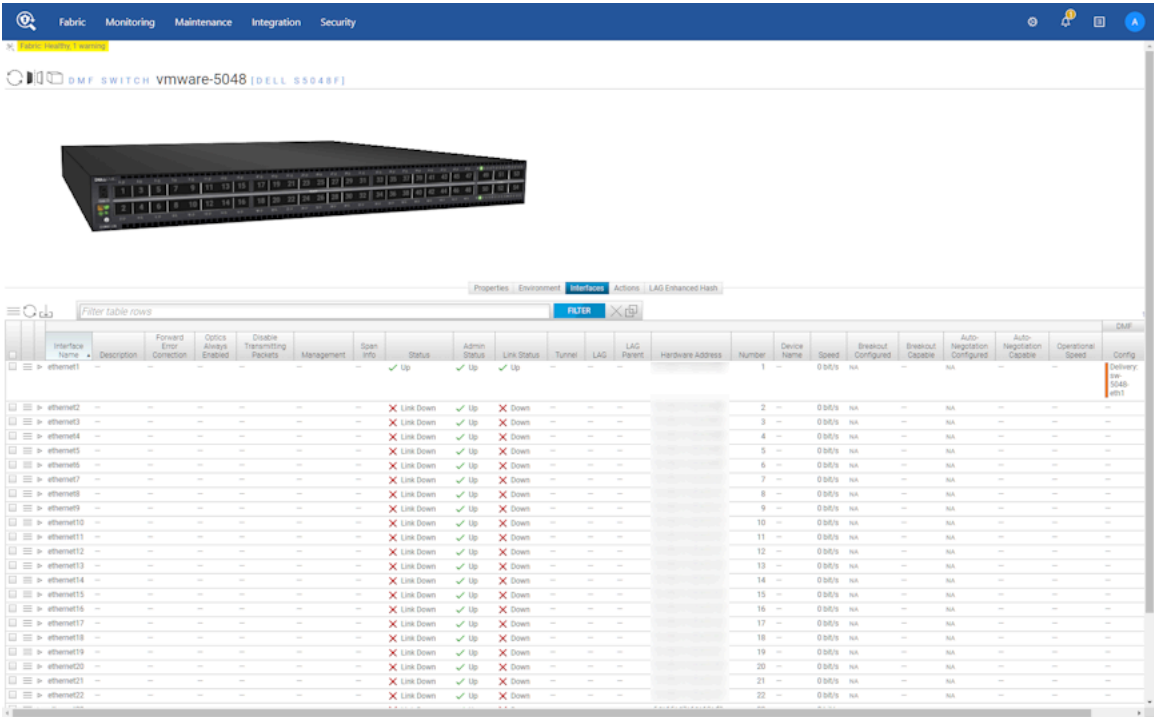
**Figure 11-63: Controller GUI Showing Inventory of Switches**

| Fabric Monitoring Maintenance Integration Security           |           |                   |              |                  |                                       |        |           |                         |                         |  |
|--------------------------------------------------------------|-----------|-------------------|--------------|------------------|---------------------------------------|--------|-----------|-------------------------|-------------------------|--|
| Fabric Summary: Healthy, Warnings: 1                         |           |                   |              |                  |                                       |        |           |                         |                         |  |
| Switches                                                     |           |                   |              |                  |                                       |        |           |                         |                         |  |
| Manage switches connected to DMF Controller                  |           |                   |              |                  |                                       |        |           |                         |                         |  |
| Switches IP Address Allocation TCAM Utilization and Capacity |           |                   |              |                  |                                       |        |           |                         |                         |  |
| Filter table rows                                            |           |                   |              |                  |                                       |        |           |                         |                         |  |
| Name                                                         | Connected | Quarantine Reason | Admin Status | Zero-touch State | Management Interface                  | Domain | Priority1 | PTP Source IPv4 Address | PTP Source IPv6 Address |  |
| vmware-5048                                                  | ✓         | —                 | ✓ Up         | Ok               | —                                     | —      | —         | —                       | —                       |  |
| vmware-7726                                                  | ✗         | —                 | ✓ Up         | —                | Prefer Dedicated Management Interface | —      | —         | —                       | —                       |  |
| vmware-s4048                                                 | ✓         | —                 | ✓ Up         | Reloading        | —                                     | —      | —         | —                       | —                       |  |
| vmware-z9100                                                 | ✓         | —                 | ✓ Up         | Ok               | —                                     | —      | —         | —                       | —                       |  |

Jan 3, 2024, 10:02:36pm UTC

3. Click on the **Interfaces** tab.

**Figure 11-64: Controller GUI Showing Switch Interfaces**

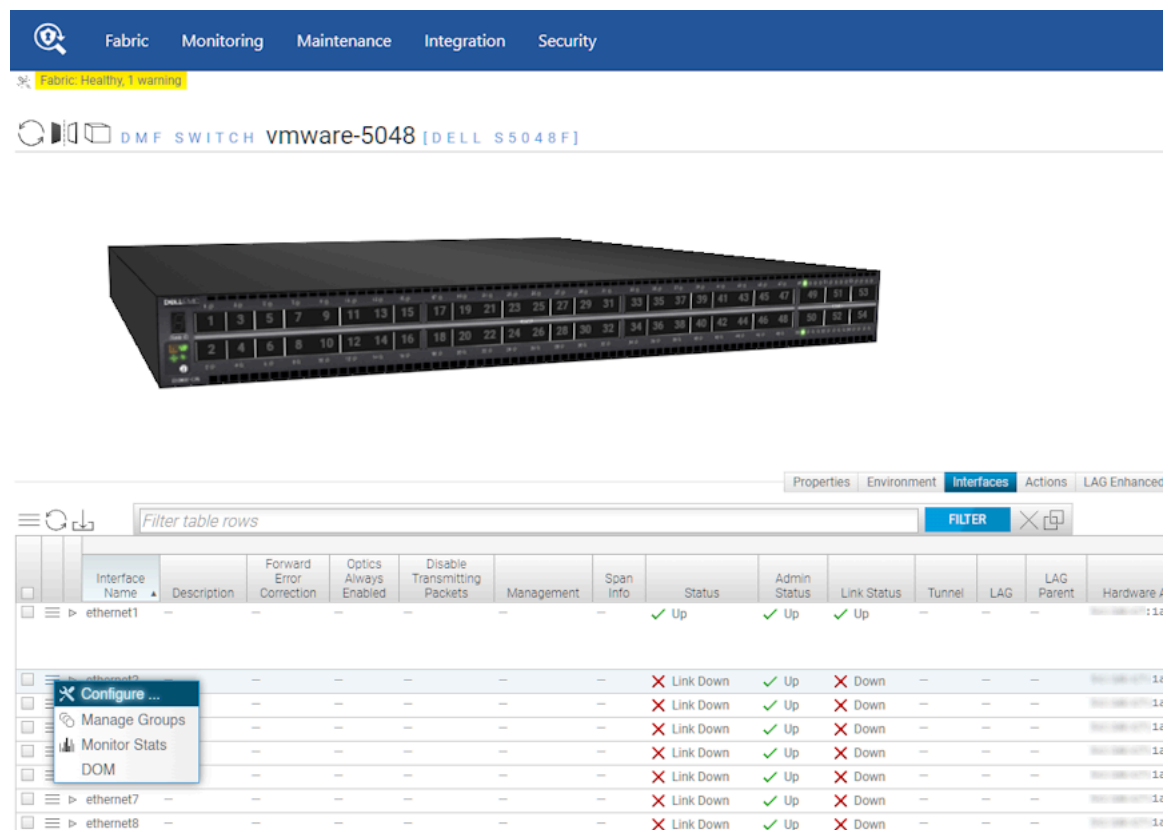


The screenshot shows the Controller GUI with the 'Switch Interfaces' tab selected. The interface displays a 3D model of a switch and a table of interface configurations. The table has columns for Interface Name, Description, Forward Error Correction, Gigabit Always Enabled, Disable Transmitting Packets, Management, Span Info, Status, Admin Status, Link Status, Tunnel, LAG, LAG Parent, Hardware Address, Number, Device Name, Speed, Breakout Configured, Breakout Capable, Auto-Negotiation Configured, Auto-Negotiation Capable, Operational Speed, and DMF.

| Interface Name | Description | Forward Error Correction | Gigabit Always Enabled | Disable Transmitting Packets | Management | Span Info | Status    | Admin Status | Link Status | Tunnel | LAG | LAG Parent | Hardware Address | Number | Device Name | Speed  | Breakout Configured | Breakout Capable | Auto-Negotiation Configured | Auto-Negotiation Capable | Operational Speed | DMF    |
|----------------|-------------|--------------------------|------------------------|------------------------------|------------|-----------|-----------|--------------|-------------|--------|-----|------------|------------------|--------|-------------|--------|---------------------|------------------|-----------------------------|--------------------------|-------------------|--------|
| ethernet1      |             |                          |                        |                              |            |           | Up        | Up           | Up          |        |     |            |                  | 1      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   | Config |
| ethernet2      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 2      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet3      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 3      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet4      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 4      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet5      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 5      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet6      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 6      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet7      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 7      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet8      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 8      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet9      |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 9      |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet10     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 10     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet11     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 11     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet12     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 12     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet13     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 13     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet14     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 14     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet15     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 15     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet16     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 16     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet17     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 17     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet18     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 18     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet19     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 19     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet20     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 20     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet21     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 21     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |
| ethernet22     |             |                          |                        |                              |            |           | Link Down | Up           | Down        |        |     |            |                  | 22     |             | 0.0G/s | NA                  |                  | NA                          |                          |                   |        |

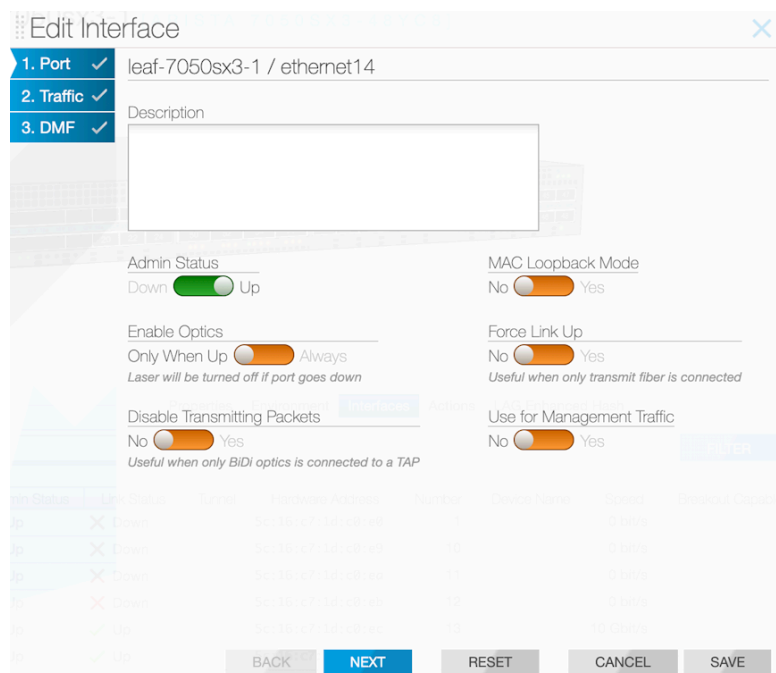
- Identify the interface to be configured as the standby management interface.

**Figure 11-65: Controller GUI Showing Configure Knob**



- Click on the **Menu** button to the left of the identified interface, then click on **Configure**.

**Figure 11-66: Controller GUI Showing Interface Settings**



6. Set **Use for Management Traffic** to **Yes**. This action configures the interface to the standby management role.

**Figure 11-67: Use for Management Traffic**

Admin Status: Down ☒ Up

MAC Loopback Mode: No ☒ Yes

Enable Optics: Only When Up ☒ Always

Force Link Up: No ☒ Yes

Disable Transmitting Packets: No ☒ Yes

Use for Management Traffic: No ☒ Yes

Buttons: Link Down, Link Up, BACK, NEXT, RESET, CANCEL, SAVE

7. Click **Save**.

### 11.20.4 Management Interface Selection Using the GUI

By default, the dedicated management interface serves as the management port, with the front panel data port acting as a 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 returns, it becomes the active management port.
- When the management network is undependable, this can lead to switch disconnects.

The **Management Interface** choice dictates what happens when the management interface returns after a failover. Make this selection using the GUI or the [CLI](#).

Select **Fabric > Switches**.

**Figure 11-68: Fabric Switches**

DMF: Fabric > Switches

Switch Health: Healthy

Policy Health: Healthy

Smart Node Health: Healthy

Buttons: Show Warnings (1), Refresh, No Actions Required

Click on the switch name to be configured.

**Figure 11-69: Switch Inventory**

The screenshot shows the 'Switches' section of the network management interface. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. Below the navigation bar, a status bar indicates 'Fabric Summary: Healthy, Warnings: 1'. The main heading is 'Switches', with a subtitle 'Manage switches connected to DMF Controller'. Below this, there are tabs for 'Switches', 'IP Address Allocation', and 'TCAM Utilization and Capacity'. The 'Switches' tab is active, displaying a table of switches. The table has columns for Name, Connected, Quarantine Reason, Admin Status, Zerotouch State, Management Interface, Domain, Priority1, PTP Source IPv4 Address, and Source IPv4 Address. The switches listed are vmware-5048, vmware-7726, vmware-s4048, and vmware-z9100. The vmware-7726 switch is marked as 'Down' with a red 'X' in the 'Connected' column and a note 'Prefer Dedicated Management Interface' in the 'Management Interface' column. The timestamp at the bottom left is 'Jan 3, 2024, 10:02:36pm UTC'.

| Name         | Connected | Quarantine Reason | Admin Status | Zerotouch State | Management Interface                  | Domain | Priority1 | PTP Source IPv4 Address | Source IPv4 Address |
|--------------|-----------|-------------------|--------------|-----------------|---------------------------------------|--------|-----------|-------------------------|---------------------|
| vmware-5048  | ✓         | —                 | ✓ Up         | Ok              | —                                     | —      | —         | —                       | —                   |
| vmware-7726  | ✗         | —                 | ✓ Up         | —               | Prefer Dedicated Management Interface | —      | —         | —                       | —                   |
| vmware-s4048 | ✓         | —                 | ✓ Up         | Reloading       | —                                     | —      | —         | —                       | —                   |
| vmware-z9100 | ✓         | —                 | ✓ Up         | Ok              | —                                     | —      | —         | —                       | —                   |

Select the **Actions** tab.

**Figure 11-70: Switch Actions**

The screenshot shows the 'Switch Actions' page for the switch 'vmware-z9100' (DELL Z9100). The top navigation bar is the same as in Figure 11-69. Below the navigation bar, a status bar indicates 'Fabric Summary: Healthy, 1 warning'. The main heading is 'DMF SWITCH vmware-z9100 [DELL Z9100]'. Below this, there is a large image of the switch hardware. At the bottom, there is a tabbed interface with tabs for 'Properties', 'Environment', 'Interfaces', 'Actions', and 'LAG Enhanced Hash'. The 'Actions' tab is active, displaying a list of actions: 'Configure Switch', 'Clear Switch Config', 'Create Tunnel', 'Manage Firmware', 'Provision Interfaces', 'Reset Connection', 'Shutdown', 'Reboot', and 'Beacon'.

Click the **Configure Switch** icon and choose the required **Management Interface** setting.

**Figure 11-71: Controller GUI Showing Dual Management Settings**

The screenshot shows the 'Configure Switch' window with a sidebar on the left containing a list of configuration tabs: 1. Info, 2. PTP, 3. Clock, 4. SNMP, 5. SNMP Traps, 6. Logging, 7. TACACS, 8. sFlow, and 9. LAG Enhanced Hash. The 'Info' tab is selected. The main content area includes the following fields and controls:

- Name \***: A text field containing 'vmware-z9100'.
- Description**: An empty text area.
- MAC Address**: A dropdown menu showing '2b:c0'. Below it is a note: '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.'
- Admin Status \***: A toggle switch labeled 'Down' and 'Up', currently set to 'Up'.
- Management Interface \***: A dropdown menu showing 'Prefer Dedicated Management Interface'. Below it is a detailed note: 'If the dedicated management interface is up, always prefer it for management. If management is currently using the front panel management interface and the dedicated management interface transitions from down to up, this means the interface used for management traffic will be changed to the dedicated management interface even if the front panel interface is still up.'
- Switch IP Allocation (IPAM) is disabled**: A blue banner.
- IP Assignment Type**: A dropdown menu.
- IP Address**: A text field showing '- IPv4 Address -' with a help icon.
- Tunnel Encapsulation Type \***: A dropdown menu showing 'GRE'.
- Change Password**: A toggle switch labeled 'No' and 'Yes', currently set to 'No'.

At the bottom right, there is an image of a network switch. At the bottom of the window, there are five buttons: BACK, NEXT, RESET, CANCEL, and SUBMIT.

If selecting **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, putting the front panel data port in an admin down state.

If selecting **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.

### 11.20.5 Management Interface Selection Using the CLI

By default, the dedicated management interface serves as the management port, with the front panel data port acting as a 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 returns, it becomes the active management port.



When the management network is undependable, this can lead to switch disconnects. The **management interface** selection choice dictates what happens when the management interface returns after a failover.

```
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 selecting **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, putting the front panel data port in an admin down state.

If selecting **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.

## 11.20.6 Switch Fabric Management Redundancy Status

To check the status of all switches configured with dual management as well as the interface that is being actively used for management, enter the following command in the CLI:

```
Controller-1# show switch all mgmt-stats
```

## 11.20.7 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 link and the data plane interface as the secondary link.
- The bandwidth of the data-plane interface is limited regardless of the physical speed of the interface. Arista Networks recommends immediate remediation when the **oma1** link fails.

## 11.21 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 a planned management switch upgrade. During a planned management switch upgrade, DANZ Monitoring Fabric (DMF) switches disconnect from the Controller, and DMF policies are reprogrammed, disrupting traffic forwarding to tools. Enabling this feature before starting a management switch upgrade will not disrupt the existing DMF policies when DMF switches disconnect from the 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 the 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 the 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 the running config to a snapshot will not include the **system control-plane-lockdown enable** command.
- The CLI prompt will start with the prefix **LOCKDOWN** when this feature is enabled.
- Link up/down and other events 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, traffic is hashed to other members automatically in Controller lockdown mode if one LAG member fails. 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)#
```

## 11.22 CPU Queue Stats and Debug Counters

Switch Light OS (SWL) switches can now report their CPU queue statistics and debug counters. To view these statistics, use the DANZ Monitoring Fabric (DMF) Controller CLI. DMF exports the statistics to any connected DMF Analytics Node.

The CPU queue statistics provide visibility into the different queues that the switch uses to prioritize packets needing to be processed by the CPU. Higher-priority traffic is assigned to higher-priority queues.

The SWL debug counters, while not strictly limited to packet processing, include information related to the Packet-In Multiplexing Unit (PIMU). The PIMU performs software-based rate limiting and acts as a second layer of protection for the CPU, allowing the switch to prioritize specific traffic.



**Note:** The feature runs on all SWL switches supported by DMF.

### Configuration

These statistics are collected automatically and do not require any additional configuration to enable.

To export statistics, configure a DMF Analytics Node. Please refer to the DMF User Guide for help configuring an Analytics Node.

## Show Commands

### Showing the CPU Queue Statistics

The following command shows the statistics for the CPU queues on a single switch.

```
controller-1> show switch FILTER-SWITCH-1 queue cpu
```

| # | Switch          | OF Port | Queue ID | Type      | Tx Packets | Tx Bytes  | Tx Drops | Usage                             |
|---|-----------------|---------|----------|-----------|------------|-----------|----------|-----------------------------------|
| 1 | FILTER-SWITCH-1 | local   | 0        | multicast | 830886     | 164100990 | 0        | lldp, l3-delivery-arp, tunnel-arp |
| 2 | FILTER-SWITCH-1 | local   | 1        | multicast | 0          | 0         | 0        | l3-filter-arp, analytics          |
| 3 | FILTER-SWITCH-1 | local   | 2        | multicast | 0          | 0         | 0        |                                   |
| 4 | FILTER-SWITCH-1 | local   | 3        | multicast | 0          | 0         | 0        |                                   |
| 5 | FILTER-SWITCH-1 | local   | 4        | multicast | 0          | 0         | 0        | sflow                             |
| 6 | FILTER-SWITCH-1 | local   | 5        | multicast | 0          | 0         | 0        |                                   |
| 7 | FILTER-SWITCH-1 | local   | 6        | multicast | 0          | 0         | 0        | l3-filter-icmp                    |

There are a few things to note about this output:

- The CPU's logical port is also known as the **local port**.
- The counter values shown are based on the last time the statistics were cleared.
- Different CPU queues may be used for various types of traffic. The **Usage** column displays the traffic that an individual queue is handling. Not every CPU queue is used.

The **details** token can be added to view more information. This includes the absolute (or raw) counter values, the last updated time, and the last cleared time.

### Showing the Debug Counters

The following command shows all of the debug counters for a single switch:

```
controller-1> show switch FILTER-SWITCH-1 debug-counters
```

| #  | Switch          | Name                            | Value   | Description                                 |
|----|-----------------|---------------------------------|---------|---------------------------------------------|
| 1  | FILTER-SWITCH-1 | arpra.total_in_packets          | 1183182 | Packet-ins recv'd by arpra                  |
| 2  | FILTER-SWITCH-1 | debug_counter.register          | 79      | Number of calls to debug_counter_register   |
| 3  | FILTER-SWITCH-1 | debug_counter.unregister        | 21      | Number of calls to debug_counter_unregister |
| 4  | FILTER-SWITCH-1 | pdua.total_pkt_in_cnt           | 1183182 | Packet-ins recv'd by pdua                   |
| 5  | FILTER-SWITCH-1 | pimu.hi.drop                    | 8       | Packets dropped                             |
| 6  | FILTER-SWITCH-1 | pimu.hi.forward                 | 1183182 | Packets forwarded                           |
| 7  | FILTER-SWITCH-1 | pimu.hi.invoke                  | 1183190 | Rate limiter invoked                        |
| 8  | FILTER-SWITCH-1 | sflowa.counter_request          | 9325983 | Counter requests polled by sflowa           |
| 9  | FILTER-SWITCH-1 | sflowa.packet_out               | 7883772 | Sflow datagrams sent by sflowa              |
| 10 | FILTER-SWITCH-1 | sflowa.port_features_update     | 22      | Port features updated by sflowa             |
| 11 | FILTER-SWITCH-1 | sflowa.port_status_notification | 428     | Port status notif's recv'd by sflowa        |

The counter values shown are based on the last time the statistics were cleared.

Add the **name** or the **ID** token and a debug counter name or ID to filter the output.

Add the **details** token to view more information. This includes the debug counter ID, the absolute (or raw) counter values, the last updated time, and the last cleared time.

## Clear Commands

### Clearing the Debug Counters

The following command will clear all of the debug counters for a single switch:

```
controller-1# clear statistics debug-counters
```

### Clearing all Statistics

To clear both the CPU queue stats and the debug counters for every switch, use the following command:

```
controller-1# clear statistics
```



**Note:** This command is not only limited to switches. It will clear any clearable statistics for every device.

## Analytics Export

The following statistics are automatically exported to a connected Analytics Node:

- CPU queue statistics for every switch.



**Note:** This does not include the statistics for queues associated with physical switch interfaces.

- The PIMU-related debug counters. These are debug counters whose name begins with `pimu`. No other debug counters are exported.

DMF exports these statistics once every minute.



**Note:** The exported CPU queue statistics will include port number -2, which refers to the switch CPU's logical port.

### 11.22.1 Troubleshooting

Use the details with the new show commands to provide more information about the statistics. This information includes timestamps showing statistics collection time and the last time the statistics were cleared.

Use the `redis-cli` command to query the Redis server on the Analytics Node from the Bash shell on the DMF Controller to view the statistics successfully exported to the Analytics Node.

The following command queries for the last ten exported debug counters:

```
redis-cli -h analytics-ip -p 6379 LRANGE switch-debug-counters -10 -1
```

Likewise, to query for the last ten exported CPU queue stats:

```
redis-cli -h analytics-ip -p 6379 LRANGE switch-queue-stats -10 -1
```

### 11.22.2 Limitations

- Only the CPU queue stats are exported to the Analytics Node. Physical interface queue stats are not exported.
- Only the PIMU-related debug counters are exported to the Analytics Node. No other debug counters are exported.
- Only SWL switches are currently supported. EOS switches are not supported.

## 11.23 Egress Filtering

Before the DANZ Monitoring Fabric (DMF) 8.6 release, DMF performed filtering at the filter port based on policy match rules. Thus, the system delivered the same traffic to all policy deliveries and tools associated with those delivery ports.

**Egress Filtering** introduces an option to send different traffic to each tool attached to the policy's delivery setting. It provides additional filtering at the delivery ports based on the egress filtering rules specified at the interface.

DMF supports egress filtering on the delivery and recorder node interfaces and supports configuring IPv4 and IPv6 rules on the same interface. Only packets with an IPv4 header are subject to the rules associated with the IPv4 token, while packets with an IPv6 header are only subject to the rules associated with the IPv6 token.

Egress Filtering applies to switches running SWL OS and does not apply to Extensible Operating System (EOS) switches.

## 11.23.1 Configuring Egress Filtering using the CLI

### CLI Configuration

The egress filtering feature is configurable at the interface level. To enable it, run the **egress-filtering** command from the **config-switch-if** submode.

```
(config)# switch DCS-7050SX3-48YC8
(config-switch)# interface ethernet18
(config-switch-if)# egress-filtering
(config-switch-if-egress-filtering)#
```

In the **config-switch-if-egress-filtering** submode, enter the rule's sequence number. The number represents the sequence in which the rules are applied. The lowest sequence number will have the highest priority.



**Tip:** Leave gaps between the sequence numbers so that new rules can be added in the middle later, if necessary.

```
(config-switch-if-egress-filtering)# 1
allow Forward traffic matching this rule
drop Drop traffic matching this rule
```

After the sequence number, specify the action of the rule. It can be either **drop** or **allow**.

```
(config-switch-if-egress-filtering)# 1 allow
any ipv4 ipv6
```

After specifying the action, enter the rule target traffic type like **IPv4**, **IPv6**, or **any**.

```
(config-switch-if-egress-filtering)# 1 allow
any ipv4 ipv6
```

### Any Traffic

The following illustrates a rule to allow all traffic on the interface, in this case, **ethernet18**.

```
(config-switch-if-egress-filtering)# 1 allow any
```

And a rule to drop all traffic on the interface, in this case, **ethernet18**.

```
(config-switch-if-egress-filtering)# 1 drop any
```

---

## IPv4 Traffic

To **allow** or **drop** all IPv4 traffic on an interface, use the following commands:

### Drop

```
(config-switch-if-egress-filtering)# 1 drop ipv4
(config-switch-if-egress-filtering-ipv4)#
```

### Allow

```
(config-switch-if-egress-filtering)# 1 allow ipv4
(config-switch-if-egress-filtering-ipv4)#
```

The following other options are available in the submode **config-switch-if-egress-filtering-ipv4**.

DMF supports the following qualifiers for IPv4 traffic filtering along with the IP address, port, and VLAN ranges.

```
(config-switch-if-egress-filtering-ipv4)#
dst-ip dst-port inner-vlan ip-proto outer-vlan-range src-ip-range
src-port-range
dst-ip-range dst-port-range inner-vlan-range outer-vlan
src-ip src-port
```

The following are examples of using the different options.

```
(config-switch-if-egress-filtering-ipv4)# dst-ip 12.123.123.39
(config-switch-if-egress-filtering-ipv4)# ip-proto 6
(config-switch-if-egress-filtering-ipv4)# dst-port 13
(config-switch-if-egress-filtering-ipv4)# src-port-range min 12 max 23
(config-switch-if-egress-filtering-ipv4)# inner-vlan 23
(config-switch-if-egress-filtering-ipv4)# outer-vlan 45
(config-switch-if-egress-filtering-ipv4)# src-ip 12.123.145.39
```

## IPv6 Traffic

To **allow** or **drop** all IPv6 traffic on an interface, use the following commands:

### Drop

```
(config-switch-if-egress-filtering)# 1 drop ipv6
(config-switch-if-egress-filtering-ipv6)#
```

### Allow

```
(config-switch-if-egress-filtering)# 1 allow ipv6
(config-switch-if-egress-filtering-ipv6)#
```

The following options are available in the submode **config-switch-if-egress-filtering-ipv6** for IPv6 traffic filtering.

```
(config-switch-if-egress-filtering-ipv6)#
dst-port inner-vlan ip-proto outer-vlan-range src-
port-range
dst-port-range inner-vlan-range outer-vlan src-port
```

## Show Commands

The following show command displays the Egress Filtering enabled device name, the information about the specified rule under the **Entry key** column, and the rule's action under the **Entry value** column. DMF uses the table to communicate the egress filtering rules with the device. The data in this table is primarily intended for debugging and communication purposes.

```
show switch all table egress-flow-1
Egress-flow-1 Device name Entry key
 Entry value

-|-----|-----|-----|
-|-----|-----|-----|
1 0 DCS-7050SX3-48YC8 Priority(1000), Port(13), EthType(2048),
 Ipv4Src(12.123.123.12) Name(__Rule1__), Data([0, 0, 0, 0]), NoDrop()
2 1 DCS-7050SX3-48YC8 Priority(0), Port(13)
 Name(__Rule0__), Data([0, 0, 0, 0]), Drop()
```

If any egress filtering warnings are present, they can be seen by running the **show fabric warnings egress-filtering-warnings** command. The output lists the switch name and the interface name on which an egress filtering warning is present, with a detailed message.

## Configuration Validation Messages

The following are examples of validation failure messages and their potential causes.

A validation exception occurs when configuring an egress filtering rule without specifying **EtherType**.

```
Validation failed: EtherType is mandatory for egress filtering rule
```

Similarly, there is another configuration validation for action, which is mandatory for egress filtering rules. Each rule can have a maximum of two ranges, and when exceeded, a validation failure occurs.

```
Validation failed: A rule cannot contain more than 2 configured ranges
```

DMF does not support configuring individual port values and port ranges of the same qualifier in the same rule, and configuring a source port and its range results in a validation failure.

```
Validation failed: Source port and its ranges are not supported together
```

A validation failure occurs if any specified ranges have a higher minimum value than the maximum value. For example, the specified inner VLAN minimum exceeds the maximum value. Similarly, validation failures may occur for ranges.

```
Validation failed: Inner VLAN min cannot be greater than inner VLAN max
```

The **ip-proto** setting is mandatory when specifying any port number, source, or destination. Specifying the source or destination port without **ip-proto** causes a validation failure.

```
Validation failed: IP protocol number is mandatory for source port
Validation failed: IP protocol number is mandatory for destination port
```

A validation failure occurs if any unsupported IP protocol number is a source or destination port specified without **ip-proto**. DMF only supports TCP(6), UDP(17), and SCTP(132) protocol numbers for port qualifiers.

```
Validation failed: IP protocol number <protocol number> is unsupported for
source port
```


Validation failed: IP protocol number <protocol number> is unsupported for destination port

### 11.23.2 Configuring Egress Filtering using the GUI

Perform the following steps to configure Egress Filtering.

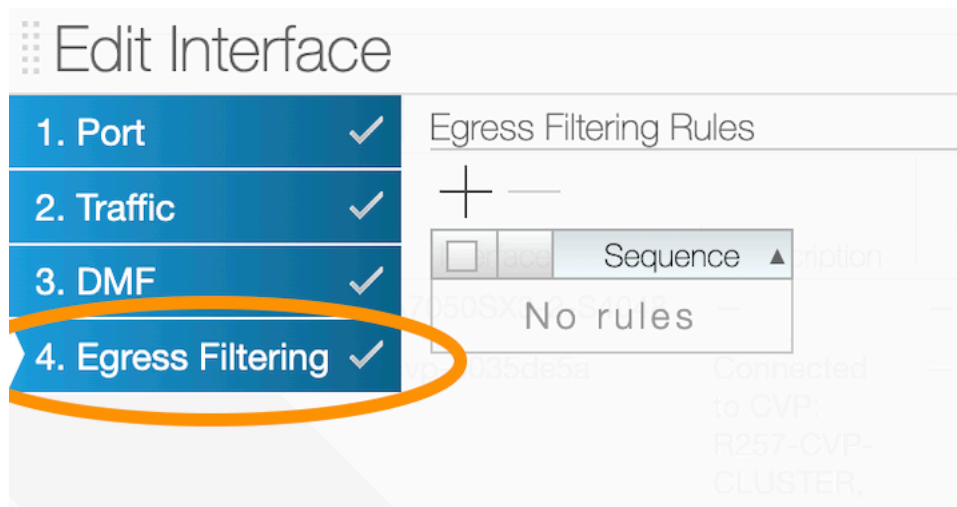
1. Navigate to the **Fabric > Interfaces** page.
2. Select **Configure ...** from the menu.

**Figure 11-72: Fabric > Interfaces > Configure**



3. Select **Egress Filtering** to add the egress filtering rules.

**Figure 11-73: Edit Interface - Egress Filtering**





4. To add a new rule, select the **+** icon to add a new **Egress Filtering Rule**.

**Figure 11-74: Add New Rule**

**Configure Egress Filtering Rule**

1. Traffic ✓

Sequence \*  
1

Action \*  
Allow

Ethertype \*  
Any

IP Protocol  
- Select By Name - - Dec -

BACK NEXT RESET CANCEL APPEND

Enter the primary information for the rule, including:

- **Sequence**
- **Action:** Allow (default), Drop
- **Ethertype:** Any (default), IPv4, IPv6
- **IP Protocol**

When **Ethertype** is set to **IPv4** or **IPv6**, the **Source**, **Destination** and **VLANs** steps are enabled.

**Figure 11-75: Configure Egress Filtering Rule**

**Configure Egress Filtering Rule**

1. Traffic ✓

2. Source ✓

3. Destination ✓

4. VLANs ✓

Sequence \*  
1

Action \*  
Allow

Ethertype \*  
IPv4

IP Protocol  
- Select By Name - - Dec -

BACK NEXT RESET CANCEL APPEND

In **Source**, enter the following information:

- **Source IP Address Type:** Single, Range
- **Source Ports:** None, Single, Range

Similarly, in **Destination**, enter the destination IP Address and ports.



**Note:** Specify the **IP Protocol** under **Traffic** when configuring source or destination ports.

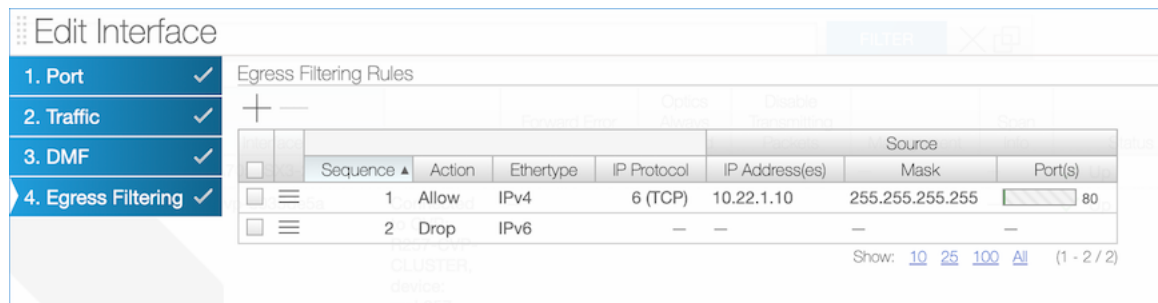
In **VLANs**, enter the following information:

- **Inner VLANs:** None, Single, Range
- **Outer VLANs:** None, Single, Range

5. Select **Append** to add the rule.

Use the **+** and **-** icons to add or remove the rules, as required.


**Figure 11-76: Edit Interface**



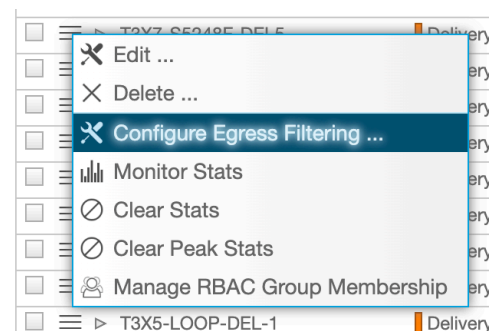
6. Select **Save** to complete the interface configuration.

The following UI menus are available to configure Egress Filtering:


- In **Fabric > Switches > Switch Detail > Interfaces**, select **Configure**.



- In **Monitoring > Interfaces**.



- In **Recorder Node Interfaces Monitoring > Recorder Nodes**.



### 11.23.3 Syslog Messages

There are no Syslog messages relevant to the Egress Filtering feature.

### 11.23.4 Troubleshooting

When a tool connected to a delivery interface configured with egress filtering rules receives an unexpected packet or does not receive the expected packet, use the following steps to troubleshoot the issue.

1. Check the **show running-config** command output to see if the egress filtering rules are configured correctly under that particular interface.
2. Verify the **show switch switch-name table egress-flow-1** command output.

It will display the port number of the interface for the configured egress filtering rules, its qualifiers as Entry key, the Entry value action of Drop or NoDrop, and a default drop rule for that port number with priority 0.

```
show sw all table egress-flow-1
Egress-flow-1 Device name Entry key
 Entry value
-----|-----|-----
1 0 DCS-7050SX3-48YC8 Priority(1000), Port(13), EthType(2048),
 Ipv4Src(12.123.123.12) Name(__Rule1__), Data([0, 0, 0, 0]), NoDrop()
2 1 DCS-7050SX3-48YC8 Priority(0), Port(13)
 Name(__Rule0__), Data([0, 0, 0, 0]), Drop()
```

3. Use the following command to verify the same information from a switch (e.g., DCS-7050SX3-48YC8).

```
root@DCS-7050SX3-48YC8:~# ofad-ctl gt egr_flow1
GENTABLE : egr_flow1
GENTABLE ID : 0x0019
Table count: matched/lookup : 0/0
Entry count/limit : 2/1024
guaranteed max: 512, potential max: 1024
priority 0 out_port 13 drop true 0p/0b eid 17
priority 1000 out_port 13 eth_type 0x800/0xffff ipv4_src 12.123.123.12
/255.255.255.255 drop false 0p/0b eid 22
```

4. Use the **show fabric warnings egress-filtering-warning** command to view any egress filtering warnings.

```
(config)# show fabric warnings egress-filtering-warning
~~~~~ Egress filtering warnings ~~~~~
~~~~~
Switch IF Name Warning message
```

| 1 | DCS-7050SX3-48YC8 | ethernet18 | Supported only on delivery or recorder node interfaces    |
|---|-------------------|------------|-----------------------------------------------------------|
| 2 | DCS-7280SR-48C6   | Ethernet8  | Egress filtering feature is not supported on EOS switches |

### 11.23.5 Limitations

- Egress filtering supports only 500 rules per interface. A validation failure occurs when exceeding this limit.

```
Validation failed: Only 500 egress filtering rules are supported per
interface
```

- DMF does not support egress filtering on MLAG delivery interfaces.
- IPv6 IP address filtering is not allowed, and a validation failure occurs.

```
Validation failed: IPv6 destination IP address is not supported
Validation failed: IPv6 source IP address is not supported
```

## Advanced Policy Configuration

This chapter describes advanced features and use cases for DANZ Monitoring Fabric (DMF) policies.

### 12.1 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. Keywords in a match rule not entered in the correct order results in the following message:

```
Error: Unexpected additional arguments ...
```

#### 12.1.1 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**, specify the mask in either CIDR notation (for example, /24) or dotted-decimal notation (**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 allows a 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**

DANZ Monitoring Fabric (DMF) supports matching on user-defined L3/L4 offsets instead of matching on these criteria. However, it is not possible to use both matching packet methods in the same DMF. Switching between these match modes may cause policies defined under the previous mode to fail.

Apply match rules 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-protocol> IP Protocol
```



**Warning:** Matching on untagged packets cannot be applied to DMF policies when in *push-per-policy* mode.

DMF uses a logical AND if a policy match rule has multiple fields. 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
```

DMF uses a logical OR when configuring two different match rules. For example, the following matches if the 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 configuring a match on TCP flags, the current GUI workflow also sets the hex value of the TCP flags for the mask attribute. When configuring 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.

## 12.1.2 Match-except Rules

The following summarizes match-except rules with examples which 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.
  - Example - Permit src-ip network, except ip-address:

```
1 match ip src-ip 172.16.0.0/16 except-src-ip 172.16.0.1
```

- Example - 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.
  - Example - 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
```

- Example - 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.

- Example - Incorrect match rule:

```
1 match icmp except-src-ip 192.168.1.10
```

- Example - 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.

- Example - Deny a subnet:

```
1 match ip dst-ip 172.16.0.0/16 except-dst-ip 172.16.0.0/18
```

- Example - Deny an IP Address:

```
1 match ip dst-ip 172.16.0.0/16 except-dst-ip 172.16.0.1
```

### 12.1.3 Matching with IPv6 Addresses

The value of the EtherType field determines whether the src-ip field to match is IPv4 or IPv6. The DANZ Monitoring Fabric (DMF) Controller displays an error if there is a mismatch between the EtherType and the IP address format.

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. Specify an address/mask or a group; DMF does not support ranges 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, write IPv6 address **1050:0000:0000:0000:0005:0600:300c:326b** as **1050:0:0:0:5:600:300c:326b**.
- Double Colon: Specify IPv6 addresses using double colons (::) instead of a series of zeros. For example, write IPv6 address **ff06:0:0:0:0:0:c3** as **ff06::c3**. Double colons may be used only once in an IP address.

DMF does not support the IPv4 address embedded in the IPv6 address format. For example, neither **0:0:0:0:0:0:101.45.75.219** nor **::101.45.75.219** can be used.

Both IPv4 and IPv6 masks must be in CIDR format. For example, **FFFF:FFFF:FFFF:FFFF:0:0:0:0** is valid in DMF, but **FFFF:0:0:FFFF:FFFF:0:0:0:0** is not a valid mask.

Both the colon-separated hexadecimal representation and the CIDR-style mask format are supported. The following example illustrates 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
```

## 12.1.4 Port and VLAN Range Matches

DANZ Monitoring Fabric (DMF) policy supports matching on source and destination port ranges with optimized hardware resource utilization. DMF uses efficient masking algorithms 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.

Provide the IP protocol information when using source and destination port ranges, fully 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.

Specify either **src-port-range** or **dst-port-range** or both in each match rule, as illustrated 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 supports matches for the VLAN ID range with optimized hardware resource utilization. Combining a VLAN ID range with a source or destination port range is supported, but not using all three ranges in a single match. The following example illustrates 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

```

### 12.1.5 User Defined Filters

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:** The DANZ Monitoring Fabric (DMF) Controller must be in push-per-policy mode for a user-defined filter to work accurately.

Selecting the **User Defined Offsets** option when the L3-L4 Offset Match switching mode is not enabled, the system displays a message to enable the correct match mode.

After enabling the L3-L4 Offset Match mode and selecting **Monitoring > User Defined Offsets**, DMF displays a table listing the currently defined offsets.



**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 the user can define the matching criteria. There are three options: 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 matching bytes, either 2 or 4 bytes.
- **Value:** The matching value of the specified length 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.

Configure each switch with a maximum of eight different offsets matching two bytes each, used in a single policy or 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

```

Enter the **show user-defined-offset** command to display the values configured in the user-defined-offset table.

```
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#
```

DMF supports user-defined filtering 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 SWL Switch | Trident 3 SWL Switch | EOS Switches  |
|-------------------------------|------------------------|----------------------|---------------|
| Total UDF Length              | 16 bytes               | 12 bytes             | 12 bytes      |
| Minimum Chunk Size            | 2 bytes                | 2 bytes              | 2 bytes       |
| Packet Start (Layer 2 Anchor) | 8 offsets              | 2 offsets            | 6 offsets     |
| Layer 3 Anchor                | 8 offsets              | 6 offsets            | 6 offsets     |
| Layer 4 Anchor                | 8 offsets              | 6 offsets            | 6 offsets     |
| Layer 2 Offset Range          | 0 - 126 bytes          | 0 - 62 bytes         | 0 - 126 bytes |
| Layer 3 Offset Range          | 0 - 114 bytes          | 0 - 112 bytes        | 0 - 114 bytes |
| Layer 4 Offset Range          | 0 - 96 bytes           | 0 - 112 bytes        | 0 - 96 bytes  |



**Note:** Please refer to the **DMF Hardware Compatibility List** for a complete list of supported switches and their corresponding Network ASIC types (Trident 3, Trident 2, etc).

## 12.2 Using the Filter and Delivery Role with MAC Loopback for a Two-stage Policy

Use the Filter and Delivery role 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).

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, simplifying monitoring deployment 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.


Switch hardware limits limit the total number of policies in the fabric. Using a Filter and Delivery role with a MAC loopback can also help eliminate 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. This command enables the physical interface without requiring a physical connection to the interface. Use a software loopback interface for copying traffic in any scenario where a physical loopback is required.

The user can also assign the Filter and Delivery role to a software loopback interface, which allows the use of a single interface for copying traffic to multiple destination interfaces. When assigning this role to an interface in loopback mode, 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 12-1: Using Software Loopback Interfaces to Avoid Overlapping Policies**



Use this configuration to copy different types of traffic from a single filter interface (**F1**) to four delivery interfaces (**D1** to **D4**). Assign the Filter and Delivery role to the software loopback interfaces (**LFD1** through **LFD4**) using just four physical interfaces. Physical loopbacks would require twice as many interfaces.

### Considerations

1. The SFP decides the Mac loopback speed. DMF uses the max port speed if there is no SFP (i.e., an empty port).
2. The port speed configuration (if any) will not impact the Mac loopback speed. The Mac loopback speed is set based on the SFP or the max port speed if there is no SFP.
3. The Rate-limit option limits the Mac loopback traffic at Rx side.



**Note:** When using a switch with the T2 chip, the 40G port Mac loopback is limited to the 10G speed.

## 12.2.1 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, perform the following steps:

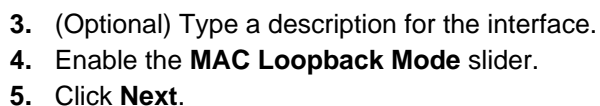
1. Display the available interfaces by selecting **Fabric > Interfaces**.  
The system displays the **Interfaces** page, which lists the interfaces connected to the DANZ Monitoring Fabric (DMF) fabric.

**Figure 12-2: Fabric Interfaces**

| Switch     | Switch DPID | Interface Name | Description | Forward Error Correction | Optics Always Enabled | Disable Transmitting Packets | Management | Spin Info | Status    | Tunnel | LAG | LAG Parent | Device Name    | Speed   | Breakout Configured | Auto-Negotiation Configured | Operational Speed | Config | Groups |
|------------|-------------|----------------|-------------|--------------------------|-----------------------|------------------------------|------------|-----------|-----------|--------|-----|------------|----------------|---------|---------------------|-----------------------------|-------------------|--------|--------|
| ethsws0048 |             | ethsws0048     |             |                          |                       |                              |            |           | Up        |        |     |            | DMF-SCALE-ESK1 | 10 Gb/s | NA                  | NA                          |                   |        |        |
| ethsws0100 |             | ethsws0100     |             |                          |                       |                              |            |           | Link Down |        |     |            | DMF-SCALE-ESK1 | 10 Gb/s | NA                  | NA                          |                   |        |        |
| ethsws0048 |             | ethsws0048     |             |                          |                       |                              |            |           | Up        |        |     |            | DMF-SCALE-ESK2 | 10 Gb/s | NA                  | NA                          |                   |        |        |
| ethsws0100 |             | ethsws0100     |             |                          |                       |                              |            |           | Link Down |        |     |            | DMF-SCALE-ESK2 | 10 Gb/s | NA                  | NA                          |                   |        |        |

2. Click the **Menu** control for the interface to use and select **Configure** from the pull-down menu.

**Figure 12-3: Fabric > Interfaces > Edit Interface > Port**



**Figure 12-4: Fabric > Interfaces > Edit Interface > Traffic**



6. (Optional) Configure **Rate Limiting**, if required, and click **Next**.

**Figure 12-5: Fabric > Interfaces > Edit Interface > DMF**

**Edit Interface**

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

1. Port ✓  
2. Traffic ✓  
3. DMF ✓

Auto VLAN Mode: push-per-policy  
Auto Strip VLAN: ✓ Enabled

DMF Mode:  
☐ None  
☐ Filter  
☒ Delivery  
☐ Filter and Delivery  
☐ Service  
☐ PTP

Interface Name \*  
TOOL-PORT-1

IP Address  
- IPv4 Address -  
For receiving IP datagram traffic

VLAN Stripping at Interface  
Default

Next Hop IP  
- IPv4 Address - / - Prefix -  
- Subnet Mask -  
Next hop IP that's reachable via delivery interface with local mask

ARP Interval  
5 seconds

VLAN Preservation  
No Yes

Preserve User Configured VLANs  
No Yes

VLAN Preserve  
+ -  
VLAN  
No VLANs

BACK NEXT RESET CANCEL SAVE

7. Enable the **Filter and Delivery** radio button.  
Optionally enable the **Rewrite VLAN** feature.



**Note:** The rewrite VLAN ID feature cannot be used with tunneling.

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

## 12.2.2 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$5niTlgPm$Jc24qOMF.hxNPI20DvnKaFzKYD6lIo59IMp3O4xIdwVTu2hx0s8Djpvz9xXAXXndiSkKe5jH.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 illustrates using five policies to implement this use case without creating overlapping policies. Otherwise, sixteen overlapping policies would be created without using the loopback interfaces to copy the traffic to separate filter interfaces.

```

! 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

```

Use the **show policy** command to verify the policy configuration.

## 12.3 Rate Limiting Traffic to Delivery Interfaces

The option exists to 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 section [Using the CLI to Rate Limit the Packets on a VXLAN Tunnel](#).

Use kbps to configure the rate-limit for the regular delivery interface. Arista Networks recommends configuring the rate limit in multiples of **64** kbps.

### 12.3.1 Rate Limiting Using the GUI

To use the GUI to set the rate limit for an interface, perform 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 12-6: Setting the Rate Limit for an Interface**

4. Enable the **Rate Limit** checkbox.
5. Use the number spinner to set the number of Kbps traffic limit.
6. Click **Save**.

### 12.3.2 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)#
```

Configure the rate limit for each member interface to rate limit a port channel. Configure individual rate limits for each member interface if the port channel has two member interfaces.

```
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, the Rate Limit column does not show the configured value for LAG and tunnel interfaces.

```
controller-1# show topology
----- Delivery Interface(s) -----
IF Switch IF Name State Speed Rate Limit

1 to-switch-scale-2 dnf-switch-2 ethernet18 up 10Gbps -
2 tunnel2switch dnf-switch-2 tunnel2switch up - -
3 tohostdelivery dnf-switch-1 ethernet11 up 10Gbps 2500Gbps
4 delivery-loop1 dnf-switch-1 ethernet15 up 10Gbps -
5 test-filter-delivery dnf-switch-1 ethernet18 up 10Gbps -
controller-1# show interface-names
----- Filter Interface(s) -----
IF Switch IF 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 vo-filter-17 DMP-FILTER-SWITCH-1 ethernet7 rx up 10Gbps 0 True anet-www-host1
4 vo-filter-18 DMP-FILTER-SWITCH-1 ethernet8 rx up 10Gbps 0 True anet-www-host2
5 vo-filter-19 DMP-FILTER-SWITCH-1 ethernet9 rx up 10Gbps 0 True anet-www-host3
----- Delivery Interface(s) -----
IF Switch IF Name Dir State Speed Rate Limit Strip Pwly Vlan Strip Customer Vlan IP Address ASD Required 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 BITTLOW-DELIVERY-PORT DMP-DELIVERY-SWITCH-1 ethernet4 tx up 10Gbps 100Mbps True False 172.43.75.1 True app-arista-5 Connected to Eth4 on DMP-DELIVERY-SWITCH /DMP DogFood
4 TOOL-PORT-01 DMP-DELIVERY-SWITCH-1 ethernet1 tx up 10Gbps - 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) -----
IF Switch IF Name Dir State Speed Connected Device

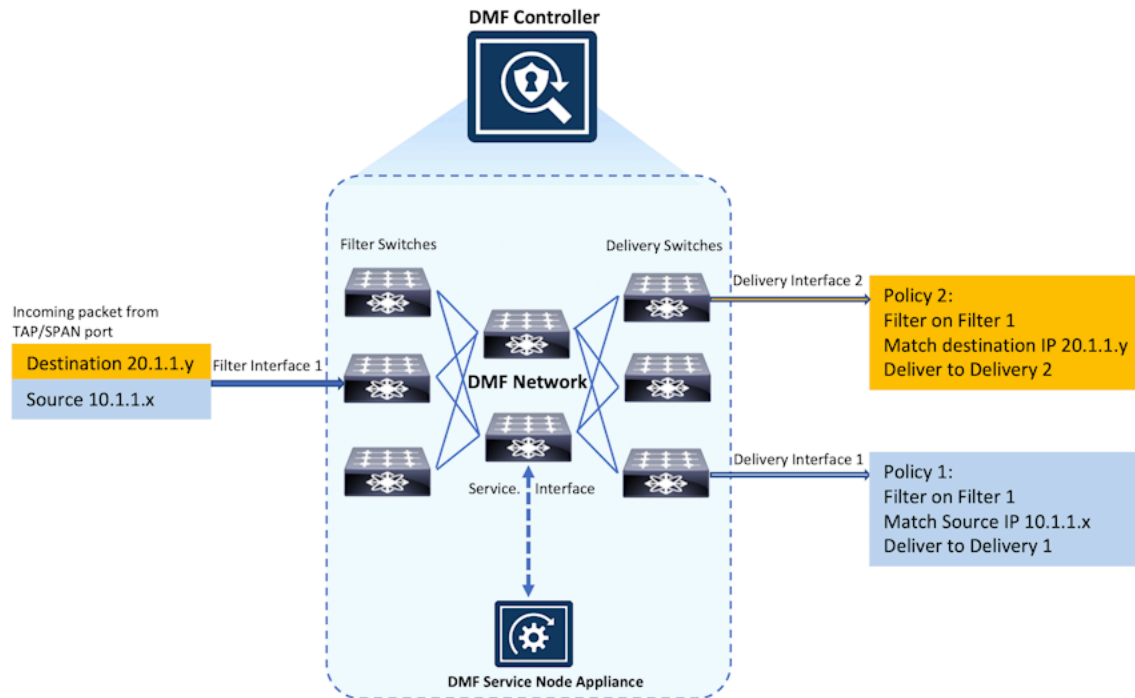
1 DMP-CORE-SWITCH-1-ethernet13:1-to-managed-service DMP-CORE-SWITCH-1 ethernet13:1 tx up 10Gbps DMP-DOCF000-EP-DN-1 1
2 DMP-CORE-SWITCH-1-ethernet13:1-to-managed-service DMP-CORE-SWITCH-1 ethernet13:1 rx up 10Gbps DMP-DOCF000-EP-DN-1 1
3 DMP-CORE-SWITCH-1-ethernet13:2-to-managed-service DMP-CORE-SWITCH-1 ethernet13:2 tx up 10Gbps DMP-DOCF000-EP-DN-1 2
4 DMP-CORE-SWITCH-1-ethernet13:2-to-managed-service DMP-CORE-SWITCH-1 ethernet13:2 rx up 10Gbps DMP-DOCF000-EP-DN-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#
```



## 12.4 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 12-7: 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 **20.1.1.5/24**, the packets are copied and forwarded to both **D1** and **D2**. Enabled by default, the DANZ Monitoring Fabric (DMF) policy overlap feature causes this behavior.

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 get 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.

The **overlap-limit-strict** command, enabled by default, strictly limits the number of overlapping policies to the maximum configured using the **overlap-policy-limit** command. For example, the operation fails with a validation error when setting the maximum number of overlapping policies to four (the default) and attempting to create a fifth policy using the same filter interface. To disable strict enforcement, use the **no overlap-limit-strict** command.



**Note:** The **overlap-strict-limit** command is disabled and must be manually enabled to enforce configurable policy limits.

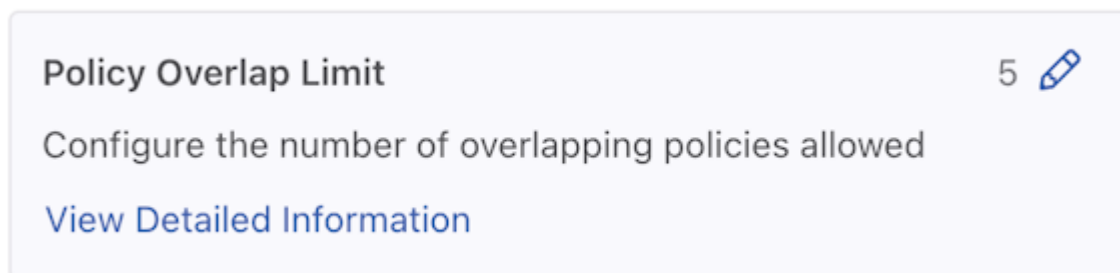
### 12.4.1 Configuring the Policy Overlap Limit Using the GUI

#### Policy Overlap Limit

Perform the following steps to configure the Policy Overlap Limit.

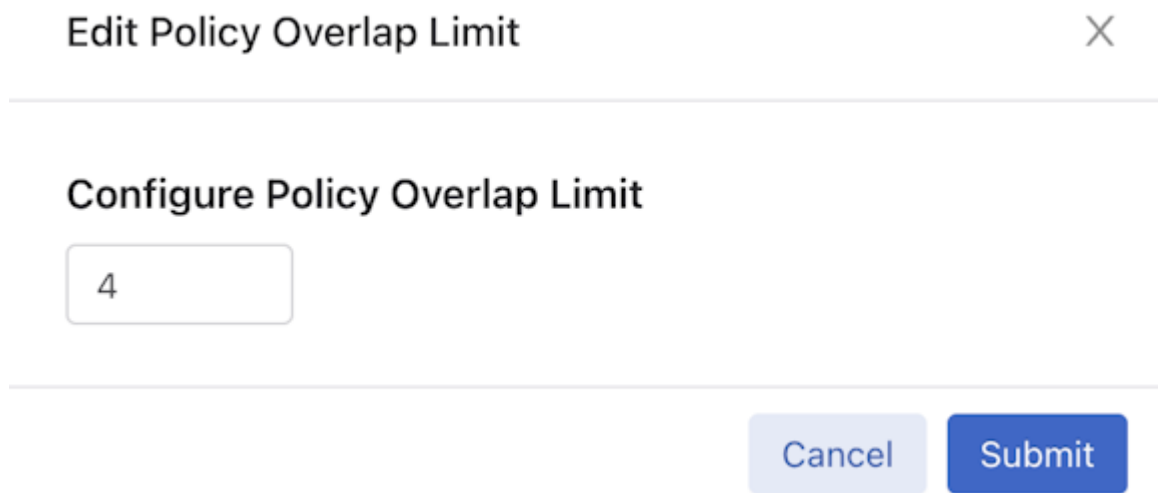
1. Control the configuration of this feature using the **Edit** icon by locating the corresponding card and clicking on the **pencil icon**.

**Figure 12-8: Policy Overlap Limit**



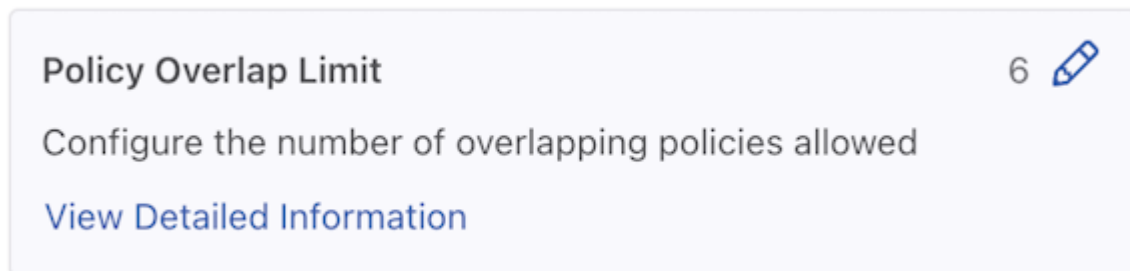
2. A configuration edit dialogue window pops up, displaying the corresponding prompt message. By default, the Policy Overlap Limit is 4.

**Figure 12-9: 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 12-10: Policy Overlap Limit Change Success**



## 12.4.2 Configuring the Overlapping Policy Limit Using the CLI

By default, the number of overlapping policies allowed is **four**. The maximum number to configure for overlapping policies is **ten**. Set the overlap policy limit to **zero** to disable the overlapping policy feature.

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:** When setting the Policy Overlap Limit to zero, ensure the policies do not overlap. If active policies overlap after disabling this feature, the forwarding result may be unpredictable.

## 12.4.3 Using the CLI to View Overlapping Policies

Enter the **show policy** command to view statistics for dynamic (overlapping) policies. If an overlapping policy appears 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

12.4.4 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 setting the maximum number of overlapping policies to 4 (the default) and users create a fifth policy using the same filter interface, the operation fails with a validation error. From the DANZ Monitoring Fabric (DMF) Features page, proceed to the Configuring the Policy Overlap Limit Strict feature card.

- 1. Select the **Policy Overlap Limit Strict** card.


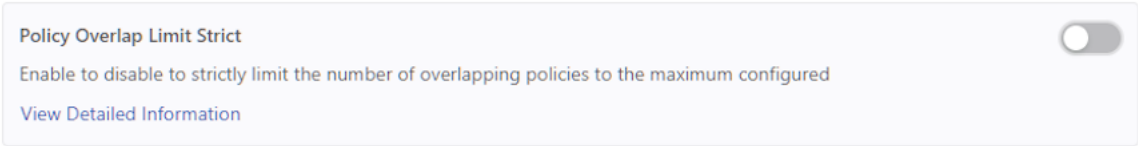
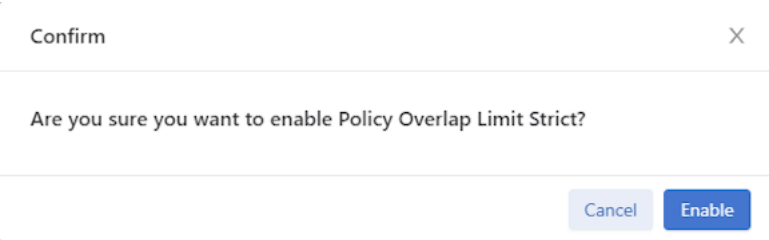
**Note:** The Policy Overlap Limit Strict option is **enabled** by default. The following steps guide if the Policy Overlap Limit Strict option is disabled.

Figure 12-11: Policy Overlap Limit Strict Disabled



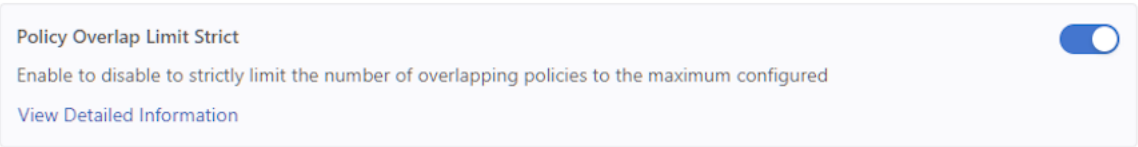
- 2. Toggle the Policy Overlap Limit Strict switch to **On**.
- 3. Confirm the activation by clicking **Enable** or **Cancel** to return to the DMF Features page.

Figure 12-12: Enable Policy Overlap Limit Strict



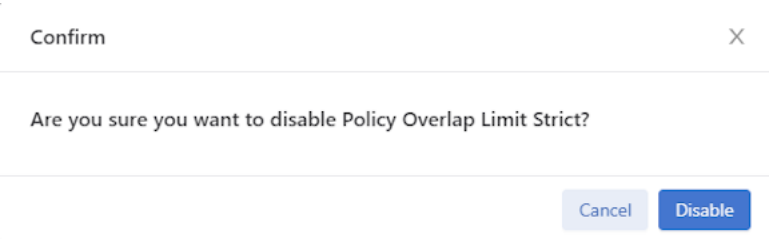
- 4. Retain Configuring the Policy Overlap Limit Strict is running.

Figure 12-13: Policy Overlap Limit Strict Enabled



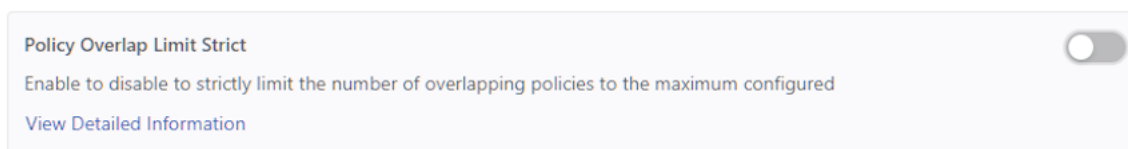
- 5. To disable the feature, toggle the Policy Overlap Limit Strict switch to **Off**. Click **Disable** and confirm.

Figure 12-14: Disable Policy Overlap Limit Strict



The feature card updates with the status.

**Figure 12-15: Policy Overlap Limit Strict Disabled**



### 12.4.5 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 setting the maximum number of overlapping policies 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
```

### 12.4.6 Exclude Inactive Policies from Overlap Limit Calculation

Previously, DANZ Monitoring Fabric (DMF) calculated the overlap policy limit by determining how many policies use the same filter interface, irrespective of whether the policies are active or inactive. By default, the overlap policy limit is 4, and the maximum is 10.

Suppose the limit is 4, and a user attempts to create a 5th policy using the same filter interface (**f1**). DMF throws the following error message: **Error: Validation failed: Filter interfaces used in more than 4 policies: f1.**

This count will include filter interfaces used in **active** or **inactive** policies.

The DMF policy overlap calculation excludes inactive policies when using the **inactive** command in policy configuration. For example, using the same policy limit settings described above, DMF supports creating a 5th policy using the same filter interface (**f1**) by first putting the 5th policy in an inactive state using the **inactive** command under policy configuration.



**Note:** This feature applies to switches running SWL OS and EOS.

#### Global Configuration Example

1. Select a switch and enter the config mode using the following command:

```
(config)# switch core1
```

2. Select an interface on the switch used as the filter-interface, as shown in the following example.

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

3. Create a filter interface, for example, **f1**, using the following command:

```
(config-switch-if)# role filter interface-name f1
```

4. Repeat the process to create the delivery interfaces.

```
(config-switch)# interface ethernet2
(config-switch-if)# role delivery interface-name d1
(config-switch)# interface ethernet3
(config-switch-if)# role delivery interface-name d2
(config-switch)# interface ethernet4
(config-switch-if)# role delivery interface-name d3
```

5. Set a `max overlap-policy-limit` value, for example, 2.

```
(config-switch)# overlap-policy-limit 2
```

6. Create overlap policies using the same `filter-interface`, in this example, *f1*.

```
(config-switch)# policy p1
(config-policy)# filter-interface f1
(config-policy)# delivery-interface d1
(config-policy)# action forward
(config-policy)# 1 match any

(config-switch)# policy p2
(config-policy)# filter-interface f1
(config-policy)# delivery-interface d2
(config-policy)# action forward
(config-policy)# 1 match any
```

7. Since the `overlap-policy-limit` value is **2**, the third overlap policy will not allow the use of the same filter interface **f1** in the third policy, **p3**. DMF throws a validation error.

```
(config-switch)# policy p3
(config-policy)# delivery-interface d3
(config-policy)# action forward
(config-policy)# filter-interface f1
Error: Validation failed: Filter interfaces used in more than 2 policies: f1
```

## Show Commands

The following command example displays the configured policies listing two overlap policies `p1` o `p2`.

```
(config-policy)# show policy
```

| # | Policy Name | Action  | Runtime Status               | Type       | Priority | Overlap | Priority | Push VLAN | Filter | BW | (truncated...) |
|---|-------------|---------|------------------------------|------------|----------|---------|----------|-----------|--------|----|----------------|
| 1 | p1_o_p2     | forward | A component policy failed    | Dynamic    | 100      | 1       |          | 3         | 10Gbps |    | (truncated...) |
| 2 | p1          | forward | all delivery interfaces down | Configured | 100      | 0       |          | 1         | 10Gbps |    | (truncated...) |
| 3 | p2          | forward | all delivery interfaces down | Configured | 100      | 0       |          | 2         | 10Gbps |    | (truncated...) |
| 4 | p3          | forward | inactive                     | Configured | 100      | 0       |          | 4         | -      |    | (truncated...) |

To add filter interface **f1** to the third overlap policy, **p3**, set the policy to **inactive**.

```
(config-switch)# policy p3
(config-policy)# inactive
(config-policy)# filter-interface f1
```

This results in an inactive policy **p3** being configured with filter interface **f1**. Use the **show running-config policy** command to view the status.

```
(config-policy)# show running-config policy
! policy
policy p1
action forward
delivery-interface d1
filter-interface f1
1 match any
```

```
policy p3
action forward
delivery-interface d3
filter-interface f1
inactive
1 match any
```

## 12.5 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.

### 12.5.1 Viewing Policy Flows

The **show policy-flow** command lists all the flows installed by the DANZ Monitoring Fabric (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 configured policy name. The packet and byte count is affiliated 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-----

```

### 12.5.2 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-l# show policy p1 drop
~ ~ ~ ~ ~
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 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 0
~ Core Interface(s) Drops ~
None.
~ Service Interface(s) Drops ~
None.
```



## 12.6 Using Rule Groups

DANZ Monitoring Fabric (DMF) supports using 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**.

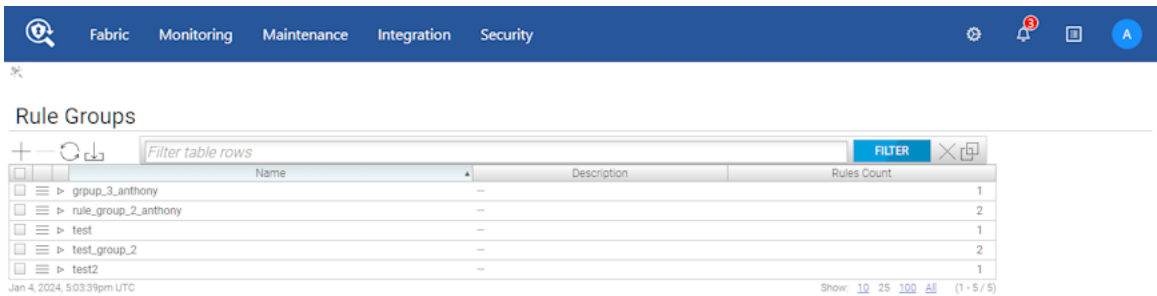
Identify only a single IP address group for a specific policy match rule. Address lists with both **src-ip** and **dst-ip** options cannot be used in the same match rule.

### 12.6.1 Using the GUI to Configure Rule Groups

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

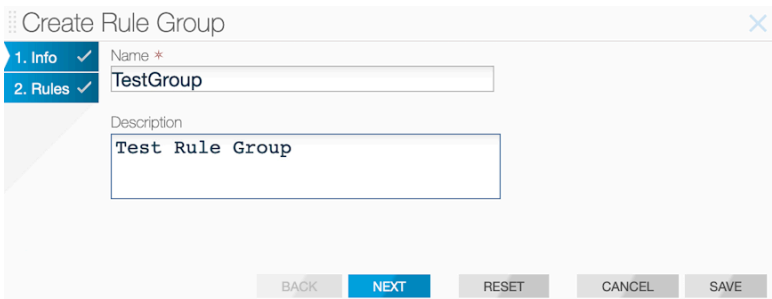
- 1. Select the **Monitoring > Rule Groups** option.

Figure 12-16: 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 a preferred name for the rule group and, optionally, a description.

Figure 12-17: Creating Rule Groups: Enter a Rule Group Name and Description



- 4. Click **NEXT** to add specific rules to the rule group.

5. In this pop-up section, add predefined rules by clicking on the options provided. In the example below, add a rule to match all IPv4 traffic by clicking on **IPv4**.

**Figure 12-18: Creating Rule Groups: Add a Predefined Rule to the Rule Group**

**Create Rule Group**

1. Info ✓

2. Rules ✓

Match Rule Shortcuts

Any ARP IPv4 ICMP SCTP TCP UDP IPv6 SCTP6 TCP6 UDP6  
 HTTP Source HTTP Dest SMTP Source SMTP Dest SMTP(SSL) Source  
 SMTP(SSL) Dest HTTPS Source HTTPS Dest SSH Source SSH Dest

Rules

+ -

| Sequence | Any Traffic | Ethertype   |
|----------|-------------|-------------|
| 1        | —           | 2048 (IPv4) |

Show: 10 25 100 All

BACK NEXT RESET CANCEL SAVE

6. As an alternative to the previous step, add custom rules by clicking the + sign under **Rules** and adding the necessary fields in the new pop-up screen.

**Figure 12-19: Creating Rule Groups: Add Custom Rules to the Rule Group**

**Configure Rule Group Rule**

1. Traffic ✓ A rule should include at least one populated field.

2. Source ✓

3. Destination ✓

4. Offset Match ✓

Sequence \* 2

☐ Match All Traffic

Ethertype

- Select By Name - - Decimal -

IP Protocol

- Select By Name - - Dec -

IP DSCP 0 - 63

IP Fragment Unspecified

VLAN(s) \*

Any

BACK NEXT RESET CANCEL APPEND

7. Complete the dialog that appears to assign a descriptive name to the rule group.
8. Add this rule group to DANZ Monitoring Fabric (DMF) policies as a match condition.

## 12.6.2 Using the CLI to Configure Interface Groups

The following example describes configuring 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 DANZ Monitoring Fabric (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 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 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
controller-1(config-filter-interface-group)#
```

## 12.7 PTP Timestamping

DANZ Monitoring Fabric (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.

- Connect a switch with a filter interface 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 configuring a policy to use timestamping, any packet matching on this policy will get its source MAC address rewritten with a timestamp value. The same holds true for any overlapping policy that carries traffic belonging to a user policy with timestamp enabled.
- The following options are available to configure a switch in boundary mode:
  - **domain:** Value for data plane PTP domain (0-255) (optional)
  - **Priority1:** Value of priority1 data plane 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)
- The 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). The default is 5 (optional)
  - **Sync Message Interval:** Set ptp sync message interval between messages (-7,3). The 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 a valid PTP config to get its packets timestamped.

### 12.7.1 Platform Compatibility

DANZ Monitoring Fabric (DMF) supports the timestamping feature 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) ... 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
2 S2 (00:00:cc:1a:a3:91:a7:6c) ...
2 S3 (00:00:cc:1a:a3:c0:94:3e) ...
```



**Note:** The CLI output example above is truncated for illustrative purposes. The actual output will differ.

### 12.7.2 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. The assumption is the network exists with a PTP domain and all necessary components.

### 12.7.3 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). If so, configure it can 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). If so, configure it 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
```

## 12.7.4 Configuring PTP Timestamping using the GUI

### Global Configuration

To view or edit the global PTP configuration, navigate to the DANZ Monitoring Fabric (DMF) Features page by clicking the gear icon.

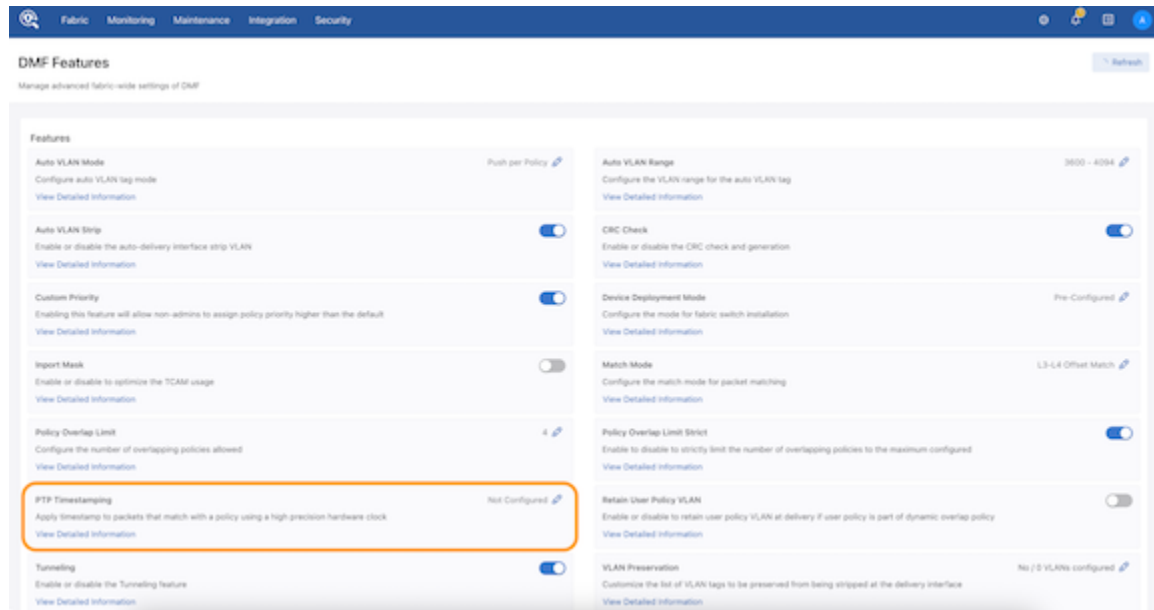
**Figure 12-20: DMF Menu Gear Icon**



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 12-21: DMF Features Page**



**Figure 12-22: 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 the global PTP configuration for a selected switch while configuring or editing a switch configuration (under the PTP step) using the **Monitoring > Switches** page.

**Figure 12-23: Configure Switch**

### PTP Interface Configuration

Configure a PTP Interface on the **Monitoring > Interfaces** page.

**Figure 12-24: Create Interface**

### 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 **PTP Timestamping** toggle to enable or disable timestamping.

**Figure 12-25: Create Timestamping Policy**

The screenshot shows the 'Create Policy' window. At the top right are icons for zooming and closing. The 'Policy Details' section includes a toggle for 'Enable Advanced Options' and a 'Collapse' link. Below this are input fields for Name, Description, Action (a dropdown menu showing 'Forward'), Push VLAN (a field with '##'), Priority (a field with '100'), and an 'Active' toggle switch. The 'Scheduling' section has dropdowns for 'Automatic', 'Always', and a 'PTP Timestamping' toggle switch. The 'Port Selection' section has a 'Match Traffic' dropdown menu showing 'No Switch Confi...'. Below this are three columns: 'Traffic Sources', 'Services', and 'Destination Tools', each with a 'Create' link and an 'Add Port(s)' button. At the bottom, a message states 'Match Traffic: Allow All Traffic (Default) has been selected'. There are 'Cancel' and 'Create Policy' buttons at the bottom right.

## 12.7.5 PTP Interface Configuration

A switch that syncs its hardware clock using PTP requires a physical front panel interface 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
```

A 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. Configure the master switch to communicate PTP messages with or without a vlan tag. Use the trunk-mode with the appropriate ptp vlan when configuring the neighbor similarly. If the neighbor's interface is in switch-port access mode or routed mode, use either of these to match it on the filter switch.

Other fields are optional, using default values 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, 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.

## 12.7.6 Policy Configuration for Timestamping

DANZ Monitoring Fabric (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
```

## 12.7.7 L2GRE Encapsulation of Packets with Arista Timestamp Headers

L2GRE encapsulation of packets with Arista timestamp headers is an extension of an existing feature allowing DANZ Monitoring Fabric (DMF) to use intra-fabric L2GRE tunnels.

These tunnels enable forwarding unmodified production network packets over intermediate L3 networks used by DMF, which can now forward packets with Arista Networks Timestamp headers across L2GRE tunnels defined on EOS switches.

When a PTP header-based timestamping capable filter switch is in a remote location, and the remote filter switches connect to the centralized tool farm via L2GRE tunnels, timestamp the packets using PTP timestamping. These packets are encapsulated in an L2GRE header and sent to the core switch via the L2GRE tunnel. The timestamped packets will properly decapsulate at the remote end and be forwarded to the destination tools.



**Note:** DMF only supports PTP timestamping with L2GRE encapsulation with the header-based timestamping feature.

Please refer to the [Tunneling Between Data Centers](#) section on using L2GRE tunnels in DMF.

---

## 12.7.8 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 the 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 view the timestamping status for a given 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 p1 unspecified action inactive Configured 100 0
1 1 - - - True

```

---

## 12.7.9 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
```

## 12.7.10 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. An unconfigured PTP interface or the interface is inactive.
3. The PTP switch configuration or PTP interface configuration is invalid or incomplete.
4. Configuring the PTP interface on a logical port (Lag or 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.

## 12.7.11 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 unrecoverable 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 DANZ Monitoring Fabric (DMF) Recorder Node accepts Stenographer queries using a syntax based on the Berkeley Packet Filter (BPF) syntax. When entering a malformed BPF string, the recorder node will respond with an error. The entire BPF grammar is not supported, but query strings can be composed using the predicates in the following table.

Table 14: 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:** Arista Networks recommends always using a specific time range 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 for 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 for 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 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 to the DANZ Monitoring Fabric (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 connected with the Recorder. The DANZ Monitoring Fabric (DMF) Recorder Node verifies at the DMF Controller if the provided username and password are valid and have sufficient privileges to use the Recorder REST API.
- Authentication tokens: The DMF Recorder Node REST API accepts revocable authentication tokens 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 can 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 the [Stenographer Reference for DMF Recorder Node](#). The DMF Recorder Node provides a REST API so clients can look up packets and metadata. The REST server runs securely (HTTPS) on TCP port **443**.

### B.1 Authentication

Clients must either authenticate using a valid DANZ Monitoring Fabric (DMF) Controller username and password over HTTP Basic or with an authentication token configured on the DMF controller specifically for DMF Recorder Node REST API authentication.

#### B.1.1 Basic HTTP Authentication

Use a valid DANZ Monitoring Fabric (DMF) Controller username and password 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 authorizes the user for the invoked recorder node REST endpoint.

In the following example, the query uses 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

An authentication token is primarily designed for third-party applications and automation scripts where creating an account or storing a username and password is undesirable. This method can also allow a DANZ Monitoring Fabric (DMF) Recorder Node access if the management network connection to the Controller is disrupted.

To create an authentication token, log in to the DMF Controller associated with the recorder node, then perform 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 (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 if it is lost or compromised.)

The Controller stores the token hash and the name assigned, viewable 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 illustrates 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**.

Only include the plaintext authentication token, not the token hash, saved in the controller running configuration. If the plaintext token is unknown, revoke access for the token and generate a new one. Note the plaintext value displayed after generating the new token.

## B.2 DMF Recorder Node API Headers

The supported REST API HTTP header entries are listed in the following table.

**Table 15: DANZ Monitoring Fabric (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 DANZ Monitoring Fabric (DMF) Recorder Node REST APIs are listed below.

### B.3.1 Ready

**/ready**

- Description: Is the DANZ Monitoring Fabric (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

**/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

**/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

**/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

`/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 Analysis Type

`/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 Stenographer Real-time Boolean

`/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 Integer

`/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

`/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

**/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 Integer

**/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

**/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

**/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

#### `/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

#### `/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

#### `/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

## DMF Controller in Microsoft Azure

---

The DANZ Monitoring Fabric (DMF) Controller in Azure feature supports the operation of the Arista Networks DMF Controller on the Microsoft Azure platform and uses the Azure CLI or the Azure portal to launch the Virtual Machine (VM) running the DMF Controller.

The DMF Controller in Azure feature enables the registration of VM deployments in Azure and supports **auto-firstboot** using Azure **userData** or **customData**.

### C.1 Configuration

Configure Azure VMs **auto-firstboot** using **customData** or **userData**. There is no data merging from these sources, so provide the data via **customData** or **userData**, but not both.

Arista Networks recommends using **customData** as it provides a better security posture because it is available only during VM provisioning and requires **sudo** access to mount the virtual CDROM.

**userData** is less secure because it is available via Instance Metadata Service (IMDS) after provisioning and can be queried from the VM without any authorization restrictions.


If **sshKey** is configured for the admin account during Azure VM provisioning along with **auto-firstboot** parameters, then it is also configured for the **admin** user of DMF controllers.

The following table lists details of the firstboot parameters for the **auto-firstboot** configuration.

#### Firstboot Parameters - Required Parameters

| Key               | Description                                                                                                                                          | Valid Values |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| admin_password    | This is the password to set for the admin user. When joining an existing cluster node this will be the admin-password for the existing cluster node. | string       |
| recovery_password | This is the password to set for the recovery user.                                                                                                   | string       |

#### Additional Parameters

| Key                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                      | Required | Valid Values      | Default Value                                   |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-------------------|-------------------------------------------------|
| hostname            | This is the hostname to set for the appliance.                                                                                                                                                                                                                                                                                                                                                                                                   | no       | string            | Configured from Azure Instance Metadata Service |
| cluster_name        | This is the name to set for the cluster.                                                                                                                                                                                                                                                                                                                                                                                                         | no       | string            | Azure-DMF-Cluster                               |
| cluster_to_join     | <p>This is the IP which firstboot will use to join an existing cluster. Omitting this parameter implies that the firstboot will create a new cluster.</p> <div>  <p><b>Note:</b> If this parameter is present ntp-servers, cluster-name, and cluster-description will be ignored. The existing cluster node will provide these values after joining.</p> </div> | no       | IP Address String |                                                 |
| cluster_description | This is the description to set for the cluster.                                                                                                                                                                                                                                                                                                                                                                                                  | no       | string            |                                                 |

## Networking Parameters

| Key                | Description                                                                            | Required                             | Valid Values                                                                                                                                           | Default Value |
|--------------------|----------------------------------------------------------------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| ip_stack           | What IP protocols to set up for the appliance management NIC.                          | no                                   | enum: ipv4, ipv6, dual-stack                                                                                                                           | ipv4          |
| ipv4_method        | Setup IPv4 for the appliance management NIC.                                           | no                                   | enum: auto, manual                                                                                                                                     | auto          |
| ipv4_address       | The static IPv4 address used for the appliance management NIC.                         | only if ipv4-method is set to manual | IPv4 Address String                                                                                                                                    |               |
| ipv4_prefix_length | The prefix length for the IPv4 address subnet to use for the appliance management NIC. | only if ipv4-method is set to manual | 0..32                                                                                                                                                  |               |
| ipv4_gateway       | The static IPv4 gateway to use for the appliance management NIC.                       | no                                   | IPv4 Address String                                                                                                                                    |               |
| ipv6_method        | Set up IPv6 for the appliance management NIC.                                          | no                                   | enum: auto, manual                                                                                                                                     | auto          |
| ipv6_address       | The static IPv6 address to use for the appliance management NIC.                       | only if ipv6-method is set to manual | IPv6 Address String                                                                                                                                    |               |
| ipv6_prefix_length | The prefix length for the IPv6 address subnet to use for the appliance management NIC. | only if ipv6-method is set to manual | 0..128                                                                                                                                                 |               |
| ipv6_gateway       | The static IPv6 gateway to use for the appliance management NIC.                       | no                                   | IPv6 Address String                                                                                                                                    |               |
| dns_servers        | The DNS servers for the cluster to use                                                 | no                                   | List of IP address strings                                                                                                                             |               |
| dns_search_domains | The DNS search domains for the cluster to use.                                         | no                                   | List of the host names or FQDN strings                                                                                                                 |               |
| ntp_servers        | The NTP servers for the cluster to use.                                                | no                                   | List of the host names of FQDN strings<br>0.bigswitch.pool.ntp.org<br>1.bigswitch.pool.ntp.org<br>2.bigswitch.pool.ntp.org<br>3.bigswitch.pool.ntp.org |               |

## Examples

```
{
 "admin_password": "admin_user_password",
 "recovery_password": "recovery_user_password"
```

```
}
```

## Full List of Parameters

```
{
 "admin-password": "admin_user_password",
 "recovery_password": "recovery_user_password",
 "hostname": "hostname",
 "cluster_name": "cluster name",
 "cluster_description": "cluster description",
 "ip_stack": "dual-stack",
 "ipv4_method": "manual",
 "ipv4_address": "10.0.0.3",
 "ipv4_prefix-length": "24",
 "ipv4_gateway": "10.0.0.1",
 "ipv6_method": "manual",
 "ipv6_address": "be:ee::1",
 "ipv6_prefix-length": "64",
 "ipv6_gateway": "be:ee::100",
 "dns_servers": [
 "10.0.0.101",
 "10.0.0.102"
],
 "dns_search_domains": [
 "dns-search1.com",
 "dns-search2.com"
],
 "ntp_servers": [
 "1.ntp.server.com",
 "2.ntp.server.com"
]
}
```

## C.2 Syslog Messages

- There are three possible failure modes:
  - VM fails Azure registration.
  - **auto-firstboot** fails due to a transient error or bug.
  - **auto-firstboot** parameter validation fails.
- These failures can be debugged by accessing the **firstboot** logs after manually booting the VM or logging via the **recovery** user on Azure serial console:
  - Azure DMF Controller VMs can be accessed via **ssh** login after successful **firstboot**:

```
dmf-controller-0-vm> enable; configure;
dmf-controller-0-vm> show logging syslog | grep 'floodlight-autofirstboot'
```

- For debugging parameter validation errors, access the parameter validation results:

```
dmf-controller-0-vm> show firstboot parameter-validation
```

- Accessing logs via the **recovery** user on Azure serial console. The following output is an example log for missing a required **firstboot** parameter – **admin\_password**:

```
Log in as 'admin' to configure
controller login: recovery
```



```

recovery@controller:~$ cat /var/log/floodlight/firstboot/firstboot.log
...
2024-06-17 17:09:09,982 autofirstboot: CRITICAL [main] Uncaught exception
Traceback (most recent call last):
File "/usr/bin/floodlight-autofirstboot", line 11, in <module>
load_entry_point('firstboot==0.1.0', 'console_scripts', 'floodlight-a
utofirstboot') ()
File "/usr/share/floodlight/firstboot/firstboot/autofirstboot.py", line 93,
in main
params, plugin = get_params(plugins)
File "/usr/share/floodlight/firstboot/firstboot/autofirstboot.py", line 44,
in get_params
params = plugin.get_firstboot_params()
File "/usr/share/floodlight/firstboot/firstboot/cloud_plugins/azure.py",
line 75, in get_firstboot_params
return FirstbootParams(**firstboot_param_dict)
TypeError: __init__() missing 1 required positional argument: 'admin_passwo
rd'

```

### C.3 Troubleshooting

- If a DMF Controller VM cannot be accessed via ssh login, the auto-firstboot has probably failed.
- The DMF Controller VMs must be recreated on Azure for any transient failure for VM registration with Azure or for auto-firstboot to occur.
- The DMF Controller VMs can also be configured manually for firstboot via the Azure serial console.

### C.4 Limitations

The following limitations apply to the DANZ Monitoring Fabric (DMF) Controller in Microsoft Azure.

- There is no support for any features specific to Azure-optimized Ubuntu Linux, including Accelerated Networking.
- The DMF Controllers in Azure are only supported on Gen-1 VMs.
- The DMF Controllers in Azure do not support adding the virtual IP address for the cluster.
- There is no support for capture interfaces in Azure.
- DMF ignores the Azure username and password fields.
- There is no support for static IP address assignment that differs from what is configured on the Azure NIC.
- The DMF Controllers are rebooted if the static IP on the NIC is updated.
- Switches are supported in L3 ZTN mode only.

### C.5 Resources

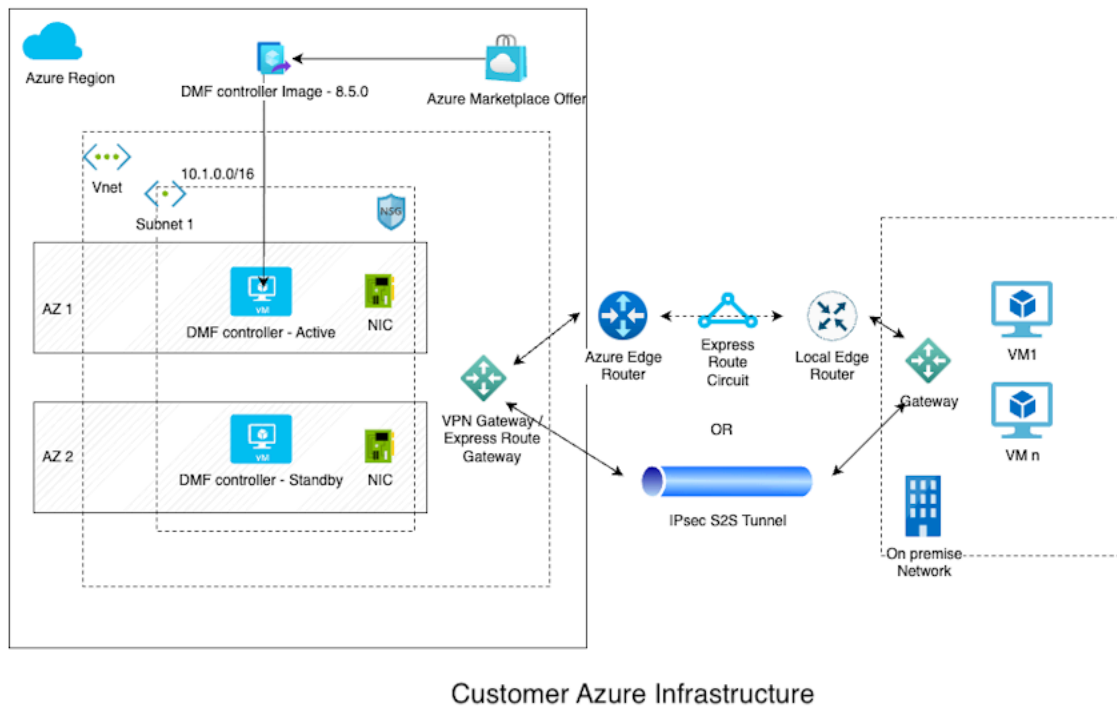
Please refer to the following resources for more information.

- Azure user data details :<https://learn.microsoft.com/en-us/azure/virtual-machines/user-data>
- Azure custom data details:<https://learn.microsoft.com/en-us/azure/virtual-machines/custom-data>
- Azure serial console details:<https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/windows/serial-console-overview>
- Azure Gen1 vs Gen2 VMs:<https://learn.microsoft.com/en-us/azure/virtual-machines/generation-2>

- Azure optimized Ubuntu Linux features: <https://ubuntu.com/blog/microsoft-and-canonical-increase-velocity-with-azure-tailored-kernel>
- Azure NIC assignment behavior: <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/reset-network-interface-azure-linux-vm>
- DMF hardware and software requirements: <https://www.arista.com/en/hcl-dmf/hcl-supported-software-controller-hardware-and-switch-hardware>

## C.6 Diagrams

Figure C-1: Customer Azure Infrastructure



## DMF Controllers in Google Cloud VMware Engine

The DANZ Monitoring Fabric (DMF) Controller in Google Cloud VMware Engine (GCVE) supports the operation of the Arista Networks DMF Controller on the GCVE platform and uses the vCenter portal to launch the Virtual Machine (VM) running the DMF Controller.

The DMF Controller in GCVE enables the registration of VM deployments in the vCenter environment and supports **auto-firstboot** when deployed using the **Deploy OVF Template** wizard in the vCenter.

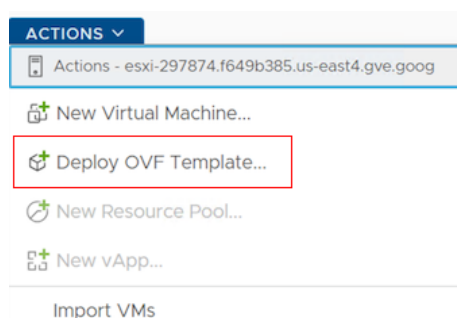
### Configuration

GCVE portal provides the vCenter server IP address with the credential details. Log in to the vCenter using these credentials to deploy the Controller VM using the Deploy OVF Template wizard in the vCenter.

The following describes using the wizard to deploy the Controller VM.

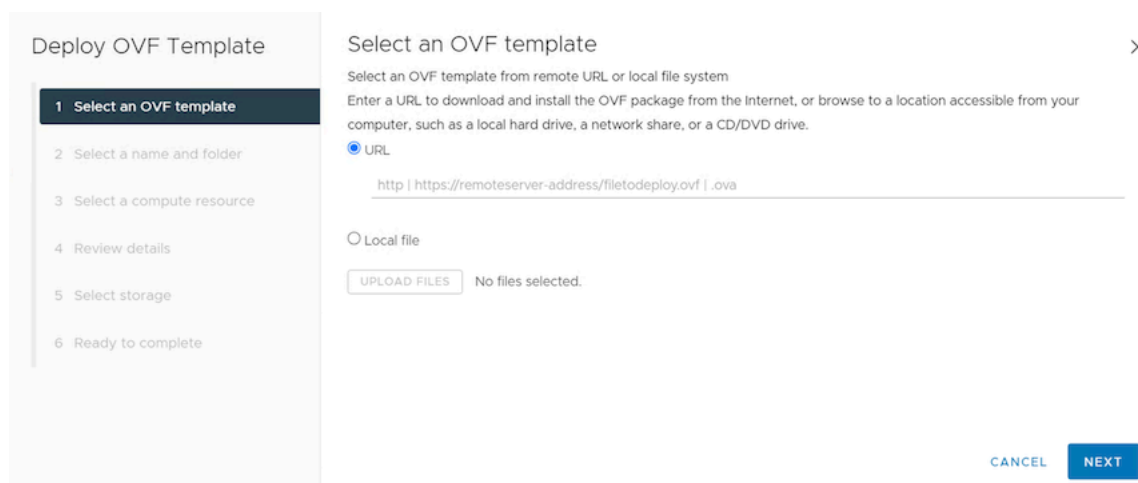
1. In VMware vCenter, navigate to the **Actions** menu and select **Deploy OVF Template**.

**Figure D-1: Deploy OVF Template**



2. The **Deploy OVF Template** wizard appears.

**Figure D-2: Deploy OVF Template Wizard**



3. Enter the **URL** for the OVA file. If the OVA file is available locally on the machine, choose the **Local file** option and select **Upload Files**.

**Figure D-3: Enter URL or Local File**

The screenshot shows the 'Deploy OVF Template' dialog box. On the left is a sidebar with a list of steps: 1. Select an OVF template (highlighted), 2. Select a name and folder, 3. Select a compute resource, 4. Review details, 5. Select storage, and 6. Ready to complete. The main area is titled 'Select an OVF template' and contains instructions: 'Select an OVF template from remote URL or local file system. Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.' There are two radio buttons: 'URL' and 'Local file' (which is selected). Below the 'URL' radio button is a text input field containing 'http | https://remote-server-address/filetoinstall.ovf|.ova'. Below the 'Local file' radio button is a blue 'UPLOAD FILES' button followed by the text 'controller-bmf-8.6.0-alpha.ova'. At the bottom right are 'CANCEL' and 'NEXT' buttons.

4. Provide the **Virtual machine name** and select its **location**.

**Figure D-4: Virtual Machine Name**

The screenshot shows the 'Deploy OVF Template' dialog box at step 2, 'Select a name and folder'. The sidebar on the left shows step 2 highlighted. The main area is titled 'Select a name and folder' and contains the instruction: 'Specify a unique name and target location'. There is a label 'Virtual machine name:' followed by a text input field containing 'controller-bmf-8.6.0-alpha'. Below this is the instruction 'Select a location for the virtual machine.' followed by a tree view. The tree view shows a folder icon and the text 'vcsa-297877.f649b385.us-east4.gve.goog'. Below this is a sub-item 'Datacenter' which is selected and highlighted. At the bottom right are 'CANCEL', 'BACK', and 'NEXT' buttons.

5. Select the **destination compute resource**.

**Figure D-5: Compute Resource**

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a compute resource

Select the destination compute resource for this operation

▼ Datacenter

▼ dmf-qa

esxi-297874.f649b385.us-east4.gve.goog

> HCX Management

> Workload

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

6. **Review details** summarize the OVA file details.

**Figure D-6: Review Details**

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Select networks

7 Customize template

8 Ready to complete

Review details

Verify the template details.

|               |                                                            |
|---------------|------------------------------------------------------------|
| Publisher     | No certificate present                                     |
| Product       | DANZ Monitoring Fabric                                     |
| Version       | 8.6.0-alpha                                                |
| Vendor        | <a href="#">Arista Networks, Inc</a>                       |
| Download size | 17.8 GB                                                    |
| Size on disk  | 19.9 GB (thin provisioned)<br>400.0 GB (thick provisioned) |

CANCEL

BACK

NEXT

7. Provide the **storage details** for the VM.

**Figure D-7: Storage Details**

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage**
- 6 Select networks
- 7 Customize template
- 8 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format

VM Storage Policy

☐ Disable Storage DRS for this

| Name          | Storage Cor | Capacity | Provisioner | Free    | Type | Placement |
|---------------|-------------|----------|-------------|---------|------|-----------|
| vsanDatastore | --          | 17.47 TB | 7.34 TB     | 15.6 TB | vSAN | Local     |

Compatibility

CANCEL BACK NEXT

8. As illustrated in the following **Customize template** sections, **autofirstboot** parameters appear. Use these parameters to configure the VM to make it ready after the OVF deployment. There are 19 settings; however, some are mandatory, while others are optional.

**Figure D-8: Customize Template Auto First Boot – 19 Settings**

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Select networks
- 7 Customize template**
- 8 Ready to complete

Customize template

Customize the deployment properties of this software solution.

All properties have valid values

autofirstboot 19 settings

use-autofirstboot Setting this to true will enable autofirstboot with the properties below.

admin-password This is the password to set for the admin user. When joining an existing cluster node this will be the admin-password for the existing cluster node.

recovery-password This is the password to set for the recovery user.

hostname This is the hostname to set for the appliance.

cluster-name This is the name to set for the cluster.

CANCEL BACK NEXT

**Figure D-9: Customize Template Auto First Boot**

Deploy OVF Template

1 Select an OVF template  
2 Select a name and folder  
3 Select a compute resource  
4 Review details  
5 Select storage  
6 Select networks  
7 **Customize template**  
8 Ready to complete

Customize template

cluster-name

This is the name to set for the cluster.

cluster-description

This is the description to set for the cluster.

cluster-to-join

This is the IP which firstboot will use to join an existing cluster. Omitting this parameter implies that the firstboot will create a new cluster. Note: If this parameter is present ntp-servers, cluster-name and cluster-description will be ignored. These values will be provided by the existing cluster node after joining.

ip-stack

What IP protocols should be setup for the appliance management NIC.  

ipv4

ipv4-method

How to setup IPv4 for the appliance management NIC.  

auto

ipv4-address

The static IPv4 address to use for the appliance management NIC.

CANCEL

BACK

NEXT

**Figure D-10: Customize Template Auto First Boot**

Deploy OVF Template

1 Select an OVF template  
2 Select a name and folder  
3 Select a compute resource  
4 Review details  
5 Select storage  
6 Select networks  
7 **Customize template**  
8 Ready to complete

Customize template

ipv4-address

The static IPv4 address to use for the appliance management NIC.

ipv4-prefix-length

The prefix length for the IPv4 address subnet to use for the appliance management NIC.

ipv4-gateway

The static IPv4 gateway to use for the appliance management NIC.

ipv6-method

How to setup IPv6 for the appliance management NIC.  

auto

ipv6-address

The static IPv6 address to use for the appliance management NIC.

ipv6-prefix-length

The prefix length for the IPv6 address subnet to use for the appliance management NIC.

ipv6-gateway

The static IPv6 gateway to use for the appliance management NIC.

CANCEL

BACK

NEXT

**Figure D-11: Customize Template Auto First Boot**

### Deploy OVF Template

- Select an OVF template
- Select a name and folder
- Select a compute resource
- Review details
- Select storage
- Select networks
- Customize template**
- Ready to complete

### Customize template

How to setup IPv6 for the appliance management NIC.

auto

IPv6 address

The static IPv6 address to use for the appliance management NIC.

IPv6 prefix length

The prefix length for the IPv6 address subnet to use for the appliance management NIC.

IPv6 gateway

The static IPv6 gateway to use for the appliance management NIC.

DNS servers

The DNS servers for the cluster to use.

DNS search domains

The DNS search domains for the cluster to use.

NTP servers

The NTP servers for the cluster to use.

CANCEL

BACK

NEXT

- Review the **Ready to complete** firstboot parameters and select **Finish** to enter the configuration. If changes are required, select **Back** and make those changes.

**Figure D-12: Ready To Complete**

### Deploy OVF Template

- Select an OVF template
- Select a name and folder
- Select a compute resource
- Review details
- Select storage
- Select networks
- Customize template
- Ready to complete**

### Ready to complete

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP allocation settings |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| IP protocol            | IPv4                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| IP allocation          | Static - Manual                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Properties             | <pre> use-autofirstboot = True admin-password = adminadmin recovery-password = adminadmin hostname = dmif-gcve-qa-03.bsn.sjc.aristanetworks.com cluster-name = dmif-gcve-cluster cluster-description = cluster-to-join = ip-stack = ipv4 ipv4-method = manual ipv4-address = 10.49.102.6 ipv4-prefix-length = 26 ipv4-gateway = 10.49.102.1 ipv6-method = auto ipv6-address = ipv6-prefix-length = ipv6-gateway = dns-servers = dns-search-domains = qa.bsn.sjc.aristanetworks.com ntp-servers = 0.bigswitch.pool.ntp.org                     </pre> |

CANCEL

BACK

FINISH

## D.1 Firstboot Parameters

The following table lists the firstboot parameters for the **auto-firstboot** configuration.



**Required Parameters**

| Key               | Description                                                                                                                                          | Valid Values |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| admin_password    | This is the password to set for the admin user. When joining an existing cluster node this will be the admin-password for the existing cluster node. | string       |
| recovery_password | This is the password to set for the recovery user.                                                                                                   | string       |

**Additional Parameters**

| Key                 | Description                                                                                                                                                                                                                                                                                                                  | Required | Valid Values      | Default Value |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-------------------|---------------|
| hostname            | This is the hostname to set for the appliance.                                                                                                                                                                                                                                                                               | no       | string            |               |
| cluster_name        | This is the name to set for the cluster.                                                                                                                                                                                                                                                                                     | no       | string            |               |
| cluster_to_join     | This is the IP which firstboot will use to join an existing cluster. Omitting this parameter implies that the firstboot will create a new cluster. Note: If this parameter is present ntp-servers, cluster-name, and cluster-description will be ignored. The existing cluster node will provide these values after joining. | no       | IP Address String |               |
| cluster_description | This is the description to set for the cluster.                                                                                                                                                                                                                                                                              | no       | string            |               |

## Networking Parameters

| Key                | Description                                                                            | Required?                            | Valid Values                                                                                                                                           | Default Value |
|--------------------|----------------------------------------------------------------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| ip_stack           | What IP protocols should be set up for the appliance management NIC.                   | no                                   | enum: ipv4, ipv6, dual-stack                                                                                                                           | ipv4          |
| ipv4_method        | How to setup IPv4 for the appliance management NIC.                                    | no                                   | enum: auto, manual                                                                                                                                     | auto          |
| ipv4_address       | The static IPv4 address used for the appliance management NIC.                         | only if ipv4-method is set to manual | IPv4 Address String                                                                                                                                    |               |
| ipv4_prefix_length | The prefix length for the IPv4 address subnet to use for the appliance management NIC. | only if ipv4-method is set to manual | 0..32                                                                                                                                                  |               |
| ipv4_gateway       | The static IPv4 gateway to use for the appliance management NIC.                       | no                                   | IPv4 Address String                                                                                                                                    |               |
| ipv6_method        | How to set up IPv6 for the appliance management NIC.                                   | no                                   | enum: auto, manual                                                                                                                                     | auto          |
| ipv6_address       | The static IPv6 address to use for the appliance management NIC.                       | only if ipv6-method is set to manual | IPv6 Address String                                                                                                                                    |               |
| ipv6_prefix_length | The prefix length for the IPv6 address subnet to use for the appliance management NIC. | only if ipv6-method is set to manual | 0..128                                                                                                                                                 |               |
| ipv6_gateway       | The static IPv6 gateway to use for the appliance management NIC.                       | no                                   | IPv6 Address String                                                                                                                                    |               |
| dns_servers        | The DNS servers for the cluster to use                                                 | no                                   | List of IP address strings                                                                                                                             |               |
| dns_search_domains | The DNS search domains for the cluster to use.                                         | no                                   | List of the hostnames or FQDN strings                                                                                                                  |               |
| ntp_servers        | The NTP servers for the cluster to use.                                                | no                                   | List of the hostnames of FQDN strings:<br>0.bigswitch.pool.ntp.org<br>1.bigswitch.pool.ntp.org<br>2.bigswitch.pool.ntp.org<br>3.bigswitch.pool.ntp.org |               |

## D.2 Post Installation Verification

1. Access a DMF Controller VM using an SSH login session.
2. Verify the Controller clusters have formed using the **show controller** command.

```
GCVE-DMF-860-1# show controller
Cluster Name : GCVE-DMF-860-1
Cluster Description : GCVE DMF CONTROLLER
Cluster Virtual IP : 10.100.189.231
Redundancy Status : redundant
Last Role Change Time : 2024-05-21 14:23:38.917000 PDT
Failover Reason : Changed connection state: connected to node 21773
Device deployment mode : pre-configured
Cluster Uptime : 1 day, 3 hours
Hostname @ State Uptime
-|-----|-----|-----|
1 GCVE-DMF-860-1 * active 5 hours, 6 minutes
2 GCVE-DMF-860-2 standby 15 minutes, 11 secs
GCVE-DMF-860-1#
```

3. Use the **show version** command to verify the installed version is correct.

```
GCVE-DMF-860-1# show version
Controller Version : DANZ Monitoring Fabric 8.6.0 (bmf/dmf-8.6.0 #10)
GCVE-DMF-860-1#
```

## D.3 Troubleshooting

The following are possible failure modes:

- **auto-firstboot** fails due to a transient error or bug.
- **auto-firstboot** parameter validation fails.

To debug these failures, connect to the VM using the console, log in as user `recovery`, and use the following command to review the **firstboot** logs.

```
recovery@dmf-controller-0-vm:~$ less /var/log/floodlight/firstboot/firstboot.log
```

To debug parameter validation errors, access the parameter validation results using the following command.

```
recovery@dmf-controller-0-vm:~$ less /var/lib/floodlight/firstboot/validation-results.json
```

## D.4 Limitations

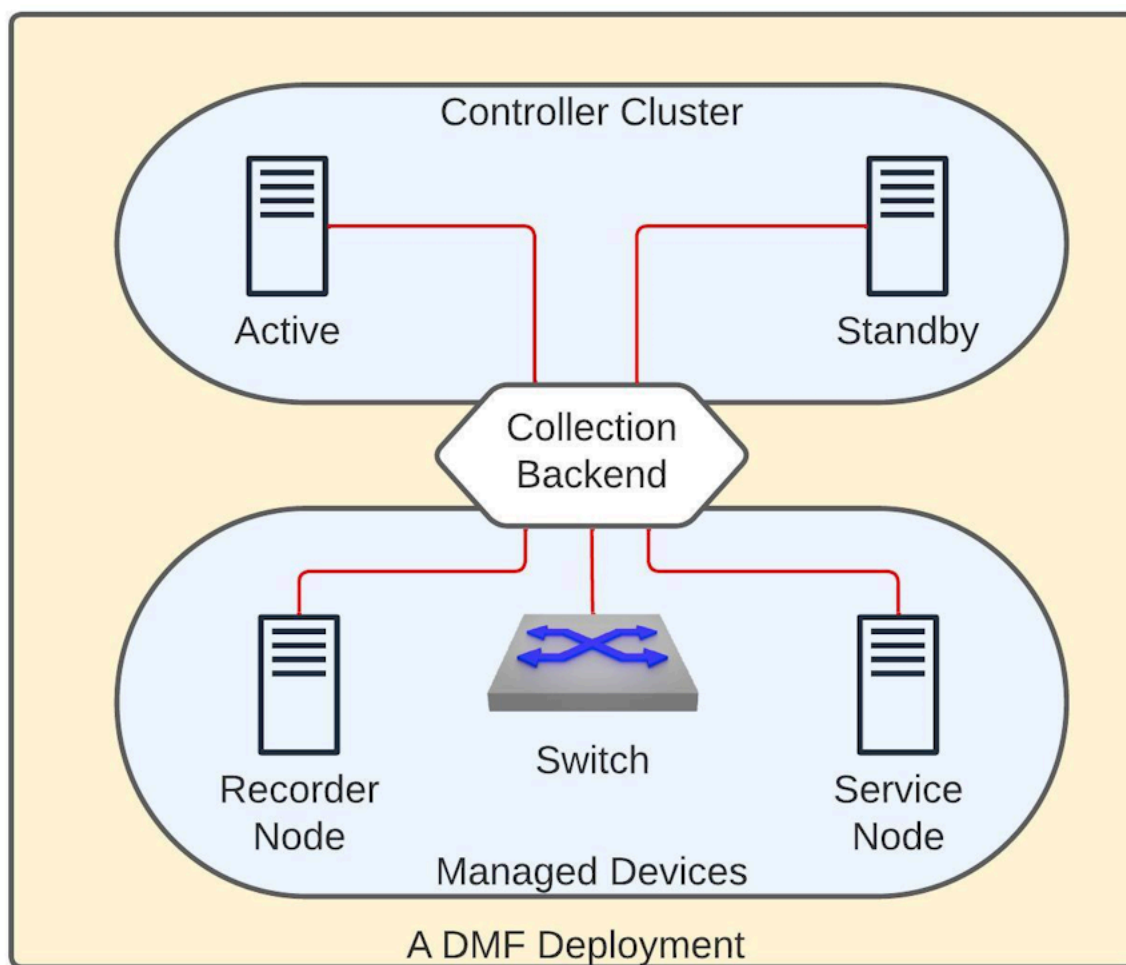
- There is no support for capture interfaces in GCVE.

## Telemetry Collector

A DANZ Monitoring Fabric (DMF) consists of a pair of Controllers, switches, and managed appliances. The Telemetry Collector feature centrally retrieves the infrastructure metrics (interface counters, CPU usage, etc.) affiliated with all these devices from the controllers using its REST API.

### Deployment

Figure E-1: DMF Deployment



The diagram above shows a DMF deployment with an Active/Standby Controller cluster. In this environment, each Controller collects metrics from all the devices it manages and the controller nodes. Each Controller establishes a gNMI connection to all the devices and the other Controller in a fabric to collect telemetry streams. gNMI is a gRPC-based protocol to configure and access states on network devices. The REST API exposes this information in the `/api/v1/data/controller/telemetry/data` subtree. The Controller exposes several metrics about CPU, memory, disk utilization, interface counters, and sensor states of all devices in a fabric. All metrics are fetched at a 10-second frequency except those associated with the sensors, which are collected every minute.

## Configuration

No additional configuration is necessary on the DMF Controller to enable the metric collection. However, to read these metrics, the user must be an admin user or configured with the privilege, category:TELEMETRY. To set the telemetry permission for a custom group and associate the user with this group, use the following commands:

```
dmf-controller(config)# group group_name
dmf-controller(config)# permission category:TELEMETRY privilege read-only
dmf-controller(config)# associate user username
```

## Connection Status

The Controller uses the gNMI protocol to collect telemetry data from all devices. DMF reports the status of these connections in both the REST API and the CLI.

### REST API

The REST API subtree `/api/v1/data/controller/telemetry/connection` reports the telemetry connection state. The API Schema browser of the GUI provides more details.

### CLI

The following **show** commands display the connection details. All these commands support filtering the output with a device name.

```
show telemetry connection device-name | all
show telemetry connection device-name | all details
show telemetry connection device-name | all last-failure
```

The **show telemetry connection device-name all** command shows the state and the latest state change of the connection between the Controller and the devices.

```
dmf-controller# show telemetry connection all
Name State Last state change
-|-----|-----|-----|-----|
1 c2 ready 2023-11-10 13:10:02.718000 UTC
2 core2 ready 2023-11-10 13:22:56.311000 UTC
<snip>
```

The **show telemetry connection device-name all** command displays the state and the latest state change of the connection between the Controller and the devices.

```
dmf-controller# show telemetry connection all details
Name State Last state change Target Connection type Last message time
-|-----|-----|-----|-----|-----|-----|
1 core2 ready 2023-11-10 13:22:56.311000 UTC 10.243.255.102:6030 clear-text 2023-11-10 13:59:35.437000 UTC
<snip>
```

The **show telemetry connection device-name all last-failure** command displays more details about a connection failure. The time of the latest failure and the potential reason appear in the output. If the connection is still in the failed state, this output also shows when the next reconnection will be attempted.

```
dmf-controller# show telemetry connection all last-failure
Name Fail time Fail type Root cause
Next retry in
-|-----|-----|-----|-----|
---|-----|
1 core2 2023-11-10 13:19:34.237000 UTC unavailable UNAVAILABLE: io
exception 0
```

---

<snip>

## Limitations

- Software interfaces (for example, loopback, bond, and management) do not report counters for broadcast and unicast packets.
- The reported interface names are the raw physical interface name (e.g., `et1`) rather than the user-configured name associated with the role of an interface (e.g., `filter1`).
- Resetting the interface counter does not affect the counter values stored at the `/telemetry/data` path. The value monotonically increases and corresponds to the total count since the device was last powered up. This value only gets reset when rebooting the device.

## Usage Notes

- DMF uses the configured name of a managed device (e.g., switch, recorder node, etc.) on the Controller as the value of the key name for the node device for all the metrics corresponding to it. In the case of a Controller, DMF uses the configured hostname as the key. Thus, these names must be unique in a specific DMF deployment.
- The possibility exists that metrics are not collected from a device for a short period. This data gap may happen when rebooting the device or when the Controllers experience a failover event.
- If the gNMI connection between the Controller and the device is interrupted, the Controller attempts a new connection after 2 minutes. The retry timeout for a subsequent connection attempt increases exponentially and can go up to 120 minutes. Upon a successful reconnection, this timeout value resets to 2 minutes.
- There might be gNMI warning messages in the floodlight log during certain events, e.g., when first adding a device or it is reloading. Ignore these messages.
- This feature enables an OpenConfig agent on switches running EOS to collect telemetry.

## Telemetry Availability

As a DMF fabric consists of different types of devices, the metrics of each vary. The following outlines the metrics collected from each device type by the Controller and typically made available over its REST API. However, some specific platforms or hardware might not report a particular metric. For brevity, the following list mentions the leaves that can correspond to a metric. For more details, use the API Schema browser of the GUI.

```
telemetry
+-- data
 +-- device
 +-- interface
 | +-- oper-status Ctrl, SWL, EOS, SN, RN
 | +-- counters
 | +-- in-octets Ctrl, SWL, EOS, SN, RN
 | +-- in-pkts Ctrl, SWL, EOS, SN, RN
 | +-- in-unicast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- in-broadcast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- in-multicast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- in-discards Ctrl, SWL, EOS, SN, RN
 | +-- in-errors Ctrl, SWL, EOS, SN, RN
 | +-- in-fcs-errors Ctrl, SWL, EOS, SN, RN
 | +-- out-octets Ctrl, SWL, EOS, SN, RN
 | +-- out-pkts Ctrl, SWL, EOS, SN, RN
 | +-- out-unicast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- out-broadcast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- out-multicast-pkts Ctrl, SWL, EOS, SN, RN
 | +-- out-discards Ctrl, SWL, EOS, SN, RN
 | +-- out-errors Ctrl, SWL, EOS, SN, RN
 +-- cpu
 | +-- utilization Ctrl, SWL, EOS, SN, RN
```

```

+-- memory
| +-- total Ctrl, SWL, SN, RN
| +-- available Ctrl, SWL, EOS, SN, RN
| +-- utilized Ctrl, SWL, EOS, SN, RN
+-- sensor
| +-- fan
| | +-- oper-status Ctrl, SWL, EOS, SN, RN
| | +-- rpm Ctrl, SWL, EOS, SN, RN
| | +-- speed SWL
| +-- power-supply
| | +-- oper-status Ctrl, SWL, EOS, SN, RN
| | +-- capacity EOS
| | +-- input-current Ctrl, SWL, EOS, SN, RN
| | +-- output-current SWL, EOS
| | +-- input-voltage Ctrl, SWL, EOS, SN, RN
| | +-- output-voltage SWL, EOS
| | +-- input-power Ctrl, SWL, SN, RN
| | +-- output-power SWL, EOS
| +-- thermal
| +-- oper-status Ctrl, SWL, SN, RN
| +-- temperature Ctrl, SWL, EOS, SN, RN
+-- mount-point
| +-- size Ctrl, SWL, SN, RN
| +-- available Ctrl, SWL, SN, RN
| +-- utilized Ctrl, SWL, SN, RN
| +-- usage-percentage Ctrl, SWL, SN, RN
+-- control-group
+-- memory Ctrl, SWL, SN, RN
+-- cpu Ctrl, SWL, SN, RN

```

\* Ctrl = Controller, SWL = A switch running SwitchLight OS,  
EOS = A switch running Arista EOS, SN = Service Node, RN = Recorder Node

### DMF GUI REST API Inspector

---

#### Overview

DANZ Monitoring Fabric (DMF) 8.6 introduces a newly designed API Inspector available on all DMF UI pages. Previously, the former API Inspector was accessible by selecting the Dragonfly icon, and it was only available on some pages.

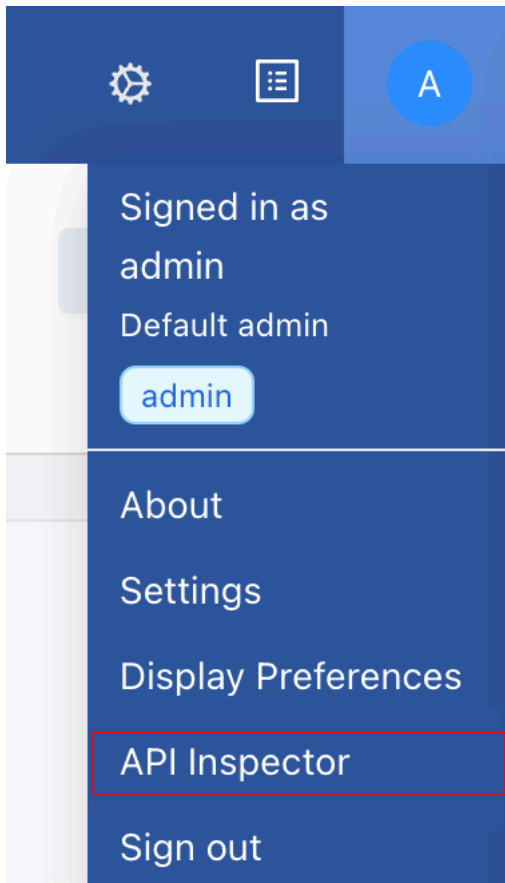
#### Launch the API Inspector

1. Hover over and select the **User** icon to see the drop-down menu.
2. Select **API Inspector** to open the API Inspector table.
3. The REST API Inspector table appears between the page content and the navigation bar. It will remain open when switching between pages.

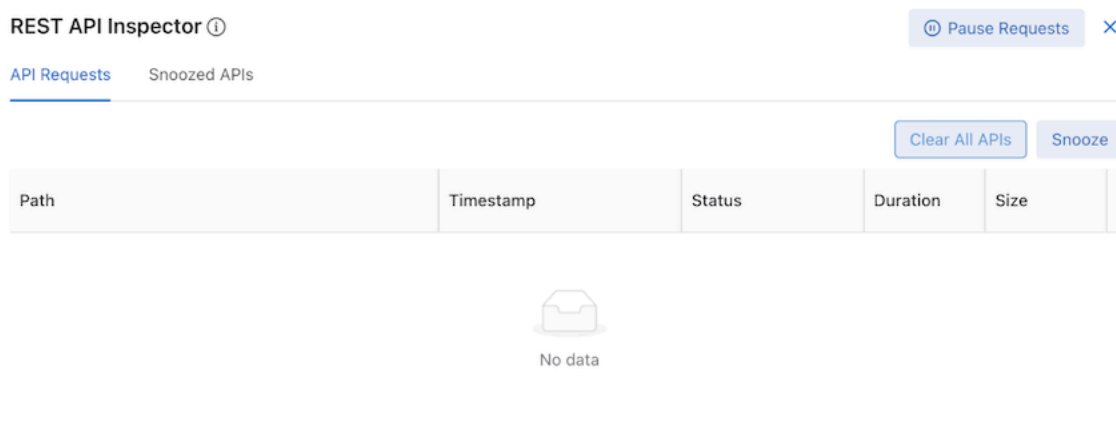


4. To close it, select **API Inspector** or the **x** icon.

**Figure F-1: Menu > API Inspector**



**Figure F-2: REST API Inspector**



### API Requests Tab

**API Requests** contains a table displaying all recent API requests and several utility buttons.

**Figure F-3: API Requests Tab**

REST API Inspector ⓘ

API Requests

Snoozed APIs

Clear All APIs

Snooze

| Path                                                                            | Timestamp              | Status | Duration | Size    |  |
|---------------------------------------------------------------------------------|------------------------|--------|----------|---------|--|
| GET /api/v1/data/controller/applications/dmf/topology/core-interface            | 4/22/2024, 11:55:00 AM | 200    | 3s 348ms | 14.8 kB |  |
| GET /api/v1/data/controller/applications/dmf/topology/filter-interface          | 4/22/2024, 11:55:00 AM | 200    | 3s 348ms | 2.0 B   |  |
| GET /api/v1/data/controller/applications/dmf/topology/packet-recorder-interface | 4/22/2024, 11:55:00 AM | 200    | 3s 348ms | 2.0 B   |  |
| GET /api/v1/data/controller/applications/dmf/topology/service-interface         | 4/22/2024, 11:55:00 AM | 200    | 3s 348ms | 2.0 B   |  |
| GET /api/v1/data/controller/applications/dmf/topology/delivery-interface        | 4/22/2024, 11:55:00 AM | 200    | 3s 347ms | 2.0 B   |  |
| GET /api/v1/data/controller/core/aaa/session-config                             | 4/22/2024, 11:55:00 AM | 200    | 2s 579ms | 29.0 B  |  |
| POST /api/v1/rpc/controller/core/aaa/permission-check                           | 4/22/2024, 11:55:00 AM | 200    | 2s 709ms | 1.0 kB  |  |
| GET /api/v1/data/controller/applications/dmf/op-global-setting                  | 4/22/2024, 11:54:59 AM | 200    | 2s 214ms | 38.0 B  |  |

### API Requests Table

Each row contains specific API request information:

- **Path** – HTTP request method and URL
- **Timestamp** – date and approximate time
- **Status** – HTTP response codes
- **Duration** – time taken to receive a response
- **Size** – response body size

The table logs the latest 500 API requests and automatically updates when new API calls are received.

### Search



The table's results are searchable. Selecting the **magnifying glass** icon displays a **Type to Search** line in the first row.

**Figure F-4: Search by Column**

| Path           | Timestamp      | Status         | Duration   | Size       |  |
|----------------|----------------|----------------|------------|------------|--|
| Type to Search | Type to Search | Type to Search | Type to... | Type to... |  |

The table supports search by Path, HTTP Method, and Status column contexts.

### Interaction

Selecting a URL expands specific API requests and enters the **Detailed API Request View**.

### Pause Requests Button

## ⏸ Pause Requests

Use **Pause Requests** to stop tracking new API requests. An alert message appears next to the REST API Inspector title.

## ▶ Resume Requests

During the pause, the button changes to **Resume Requests**. To resume tracking API requests, select **Resume Requests**. The table updates accordingly.



**Note:** Any requests made during the pausing state are not retrieved.

Figure F-5: Paused Requests

## REST API Inspector ⓘ

Requests has been paused.

### API Requests

### Snoozed APIs

#### Clear All APIs Button

## Clear All APIs

To clear the entire table, select **Clear All APIs**.

Figure F-6: Cleared Table

| REST API Inspector ⓘ        |           | ⏸ Pause Requests ×      |          |      |
|-----------------------------|-----------|-------------------------|----------|------|
| API Requests   Snoozed APIs |           |                         |          |      |
|                             |           | Clear All APIs   Snooze |          |      |
| Path                        | Timestamp | Status                  | Duration | Size |
| <br>No data                 |           |                         |          |      |

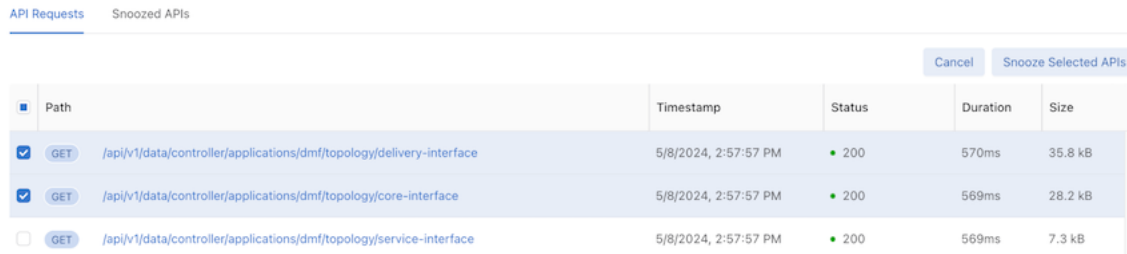
The Snoozed API table is not affected. API Inspector continues receiving and displaying new requests in the table.

#### Snooze Button

## Snooze

Using **Snooze** makes the table rows selectable, and the Snooze button changes to **Snooze Selected APIs**. Clicking **Cancel** or **Snooze Selected APIs** returns the table entries to an unselected state.

**Figure F-7: API Requests**



| <input type="checkbox"/>            | Path                                                                     | Timestamp            | Status | Duration | Size    |
|-------------------------------------|--------------------------------------------------------------------------|----------------------|--------|----------|---------|
| <input checked="" type="checkbox"/> | GET /api/v1/data/controller/applications/dmf/topology/delivery-interface | 5/8/2024, 2:57:57 PM | 200    | 570ms    | 35.8 kB |
| <input checked="" type="checkbox"/> | GET /api/v1/data/controller/applications/dmf/topology/core-interface     | 5/8/2024, 2:57:57 PM | 200    | 569ms    | 28.2 kB |
| <input type="checkbox"/>            | GET /api/v1/data/controller/applications/dmf/topology/service-interface  | 5/8/2024, 2:57:57 PM | 200    | 569ms    | 7.3 kB  |

During the selectable state, select specific API Requests and select **Snooze Selected APIs** to prevent them from appearing in the table.

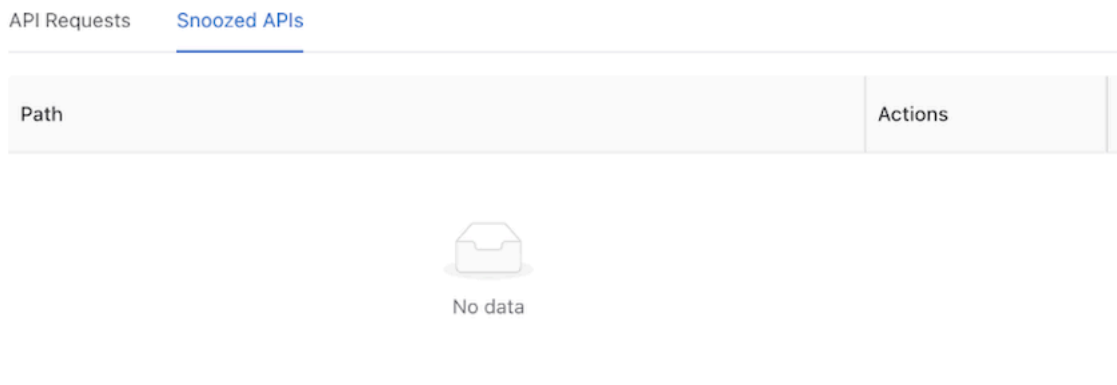
### Snoozed APIs Tab


Snoozed APIs display all snoozed API requests.

### Snoozed APIs Table

The snoozed APIs table is empty by default.








**Figure F-8: Snoozed APIs Default Table**



| Path                                                                                           | Actions |
|------------------------------------------------------------------------------------------------|---------|
| <br>No data |         |

The table displays snoozed APIs. However, unlike the **API Requests** table, the **Snoozed API** table only distinguishes API requests by URL and HTTP Method. Requests with the same URL only appear once in the Snoozed API Table.

**Figure F-9: Snoozed APIs**


| API Requests                                                                                                                       |  | Snoozed APIs (3)                                                                   |                                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------|--|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Path                                                                                                                               |  | Actions                                                                            |  |
|  /api/v1/data/controller/core/aaa/group           |  |  |                                                                                     |
|  /api/v1/data/controller/core/aaa/local-user      |  |  |                                                                                     |
|  /api/v1/rpc/controller/core/aaa/permission-check |  |  |                                                                                     |

### Search



The table's results are searchable. Selecting the **magnifying glass** icon displays a **Type to Search** line in the first row.

**Figure F-10: Type to Search**

| Path           | Actions    |  |
|----------------|------------|---------------------------------------------------------------------------------------|
| Type to Search | Type to... |                                                                                       |

The **Snoozed APIs** table supports search by Path columns contexts only.

### Unmute Button

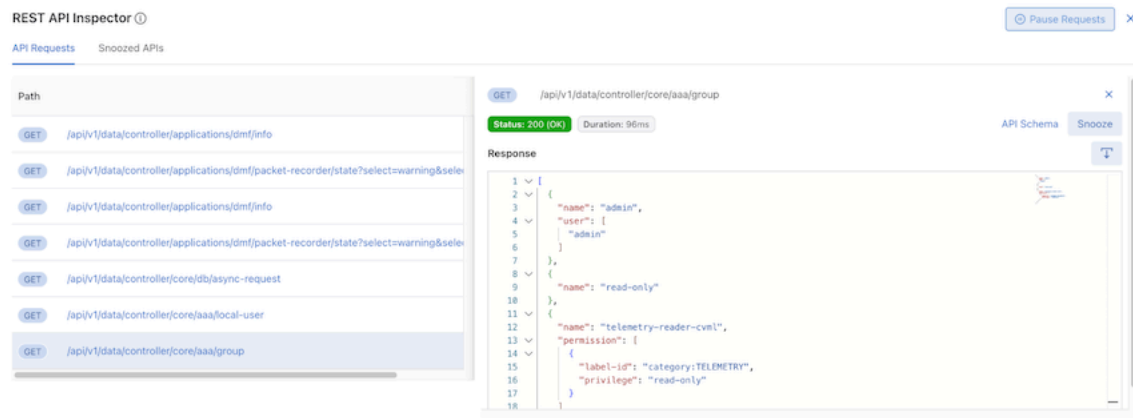


**Unmute** resumes tracking the API request and updates the table. Any requests with the same URL that occur going forward are added to the API requests table.

## Detail API Requests View

Selecting a **URL** in the API Requests table expands that specific API request and enters the **Detail API Requests** view.

**Figure F-11: Detailed API Requests View**



The UI displays a truncated API Requests table and a detailed view; selecting another **URL** changes the detailed information in the detailed view.

Information on the selected API request includes:

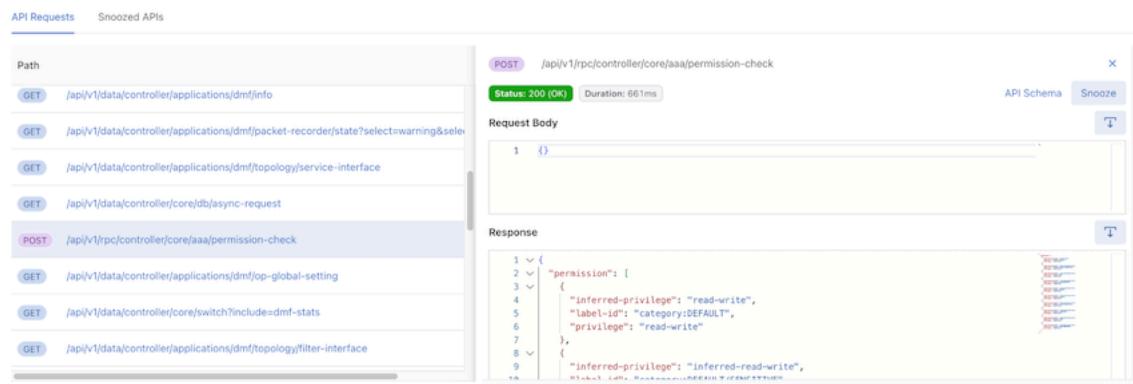
- **HTTP methods** – show the HTTP methods
- **Response status** – show the response status and response text
- **Duration** – show the duration of the API request
- **API Schema** link button – Open a new tab to the API Schema Browser Page
- **Snooze** – Snooze this selected API request (same as **Snooze** Button )
- **Response** – Show all the response information of this selected API request
- **Download** Icon – Download response information into a .json file.



**Note:** For PATCH / POST / PUT requests, another field, **Request Body**, shows the request body information.

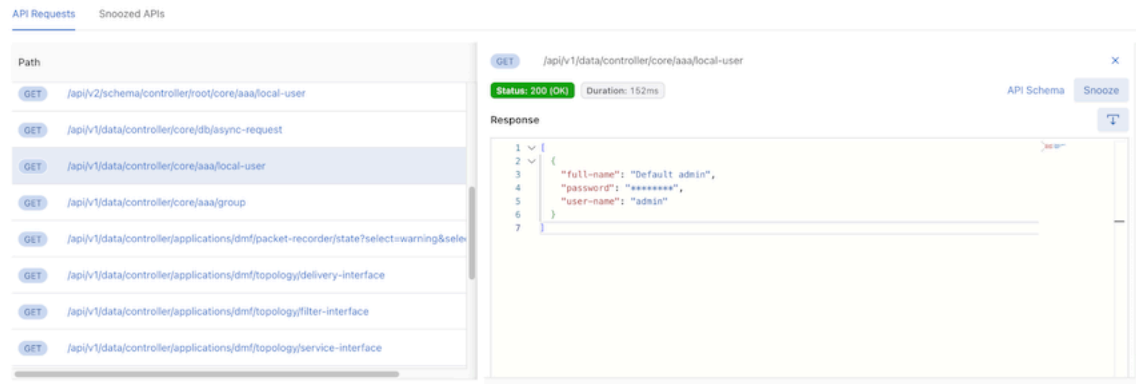
- Use the **download** icon to download the request body information into a .json file.

**Figure F-12: Download JSON File**



- **API Inspector** obfuscates sensitive values such as password, hash ID, hash password, auth session, encode result, and secret by replacing them with **\*\*\*\*\***.

**Figure F-13: Obfuscated Sensitive Data**



## Configuring Third-party Services

### G.1 Services in the DANZ Monitoring Fabric

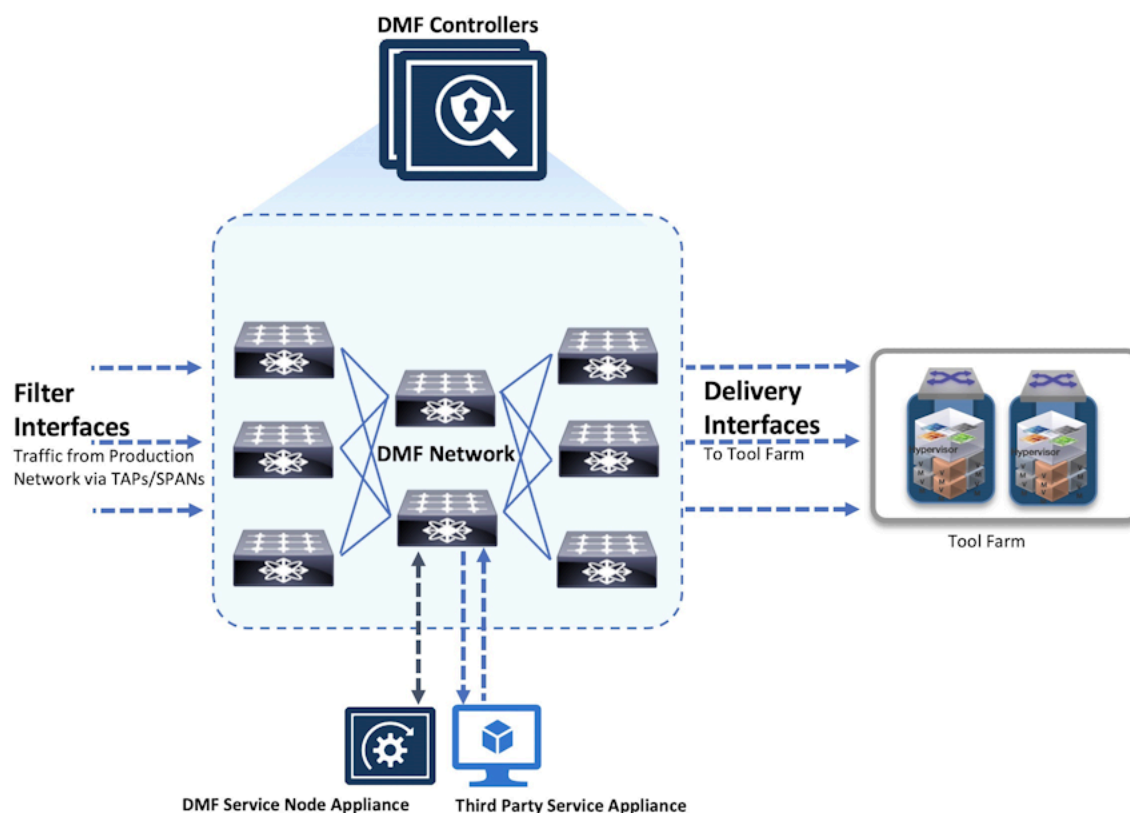
Services in the DANZ Monitoring Fabric (DMF) refer to packet modification operations provided by third-party network packet brokers (NPBs), referred to as service nodes. Services can include operations that refine or modify the data stream delivered to analysis tools.

Each service instance is assigned a numeric identifier because multiple services can be specified for a given policy. Services are applied sequentially, applying a service with a lower sequence number first.

Service nodes are optional devices that process interesting traffic before forwarding it 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 used with the service.
- Use the DMF interface names to create a service node and add pre-service and post-service interfaces.

**Figure G-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 that 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 policy includes a service, it is only optional if defined explicitly as optional. If not defined as optional in the policy, packet forwarding does not occur when the service is unavailable. For example, configuring the packet-slicing service as optional and a pre-service or post-service interface assigned to that service node is down, 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 unavailable for the time-stamping service, this policy does not forward packets to the delivery interfaces.

Configure all the service interfaces before creating 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, and 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. Use a Link Access Group (LAG) as a pre-service or a post-service interface.



**Note:** Arista Networks strongly recommends configuring the post-service and pre-service interfaces on the same switch for any DMF service.

### G.1.1 Using the GUI to Configure a DMF Unmanaged Service

To create a DANZ Monitoring Fabric (DMF) unmanaged service, perform the following steps:

1. Select **Monitoring > Services**.  
The system displays the following table:

**Figure G-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 |
|-------------------|-------------|-----------------------------|------------------------------|--------------|------------------------|-------------------------|-----------------------|------------------------|
| anthony_service_1 |             | 1                           | --                           | 0            | --                     | --                      | --                    | --                     |
| anthony_service_2 |             | 1                           | --                           | 0            | --                     | --                      | --                    | --                     |

- This table lists the services configured for the DMF. Add, delete, or modify existing services as required.
2. To create a new service, click the provision control (+) in the table.

The system displays the following dialog:

**Figure G-3: Create Service Dialog: Info**



The 'Create Service' dialog is shown with the 'Info' tab selected. The 'Name' field contains 'THIRD-PARTY-SERVICE-1' and the 'Description' field is empty. The 'NEXT' button is highlighted in blue.

| Step                       | Status |
|----------------------------|--------|
| 1. Info                    | ✓      |
| 2. Pre-Service Interfaces  | ✓      |
| 3. Post-Service Interfaces | ✓      |

Name \*  
THIRD-PARTY-SERVICE-1

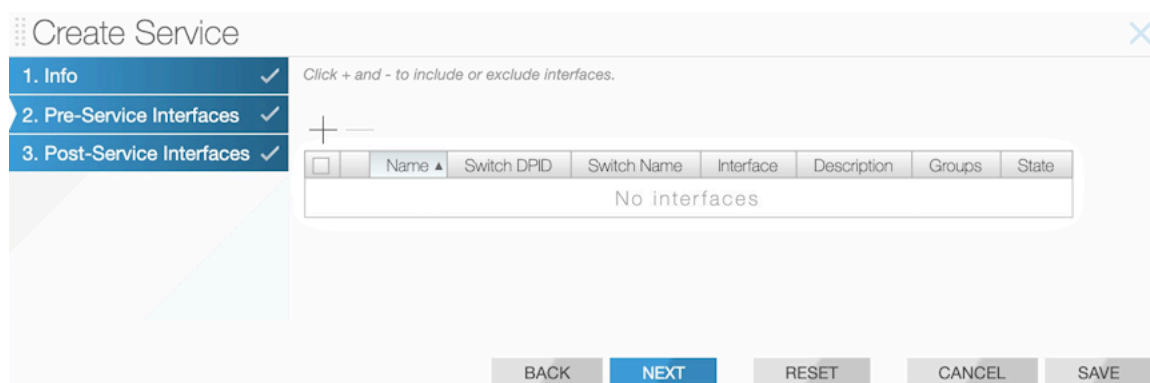
Description

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 G-4: Create Service Dialog: Pre-service Interfaces**



The 'Create Service' dialog is shown with the 'Pre-Service Interfaces' tab selected. A table lists the interfaces assigned as pre-service interfaces. The table is currently empty, showing 'No interfaces'. The 'NEXT' button is highlighted in blue.

Click + and - to include or exclude interfaces.

+ -

|               | Name | Switch DPID | Switch Name | Interface | Description | Groups | State |
|---------------|------|-------------|-------------|-----------|-------------|--------|-------|
| No interfaces |      |             |             |           |             |        |       |

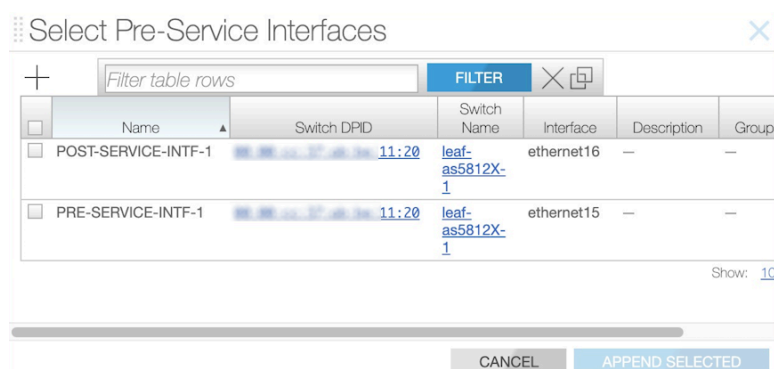
BACK NEXT RESET CANCEL SAVE

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 (+) in the table.

The system displays the following dialog:

**Figure G-5: Select Pre-service Interfaces**



The 'Select Pre-Service Interfaces' dialog is shown. It features a table with columns for Name, Switch DPID, Switch Name, Interface, Description, and Groups. Two interfaces are listed: 'POST-SERVICE-INTF-1' and 'PRE-SERVICE-INTF-1'. The 'PRE-SERVICE-INTF-1' interface is selected. The 'APPEND SELECTED' button is highlighted in blue.

+ Filter table rows FILTER

|                                     | Name                | Switch DPID | Switch Name    | Interface  | Description | Groups |
|-------------------------------------|---------------------|-------------|----------------|------------|-------------|--------|
| <input type="checkbox"/>            | POST-SERVICE-INTF-1 | 11:20       | leaf-as5812X-1 | ethernet16 | —           | —      |
| <input checked="" type="checkbox"/> | PRE-SERVICE-INTF-1  | 11:20       | leaf-as5812X-1 | ethernet15 | —           | —      |

Show: 10

CANCEL APPEND SELECTED

This table lists the interfaces available for assignment as pre-service interfaces. To configure a new interface, click the provision control (+) in 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**.

- On page two of the **Create Service Interface** dialog, click **Next**.  
The system displays the following dialog:

**Figure G-6: Create Service Dialog: Post-service Interfaces**

This table lists the interfaces assigned as post-service interfaces for the current service.

- To add a post-service interface, click the provision control (+) in the table.

The system displays the following dialog.

**Figure G-7: Select Post-service Interfaces**

This table lists the interfaces available for assignment as post-service interfaces.

To configure a new interface, click the provision control (+) in the table. The system displays a dialog for adding a service interface, as described in the **Configuring DMF Unmanaged Services** section.

- Enable the checkbox for one or more interfaces to assign as a post-service interface for the current service and click **Append Selected**.
- Click **Save** on page three of the **Create Service Dialog**.

## G.1.2 Using the CLI to Configure a DMF Unmanaged Service

In the DANZ Monitoring Fabric (DMF), third-party tools that provide packet manipulation services, such as time stamping and packet slicing, are called DMF Unmanaged Services. These optional devices process traffic from filter interfaces before being forwarded to delivery interfaces.



**Note:** After adding a service to a policy, it is no longer optional unless specifically defining it as optional. If not defined as optional, the policy does not forward packets if the service is unavailable.

To configure an unmanaged service using the CLI, perform the following steps:

- 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 configured services 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 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

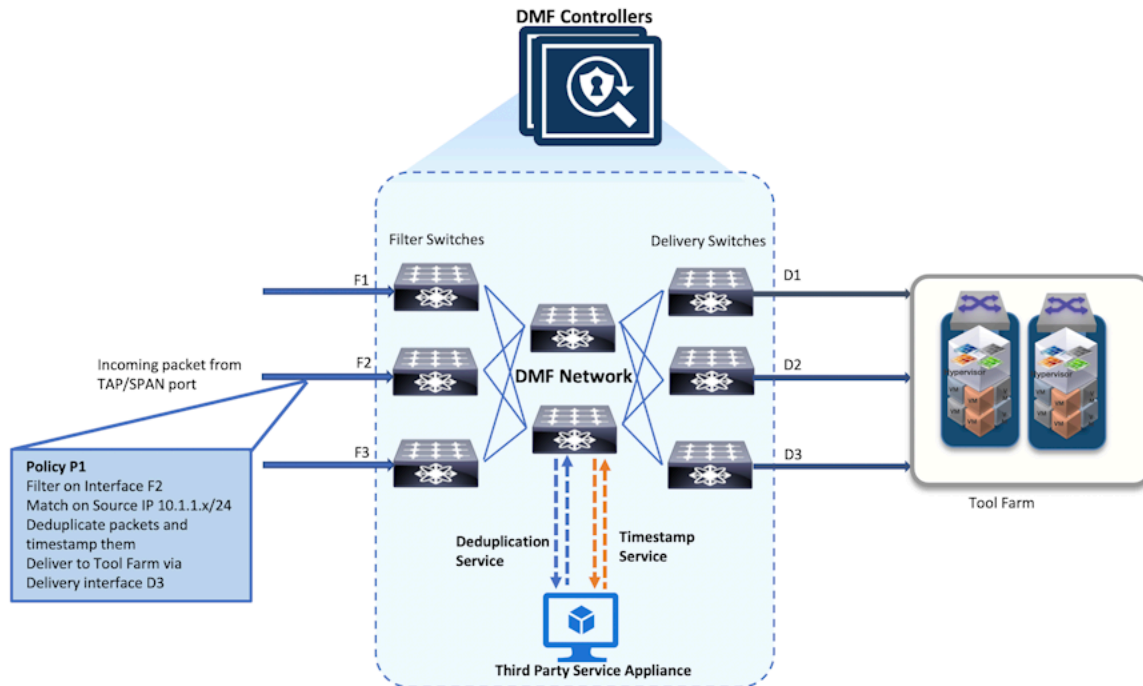
1 THIRD-PARTY-SERVICE-1 10Gbps 10Gbps - -
~~~~~ Post-groups of Service Names ~~~~~
# Service Name      Dmf name
-----
1 THIRD-PARTY-SERVICE-1 POST-SERVICE-INTF-1
~~~~~ Pre-groups of Service Names ~~~~~
Service Name Dmf name

1 THIRD-PARTY-SERVICE-1 PRE-SERVICE-INTF-1
```

## G.2 Service Insertion and Chaining in a DMF Policy

To configure a DANZ Monitoring Fabric (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 G-8: Service Insertion and Chaining**



Because a given policy can specify multiple services, set a sequence number for each service instance so the services are applied in order for the policy traffic. A lower sequence number applies the service 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 or post-service interfaces for the packet-slicing service nodes are down, then this service is skipped if 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.
```

Enter the **show policy** command from any mode to display the run time services being applied.

## References

---

### H.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*