

ARISTA

Deployment Guide

DANZ Monitoring Fabric (DMF)

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 2025 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: Getting Started..... 1**
 - 1.1 DMF Installation Prerequisites.....1
 - 1.2 Downloading Software.....1
 - 1.3 Where to Start.....2
 - 1.4 Documentation Summary.....2
 - 1.5 DMF Software and Hardware.....3
 - 1.6 DANZ Monitoring Fabric Quick Start.....4

- Chapter 2: Installing and Configuring the DMF Controller..... 6**
 - 2.1 Connecting to the Controller..... 6
 - 2.1.1 Connecting to the Controller Appliance Using a Terminal Server..... 6
 - 2.2 Configuring the Active Controller Using the First Boot Script..... 7
 - 2.3 Configuring the Standby Controller..... 10
 - 2.3.1 Joining Standby Controller to Existing Cluster..... 10
 - 2.3.2 Moving existing Standby Controller to different IPv4 subnet..... 13
 - 2.3.3 Moving an Existing Standby Controller to a Different Controller Cluster..... 14
 - 2.4 Accessing the DANZ Monitoring Fabric Controller..... 14
 - 2.4.1 Using the DANZ Monitoring Fabric CLI..... 14
 - 2.4.2 Capturing CLI Command Output..... 15
 - 2.4.3 Using the DANZ Monitoring Fabric GUI..... 16
 - 2.5 Managing the DMF Controller Cluster..... 18
 - 2.5.1 Verifying Cluster Configuration..... 18
 - 2.5.2 Configuring the Cluster Virtual IP..... 18
 - 2.5.3 Setting the Time Zone..... 19
 - 2.5.4 Viewing Controller and Cluster Status..... 21
 - 2.5.5 Saving and Restoring Controller Configuration..... 25
 - 2.5.6 Reloading and Rebooting a Controller Using the DMF UI..... 26
 - 2.6 Copying Files Between a Workstation and a DMF Controller..... 27
 - 2.6.1 Merge and Replace Parameters of Copy Config..... 29
 - 2.6.2 Snapshot File Management Using REST API..... 33
 - 2.6.3 Convert a Text-based running-config into a Snapshot..... 34
 - 2.7 Managing DMF Sessions..... 39
 - 2.8 Managing and Viewing Logs..... 41
 - 2.8.1 Sending Logs to a Remote Syslog Server..... 41
 - 2.8.2 Using the GUI to Configure Remote Logging..... 41
 - 2.8.3 Using the CLI to Configure Remote Logging..... 42
 - 2.8.4 Viewing Log Files on the Controller..... 43
 - 2.8.5 Administrative Activity Logs..... 43
 - 2.9 REST API Logging..... 44
 - 2.9.1 Restricting Size of REST API Body..... 45
 - 2.10 Syslog Over TLS..... 45
 - 2.10.1 Overview..... 45
 - 2.10.2 Configuration..... 46
 - 2.11 Creating a Support Bundle..... 47
 - 2.12 NIST 800-63b Password Compliance..... 49
 - 2.12.1 Configuration..... 49
 - 2.13 Custom Password Compliance..... 50

2.14 Switch/Managed Appliance Management Interfaces Not Mirroring Controller Management Interface ACLs.....	51
2.14.1 Configuration using the CLI.....	52
2.14.2 CLI Show Commands.....	54
2.14.3 Limitations.....	55
2.15 Recovery Procedure.....	55
2.15.1 Recovery from a Single Controller Failure.....	55
2.15.2 Recovery from a Dual Controller Failure.....	55
Chapter 3: DANZ Monitoring Fabric Deployment Topologies.....	57
3.1 DANZ Monitoring Fabric Topologies.....	57
3.1.1 Single Switch Topology.....	57
3.1.2 Two-Tier Topology.....	58
3.1.3 Three-Tier Any-Tap-to-Any-Tool Topology.....	58
Chapter 4: Installing DMF Switches.....	60
4.1 HTTPS Support for Controller Hosted URLs using ZTP.....	60
4.2 Zero Touch Fabric Provisioning Modes.....	60
4.3 Using L2 ZTF (Auto-Discovery) Provisioning Mode.....	61
4.3.1 Requirements.....	62
4.3.2 Switch Installation Procedure.....	62
4.3.3 Arista Switch Installation Procedure for 7050X Series and 7260X Series.....	66
4.3.4 Allocating IPv4 Addresses to Fabric Switches.....	68
4.3.5 Using the GUI to Allocate IPv4 Addresses.....	69
4.3.6 Using the CLI to Allocate IPv4 Addresses with IPAM.....	71
4.3.7 Assigning Static IPv4 Addresses.....	72
4.3.8 Using the GUI to Assign a Static IPv4 Address.....	72
4.3.9 Using the CLI to Assign a Static IPv4 Address.....	74
4.3.10 Static Address Assignment Troubleshooting and Limitations.....	75
4.4 Using L3 ZTN (Pre-Configured) Switch Provisioning Mode.....	77
4.4.1 Installing a Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	77
4.4.2 Installing Arista 7050X and 7260X Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	80
4.4.3 Installing Arista 7280R Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	84
4.4.4 Installing Arista 7800R3 Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	88
4.4.5 Configuring the Switch Static IP and Controller IP in Interactive ZTF Mode.....	94
4.4.6 Installing Arista 7050X and 7260X Series using DHCP with bootfile-name option.....	95
4.4.7 Installing Arista 7280R Series using DHCP with bootfile-name option.....	96
4.4.8 Using DHCP with Default URL for Switch Installation in Preconfigured Provisioning Mode.....	98
4.5 Registering a Switch After Initial Deployment.....	99
4.5.1 Using the GUI to Register a Switch.....	99
4.5.2 Using the CLI to Register a Switch.....	101
4.6 Changing the ZTF Mode After Deployment.....	102
4.6.1 Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode.....	102
4.6.2 Using the GUI to Change the Switch Provisioning Mode.....	102
4.6.3 Using the CLI to Change the Switch Provisioning Mode.....	103
4.6.4 Changing to Layer 3 ZTF (Preconfigured) Mode.....	104
4.6.5 Changing to Layer 2 ZTF (Auto-Discovery) Mode.....	105
4.6.6 System Reinstall for an EOS Switch.....	105
4.7 SKU Reporting for EOS Switches.....	106
4.7.1 Using the CLI to Configure SKU Reporting for EOS Switches.....	107

4.7.2 Using the GUI to Configure SKU Reporting for EOS Switches..... 107

Chapter 5: Managing Switches and Interfaces..... 112

5.1 Connecting Directly to a Switch..... 112
 5.1.1 Manually Configuring Enhanced Hashing for Load Distribution..... 112
 5.1.2 Configuring Enhanced Hashing..... 113
 5.1.3 Symmetric Load Balancing..... 113
 5.1.4 GTP Hashing..... 113
 5.2 Overriding the Default Switch Configuration..... 114
 5.3 Configuring Switch Interfaces..... 116
 5.3.1 Forward Error Correction (FEC)..... 118
 5.3.2 Autonegotiation..... 119
 5.3.3 Manually Setting the Interface Speed..... 120
 5.3.4 Using Breakout Cables..... 120
 5.3.5 Verifying Switch Configuration..... 122
 5.3.6 Disabling the TX Direction on an Interface..... 123

Chapter 6: DMF Upgrade Procedures..... 124

6.1 Upgrading the Controller..... 124
 6.1.1 Upgrade Procedure Summary..... 125
 6.1.2 Upgrade Options..... 125
 6.2 Copying the ISO Image File to the Controller..... 126
 6.3 Staging the Upgrade..... 127
 6.4 Launching the Upgrade..... 128
 6.5 DMF Switch Automatic Upgrade..... 128
 6.6 Verifying the Upgrade..... 129
 6.7 Verify Persistent IPAM Assigned IP Addresses after Upgrade..... 129
 6.8 Rolling Back an Upgrade..... 129

Chapter 7: Installing and Upgrading the DMF Service Node..... 132

7.1 Overview..... 132
 7.2 Connecting the Service Node..... 133
 7.3 Service Node Setup and Initial Configuration..... 134
 7.4 Creating Support Bundle on Service Node..... 137
 7.5 Upgrading Service Node Software From Release 7.x.x..... 138

Chapter 8: Installing and Configuring the DMF Recorder Node..... 140

8.1 Overview..... 140
 8.2 DMF Recorder Installation Procedure..... 142
 8.3 Initial Configuration - GUI..... 145
 8.4 Initial Configuration - CLI..... 147
 8.5 Changing the Recorder Node Default Configuration..... 148

Chapter 9: DHCPv4-Based First-boot for DMF Controller and Managed Appliances..... 151

9.1 Introduction..... 151
 9.2 Prepare Services (TFTP/NFS) for PAVE/REPAVE Operation..... 151
 9.3 DHCPv4-Based First-boot..... 154
 9.3.1 Auto Installation of Images..... 154
 9.3.2 Auto Configuration of First-boot..... 156

9.4 Assumption and Trust Model.....	158
Chapter 10: Managing SNMP.....	159
10.1 SNMP Overview.....	159
10.2 Using the DMF GUI to Configure SNMP.....	159
10.2.1 Configuring SNMP Traps.....	165
10.2.2 Configuring the System Name.....	170
10.3 Using the CLI to Configure SNMP.....	171
10.3.1 Configuring SNMP Access to the Controller.....	171
10.3.2 Configuring SNMP Access to the Analytics Node.....	172
10.3.3 Identifying the SNMP Trap Receiver.....	172
10.3.4 Configuring SNMP Settings.....	173
10.3.5 Configuring SNMP Switch Trap Thresholds.....	174
10.3.6 SNMP Traps for DMF Service Node Appliance.....	175
10.3.7 Managing the SNMPv3 Engine ID for Trap Receivers.....	175
10.3.8 Configuring SNMPv3 Users.....	176
10.3.9 SNMPv3 Command Examples.....	177
10.4 Configuring SNMP on a Specific Switch.....	178
10.4.1 Using the GUI to Configure SNMP on a Specific Switch.....	178
10.4.2 Using the CLI to Configure SNMP on a Specific Switch.....	180
10.5 SNMP Clear Trap.....	183
10.6 SNMP Trap Generation for Packet Drops and Link Saturation.....	184
10.6.1 Using the CLI to Configure the SNMP Traps.....	189
Chapter 11: Using Authentication, Authorization, and Accounting.....	192
11.1 Overview.....	192
11.2 Using Local Groups and Users to Manage DMF Controller Access.....	193
11.2.1 Viewing Existing Groups and Users.....	193
11.2.2 Using the GUI to Manage User Accounts & Security Groups.....	194
11.2.3 Using the CLI to Manage Groups and User Accounts.....	197
11.2.4 Changing User Passwords.....	198
11.2.5 Password Reset.....	199
11.2.6 Authentication with a User Token and REST API.....	200
11.3 Configuring AAA to Manage DMF Controller Access.....	201
11.3.1 Using the GUI - Enabling Remote AAA Services.....	201
11.3.2 Using the CLI - Enabling Remote AAA Services.....	205
11.4 Time-based User Lockout.....	205
11.5 Using TACACS+ to Control Access to the DMF Controller.....	206
11.5.1 Using the GUI to Add a TACACS+ Server.....	208
11.5.2 Using the CLI to Enable Remote Authentication and Authorization on the DMF Fabric Controller.....	209
11.5.3 Using the CLI to Add a TACACS+ Server.....	209
11.6 Setting up a Tac_Plus Server.....	210
11.6.1 Using the Same Credentials for DMF and Other Devices.....	211
11.6.2 RBAC-Based Configuration for Non-Default Group User.....	211
11.7 Using RADIUS for Managing Access to the DMF Controller.....	212
11.7.1 Using the GUI to Add a RADIUS Server.....	213
11.7.2 Using the CLI to Add a RADIUS Server.....	214
11.7.3 Setting up a FreeRADIUS Server.....	214
11.8 Custom admin role.....	216
11.8.1 Categories.....	216
11.8.2 Permissions & Privileges.....	221
11.8.3 Category Identification.....	222
11.8.4 Group Management.....	222

11.8.5 Remote Users.....	223
11.8.6 Commonly used profiles.....	223
11.8.7 Commonly used Category-feature Matrix.....	225
Chapter 12: Management and Control Plane Security.....	227
12.1 Management Plane Security.....	227
12.2 Importing the Controller Private Key and Certificate.....	227
12.3 Using Certificates Signed by a CA for GUI Access to the Controller.....	228
12.3.1 Replacing the Certificate.....	230
12.4 Managing the Controller HTTP and SSH Ciphers, Protocols, and Data Integrity Algorithms.....	231
12.4.1 Configuring HTTP Ciphers.....	231
12.4.2 Configuring HTTP Protocols.....	231
12.4.3 Configuring SSH Ciphers.....	232
12.4.4 Configuring SSH Data Integrity Algorithms.....	232
12.4.5 Changes to Supported MACs/Ciphers/SSH Keys.....	232
12.5 Inherit MAC and Cipher Configuration.....	234
12.5.1 Using the CLI to Configure SSH and HTTPS.....	234
12.5.2 Verify the Cryptographic Configuration.....	235
12.5.3 Limitations.....	236
12.6 Protocol Access Required to the DMF Controller.....	237
12.6.1 Management Plane Access.....	237
12.6.2 Control Plane Access for DMF Controller.....	238
12.6.3 Protocol Access Required to the DMF Controller - Sync.....	241
12.6.4 Control Plane Access for DMF Switches.....	242
12.6.5 Control Plane Access for DMF Service Node.....	243
12.6.6 Control Plane Access for DMF Recorder Node.....	244
12.6.7 Control Plane Access for Analytics Node.....	245
Chapter 13: Using iDRAC with a Dell R430, R630, or R730 Server.....	249
13.1 Setting Up and Configuring iDRAC.....	249
13.2 Using iDRAC to Install the DMF Controller or DMF Service Node Image.....	260
Chapter 14: Using iDRAC with a Dell R440 or R740 Server.....	267
14.1 Setting Up and Configuring iDRAC.....	267
14.2 Using iDRAC to Install a DMF Controller, Analytics, or Recorder Software Image.....	275
Chapter 15: Switch CPLD Upgrade Procedure.....	282
Chapter 16: Switch ONIE Upgrade Procedure.....	283
Appendix A: Removing an Existing OS from a Switch.....	284
A.1 Reverting from DMF (Switch Light OS) to EOS - 7050X3 and 7260CX3.....	284
A.2 Reverting from DMF (EOS) to EOS - 7280R.....	285
A.3 Removing the existing OS from a Switch.....	286
Appendix B: Creating a USB Boot Image.....	289
B.1 Creating the USB Boot Drive with MacOS X.....	289
B.2 Creating the USB Boot Image with Linux.....	290
B.3 Creating a USB Boot Image Using Windows.....	291

Appendix C: Installing a Controller VM.....	296
C.1 General Requirements.....	296
C.2 Installing on VMware ESXi/vSphere.....	296
C.2.1 Prerequisites.....	296
C.2.2 VM Installation.....	296
C.2.3 vMotion support for Virtual Controller.....	297
C.3 Installing on Ubuntu KVM.....	297
C.3.1 Prerequisites.....	297
C.3.2 VM Installation.....	297
Appendix D: Erasing DMF Appliance.....	299
D.1 Using the Dell Lifecycle Controller.....	299
Appendix E: Reforming Controller HA Cluster.....	308
E.1 Controller Cluster Recovery.....	308
Appendix F: DMF Controller in Microsoft Azure.....	309
F.1 Configuration.....	310
F.2 Limitations.....	314
F.3 Resources.....	314
F.4 Syslog Messages.....	314
F.5 Troubleshooting.....	315
Appendix G: Prometheus Endpoint Support for Infrastructure Metrics.....	316
Appendix H: References.....	323
H.1 Related Documents.....	323

Getting Started

This chapter describes the software and hardware requirements and summarizes the steps for deploying the DANZ Monitoring Fabric (DMF) Controller.

1.1 DMF Installation Prerequisites

Follow the prerequisite steps below:

Table 1: DANZ Monitoring Fabric (DMF) Prerequisites

Task	Description	Comments
1.	Connect the management port of the DMF switches to the management network switch.	
2.	Connect the DMF switch console port to the console server. The default baud rate is 9600 for Arista Networks switches. The default baud rate is 115200 for ONIE-enabled switches.	
3.	Connect the DMF Controller, Service Node, Recorder Node, and Analytics Node Ethernet port to the management network switch.	
4.	For the DMF Controller, Service Node, Recorder Node, and Analytics Node, connect the iDRAC port to the management network (if possible, to a different subnet). Assign an iDRAC IP and validate the DMF Controller, Service Node, Recorder Node and, Analytics Node iDRAC are reachable. The default password for iDRAC should be at the bottom of the Service Tag.	
5.	Connect the DMF Service Node and Recorder Node data ports to the DMF switch per customer-designed topology. Refer to Overview for Service Node data ports and Overview for Recorder Node data port.	
6.	On the management network switch, ensure that <code>spanning-tree mode edge</code> or <code>spanning-tree mode portfast</code> is enabled for all the ports to which the DMF appliances and switches are connected.	
7.	On management network switch ensure that IGMP snooping is disabled.	
8.	On the management network switch ensure IPv6 communication in the L2 domain.	
9.	Verify reachability to NTP, DNS, and default gateway IP address for the management network.	
10.	For L3 deployment (where Controller, DMF switches, DMF Service Node, or DMF Recorder Node are in different subnets), ensure the correct ports on the firewall are open. Refer to Management Plane Security (Protocol Access Required for the DMF Controller) for a list of ports on the firewall to open.	

1.2 Downloading Software

The Arista Support Portal provides links to software packages and documentation for DANZ Monitoring Fabric (DMF) products.

To download the DMF software or documentation, click the **DMF (BMF) / CCF (BCF)** tab and expand the **DANZ Monitoring Fabric (DMF) - Big Monitoring Fabric (BMF)** folder tree.

Download the appropriate release documents from the **DMF - BMF Product Documentation** folder tree.

To download software, expand the **DMF - BMF Software Downloads** folder tree and locate the required software from either **DMF - BMF Latest Recommended Release**, **DMF - BMF Prior Recommended Release**, or **DMF - BMF All Releases** folder trees.

Figure 1-1: Downloading Software



1.3 Where to Start

The following table identifies the documents that provide information and procedures for common tasks when deploying DANZ Monitoring Fabric (DMF).

Task	Document
1. Identify new features and any software or upgrade issues for the current release.	<i>Release Notes</i>
2. Verify the compatibility of any hardware components used in your deployment.	<i>Hardware Compatibility List</i>
3. Explore features and identify the configuration required.	<i>DMF User Guide</i> and <i>Analytics User Guide</i>
4. Other tasks as needed.	See Documentation Summary

 **Note:** The Documentation Summary identifies other documents to help perform tasks after initial deployment.

1.4 Documentation Summary

The following documents help perform tasks after initial deployment:

- **Analytics User Guide** - Configuring and using Arista Analytics.

- **Deployment Guide** - Procedures for installation, upgrade, and initial DMF Controller configuration common to all DMF features and use cases.
- **Hardware Compatibility List** - Lists software and hardware tested for interoperability with the current DMF release.
- **Hardware Guide** - Describes the LEDs and ports for DMF supported switch and appliance hardware.
- **DMF User Guide** - Configuring and using DANZ Monitoring Fabric features.
- **Release Notes** - New features, upgrade issues, open and resolved issues in the current release, SW behavior changes, and known limitations.
- **REST API Guide** - General Guidelines using REST API and a list of REST API calls.
- **SNMP MIB Reference Guide** - Supported MIBs.
- **Verified Scale** - Tested and verified scaling for DMF.

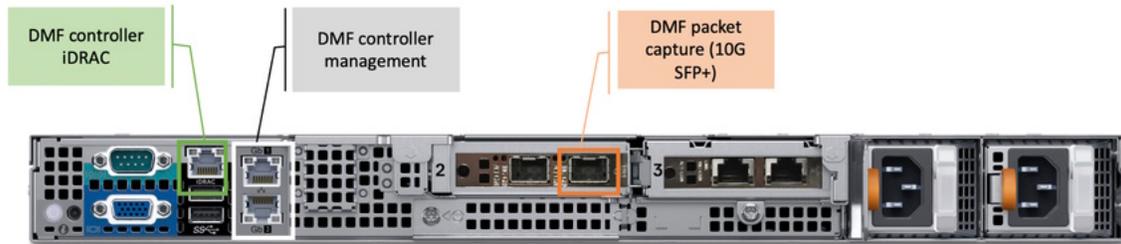
1.5 DMF Software and Hardware

This section lists the Controller and switch requirements for deploying the DANZ Monitoring Fabric (DMF).

- DMF Controller:
 - Controller Software.
 - Controller Hardware Appliance or KVM/ESX Virtual Machine (For VM requirements, refer to [Installing a Controller VM](#)).
- Switch Requirements
 - Refer to the **DANZ Monitoring Fabric 8.4 Hardware Compatibility List** for the supported switch hardware.
 - Switch software.

The figure below shows the interfaces provided by the DMF Controller hardware.

Figure 1-2: DMF Controller Hardware Appliance



The following table describes the interfaces:

Table 2: DMF Controller Interfaces

Port Set	Function	Where It is Connected
Controller Management Ports 2 x 1G white-labeled ports	Administrative access to the DMF Controller using CLI, GUI, or REST. Connect both ports for redundancy.	Customer management network.
Controller iDRAC green-labeled port	For remote management of DMF Controller Appliance.	Connect green-labeled 1Gb port to DMF management network.
Packet Capture Ports-10Gb First 10G orange-labeled ports	For Packet Capturing Feature on DMF Controller Appliance.	Connect the first orange-labeled 10Gb port to DMF fabric switches.

1.6 DANZ Monitoring Fabric Quick Start

Configuration	CLI Example	GUI Option	Documentation
Active and Standby DANZ Monitoring Fabric (DMF) Controllers.	Complete the First boot script. # show controller	Controller View	Installing and Configuring the DMF Controller
Monitoring Fabric Switches <ul style="list-style-type: none"> Name and MAC address. 	# switch switch-1 mac <mac-address>	Fabric > Switches	Installing DMF Switches
DMF Service Node <ul style="list-style-type: none"> Name and MAC address. Connect it to the DMF fabric and management network. First boot and initial setup. Managed services. 	# service-node <name> mac <mac-address>	Fabric > Switches Monitoring > Managed Services	Installing and Upgrading the DMF Service Node
DMF Recorder Node <ul style="list-style-type: none"> Register name and MAC address. Define recorder instance and interface. Assign a recorder interface to policy. 	# recorder-node device <recorder-name> mac <mac-address> # recorder-fabric interface <name> recorder-interface switch <switch name> <interface used for recording>	Monitoring > Recorder Node Monitoring > Policies	Installing and Configuring the DMF Recorder Node
Out-of-Band policies <ul style="list-style-type: none"> Assign <i>filter</i> interface role to port connected to SPAN or Tap interfaces. Assign <i>delivery</i> interface role to ports connected to tools. Match interesting traffic and forward traffic to tools connected to delivery ports. 	# switch switch-1 interface ethernet1 role filter interface-name span-tap-1 # switch switch-1 interface ethernet2 role delivery interface-name tool-1 # policy p1 10 match tcp action forward filter-interface span-tap-1 delivery-interface tool-1	Monitoring > Interfaces Monitoring > Policies	DANZ Monitoring Fabric 8.4 User Guide
Arista Analytics	Complete the First boot script and perform initial configuration	Settings > Analytics Configuration	Arista Analytics 8.4 User Guide

Installing and Configuring the DMF Controller

This chapter describes the basic procedure for installing the DANZ Monitoring Fabric (DMF) Controller software.

2.1 Connecting to the Controller

The DANZ Monitoring Fabric (DMF) Controller can be deployed as a hardware appliance, a one-RU hardware device containing preloaded software, or as a Virtual Machine (VM).

Complete the initial setup of the Controller in any of the following ways:

- Attach a monitor and a keyboard directly to the appliance's VGA and USB ports.
- Use a terminal server to connect to the appliance console port and telnet to the terminal server.
- When deploying the Controller as a VM, connect to the VM console.



Note: Arista Networks recommends having an iDRAC connection to the Controller, Service Node, Analytics Node, and Recorder Node appliances. This connection helps in easy troubleshooting of issues. Refer to the chapter [Using iDRAC with a Dell R440 or R740 Server](#) later in this guide for more details.

2.1.1 Connecting to the Controller Appliance Using a Terminal Server

After you connect the DANZ Monitoring Fabric (DMF) Controller appliance serial console to a terminal server, you can use telnet (or SSH if supported by your terminal server) to connect to the hardware appliance through the terminal server.

To connect the serial connection on a Controller appliance to a terminal server, complete the following steps:

1. Obtain a serial console cable with a DB-9 connector on one end and an RJ-45 connector on the other.
2. Connect the DB-9 connector to the hardware appliance DB-9 port (not the VGA port) and the RJ45 to the terminal server.
3. Configure the terminal server port baud rate to **115200**.
4. Set the port baud rate of **115200** on the terminal server.



Note:

- On some terminal servers, the saved configuration must be reloaded after changing the baud rate.
- You should now be able to use Telnet or SSH to access the hardware appliance serial console through the terminal server.

2.2 Configuring the Active Controller Using the First Boot Script

After connecting to the Controller appliance to start the initial setup, the system runs the First Boot script. The following configuration information is required to configure the DANZ Monitoring Fabric (DMF) Controller:

- IP address for the active and standby Controller
- The subnet mask and the default gateway IP address
- NTP server IP address
- Host name
- (Optional) DNS server IP address
- (Optional) DNS search domain name



Note: The default serial console rate on DMF hardware appliances for releases before **6.1** was **9600**. When upgrading an existing Controller appliance with a serial interface set to **9600** baud, change the terminal setting to **115200** after the upgrade.

To perform the initial configuration of the DMF Controller, complete the following steps:

1. Press **Enter** to begin the installation.

```
This product is governed by an End User License Agreement (EULA) .
You must accept this EULA to continue using this product.
You can view this EULA from our website at:
https://www.arista.com/en/eula
Do you accept the EULA for this product? (Yes/No) [Yes] > Yes
Running system pre-check
Finished system pre-check
Starting first-time setup
Local Node Configuration
-----
```

2. To accept the agreement, type the command: **Yes**.
Refer to Appendix B: Creating a USB Boot Image to make a bootable USB drive.



Note: When entering the command: **No**, the system prompts the user to log out. Accepting the EULA is mandatory to install the product.

Enter the command: **Yes** to accept the EULA. The system prompts for a password to be used for emergency recovery of the Controller node, as shown below.

```
Emergency recovery user password >
```

3. Set the recovery user password and confirm the entry.

```
Local Node Configuration
-----
Emergency recovery user password>
Emergency recovery user password (retype to confirm)>
Hostname> controller-1
```

4. Choose the IP forwarding mode when prompted.

```
[1] IPv4 only
[2] IPv6 only
[3] IPv4 and IPv6
>3
```

5. Choose the method for assigning IPv4 addresses.

```
Management network (IPv4) options:
[1] Manual
[2] Automatic via DHCP
[1] >
```



Note: When using [2] **Automatic via DHCP**, Arista Networks recommends reserving the DHCP address for the MAC address of the Controller's management port.

6. Choose the method for assigning IPv6 addresses.

```
Management network (IPv6) options:
[1] Automatic via SLAAC & DHCPv6
[2] Manual
[1] >
```

The DMF Controller Appliance allows IPv6 Stateless Address Auto-configuration for Controller IPv6 management.



Note: Support for Jumbo Ethernet frames is enabled by default.

When using [1] in **Step 6** above, the system prompts the user to enter the IPv4 address and related configuration information for the Controller node.

```
IPv4 address [0.0.0.0/0] > 10.106.8.2/24
IPv4 gateway (Optional) > 10.106.8.1
DNS server 1 (Optional) > 10.3.0.4
DNS server 2 (Optional) > 10.1.5.200
DNS search domain (Optional) > qa.aristanetworks.com
```

The IPv4 address configuration is applied to the Controller Management Interface (white-labeled ports) on the Controller appliance.

7. Enter the IP address and optional information regarding the DNS server, as shown in the example.



Note: Entering the IP address followed by a slash (/) and the number of bits in the subnet mask, the system does not prompt for the CIDR.

The system prompts the user to create a new cluster or join an existing cluster.

```
DNS search domain (Optional) > bsn.sjc.aristanetworks.com
Clustering
-----
Controller cluster options:
[1] Start a new cluster
[2] Join an existing cluster
> 1
```

8. Type 1 when configuring the active Controller.

The system prompts the user to enter the cluster name and description and to set the password.

```
> 1
Cluster name > dmf-cluster
Cluster description (Optional) >
Cluster administrator password >
Cluster administrator password (retype to confirm) >
```

9. Enter a name for the cluster and an optional description, and set the password for administrative access to the cluster.

The system then prompts for the name of each NTP server to use to set the system time required for synchronizing between nodes in the cluster and fabric switches.

```
System Time
-----
Default NTP servers:
- 0.bigswitch.pool.ntp.org
- 1.bigswitch.pool.ntp.org
- 2.bigswitch.pool.ntp.org
- 3.bigswitch.pool.ntp.org
NTP server options:
[1] Use default NTP servers
[2] Use custom NTP servers
[1] > 1
```



Note: Fabric switches and nodes obtain their NTP service from the active Controller.

10. Enter the IP address or fully qualified domain name of the NTP server for the network.

Completing the initial configuration without specifying the NTP server is possible, but deploying the DMF Controller in a production environment requires a functioning NTP configuration.

The system completes the configuration and displays the First Boot Menu.

```
[ 1] Apply settings
[ 2] Reset and start over
[ 3] Update Recovery password (*****)
[ 4] Update Hostname (controller-1)
[ 5] Update IP Option (IPv4 and IPv6)
[ 6] Update IPv6 Method (Automatic via SLAAC & DHCPv6)
[ 7] Update IPv4 Address (10.106.8.2/24)
[ 8] Update IPv4 Gateway (10.106.8.3)
[ 9] Update DNS Server 1 (10.3.0.4)
[10] Update DNS Server 2 (10.1.5.200)
[11] Update DNS Search Domain (qa.aristanetworks.com)
[12] Update Cluster Option (Start a new cluster)
[13] Update Cluster Name (dmf-cluster)
[14] Update Cluster Description (<none>)
[15] Update Admin Password (*****)
[16] Update NTP Option (Use default NTP servers)
[1] >
```

11. To apply the configuration, type 1.



Note: To change any previous configurations, type the appropriate menu option, make the required changes, and then type 1 to apply the settings when the system returns to this menu.

After entering the option to apply the settings, the system applies the settings, starts the DMF Controller, and displays the Controller banner and login prompt.

The system applies the settings and displays a series of prompts as each process completes. When successful, the system displays the following prompts.

```
[Stage 1] Initializing system
[Stage 2] Configuring controller
Waiting for network configuration
IP address on em1 is 10.106.8.3
Generating cryptographic keys
[Stage 3] Configuring system time
Initializing the system time by polling the NTP servers:
0.bigswitch.pool.ntp.org
1.bigswitch.pool.ntp.org
2.bigswitch.pool.ntp.org
3.bigswitch.pool.ntp.org
```

```
[Stage 4] Configuring cluster
Cluster configured successfully.
Current node ID is 10249
All cluster nodes:
Node 10249: fe80::1618:77ff:fe67:3f0c:6642
First-time setup is complete!
Press enter to continue >
DANZ Monitoring Fabric 8.4.0 (dmf-8.4.0 #11)
Log in as 'admin' to configure
controller-1 login:
```

To login to the Controller, use the account name `admin` and the password set for the cluster during installation. Optionally, log in to the active Controller using SSH with the assigned IP address or use a web browser to connect to the IP address using HTTPS.

```
controller login: admin
admin@10.106.8.2's password:
Login: admin, on Thu 2020-11-19 21:39:28 UTC, from 10.95.66.14
Last login: on Thu 2020-11-19 21:39:05 UTC, from 10.95.66.14
controller-1>
```

2.3 Configuring the Standby Controller

2.3.1 Joining Standby Controller to Existing Cluster

Arista Networks recommends deploying the DANZ Monitoring Fabric (DMF) in a two-node cluster for operational resilience. Set up the first (active) cluster node described above. Use the steps described here to configure a second Controller node and join it to the cluster as a standby node.

An active and standby Controller can be deployed in different L3 subnets, changing how Controllers communicate. Active and standby Controllers communicate using each Controller's management **IPv4** address. This configuration enables the provisioning of active and standby Controllers in different IP subnets (L3 domains) or on the same subnet (L2 domain).



Note: Both nodes in a cluster must be running the same version of DMF Controller software. The standby node will not join the cluster if it is not running the same version as the active node. The maximum latency between active and standby Controllers should be less than 300ms.

Powering on the hardware appliance or the VM with the pre-installed DMF software prompts the user to log in as `admin` for the first-time configuration.

1. Log in to the second appliance using the `admin` account.



Note: Ports on the Controller-facing management-switch port must come up within 5 seconds.

Powering on the hardware appliance for the active Controller prompts the user to press **Enter** to begin the installation.

2. Press **Enter** to begin the installation.
3. Log in to the standby Controller using the `admin` account.

Use the default account (**admin**) without a password when the system is booting from the factory default image. The system displays the following prompt to accept the End User License Agreement (EULA).

```
This product is governed by an End User License Agreement (EULA).
You must accept this EULA to continue using this product.
You can view this EULA from our website at:
https://www.arista.com/en/eula
Do you accept the EULA for this product? (Yes/No) [Yes] > Yes
To read the agreement, type ``View``
```

4. Accept the agreement.



Note: When entering the command: **No**, the system prompts the user to log out. Accepting the EULA is mandatory to install the product.

After typing **Yes** to accept the EULA, the system prompts for a password to be used for emergency recovery of the Controller node, as shown below.

To read the agreement, type **View**.

To accept the agreement, type **Yes**.

```
Starting first-time setup
Local Node Configuration
-----
Password for emergency recovery user >
Retype Password for emergency recovery user >
```

5. Set the recovery user password and confirm the entry.

```
Local Node Configuration
-----
Emergency recovery user password >
Emergency recovery user password (retype to confirm) >
Hostname > controller-2
```

6. Choose the IP forwarding mode, when prompted.

```
Management network options:
[1] IPv4 only
[2] IPv6 only
[3] IPv4 and IPv6
> 3
```

7. Choose the method for assigning IPv4 addresses.

```
Management network (IPv4) options:
[1] Manual
[2] Automatic via DHCP
[1] >
```



Note: When using **[2] Automatic via DHCP**, Arista Networks recommends reserving the DHCP address for the MAC address of the Controller's management port.

8. Choose the method for assigning IPv6 addresses.

```
Management network (IPv6) options:
1] Automatic via SLAAC & DHCPv6
2] Manual
1] >
```

The DMF Controller appliance allows IPv6 Stateless Address Auto-configuration for Controller IPv6 management.

9. (Optional) When using [1] **Manual** in Step 7, the system prompts the user to enter the IPv4 address and related configuration information for the Controller node.

Enter the IP address and optional information regarding the DNS server, as shown in this example.



Note: Entering the IP address followed by a slash (/) and the number of bits in the subnet mask, the system may not prompt for the CIDR.

```
IPv4 address [0.0.0.0/0] > 10.106.8.3/24
IPv4 gateway (Optional) > 10.106.8.1
DNS server 1 (Optional) > 10.3.0.4
DNS server 2 (Optional) > 10.1.5.200
DNS search domain (Optional) > qa.aristanetworks.com
```



Note: The IP address configuration is applied to the Controller Management Interface on the Controller appliance.

10. The system prompts the user wants to create a new cluster or join an existing cluster.

```
Controller Clustering
-----
Controller cluster options:
[1] Start a new cluster
[2] Join an existing cluster
> 2
```

Type **2** to join the standby Controller to the existing cluster.

11. The system prompts the user for the IPv4 address of the active Controller for the cluster the standby will join.

```
Existing Controller IP > 10.106.8.2
```

The user is prompted to enter the IP address of the active Controller for the cluster.

12. The user is prompted to enter and confirm the cluster password.

```
Cluster administrator password >
Cluster administrator password (retype to confirm) > Menu
```

13. The system displays the First Boot Menu and all provided details for final verification. Press **1** to apply the settings and start the process of the standby Controller joining the DMF cluster.

```
:: Please choose an option:
[ 1] Apply settings
[ 2] Reset and start over
[ 3] Update Recovery Password
[ 4] Update Hostname
[ 5] Update IP Option
[ 6] Update IPv4 Address
[ 7] Update IPv4 Gateway
[ 8] Update DNS Server 1
[ 9] Update DNS Server 2
[10] Update DNS Search Domain
[11] Update Cluster Option
[12] Update Cluster to Join
[13] Update Admin Password
[1] > 1
```

14. The system advises the current DMF cluster does not have **secure control** enabled. Enable this optional feature later, as required. Press **1** to resume connecting the standby Controller with the active Controller in the DMF cluster.

```
[Stage 1] Initializing system
[Stage 2] Configuring local node
Waiting for network configuration
IP address on ens4 is 10.2.0.185
Generating cryptographic keys
Please verify that:
Secure control plane is NOT configured.
You can verify the above by running "show secure control plane"
on the existing controller 10.2.1.90.
Options:
[1] Continue connecting (the above info is correct)
[2] Cancel and review parameters
> 1
[Stage 3] Configuring system time
Initializing the system time by polling the NTP servers:
0.bigswitch.pool.ntp.org
1.bigswitch.pool.ntp.org
2.bigswitch.pool.ntp.org
3.bigswitch.pool.ntp.org
[Stage 4] Configuring cluster
Cluster configured successfully.
Current node ID is 5302
All cluster nodes:
Node 5302: [fe80::5054:ff:fe6a:dd0f]:6642
Node 15674: [fe80::5054:ff:fec4:d1b8]:6642
First-time setup is complete!
Press enter to continue >
DANZ Monitoring Fabric 8.0.0 (dmf-8.0.0 #62)
Log in as 'admin' to configure
```



Note: It is possible to connect to the DMF Controller using the IP address assigned to either the active or standby Controller. However, you can make configuration changes only when connected to the active Controller. In addition, statistics and other information will be more accurate and up-to-date when viewed on the active Controller.

```
DANZ Monitoring Fabric 8.4.0 (dmf-8.4.0 #11)
Log in as 'admin' to configure
admin@10.106.8.3's password:
Login: admin, on Fri 2020-11-20 05:06:19 UTC, from 10.95.66.14
Last login: on Fri 2020-11-20 05:06:09 UTC, from 10.95.66.14
===== WARNING: STANDBY CONTROLLER =====
This controller node is in standby mode.
This session should only be used for for troubleshooting.
Log in to the active to make configuration changes
and to access up-to-date operational data.
=====
standby controller-2>
```

2.3.2 Moving existing Standby Controller to different IPv4 subnet

To move an existing standby Controller to a different subnet, the standby Controller management IPv4 address must be changed before moving the appliance to a different subnet. Before making any changes, ensure the underlying network configuration is set up correctly for connectivity.

1. Change the management IP address on the standby Controller. There are two methods to change the Controller management IP address.

Method A: Using iDRAC or Controller console, remove the existing management IP address and add the new IP address. Go to **config->>local node->>interface management->>ipv4** to change the Controller's management IP address.

```
Controller-2(config-local-if-ipv4)# no ip 192.168.39.44/24 gateway 192.168.39.1
```

```
192.168.39.44: currently in use: port 22: 192.168.39.1:54978, 192.168.39.1:54979
192.168.39.44: proceed ("y" or "yes" to continue): y
DMF-Controller-2(config-local-if-ipv4)# ip 192.168.39.45/24 gateway 192.168.39.1
```

Method B: Use the REST API to replace the Controller's management IP address instead of removing and adding a new management IP address. When using this method, there is no need to use Controller iDRAC or console access to change the management IP address. You can use the REST API and parameters below to replace the Controller's management IP address:

```
/api/v1/data/controller/os/config/local/network/interface[name="management "]/ipv4/address' -X
PUT -d '{"ip-address": "192.168.39.45", "prefix": 24, "gateway": "192.168.39.1"}'
```

2. Move the standby Controller to the new subnet.

2.3.3 Moving an Existing Standby Controller to a Different Controller Cluster

Follow the procedure below to move an existing standby Controller to a different Controller cluster.

1. Remove the standby Controller from the existing cluster.
2. Connect to the Controller using iDRAC or Controller console.
3. Run the `boot factory-default` command.

```
Controller-2# boot factory-default
```

4. Perform the first-boot operation specified in the section [Joining Standby Controller to Existing Cluster](#) ** if joining an existing new Controller cluster or [Configuring the Active Controller Using the First Boot Script](#) to create a new Controller cluster.

2.4 Accessing the DANZ Monitoring Fabric Controller

This section describes connecting to the DANZ Monitoring Fabric (DMF) Controller.

To access the active DMF Controller, use the IP address of the active Controller. If configured, use the cluster's Virtual IP (VIP) as described in the [Configuring the Cluster Virtual IP](#) section.

Refer to the [Using Local Groups and Users to Manage DMF Controller Access](#) section to manage administrative user accounts and passwords.

2.4.1 Using the DANZ Monitoring Fabric CLI

Once the DANZ Monitoring Fabric (DMF) Controllers are up and running, log in to the DMF Controller using the VM local console or SSH.

DMF divides CLI commands into modes and sub-modes, which restrict commands to the appropriate context. The main modes 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 are used after entering the configured command. Access sub-modes from this mode.

Enter sub-modes from config mode to provision specific monitoring fabric objects. For example, the `switch` command changes the CLI prompt to **(config-switch)#** to configure the switch identified by the switch DPID or alias.

After logging in, the CLI appears in the login mode where the default prompt is the system name followed by a greater than sign (>), as in the following example.

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

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

To return to the enable mode, enter the `end` command, as shown below.

```
controller-1 (config) # end
controller1#
```

To view a list of the commands available from the current mode or submode, enter the `help` command. To view detailed online help for the command, enter the `help` command followed by the command name.

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

2.4.2 Capturing CLI Command Output

Pipe commands through other commands like `grep` for analysis and troubleshooting DANZ Monitoring Fabric (DMF) operations. View the contents of the output files by entering the `show running-config` command, as in the example below.

```
controller-1> show running-config | grep <service unmanaged-service TSTS>
post-service pst2
pre-service pst
```

Copy files to an external FTP server, as in the example below.

```
copy running-config scp://<username@scp_server>///<file>
```

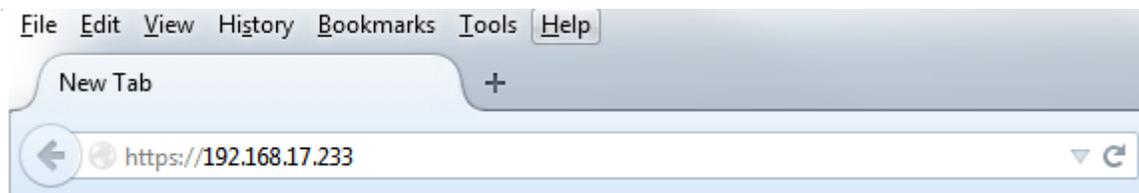
2.4.3 Using the DANZ Monitoring Fabric GUI

The DANZ Monitoring Fabric (DMF) GUI performs similar CLI operations using a graphic user interface instead of text commands and options. DMF supports the use of the GUI with recent versions of any of the following supported browsers:

- Firefox
- Chrome
- Internet Explorer
- Safari
- Microsoft Edge

To connect to the DMF GUI, use the IP address assigned to the DMF Controller. The following figure shows a browser connecting to the GUI using HTTPS at the IP address **192.168.17.233**.

Figure 2-1: Accessing the DANZ Monitoring Fabric GUI



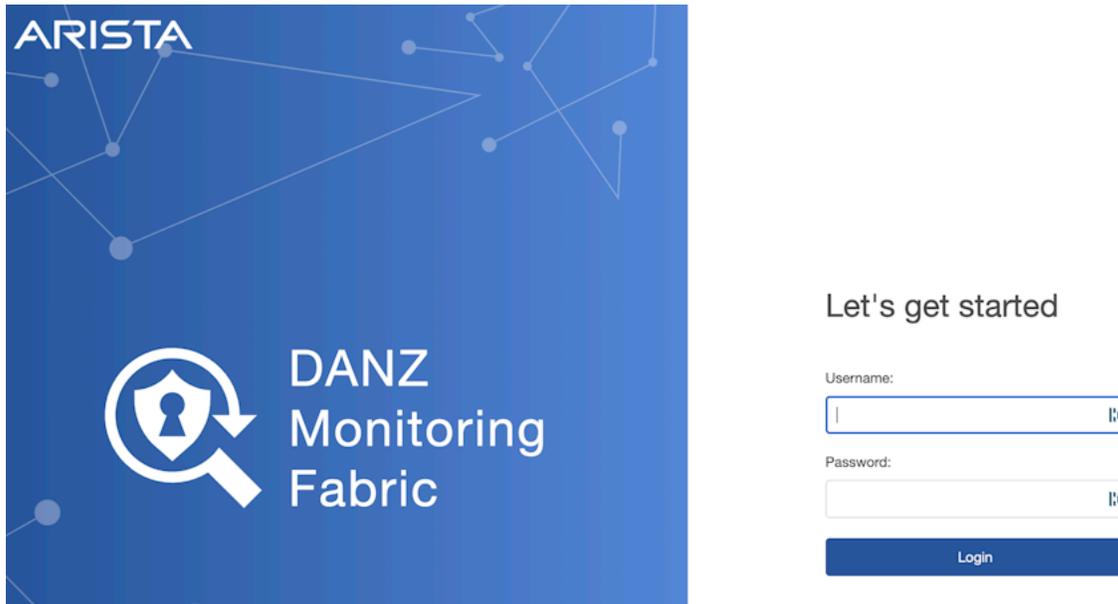
When connecting to the Controller for the first time, a prompt requesting a security exception may appear because the Controller HTTPS server uses an unknown (self-signed) certificate authority.



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

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

Figure 2-2: DANZ Monitoring Fabric GUI Login Prompt



Use the admin username and password configured for DMF 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.

The main menu for the DMF GUI appears in the following figure.

Figure 2-3: 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.

The newly designed dashboard displays information about the Controller, including switches, interfaces, policies, and Smart Nodes.

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

2.5 Managing the DMF Controller Cluster

This section describes configuring settings for a DANZ Monitoring Fabric (DMF) cluster or both active and standby Controllers. Most configuration occurs on the active Controller, which synchronizes the standby Controller.

2.5.1 Verifying Cluster Configuration

A DANZ Monitoring Fabric (DMF) Out-of-Band HA cluster consists of two Controller nodes, one active and one standby. Keep the following conditions in mind when configuring the cluster:

- Both active and standby must be in the same IP subnet.
- Firewall configurations are separate for active and standby, so manually keep the configuration consistent between the two nodes.
- NTP service is required to establish a cluster. Starting with **DMF 7.0.0**, the active Controller provides the NTP service for the cluster and connected switches. The Controller should sync the time from an external NTP server for this service to work.
- When SNMP service is enabled, it must be manually configured to be consistent on both nodes. To verify the cluster state, use the following commands:
- Enter the **show controller details** command from either the active or standby Controller.

```
controller-1(config)# show controller details
Cluster Name : dmf-cluster
Cluster UID : 8ef968f80bd72d20d30df3bc4cb6b271a44de970
Cluster Virtual IP : 10.106.8.4
Redundancy Status : redundant
Redundancy Description : Cluster is Redundant
Last Role Change Time : 2020-11-19 18:12:49.699000 UTC
Cluster Uptime : 3 weeks, 1 day
# IP @ Node Id Domain Id State Status Uptime
- |-----|-----|-----|-----|-----|
1 10.106.8.3 5784 1 standby connected 11 hours, 6 minutes
2 10.106.8.2 * 25377 1 active connected 11 hours, 11 minutes
-----|-----|-----|-----|-----|
# New Active Time completed Node Reason Description
- |-----|-----|-----|-----|-----|
1 25377 2020-11-19 18:12:38.299000 UTC 25377 cluster-config-change Changed connection state: cluster configuration changed
controller-1(config)#
```

- Enter the **show controller** command to display the current Controller roles from either the active or standby node, as in the following example.

```
controller-1(config)# show controller
Cluster Name : dmf-cluster
Cluster Virtual IP : 10.106.8.4
Redundancy Status : redundant
Last Role Change Time : 2020-11-19 18:12:49.699000 UTC
Failover Reason : Changed connection state: cluster configuration changed
Cluster Uptime : 3 weeks, 1 day
# IP @ State Uptime
- |-----|-----|-----|-----|
1 10.106.8.3 standby 11 hours, 24 minutes
2 10.106.8.2 * active 11 hours, 29 minutes
controller-1(config)#
```

2.5.2 Configuring the Cluster Virtual IP

Setting up the Virtual IP (VIP) for the cluster is a best practice for connecting to the management port of the active node using an IP address that does not change even if the active Controller fails over and the role of the standby Controller changes to active.



Note: VIP will not work if active and standby Controllers are in different IPv4 subnets. The active and standby Controllers have to be in the same L2 domain.

On the active Controller, enter the `virtual-ip` command from the **config-controller** submode.

```
controller-1> enable
controller-1# config
controller-1(config)# controller
controller-1(config-controller)# virtual-ip 10.106.8.4
controller-1(config-controller)#
```

Verify the VIP by entering the `show controller` command.

```
controller-1(config)# show controller
Cluster Name : dmf-cluster
Cluster Virtual IP : 10.106.8.4
Redundancy Status : redundant
Last Role Change Time : 2020-11-19 18:12:49.699000 UTC
Failover Reason : Changed connection state: cluster configuration changed
Cluster Uptime : 3 weeks, 1 day
# IP      @ State  Uptime
- |-----| - |-----|-----|
1 10.106.8.3   standby 11 hours, 24 minutes
2 10.106.8.2 * active 11 hourmvs, 29 minutes
controller-1(config)#
```



Note: Make sure you use a unique IP address for the VIP of the cluster. If you use the IP of the standby Controller by mistake, the nodes will form separate clusters. To resolve this problem, assign a unique IP address for the VIP.

2.5.3 Setting the Time Zone

It is essential that the switches, Controllers, hypervisors, VMs, and management systems in the DANZ Monitoring Fabric (DMF) have synchronized system clocks and all use the same time zones.



Note: The hypervisor and the virtual machine running the Controller using different time zones may cause problems with log files or access to the DMF GUI.

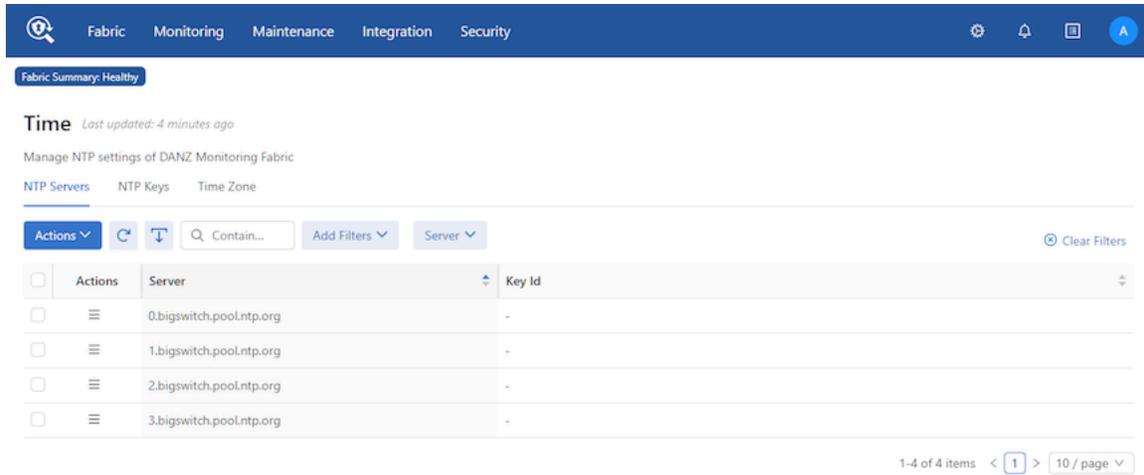
To view or change the current time zone on the DMF Controller, complete the following steps.

GUI Procedure

Select Maintenance > Time from the main menu.

The dashboard displays the last updated time on the system, information about the configured NTP server, and provides an option for refreshing the NTP synchronization with the NTP server running on the DMF Master Controller.

Figure 2-4: Maintenance Time

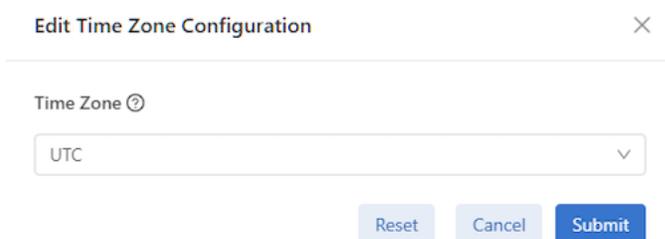


Time configuration tabs and features include:

- NTP Servers
- NTP Keys
- Time Zone

Select the **Time Zone** tab and click **Edit**.

Figure 2-5: Edit Time Zone Configuration

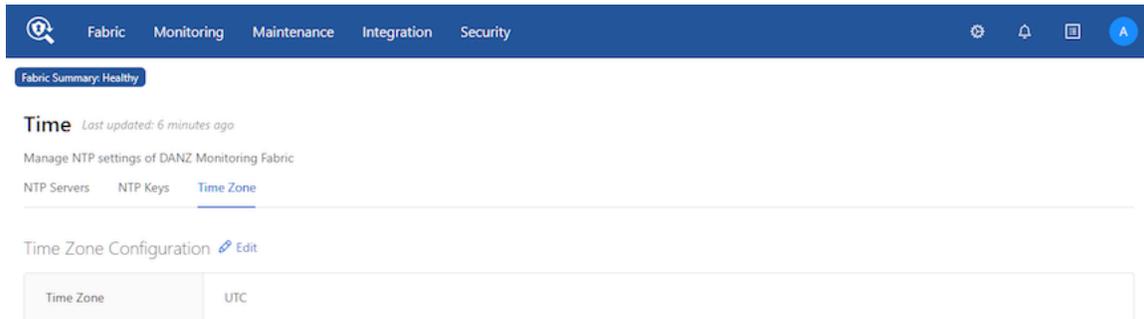


Select the desired time zone from the drop-down list and click **Submit**.

i **Tip:** Use the **Reset** button to clear a time zone entry and use the default UTC.

The dashboard displays the selected time zone.

Figure 2-6: Time Zone



CLI Procedure

To set the time zone for the Controller from the config-controller submode, use the `set timezone` command:

```
[no] ntp time-zone <time-zone>
```

Replace **time-zone** with the specific string for the desired time zone. To see a list of supported values, press **Tab** after the `clock set` command. Certain values, such as **America/**, are followed by a slash (/). These values must be followed by a keyword for the specific time zone. View the supported time zones by pressing **Tab** again after selecting the first keyword.

For example, the following commands set the time zone for the current Controller to Pacific Standard Time (PST):

```
controller-1 (config)# ntp time-zone America/Los_Angeles
Warning: Active REST API sessions will not be informed of updates to time-zone.
Warning: Please logout and login to any other active CLI sessions to
Warning: update the time-zone for those sessions.
controller-1 (config)#
```



Note: Changes in time zone are logged to the `floodlight.log`.

The following command removes the manually configured time zone setting and resets the Controller to the default (UTC).

```
controller-1(config-controller)# no ntp time-zone
```

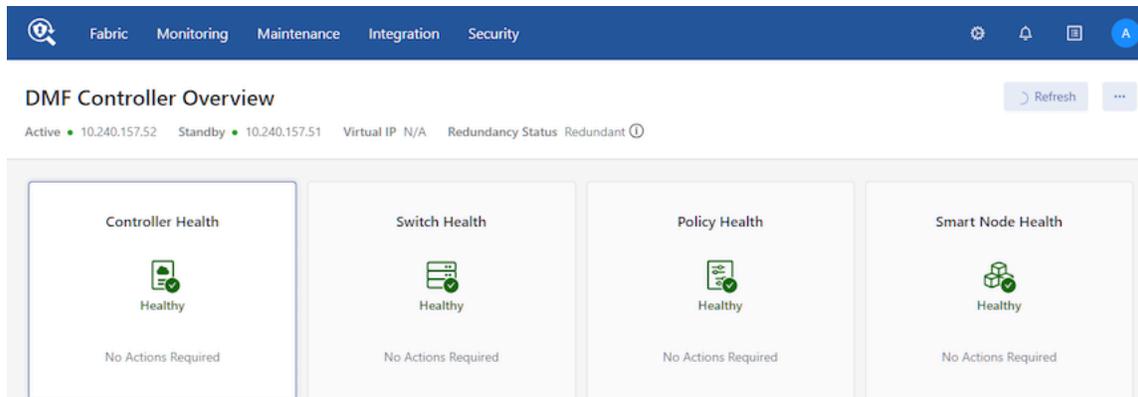


Warning: Manually setting the clock for the DMF Controller using the `clock set` command may affect database reconciliation between the nodes in a Controller cluster. Because time stamps are used to identify records that should be updated, Arista Networks recommends adjusting any time skew using the `ntpdate` command. Using an NTP server ensures accurate synchronization between the Controller clocks.

2.5.4 Viewing Controller and Cluster Status

To view the overall Controller status, click the DANZ Monitoring Fabric (DMF) logo in the left corner of the GUI Main menu. The system displays the DMF GUI landing page, shown in the following figure.

Figure 2-7: Cluster and Controller Status



This page also provides configuration options that let you view and change fabric-wide settings.

To view more details about the DMF cluster, select **Maintenance > Cluster**. Cluster configuration tabs and features include:

- Cluster
 - Cluster Information
 - Node Information
 - Node Config State
 - Node Operational State
 - Failover History
- Local Node
 - Local Node Interface - Management
 - Interfaces
 - Addresses of Interfaces

- DNS Servers

Figure 2-8: Cluster Information - Node Information

The screenshot shows the DMF Controller interface with a navigation bar at the top containing 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. The main content area is titled 'Cluster' and includes a 'Refresh' button. Below the title, there are tabs for 'Cluster' and 'Local Node'. The 'Cluster Information' section displays the following details:

- Cluster Name: cliff-dmf
- Virtual IP: -
- Last Role Change Time: Jan 31, 2024 10:49AM
- Config Sync Status: In Sync
- UID: c9e8f5ceb0b5a110b2dad7d79527dc9bed75794a
- Cluster Description:
- Redundancy Status: Redundant
- Cluster Uptime: 1 month 12 days 23 hours 23 minutes
- Operational Sync Status: In Sync

The 'Node Information' section contains a table with the following data:

Host Name	IP Address	Node ID	Role	Status	Uptime	Actions
cliff-dmf-1	10.240.157.51	20296	Standby	Connected	1 month 12 days 23 hours 18 minutes	CU Web
cliff-dmf-2	10.240.157.52	18030	Active	Connected	1 month 12 days 23 hours 23 minutes	CU Web

Figure 2-9: Cluster - Node Config State, Operational State, Failover History

The screenshot displays three sections related to cluster nodes:

Node Config State

Node ID	Delay From Active	Digest Value	Role	Update Time
20296	38 milliseconds	8316DCFBF6B812DCE43D0F2F8672DDD7E3566ED	Standby	Mar 15, 2024 10:35AM
18030	-	8316DCFBF6B812DCE43D0F2F8672DDD7E3566ED	Active	Mar 15, 2024 10:35AM

Node Operational State

Node ID	Delay From Active	Digest Value	Role	Snapshot ID	Update Time
20296	37 milliseconds	8316DCFBF6B812DCE43D0F2F8672DDD7E3566ED	Standby	[18030,32]	Mar 15, 2024 10:35AM
18030	-	8316DCFBF6B812DCE43D0F2F8672DDD7E3566ED	Active	[18030,32]	Mar 15, 2024 10:35AM

Failover History

Active Node	Description	Reason	Node ID	Time Completed
18030	Changed connection state: cluster configuration changed	Cluster Config Change	18030	Jan 31, 2024 10:49AM

Select the **Local Node** tab to view local node details.

Figure 2-10: Cluster - Local Node

Local Node Interface - Management

Host Name: diff-dmf-2 [Edit](#)

IPv4 Configuration: Method Manual, Address 10.240.157.52/25 gateway: 10.240.157.1, DNS Server 10.240.48.6 and 10.242.35.6, DNS Search Path qa.aaa.com,aaa.com

IPv6 Configuration: Method Manual, DNS Server 0 servers, DNS Search Path 0 search paths

Interface	Hardware Address	Permanent Hardware Address	Master	Operstate	Carrier	Bond Mode	Bond Role
bond0	08:92:04:9a:17:3a	-	-	Up	Up	Active Backup	-
eno8303	08:92:04:9a:17:3a	08:92:04:9a:17:3a	bond0	Up	Up	-	Active
eno8403	08:92:04:9a:17:3a	08:92:04:9a:17:3b	bond0	Down	Down	-	Backup
ens1f0	6cfe54:5b4a:a0	6cfe54:5b4a:a0	-	Down	Down	-	-
ens1f1	6cfe54:5b4a:a1	6cfe54:5b4a:a1	-	Down	Down	-	-
ens3f0	b483:51:06:0bae	b483:51:06:0bae	-	Down	Down	-	-
ens3f1	b483:51:06:0baf	b483:51:06:0baf	-	Down	Down	-	-

Addresses of Interfaces

Interface	IP CIDR
bond0	10.240.157.52/25 fe80:0:0:a92:4ff:fe9a:173a/64
eno8303	-
eno8403	-
ens1f0	-
ens1f1	-
ens3f0	-
ens3f1	-

Routes

Destination	Gateway	Interface
0.0.0.0	10.240.157.1	bond0
10.240.157.0/25	-	bond0
fe80:0:0:0:0:0:0:0/64	-	bond0

DNS Servers

CLI Procedure

View the overall Controller status by entering the **show controller** command.

```

controller-1(config)# show controller details
Cluster Name : dmf-cluster
Cluster UID : 8ef968f80bd72d20d30df3bc4cb6b271a44de970
Cluster Virtual IP : 10.106.8.4
Redundancy Status : redundant
Redundancy Description : Cluster is Redundant
Last Role Change Time : 2020-11-19 18:12:49.699000 UTC
Cluster Uptime : 3 weeks, 1 day
# IP @ Node Id Domain Id State Status Uptime
- |-----|-----|-----|-----|-----|-----|
1 10.106.8.3 5784 1 standby connected 11 hours, 6 minutes
2 10.106.8.2 * 25377 1 active connected 11 hours, 11 minutes
-----|-----|-----|-----|-----|-----|
# New Active Time completed Node Reason Description
- |-----|-----|-----|-----|-----|
1 25377 2020-11-19 18:12:38.299000 UTC 25377 cluster-config-change Changed connection state:
cluster configuration changed
controller-1(config)#

```

Use the following commands to modify the Controller configuration:

- **access-control1**: Configure access control of the Controller

- **cluster-name**: Configure cluster name
- **description**: Configure cluster description
- **virtual-ip**: Configure management virtual IP

To modify the hostname, IPv4 or v6 addresses, or SNMP server configuration, use the following commands:

- **hostname**: Configure hostname for this host
- **interface**: Configure controller network interface
- **snmp-server**: Configure local SNMP attributes

2.5.5 Saving and Restoring Controller Configuration

A properly configured Controller requires no regular maintenance. However, Arista Networks recommends periodically saving a copy of the Controller running-config to a location outside the Controller. Use the **copy** command from the Controller to an FTP server.

```
# copy running-config scp://admin:admin@myserver/configs
```

The format for the remote server username, password, and path name is as follows: **scp://<user>:<password>@<host>/<path_to_saved_config>**

To restore the Controller configuration from a remote server, remove the standby Controller, then boot factory default on the active node to start a new cluster.

CLI Procedure

1. Remove the standby node if present.

```
# show controller
Cluster Name : desktop-DMF
Cluster Virtual IP : 10.100.6.4
Redundancy Status : redundant
Last Role Change Time : 2018-10-01 10:43:51.582000 CDT
Failover Reason : Changed connection state: connected to node 15444
Cluster Uptime : 3 months, 2 weeks
# IP @ State Uptime
-|-----|-----|-----|-----|
1 10.100.6.3 * active 4 days, 2 hours
2 10.100.6.2 standby 4 days, 2 hours
# system remove node 10.100.6.2
```

2. Reboot the DANZ Monitoring Fabric (DMF) Controller to the factory default by entering the following commands:

```
controller-1> enable
controller-1# boot factory-default
Re-setting controller to factory defaults...
Warning: This will reset your controller to factory-default state and reboot it.
You will lose all node/controller configuration and the logs
```

3. Confirm the operation and enter the administrator password.

```
Do you want to continue [no]? yes ...
Removing existing log files ... Resetting system state ...
Current default time zone: 'Etc/UTC' ...
Enter NEW admin password:
...
boot: ...
localhost login:
```

- Rerun the first-time setup again to reconfigure the Controller.
- Copy the saved running config from the remote server to the DMF Controller by entering the following command.

```
# copy scp://<user>:<password>@<host>/<path_to_saved_config> running-config
```

- Restore the cluster.
Remove the standby Controller, then boot factory default on the active node to start a new cluster.



Note: With the introduction of the Command Line Interface (CLI) to Managed Appliances (Service Node / Recorder Node), the CLI command `boot factory-default` has been extended to launch the first-boot setup process.

2.5.6 Reloading and Rebooting a Controller Using the DMF UI

To use the **Reload** or **Reboot** functions, navigate to the DANZ Monitoring Fabric (DMF) **Maintenance -> Cluster** page. **Reload** and **Reboot** appear under the Actions column in the Node Information table.



Note: As a precaution, DMF disables these functions by default. To enable them, click **Unlock Power Actions**.

To use the **Reload** and **Reboot** features, a logged-in DMF user requires read/write sysops permissions on the active cluster node.

For an active/standby cluster configuration, the user can reload or reboot the node.

When the active node is reloaded or rebooted, the standby node becomes the new active node. When the user logs in to the Virtual IP (VIP) address of an active/standby cluster configuration, the user can reload or reboot the active node. The VIP logs in to the new active node without the user changing IP addresses.

The new standby node will not have the Reload or Reboot functionality unless the user directly logs in to that node.

Figure 2-11: Cluster Page with Reload and Reboot Actions

The screenshot shows the DMF Cluster page. At the top, there's a 'Cluster' header and a 'Refresh' button. Below that, there's a sub-header 'View and manage cluster and local node information for DANZ Monitoring Fabric'. The page is divided into two tabs: 'Cluster' (selected) and 'Local Node'. Under the 'Cluster' tab, there's a 'Cluster Information' section with details like Cluster Name, Virtual IP, Last Role Change Time, Config Sync Status, and UD. To the right, there's a 'Cluster Description' and 'Operational Sync Status'. Below this is a 'Node Information' table with columns for Host Name, IP Address, Node ID, Role, Status, Uptime, and Actions. The table contains one row for 'je-condescendingbranch-horry-c1' which is 'Currently logged in'. The Actions column for this row includes buttons for 'CLI', 'View', 'Reload', and 'Reboot'. There is also a 'Lock Power Actions' button on the right side of the table.

Workflow

The following information describes the expected behavior of the **Reload** and **Reboot** functions based on the type of configuration:

Action	Single Node
Reload	Click the submit button -> Display submit button loading -> Redirects to login page.
Reboot	Click the submit button -> Display submit button loading -> Redirects to the login page.

Action	Two Node Cluster without Virtual IP (VIP)	
	Log into Active	Log into Standby
Reload	Click the submit button -> Display submit button loading -> Redirect to the login page (<u>it doesn't redirect because the session is still active</u>). This node now becomes the standby node).	Click the submit button -> Display submit button loading -> Redirect to the login page (<u>it doesn't redirect because the session is still active</u> . This node will remain the standby node).
Reboot	Same as above.	Same as above.

Action	Two Node Cluster with Virtual IP (VIP)		
	Log into VIP	Log into Active	Log into Standby
Reload	Click the submit button -> Display submit button loading -> Redirect to the login page. The active and standby nodes are swapped. The VIP has access to the new active node.	Same behavior as active node without VIP.	Same behavior as the standby node without VIP.
Reboot	Same as above.	Same behavior as active node without VIP.	Same behavior as the standby node without VIP.

UI Notifications

The standby banner in the UI displays a link to the active node. This link appears when a user is logged in to the standby node.

Limitations

1. Executing the power action on a node is only possible if the user logs in to that node. For example, if the user logs in to the active node, they will only be able to execute power actions on the active node; they will have to log in to the standby node to execute power actions on the standby node.
2. The CLI supports a third power action: **shutdown**. The UI does not include shutdown support because of the previous limitation. When executing a shutdown on the currently logged-in node, the user loses connection with the UI and cannot view the updated status of the node.
3. The CLI supports the **reload** and **reboot** functionalities for both the active and the standby nodes.

2.6 Copying Files Between a Workstation and a DMF Controller

Use the `SCP` command followed by the keywords shown below on a workstation connected to the DANZ Monitoring Fabric (DMF) Controller management network to copy the following types of files to the Controller:

- Certificate (`//cert`)
- Private key (`//private-key/<name>`)
- Running configuration (`//running-config`)
- Snapshots (`//snapshot`)
- Controller image files (`//image`)

```
Copying into //snapshot on the controller overwrites the current running-config except the local
node configuration
Copying into //running-config merges the destination configuration to the running-config on the
controller
```

To copy, use the following syntax:

```
scp -O <filename> admin@<controller-ip>://<keyword>
```



Note: The `-O` option forces the use of the legacy SCP protocol for file transfers instead of the SFTP protocol.

For example, the following command copies a Controller ISO image file from a workstation to the image file partition on the Controller to install it.

```
scp -O DMF-8.0.0-Controller-Appliance-2020-12-21.iso admin@10.2.1.183://image
```

This example copies the DMF 8.0.0 ISO file from the local workstation or server to the image partition of the DMF Controller running at **10.2.1.183**.

Use any user account belonging to the admin group. Replace the **admin** username in the example above with any other admin-group user, as in the following example.

```
c:\>pscp -P 22 -scp DMF-8.0.0-Controller-Appliance-2020-12-21.iso admin-upgrade@10.2.1.183://image
```

This example uses the user account **admin-upgrade**, which should be a member of the admin group. Use the `PSCP` command on a Windows workstation to copy the file to the Controller.

```
c:\>pscp -P 22 -scp <filename> admin@<controller-ip>://<keyword>
```

Use the `SCP` command to get the following files from the Controller and copy them to the local file system of the server or workstation.

- Running configuration (`copy running-config <dest>`)
- Snapshots (`copy snapshot:// <dest>`)

Use the following commands to copy running-config/snapshot to the remote location:

```
controller-1# copy snapshot:// scp://root@10.100.40.10://anet/DMF-720.snp
controller-1# copy running-config scp://root@10.100.40.10://anet/DMF-720.cfg
```

2.6.1 Merge and Replace Parameters of Copy Config

When copying a *text-based* configuration file with the `copy <text config> running-config` command, it is possible to specify either of two parameters: `merge` or `replace`.



Note: The default parameter is `merge`.

Replace:

The `replace` option applies a text-based configuration as a full override of the existing one. In other words, it erases the old configuration and applies the new one to a blank slate.

For example, if `current_policy.txt` is the current running-config on the Controller and `policy_modified.txt` is a configuration file with modified NTP time zone, NTP servers, recorder-node name, and other policy changes, it is possible first to verify the differences and then to replace the current configuration with a modified configuration file by using the following commands:

```
dmf-controller-1# compare file://current_policy.txt file://policy_modified.txt
3c3
< ! Current Time: 2024-06-03 22:48:14 UTC
---
> ! Current Time: 2024-06-03 22:44:04 UTC
8,12c8,10
< ntp time-zone America/Los_Angeles
< ntp server ntp1.aristanetworks.com
< ntp server ntp2.aristanetworks.com
< ntp server ntp3.aristanetworks.com
< ntp server ntp4.aristanetworks.com
---
> ntp time-zone UTC
> ntp server time.google.com
> ntp server time.windows.com
79c77
< recorder-node device rn1
---
> recorder-node device rn-modified
103c101
< description testone
---
> description testone-modified
134,135c132,133
< delivery-interface leaf-1-eth3
< filter-interface ixia-5-6
---
> delivery-interface leaf-1-eth3-modified
> filter-interface ixia-5-6-modified

dmf-controller-1# copy file://policy_modified.txt running-config replace
Applied Command: 4: ntp time-zone UTC:
Warning: Active REST API sessions will not be informed of updates to time-zone.
Applied Command: 4: ntp time-zone UTC:
Warning: Please logout and login to any other active CLI sessions to
Applied Command: 4: ntp time-zone UTC:
Warning: update the time-zone for those sessions.
Lines applied: 142, No Errors, Warnings: 3; Commit completed
```

```

dmf-controller-1# show run ntp

! ntp
ntp server time.google.com
ntp server time.windows.com

dmf-controller-1# show run ntp details

! ntp
ntp time-zone UTC
ntp server time.google.com
ntp server time.windows.com

dmf-controller-1# show run recorder-node

! recorder-node
recorder-node device rn-modified
  mac 52:54:00:59:57:b0

dmf-controller-1# show run policy | grep modified
description testone-modified
delivery-interface leaf-1-eth3-modified
filter-interface ixia-5-6-modified

```

Merge:

The `merge` option applies a text-based configuration *additively* to an existing one. In other words, it adds a new configuration *delta* to the current configuration. In the case of the previous example, using the `merge` option would have failed because it would have added a conflicting Recorder Node with a duplicate MAC address, as shown below:

```

dmf-controller-1# copy file://policy_modified.txt running-config merge
Applied Command: 4: ntp time-zone UTC:
Warning: Active REST API sessions will not be informed of updates to time-zone.
Applied Command: 4: ntp time-zone UTC:
Warning: Please logout and login to any other active CLI sessions to
Applied Command: 4: ntp time-zone UTC:
Warning: update the time-zone for those sessions.
Error: Validation failed: Multiple devices (switch/service) specify 52:54:00:59:57:b0
as mac address
Lines applied: 142, No Errors, Warnings: 0; Commit: failed, changes rolled back

dmf-controller-1# show run recorder-node

! recorder-node
recorder-node device rn1
  mac 52:54:00:59:57:b0

```

If you wanted to apply just a one-line configuration change (let's say, change only the NTP time zone), the one-line change would work with the `merge` option but not with the `replace` option, since `replace` requires a full valid configuration to be applied, as shown below:

```

dmf-controller-1# show run ntp details

! ntp
ntp time-zone UTC
ntp server ntp1.aristanetworks.com
ntp server ntp2.aristanetworks.com
ntp server ntp3.aristanetworks.com

```



```
*****
*****Error: Validation failed: deletion of predefined admin group is not allowed.
Lines applied: 1, No Errors, Warnings: 0; Commit: failed, changes rolled back
```

Instead, using the (default) merge option to incrementally add a line would succeed, as shown below:

```
dmf-controller-1# show run ntp details

! ntp
ntp time-zone UTC
ntp server ntp1.aristanetworks.com
ntp server ntp2.aristanetworks.com
ntp server ntp3.aristanetworks.com
ntp server ntp4.aristanetworks.com

dmf-controller-1# show file test.txt
ntp time-zone America/Los_Angeles

dmf-controller-1# copy file://test.txt running-config
300.0% |*****
*****
*****
*****
*****
*****
*****Applied Command: 3: ntp time-zone America/Los_Angeles:
300.0% |*****
*****
*****
*****
*****
*****Warning: Active REST API sessions will not be informed of updates to time-zone.
300.0% |*****
*****
*****
*****
*****Applied Command: 3: ntp time-zone America/Los_Angeles:
300.0% |*****
*****
*****
*****
*****Warning: Please logout and login to any other active CLI sessions to
300.0% |*****
*****
*****
*****
*****Applied Command: 3: ntp time-zone America/Los_Angeles:
300.0% |*****
*****
```

Read:

View the contents of the snapshot named **testsnap**. The API action isn't quite the same as the CLI command, since the CLI gives the snapshot contents translated into running-config style CLI commands whereas the API produces the raw, untranslated snapshot data.



Note: Using this API, the user can save the snapshot to a local file on a server. This is equivalent to **copy snapshot://<filename> scp://<remote file location>**.

To dump the content of snapshot file.

```
curl -g -k -H "Cookie: session_cookie=<session-cookie>" https://<Controller IP>:8443/api/v1/snapshot/testsnap/data
```

To save the content of snapshot to another file.

```
curl -g -k -H "Cookie: session_cookie=<session-cookie>" https://<Controller IP>:8443/api/v1/snapshot/testsnap/data --output testfile.snp
```

Above curl example reads the **testsnap** snapshot file from **Controller IP** and writes to a file named **testfile.snp**.

Apply:

Apply the snapshot named **testsnap** to the controller.

```
curl -g -k -H "Cookie: session_cookie=<session-cookie>" https://<Controller IP>:8443/api/v1/rpc/controller/snapshot/apply -d '{"name": "snapshot://testsnap"}'
```

List:

List all snapshots on a controller.

```
curl -g -k -H "Cookie: session_cookie=<session-cookie>" https://<Controller IP>:8443/api/v1/data/controller/snapshot/available
```

Delete:

Delete a snapshot named **testsnap** from the controller.

```
curl -g -k -H "Cookie: session_cookie=<session-cookie>" 'https://<Controller IP>:8443/api/v1/data/controller/snapshot/available[name="testsnap"]'
```

2.6.3 Convert a Text-based running-config into a Snapshot

A keyword is added to the **copy running-config snapshot://sample** command to convert **text running-config** commands into a JSON snapshot.

Use the keyword **transaction** to perform the conversion. The keyword can also create a snapshot with specific commands included.

While similar, the following workflows describe two applications of the `copy running-config snapshot://sample` command and the new `transaction` keyword.

Workflow - Create a Snapshot

Create a snapshot using the following command:

```
> copy file://text-commands snapshot://new-snapshot
```

Use this choice to convert a collection of text commands into a working snapshot. The resulting snapshot has several advantages over the collection of text commands:

- Snapshots have version details as part of the file format, allowing the configuration to be rewritten based on changes in the new versions of the product.
- REST APIs post the configuration in large chunks, applying snapshots more quickly. A single text command typically updates a specific amount of configuration parameters while writing the resulting modification to the persistent storage.

The conversion process will:

- Create a new transaction with all the configuration parameters removed.
- Replay each command and apply the configuration changes to the active transaction.
- Invoke the snapshot REST API, which builds a snapshot from the current transaction.
- Delete the current transaction, preventing any of the applied configuration changes from the replayed command from becoming active.



Note: The conversion requires that the syntax of the text command to be replayed works with the currently supported syntax where it will be applied.

Workflow - Create a Snapshot containing a Specific Configuration

Manually create a snapshot that contains a specific configuration using the following steps.

1. Enter the configuration mode that supports changes to the configuration.
2. Create a new transaction using an empty or a current configuration by running one of the following command options.

```
# begin transaction erase
```

```
# begin transaction append
```

Configuration changes applied while the transaction is active are performed against the transaction but do not update the configuration until the transaction is committed (using the `commit transaction` command). Several validation checks are postponed until after committing the changes. The commit does not post if validation errors are present.



Note: In this sample workflow, the objective is to create a new snapshot and not to update the system's configuration.

Various changes to the configuration include:

- Add a new configuration, for example, new switches, new policies, new users, etc.
- Modify the configuration.
- Delete the configuration.
- The local configuration should not be changed, as transactions do not currently manage it. The local system configuration (for example, the hostname) is updated immediately.

The new **transaction** keyword can be used with the **copy** command, requesting that the configuration within the transaction be copied or applied to the snapshot and not to the current system configuration. For example, using the following command:

```
# copy running-config snapshot//:sample transaction
```

Delete the transaction using the following command:

```
# delete transaction
```



Note: The snapshot is updated while the system configuration is not updated.

To check the active transaction on the system, use the following command:

```
# show transaction
```

Sample Sessions

The examples below will familiarize the reader with converting a text-based running-config into a snapshot.

Example One

Convert a collection of text commands into a working snapshot using the following command:

```
> copy file://text-commands snapshot://new-snapshot
Controller1(config)# show file
# Name      Size Created
-|-----|----|-----|
1 textcommands      3560 2023-08-15 21:37:44 UTC
Controller1(config)# copy file://textcommands snapshot://snap_textcommands
Lines applied: 175, snap_textcommands: Snapshot updated
Controller1(config)# show snapshot snap_textcommands
```



Note: Some command options listed below are hidden in the CLI since they are in beta in DMF 8.4.0.

Example Two

Manually create a snapshot containing a specific configuration to be added to the existing running configuration. The commands **begin transaction** or **begin transaction append** add the commands executed on the transaction to the existing running configuration.

```

Controller1(config)# begin transaction
id : 5gMoJ6uu7mcA3j3Gs5fGyLRT8ZJW7CI9
Controller1(config)#
Controller1(config)#
Controller1(config)# policy policy15
Controller1(config-policy)# action forward
Controller1(config-policy)#
Controller1(config-policy)# exit
Controller1(config)# copy running-config snapshot://snap_policy15 transaction
Controller1(config)#
Controller1(config)# delete transaction
Controller1(config)# show snapshot snap_policy15
!
! Saved-Config snap_policy15
! Saved-Config version: 1.0
! Version: 8.4.0
! Appliance: bmf/dmf-8.4.x
! Build-Number 133
!
version 1.0
! ntp
ntp server ntp.aristanetworks.com
! snmp-server
snmp-server enable traps
snmp-server community ro 7 02161159070f0c
! local
local node
hostname Controller1
interface management
!
ipv4
ip 10.93.194.145/22 gateway 10.93.192.1
method manual
!
ipv6
method manual
! deployment-mode
deployment-mode pre-configured
! user
user admin
full-name 'Default admin'
hashed-password method=PBKDF2WithHmacSHA512,salt=qV-1YyqWIZsYc_SK1ajniQ,rounds=25000,ph=true,0
vtPyup3h5JThGFLff-1zw-42-BV7tG7Sm99ROT1OmZCZj1zcWLJj9Lc28mxkQI1-assfW2e-OPDhZbu9qcE2Q
! group
group admin
associate user admin
group read-only
! aaa
aaa accounting exec default start-stop local
! controller
controller
cluster-name dmf204
virtual-ip 10.93.194.147
access-control
!
access-list api
10 permit from 10.93.194.145/32
15 permit from 10.93.194.146/32
!
access-list gui
1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list ntp

```

```

1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list snmp
1 permit from 0.0.0.0/0
!
access-list ssh
1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list sync
1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list vce-api
1 permit from ::/0
2 permit from 0.0.0.0/0
! auto-vlan-mode
auto-vlan-mode push-per-filter
! service-node
service-node dm204-sn
mac e4:43:4b:48:58:ac
! switch
switch gt160
mac c4:ca:2b:47:97:bf
admin hashed-password $6$ppX0yA92$0hibVW63R0t1T3f7NRUFxPEWUb4b6414dTGEayrrXcw5or/ZDxm/Z
NvotQ7AQfVMO7OZ1I.yDLwrnlVXrZkv3.
!
interface Ethernet1
speed 10G
!
interface Ethernet5
speed 25G
switch hs160
mac c4:ca:2b:b7:44:83
admin hashed-password $6$McgvJd94$vrXDNkr2OSz3kiZSYPFcfuIbcaBuIcoC7ywlVeFFd7oAgLnleVcV6
NyEFZnykje4ILUjmJPWdWeu3LaF4sWzd/
!
interface Ethernet4
speed 10G
!
interface Ethernet47
role delivery interface-name veos2-delivery
!
interface Ethernet48
loopback-mode mac
speed 10G
role both-filter-and-delivery interface-name veos1-filter strip-no-vlan
switch smv160
mac 2c:dd:e9:4e:5e:f5
admin hashed-password $6$RyahYdXx$bUXeQCZ1bHNLcJBA9ZmoH/RmErwpDXvJE20UnEXKoLQffodjsyIln
ZlnG54X5Cq5qgb6uTGXs1TMYkqBWurLh1
!
interface Ethernet31/1
rate-limit 100
speed 10G
role delivery interface-name veos6-delivery ip-address 10.0.1.11 nexthop-ip 10.0.1.10
255.255.255.0
! crypto
crypto
!
http
cipher 1 ECDHE-ECDSA-AES128-GCM-SHA256
!
ssh
cipher 1 aes192-ctr
mac 1 hmac-shal
! policy
policy policy15
action forward

```

Example Two

Manually create a snapshot that contains a specific configuration.

```

Controller1(config)# begin transaction erase
id : GAOGEuLS26I67bqJ7J2NcpOtORfflUn_
Controller1(config)#
Controller1(config)# policy policy16
Controller1(config-policy)# action forward
Controller1(config-policy)#
Controller1(config-policy)# exit
Controller1(config)#
Controller1(config)# copy running-config snapshot://snap_policy16 transaction
Controller1(config)#
Controller1(config)#
Controller1(config)# delete transaction
Controller1(config)#
Controller1(config)# show snapshot snap_policy16
!
! Saved-Config snap_policy16
! Saved-Config version: 1.0
! Version: 8.4.0
! Appliance: bmf/dmf-8.4.x
! Build-Number 133
!
version 1.0
! local
local node
hostname Controller1
interface management
!
ipv4
ip 10.93.194.145/22 gateway 10.93.192.1
method manual
!
ipv6
method manual
! policy
policy policy16
action forward

```

Limitations

- The text-command-to-snapshot conversion process requires that the syntax of the text command to be replayed works (i.e., be compatible) with the currently supported syntax where it is getting applied.
- Only a global (i.e., cluster-wide) configuration can be managed with snapshots and transactions. View the local (non-global) configuration with the **show running-config local** command.
- An error is displayed if the **copy running-config snapshot://sample transaction** command is performed without starting a new transaction.

2.7 Managing DMF Sessions

Complete the following steps to view and configure settings for remote sessions established to the DANZ Monitoring Fabric (DMF) Controller, switches, and managed appliances. Security Sessions configuration tabs and features include:

- Configuration
 - Edit
- Sessions
 - Actions - Delete Security Sessions

1. Select **Security > Sessions** from the DMF main menu.

Figure 2-12: Security Sessions

The screenshot shows the 'Sessions' configuration page. At the top, there is a navigation bar with 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. Below the navigation bar, there is a 'Fabric Summary: Healthy' indicator. The main heading is 'Sessions' with a sub-heading 'Last updated: 18 minutes ago'. Below this, there is a description: 'Manage active sessions of DANZ Monitoring Fabric.' There is a 'Configuration' section with an 'Edit' link. Below the configuration section, there is a table with two rows: 'Expire After Access' with a value of '2 hours' and 'Concurrent Limit' with a value of '-'. To the right of this table is a yellow 'Note' box with the text: 'All users should make sure to log out when finished to avoid access being blocked. If the number of existing sessions is equal to the limit configured, then no new sessions are allowed.' Below the configuration section, there is a 'Sessions' section with a description: 'View all the sessions currently established to the controller.' There is a search bar with 'Contains...' and 'Add Filters' buttons. Below the search bar, there is a table with columns: 'Actions', 'Hash Id', 'User Name', 'Full Name', 'Groups', 'Last Address', 'Created', and 'Last Used'. The table contains one row with the following data: '4846cd55aa9f7b8ec0924cf6552fdaf...', 'admin', 'Default admin', '10.95.66.186', 'Mar 18, 2024 09:32:57:065am', and 'Mar 18, 2024 09:32:57:241am'. At the bottom right of the table, there is a pagination control showing '1-1 of 1 items' and '10 / page'.

The page displays the sessions currently established. To forcibly end a session, select **Delete Security Sessions** from the **Actions** control for the session.



Note: All users should log out when finished to avoid blocked access. No new sessions are allowed if the number of existing sessions equals the limit configured.

2. Under Configuration, click **Edit** to configure the default settings for remote connections to the Controller.

Figure 2-13: Security Sessions Configurations

The screenshot shows the 'Edit Session Configuration' dialog box. It has a title bar with 'Edit Session Configuration' and a close button. Below the title bar, there are two sections: 'Expire After Access' and 'Concurrent Limit'. The 'Expire After Access' section has a text input field with the value '120' and the label 'minutes'. The 'Concurrent Limit' section has an empty text input field. At the bottom of the dialog, there are three buttons: 'Reset', 'Cancel', and 'Submit'.

This dialog lets you set an expiration time for sessions and the maximum number of sessions allowed for each user.



Tip: Use the **Reset** button to clear the settings and return to the default values, if required.

3. Make any changes required and click **Submit**.

CLI Commands

To limit the number of concurrent sessions allowed for each user, use the following command:

```
aaa concurrent-limit <number>
```

For example, the following command limits the number of concurrent sessions to 5 for each user account.

```
controller-1(config)# aaa concurrent-limit 5
```

This limit causes the sixth session connection attempt by the same user account to fail. The excess oldest sessions are closed if more than five are already open.

This limit applies to sessions established through the GUI, CLI, or REST API, whether directed to the DANZ Monitoring Fabric switches, managed appliances, active or standby Controller, and the cluster virtual IP address.



Note: All users should log out when finished to avoid blocked access. No new sessions are allowed if the number of existing sessions equals the limit configured.

To set the expiration time for an AAA session, use the following command from config mode:

```
[no] aaa session-expiration <minutes>
```

Replace **minutes** with an integer specifying the number of minutes for the session expiration. For example, the following command sets the AAAS session expiration to **10** minutes:

```
controller-1(config)# aaa session-expiration 10
```

2.8 Managing and Viewing Logs

By default, the DANZ Monitoring Fabric (DMF) fabric switches send syslog messages to both the active and standby Controllers. Syslog messages are disk persistent and only removed based on time and rotation policy.

After configuring an external syslog server, the Controllers send the syslog messages to the external server and keep a local copy of the syslog messages on the Controller. When configuring multiple external syslog servers, DMF sends the syslog messages to every server. Physical switch logs can be sent directly to an external syslog server instead of sending the logs to the DMF Controller.

2.8.1 Sending Logs to a Remote Syslog Server

With the external Syslog server configured and the logging switch-remote option enabled, the fabric switches send Syslog messages only to the configured external Syslog servers but not to the Controllers. When the logging switch-remote option is enabled, the Controller does not have a local copy of the switch logs.

The Controllers do not share their syslog with the other Controller in the cluster. In other words, the active Controller does not send its syslog messages to the standby Controller, and the standby Controller does not share its syslog messages with the active Controller. Access the individual Controller logs from either the active or standby node.

2.8.2 Using the GUI to Configure Remote Logging

To manage logging, complete the following steps:

1. Select **Maintenance > Logging** from the main menu and click **Remote Logging**.

Figure 2-14: Maintenance Logging

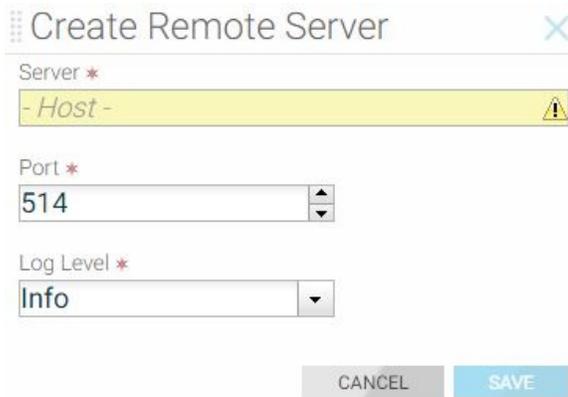


2. To enable remote logging, identify a remote syslog server by clicking the **Provision control (+)** in the Remote Servers table.



Note: Enabling remote logging causes syslog messages to be sent directly from the fabric switches to the specified server, bypassing the Controller. As a result, the Switch Light log files will not be available on the local Controller for analysis using the Analytics option.

Figure 2-15: Create a Remote Server



3. Enter the IP address of the remote server. If you are not using the default port (**514**), specify the port to use and click **Save**. The server appears on the Logging page.

2.8.3 Using the CLI to Configure Remote Logging

To configure the syslog server for a Controller node, enter the logging remote command using the following syntax:

```
..code-block:: none
```

```
logging remote <ip-address>
```

For example, the following command sets the syslog server address to **192.168.100.1**:

```
controller-1(config)# logging remote 192.168.100.1
```

This example exports the syslog entries from the Controller node to the syslog server at IP address **192.168.100.1**.

2.8.4 Viewing Log Files on the Controller

Use the **show logging** command to view the log files for the different components of the DANZ Monitoring Fabric (DMF) Controller. The options for this command are as follows:

- **audit**: Show audit file contents
- **controller**: Show log contents for the Controller floodlight process
- **remote**: Show remote logs
- **switch *switch***: Show logs for the specified switch
- **syslog**: Show syslog file contents
- **web-access**: Show content of the web server access log
- **web-error**: Show content of the web server error log

For example, the following command shows the logs for the DMF Controller.

```
controller-1> show logging controller
floodlight: WARN [MdnsResponder:Network Queue Processing Thread] ZTN4093: 1CC38M2._gm_idrac._tcp.local.
2018-03-26T04:03:17.900-07:00 invalid/unrecognized SRV name 1CC38M2._gm_idrac._tcp.local.
floodlight: WARN [MdnsResponder:Network Queue Processing Thread] ZTN4093: 1CC38M2._gm_idrac._tcp.local.
2018-03-26T04:03:17.900-07:00 invalid/unrecognized SRV name 1CC38M2._gm_idrac._tcp.local.
floodlight: WARN [MdnsResponder:Network Queue Processing Thread] ZTN4093: 1CC38M2._gm_idrac._tcp.local.
2018-03-26T04:03:17.900-07:00 invalid/unrecognized SRV name 1CC38M2._gm_idrac._tcp.local.
...
controller-1>
```

2.8.5 Administrative Activity Logs

The DANZ Monitoring Fabric (DMF) Controller logs user activities to the floodlight log file, logging the following events:

- CLI commands entered
- Login and logout events
- Queries to the REST server
- RPC message summaries between components in the Controller

CLI Commands

Use **grep** with the **show logging** command to see the local accounting logs, as in the following example:

```
controller-1# show logging controller | grep "cmd_args"
Sep 4 21:23:10 BT2 floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-3] AUDIT EVENT: bigcli.
command application_id=bigcli cmd_args=enable
Sep 4 21:23:10 BT2 floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-4] AUDIT EVENT: bigcli.
```

```

command application_id=bigcli cmd_args=configure
Sep 4 21:23:16 BT2_floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-3] AUDIT EVENT: bigcli.
command application_id=bigcli cmd_args=bigtap policy policy3
Sep 4 21:23:20 BT2_floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-6] AUDIT EVENT: bigcli.
command application_id=bigcli cmd_args=bigchain chain hohoh
Sep 4 21:23:22 BT2_floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-3] AUDIT EVENT: bigcli.
command application_id=bigcli cmd_args=ext
Sep 4 21:23:24 BT2_floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-3] AUDIT EVENT: bigcli.
command application_id=bigcli cmd_args=configure
Sep 4 21:23:30 BT2_floodlight: INFO [net.bigdb.auth.AuditServer:Dispatcher-5] AUDIT EVENT: bigcli.
command

```

To modify the accounting configuration, use the following commands. To start local accounting only, enter the following command:

```

controller-1(config)# aaa accounting exec default start-stop local

```

To start remote accounting only, enter the following command:

```

controller-1(config)# aaa accounting exec default start-stop {local [group tacacs+] | [group
radius]}

```

To view audit records, enter the following command if local accounting is enabled:

```

controller-1# show logging controller | grep AUDIT

```

If accounting is remote only, consult the administrator for the TACACS+ server for more information.

2.9 REST API Logging

Starting with *the DANZ Monitoring Fabric 8.4* release, the REST API body can be logged into the `audit.log` file. Before the 8.4 release, REST API calls from GUI or REST clients were logged in the `audit.log` file without the request's data (or) body. In this release, the data and body of the REST API call can be logged.

To enable the audit logging, use the configuration below:

```

aaa audit-logging log-request-leaf-values record-all-request-values

```

```

controller-1# config
controller-1(config)# aaa audit-logging log-request-leaf-values record-all-request-values
controller-1(config)#

```

To disable, use the `no` form of the command:

```

no aaa audit-logging log-request-leaf-values record-all-request-values

```

Example of `audit.log` entry showing SNMP configuration from GUI:

```

2022-11-16T12:18:05.106-08:00 floodlight: INFO LOCLAUD1001: AUDIT EVENT: DB_QUERY auth-
description="session/9d0a66315f7d9e0df8f2478fe7c0b3d77ceec25e865e4e135f0f9e28237570b70"
user="admin" remote-address="fd7a:629f:52a4:20d0:1ca8:28ed:6f59:cd47" session-id=
"9d0a66315f7d9e0df8f2478fe7c0b3d77ceec25e865e4e135f0f9e28237570b70" operation="REPLACE"
uri="https://[fdfd:5c41:712d:d080:5054:ff:fe57:5dba]/api/v1/data/controller/os/config/
global/snmp" http-method="PUT" request-leaf-values={"contact":"Arista","location":"HQ",

```

```
"trap-enabled":false,"user[name=\"cvml\"]/auth-passphrase\":\"AUTH-STRING\",\"user[name=\"cvml\"]/name\":\"cvml\",\"user[name=\"cvml\"]/priv-passphrase\":\"PRIV-STRING\",\"user[name=\"cvml\"]/priv-protocol\":\"aes\"} \"code=204"
```

To view the configuration:

```
controller-# show run aaa
! aaa
aaa accounting exec default start-stop local
aaa audit-logging log-request-leaf-values record-all-request-values
controller-1#
```

2.9.1 Restricting Size of REST API Body

Users can restrict the maximum size of the REST API body being passed to a rest call using the following command. Unless configured, the **max-body-size** is **9223372036854775807** bytes.

```
rest-api max-body-size <>
```

In the example below, the configuration restricts the body of the REST API call to 16K bytes. The REST API call is rejected if the body size exceeds 16K bytes. This feature can help remediate Denial of Service (DOS) attacks where an intruder tries to send a very large configuration repeatedly, attempting to keep the Controller busy parsing the information before applying. As needed, enable this configuration to the attack vector quickly before trying to find the culprits.

```
controller-1(config)# rest-api max-body-size
<max-body-size> Integer between 16384 to max integer size
controller-1(config)# rest-api max-body-size
controller-1(config)# rest-api max-body-size 16384
controller-1(config)#
```

To disable, use the **no** form of the command:

```
no rest-api max-body-size <>
```

2.10 Syslog Over TLS

This section describes how Syslog over TLS is implemented in DANZ Monitoring Fabric (DMF) and the configuration required to implement this feature.

2.10.1 Overview

Transport Layer Security (TLS) is used in DANZ Monitoring Fabric (DMF) to secure syslog messages, which may originate or traverse a non-secure network in transit to the syslog server. Using TLS helps mitigate the following primary threats:

- **Impersonation:** An unauthorized sender may send messages to an authorized receiver, or an unauthorized receiver may try to deceive a sender.
- **Modification:** An attacker may modify a syslog message in transit to the receiver.
- **Disclosure:** An unauthorized entity could examine the contents of a syslog message. TLS, when used as a secure transport, reduces these threats by providing the following functionality.

- Authentication counters impersonation.
- Integrity-checking counters modifications to a message on a hop-by-hop basis.
- Confidentiality works to prevent disclosure of message contents.

Starting from DMF Release 8.4, syslog over TLS is supported only for communication from Controllers to remote syslog servers. Switches and managed appliances such as Recorder Nodes and Service Nodes support plain unencrypted UDP-based syslog messages. To enable syslog over TLS on the Controller, refer to the next section.

2.10.2 Configuration

To enable TLS for syslog messaging, complete the following steps:



Note: The steps below are needed for mutual authentication between client and server. For non-mutual cases, users can skip importing the Controller's certificate and a private key and configuring them.

1. Generate the certificates for the remote syslog servers and the DANZ Monitoring Fabric (DMF) Controller using the same trusted public or private CA.
2. Copy the trusted CA root certificate, the Controller's certificate, and the Controller's private key to the DMF Controller.

```
controller-1# copy scp://user@x.x.x.x:/.../cacert.pem cert://
controller-1# copy scp://user@x.x.x.x:/.../dmf-controller-cert.pem cert://
controller-1# copy scp://user@x.x.x.x:/.../dmf-controller-privkey.pem private-key://dmf-privkey
```

3. Copy the CA certificate and the syslog server certificates and keys to the respective syslog servers. TLS support should be enabled on the servers by following the prescribed steps for the syslog applications that they are running.
4. On the active DMF Controller, identify the CA certificate to use with remote syslog servers.

```
controller-1(config)# logging secure ca <capath>
```



Note: In the example in **Step 2**, it is `cacert.pem`.

5. Identify a certificate for mutual authentication.

```
controller-1(config)# logging secure cert <certpath>
```



Note: In the example in **Step 2**, it is `dmf-controller-cert.pem`.

6. On the active DMF Controller, identify the key for mutual authentication.

```
controller-1(config)# logging secure private-key <keypath>
```



Note: In the example in **Step 2**, it is `dmf-privkey`, the name given to the key during the copy procedure.

7. On the active DMF Controller, enable secure logging.

```
controller-1(config)# logging secure tls
```

8. To view the syslog messages, enter the following command.

```
controller-1# show logging
```



Note: Syslog applications typically send logs to `/var/log/messages`, `/var/log/dmesg`, and `/var/log/journal` by default.

2.11 Creating a Support Bundle

A collection of running configuration and log files is critical to understanding the fabric behavior when the fabric is in a faulty state.

The DANZ Monitoring Fabric (DMF) CLI provides the commands to automate the collecting, archiving, and uploading of critical data. These commands cover all devices of the DMF fabric, such as Controllers, switches, DMF Service Node, and DMF Recorder Node.

The following are the commands to configure the Support Bundle data upload:

```
controller-1> enable
controller-1# configure
controller-1(config)# service
controller-1(config-service)# support auto-upload
Enabled diagnostic data bundle upload
Use "diagnose upload support" to verify upload server connectivity
```

To check if the auto-upload is enabled:

```
controller-1# show run service
! service
service
support auto-upload
controller-1#
```

The following is the command to launch the Support Bundle collection. Once the support bundle is collected, it will automatically be uploaded, as seen in the output. Provide the bundle ID to support personnel.

```
controller-1# support
Generating diagnostic data bundle for technical support. This may take several minutes...
Support Bundle ID: SGPVW-BZ3MM
Switchlight collection completed after 14.2s. Collected 1 switch (8.56 MB)
Local cli collection completed after 38.2s. Collected 75 commands (0.32 MB)
Local rest collection completed after 0.1s. Collected 3 endpoints (0.43 MB)
Local bash collection completed after 10.0s. Collected 127 commands (4.74 MB)
Local file collection completed after 15.5s. Collected 39 paths (1753.31 MB)
Cluster collection completed after 138.2s. Collected 1 node (1764.50 MB)
Collection completed. Signing and compressing bundle...
Support bundle created successfully
00:04:03: Completed
Generated Support Bundle Information:
Name : anet-support--DMF-Controller--2020-11-24--18-31-39Z--SGPVW-BZ3MM.tar.gz
Size : 893MB
File System Path : /var/lib/floodlight/support/anet-support--DMF-Controller--2020-11-24--18-31-39Z--SGPVW-BZ3MM.tar.gz
Url : https://10.2.1.103:8443/api/v1/support/anet-support--DMF-Controller--2020-11-24--
```

```

18-31-39Z--SGPVW-BZ3MM.tar.gz
Bundle id : SGPVW-BZ3MM
Auto-uploading support anet-support--DMF-Controller--2020-11-24--18-31-39Z--SGPVW-BZ3MM.tar.gz
Transfer complete, finalizing upload
Please provide the bundle ID SGPVW-BZ3MM to your support representative.
00:00:48: Completed
controller-1#

```

The **show support** command shows the status of the automatic upload.

```

controller-1# show support
# Bundle          Bundle id      Size  Last modified      Upload status
-----|-----|-----|-----|-----|
1 anet-support--DMF-Controller--2020-11-24--18-31-39Z--SGPVW-BZ3MM.tar.gz  SGPVW-BZ3MM  893MB  2020-11-24 18:35:46.400000 UTC  upload-completed

```

Use the **diagnose upload support** command to verify the reachability and health of the server used to receive the support bundle. Below is an example output of checks performed when running the command. When a support bundle upload fails, use the command to identify the possible causes.

```

controller-1# diagnose upload support
Upload server version: diagus-master-43
Upload diagnostics completed successfully
00:00:02: Completed
Check : Resolving upload server hostname
Outcome : ok
Check : Communicating with upload server diagnostics endpoint
Outcome : ok
Check : Upload server healthcheck status
Outcome : ok
Check : Upload server trusts authority key
Outcome : ok
Check : Signature verification test
Outcome : ok
Check : Resolving objectstore-accelerate hostname
Outcome : ok
Check : Resolving objectstore-direct hostname
Outcome : ok
Check : Communicating with objectstore-accelerate
Outcome : ok
Check : Communicating with objectstore-direct
Outcome : ok
controller-1#

```

After resolving the issues, retry the upload using the **upload support *support bundle file name*** command. Use the same command if auto-upload was not configured before taking the support bundle.

```

controller-1# upload support anet-support--DMF-Controller--2020-11-24--18-31-39Z--SGPVW-BZ3MM.tar.gz

```

The software provides the following options to reduce the amount of log entries collected in the support bundle:

```

controller-1# support
logs-since-days      Generate diagnostic data bundle
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-jfr-dump        Skip creating JFR information during 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.
switch               Copy a switch core dump
<Command-end>       <cr>: Return to execute command
;                   command separator
|                   pipe to command

```

```
> redirect
controller-1#
```

2.12 NIST 800-63b Password Compliance

This feature activates password compliance for local accounts on DANZ Monitoring Fabric (DMF) devices (Controller, switches, DMF Service Node, Arista Analytics Node, and DMF Recorder Node). The NIST 800-63b feature enforces that any newly chosen password fulfills the following requirements:

- The password needs to be at least 8 characters long.
- The password is not a known compromised password.



Note: Enabling NIST 800-63b compliance only enforces that password changes in the future will enforce the rules; it does not affect existing passwords. Updating all passwords after enabling NIST-800-63b compliance is strongly recommended.

2.12.1 Configuration

The NIST-800-63b compliance mode needs to be set separately on the Controller and the Arista Analytics Node to enforce password compliance for the entire DANZ Monitoring Fabric (DMF) cluster.

Enable NIST 800-63b password compliance

```
Controller(config)# aaa authentication password-compliance nist-800-63b
Warning: A password compliance check has been enabled. This enforces compliance
Warning: rules for all newly chosen passwords, but it doesn't retroactively
Warning: apply to existing passwords. Please choose new passwords for
Warning: all local users, managed appliances, switches, and the recovery user.
```

Disable NIST 800-63b password compliance (non-FIPS)

```
Controller(config)# no aaa authentication password-compliance
```

FIPS versions always enforce NIST 800-63b password compliance by default unless explicitly configured not to do so.

Disable NIST 800-63b password compliance (FIPS)

```
Controller(config)# aaa authentication password-compliance no-check
```



Note: The NIST-800-63b compliance mode must be set separately on the Controller and the Arista Analytics Node to enforce password compliance for the entire DMF cluster.

View the NIST 800-63b password compliance configuration

```
Controller(config)# show running-config aaa
! aaa
aaa authentication password-compliance nist-800-63b
```



Note: Upon activation of NIST 800-63b password compliance, updating local account passwords for DMF devices is recommended. Refer to the commands below to update the passwords.

Controller admin password update

```
Controller# configure
Controller(config)# user admin
Controller(config-user)# password <nist-compliant-password>
```

Service Node admin password update

```
Controller(config)# service-node <service-node-name>
Controller(config-service-node)# admin password <nist-compliant-password>
```

Recorder Node admin password update

```
Controller(config)# recorder-node <recorder-name>
Controller(config-packet-recorder)# admin password <nist-compliant-password>
```

Switch admin password update

```
Controller(config)# switch <switch-name>
Controller(config-switch)# admin password <nist-compliant-password>
```



Note: After upgrading to **DMF-7.2.1** and above, existing user credentials continue to work. In the case of a FIPS image, password compliance is enabled by default, and this does not affect existing user accounts. Changing the recovery and local account passwords is recommended after upgrading from the previous FIPS image. In the first time installation of a FIPS image, the first-boot process enforces the NIST compliance standards for the recovery and admin accounts.

2.13 Custom Password Compliance

This feature activates password compliance for local accounts on DANZ Monitoring Fabric (DMF) devices (Controller, switches, DMF Service Node, Arista Analytics Node, and DMF Recorder Node).

DMF supports custom password requirements for local users.

To enable the custom password compliance method, use the following command:

```
aaa authentication password-compliance custom-check
```

```
controller-1(config)# aaa authentication password-compliance custom-check
Warning: A password compliance check has been enabled. This enforces compliance
Warning: rules for all newly chosen passwords, but it doesn't retroactively
Warning: apply to existing passwords. Please choose new passwords for
Warning: all local users, managed appliances, switches, and the recovery user.
controller-1(config)#
```

To disable, use the **no** form of the command:

```
no aaa authentication password-compliance custom
```



Note: Enabling the custom password compliance does not apply to already configured user passwords. Arista Networks recommends resetting or reconfiguring the user password to adhere to the password requirements. Without configuring the custom password compliance, configuring the requirements is not effective. To use the feature configure the *custom* compliance method and set the password requirements.

Once the custom password compliance is enabled, configure the requirements for a password using the following commands:

```
controller-1(config)# aaa authentication password-requirement
max-repeated-characters the maximum number of repeated characters allowed
max-sequential-characters the maximum number of sequential characters allowed
minimum-length the minimum required length for passwords
minimum-lowercase-letter the minimum number of lowercase characters required
minimum-numbers the minimum number of numerical characters required
minimum-special-characters the minimum number of special characters required
minimum-uppercase-letter the minimum number of uppercase characters required
reject-known-exposed-passwords check the password against known exposed passwords
controller-1(config)# aaa authentication password-requirement
```

For example, to set a password that requires 10 minimum characters, 1 minimum number, and 2 maximum repeated characters, do the following:

```
controller-1(config)# aaa authentication password-requirement minimum-length 10
controller-1(config)# aaa authentication password-requirement minimum-numbers 1
controller-1(config)# aaa authentication password-requirement max-repeated-characters 2
```

Configuring or resetting a password that does not meet the requirements results in the following error message:

```
controller-1# conf
controller-1(config)# user customPW
controller-1(config-user)# password admin
Error: the password needs to be at least 10 characters long
controller-1(config-user)#
```

To view the configured password compliance and requirements, use the **show run aaa authentication** command:

```
controller-1# show run aaa authentication
! aaa
aaa authentication password-compliance custom-check
aaa authentication password-requirement max-repeated-characters 2
aaa authentication password-requirement minimum-length 10
aaa authentication password-requirement minimum-numbers 1
controller-1#
```

To remove the requirements configured, use **no** form of the commands:

```
controller-1(config)# no aaa authentication password-requirement minimum-length 10
controller-1(config)# no aaa authentication password-requirement minimum-numbers 1
controller-1(config)# no aaa authentication password-requirement max-repeated-characters 2
```

2.14 Switch/Managed Appliance Management Interfaces Not Mirroring Controller Management Interface ACLs

Use this feature to configure Access Control Lists (ACLs) on a managed device that do not directly reflect the ACLs configured on the Controller.

Specifically, a user can override the user-configured ACLs on the Controller (generally inherited by the managed devices) so that ACLs allowing *specific types* of traffic from the Controller only are pushed to managed devices.

The user performs this action on a per-managed-device basis or globally for all managed devices on the CLI. The Controller and analytics node are excluded from receiving this configuration type (when performed globally).

The feature introduces a new CLI mode to address this configuration type on the Controller. That is, the configuration used for pushing to all managed devices exclusively (excluding the Controller) unless overrides exist.

A managed device is a device whose life cycle ZTN manages.

The total set of managed devices is as follows:

- Managed Appliances
 - Service Node
 - Recorder Node
- Switches
 - SWL
 - EOS



Note: The analytics node is not in this set since the ZTN mechanisms on the Controller do not manage its life cycle.

2.14.1 Configuration using the CLI

Configure the following on the Controller to enforce Intra-Fabric Only Access (IFOA) for the API service (i.e., port 8443) on all managed devices.

```
C1> en
C1# conf
C1(config)# managed-devices
C1(config-managed-devices)# access-control
C1(config-managed-devices-access)# service api
C1(config-managed-devices-access-service)# intra-fabric-only-access
Reminder: IP address/method of the management interface cannot be changed,
when a service has intra-fabric only access enforced.
```



Note: A warning is displayed outlining that the management interface's IP address cannot be changed once any of the services have IFOA enforced for any managed devices. To change the management interface's IP address, disable IFOA for all services.

Several services can have IFOA enforced on them. The list of services and their corresponding ports are shown in the table below. An **X** mark means enforcing IFOA on that port on that managed device type is not supported.

Conversely, a **✓** mark means enforcing IFOA on that port on that managed device type is supported. There may be some information beside the **✓** mark indicating what runs on that port on the managed device.

Protocol	Service Node	Recorder Node	SWL	EOS
SSH (22, TCP)	✓	✓	✓	✓
WEB (80, TCP)	X	X	✓ (SLRest, plain-text)	✓ (cAPI)
HTTPS (443, TCP)	X	✓ (nginx reverse proxies to 1234, for the stenographer)	X (SLRest, encrypted)	✓
API (8443, TCP)	✓ (BigDB)	✓ (BigDB)	✓ (BigDB)	X

To override this global/default config (i.e., config applied to all managed devices) for a specific managed device, use the following config and push it to the Controller.

```
C1(config)# switch switch-name
C1(config-switch)# access-control override-global
C1(config-switch)# access-control
C1(config-switch-access)# service api
C1(config-switch-access-service)# no intra-fabric-only-access
```

As illustrated below, push a similar configuration for the managed appliances, i.e., the Recorder and Service nodes.

Recorder Node

```
C1(config)# recorder-node device rn1
C1(config-recorder-node)#
C1(config-recorder-node)# access-control override-global
C1(config-recorder-node)# access-control
C1(config-recorder-node-access)# service api
C1(config-recorder-node-access-service)# no intra-fabric-only-access
```

Service Node

```
C1(config)# service-node sn1
C1(config-service-node)#
C1(config-service-node)# access-control override-global
C1(config-service-node)# access-control
C1(config-service-node-access)# service api
C1(config-service-node-access-service)# no intra-fabric-only-access
```

It is also possible to push a configuration that does not override the entire configuration under managed devices but instead merges with it on a per-service basis, for example:

```
C1(config)# switch core1
C1(config-switch)# access-control merge-global
C1(config-switch)# access-control
C1(config-switch-access)# service api
C1(config-switch-access-service)# no intra-fabric-only-access
```

This action will merge the global/default config specified under the `config-managed-devices` CLI submode with the config set for this specific managed device (in this case, the device is a *switch*, and its name on the Controller is *core1*).

2.14.2 CLI Show Commands

There are several helpful `show` commands.

When merging the global/default access-control configuration with the device-specific configuration, understanding the effective configuration (the configuration used in generating the appropriate ACLs) may not be obvious. To see the **effective** configuration for a specific device, perform the following command:

```
C1(config)# show effective-config switch core1
! switch
switch core1
access-control
!
service api
intra-fabric-only-access
```

While displaying the managed device's effective configuration, check the **running-config** generated by ZTN (the configuration sent to the device), confirming the configuration pushed to the managed device.

```
C1(config)# show service-node rn1 running-config
.
.
.
interface mal acl subnet 10.243.254.20/32 proto tcp port 8443 accept
interface mal acl subnet fe80::5054:ff:fe80:b844/128 proto tcp port 8443 accept
interface mal acl subnet 0.0.0.0/0 proto tcp port 8443 drop
interface mal acl subnet ::/0 proto tcp port 8443 drop
interface mal acl subnet ::/0 proto udp port 161 accept
interface mal acl subnet 0.0.0.0/0 proto udp port 161 accept
interface mal acl subnet 0.0.0.0/0 proto udp port 161 drop
interface mal acl subnet ::/0 proto udp port 161 drop
interface mal acl subnet 10.243.254.20/32 proto tcp port 22 accept
interface mal acl subnet fe80::5054:ff:fe80:b844/128 proto tcp port 22 accept
interface mal acl subnet ::/0 proto tcp port 22 accept
interface mal acl subnet 0.0.0.0/0 proto tcp port 22 accept
interface mal acl subnet 0.0.0.0/0 proto tcp port 22 drop
interface mal acl subnet ::/0 proto tcp port 22 drop
interface mal acl default accept
.
.
.
```



Note: For API (port 8443), the system only pushes ACLs that permit IPv4/LLv6 traffic from the Controller and drops everything else. Other ACLs (SSH/SNMP) are not generated from the managed devices' access control configuration on the CLI. They are generated from access-rules configuration for the Controller that gets used or inherited by the managed devices.

Note: The same show command exists for Recorder Nodes and Service Nodes, i.e.:



Recorder Nodes

- `show effective-config recorder-node rn1`
- `show recorder-node rn1 running-config`

Service Nodes

- `show effective-config service-node rn1`
- `show service-node rn1 running-config`

2.14.3 Limitations

The main limitation of this feature is the inability to change the management interface's IP address (on the CLI) once enforcing IFOA for any of the services on any managed devices so that the Controller does not inadvertently get locked out from the managed devices.



Note: Changing the management IP address via the REST API without getting blocked is possible. However, Arista Networks advises against doing so when enforcing IFOA for any service on any managed device.

2.15 Recovery Procedure

This section describes the recovery procedure when one or both Controllers go down.

2.15.1 Recovery from a Single Controller Failure**Procedure**

1. Log in to the remaining (online) Controller and enter the `system remove-node failed controller` command.
2. Start a new Controller and complete the first boot process.
3. When prompted, join the existing/remaining Controller (the configuration syncs to the new Controller as soon as it joins the cluster).



Note: This step restores only the running configuration. It does not restore user files from the failed Controller.

2.15.2 Recovery from a Dual Controller Failure

The following procedure assumes that the Controllers have the same IP addresses as the failed Controllers.

Procedure

1. Archive the current running configuration as a database snapshot and save it to the SCP server. Enter the following command from enable mode.

```
controller-1# copy running-config snapshot://current.snp
controller-1# copy snapshot://current.snp scp://anetadmin@10.240.88.130/home/anetadmin/
controller.snp
anetadmin@10.240.88.130's password:
controller.snp 5.05KB -00:00
controller-1#
```

2. Install the first Controller and complete the first boot process.

3. Restore the archived version of the running configuration by entering the following command from enable mode.

```
new-controller-1# copy scp://anetadmin@10.240.88.130/home/anetadmin/controller.snp
snapshot://controller.snp
anetadmin@10.240.88.130's password:
controller.snp
5.05KB - 00:00
new-controller-1# copy snapshot://controller. running-config
new-controller-1#
```



Note: This step only restores the running configuration. It does not restore all user files from the previously running Controller.

4. Install the second Controller and complete the first boot process.
5. When prompted, join the first Controller to form a cluster.
6. The configuration between Controllers synchronizes as soon as the standby Controller joins the cluster.



Note: This step only restores the running configuration. It does not restore user files from the failed Controller.

7. If the new Controllers are running a different version than the previous set of Controllers, reboot the switches to obtain a compatible switch image from the Controller. Enter the following command from enable mode:

```
new-controller-1# system reboot switch switch
```

DANZ Monitoring Fabric Deployment Topologies

This chapter describes the different topologies for out-of-band deployment of DANZ Monitoring Fabric.

3.1 DANZ Monitoring Fabric Topologies

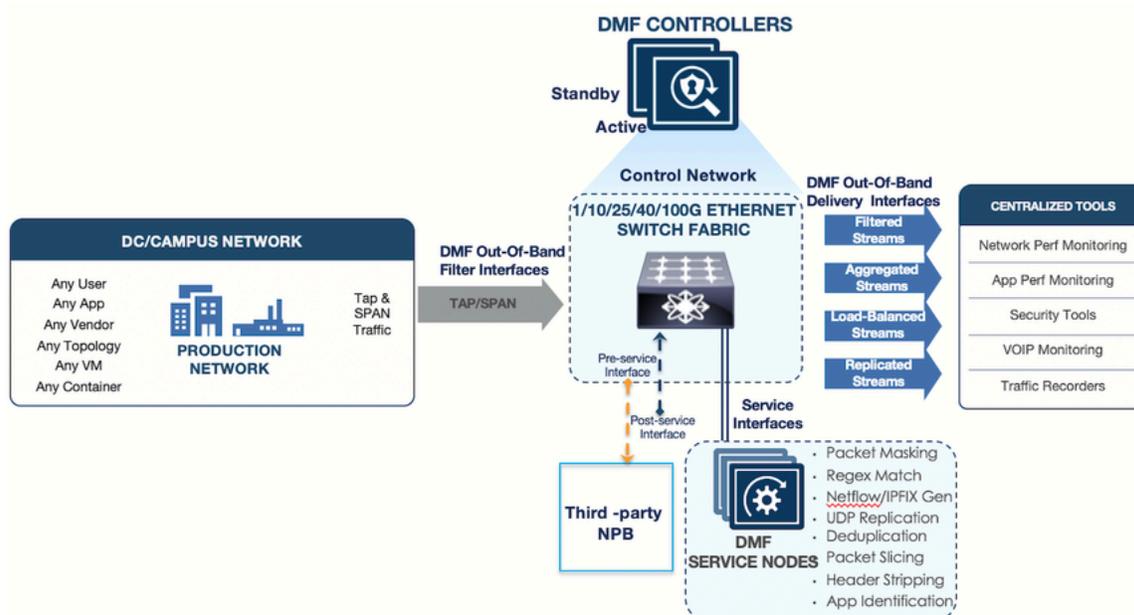
DANZ Monitoring Fabric (DMF) supports the implementation of out-of-band monitoring using a single switch or with many switches for a scalable, high-availability topology. This topology supports thousands of filter and delivery ports.

This section provides a summary of the recommended deployments.

3.1.1 Single Switch Topology

The single switch topology is the most basic design for small-scale environments, where a single switch provides enough filter interfaces, delivery interfaces, and optional service interfaces for connecting to NPBs for various packet manipulation operations, such as time stamping and packet slicing.

Figure 3-1: DANZ Monitoring Fabric (DMF) Single Switch Topology



This design option is most useful in the following scenarios:

- The environment does not need to scale beyond the interfaces provided by a single DMF Out-of-Band (OOB) switch.

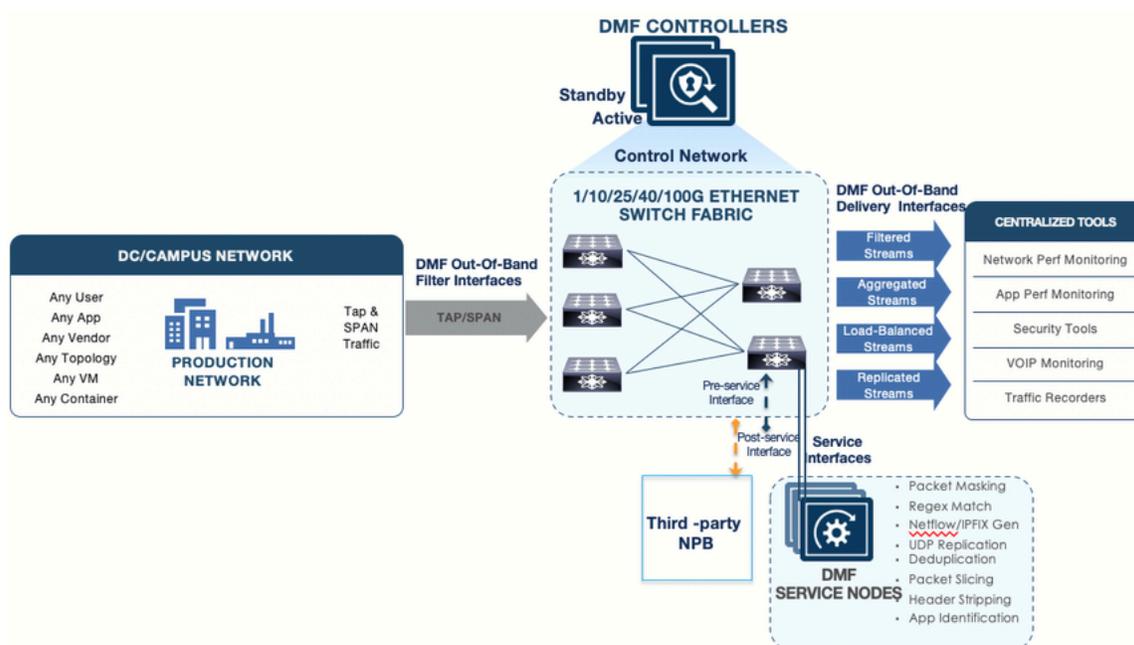
- A single switch topology improves cable management when filter and delivery ports are physically dispersed throughout the data center.

3.1.2 Two-Tier Topology

A two-tier design, shown in the figure below, is most useful in the following scenarios:

- Medium-to-high port scalability requirements.
- Production network TAPs are dispersed across the data center and require aggregation.
- Tools are physically consolidated, and the traffic needs to be aggregated.

Figure 3-2: DANZ Monitoring Fabric (DMF) Two-Tier Topology



When deploying this topology, ensure the following requirements:

- Only use core links between monitoring switches in different tiers. Depending on port availability and the bandwidth requirements, the core links can be 10, 25, 40, or 100 GbE.
- Avoid connecting links between filter switches to help ensure efficient path computation.
- Connect at least two links between tiers for link redundancy. The total number of physical links between the tiers will vary according to the amount of oversubscription in the design.
- Service nodes should be connected to the delivery switch to send the aggregated traffic to the service nodes (NPBs).

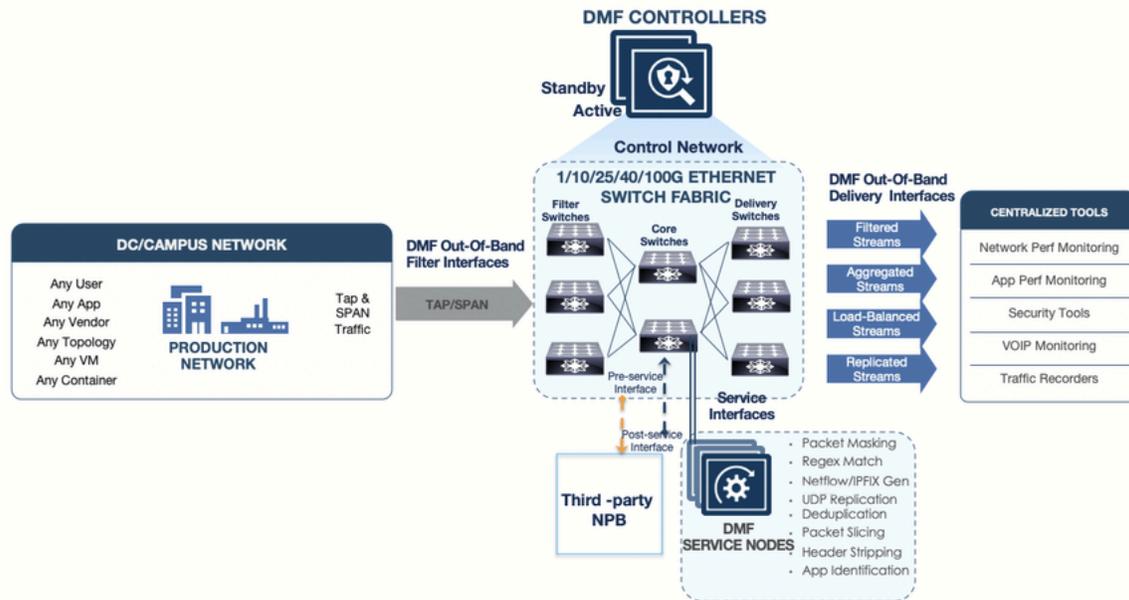
3.1.3 Three-Tier Any-Tap-to-Any-Tool Topology

A three-tier design, as shown in the figure below, is most useful in the following scenarios:

- Large-scale deployments where hundreds of TAP ports are installed across the data center.

- Traffic from TAPs must be aggregated, and the aggregated traffic forwarded to analysis tools in different locations. This design provides any-TAP-to-any-tool connectivity.

Figure 3-3: DANZ Monitoring Fabric (DMF) Three-Tier (Any-TAP-to-Any-Tool) Topology



With this topology, ensure the following requirements are satisfied:

- 40 or 100 GbE links are recommended to the core switches.
- Connect each filter switch to at least two core tier switches for redundancy. Depending on the amount of oversubscription in the design and on port availability, more ports can be connected between tiers.
- Service nodes should be connected to a core (aggregation) switch so that aggregated traffic can be delivered to the service nodes (NPBs).
- Avoid connecting links between filter switches to help ensure efficient path computation.

4.2 Zero Touch Fabric Provisioning Modes

Complete the fabric switch installation using one of the following two modes:

1. **Layer 2 Zero Touch Fabric (L2ZTF, Auto-discovery switch provisioning mode):** In this mode (which was the default up to DMF release 8.4), the 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). If the fabric switches need IPv4 addresses to communicate with SNMP or other external services, configure IPAM, which provides the Controller with a range of IPv4 addresses to allocate to the fabric switches.
2. **Layer 3 Zero Touch Fabric (L3ZTF, Pre-configured switch provisioning mode):** In this mode, which is the default starting from DMF release 8.5, when fabric switches are in a different Layer 2 network from the Controller, log in to each switch individually to configure network information and download the ZTF installer. Subsequently, the switch automatically downloads Switch Light OS from the Controller. This mode requires communication between the Controller and the fabric switches to occur using IPv4 addresses, and no IPAM configuration is required.

Note: For both switch installation modes, you must enter the following commands for every switch:



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

The following table summarizes the requirements for installation using each mode:

Requirements	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-only
Refer to this section	Using L2 ZTF (Auto-Discovery) Provisioning Mode	Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode

Install all the fabric switches in a single fabric using the same mode. If there are any fabric switches in a different IP subnet than the Controller, DANZ Monitoring Fabric (DMF) **requires** using Layer 3 mode to install 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.

4.3 Using L2 ZTF (Auto-Discovery) Provisioning Mode

Layer 2 Zero Touch Fabric (L2 ZTF) is used to provision and install DANZ Monitoring Fabric (DMF) switches that are in the same Layer 2 management network as the DMF Controllers. ZTF uses the Open Network Install Environment (ONIE) boot loader to automate switch installation and configuration. Supported fabric switches are shipped with an ONIE network-enabled boot image. Refer to the **DANZ Monitoring Fabric 8.5 Hardware Compatibility List** for a list of supported monitoring fabric switches. During switch boot up, each switch gets the Switch Light OS software from the DMF Controller.



Note: DMF supports deploying switches from different vendors in the same fabric. However, using cables from different vendors to connect switches is not supported. Optics are required to interconnect switches from different vendors. See the Hardware Compatibility List for details on supported optics for each switch platform.



Note: If a switch is in a different subnet than the Controller, refer to the [Changing to Layer 3 \(Pre-Configured\) Switch Provisioning Mode](#) section for details about how to install it.

4.3.1 Requirements

Consider the following and perform each item when using ZTF to install fabric switches:

- The **DANZ Monitoring Fabric 8.5 Hardware Compatibility List** lists the supported fabric switches.
- Connect the management Ethernet interface of each physical switch to the management network and power it up.
- Connect the DANZ Monitoring Fabric (DMF) Controller appliance management interface to the same Layer 2 management network as the management Ethernet interface of every physical switch.
- When upgrading switches from a previous deployment, ensure the Switch Light OS image is compatible with your Controller version.
- Designate a range of IPv4 addresses to be assigned using IPAM when switches must communicate with SNMP, NTP, syslog, or other IPv4 services.



Note: DMF implements ZTF using the IPv6 link-local address, which is auto-generated.



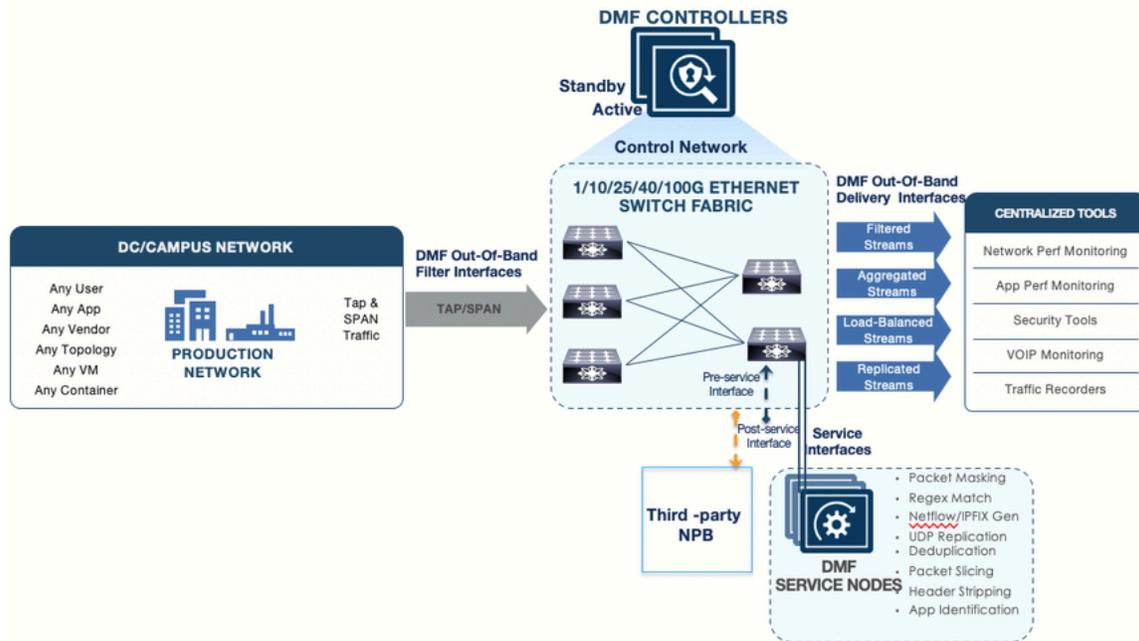
Warning: If a switch has an installed operating system, follow the instructions provided in the [Removing the existing OS from a Switch](#) topic before connecting the switch to a DMF Controller.

4.3.2 Switch Installation Procedure

This section describes using ZTF to perform a fresh installation of one or more supported DANZ Monitoring Fabric (DMF) switches using the Switch Light OS included with the Controller software.

The following figure illustrates a two-tier DMF out-of-band deployment with two switches connected to SPAN or TAP ports in a production network and a third switch connected to monitoring and analysis tools.

Figure 4-1: DMF Two Tier Out-of-Band Deployment



As shown in the illustration, services can be provided by the DMF Service Node Appliance or a third-party Network Packet Broker (NPB).

To use ZTF to bring up a DMF switch in this deployment, complete the following steps:

1. Take note of the MAC address of the switch.

Note: The MAC address is usually printed on the top surface of the switch.

If it is not possible to view the switch label (it will not be visible for already racked switches), obtain the MAC address from the local console of each switch by entering one of the following commands.



- In uboot mode, enter `printenv ethaddr`.
- In ONIE mode, enter `ifconfig` to display the ONIE prompt, and type the following at the command prompt:

```
``=> run onie_bootcmd``
```

2. Register each switch on the Controller by entering the following commands.



Note: A given switch can be connected to only one Controller cluster.

```
controller-1(config)# switch Filter-1
controller-1(config-switch)# mac 70:72:cf:bc:c4:c4
controller-1(config-switch)#

controller-1(config)# switch Filter-2
controller-1(config-switch)# mac 70:72:cf:ea:1b:bb
controller-1(config-switch)#
```

```
controller-1(config)# switch Delivery-1
controller-1(config-switch)# mac 70:72:cf:bd:54:24
controller-1(config-switch)#
```

Note: Verify that the switch provisioning mode is configured for auto-discovery.



Auto-discovery was the default mode of operation up to DMF release 8.4, pre-configured mode is the default starting from DMF release 8.5.0. After entering the show **running-config** command, **deployment-mode auto-discovery** should appear.

3. Turn on or restart the fabric switch.
4. Initiate the ONIE request on the switch.
 - a. On the GNU GRUB menu, select **ONIE**.

To get to the ONIE mode, during the reboot countdown, press any key when you see the prompt: “**Hit any key to stop autoboot: 0**”. The following command launches the ONIE install mode:

```
=> run onie_bootcmd
GNU GRUB version 2.02~beta2+e4a1fe391
+-----+
| Switch Light OS |
| *ONIE |
| |
| |
| |
| |
+-----+
Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, #e' to edit the commands
before booting or #c' for a command-line.
```

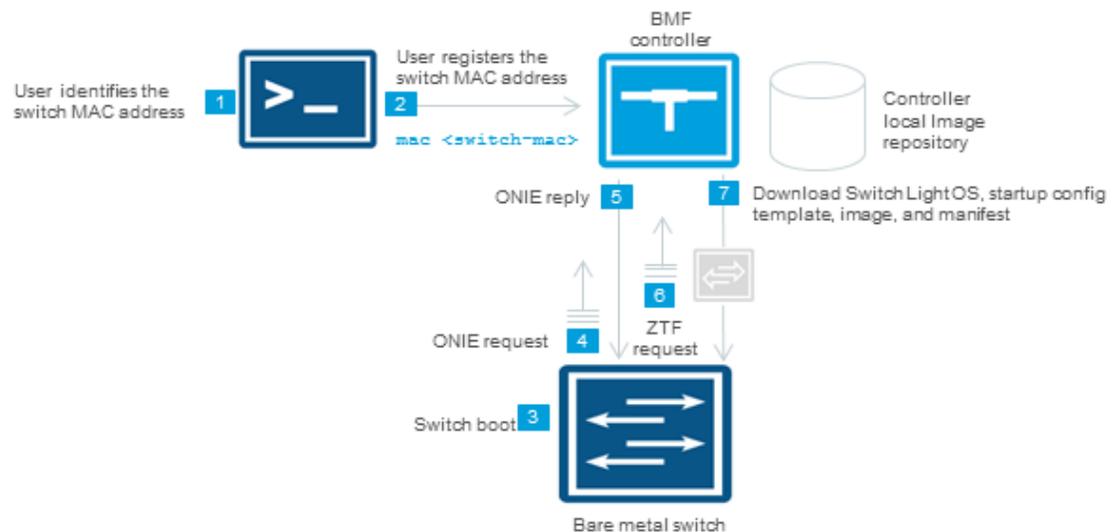
- b. Select **ONIE: Install OS**.

This selection puts the switch into the installer mode, and the rest of the process runs automatically, beginning with the following message.

```
### START - switch console output for installer :
ONIE: OS Install Mode ...
...
```

The switch discovers the Controller on the management plane as part of the installation process that occurs when it starts. The following figure illustrates the process after registering a connected fabric switch on the Controller.

Figure 4-2: Switch Boot and ZTF Configuration



When powering on a compatible switch, it boots using the U-boot service, which is pre-installed, and this starts the ONIE loader.

The remainder of the steps in the installation and configuration process happen automatically after powering on the registered and physically connected switch. These automated steps are shown in the figure above and are summarized below.



Note: Steps 5 through 7 below are performed automatically during the ZTF process and typically require no user intervention. If a switch is in a different subnet than the Controller, refer to the **Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode** section for details about how to install it.

5. The ONIE loader generates an IPv6 neighbor discovery message on the local network segment.
6. Because the MAC address is already registered in **Step 2** of the procedure described above, the Controller responds to the ONIE request from that switch and instructs the switch to download the Switch Light OS loader and begin the installation.
7. After installing the Switch Light OS loader and rebooting, the loader broadcasts a ZTF request.
8. (Not illustrated) The ZTF server (the active DMF Controller) sends the Switch Light OS image, manifest, and startup-config, or a URL to the location where the switch can download it.

The switch downloads its startup-config from the Controller, which includes the following configuration information:

- Hostname
- Switch MAC address
- Controller IP addresses
- NTP, logging, and SNMP configuration mirrored from the Controller configuration

9. When all the switches are powered up, enter the `show switch` command to verify the switch connectivity to the Controller. For example:

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

4.3.3 Arista Switch Installation Procedure for 7050X Series and 7260X Series

The initial installation of Switch Light OS on the Arista switch platforms must be performed manually in DANZ Monitoring Fabric (DMF). The initial boot process cannot be performed automatically in the same L2 domain like existing DMF-supported switches that are running ONIE.

The initial installation of Switch Light OS on the Arista platforms is accomplished by dropping it into the Aboot shell interface at boot time and telling it to boot the Switch Light switch image. This operation will install Switch Light on the system. This is a one-time extra step needed during the first installation of Switch Light OS in DMF. The boxes will subsequently boot as expected under Switch Light.

This procedure is also required for any Arista switches currently running EOS. Perform the following steps for the Arista switch to boot from the DMF Controller:

1. Attach a console cable to the Arista switch and then turn on or reboot the switch.
2. Interrupt the boot process with **Control-C** to drop into the Aboot shell.

```
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0)
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
```

3. Configure an IP address for the switch's ma1 management interface. Either configure ma1 statically or use DHCP.



Note: To use DHCP, type `udhcpc -i ma1`.

```
Aboot# udhcpc -i ma1
udhcpc (v1.18.1) started
Sending discover...
Sending discover...
Sending select for 10.6.3.237...
Lease of 10.6.3.237 obtained, lease time 534
```

To configure ma1 statically, use the following command. Add the gateway IP address using the `route add` or `ip route add` command.

```
Aboot# ifconfig ma1 172.24.210.61 netmask 255.255.252.0
Aboot# route add default gateway 172.24.208.1 ma1
OR
Aboot# ip route add default via 172.24.208.1
```

- Identify the MAC address of the ma1 management interface of the switch. Use the `ifconfig -a` command. The HWaddr is the MAC address. A label on the rear of the switch contains the MAC address.



Note: On switches with **Aboot version 6.1.x**, the HWAddr for interface ma1 can display **00:10:18:00:00:00** if there is any delay in entering the Aboot shell. To avoid this, press **Control-C** exactly at the prompt and not later. If there is a delay in entering the Aboot shell, do not use the MAC address **00:10:18:00:00:00** in the next step. Instead, use the MAC address printed on the rear of the switch or the System MAC address from the show version output in EOS.

```
Aboot# ifconfig -a
lo Link encap:Local Loopback
LOOPBACK MTU:65536 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
ma1 Link encap:Ethernet HWaddr C0:D6:82:18:00:3C
inet addr:172.24.210.61 Bcast:172.24.211.255 Mask:255.255.252.0
inet6 addr: fe80::c2d6:82ff:fe18:3c/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:198 errors:0 dropped:0 overruns:0 frame:0
TX packets:24 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:46514 (45.4 KiB) TX bytes:2258 (2.2 KiB)
Interrupt:37
```

- On a separate terminal window, log into the DMF Controller to configure the name of the Arista switch and its MAC address. In the following example, the switch is assigned the name **filter-1**.

```
DMF(config)# switch filter-1
DMF(config-switch)# mac c0:d6:82:18:00:3c
```

- On the DMF Controller, issue the command `show switch-image url` to obtain the Switch Light boot image URL. The required URL is **Update-amd64**. Please note this URL; the following steps need it.

```
DMF(config-switch)# show switch-image url
Install-amd64 : http://172.24.210.21/switchlight/install-amd64
Install-powerpc: http://172.24.210.21/switchlight/install-powerpc
Update-amd64 : http://172.24.210.21/switchlight/amd64
Update-powerpc : http://172.24.210.21/switchlight/powerpc
```

- Return to the switch Aboot shell and boot with the URL of the Switch Light image. The command is `boot url`. The following is the complete console log of a successful Switch Light OS installation in L2 ZTN mode.

```
Aboot# boot http://172.24.210.21/switchlight/amd64
Downloading http://172.24.210.21/switchlight/amd64
Connecting to 172.24.210.21 (172.24.210.21:80)
swi 100% |*****| 316M 0:00:00 ETA
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
4444.69: Running SwitchLight install...
...
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
Booting flash:aboot-chainloader.swi
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
11.24: SKU: DCS-7050SX3-48YC8
11.24: DCS-7050SX3-48YC8: kernel=kernel-4.9-lts-x86_64-all args=console=ttyS0,9600n8
platform=woodpecker scd.lpc_irq=13 scd.lpc_res_addr=0xf00000 scd.lpc_res_size=0x100000
sid=Marysville onl_mnt=/dev/mmcblk0p1 tsc=reliable pcie_ports=native reboot=p pti=off
reassign_prefmem amd iommu_dump=1 platform=x86-64-arista-7050sx3-48yc8-r0
11.24: Loading kernel and initrd...
+ kexec --load --command-line 'console=ttyS0,9600n8 platform=woodpecker scd.lpc_irq=13 scd.
lpc_res_addr=0xf00000 scd.lpc_res_size=0x100000 sid[ 11.969258] kexec_core: Starting
new kernel
=Marysville onl_mnt=/dev/mmcblk0p1 tsc=reliable pcie_ports=native reboot=p pti=off
```

```

reassign_prefmem amd_iommu_dump=1 quiet=1 onl_platform=x86-64-arista-7050sx3-48yc8-r0
onl_sku=DCS-7050SX3-48YC8'--initrd /mnt/flash/onl/boot/x86-64-arista-7050sx3-48yc8-r0.
initrng: Linux Random Number Generator (RNG) early init
found interface name alias eth0 --> mal
remapping interface eth0 --> mal
No dynamic mount operations in unified mode.
No dynamic mount operations in unified mode.
INFO:PKI:Using existing private key.
INFO:PKI:Using existing certificate.
Setting up mal as bonded interface...
mal is now [ omal ]
*****
*
* Switch Light OS Loader
*
* Version: SWL-OS-DMF-8.0.0(0)
* Id: 2020-08-27.14:06-dff2d80
*
* Platform: x86-64-arista-7050sx3-48yc8-r0
* mal: c0:d6:82:18:00:3c
*
*****
[ boot-config ]
NETDEV=mal
NETAUTO=up
BOOTMODE=ztn
ZTNMODE=deferred
Press Control-C now to enter the interactive loader shell.
[ Starting Autoboot ]
[ Configuring Interfaces ]
Waiting for link on mal...
mal: up
[ BOOTMODE is ztn. ]
...
SLREST port is not ready...
Saving switch default settings...done.
Loading ZTN startup-config.....done.
Saving last startup-config.....done.
Stopping watchdog keepalive daemon...
Starting watchdog daemon...
Switch Light OS SWL-OS-DMF-8.0.0(0), 2020-08-27.14:06-dff2d80
filter-1 login:

```

- Once the switch has booted up successfully, the Arista switch will appear as **connected** under the State column when viewed from the DMF Controller using the **show switch DMF-F1** command. For example:

```

DMF-CTRL(config)# show switch DMF-F1
# Switch Name IP Address State Pipeline Mode
- |-----|-----|-----|-----|
1 DMF-F1 fe80::7272:cfff:febd:dcbc%9 connected l3-14-match-push-vlan
DMF-CTRL(config)#

```

4.3.4 Allocating IPv4 Addresses to Fabric Switches

When using L2 ZTF, the DANZ Monitoring Fabric (DMF) Controllers and fabric switches use link-local IPv6 for communication. To enable switches to communicate with external (IPv4) services, configure IP address management (IPAM), which assigns IPv4 addresses to the switches in the fabric from a configured pool of addresses. This configuration enables a fabric switch in L2-ZTN mode to communicate with external services such as NTP, SNMP, and Syslog.

No IPv4 address is required for the switch to interact with the Controller for time synchronization (NTP) and logging (syslog).



Note: This procedure applies only to Layer 2 ZTF. Attempting to configure IPAM when the provisioning mode is set to Layer 3 (Preconfigured) results in the display of an error message.

Static IP Addresses

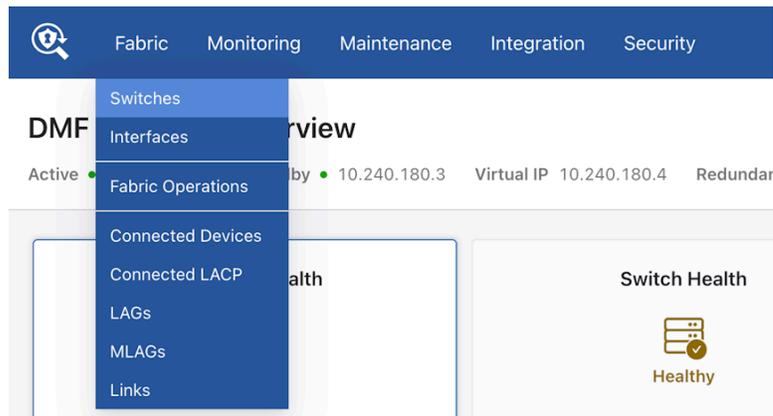
Static IPv4 addresses can be configured on switches in a fabric managed by IPAM. When IPAM is enabled, IPAM will automatically assign IPv4 addresses to switches on which a static IPv4 address has not been configured as long as there are allocated IP addresses available. DMF preserves both automatically and statically allocated IP addresses in the event of a reboot or Controller failure.

4.3.5 Using the GUI to Allocate IPv4 Addresses

To allocate a pool of IPv4 addresses for IPAM to assign to fabric switches, complete the following steps:

1. Select **Fabric > Switches** from the main menu and click on the **IP Address Allocation** tab.

Figure 4-3: Accessing the Switches Page



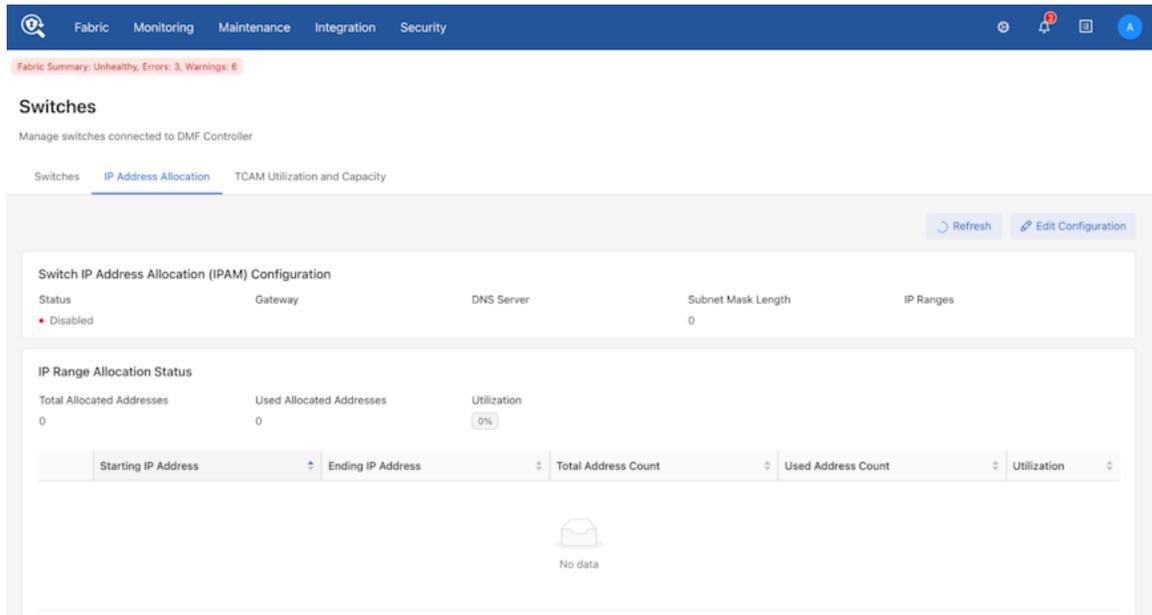
The **Switches** page lists the switches connected to the DANZ Monitoring Fabric (DMF) Controller, and the **IP Address Allocation** tab provides controls for configuring a pool of IPv4 addresses for IPAM assignment to the fabric switches.

Figure 4-4: Switches

Name	Connected	Quarantine Reason	Admin Status	Zero-touch State	Management Interface	Domain	Priority	PTP Source IPv4 Address	Source IPv4 Address	Connected Since	Connection Time	Management IP Address(es)	Associated IP Address	Port	Interface Count	Flow Count
cliff-core-1	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:54:24pm UTC	1M 2w 2d	fe80::10240:142:11	--	42236	33	NA
cliff-core-2	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:54:01pm UTC	1M 2w 2d	fe80::10107:118:38	--	33476	33	NA
cliff-delivery-1	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:53:57pm UTC	1M 2w 2d	fe80::10240:153:132	--	35002	56	NA
cliff-delivery-2	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:53:39pm UTC	1M 2w 2d	fe80::10240:153:131	--	42436	56	NA
cliff-finec-1	✗	--	✓ Up	Reloading	Prefer Dedicated Management Interface	--	--	--	--	--	--	--	--	--	--	NA
cliff-finec-2	✗	--	✓ Up	--	Prefer Dedicated Management Interface	--	--	--	--	--	--	--	--	--	--	NA

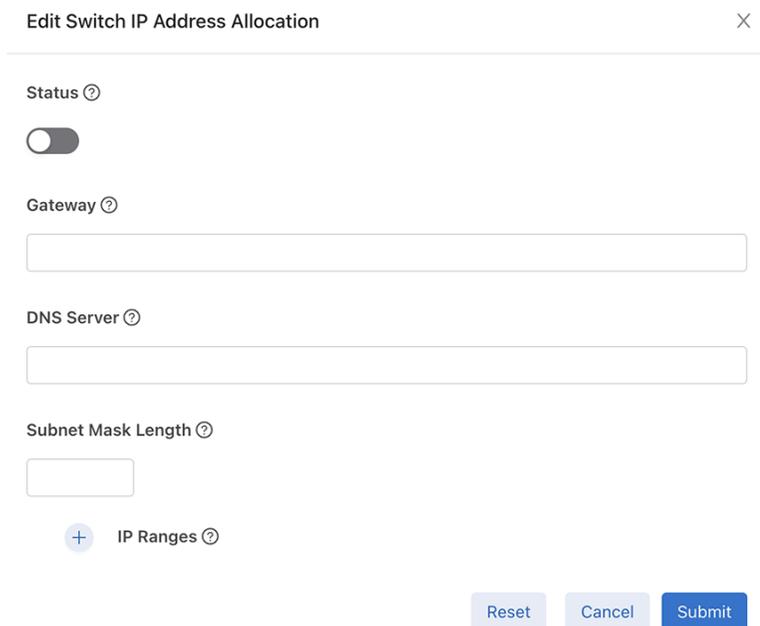
2. Click the **Edit Configuration** control at the top of the **IP Address Allocation** tab.

Figure 4-5: IP Address Allocation Tab



3. In the **Edit Switch IP Address Allocation** dialog, enable IPAM using the **Status** switch, and specify the Gateway, DNS Server, and Subnet Mask Length.

Figure 4-6: Edit Switch IP Address Allocation Dialog



4. Click the **Provision (+)** button next to **IP Ranges** to add IP ranges to be used for IPAM assignment. When finished adding IP address ranges to the pool, click **Submit**.

- To view the operational state of IPAM in the fabric and confirm the address range, select **Fabric > Switches** from the main menu and click on the **IP Address Allocation** tab. To view the IP address assignment for each switch, click on the **+** icon.

Figure 4-7: Viewing Operational State

The screenshot shows the IPAM configuration page. The 'Switch IP Address Allocation (IPAM) Configuration' section displays the following details:

Status	Gateway	DNS Server	Subnet Mask Length	IP Ranges
Enabled	10.0.0.1	8.8.8.8	21	10.0.0.2 - 10.0.0.255

The 'IP Range Allocation Status' section shows:

Total Allocated Addresses	Used Allocated Addresses	Utilization
254	6	2%

Below this, a table lists the IP address allocation details:

Starting IP Address	Ending IP Address	Total Address Count	Used Address Count	Utilization
10.0.0.2	10.0.0.255	254	6	2%

Figure 4-8: Viewing IP Address Assignments

Starting IP Address	Ending IP Address	Total Address Count	Used Address Count	Utilization
10.0.0.2	10.0.0.255	254	6	2%

IP Address Allocation (10.0.0.2 - 10.0.0.255)	
Switch Name	IP Address
C1	10.0.0.2
C2	10.0.0.3
D1	10.0.0.4
D2	10.0.0.5
F1	10.0.0.6
F2	10.0.0.7

4.3.6 Using the CLI to Allocate IPv4 Addresses with IPAM

To allocate a pool of IPv4 addresses for IPAM to assign to fabric switches and configure the DNS server, default gateway, and subnet mask length, complete the following steps:

- Enter the **config-ipam-switch** submode using the **ipam switch** command.

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

- Identify the DNS server IP address to be used by the fabric switches using the **dns-server** command.

```
controller-1(config-ipam-switch)# dns-server 192.168.1.1
```

- Identify the default gateway server IP address to be used by the fabric switches using the **gateway** command.

```
controller-1(config-ipam-switch)# gateway 192.168.1.1
```

- Identify the range of IP addresses to be used by the fabric switches via the **ip-range** command. For multiple non-contiguous address ranges, repeat the command as needed.

```
controller-1(config-ipam-switch)# ip-range 192.168.1.100 192.168.1.200
controller-1(config-ipam-switch)# ip-range 192.168.3.100 192.168.3.200
```



Note: This example allocates 100 addresses in the subnetwork 192.168.1.0 and 100 addresses in the subnetwork 192.168.3.0. To view the IP addresses allocated by IPAM, enter the **show ipam switch** command.

5. Set the length of the subnet mask using the **subnet-mask-length** command. This mask length applies to all configured address ranges.

```
controller-1(config-ipam-switch)# subnet-mask-length 21
```

6. Enable IPAM IPv4 address allocation to the fabric switches using the **allocate** command.

```
controller-1(config-ipam-switch)# allocate
```

7. To display the static or automatic IP address configuration for all switches, use the **show running-config switch** command. The **ip-address** fields contain the newly introduced values.

```
controller-1(config-ipam-switch)# show running-config switch  
  
! switch  
switch core1  
  ip-address auto 10.0.0.2  
  mac 52:54:00:57:c9:3b  
  
switch delivery1  
  ip-address auto 10.0.0.3  
  mac 52:54:00:c2:9c:24  
  
switch filter1  
  ip-address auto 10.0.0.4  
  mac 52:54:00:ab:3a:e6
```



Note: The address displayed does not necessarily mean that the specified IP address has been assigned to the switch. For an address to be automatically assigned, IPAM must be enabled and a corresponding IP range must be defined in the IPAM configuration. In the case of a static IP address, the IP address should exist within the defined IPAM subnet. Use the **show ipam switch** command to display actual allocations rather than the configuration.

4.3.7 Assigning Static IPv4 Addresses

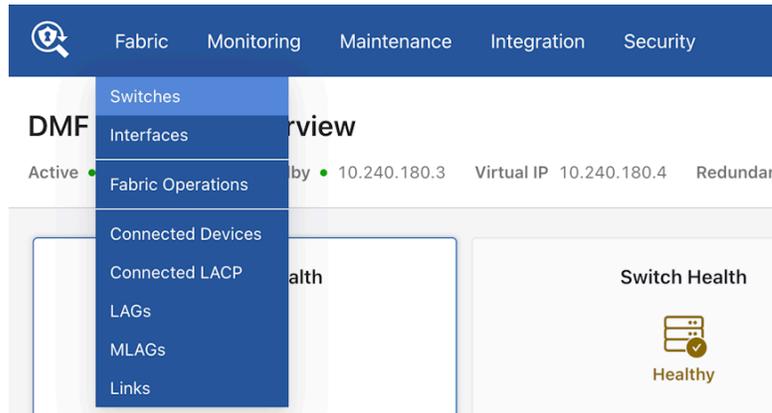
If needed, a static IPv4 address can be assigned to a switch in a fabric managed by IPAM. Removing the assigned address will return the switch to IPAM address management, and an IPv4 address will be assigned to it from the allocated pool if one is available.

4.3.8 Using the GUI to Assign a Static IPv4 Address

To assign a static IPv4 address to a fabric switch, complete the following steps:

1. Select **Fabric > Switches** from the main menu and select the **IP Address Allocation** tab.

Figure 4-9: Accessing the Switches Page



2. On the **Switches** tab, click the menu icon next to the switch to which the address will be assigned, and select **Configure** from the menu.

Figure 4-10: Switches

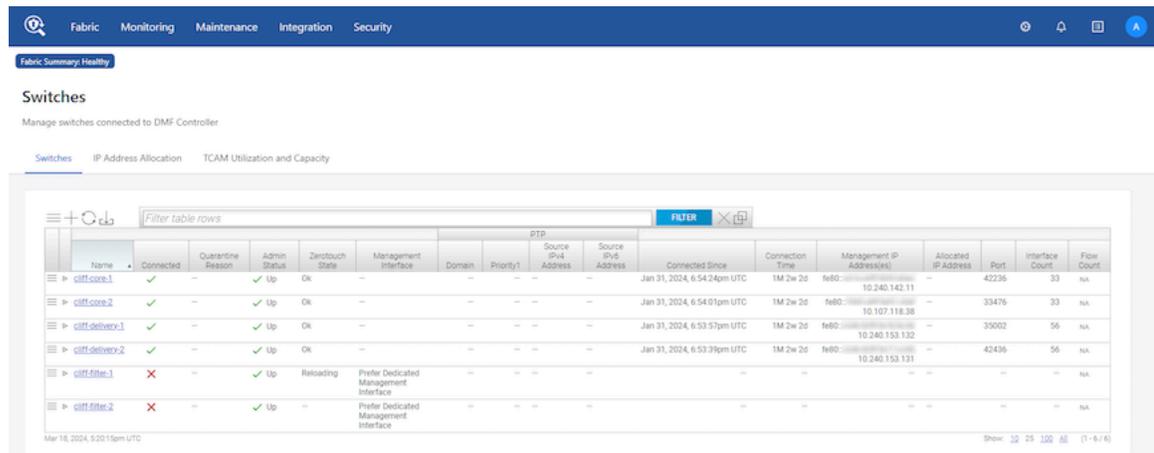
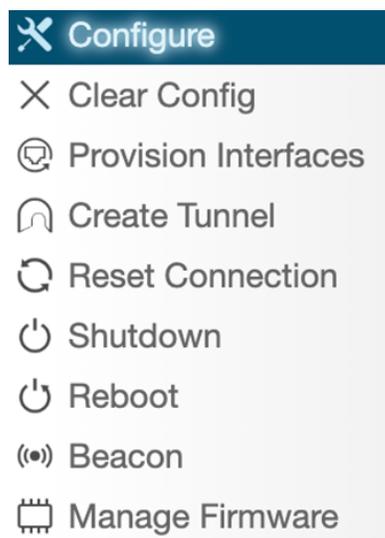


Figure 4-11: Configure



- In the **Configure Switch** dialog, set the **IP Assignment Type** to `Manual` and enter the desired IPv4 address for the switch.

Figure 4-12: Configure Switch



Note: The assigned static address must be in the defined IPAM subnet.

- Click **Submit** to assign the configured IPv4 address to the switch.
- To confirm the IP address assignment for the switch, select **Fabric > Switches** from the main menu and click on the **IP Address Allocation** tab, then click on the **+** icon to display IP addresses for all switches.

Figure 4-13: Viewing IP Address Assignments

Starting IP Address	Ending IP Address	Total Address Count	Used Address Count	Utilization
10.0.0.2	10.0.0.255	254	6	2%

IP Address Allocation (10.0.0.2 - 10.0.0.255)	
Switch Name	IP Address
C1	10.0.0.2
C2	10.0.0.3
D1	10.0.0.4
D2	10.0.0.5
F1	10.0.0.6
F2	10.0.0.7

4.3.9 Using the CLI to Assign a Static IPv4 Address

When IPAM is enabled and there are enough IPv4 addresses allocated for assignment, it automatically assigns IPv4 addresses to each switch in the fabric. The `show running-config` command displays those addresses as `auto` as shown below:

```
controller-1(config)# show running-config switch core1

! switch
switch core1
 ip-address auto 10.0.0.2
 mac 00:53:00:57:c9:3b
```

To assign a static IPv4 address to a fabric switch instead, complete the following steps:

1. Enter switch configuration mode for the switch to which a static address will be assigned using the **switch** command.

```
controller-1(config)# switch core1
controller-1(config-switch)#
```

2. Set the IP address to `static` and configure the address using the **ip-address static** command.



Note: The assigned IPv4 address must be within the defined IPAM subnet.

```
controller-1(config-switch)# ip-address static 10.0.0.3
controller-1(config-switch)#
```

3. To confirm the address assignment, use the **show this** and **show ipam switch** commands.

```
controller-1(config-switch)# show this

! switch
switch core1
  ip-address static 10.0.0.3
  mac 00:53:00:57:c9:3b

controller-1(config-ipam-switch)# show ipam switch
~~~~~ Allocated IP Addresses ~~~~~
# Start IP End IP      Count Used Switch      Allocated IP
-|-----|-----|-----|----|-----|-----|
1 10.0.0.2 10.0.0.255 254   3   delivery1 10.0.0.2
2 10.0.0.2
3 10.0.0.2
                               filter1    10.0.0.4
```

4. To remove a static IPv4 address assignment from an IPAM-managed fabric switch (restoring it to `auto`), use the **no ip-address static** command.

```
controller-1(config)# switch core1
controller-1(config-switch)# no ip-address static 10.0.0.3
controller-1(config-switch)# show this

! switch
switch core1
  ip-address auto 10.0.0.3
  mac 52:54:00:01:b4:18
```



Note: The switch IP address may not change when the static address is removed. In the previous example, DMF removed the static IP address 10.0.0.3; however, since IPAM is still enabled, a new IP address was immediately assigned from the IPAM IP range. The IP address 10.0.0.3 was the first available IP address in the pool and was therefore assigned to the switch. The address remains the same in this case, but the status changes to `auto`.

4.3.10 Static Address Assignment Troubleshooting and Limitations

Syslog Messages and Tracing

There are no syslog messages associated with this feature. To gain insight into the IPAM IP allocation process, enable tracing logs.



Note: In some cases, enabling tracing can seriously impact switch performance. Please use it cautiously and seek advice from an Arista Networks representative before enabling tracing in any production environment.

Enable more detailed tracing logs using the following commands:

```
controller-1(config)#logging level org.projectfloodlight.core.ipalloc trace
controller-1(config)#logging level org.projectfloodlight.zerotouch.startupconfig trace
controller-1(config)#show logging controller | grep Ipam
```

Locate relevant log output in the floodlight syslog at `/var/log/floodlight/floodlight.log`.

Troubleshooting

When troubleshooting, be sure to use the `show ipam switch` command to display actual IP address allocations instead of the `show running-config ipam switch` command, which only displays the configuration and last used `auto` IP addresses of the switches.

If a switch does not receive the expected IP address allocation, ensure that:

- IPAM is enabled.
 - Make sure that the `allocate` field is present in the IPAM configuration. If not, configure the deployment mode as shown below:
- Make sure that the ZTN deployment mode is set to `auto-discovery`. If not, configure the deployment mode as shown below:

```
controller-1(config)#ipam switch
controller-1(config-ipam-switch)#allocate
```

```
controller-1(config)deployment-mode auto-discovery
```

- For a configured `s` IP address, make sure that the address is in the defined IPAM subnet.
- For an `auto` IP address, make sure that there are enough IP addresses in the defined IP address ranges for all switches.



Note: An `auto` IP address in a switch configuration does not necessarily result in an actual IP address assignment and stays there even if the allocated IP address is unavailable, e.g., if IPAM is disabled or the corresponding IP address range is removed. If a configuration change (such as enabling IPAM or adding an IP address range that includes the allocated address), the switch will get this IP address, which was the last automatically allocated IP address, to promote IP stability.

- You may cross-check the applied IP address on the running configuration of the switch as shown below:

```
controller-1#show switch core1 running-config | grep ip-address
sw1 interface ip-address 10.0.0.3 prefix 21
```

- An alternative way to cross-check the IP address of the switch is to connect to it and then execute the `ifconfig` command to view the IP address of the interface:

```
> connect switch core1
Switch Light OS SWL-OS-DMF-8.5.x(0), 2023-12-01.02:24-4be6844
Linux core1 4.19.296-OpenNetworkLinux #1 SMP Fri Dec 1 02:35:57 UTC 2023 x86_64

SwitchLight 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@core1:~# ifconfig -a
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
       inet6 2001:0DB8:0:1:5054:ff:fe59:b9b6 prefixlen 64 scopeid 0x20<link>
...
mal: flags=5187<UP,BROADCAST,RUNNING,MASTER,MULTICAST> mtu 1500
       inet 10.0.0.2 netmask 255.255.248.0 broadcast 0.0.0.0
...

```

Limitations

- Note that the `show switch` command does not display IPv4 addresses. To cross-check the assigned IPv4 number, examine the running config of the switch. (See *Troubleshooting* section.)
- In order to enable IPAM, ZTN deployment mode must be configured as `auto-discovery`.
- IPAM can only manage switch IP addresses in a single subnet, but multiple IP address ranges can be defined in that subnet.

4.4 Using L3 ZTN (Pre-Configured) Switch Provisioning Mode

When a switch is in a different subnet than the Controller, configure the network information in the switch, which enables the switch loader to download the correct Switch Light OS image from the Controller and install it on the switch via the ZTN process.



Note: DANZ Monitoring Fabric (DMF) supports deploying switches from different vendors in the same fabric. However, using cables from different vendors to connect switches is not supported. Optics are required to interconnect switches from different vendors. See the Hardware Compatibility List for details on supported optics for each switch platform.

Ensure the `deployment-mode pre-configured` command is entered on the DMF Controller to enable Layer 3 ZTF.



Note: In Layer 3 mode, ZTF uses TCP port **8843** for communication between the Controller and switches. This port must be allowed on the Controller and on any devices connecting the Controller to the fabric switches.

4.4.1 Installing a Switch Using L3 ZTF (Preconfigured) Provisioning Mode

To install a switch with the DANZ Monitoring Fabric (DMF) provisioning mode set to pre-configured, complete the following steps:

1. Confirm that the switch has ONIE installed.



Note: Power on the switch and connect to the switch console using the default baud rate. The default baud rate is **9600** for switches by Arista Networks. The default baud rate is **115200** for ONIE-enabled switches.

The supported switches, listed in the ***DANZ Monitoring Fabric 8.5 Hardware Compatibility List***, come with ONIE installed by the manufacturer.

2. Verify that the switch management port is connected to the management network with an IP address assigned either manually or from a DHCP server. If an IP address in **Step 2a** is not assigned, follow **Steps 2b** and **2c**. If an IP address is already set in **Step 2a**, skip to **Step 3**.

- a. Check IP addressing on the management port.

```
ONIE:/ # ifconfig
eth0 Link encap:Ethernet HWaddr 90:B1:1C:F4:CB:A9
inet addr:10.240.130.96 Bcast:10.240.130.127 Mask:255.255.255.128
inet6 addr: fe80::92b1:1cff:fef4:cba9/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:5550 errors:0 dropped:0 overruns:0 frame:0
TX packets:3937 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:728610 (711.5 KiB) TX bytes:563219 (550.0 KiB)
Interrupt:21 Memory:ff300000-ff320000
ONIE:/ #
```

- b. Manually assign an IP address to the management port.

```
ONIE:/ # ifconfig eth0 192.168.10.10 netmask 255.255.255.0
```

- c. Add the default gateway.

```
ONIE:/ # route add default gw 192.168.10.1 eth0
ONIE:/ # netstat -arn
Kernel IP routing table
Destination Gateway Genmask Flags MSS Window irtt Iface
0.0.0.0 192.168.10.1 0.0.0.0 UG 0 0 0 eth0
192.168.10.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
ONIE:/ #
```



Note: To verify or change the IP address and gateway, use the `ifconfig` command.

3. On a separate terminal window, log into the DMF Controller to configure the name of the Arista switch and its MAC address. In this example, the switch is assigned the name **s1-s6100**.

```
DMF(config)# switch s1-s6100
DMF(config-switch)# mac 4c:76:25:f6:c8:80
```

4. On the DMF Controller, identify the switch image URL for the installation files by entering the `show switch-image url` command, as in the following example.

```
controller-1# show switch-image url
Install-amd64 : http://10.8.25.31/switchlight/install-amd64
Install-powerpc: http://10.8.25.31/switchlight/install-powerpc
Update-amd64 : http://10.8.25.31/switchlight/amd64
Update-powerpc : http://10.8.25.31/switchlight/powerpc
Update-powerpc : http://10.8.34.1/switchlight/powerpc
```

For example, for a switch using amd-64 hardware, the URL would be `http://10.8.25.31/switchlight/install-amd64`.

5. For AMD-64 platform boxes (for example, Dell S6100):

- a. Reboot and select **ONIE** from the GNU/GRUB menu.
- b. To get to the ONIE mode and during the reboot countdown, press any key when the prompt: Hit any key to stop autoboot: 0 appears. The following command enters the ONIE install mode: `run onie_bootcmd`
- c. Type `onie-discovery-stop` at the installer mode in ONIE.

```
ONIE:/ # onie-discovery-stop
```

6. Verify the switch can ping the Controller IP address.

```
ONIE:/ # ping 10.8.25.31
PING 10.8.25.31 (10.8.25.31): 56 data bytes
64 bytes from 10.8.25.31: seq=0 ttl=63 time=0.398 ms
64 bytes from 10.8.25.31: seq=1 ttl=63 time=0.434 ms
```

7. Use the image-url path from the Controller for the platform amd-64.

```
ONIE:/ # onie-nos-install http://10.8.25.31/switchlight/install-amd64

discover: Rescue mode detected. No discover stopped.
```

Note: The switch now goes into ZTN discover.

```
INFO:PKI:Using existing certificate.
*****
*
* Switch Light OS Loader
...
Press Control-C now to enter the interactive loader shell.
^C
Welcome to the shell.
Type 'help' for command help.
```

8. Enter **Control-C** to drop into the loader mode.

```
Press Control-C now to enter the interactive loader shell.
^C
Welcome to the shell.
Type 'help' for command help.
loader#
```

9. Enter **zcs** to drop into **ztn-config** mode.

```
loader# zcs
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
(ztn-config)
(ztn-config) help
Documented commands (type help <topic>):
=====
controller debug dns help interface reboot setup show version
Undocumented commands:
=====
EOF exit quit
(ztn-config)
```

10. Configure the static IP address on the switch.

```
(ztn-config) interface ma1 ip-address 10.8.67.135/24 gateway 10.8.67.1
```

 **Note:** Attempting to configure IPAM when the provisioning mode is set to Layer 3 (Preconfigured) results in an error message.

11. Configure the IP address of the DMF Controller (ZTN server).

```
(ztn-config) controller set 10.8.25.31,10.8.25.32
```

12. Set the DNS information.

```
(ztn-config) dns server 10.3.0.4
(ztn-config) dns domain qa.arista.com
```

13. Enter the **show** command to verify the configuration.

```
(ztn-config) show
```

```
IP Option: Static
IP Address: 10.8.67.135
Netmask: 255.255.255.0
Gateway: 10.8.67.1
DNS Server: 10.3.0.4
DNS Domain: qa.arista.com
Controllers: 10.8.25.32,10.8.25.31
(ztn-config)
```

14. Reboot the switch to activate the configuration.

```
(ztn-config) reboot
```



Note: The switch now automatically downloads the SWI image from the Controller and boots up with the image.

```
##### Start - Switch Console output in DMF-L3 ZTN mode
...
Switch Light OS SWL-OS-DMF-7.0.0(0), 2018-01.21.00:19-784f432
##### END - Switch Console output in DMF-L3 ZTN mode
```



Note: When the final messages appear, the installation is completed.

15. To verify the installation, enter the following **show** commands on the Controller.

```
controller-1(config)# show switch s1-s6100 details
# Switch Name Mac address Switch DPID State
- |-----|-----|-----|-----|
1 s1-s6100 4c:76:25:f6:c8:80 (Dell) 00:00:4c:76:25:f6:c8:80 connected
Quarantine reason IP Address TCP Port Connected Since Pipeline Mode
|-----|-----|-----|-----|
10.8.67.135 38018 2018-10.01 11:27:48.663000 PST 1314-push-vlan
```

```
controller-1 (config)# show switch s1-s6100 zerotouch
Device : 4c:76:25:f6:c8:80 (Dell)
Zerotouch state : ok
Name : s1-s6100
Reload pending : False
Platform : x86-64-dell-s6100-c2538-r0
Serial number : CNOWKFYN7793164P0004
Ip address : 10.8.67.135
Dpid : 00:00:4c:76:25:f6:c8:80
Last update : 2018-10.01 15:06:13.619000 PST
Controller address : 10.8.25.31
controller-1 (config)#
```

16. Log in to the switch and enter the **show** command to display the switch status when the following (ztn-config) prompt appears.

```
Switch Light OS SWL-OS-DMF-6.3.0(0), 2017-01-24.00:19-784f432
s1-s6100 login: admin
Password:
Linux s1-s6100 3.16.39-OpenNetworkLinux #1 SMP Tue Jan 24 08:38:28 UTC 2017 x86_64
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
(ztn-config) show
IP Option : Static
IP Address : 10.8.67.135
Netmask : 255.255.255.0
Gateway : 10.8.67.1
Controllers: 10.8.25.33,10.8.25.32,10.8.25.31
```

4.4.2 Installing Arista 7050X and 7260X Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode

Perform the initial installation of Switch Light OS on the Arista platforms manually in the DANZ Monitoring Fabric (DMF), similarly to the requirements for ONIE-enabled switches. Use this guide for DMF Controllers configured for L3 ZTN mode.



Note: L3 ZTN means the Controller is set up for deployment-mode pre-configured.

The initial installation of Switch Light OS on the Arista switch is accomplished by dropping it to the Aboot shell interface at boot and telling it to boot the Switch Light switch image. This operation installs Switch Light on the system. This is a one-time extra step needed during the first installation of Switch Light in DMF. The boxes will subsequently boot as expected under Switch Light.



Note: This procedure is also required for any Arista switches running EOS.

Procedure

1. Attach a console cable to the Arista switch and power on or reboot the switch.
2. Interrupt the boot process with **Control-C** to drop into the Aboot shell.

```
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
```

3. Configure an IP address for the switch's **ma1** management interface. Configure ma1 statically or use DHCP.
4. To use DHCP, type **udhcpc -i ma1**.

```
Aboot# udhcpc -i ma1
udhcpc (v1.18.1) started
Sending discover...
Sending discover...
Sending select for 10.6.3.237...
Lease of 10.6.3.237 obtained, lease time 534
```

5. To configure ma1 statically.

```
Aboot# ifconfig ma1 172.24.210.61 netmask 255.255.252.0
Aboot# route add default gateway 172.24.208.1 ma1
OR
Aboot# ip route add default via 172.24.208.1
Tip: Please use the "route" command to verify the correct routing table.
```

6. Identify the MAC address of the ma1 management interface of the switch. Use the **ifconfig -a** command. The HWaddr is the MAC address. A label on the rear of the switch contains the MAC address.

```
Aboot# ifconfig -a
lo Link encap:Local Loopback
LOOPBACK MTU:65536 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1
```

```

RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
mal Link encap:Ethernet HWaddr C0:D6:82:18:00:3C
inet addr:172.24.210.61 Bcast:172.24.211.255 Mask:255.255.252.0
inet6 addr: fe80::c2d6:82ff:fe18:3c/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:198 errors:0 dropped:0 overruns:0 frame:0
TX packets:24 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:46514 (45.4 KiB) TX bytes:2258 (2.2 KiB)
Interrupt:37

```

7. On a separate terminal window, log into the DMF Controller to configure the name of the Arista switch and its MAC address. In this example, the switch is assigned the name *filter-1*.

```

DMF(config)# switch filter-1
DMF(config-switch)# mac c0:d6:82:18:00:3c

```

8. On the DMF Controller, issue the command `show switch-image url` to obtain the Switch Light boot image URL. The URL that is needed is *Update-amd64*. Please note this URL; the following steps need it. Do **not** use the other URLs.

```

DMF(config-switch)# show switch-image url
Install-amd64 : http://172.24.210.21/switchlight/install-amd64
Install-powerpc: http://172.24.210.21/switchlight/install-powerpc
Update-amd64 : http://172.24.210.21/switchlight/amd64
Update-powerpc : http://172.24.210.21/switchlight/powerpc

```

9. Back on to the switch About shell, boot with the URL of the Switch Light image. The command is `boot url`.

```

About# boot http://172.24.210.21/switchlight/amd64
Downloading http://172.24.210.21/switchlight/amd64
Connecting to 172.24.210.21 (172.24.210.21:80)
swi 100% |*****| 316M 0:00:00 ETA
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
93.50: Running SwitchLight install...
Archive: /tmp/swi
..
..

```

10. Interrupt the Switch Light OS Loader with **Control-C** to drop into the interactive loader shell.

```

No dynamic mount operations in unified mode.
No dynamic mount operations in unified mode.
INFO:PKI:Using existing private key.
INFO:PKI:Using existing certificate.
Setting up mal as bonded interface...
mal is now [ omal ]
*****
*
* Switch Light OS Loader
*
* Version: SWL-OS-DMF-8.0.0(0)
* Id: 2020-08-27.14:06-dff2d80
*
* Platform: x86-64-arista-7050sx3-48yc8-r0
* mal: c0:d6:82:18:00:3c
*
*****
[ boot-config ]
NETDEV=mal
NETAUTO=up
BOOTMODE=ztn
ZTNMODE=deferred
Press Control-C now to enter the interactive loader shell.
^C
Welcome to the shell.
Type 'help' for command help.
loader#

```

11. Type **zcs** to start the Switch Light ZTN manual configuration.

```
loader# zcs
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
(ztn-config)
```

12. Type **setup** followed by **Enter** to begin interactive setup.

```
(ztn-config) setup
You are now running the interactive setup.
Press Enter to continue...
```

13. Configure the IP address of the management interface on the switch. Choose DHCP or Static. Optionally configure the DNS parameters. The following example uses a static IP address.

```
Please choose an IP option:
(DHCP/Static)? Static
Please provide static IP settings:
IP Address: 172.24.210.64
Netmask: 255.255.252.0
Gateway: 172.24.208.1
Do you want to configure DNS settings?
(Yes/No)? yes
DNS Server: 10.3.0.4
DNS Domain: arista.com
```

14. Configure the IP addresses of the primary and secondary DMF Controller. In this example, there is no secondary Controller.

```
Please provide the IP address of the controller:
Controller IP: 172.24.210.21
Do you have a second controller?
(Yes/No)? No
```

15. Review the manual ZTN configuration. Type **Yes** to complete the Switch Light ZTN manual configuration.

```
Configuration Summary:
IP Option: Static
IP Address: 172.24.210.64
Netmask: 255.255.252.0
Gateway: 172.24.208.1
DNS Server: 10.3.0.4
DNS Domain: arista.com
Controller IP: 172.24.210.21
Please confirm that the above settings are correct:
(Yes/Reset)?
(Yes/Reset)? Yes
Interactive setup completed successfully.
(ztn-config)
```

16. Type **reboot**. The Arista switch will reboot and complete the ZTN process using the previously configured boot parameters.

```
(ztn-config) reboot
Proceed with reboot [confirm]
Requesting system reboot
[ 2115.237162] reboot: Restarting system
coreboot-coreboot-unknown-Aboot-norcal9-9.0.3-4core-14223577 Wed Nov 13 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxC0 STATUS: 0x800
...
...
*****
*
* Switch Light OS Loader
*
* Version: SWL-OS-DMF-8.0.0(0)
* Id: 2020-08-27.14:06-dff2d80
*
* Platform: x86-64-arista-7050sx3-48yc8-r0
* mal: c0:d6:82:18:00:3c
```

```

*
*****
[ boot-config ]
ZTNSERVERS=172.24.210.21
NETGW=172.24.208.1
NETDEV=mal
NETDOMAIN=arista.com
BOOTMODE=ztn
NETMASK=255.255.252.0
NETIP=172.24.210.64
ZTNMODE=deferred
NETDNS=10.3.0.4
Press Control-C now to enter the interactive loader shell.
[ Starting Autoboot ]
[ Configuring Interfaces ]
[ BOOTMODE is ztn. ]
....
Saving switch default settings...done.
Loading ZTN startup-config.....done.
Saving last startup-config.....done.
Stopping watchdog keepalive daemon....
Starting watchdog daemon....
Switch Light OS SWL-OS-DMF-8.0.0(0), 2020-08-27.14:06-dff2d80
filter-1 login:

```

- Once the switch has booted up successfully, the Arista switch will appear as connected under the State column when viewed from the DMF Controller using the `show switch` command. For example:

```

DMF(config)# show switch
# Switch Name  IP Address      State           Pipeline Mode
- |-----|-----|-----|-----|
1 filter-1     172.24.210.64  connected      13-14-push-vlan

```

4.4.3 Installing Arista 7280R Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode

DANZ Monitoring Fabric (DMF) 8.1.0 is the first release to support the Arista 7280R, 7280R2, and 7280R3 series of switches. These switches use Arista Networks EOS operating system instead of running Switch Light OS while deployed in the DANZ Monitoring Fabric.

For all the 7280 series of switches, configure the Controller with deployment mode pre-configured (L3 ZTN). This is the default mode of operation starting in the DMF 8.5.0 release.



Note: L3 ZTN means the Controller is set up for deployment mode pre-configured.

This is a one-time setup needed to load the DMF-compatible EOS image. When set up, the next Controller upgrade will also automatically upgrade the switches.

Perform these steps on the 7280R Series switch to boot from the DMF Controller.

- On the DMF Controller, issue the command `show switch-image url` to obtain the URL for the boot image. The URL that is needed is `update-aristaeos`. Make a note of this URL, which will be required later when copying the image to the switch. Do **not** use the other URLs.

```

DMF(config)# show switch-image url
# File          Url
-|-----|-----|
1 arista-ztp-install-script http://172.24.210.21/switchlight/arista-ztp-install-script
2 arista-ztp-install-script https://172.24.210.21/switchlight/arista-ztp-install-script
3 install-amd64 http://172.24.210.21/switchlight/install-amd64
4 install-amd64 https://172.24.210.21/switchlight/install-amd64
5 update-amd64 http://172.24.210.21/switchlight/amd64
6 update-amd64 https://172.24.210.21/switchlight/amd64

```

```
7 update-aristaeos http://172.24.210.21/eos/x86_64
8 update-aristaeos https://172.24.210.21/eos/x86_64
```

2. Attach a console connection to the Arista 7280R Series switch and turn on or reboot the switch.
3. Interrupt the boot process with **Control-C** to drop into the Aboot shell.

```
Warning - AGESA callout: platform_PcieSlotResetControl not supported
agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0)
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
```

4. Configure an IP address for the switch's ma1 management interface. Configure ma1 statically or use DHCP.
5. To use DHCP, type `udhcpc -i ma1`.

```
Aboot# udhcpc -i ma1
udhcpc: started, v1.30.1
udhcpc: sending discover
udhcpc: sending select for 172.24.208.123
udhcpc: lease of 172.24.208.123 obtained, lease time 86400
```

6. To configure ma1 statically.

```
Aboot# ifconfig ma1 172.24.210.61 netmask 255.255.252.0
Aboot# route add default gateway 172.24.208.1 ma1
Tip: Please use the "route" command to verify the correct routing table.
```

7. Change directory to `/mnt/flash/` on the switch.

```
Aboot# cd /mnt/flash
Aboot#
Aboot# ls -l
-rw-rw-r-- 1 root 88 2541 Apr 29 19:29 AsuFastPktTransmit.log
drwxrwxr-x 2 root 88 4096 Oct 31 01:06 Fossil
-rw-rw-r-- 1 root 88 1562 Apr 29 19:29 SsuRestore.log
-rw-rw-r-- 1 root 88 1562 Apr 29 19:29 SsuRestoreLegacy.log
-rw-rwx--- 1 root 88 47 Apr 29 19:37 boot-config
-rw-rw-r-- 1 root 88 4 Apr 29 18:19 config_match
drwxrwx--- 3 root 88 4096 Apr 29 19:38 debug
drwxrwxr-x 2 root 88 4096 Oct 31 01:06 fastpkttx.backup
drwxrwx--- 2 root 88 16384 Oct 31 01:04 lost+found
drwxrwxr-x 3 root 88 4096 Apr 29 19:36 persist
drwxrwxr-x 3 root 88 4096 Oct 31 01:20 schedule
-rw-rw-r-- 1 root 88 0 Oct 31 01:20 startup-config
-rw-rw-r-- 1 root 88 0 Apr 29 19:30 zerotouch-config
```

8. Use `wget` to copy the Arista EOS image from the DMF Controller. Use the `update-aristaeos` URL from the **Step 1**. The command is `wget url`.

```
Aboot# wget http://172.24.210.21/eos/x86_64
Connecting to 172.24.210.21 (172.24.210.21:80)
x86_64 100% |*****| 933M 0:00:00 ETA
```

9. Edit the `/mnt/flash/boot-config` file. Boot using the newly downloaded EOS image from the DMF Controller.

```
SWI=flash:/x86_64
```

10. Verify the boot-config was saved.

```
Aboot# cat /mnt/flash/boot-config
SWI=flash:/x86_64
```

11. Reboot the system. Type: **reboot**.

```
Abboot# reboot
Abboot# [ 1096.250482] sysrq: SysRq : Remount R/O
Requesting system reboot
Restarting system
coreboot-coreboot-unknown-Abboot-norcal9-9.0.3-4core-14223577 Wed Nov 13 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxCO STATUS: 0x800
BIT11
agesawrapper_amdinitreset() entry
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'AGESA'
CBFS: Found @ offset dffdc0 size 71786
Fch OEM config in INIT RESET Done
coreboot-coreboot-unknown-Abboot-norcal9-9.0.3-4core-14223577 Wed Nov 13 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxCO STATUS: 0x80800
DoReset BIT11
agesawrapper_amdinitreset() entry
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'AGESA'
CBFS: Found @ offset dffdc0 size 71786
Fch OEM config in INIT RESET Done
agesawrapper_amdinitreset() returned AGESA_SUCCESS
agesawrapper_amdinitearly() entry
Warning - AGESA callout: platform_PcieSlotResetControl not supported
agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Abboot 9.0.3-4core-14223577
Press Control-C now to enter Abboot shell
Booting flash:/x86_64
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
[ 12.590004] kexec_core: Starting new kernel
[ 0.972580] Running e2fsck on: /mnt/flash
[ 4.216655] Running e2fsck on: /mnt/crash
Switching rootfs
starting version 219
Welcome to Arista Networks EOS 4.26.0FX-DMF
New seat seat0.
Starting ProcMgr: Removing all files in all subdirs of /etc/ProcMgr.d/run
[ OK ]
Starting EOS initialization stage 1: [ OK ]
Starting NorCal initialization: [ OK ]
Starting EOS initialization stage 2: [ OK ]
Completing EOS initialization (press ESC to skip): [ OK ]
Model: DCS-7280CR3-32P4
Serial Number: JPE20383403
System RAM: 8147180 kB
Flash Memory size: 7.1G
Apr 29 19:59:55 localhost SandFapNi: %AGENT-6-INITIALIZED: Agent 'SandFapNi-FixedSystem'
initialized; pid=3054
Apr 29 19:59:55 localhost PowerManager: %PWRMGMT-4-INPUT_POWER_LOSS: PowerSupply1 has lost
input power.
Apr 29 19:59:55 localhost SandMcast: %AGENT-6-INITIALIZED: Agent 'SandMcast' initialized;
pid=3051
No startup-config was found.
The device is in Zero Touch Provisioning mode and is attempting to
download the startup-config from a remote system. The device will not
be fully functional until either a valid startup-config is downloaded
from a remote system or Zero Touch Provisioning is cancelled.
To cancel Zero Touch Provisioning, login as admin and type
'zerotouch cancel' at the CLI. Alternatively, to disable Zero Touch
Provisioning permanently, type 'zerotouch disable' at the CLI.
Note: The device will reload when these commands are issued.
localhost login:
```

12. Log in as admin. From the enable mode, turn off Zero Touch Provisioning. The system will reload again.
Type: **zerotouch cancel**.

```
localhost> en
localhost# zerotouch cancel
Apr 29 20:00:12 localhost ZeroTouch: %ZTP-6-CANCEL: Cancelling Zero Touch Provisioning
Apr 29 20:00:12 localhost ZeroTouch: %ZTP-6-RELOAD: Rebooting the system
localhost# Flushing AAA accounting queue: [ OK ]
Restarting system
[20:00:14] watchdog punch .
[20:00:14] watchdog punch .
[20:00:15] watchdog punch .
[20:00:16] watchdog punch .
[ 176.580458] sysrq: Remount R/O
[20:00:16] watchdog punch .
[20:00:16] watchdog punch .
[20:00:17] watchdog punch .
[20:00:18] watchdog punch .
coreboot-coreboot-unknown-Aboot-norcal9-9.0.3-4core-14223577 Wed Nov 13 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxCO STATUS: 0x800
BIT11
agesawrapper_amdinitreset() entry
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'AGESA'
CBFS: Found @ offset dffdc0 size 71786
Fch OEM config in INIT RESET Done
coreboot-coreboot-unknown-Aboot-norcal9-9.0.3-4core-14223577 Wed Nov 13 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxCO STATUS: 0x80800
DoReset BIT11
agesawrapper_amdinitreset() entry
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'AGESA'
CBFS: Found @ offset dffdc0 size 71786
Fch OEM config in INIT RESET Done
agesawrapper_amdinitreset() returned AGESA_SUCCESS
agesawrapper_amdinitearly() entry
Warning - AGESA callout: platform_PcieSlotResetControl not supported
agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0]
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
Booting flash:/x86_64
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
[ 12.512179] kexec_core: Starting new kernel
[ 0.976275] Running e2fsck on: /mnt/flash
[ 4.186065] Running e2fsck on: /mnt/crash
Switching rootfs
starting version 219
Welcome to Arista Networks EOS 4.26.0FX-DMF
New seat seat0.
Starting ProcMgr: Removing all files in all subdirs of /etc/ProcMgr.d/run
[ OK ]
Starting EOS initialization stage 1: [ OK ]
Starting NorCal initialization: [ OK ]
Starting EOS initialization stage 2: [ OK ]
Completing EOS initialization (press ESC to skip): [ OK ]
Model: DCS-7280CR3-32P4
Serial Number: JPE20383403
System RAM: 8147180 kB
Flash Memory size: 7.1G
localhost login:
```

13. Log in to the switch as admin. Get the System MAC address of the switch from the `show version` command's output. Type: `show version`. In this example, the System MAC address is **d4af.f754.195b**. Obtain the MAC address located on the ID label on the rear of the 7280R Series switch.

```
localhost login: admin
Output to this terminal is being recorded for diagnostic purposes.
Note that only output that is visible on the console is recorded.
localhost> show version
Arista DCS-7280CR3-32P4-F
Hardware version: 12.25
Serial number: JPE20383403
Hardware MAC address: d4af.f754.195b
System MAC address: d4af.f754.195b
Software image version: 4.26.0FX-DMF-21985293.4260FXDMF (engineering build)
Architecture: x86_64
Internal build version: 4.26.0FX-DMF-21985293.4260FXDMF
Internal build ID: 568674e7-5c84-4fc6-8a42-8fb55d2fa639
Uptime: 0 weeks, 0 days, 0 hours and 27 minutes
Total memory: 8147180 kB
Free memory: 6097456 kB
```

14. On the DMF Controller, assign a switch name and configure the System MAC address noted from the previous step. The MAC address must be entered in colon format, such as **d4:af:f7:54:19:5b**.

```
DMF(config)# switch filter-1
DMF(config-switch)# mac d4:af:f7:54:19:5b
```

15. On the switch console, configure the IP address of the management interface.

```
localhost(config)# interface Management1
localhost(config-if-Mal)# ip address 172.24.210.89/22
```

16. Configure a route to the gateway using the following command.

```
localhost(config)# ip route 0.0.0.0/0 172.24.208.1
```

17. On the switch console, configure the IP address(es) of the DMF Controllers. For dual Controllers, the syntax is: `controller address controller#1 controller#2`. In this example, there is only one Controller. ZTN configuration download will begin after configuring this part.

```
localhost(config-if-Mal)# management dmf
localhost(config-mgmt-dmf)# controller address 172.24.210.21
localhost(config-mgmt-dmf)# no disabled
```

18. Verify the switch is connected to the DMF Controller. From the switch console, type: `show management dmf indigo`.

```
filter-1(config)# show management dmf indigo
DMF: enabled
Indigo agent: active
TCAM profile programming status: success
Controllers:
ID      IP Address      Connection State      Connection Role
-----|-----|-----|-----|
0      172.24.210.21   connected              active
```

19. The Arista switch will appear as **connected** under the State column when viewed from the DMF Controller using the `show switch` command. For example:

```
DMF(config-switch)# show switch
# Switch Name IP Address      State      Pipeline Mode
-|-----|-----|-----|-----|
1 filter-1     172.24.210.89  connected  13-14-match-push-vlan
```

4.4.4 Installing Arista 7800R3 Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode

DANZ Monitoring Fabric (DMF) 8.6.0 is the first release to support the Arista 7800R3 modular chassis series of switches, which use the Arista Networks EOS operating system.

A chassis can have one or more line cards. From a Controller's perspective, a chassis-based switch with multiple line cards, each with its own application-specific integrated circuit (ASIC), is treated as a single switch. When connected, a chassis works like any other switch and requires no user intervention for this support to work. The Controller automatically recognizes the chassis, initiates a handshake, and reacts to any chassis events like line card addition and removal.

Show commands display the modules in each chassis slot, whether line cards or supervisors, display line card properties, and what redundancy mode is active.



Note: One important aspect of treating a chassis like a non-modular switch is that the Controller reconciles the line card ASICs' differences by picking the least capable ASIC to be used as a base for the entire chassis. This approach provides the functional requirements the Controller needs to make decisions based on the (least common denominator) ASIC type and capabilities: therefore, picking the least capable ASIC ensures that all line cards can handle any Controller decision.



Note: DMF Release 8.6.0 supports the following modular switches:

- DCS-7804-CH
- DCS-7808-CH
- DCS-7812-CH
- DCS-7816-CH



Note: These are the current limitations in the DMF support for chassis-based switches:

- The maximum number of flows supported by the chassis is 8188 (in total for the entire system).
- The SSO redundancy protocol is not supported.

For all the 7800R3 series of switches, configure the Controller with deployment mode pre-configured (L3 ZTN). This is the default mode of operation starting in the DMF 8.5.0 release.



Note: L3 ZTN means the Controller is set up for deployment mode pre-configured.

This is a one-time setup needed to load the DMF-compatible EOS image. When set up, the next Controller upgrade will also automatically upgrade the switches.

Perform these steps on a 7800R3 Series switch to boot from the DMF Controller:

1. On the DMF Controller, issue the command `show switch-image url` to obtain the URL for the boot image. The URL that is needed is the one that corresponds to file `update-aristaeos`. Make a note of this URL, which will be required later when copying the image to the switch. Do **not** use the other URLs.

```
DMF(config)# sh switch-image url
# File                               Url                               Alternative Url
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
1 arista-ztp-install-script http://10.240.189.244/switchlight/arista-ztp-install-script
2 arista-ztp-install-script https://10.240.189.244/switchlight/arista-ztp-install-script
```

```

3 install-amd64 http://10.240.189.244/switchlight/install-amd64
4 install-amd64 https://10.240.189.244/switchlight/install-amd64
5 update-amd64 http://10.240.189.244/switchlight/amd64
6 update-amd64 https://10.240.189.244/switchlight/amd64
7 update-aristaeos http://10.240.189.244/eos/x86_64
8 update-aristaeos https://10.240.189.244/eos/x86_64
9 update-aristaeos-strata http://10.240.189.244/eos/eft
10 update-aristaeos-strata https://10.240.189.244/eos/eft

```

2. Attach a console connection to the Arista 7800R3 Series switch supervisor and turn on the switch or reboot the supervisor.



Note: If there is a redundant supervisor card in the chassis, you must also repeat steps 2-11 on the redundant supervisor.

3. Interrupt the boot process with **Control-C** to drop into the Aboot shell.

```

Aboot 7.1.11-33309652
Loading AUF data

Press Control-C now to enter Aboot shell
^CWelcome to Aboot.

```

4. Configure an IP address for the switch's *ma1* management interface. Configure *ma1* statically or use DHCP.
5. To use DHCP, type `udhcpc -i ma1`.

```

Aboot# udhcpc -i ma1
udhcpc: started, v1.30.1
udhcpc: sending discover
udhcpc: sending select for 172.24.208.123
udhcpc: lease of 172.24.208.123 obtained, lease time 86400

```

6. To configure *ma1* statically.

```

Aboot# ifconfig ma1 10.240.179.44 netmask 255.255.254.0
Aboot# route add default gateway 10.240.178.1 ma1
Tip: Please use the "route" command to verify the correct routing table.

```

7. Change the directory to `/mnt/flash/` on the switch.

```

Aboot# cd /mnt/flash
Aboot#
Aboot# ls -l
-rw-rw-r-- 1 root 88 2541 Apr 29 19:29 AsuFastPktTransmit.log
drwxrwxr-x 2 root 88 4096 Oct 31 01:06 Fossil
-rw-rw-r-- 1 root 88 1562 Apr 29 19:29 SsuRestore.log
-rw-rw-r-- 1 root 88 1562 Apr 29 19:29 SsuRestoreLegacy.log
-rw-rwx--- 1 root 88 47 Apr 29 19:37 boot-config
-rw-rw-r-- 1 root 88 4 Apr 29 18:19 config_match
drwxrwx--- 3 root 88 4096 Apr 29 19:38 debug
drwxrwxr-x 2 root 88 4096 Oct 31 01:06 fastpkttx.backup
drwxrwx--- 2 root 88 16384 Oct 31 01:04 lost+found
drwxrwxr-x 3 root 88 4096 Apr 29 19:36 persist
drwxrwxr-x 3 root 88 4096 Oct 31 01:20 schedule
-rw-rw-r-- 1 root 88 0 Oct 31 01:20 startup-config
-rw-rw-r-- 1 root 88 0 Apr 29 19:30 zerotouch-config

```

8. Use `wget` to copy the Arista EOS image from the DMF Controller. Use the `update-aristaeos` URL from **Step 1**. The command is `wget url`.

```

Aboot# wget http://10.240.189.244/eos/x86_64
Connecting to 10.240.189.244 (10.240.189.244:80)
x86_64 100% |*****| 1324M 0:00:00 ETA

```

9. Edit the `/mnt/flash/boot-config` file. Boot using the newly downloaded EOS image from the DMF Controller.

```
SWI=flash:/x86_64
```

10. Verify the boot-config was saved.

```
About# cat /mnt/flash/boot-config
SWI=flash:/x86_64
```

11. Reboot the system. Type: **reboot**.

```
About# reboot
[ 760.531244] SysRq : Remount R/O
Requesting system reboot
Restarting system
[ 3.988971] Running e2fsck on: /mnt/flash
[ 4.050509] e2fsck on /mnt/flash took 0s
[ 4.149265] Running e2fsck on: /mnt/crash
[ 4.201869] e2fsck on /mnt/crash took 0s

About 7.1.11-33309652
Loading AUF data

Press Control-C now to enter About shell
Booting flash:/x86_64
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
[ 9.097555] sd 0:0:0:0: [sda] Synchronizing SCSI cache
[ 9.161536] Starting new kernel
[ 3.078215][ T438] Running e2fsck on: /mnt/flash
[ 3.145683][ T443] e2fsck on /mnt/flash took 0s
[ 3.327339][ T511] Running e2fsck on: /mnt/crash
[ 3.388905][ T516] e2fsck on /mnt/crash took 0s
Discovering SID and required optimization
  SID OtterLake
  optimization not required
Data in /mnt/flash/x86_64 differs from previous boot image on /mnt/flash.
Saving new boot image to /mnt/flash...
Mounting SWIM Filesystem
  Optimization Default root squash found
  Optimization Default all squashes found
  Mounting optimization Default
Switching rootfs[ 14.363899][ T1248] tpm tpm0: [Firmware Bug]: TPM interrupt not working,
polling instead

Welcome to Arista Networks EOS 4.31.3M
Architecture: x86_64
[ 24.727098] sh[2413]: Starting EOS initialization stage 1
Starting NorCal initialization:
Configuration version in file is 10B008C.A00A3477 and in hardware is 1090063.A00A3475,
updating...

Download status = 'Ready for download', background status = 'Idle'
Download status = 'Ready for download', background status = 'Done'
Download status = 'Ready for download', background status = 'Done'
Download status = 'Ready for download', background status = 'Done'
Download status = 'Ready for download', background status = 'Done'
Download status = 'Download complete, config. activation needed', background status = 'Idle'
File /usr/share/Microsemi/OtterLake.pmc downloaded.
Firmware version in file is 10B008C and in hardware is 1090063, updating...
Download status = 'Download complete, config. activation needed', background status = 'Idle'
Download status = 'Download complete, config. activation needed', background status = 'Done'
Download status = 'Download complete, config. activation needed', background status = 'Done'
```

```

Download status = 'Download complete, config. activation needed', background status = 'Done'
Download status = 'Download complete, config. activation needed', background status = 'Done'
Download status = 'Download complete, firmware activation needed', background status = 'Idle'
File /usr/share/Microsemi/firmware-1.B.8C.pmc downloaded.

Microsemi EEPROM Upgraded

Restarting system
[18:57:49] watchdog punch .
[18:57:49] watchdog punch .
[18:57:50] watchdog punch .
[18:57:50] watchdog punch .
[18:57:50] watchdog punch .
[18:57:51] watchdog punch .
[18:57:51] watchdog punch .
[18:57:52] watchdog punch .
'[ 3.983263] Running e2fsck on: /mnt/flash
[ 4.044660] e2fsck on /mnt/flash took 0s
[ 4.154619] Running e2fsck on: /mnt/crash
[ 4.208347] e2fsck on /mnt/crash took 0s

Aboot 7.1.11-33309652
Loading AUF data

Press Control-C now to enter Aboot shell
Booting flash:/x86_64
Secure Boot disabled, skipping check
SPI flash hardware write protection disabled
[ 9.126662] sd 0:0:0:0: [sda] Synchronizing SCSI cache
[ 9.190674] Starting new kernel
[ 3.066762][ T439] Running e2fsck on: /mnt/flash
[ 3.134327][ T444] e2fsck on /mnt/flash took 0s
[ 3.348235][ T511] Running e2fsck on: /mnt/crash
[ 3.408605][ T517] e2fsck on /mnt/crash took 0s
Discovering SID and required optimization
  SID OtterLake
  optimization not required
Mounting SWIM Filesystem
  Optimization Default root squash found
  Optimization Default all squashes found
  Mounting optimization Default
Switching rootfs[ 7.615707][ T1216] tpm tpm0: [Firmware Bug]: TPM interrupt not working,
polling instead

Welcome to Arista Networks EOS 4.31.3M
Architecture: x86_64
[ 18.256937] sh[2383]: Starting EOS initialization stage 1
Starting NorCal initialization: [ OK ]
[ 35.769095] sh[2826]: Starting EOS initialization stage 2
Completing EOS initialization (press ESC to skip): [ OK ]
Model: DCS-7800-SUP1A
Serial Number: FGN234503DR
System RAM: 65684540 kB
Flash Memory size: 219G

localhost login:

```

12. Log in as admin. From enable mode, turn off Zero Touch Provisioning, as shown below. Then the system may reload again. Type: **zerotouch cancel**.

```

localhost> enable
localhost# zerotouch cancel
Apr 29 20:00:12 localhost ZeroTouch: %ZTP-6-CANCEL: Cancelling Zero Touch Provisioning
Apr 29 20:00:12 localhost ZeroTouch: %ZTP-6-RELOAD: Rebooting the system

```

13. Log in to the switch as admin. Get the System MAC address of the switch from the `show version` command's output. Type: **show version**. In this example, the System MAC address is **e8ae.c5ff.faba**. Obtain the MAC address located on the ID label on the rear of the 7800R3 Series switch.

```

localhost login: admin

```

```

Output to this terminal is being recorded for diagnostic purposes.
Note that only output that is visible on the console is recorded.
localhost> show version
Arista DCS-7808-CH
Hardware version: 12.10
Serial number: HNN23166835
Hardware MAC address: e8ae.c5ff.faba
System MAC address: e8ae.c5ff.faba

Software image version: 4.31.3M
Architecture: x86_64
Internal build version: 4.31.3M-36737551.4313M
Internal build ID: 2e80f044-7775-4e3d-968e-fa619089a514
Image format version: 3.0
Image optimization: Default

Uptime: 3 days, 21 hours and 46 minutes
Total memory: 65684540 kB
Free memory: 57729216 kB

```

14. On the DMF Controller, assign a switch name and configure the System MAC address noted from the previous step. The MAC address must be entered in colon format, such as **e8:ae:c5:ff:fa:ba**.

```

DMF(config)# switch filter-1
DMF(config-switch)# mac e8:ae:c5:ff:fa:ba

```

15. On the switch console, configure the IP address of the management interface.

```

localhost(config)# interface management 0
localhost(config-if-Ma0)# ip address 10.240.179.44/23

```



Note: The Arista 7800R3 data center switches support configuring the management IP address on interface Management 1/1, Management 1/2, Management 2/1, or Management 2/2. However, for the 7800R3 series switch to connect with the DMF Controller, you *must* configure the management virtual IP address on interface **Management0** of the active supervisor card only.



Note: For the Controller's communication with the switch to work, you should use interface Management 1/1 on the active supervisor to connect to the out-of-band management network and, if a standby supervisor is present, you should also use interface Management 2/1 on the standby supervisor to connect to the management network. You cannot use Management 1/2 and 2/2 for communication with the Controller.

16. Configure a route to the gateway using the following command.

```

localhost(config)# ip route 0.0.0.0/0 10.240.178.1

```

17. On the switch console, configure the IP address(es) of the DMF Controllers. For dual Controllers, the syntax is: **controller address controller#1 controller#2**. ZTN configuration download begins after configuring this part.

```

localhost(config-if-Ma1)# management dmf
localhost(config-mgmt-dmf)# controller address 10.240.189.243 10.240.189.244
localhost(config-mgmt-dmf)# no disabled

```

18. Verify the switch is connected to the DMF Controller. From the switch console, type: **show management dmf indigo**.

```

filter-1(config)# show management dmf indigo
DMF: enabled
Indigo agent: active
TCAM profile programming status: success
Hardware counters status: success
Forwarding chip discovery: success

Controllers:

```

ID	IP Address	Connection State	Connection Role
0	10.240.189.243	connected	standby
1	10.240.189.244	connected	active

19. The Arista switch appears as **connected** under the State column when viewed from the DMF Controller using the **show switch** command. For example:

```
DMF(config-switch)# show switch
# Switch Name IP Address      State      Pipeline Mode
-----|-----|-----|-----|
1 filter-1    10.240.179.44 connected  l3-14-match-push-vlan
```

4.4.5 Configuring the Switch Static IP and Controller IP in Interactive ZTF Mode

To configure or change the static IP or Controller IP addressing from the **zcs** CLI, complete the following steps:

1. From the Switch Light OS prompt on the switch to be configured, enter **Control-c** to drop into the loader mode.

```
Press Control-C now to enter the interactive loader shell.
^C
Welcome to the shell.
Type 'help' for command help.
loader#
```

2. Enter **zcs** to drop to **ztn-config** mode.

```
loader# zcs
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
(ztn-config)
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
```

3. Enter the **setup** command.

```
(ztn-config) setup
You are now running the interactive setup.
Press Enter to continue...
```

4. When prompted, type **static** and enter the IP address for the switch.

 **Note:** Additional settings for the DNS server and DNS domain options are available starting with **DMF Release 6.3.1**.

```
Please choose an IP option:
(DHCP/Static)? static
Please provide static IP settings:
IP Address: 10.9.36.29
Netmask: 255.255.255.0
Gateway: 10.9.36.1
Do you want to configure DNS settings?
(Yes/No)? yes
DNS Server: 10.3.0.4
DNS Domain: 10.1.5.200
Please provide the IP address of the Controller:
Controller IP: 10.2.0.66
```

5. If there is a second Controller, type **yes** when prompted and enter the IP address of the secondary Controller.

```
Do you have a second controller?
(Yes/No)? yes
Please provide the IP address of the second controller:
Second Controller IP: 10.8.25.32
```

```

IP Option : Static
IP Address : 10.8.39.203
Netmask : 255.255.192.0
Gateway : 10.8.0.1
DNS Server : 10.3.0.4
DNS Domain : qa.arista.com
Controller IP : 10.8.25.31
Second Controller IP: 10.8.25.32
Please confirm that the above settings are correct:
(Yes/Reset)? yes
Interactive setup completed successfully.
(ztn-config) reboot
Proceed with reboot [confirm]
Terminated
Requesting systRestarting system.

```

4.4.6 Installing Arista 7050X and 7260X Series using DHCP with bootfile-name option

Starting with **DANZ Monitoring Fabric (DMF) 8.2**, installing Switch Light OS on Arista switches can be automated using an Arista ZTP boot script available on the DMF Controller. The Arista switch models that support this procedure are the 7050CX3, 7050SX3, and 7260CX3.

The Arista ZTP boot script is served to the Arista switch using DHCP's bootfile-name option (option #67). The Arista switch downloads and executes this Arista ZTP boot script during its ZTP (Zero Touch Provisioning) phase following boot. The Arista ZTP boot script copies the Switch Light OS files from the DMF Controller and configures the appropriate boot settings on the Arista switch.

Procedure

1. Connect to the Arista switch to get the system MAC access. The system MAC address is the HWaddr for the management interface (**ma1**). The MAC in this example is **2C:DD:E9:7C:84:38**.

```

localhost# show interfaces management 1
Management1 is up, line protocol is up (connected)
Hardware is Ethernet, address is 2cdd.e97c.8438 (bia 2cdd.e97c.8438)
IPv6 link-local address is fe80::2edd:e9ff:fe7c:8438/64
Address being determined by SLAAC
No IPv6 global unicast address is assigned
IP MTU 1500 bytes (default) , BW 1000000 kbit
Full-duplex, 1Gb/s, auto negotiation: on, uni-link: n/a
Up 3 hours, 50 minutes, 21 seconds
Loopback Mode : None
4 link status changes since last clear
Last clearing of "show interface" counters 3:53:44 ago
5 minutes input rate 4.71 kbps (0.0% with framing overhead), 5 packets/sec
5 minutes output rate 172 bps (0.0% with framing overhead), 0 packets/sec
65939 packets input, 8202861 bytes
Received 10799 broadcasts, 51269 multicast
0 runts, 0 giants
0 input errors, 0 CRC, 0 alignment, 0 symbol, 0 input discards
0 PAUSE input
1155 packets output, 303450 bytes
Sent 584 broadcasts, 470 multicast
0 output errors, 0 collisions
0 late collision, 0 deferred, 0 output discards
0 PAUSE output localhost#

```

2. On the DMF Controller, configure the name of the Arista switch and its MAC address. In this example, the switch is assigned the name **DMF-F1**.

```

switch DMF-F1
mac 2c:dd:e9:7c:84:38

```

3. From the DMF Controller, obtain the Arista ZTP boot script URL. Run the command **show switch-image url**. The **arista-ztp-install-script** is the URL needed.

```

DMF-CTL2(config)# show switch-image url

```

File	Url
arista-ztp-install-script	http://10.240.129.29/switchlight/arista-ztp-install-script
install-amd64	http://10.240.129.29/switchlight/install-amd64
update-amd64	http://10.240.129.29/switchlight/amd64
update-aristaeos	http://10.240.129.29/eos/x86_64

```
DMF-CTL2(config)#
```

- On the DHCP server, include the bootfile-name option. Use the URL of the Arista ZTP script from the previous step. In this example, the `/etc/dhcp/dhcpd.conf` file is from an ISC DHCP server. Bring up the DMF switches using either A) **vendor-class-identifier** or B) **switch hardware address**.



Note: The edits required depend on the specific network environment and the type of switches.

- For a large DMF deployment, use the vendor-class-identifier as shown below. **DMF switches should use a different subnet than UCN switches when using vendor-class-identifier.**

```
subnet 10.240.130.0 netmask 255.255.255.128 {
  range 10.240.130.61 10.240.130.64;
  option domain-name-servers 10.240.48.6;
  option subnet-mask 255.255.255.128;
  option routers 10.240.130.1;
  option broadcast-address 10.240.130.127;
  class "Arista"{ match if substring (option vendor-class-identifier, 0, 6) =
  "Arista";
  option bootfile-name = "http://10.240.129.29/switchlight/
  arista-ztp-install-script"; }
}
```

- If DMF switches are in the same subnet as UCN switches, use the host address on `dhcpd.config` to identify the DMF switches.

```
host 7050X3 {
  hardware ethernet 2c:dd:e9:7c:84:38;
  option bootfile-name = "http://10.240.129.29/switchlight/arista-ztp-
  install-script";
}
```

- Restart the `dhcp` server process after editing the `/etc/dhcp/dhcpd.conf` file.
- The Arista ZTP boot script will be downloaded and executed at the ZTP phase.
- On the DMF Controller, verify the switch is connected using the `show switch DMF-F1` command.

```
DMF-CTL2(config)# show switch DMF-F1
# Switch Name IP Address State Pipeline Mode
- |-----|-----|-----|-----|
1 DMF-F1 fe80::968e:d3ff:feaa:ad0e%9 connected full-match-push-vlan
DMF-CTL2(config)#
```

4.4.7 Installing Arista 7280R Series using DHCP with bootfile-name option

The installation of EOS on Arista SAND platforms (7280x) can be automated using an Arista ZTP boot script available on the DANZ Monitoring Fabric (DMF) Controller. This procedure applies to all Arista 7280R Series platforms that are supported in DMF.

The Arista ZTP boot script is served to the Arista switch using DHCP's bootfile-name option (option #67). The Arista switch downloads and executes this Arista ZTP boot script during its ZTP (Zero Touch Provisioning) phase following boot. The Arista ZTP boot script copies the EOS SWI from the DMF Controller and configures the appropriate boot settings on the Arista switch.



Note: The Controller must be set up for deployment-mode pre-configured.

Procedure

1. Connect to the Arista switch to get the system MAC address. The system MAC address is the `HWaddr` parameter of `interface mal` in `About`. The MAC address in this example is ***D4:AF:F7:54:19:5A***.

```
About# ifconfig mal
mal Link encap:Ethernet HWaddr D4:AF:F7:54:19:5A
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
Interrupt:37
```

2. On the DMF Controller, configure the name of the Arista switch and its MAC address. In this example, the switch is assigned the name ***DMF-F1***.

```
switch DMF-F1
mac d4:af:f7:54:19:5a
```

3. From the DMF Controller, obtain the Arista ZTP boot script URL. Run the command `show switch-image url`. The `arista-ztp-install-script` is the URL needed.

```
DMF-CTL2(config)# show switch-image url
File                               Url
-----|-----
arista-ztp-install-script          http://10.240.129.29/switchlight/arista-ztp-install-script
install-amd64                      http://10.240.129.29/switchlight/install-amd64
update-amd64                       http://10.240.129.29/switchlight/amd64
update-aristaeos                   http://10.240.129.29/eos/x86_64
DMF-CTL2(config)#
```

4. On the DHCP server, include the `bootfile-name` option. Use the URL of the Arista ZTP script from the previous step. In this example, the `/etc/dhcp/dhcpd.conf` file is from an ISC DHCP server. Bring up the DMF switches using either A) **vendor-class-identifier** or B) **switch hardware address**.



Note: The edits required depend on the specific network environment and the type of switches.

- a. For a large DMF deployment, use the `vendor-class-identifier` as shown below. DMF switches should use a different subnet than UCN switches when using `vendor-class-identifier`.

```
subnet 10.240.130.0 netmask 255.255.255.128 {
  range 10.240.130.61 10.240.130.64;
  option domain-name-servers 10.240.48.6;
  option subnet-mask 255.255.255.128;
  option routers 10.240.130.1;
  option broadcast-address 10.240.130.127;
  class "Arista"{ match if substring (option vendor-class-identifier, 0, 6) =
  "Arista";
  option bootfile-name = "http://10.240.129.29/switchlight/
  arista-ztp-install-script"; }
}
```

- b. If DMF switches are in the same subnet as the UCN switches, use the host address on `dhcpd.conf` to identify the DMF switches.

```
host 7050X3 {
  hardware ethernet 2c:dd:e9:7c:84:38;
  option bootfile-name = "http://10.240.129.29/switchlight/arista-ztp-
  install-script";
}
```

5. Restart the `dhcp` server process after editing the `/etc/dhcp/dhcpd.conf` file.
6. The Arista ZTP boot script will be downloaded and executed at the ZTP phase.

7. On the DMF Controller, verify the switch is connected using the `show switch DMF-F1` command.

```
DMF-CTL2(config)# show switch DMF-F1
# Switch Name IP Address State Pipeline Mode
- |-----|-----|-----|-----|
1 DMF-F1 10.240.130.62 connected full-match-push-vlan
DMF-CTL2(config)#
```

4.4.8 Using DHCP with Default URL for Switch Installation in Preconfigured Provisioning Mode

To install the Switch Light OS using a DHCP server, complete the following steps:

1. Log in to the DHCP server serving the L2 segment where the monitoring fabric switches are connected.
2. Edit the `/etc/dhcp/dhcpd.conf` file.



Note: The edits required depend on your network environment and the type of switches you are connecting to the fabric. The example below shows how the `dhcpd.conf` file might look like when using fabric switches of the same architectural type, in this case, the PowerPC architecture.

dhcpd.conf File (All Switches of the Same Type)

```
subnet 10.9.18.0 netmask 255.255.254.0 {
range 10.9.18.201 10.9.18.254;
option routers 10.9.18.1;
option domain-name-servers 10.3.0.4;
option domain-name "qa.arista.com";
option domain-search "qa.arista.com", ".com";
option default-url = "http://10.9.18.12/switchlight/install-powerpc";
filename "pxelinux.0";
next-server 10.8.0.3;
}
```

Note: The URL points to the PowerPC install image on the Controller. The example below shows a DHCP configuration file with PowerPC and AMD64-based switches.

dhcpd.conf File (Switches of Different Types)



```
class "onie-vendor-powerpc-class" {
match if substring(option vendor-class-identifier, 0, 19) = "onie_vendor:powerpc";
option default-url = "http://10.9.18.11/switchlight/install-powerpc";
}
class "onie-vendor-amd64-class" {
match if substring(option vendor-class-identifier, 0, 18) = "onie_vendor:x86_64";
option default-url = "http://10.9.18.11/switchlight/install-amd64";
}
subnet 10.9.18.0 netmask 255.255.254.0 {
pool {
allow members of "onie-vendor-powerpc-class";
range 10.9.18.120 10.9.18.150;
}
pool {
allow members of "onie-vendor-amd64-class";
range 10.9.18.151 10.9.18.200;
}
range 10.9.18.201 10.9.18.254;
option routers 10.9.18.1;
option domain-name-servers 10.3.0.4;
option domain-name "qa.arista.com";
option domain-search "qa.arista.com", "qa.arista.com";
filename "pxelinux.0";
next-server 10.8.0.3;
}
```



Note: The URLs point to the PowerPC and AMD64 install images on the Controller. In the examples above, two DHCP options are supported to deliver ZTN Controller addresses to the switch.

default-url

- If the `default-url` option is set, the address from the URL will be extracted and used for ZTN transactions.
- If the `default-url` option is used to support automatic installation via ONIE then the same setting can be used to indicate the initial Controller address against which L3 manifest transactions should be performed.

next-server

- This option is used to get the software images and configurations. In the case of DMF, this is optional as `default-url` provides for SWI and configurations.

3. Restart the `dhcp` server process on the DHCP server.
4. Restart the switch after the DHCP service restarts.

4.5 Registering a Switch After Initial Deployment

To add a switch to the fabric after initial deployment, register the name and MAC address of the switch with the active DANZ Monitoring Fabric (DMF) Controller. The switch downloads a compatible Switch Light OS image and configuration from the Controller and uses the registered switch name to refer to the switch in the CLI output and GUI displays.

4.5.1 Using the GUI to Register a Switch

Procedure

1. Select **Fabric > Switches** from the main menu.

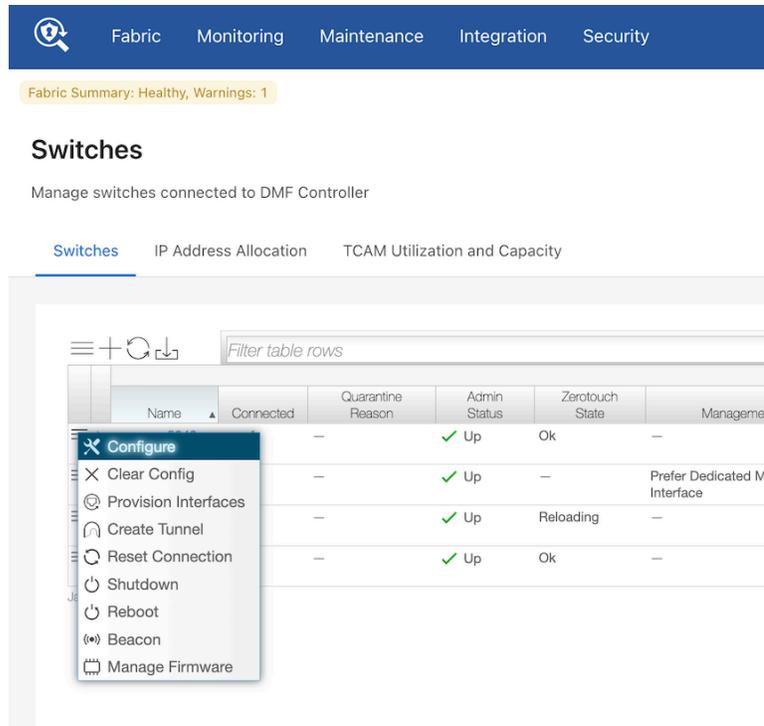
Figure 4-14: Fabric Switches Option

Name	Connected	Quarantine Reason	Admin Status	Ztn/Boot Date	Management Interface	Domain	Priority	Source Pkts Address	Source Pkts Address	Connected Time	Connection Time	Management IP Address	Allocated IP Address	Port	Interface Count	Flow Count
swname-5058	✓	---	Up	OK	---	---	---	---	---	Jan 16, 2024 4:11:25pm UTC	17h 31m	160.340.309.144.8700 10.240.186.102	---	37584	54	---
swname-7723	✗	---	Up	---	Prefer Dedicated Management Interface	---	---	---	---	---	---	---	---	---	---	---
swname-a508	✓	---	Up	Rebooting	---	---	---	---	---	Jan 16, 2024 4:12:02pm UTC	17h 30m	160.3617.46167.1604 10.240.186.103	---	55040	87	---
swname-df100	✓	---	Up	OK	---	---	---	---	---	Jan 16, 2024 4:12:02pm UTC	17h 30m	160.4476.219.1601.2040 10.240.186.104	---	36932	34	---

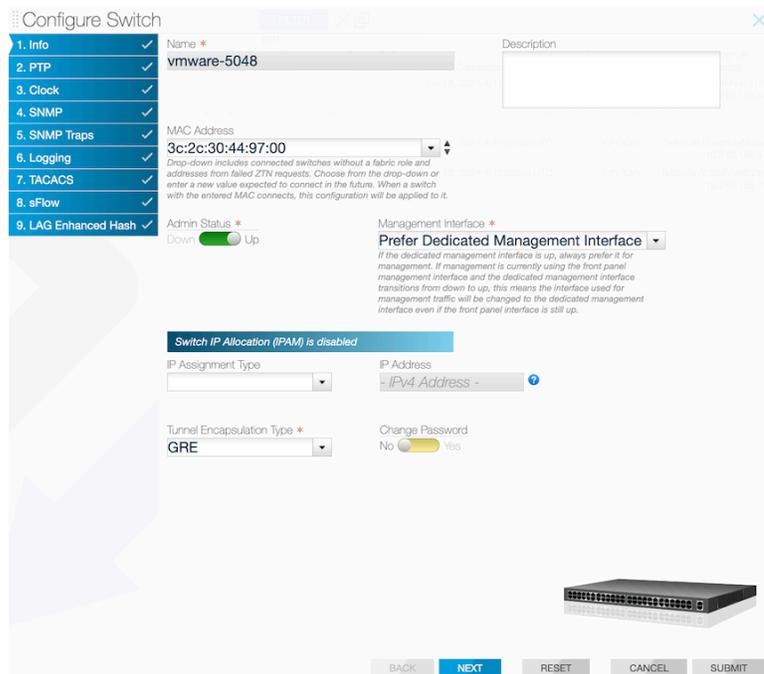
This page lists the switches connected to the DANZ Monitoring Fabric (DMF) Controller, with the current alias of the switch providing a link to the Switch View for the specified switch.

- To add an alias or change the existing alias, click the **Menu** control next to the switch name and select **Configure** from the pull-down menu that appears.

Figure 4-15: Configure Switch (Page 1)



This dialog provides the means to assign an alias and a MAC address, shut down or re-enable the switch, and change the password for direct remote connections to the switch.



3. Type the alias for the switch in the **Name** field.
4. Enter the MAC address and click **Submit**.
5. This dialog also provides access to a series of dialogues used to override the default configuration pushed from the DMF Controller to the switch. To advance to another page, click the numbered link for the page or click **Next**.

Figure 4-16: Continue Configuration

4.5.2 Using the CLI to Register a Switch

Enter the **switch *switch-name*** command to enter the **config-switch** submode, to associate the switch name with the MAC address of a physical switch. Replace ***switch-name*** with a unique alphanumeric text string. For example, the following commands assign the switch names ***core-sw-1***, ***filter-sw-1***, and ***delivery-sw-1*** to three switches:

```
controller-1 (config) # switch DMF-CORE-SWITCH-1
controller-1 (config-switch) # mac 00:00:00:00:00:09
controller-1 (config-switch) # switch DMF-FILTER-SWITCH-1
controller-1 (config-switch) # mac 00:00:00:00:00:0b
controller-1 (config-switch) # switch DMF-DELIVERY-SWITCH-1
controller-1 (config-switch) # mac 00:00:00:00:00:0e
```

To view the switches in the DANZ Monitoring Fabric (DMF), enter the **show switch** command from any mode, as in the following example:

```
controller-1> show switch
# Switch Name      IP Address          State      Pipeline Mode
- |-----|-----|-----|-----
1 bigtap-switch-1  fe80::ce37:abff:fe60:d474%2  connected  bigtap-l314-push-vlan
2 bigtap-switch-2  fe80::ce37:abff:fe60:cf8a%2  connected  bigtap-l314-push-vlan
3 bigtap-switch-3  fe80::ce37:abff:fea0:9071%2  connected  bigtap-l314-push-vlan
```

The output shows the switch alias, IP address, state, and pipeline mode.

To associate a new name with an existing switch MAC address, remove the switch registration with the **no switch** command.

```
controller-1 (config) # no switch DMF-CORE-SWITCH-1
```

To view additional details about a switch, enter the `show switch all detail` command, as in the following example:

```

controller-1> show switch all detail
# Switch Name      Mac address      Switch DPID      State      Quarantine reason  IP Address  TCP Port  Connected Since      Pipeline Mode
-----
1 DMF-CORE-SWITCH-1 cc:37:ab:a0:90:71 (Edgecore) 00:00:cc:37:ab:a0:90:71 connected 10.9.34.21 34808 2018-02-21 20:49:22.219000 PST bigtap-offset-match-push-vlan
2 DMF-DELIVERY-SWITCH-1 cc:37:ab:60:d4:74 (Edgecore) 00:00:cc:37:ab:60:d4:74 connected 10.9.34.22 43840 2018-02-21 20:58:25.046000 PST bigtap-offset-match-push-vlan
3 DMF-FILTER-SWITCH-1 8c:ea:1b:26:73:6f (Edgecore) 00:00:8c:ea:1b:26:73:6f connected 10.9.34.20 49049 2018-02-21 21:01:54.596000 PST bigtap-offset-match-push-vlan
controller-1>

```

After removing the switch registration, perform a new switch registration using the new switch name.

4.6 Changing the ZTF Mode After Deployment

4.6.1 Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode

ZTF cannot be used to install the switches when the switch management network connects the DANZ Monitoring Fabric (DMF) Controllers through a Layer 3 network. However, when a switch is in a different subnet than the Controller, manually configure the switches or use a DHCP server to download the Switch Light OS image to each fabric switch. To do this, change the switch provisioning mode to *Pre-Configured*.

If the switches and Controllers are in the same L2 broadcast domain, use the auto-discovery switch deployment mode for an L2-ZTF deployment. If the switches and Controllers are not in the same L2 broadcast domain, use the pre-configured provisioning mode to enable an L3-ZTF deployment. The entire fabric must be in a single provisioning mode; DMF only supports the auto-discovery provisioning mode if all the switches are in the same Layer 2 domain.

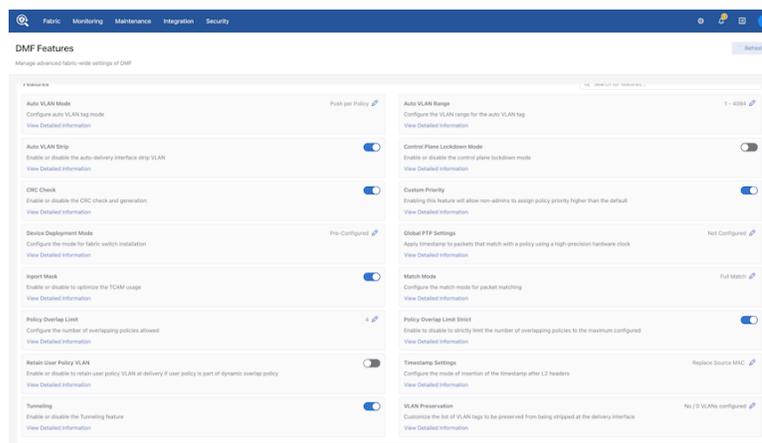
4.6.2 Using the GUI to Change the Switch Provisioning Mode

Complete the following steps to use the DANZ Monitoring Fabric (DMF) GUI to change the switch provisioning mode.

Procedure

1. Click the **DMF logo** in the GUI Main menu to display the Controller landing page.
2. Click the **Settings** control in the **Features** section on the Controller landing page.

Figure 4-17: Changing the Switch Provisioning Mode

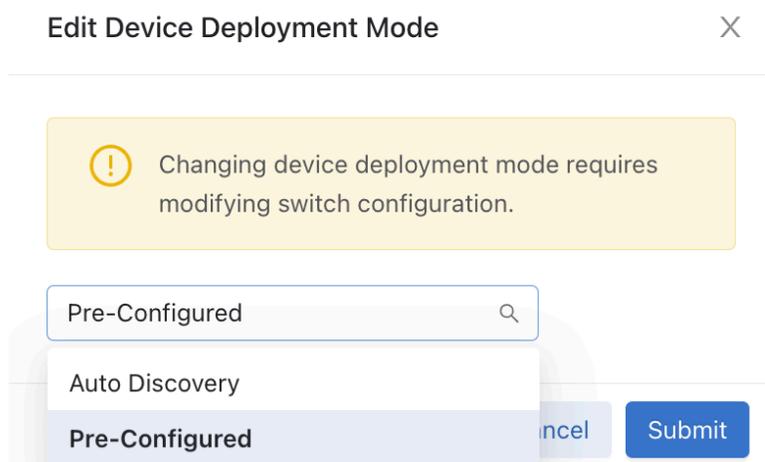




Note: For information about the other options in this section, refer to the DMF User Guide.

- Click the Settings controller to the right of the **Device Provisioning Mode** option and click **Submit**. Control the configuration of this feature using the **Edit** icon by clicking on the **pencil icon**.

Figure 4-18: Edit Device Deployment Mode



- Click on the downward arrow and choose from the drop-down options. There are two ways to modify the switch configuration. Select the **Pre-Configured** option, click on the **Submit** button, and confirm the operation when prompted. An error message is displayed if switches are already deployed using Auto Discovery mode.

4.6.3 Using the CLI to Change the Switch Provisioning Mode

Complete the following steps to use the CLI to change from Auto-discovery (L2-ZTF) Mode to preconfigured (L3-ZTF) Mode.

Procedure

- Disable IPAM by entering the following commands.



Note: IPAM is supported only in L2-ZTF mode. Disable it before moving to L3-ZTF mode. However, before disabling IPAM, remove the IPAM configuration in the **config-ipam-switch** submode.

```
controller-1(config-ipam-switch)# no dns-server
controller-1(config-ipam-switch)# no gateway
controller-1(config-ipam-switch)# no ip-range 10.8.39.81 10.8.39.90 subnet-mask-length 18
controller-1(config-ipam-switch)# (config-ipam-switch)# exit
controller-1(config-ipam-switch) (config)# no ipam switch
```

- Change the switch provisioning mode by entering the following command on the active DANZ Monitoring Fabric (DMF) Controller.

```
controller-1(config)# deployment-mode pre-configured
```

- Configure the switches using DHCP or static IP addresses.



Note: If not using DHCP, assign a static IP using the switch CLI (PCLI) on each fabric switch.

4. Configure the active and standby DMF Controller IP address on each fabric switch.
5. Reboot each switch.

If switches and Controllers are in the same L2 broadcast domain, use the auto-discovery switch deployment mode for an L2-ZTF deployment. If the switches and Controllers are not in the same L2 broadcast domain, use the pre-configured provisioning mode to enable an L3-ZTF deployment. The entire fabric must be in a single provisioning mode; DMF only supports the auto-discovery provisioning mode if all the switches are in the same Layer 2 domain.

4.6.4 Changing to Layer 3 ZTF (Preconfigured) Mode

Complete the following steps to change to Preconfigured (L3-ZTF) Mode from Auto Discovery (L2-ZTF) mode.

Procedure

1. Change the configuration of each switch management (**ma1**) IP to DHCP or assign static IP.

```
(ztn-config) interface ma1 ip-address
Set the management interface address parameters. Possibilities are:
interface ma1 ip-address dhcp
interface ma1 ip-address <ip-address>/<prefix> gateway <gateway-address>
A) Setting ma1 interface of switch to use DHCP.
(ztn-config) interface ma1 ip-address dhcp
(ztn-config) show
IP Option: DHCP
Controllers: fe80::250:56ff:fea2:df9b
(ztn-config)
B) Setting ma1 interface of switch for static IP
(ztn-config) interface ma1 ip-address 192.168.10.10/25 gateway 192.168.10.1
(ztn-config) show
IP Option: Static
IP Address: 192.168.10.10
Netmask: 255.255.255.128
Gateway: 192.168.10.1
DNS Server: None
DNS Domain: None
Controllers: fe80::250:56ff:fea2:df9b
(ztn-config)
```

2. At the **ztn-config** prompt, clear the Controller configuration and add the Controller IP (active and standby).

```
A) Clearing controller config on switch.
(ztn-config) controller clear
(ztn-config) show
IP Option: Static
IP Address: 10.240.130.13
Netmask: 255.255.255.128
Gateway: 10.240.130.1
DNS Server: None
DNS Domain: None
Controllers:
(ztn-config)
B) Adding controller IP.
(ztn-config) controller set 10.240.130.15,10.240.130.16
(ztn-config)
(ztn-config) show
IP Option: Static
IP Address: 10.240.130.13
Netmask: 255.255.255.128
Gateway: 10.240.130.1
DNS Server: None
DNS Domain: None
Controllers: 10.240.130.16,10.240.130.15
```

```
(ztn-config)
```

3. At the **ztn-config** prompt, reboot each switch.

Enter the **deployment-mode pre-configured** command on the DANZ Monitoring Fabric (DMF) Controller to enable Layer 3 mode switch installation, whether using the manual method or DHCP with the default URL method of installation.



Note: In Layer 3 mode, ZTF uses TCP port 8843 for communication between the Controller and switches. This port must be allowed on the Controller and on any devices connecting the Controller to the fabric switches.

4.6.5 Changing to Layer 2 ZTF (Auto-Discovery) Mode

Complete the following steps to change to Auto-Discovery (L2-ZTF) Mode from Preconfigured (L3-ZTF) Mode.

Procedure

1. Move the switches or Controllers into the same L2 broadcast domain if required.



Note: L2-ZTF requires all the switches and Controllers to be in the same broadcast domain.

2. Change the switch provisioning mode by entering the following command on the active DANZ Monitoring Fabric (DMF) Controller.

```
controller-1(config)# deployment-mode auto-discovery
```

3. Login to each switch. At the **ztn-config** prompt, clear the ZTF configuration.

In L2-ZTF mode, the Controllers are auto-discovered by switches.

```
(ztn-config) interface ma1 ip-address dhcp
(ztn-config) controller clear
(ztn-config) show
IP Option: DHCP
Controllers:
(ztn-config)
```

4. At the **ztn-config** prompt, reboot each switch.

4.6.6 System Reinstall for an EOS Switch

Perform a system reinstall by removing the local startup-config/zerotouch-config on the switch so the DANZ Monitoring Fabric (DMF) Controller no longer manages it.

Rebooting the switch restarts the Arista-native ZTP process and requests a fresh image from the Controller.

Use the following command to perform a system reinstall:

```
C1# system reinstall switch eos-switch-name reboot
```



Note: There are other optional parameters (such as **timeout** and **factory-default**), but they do not apply to EOS switches.

The following is an example where the switch name is **core1**.

```
C1(config)# system reinstall switch core1 reboot
system switch reinstall: "deployment-mode pre-configured"
system switch reinstall: l3-ztn currently configured
system switch reinstall: l3-ztn implies switches are remote
system switch reinstall: l3-ztn and some switches may not rejoin
reinstall may cause service interruption
system switch reinstall ("y" or "yes" to continue): y
```

An optional parameter called **reboot** forces the switch to reboot and begin the re-installation process.

CLI Show Commands

When the switch is rebooting, ZTN cannot communicate with the switch, so a **Zerotouch state error hint** and **Zerotouch state error msg** appear when using the following show command:

```
(config)# show switch core1 zerotouch
Name          : core1
Ip address    : 10.243.254.25
Last update   : 2023-06-02 07:18:35.749000 UTC
Zerotouch state: reloading
Zerotouch state error hint : Rest API Client problem
Zerotouch state error msg  : Connect to 10.243.254.25:80 [/10.243.254.25] failed: Connection
refused (Connection refused)
```

The error message changes after the switch has fully booted.

```
SM-InspiringPare-Broadwater-C1(config-crypto)# show switch core1 zerotouch
Name          : core1
Ip address    : 10.243.254.25
Last update   : 2023-06-02 07:29:34.850000 UTC
Zerotouch state: reloading
Zerotouch state error hint : Rest API Client problem
Zerotouch state error msg  : No route to host (Host unreachable)
```

At this point, the switch has booted up entirely. Still, the Controller cannot talk to the switch, as the necessary configuration is absent. Kick-start the DMF ZTN process on the switch again using the commands below:

```
(config)# management dmf
(config-mgmt-dmf)# controller address ip-address
(config-mgmt-dmf)# no disabled
```

Troubleshooting

Check the status using the command **show switch switch-name zerotouch**.

After performing the steps above for reconnecting an EOS switch, and if the state remains stuck in reloading (and there is a **Zerotouch state error hint** / **Zerotouch state error msg** output), please contact [Arista Support](#).

4.7 SKU Reporting for EOS Switches

Like SwitchLight (SWL) OS switches, EOS switches now report their SKUs to the DANZ Monitoring Fabric (DMF) Controller.

View the EOS switch SKU using the DMF Controller CLI or GUI.

4.7.1 Using the CLI to Configure SKU Reporting for EOS Switches

Run the `show fabric inventory` command from the `login` mode to view the switch SKUs from the DMF Controller CLI.

The `Switch Inventory` table enumerates the switches associated with the DMF environment. The `SKU` column indicates the SKU of each switch.

```

CONTROLLER-1> show fabric inventory
~~~~~ Controller Inventory ~~~~~
# Node Id Hostname          SKU          Serial Number
-|-----|-----|-----|-----|
1 29617  CONTROLLER-1      DCA-DM-C450  FF99R52
2 23262  CONTROLLER-2      DCA-DM-C450  3W9D3Y2

~~~~~ Switch Inventory ~~~~~
# Switch          SKU          Serial Number Manufacturer  Asic
-|-----|-----|-----|-----|-----|
1 dmf-arista-7280SR2-2 DCS-7280SR2-48YC6 JPE22123192  Arista Networks  jericho-plus
2 dmf-arista-7280CR3-1 DCS-7280CR3-32P4  JPE20383391  Arista Networks  jericho2
3 dmf-arista-7280SR3-1 DCS-7280SR3-48YC8 JPE22191168  Arista Networks  jericho2c
4 dmf-arista-7280CR3-2 DCS-7280CR3-32P4  JPE20383398  Arista Networks  jericho2
5 dmf-arista-7280SR-1  DCS-7280SR-48C6   SGD20370893  Arista Networks  jericho
6 dmf-arista-7280SR2-1 DCS-7280SR2-48YC6 JPE20476226  Arista Networks  jericho-plus

~~~~~ Recorder Node Inventory ~~~~~
# Recorder Node  SKU          Serial Number
-|-----|-----|-----|
1 DMF-RECORDER-NODE DCA-DM-RA3  FLC1RN3

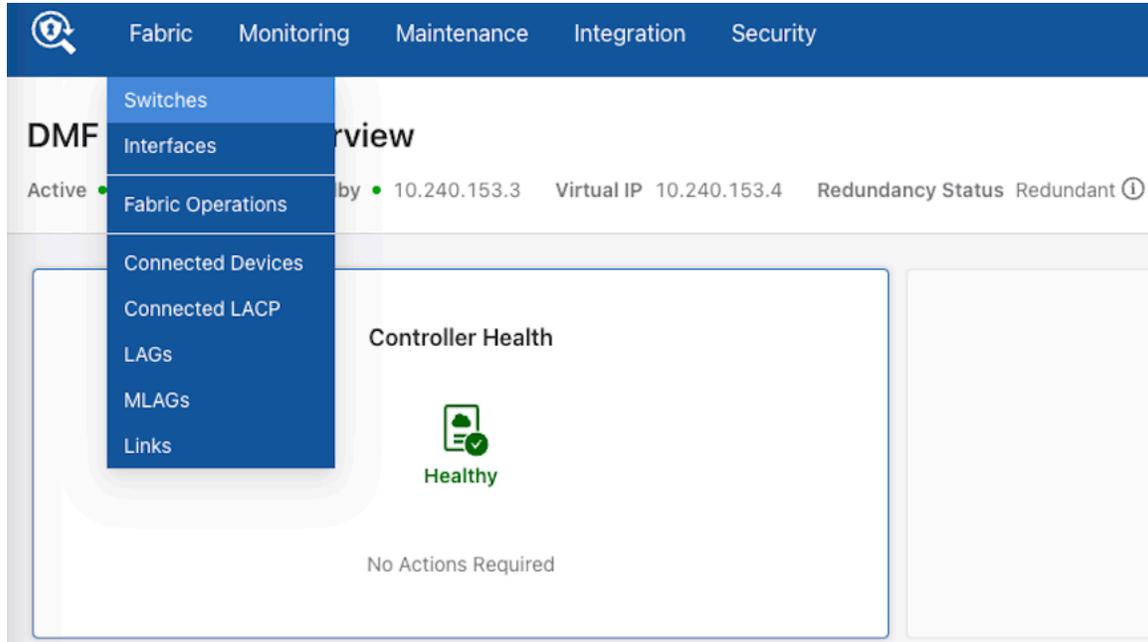
~~~~~ Service Node Inventory ~~~~~
# Service Node  SKU          Serial Number
-|-----|-----|-----|
1 DMF-SERVICE-NODE DCA-DM-SDL  GS11RN3

```

4.7.2 Using the GUI to Configure SKU Reporting for EOS Switches

To view the switch SKUs from the DMF Controller GUI, hover the mouse over the **Fabric** menu bar and select **Switches**.

Figure 4-19: Fabric > Switches



The **Switches** page loads.

Figure 4-20: Switches

The screenshot shows the 'Switches' page in the DMF Controller GUI. The page title is 'Switches' and it says 'Manage switches connected to DMF Controller'. There are tabs for 'Switches', 'IP Address Allocation', and 'TCAM Utilization and Capacity'. Below the tabs is a table of switches. The table has columns for Name, Connected, Quarantine Reason, Admin Status, Zabbix State, Management Interface, Domain, Priority, Source Pkts, and Connected Since. The table contains 6 rows of data, all with 'Connected' status and 'Up' admin status.

Name	Connected	Quarantine Reason	Admin Status	Zabbix State	Management Interface	Domain	Priority	Source Pkts	Connected Since	Connection Time	Management IP Address(es)	Allocated IP Address	Port	Interface Count	Flow Count
dmf-arlsta-7280SR-1	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:42pm UTC	4h 50m	---	56554	57	NA	
dmf-arlsta-7280SR3-1	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:36pm UTC	4h 50m	---	58998	59	NA	
dmf-arlsta-7280SR2-2	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:43pm UTC	4h 50m	---	43114	57	NA	
dmf-arlsta-7280SR2-1	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:44pm UTC	4h 50m	---	39180	57	NA	
dmf-arlsta-7280CR3-2	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:36pm UTC	4h 50m	---	59824	44	NA	
dmf-arlsta-7280CR3-1	✓	---	Up	Ok	---	---	---	---	Dec 20, 2023, 5:38:36pm UTC	4h 50m	---	58750	46	NA	

The SKUs do not appear by default but display after enabling the SKU column. To enable the column, click the menu button in the table. In the menu, select **Show/Hide Columns**.

Figure 4-21: Show/Hide Columns

The screenshot shows the DMF Controller interface. At the top, there is a navigation bar with icons for Fabric, Monitoring, Maintenance, Integration, and Security. Below this, a status bar indicates 'Fabric Summary: Healthy'. The main section is titled 'Switches' and includes the subtitle 'Manage switches connected to DMF Controller'. There are three sub-sections: 'Switches' (selected), 'IP Address Allocation', and 'TCAM Utilization and Capacity'. The 'Switches' section contains a table with the following data:

	Connected	Quarantine Reason	Admin Status	Zerotouch State
dmf-arista-7280CR3-1	✓	—	✓ Up	Ok
▶ dmf-arista-7280SR3-1	✓	—	✓ Up	Ok
▶ dmf-arista-7280SR2-2	✓	—	✓ Up	Ok
▶ dmf-arista-7280SR2-1	✓	—	✓ Up	Ok
▶ dmf-arista-7280CR3-2	✓	—	✓ Up	Ok
▶ dmf-arista-7280CR3-1	✓	—	✓ Up	Ok

A context menu is open over the table, showing options: 'Show/Hide Columns' (highlighted), 'Create Tunnel', and 'Refresh'. The 'Show/Hide Columns' option is selected, indicating that the 'SKU' column is being enabled.

Dec 20, 2023, 10:29:08pm UTC

In the dialog box, select the **SKU** checkbox and click **Save Preferences**.

Figure 4-22: Configure Table Columns

Configure Table Columns
✕

Check Columns to Show

All Columns [select](#) / [unselect](#)

Clock [select](#) / [unselect](#)

Hardware [select](#) / [unselect](#)

Logging [select](#) / [unselect](#)

sFlow [select](#) / [unselect](#)

SNMP [select](#) / [unselect](#)

SNMP Traps [select](#) / [unselect](#)

TACACS [select](#) / [unselect](#)

Tunneling [select](#) / [unselect](#)

Version Stats [select](#) / [unselect](#)

DPID

MAC

Name

Connected

Handshake State

Quarantine Reason

Admin Status

Zerotouch State

Management Interface

Manufacturer

SKU

Platform

Serial Number

Software Description

ASIC

Build

Version

Implementation

OS

Loader Version
Version of switch loader

Next Loader Version
Version of switch loader available after upgrade

CPLD
CPLD of switch

Next CPLD
CPLD of switch available after power-cycle

ONIE
ONIE version of switch

Next ONIE
ONIE version of switch available after upgrade

Password

PTP: Domain

PTP: Priority1

PTP: Source IPv4 Address

PTP: Source IPv6 Address

Clock: Time Zone

Clock: NTP Servers

SNMP: Contact

SNMP: Location

SNMP: Community

SNMP: Users

SNMP: Traps Enabled

SNMP: Trap Hosts

SNMP Traps: PSU Status

SNMP Traps: Fan Status

SNMP Traps: Link Status

SNMP Traps: Auth Failure

SNMP Traps: CPU Load Thresholds (1-Minute)

SNMP Traps: CPU Load Thresholds (5-Minute)

SNMP Traps: CPU Load Thresholds (15-Minute)

SNMP Traps: Percent Idle

SNMP Traps: Percent Utilization

SNMP Traps: Available Memory Thresholds

SNMP Traps: TCAM Utilization Thresholds

SNMP Traps: Thermal

SNMP Traps: Thermal Sensor Status

Logging: Controller Enabled

Logging: Remote Enabled

Logging: Remote Servers

TACACS: Default Timeout

TACACS: Servers

sFlow: Sample Rate (packets)

sFlow: Max Header Size

sFlow: Counter Interval (s)

LAG Enhanced Hash Configured

Connected Since

Connection Time

Management IP Address(es)

Allocated IP Address

Port

Interface Count

Flow Count

Tunneling: Encapsulation Type

Tunneling: GRE Tunnel Interface Count

Tunneling: VXLAN Tunnel Interface Count

RESTORE DEFAULTS
CLOSE
SAVE PREFERENCES

The **SKU** column appears in the **Switches** table and displays the SKU of each switch.

Figure 4-23: SKU Column

Fabric Summary: Healthy

Switches

Manage switches connected to DMF Controller

Switches IP Address Allocation TCAM Utilization and Capacity

Name	Connected	Zerotouch State	Manufacturer	SKU	Platform	Serial Number	Connected Since	Connection Time
dmf-arista-7280CR3-1	✓	Ok	Arista Networks	DCS-7280CR3-32P4	x86_64-dcs-7280cr3-32p4-eos	JPE20383391	Dec 20, 2023, 5:38:36pm UTC	4h 40m
dmf-arista-7280CR3-2	✓	Ok	Arista Networks	DCS-7280CR3-32P4	x86_64-dcs-7280cr3-32p4-eos	JPE20383398	Dec 20, 2023, 5:38:36pm UTC	4h 40m
dmf-arista-7280SR2-1	✓	Ok	Arista Networks	DCS-7280SR2-48YC6	x86_64-dcs-7280sr2-48yc6-eos	JPE20476226	Dec 20, 2023, 5:38:44pm UTC	4h 40m
dmf-arista-7280SR2-2	✓	Ok	Arista Networks	DCS-7280SR2-48YC6	x86_64-dcs-7280sr2-48yc6-eos	JPE22123192	Dec 20, 2023, 5:38:43pm UTC	4h 40m
dmf-arista-7280SR3-1	✓	Ok	Arista Networks	DCS-7280SR3-48YC8	x86_64-dcs-7280sr3-48yc8-eos	JPE22191168	Dec 20, 2023, 5:38:36pm UTC	4h 40m
dmf-arista-7280SR-1	✓	Ok	Arista Networks	DCS-7280SR-48C6	x86_64-dcs-7280sr-48c6-eos	SGD20370893	Dec 20, 2023, 5:38:42pm UTC	4h 40m

Dec 20, 2023, 10:18:54pm UTC Show: 10 25 100 All (1 - 6 / 6)

Managing Switches and Interfaces

This chapter describes how to manage switches and interfaces after installing the monitoring fabric switches.

5.1 Connecting Directly to a Switch

To install the Switch Light OS individually on a switch in a different Layer 2 domain or troubleshoot the switch, use telnet or SSH to connect to the switch.

To allow SSH to a switch, if using ZTF for installing switches from the DANZ Monitoring Fabric (DMF) Controller in the same Layer 2 domain, configure an IPAM IP address pool. Alternatively, use the `connect switch switch-name` command to connect to the switch CLI.

```
controller-1(config)# connect switch DMF-FILTER-SWITCH-1
Warning: Permanently added the RSA host key for IP address 'fe80::3617:ebff:fef2:cfc4%em1' to the
list of known hosts.
Last login: Mon Sep 18 02:32:35 2017 from fe80::46a8:42ff:fe35:29f7%ma1
SwitchLight ZTN Manual Configuration. Type help or ? to list commands.
```

After connecting to the switch, enter the `debug admin` command at the displayed ztn-config prompt.

```
(ztn-config) debug admin
DMF-FILTER-SWITCH-1>
```

This command provides access to the switch CLI prompt for directly managing or troubleshooting the switch. The following commands are available in the ZTN console:

- **controller**: adds, removes, sets, or clears the L3 ZTN Controller list
- **debug**: Special command to access the full switch CLI
- **help**: Displays CLI help
- **interface**: sets the *ma1* address parameters
- **reboot**: restarts the switch
- **setup**: performs interactive setup
- **show**: displays the current settings

5.1.1 Manually Configuring Enhanced Hashing for Load Distribution

In some scenarios, it may be desirable to manually select the bytes in each packet that are used for load distribution among the members in a LAG.

Enter the `hash-type enhanced` command from the `config-switch-lag-if` submode to manually configure enhanced hashing.

Use the `lag-enhanced-hash` command to enter the *config-switch-hash* submode and use the `hash` command to identify the values to use for load distribution.



Note: For a list of the switch platforms that support enhanced hashing (including symmetric hashing), refer to the *DANZ Monitoring Fabric Hardware Compatibility List*.

Changes in hash configuration do not affect the LAG configuration, so there is no need to reconfigure LAGs after changing the hash type.

5.1.2 Configuring Enhanced Hashing

To configure enhanced hashing, use the `lag-enhanced-hash` command to enter the *config-switch-hash* submode. Use the `hash` command to identify the hash type and the specific fields for load distribution.

```
controller-1(config-switch)# lag-enhanced-hash
controller-1(config-switch-hash)#
```

The hash command has the following syntax:

```
[no] hash gtp header-first-byte <GTP header first byte> header-first-byte-mask <GTP header first byte mask> | gtp port-match <UDP tunnel port match entry number> {dst-port <GTP tunnel UDP destination port> {and | or} src-port <GTP tunnel UDP source port> | src-port <GTP tunnel UDP source port>} | ipv4 {[dst-ip] [I4-dst-port] [I4-src-port] [protocol] [src-ip] [vlan-id]} | ipv6 {[dst-ip] [I4-dst-port] [I4-src-port] [nxt-hdr] [src-ip] [vlan-id]} | I2 [dst-mac] [eth-type] [src-mac] [vlan-id] I2gre {inner-I2 [dst-mac] [eth-type] [src-mac] [vlan-id] | inner-I3 [dst-ip] [I4-dst-port] [I4-src-port] [protocol] [src-ip] [vlan-id]} mpls {[label-1] [label-2] [label-3] [label-hi-bits] [payload-dst-ip] [payload-src-ip]} seeds { <First hash seed> [<Second hash seed>]} symmetric {enable | disable}
```

5.1.3 Symmetric Load Balancing

Enhanced hashing supports symmetric load balancing (enabled by default) for switch platforms that support this feature.



Note: DANZ Monitoring Fabric (DMF) supports symmetric hashing on specific switches for IP and Fibre Channel over Ethernet (FCoE) traffic. Symmetric hashing on MPLS traffic labels is not supported.

With symmetric load balancing, the link selected for distributing traffic in one direction is also used for traffic in the other direction. Arista Networks recommends enabling hashing on the source IP address and destination IP address for optimal symmetric behavior.



Note: In some scenarios, using Layer 4 protocol ports can improve load-balancing efficiency. However, these fields are not used by default because they cannot be used if packet fragmentation is likely to occur.

5.1.4 GTP Hashing

Generic Tunneling Protocol (GTP) hashing provides a more even distribution of GTP-encapsulated packets among the members of a port-group. When GTP hashing is enabled, DANZ Monitoring Fabric (DMF) includes the Tunnel endpoint identifier (TEID) value in the GTP packets in its hashing algorithm for outbound traffic.

This applies only to GTP user data tunneling packets (**udp port 2152**). GTP control traffic (**udp port 2123**) is not affected.

To enable hashing with Generic Tunneling Protocol (GTP), use the `hash gtp` command. This command sets enhanced hash parameters for distributing traffic on port-channel member ports for which enhanced hashing is enabled. The command syntax is as follows:

```
hash gtp port-match <port-match> {dst-port <dst-port> {and | or} src-port <src-port> | dst-port <dst-port> | src-port <src-port>}
```

The GTP command specifies the packet fields to identify GTP traffic. When enabled, DMF uses the TEID in the GTP header for hashing GTP traffic instead of using L4 ports—Configure `14-dst-port` and `14-src-port` in `hash ipv4` or `hash ipv6` for proper operation.

5.2 Overriding the Default Switch Configuration

After completing the switch installation, further switch configuration, including software upgrades, is managed from the DANZ Monitoring Fabric (DMF) Controller.



CAUTION: Make **all** configuration changes related to fabric switches using the Controller CLI or the Controller GUI, which provides DMF-level 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.

In general, the configuration options set on the DMF Controller are pushed to each connected switch, eliminating the need for box-by-box configuration. However, merging or overriding the default configuration pushed from the DMF Controller with switch-specific configuration for some parameters is possible. These parameters are as follows:

- Clock
- SNMP and SNMP Traps
- Logging
- TACACS

DMF supports two types of overriding mechanisms:

- `override-global`: Only the switch-specific configuration is applied.
- `merge-global`: The global config and switch-specific configuration are merged and then applied.

In the merge mode, the effective switch configuration is determined by the following rules:

- **Stand-alone values**: If the key only exists in one of the configurations; take it as-is in the resultant configuration, otherwise:
- If the key exists in both global and switch configurations: the value of the key from the switch-config takes precedence (over its value from the global-config).
- **Lists**: If the list only exists in one of the configurations; take that list as-is in the result configuration, otherwise:

- If it exists in both global and per-switch configuration, then merge with this rule.
- If the global and switch-specific configuration has an entry with the same key, the switch-specific list entry completely replaces the entry from the **global-config**, otherwise:
- All entries from the switch-specific configuration are appended to the **global-config** (with de-duplication). The configurations that occur as lists for the above overridable parameters are indicated below:
- ntp
 - server <- list
 - time-zone
- snmp-server
 - community <- list
 - contact
 - enable
 - host <- list
 - location
 - switch trap
 - user <- list
- logging
 - controller
 - remote
 - remote server <- list
- tacacs
 - server <- list

GUI Procedure

1. Select **Fabric > Switches** from the main menu.
The system displays the Switches page, which lists the switches connected to the DMF controller.

- To override any of the default switch configuration settings, click the **Menu** control next to the switch and select **Configure** from the pull-down menu that appears.

Figure 5-1: Configure Switch (Page 1)

This dialog provides access to a series of dialogs used to override the default configuration that is pushed from the DMF Controller to the switch.

- To advance to another page, click the numbered link for the page or click **Next**.
- After making any changes required, click **Submit**.

CLI Procedure

To override the default configuration for a specific switch, enter the **config-switch** submode for the specific switch and use the commands available, viewable by entering the **help** command or by using tab completion.

```

controller-1# config
controller-1(config) switch DMF-FILTER-SWITCH-1
controller-1(config-switch) <Tab>
admin lag-interface sflow switch-group
banner logging show
tacacs
description mac shut down
tunnel-interface
interface ntp snmp
lag-enhanced-hash role snmp-server
controller-1 (config-switch) #

```

Use the **shutdown** command to shut down a switch from the Controller, in which case all the interfaces of the switch are put in admin down mode and the switch is black-holed.

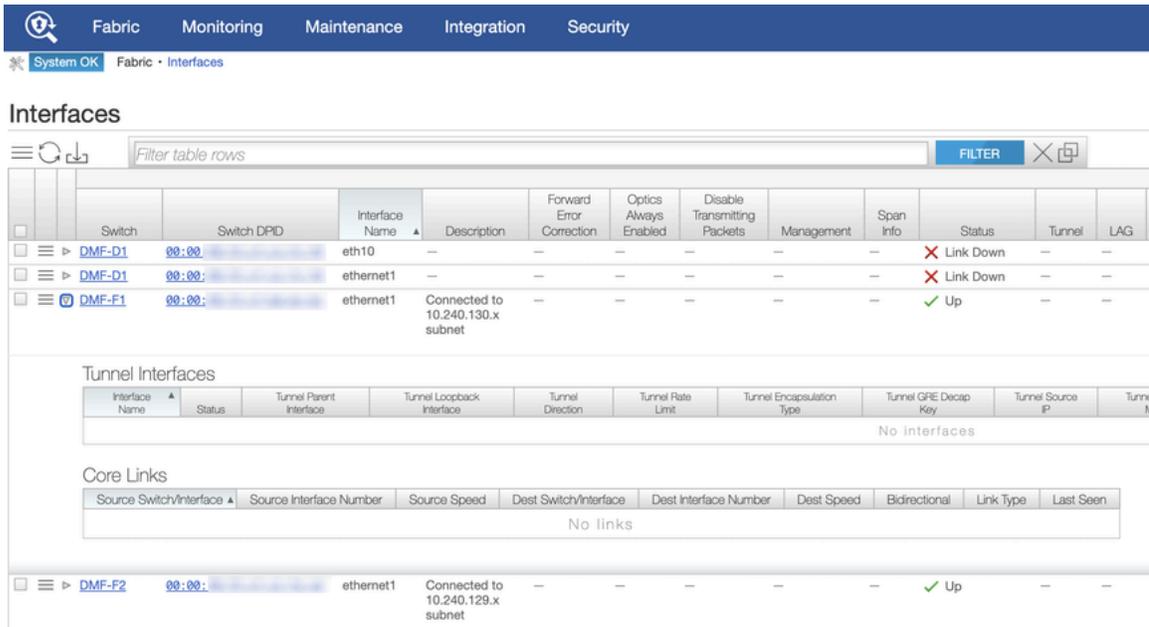
5.3 Configuring Switch Interfaces

To use the GUI to configure switch interfaces, complete the following steps.

Procedure

1. Select **Fabric > Interfaces** from the GUI Main menu.

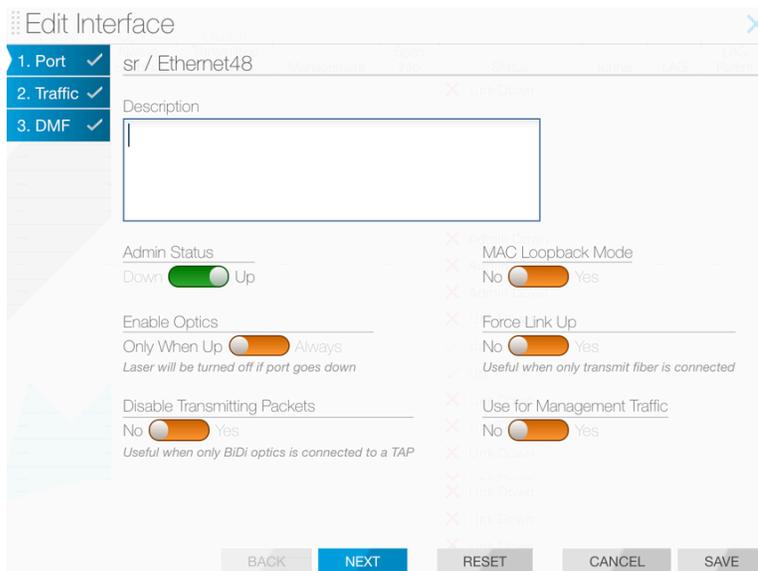
Figure 5-2: Fabric Interfaces Option



This page allows monitoring and configuring the interfaces on switches connected to the DANZ Monitoring Fabric (DMF) Controller. To display details about a specific interface, click the **Expansion** control to the left of the interface, and the entry expands to show any Tunnel Interfaces or Core Links currently using the interface.

To configure an interface, click the **Menu** control next to the interface and select **Configure** from the pull-down menu.

Figure 5-3: Configuring Interface Settings - Page 1



This dialog provides access to three pages.



Note: Page 3 is the DMF page configuration settings for interfaces to use in policies. For further information, refer to the **DANZ Monitoring Fabric User Guide**.

Page 1: The **Port** dialog provides the following options:

- **Admin Status:** Enable or disable the switch administratively.
- **Enable Optics:** Change the default to cause the optical laser to be left on after the port goes down.
- **MAC Loopback Mode:** Returns traffic to the originating interface.
- **Force Link Up:** This is useful to enable when only the transmit fiber is connected.
- **Description:** Assign a description for the interface.

Tip:

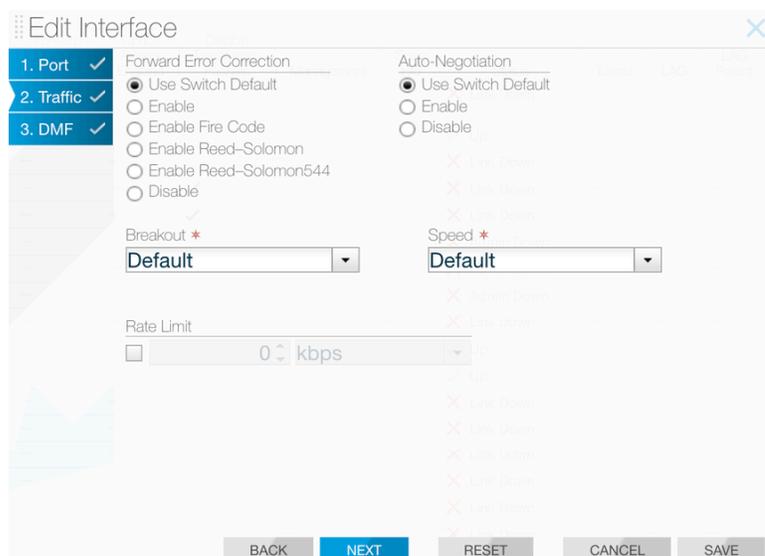


- Ideally, only apply a force link-up configuration on a delivery interface.
- This configuration allows L1 on the port to stay up, even when the optical fiber cable is connected only in the TX direction.
- This feature helps to black hole traffic if applied on links between switches.

- When finished, click **Save**. To configure traffic options, click **Next**. After clicking **Next**, the system displays the following page.

Page 2:

Figure 5-4: Configuring Interface Settings - Page 2



This page provides the following options.

- Forward Error Correction
 - Auto-Negotiation
 - Breakout
 - Speed
 - Rate Limit
- After making any changes required, click **Save**.

5.3.1 Forward Error Correction (FEC)

Use **Page 2** of the **Edit Interface** dialog to explicitly enable or disable forward error correction or to restore the default.

CLI Procedure

To use the CLI to explicitly enable or disable forward error correction or restore the default, use the following commands from **config-switch** mode:

```

controller-1(config-switch)# interface ethernet10
controller-1(config-switch-if)# forward-error-correction
disable Force disable interface forward-error-correction
enable Force enable interface forward-error-correction
enable-fire-code Force enable interface fire-code forward-error-correction
enable-reed-solomon Force enable interface reed-solomon forward-error-correction
enable-reed-solomon544 Force enable interface reed-solomon544 forward-error-correction
controller-1(config-switch-if)# forward-error-correction enable
controller-1(config-switch-if)# show this
! switch
switch DMF-FILTER-SWITCH-1
!
interface ethernet10
autoneg disable
forward-error-correction enable
controller-1(config-switch-if)# forward-error-correction disable
controller-1(config-switch-if)# show this
! switch
switch DMF-FILTER-SWITCH-1
!
interface ethernet10
autoneg disable
forward-error-correction disable
controller-1(config-switch-if

```

Replace **intf-port-list** by the interface name or port list.

The following summarizes the effect of each forward error correction keyword option:

- **disabled** – Disable if possible in the current port context.
- **enabled** – Enable if possible in the current port context.
- **enable-fire-code** – Request Fire-Code FEC (CL74) on port on port context.
- **enable-reed-solomon** – Request Reed-Solomon FEC (CL91, CL108) on port context.
- **enable-reed-solomon544** – Force Reed-Solomon544 on port context.



Note: Switch platforms with the Tomahawk ASIC cannot support Reed-Solomon FEC (CL108) on 25G interfaces. FEC options supported on 25G interfaces are limited Fire-Code FEC or FEC Disabled. This limitation does not apply to 100G interfaces.

5.3.2 Autonegotiation

The following summarizes the effect of each option:

- **autoneg enabled** – Enable if possible in the current port context.
- **autoneg disabled** – Disable if possible in the current port context.

Autonegotiation can be enabled or disabled for the following interface configurations:

-
- 100 GbE DAC in 100 GbE mode
 - 25 GbE
 - 1G-BASE-SX
 - 1G-BASE-LX

GUI Procedure

Use **Page 2** of the **Edit Interface** dialog to enable or disable autonegotiation, or to restore the default.

CLI Procedure

To use the CLI to explicitly enable or disable autonegotiation, or to restore the default, enter the following command from config-switch mode for any fabric switch:

```
controller-1(config-switch)# interface ethernet10
controller-1 (config-switch-if)# autoneg enable
controller-1 (config-switch-if)# show this
! switch
switch DMF-FILTER-SWITCH-1
!
interface ethernet10 autoneg enable
controller-1 (config-switch-if)# autoneg disable
controller-1 (config-switch-if)# show this
! switch
switch DMF-FILTER-SWITCH-1
!
interface ethernet10
autoneg disable
controller-1 (config-switch-if)#
```

Replace *intf-port-list* by the interface name or port list.

5.3.3 Manually Setting the Interface Speed

To manually set the interface speed for an interface, from *config-switch-if* submode, enter the **speed** command, which has the following syntax.

```
[no] speed [{100G | 25G | 200G | 1G | 10G | 40G | 400G | 50G}]
```

The following are the options supported.

- 100G Set interface speed to 100 Gbps
- 10G Set interface speed to 10 Gbps
- 1G Set interface speed to 1 Gbps
- 200G Set interface speed to 200 Gbps
- 25G Set interface speed to 25 Gbps
- 400G Set interface speed to 400 Gbps
- 40G Set interface speed to 40 Gbps
- 50G Set interface speed to 50 Gbps

5.3.4 Using Breakout Cables

DANZ Monitoring Fabric (DMF) supports using breakout (splitter) cables to split a single 40 GbE, 100 GbE, 200 GbE, and 400 GbE port into individual sub-interfaces.

For a list of supported switches, ports, and breakout cables, refer to the **DANZ Monitoring Fabric 8.5 Hardware Compatibility List**. The breakout cables listed in the **Hardware Compatibility List** are broken out automatically; manually entering the `breakout` command is not required.

GUI Procedure

To use the GUI to manually enable the use of multiple interfaces on a single switch port with a breakout cable, select **Fabric > Interfaces**, select **Edit** from the menu control for an interface, and use the settings on the **Traffic** page (Page 2) of the **Edit Interface** dialog.

To enable the use of multiple interfaces on a single switch port with a breakout cable, complete the following steps.

CLI Procedure

Use the `breakout mode` command to configure the breakout property for the current interface to configure the force breakout on an interface. When the config is applied, the interface, if it supports the breakout for the mode, is broken out into sub-interfaces based on the specified mode. Auto is the default option, which lets the switch automatically select the mode for the breakout. The following are the modes supported.

- 2x100G Breakout to 2 sub-interfaces of 100G each
- 2x200G Breakout to 2 sub-interfaces of 200G each
- 2x40G Breakout to 2 sub-interfaces of 40G each
- 2x50G Breakout to 2 sub-interfaces of 50G each
- 4x100G Breakout to 4 sub-interfaces of 100G each
- 4x10G Breakout to 4 sub-interfaces of 10G each
- 4x1G Breakout to 4 sub-interfaces of 1G each
- 4x25G Breakout to 4 sub-interfaces of 25G each
- 4x50G Breakout to 4 sub-interfaces of 50G each
- 8x10G Breakout to 8 sub-interfaces of 10G each
- 8x25G Breakout to 8 sub-interfaces of 25G each
- 8x50G Breakout to 8 sub-interfaces of 50G each

1. To verify the breakout-capable ports on the switch, enter the `show switch interfaces` command, which has the following syntax:

```
show switch switch-name interfaces
```

To locate breakout-capable ports from the Controller, enter the following command:

```
controller-1> show switch DMF-F1 interfaces
```

#	IF Name	MAC Address	Config	State	Adv. Features	Curr Features	Supported Features
1	ethernet1	5c:16:c7:13:d5:9f (Big Switch)	up	up	autoneg, fec, 1g, 10g, 25g	copper, autoneg, fec, 25g	copper, autoneg, fec, 1g, 10g, 25g
2	ethernet2	5c:16:c7:13:d5:a0 (Big Switch)	up	up	autoneg, fec, 1g, 10g, 25g	copper, autoneg, fec, 25g	copper, autoneg, fec, 1g, 10g, 25g
...							
49	ethernet49	5c:16:c7:13:d5:d5 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g
50	ethernet50	5c:16:c7:13:d5:d6 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g
51	ethernet51	5c:16:c7:13:d5:d7 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g
52	ethernet52	5c:16:c7:13:d5:d8 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g
53	ethernet53	5c:16:c7:13:d5:d9 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g
54	ethernet54	5c:16:c7:13:d5:d0 (Big Switch)	up	up	40g	fiber, 40g	fiber, bsn-breakout-capable, 1g, 10g, 25g, 40g, 100g

Each breakout-capable port is identified by the string **bsn-breakout-capable** in the Supported Features column. In this example, ports ethernet49 through ethernet54 are breakout-capable.

2. Enter the **enable** and **configure** command to enter **config mode**, as in the following example.

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

3. Enter the **config-switch** submode and then the **config-switch-if** submode to enter the **breakout** command, as in the following example.

```
controller-1 (config)# switch DMF-FILTER-SWITCH-1
controller-1 (config-switch)# interface ethernet54
controller-1 (config-switch-if)# breakout mode 4x10G
```

4. Enter the **show switch switch-name** interface command to verify the operation, as in the following example.

```
controller-1 (config-switch-if)# show switch DMF-FILTER-SWITCH-1 interfaces
# IF Name      MAC Address      Config State Adv. Features      Curr Features      Supported Features
-----|-----|-----|-----|-----|-----|
1 ethernet1    5c:16:c7:13:d5:9f (Big Switch) up      up      autoneg, fec, 1g, 10g, 25g copper, autoneg, fec, 25g copper, autoneg, fec, 1g, 10g, 25g
2 ethernet2    5c:16:c7:13:d5:a0 (Big Switch) up      up      autoneg, fec, 1g, 10g, 25g copper, autoneg, fec, 25g copper, autoneg, fec, 1g, 10g, 25g
...
54 ethernet54/1 5c:16:c7:13:d5:d5 (Big Switch) up      up      40g fiber, 10g fiber, 10g
55 ethernet54/2 5c:16:c7:13:d5:d6 (Big Switch) up      up      40g fiber, 10g fiber, 10g
56 ethernet54/3 5c:16:c7:13:d5:d7 (Big Switch) up      up      40g fiber, 10g fiber, 10g
57 ethernet54/4 5c:16:c7:13:d5:d8 (Big Switch) up      up      40g fiber, 10g fiber, 10g
```

5. The breakout ports are named **ethernetx/1** through **ethernetx/4**. The example output shows four 10G ports where there was previously a single 40G port (**ethernet54/1** through **ethernet54/4**).

5.3.5 Verifying Switch Configuration

GUI Procedure

Use the **Fabric > Interfaces** option to view the interfaces table, which provides information about the configuration and activity on each interface of the switches connected to the DANZ Monitoring Fabric (DMF) controller.

CLI Procedure

To view the configuration or activity for a specific interface, use the **show switch switchname interfaces** command. The **detail** option provides additional information about the interface, including the up and down counts, indicating if the interface has been flapping. The output also indicates if the interface supports breakout interfaces.

The following is an example.

```
controller-1# show switch DMF-DELIVERY-SWITCH-1 interfaces ethernet49 detail
# IF Name      MAC Address      Config State Adv. Features      Curr Features      SupportedFeatures
-----|-----|-----|-----|-----|-----|
1 ethernet49    5c:16:c7:13:d5:d8 (Big Switch) up      down    fec, 25g, 50g, 100g fec      fec,bsn-breakout-capable, 100g
```

To view the interface descriptions, use the **show switch switchname interface description** command as shown in the following example.

```
controller-1 (config-switch-if)# show switch DMF-DELIVERY-SWITCH-1 interface description
# Switch Name      IF Name      Description
-----|-----|-----|
1 DMF-DELIVERY-SWITCH-1 ethernet1    100g-to-SFO
2 DMF-DELIVERY-SWITCH-1 ethernet2    100g-to-NYC
```

5.3.6 Disabling the TX Direction on an Interface

DANZ Monitoring Fabric (DMF) supports disabling the TX direction of an interface, which is useful for example on a DMF switch whose filter interfaces are directly connected to a 40 GbE *bidirectional tap* using 40 GbE *BiDi optics*.



Note: when using bidirectional TAPs, either you need to use RX-only BiDi optics or you need to disable the TX direction on the DMF interfaces; otherwise, the optic transmits light back to the TAP and interferes with the production link. This is the case with all bidirectional TAPs: the outputs of the bidirectional TAPs should be connected to transceivers that don't transmit because there is no way for the TAP to filter out light coming back from the packet broker.

1. If a filter interface is connected to a bidirectional TAP using a BiDi optic, you should configure that filter interface with the `disable-xmit` command to avoid optical interference with the TAP:

Example of command usage:

```
dmf-controller-1(config)# switch sw-filter1
dmf-controller-1(config-switch)# interface ethernet2
dmf-controller-1(config-switch-if)# disable-xmit
```

DMF Upgrade Procedures

This chapter describes how to upgrade the DANZ Monitoring Fabric (DMF) Controller and fabric switches after initial installation.

6.1 Upgrading the Controller

The default serial console baud rate on **DMF 6.1** and later hardware appliances is **115200**. Arista recommends against using the serial interface to perform an upgrade. However, when upgrading an existing Controller appliance with a serial interface set to **9600** baud, change the terminal setting to **115200** after upgrading to **DMF 6.1.0** or later.



Note: The upgrade process checks for supported versions when upgrading. Refer to the most recent **DANZ Monitoring Fabric 8.5.0 Release Notes** for a list of supported versions for upgrade.

The DANZ Monitoring Fabric (DMF) Controller platform supports two partitions for Controller software images. The active partition contains the active image currently running on the Controller. The alternate partition can be updated without interrupting service. The image used for booting the Controller is called the boot image. In addition, the Controller has an image repository in its local file system where you can copy upgrade images. The upgrade image is verified for integrity before it is copied and rejected if it fails the checksum test.



Note: Copying an upgrade bundle to the Controller overwrites any older image currently in the image repository. When downloading a bundle identical to the image repository version, the operation fails when calculating the checksum after copying is complete.

After copying the upgrade image to the image repository, use the **upgrade stage** command to copy the upgrade image to the alternate partition and prepare the Controller for the upgrade.

The boot image remains in the active partition after copying the ISO image file to the alternate partition. After entering the **upgrade launch** command, the upgrade image is changed to the boot image and the Controller reboots.



Note: The upgrade process is not hitless.

The **upgrade launch** command copies the active state of the Controller to the alternate partition and reboots the Controller using the upgrade image. After completing the upgrade process, the alternate and active partitions are reversed, the Controller is running with the upgrade image, and the older image remains available in the alternate partition until a new image is copied into it, which will overwrite it.

As long as the original image remains in the alternate partition, use the **boot partition alternate** command to restore the Controller to its previous software and configuration.

Log files are written in a separate partition, available from either of the two images by entering the **show logging controller** command.

6.1.1 Upgrade Procedure Summary

Complete the following steps to upgrade the Controllers. The switch upgrade process, using ZTF, happens automatically (see “[DMF Switch Automatic Upgrade](#)” section for details)



Note: Only the admin user can perform an upgrade or enter the **show image** command.

Procedure

1. Log in to the active Controller using its individual IP address or the virtual IP address for the cluster.
2. From the active Controller, copy the ISO image file to the image repository on the active Controller (see the [Copying the ISO Image File to the Controller](#) section).
3. From the active Controller, enter the upgrade stage command and respond to the prompts as required (see the [Staging the Upgrade](#) section for details).
4. From the active Controller, enter the upgrade launch command and respond to the prompts as required (see the [Launching the Upgrade](#) section for details).



Note: Do not attempt to launch or stage another upgrade process until the current process is either completed or times out.

Wait for the process to complete and verify the upgrade before making any configuration changes.

5. Verify the upgrade (see the [Verifying the Upgrade](#) section).

6.1.2 Upgrade Options

The following options are available for use with the **upgrade** command.

cluster: with the cluster option, the upgrade command is executed cluster-wide. The user must log in to the active Controller to run the cluster option.

launch: Complete the upgrade process by rebooting the Controller from the alternate partition and transferring state and configuration information from the Controller to the upgraded Controller and its running-config. This keyword manages the transition from the current version to the next version, which may include rebooting all Controllers in the cluster, rebooting.

Use the following options with the **launch** keyword as required:

- **support-bundle:** Generate a support bundle for use by Arista tech support.
- **switch-timeout:** Optionally, specify the number of seconds to wait before terminating the command to a switch during the upgrade.
- **pre-launch-check:** Identifies the status of the Controller regarding readiness for upgrade.
- **stage:** Prepares the platform for the upgrade ahead of the actual upgrade process by copying the upgrade image to the alternate partition on the Controller.

6.2 Copying the ISO Image File to the Controller

The ISO image file is a software image that includes the files required to complete the upgrade of the Controller. It also contains the files for installing or upgrading the Switch Light OS image on switches. The primary component of the upgrade image is a new root file system (rootfs) for the Controller.

When copying the upgrade image, if there is not enough file space on the local file system, the system prompts to create space on the local file system. During the image copy process, if the integrity of the image cannot be verified, the system displays a message, and the copy fails. After copying the image, a warning message appears if the image is not a compatible upgrade image for the existing application, or if the image is not a newer version.



Note: To copy the image file (or other files) from the workstation command prompt, refer to the [Copying Files Between a Workstation and a DMF Controller](#) section.

Note: Starting from DMF release 8.4, due to an increased image size, when trying to copy a new image you will likely encounter the following error:

```
Error: Invalid Use: No space left, use "delete image" command to free space
```

It is therefore required to remove all images from the image partition, including the image for the currently running version, to free up space.

```
CONTROLLER-1# show cluster image
# Checksum Cluster Nodes          Product          Version Build
-|-----|-----|-----|-----|-----|
1 d301a     CONTROLLER-1, CONTROLLER-2 DANZ Monitoring Fabric 8.5.0   33
CONTROLLER-1#
```



```
CONTROLLER-1# delete cluster image all
all: delete: confirm ("y" or "yes" to continue): yes
CONTROLLER-1#
```

Note that you can also delete images on the individual Controllers like so:

```
CONTROLLER-1# show image
CONTROLLER-1# delete image all
```

Even though you have deleted all the images from the /images repository, you can still roll back to the current version after the next upgrade operation, as the current image is still present in the boot partition.

To use the Controller CLI to copy the ISO image file to the Controller Image repository from an external server, use the `copy` command, which has the following syntax:

```
controller-1# copy <source> <destination>
```

Replace **source** with a location accessible from the Controller node using one of the following protocols:

- `scp://<user>@<host>:path`: Remote server filename, path, and user name with access permissions to the file. The remote system prompts for a password if required.
- `http://`: HTTP URL including path and filename.
- `https://`: HTTPS URL including path and filename

Replace **destination** with **image://cluster**. This will download the image from the remote server to all the nodes in the cluster. For example, the following command copies the file **DMF-*<X.X.X>*-Controller-Appliance-*<date>*.iso** from `myscpserver.example.com` to the Controller alternate partition for both active and standby Controller:

```
controller-1# copy scp://admin@myscpserver.example.com:*DMF-<X.X.X>-Controller-Appliance-<date>.iso* image://cluster
```

Use the **show cluster image** command to verify that the ISO image file has successfully been copied to both the active and standby Controller, as shown in the following example:

```
controller-1# show cluster image
```

6.3 Staging the Upgrade

The **upgrade cluster stage** command applies the specified upgrade image into the alternate partition for both active and standby Controllers. This command prepares the platform for an upgrade by populating the alternate partition with the contents of the new image. This is a separate step from the launch step, where the upgraded image becomes the active image to prepare the platform ahead of the upgrade process.

To view the upgrade images currently available, use the **show upgrade image** or **show cluster image** commands.

Use the **upgrade cluster stage** command to prepare the Controllers for the upgrade. This action copies the running-config to a safe location and marks the alternate partition as the boot partition.

The system responds as shown in the following example:

```
controller-1# # upgrade cluster stage
Upgrade stage: alternate partition will be overwritten
proceed ("yes" or "y" to continue): y
```



Note: Do not attempt to launch or stage another upgrade process until the current process is either completed or times out.

At this point, enter **y** to continue the staging process. The system continues the process and displays the following prompts:

```
Upgrade stage: copying image into alternate partition
Upgrade stage: checking integrity of new partition
Upgrade stage: New Partition Ready
Upgrade stage: Upgrade Staged
```

To verify that the system is ready for upgrade, enter the **show boot partition** command, as in the following example:

```
controller-1# #show boot partition
```

This command lists the available partitions and the information about the Controller versions installed on each. In this example, the original image remains in active state and is still the boot image, meaning if a reboot occurs now, the original image is used to reboot. The upgrade image will not be used to boot until

it has been changed to the boot image, which is one of the effects of the `upgrade cluster launch` command.

To display the current status of the upgrade process, you can use the `show upgrade progress` command, as in the following example:

```
controller-1# # show upgrade progress
```



Note: Upgrade requires a successfully staged image partition. If the upgrade stage is interrupted, it is necessary to stage the image again to launch the upgrade.

6.4 Launching the Upgrade

The `upgrade cluster launch` command reboots the system using the upgrade image in the alternate partition and copies the current state from the previous Controller. The system collects information from the existing cluster, for example, the current running- config and saves this state in the staged partition, so it is available for the new cluster.

Once the running-config is applied, the existing operational state of the existing cluster is requested and then applied to the new cluster. The new cluster requests the switches from the old cluster, adjusts its local configuration to describe this new ownership, and reboots the switches.

The `upgrade cluster launch` command manages the transition from the current version to the next for both active and standby Controllers. This includes various steps, including rebooting all the Controllers in the cluster, rebooting all the switches, and upgrading the switches. After the standby node is upgraded, the roles of active and standby are reversed, so the upgraded node can assume control while the other node is upgraded and reboots.



Note: Do not attempt to launch or stage another upgrade process until the current process is either completed or times out.

Log into the active Controller via the Controller management IP or virtual IP and enter the `upgrade cluster launch` command:

```
controller-1# upgrade cluster launch
```

The system prompts to proceed with the reboot. Enter `yes`. The standby Controller reboots and displays the following messages while the active Controller waits for the standby Controller to boot up:

```
Upgrade Launched, Rebooting  
Broadcast message from root@controller (unknown) at 19:40 ...The system is going down for  
reboot  
NOW!User initiated reboot
```

While rebooting, the SSH terminal session is terminated. Reconnect after the reboot is complete.

6.5 DMF Switch Automatic Upgrade

When the fabric switch reboots, it compares its current software image, manifest, and startup-config to the information in the switch manifest that the Controller sends after the switch reboots. The switch optimizes the process by caching its last known good copies of the software image and the startup-config in its local flash memory.

The switch automatically starts the upgrade process when it reboots if the ZTF server has a software image or startup-config with a different checksum than it currently has. This check is performed every time the switch restarts.

6.6 Verifying the Upgrade

After completing the upgrade process, verify that both Controllers and the fabric switches have been upgraded to the correct DANZ Monitoring Fabric and Switch Light OS versions.

To verify the upgrade of the Controllers, log in to each Controller and enter the `show version` command, as shown in the following example:

```
controller-1># show version
```

To verify the upgrade of the switches, enter the `show switch all desc` command, as shown in the following example:

```
controller-1># show switch all desc
```

This command displays the current switch version.

6.7 Verify Persistent IPAM Assigned IP Addresses after Upgrade

IPAM-assigned IP addresses for switches are persistent across DMF upgrades.

Configuration

No additional configuration changes are required to keep IPAM-assigned IP address configuration persistent across cluster upgrades. However, this assumes that IPAM has been previously configured.

CLI Commands

Refer to the sections [Using L2 ZTF \(Auto-Discovery\) Provisioning Mode](#) through [Static Address Assignment Troubleshooting and Limitations](#) for IPAM and switch IP address configuration, which is also valid after the upgrade.

6.8 Rolling Back an Upgrade

When deciding to rollback or downgrade the Controller software after an upgrade, note that the rollback or downgrade is not hitless and will take several minutes to complete. After both Controllers are up and have joined the cluster, check each switch version and reboot each switch as needed to ensure all the switches have an image compatible with the Controller version.

To restore the system to the previous image, complete the following steps:

Procedure

1. On both the active and standby Controllers, enter the `show boot partition` command to ensure that the previous image remains in the Alternate partition.

```
controller-1># show boot partition
```

2. Reboot the active Controller node from the alternate partition by entering the `boot` command from the CLI prompt of the active Controller.

```
controller-1># boot partition alternate
Next Reboot will boot into partition 1 (/dev/vda2)
boot partition: reboot? ("y" or "yes" to continue): yes
```

Answer **yes** when prompted.

3. Reboot the standby Controller node from the alternate partition by entering the `boot` command from the CLI prompt of the standby Controller.

```
standby controller-2># boot partition alternate
Next Reboot will boot into partition 1 (/dev/vda2)
boot partition: reboot? ("y" or "yes" to continue): yes
```

Answer **yes** when prompted.

4. After both Controller nodes have restarted, reboot all switches by entering the `system reboot switch` command on the active Controller.

```
controller-1># system reboot switch all
system switch reboot all: connected switches:
00:00:70:72:cf:c7:06:bf
00:00:70:72:cf:c7:c5:f9
00:00:70:72:cf:ae:b7:38
00:00:70:72:cf:c8:f9:25
00:00:70:72:cf:c7:00:ad
00:00:70:72:cf:c7:00:63
reboot may cause service interruption
system switch reboot all ("y" or "yes" to continue): yes
Answer ``yes`` when prompted.
```

5. Wait for all the switches to reconnect.

```
controller-1># show switch
```

6. Verify that the standby Controller has rejoined the cluster with the reverted active Controller by entering the `show controller` command from both nodes.

```
controller-1># show controller
```

Any connected DMF Service Node or DMF Recorder Node must be upgraded after upgrading DMF fabric Controllers and switches.

Upgrade from **DMF Release 7.1.0** to later versions uses Zero Touch Fabric (ZTF) and occurs automatically for connected nodes after the DMF Controller is upgraded.



Note: If the DMF Recorder Node is in a different Layer 2 segment than the DMF Controller, a fresh installation of the Recorder is required to upgrade to **Release 7.1.0**.

For details about upgrading DMF Service Node software, refer to [Installing and Upgrading the DMF Service Node](#).

Installing and Upgrading the DMF Service Node

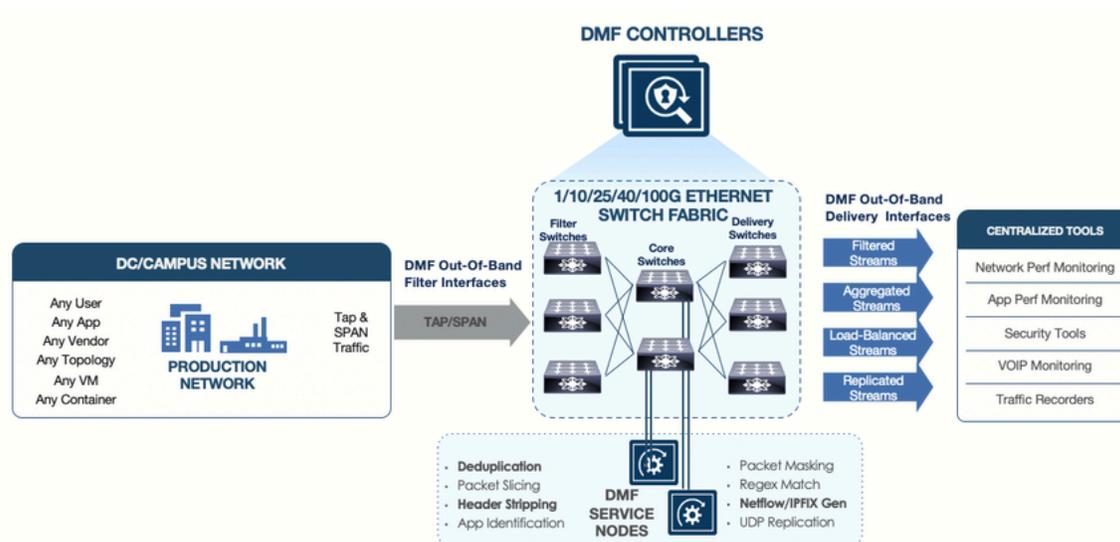
This chapter describes how to install the DANZ Monitoring Fabric DMF Service Node.

7.1 Overview

The DANZ Monitoring Fabric (DMF) Service Node provides advanced packet matching and modification capabilities for monitored traffic. The DMF Service Node is an optional component in the DANZ Monitoring Fabric, providing advanced features. The DMF Service Node offers the following services:

- Deduplication
- Header stripping
- IPFIX generation
- Packet masking
- NetFlow
- Pattern dropping
- Pattern matching
- Packet slicing
- Timestamping
- UDP replication

Figure 7-1: DMF Service Node



After the basic installation, the DMF Controller automatically detects each connected DMF Service Node, and the Controller starts managing the DMF Service Node Appliance along with the connected fabric switches.

For information about configuring and using the DMF Service Node, refer to the ***DANZ Monitoring Fabric 8.4 User Guide***.

7.2 Connecting the Service Node

Connect the DMF Service Node to a core interface on a DMF switch, which provides access to filter interfaces, where traffic is received for the DMF Service Node, and to delivery interfaces, where traffic is delivered after it is processed. The service node is connected in two ways:

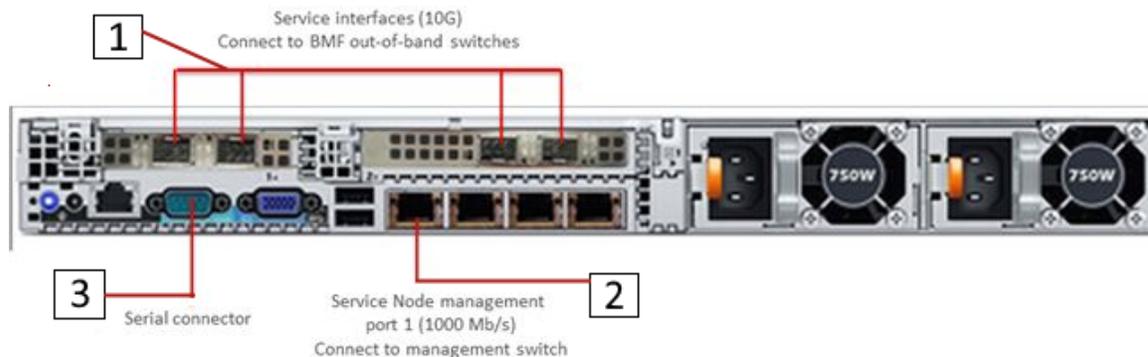
- Using the management interface (GbE) connected to a management switch, which provides control plane connectivity between the Service Node and the DMF controller.
- Using the Service Node data interfaces (SNI) which connect to DMF switches. For information on supported SNI interfaces, refer to the *DMF Hardware Guide*.



Note: Arista recommends having an iDRAC connection for the DMF Controller, DMF Service Node, Arista Analytics Node, and DMF Recorder Node appliances. This connection helps in easy troubleshooting of issues. For more details, refer to the chapter on Using iDRAC later in this guide.

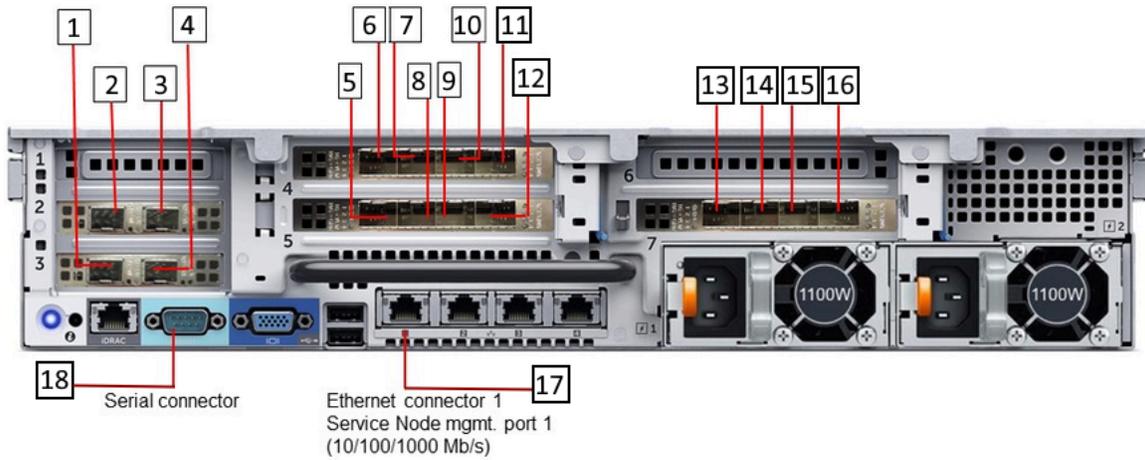
The figure below shows the interfaces provided on the 4-port DMF Service Node Appliance.

Figure 7-2: DMF Service Node (4-Port Appliance)



- 1 Service Interfaces (10G)
- 2 Service Node Management Port 1 (1000 Mb/s) - Connect to the management switch.
- 3 Serial Connector

Figure 7-3: DMF Service Node BL (16-Port Appliance)



- | | | | |
|---|--------------------------|----|--|
| 1 | Service interfaces SNI13 | 10 | Service interfaces SNI10 |
| 2 | Service interfaces SNI15 | 11 | Service interfaces SNI9 |
| 3 | Service interfaces SNI16 | 12 | Service interfaces SNI5 |
| 4 | Service interfaces SNI14 | 13 | Service interfaces SNI4 |
| 5 | Service interfaces SNI8 | 14 | Service interfaces SNI3 |
| 6 | Service interfaces SNI12 | 15 | Service interfaces SNI2 |
| 7 | Service interfaces SNI11 | 16 | Service interfaces SNI1 |
| 8 | Service interfaces SNI17 | 17 | Ethernet Connector 1 Service Node Management Port 1 (10/100/1000 Mb/s) |
| 9 | Service interfaces SNI6 | 18 | Serial Connector |

7.3 Service Node Setup and Initial Configuration

This section describes how to perform the initial setup and configuration on a new DANZ Monitoring Fabric (DMF) Service Node appliance.



Note: Disable hyperthreading on the hardware appliance before installing the Service Node software to avoid performance issues. When disabling hyperthreading after installing the Service Node software, performance will be affected, and reinstalling the software will be required to resolve the issue.

Complete the following steps to run the first boot setup, which performs the initial configuration required for a new DMF Service Node.

Procedure

1. Rack the DMF Service Node Appliance.

The appliance interfaces are on the back of the appliance, where the power cord connects. These include the following:

- Four management interfaces (10/100/1000 Mb/s): You can connect either of the two lower left interfaces (Ethernet 1 or Ethernet 2) to the network management switch.
 - One serial interface (db9).
 - Four 10-GbE SFP ports on the R640 server and 16 10-GbE ports on the R740 server. Connect these ports to a DMF switch.
2. Turn on the DMF Service Node server appliance.
 3. Log in via the serial port using the admin account name. The baud rate is **115200**. When using a terminal server to connect, ensure the baud rate on the terminal server is **115200**.
 4. When the first boot process begins, accept the End User License Agreement (EULA).

```
This product is governed by an End User License Agreement (EULA). You must accept this
EULA to continue using this product.
You can view this EULA by typing 'View', or from our website at. https://www.arista.com/en/eula
Do you accept the EULA for this product? (Yes/No/View) [Yes] > yes Running system pre-check
Finished system pre-check
Starting first-time setup
```

5. Complete the local node configuration according to the requirements of your network environment. The following is only an example. Change for your specific deployment.

```
Local Node Configuration
-----
Emergency recovery user password >
Emergency recovery user password (retype to confirm) >
Hostname > DMF-Service-Node
Management network options:
[1] IPv4 only
[2] IPv6 only
[3] IPv4 and IPv6
>1
IPv4 address [0.0.0.0/0] > 10.8.39.200/18
IPv4 gateway (Optional) > 10.8.0.1
DNS server 1 (Optional) > 10.3.0.4
DNS server 2 (Optional) > 10.1.5.200
DNS search domain (Optional) > qa.arista.com
Administrator password >
Administrator password (retype to confirm) >
Controller address if deployment mode is preconfigured (L3 ZTN) (Optional) > 10.106.6.4
```

6. If the DMF Service Node is connected to the DMF Controller by a Layer 3 device (such as a router) in preconfigured (L3 ZTN) mode, enter the active DMF Controller's IP address.



Note: Starting with **DMF Release 7.1.0**, the DMF Service Node can be installed using Zero Touch Fabric (ZTF) even if it is in a different subnet than the DMF Controller.

7. Identify the Network Time Protocol (NTP) servers.

```
System Time
-----
Default NTP servers:
- 0.bigswitch.pool.ntp.org
- 1.bigswitch.pool.ntp.org
- 2.bigswitch.pool.ntp.org
- 3.bigswitch.pool.ntp.org
NTP server options:
[1] Use default NTP servers
[2] Use custom NTP servers
[1] > 1
```

- When prompted, type **1** to apply the selected options, or type any number on the menu that is displayed to change the current setting.

```
Please choose an option:
[ 1] Apply settings
[ 2] Reset and start over
[ 3] Update Recovery Password (*****)
[ 4] Update Hostname (R740)
[ 5] Update IP Option (IPv4 only)
[ 6] Update IPv4 Address (10.106.6.7/23)
[ 7] Update IPv4 Gateway (10.106.6.1)
[ 8] Update DNS Server 1 (10.108.200.200)
[ 9] Update DNS Server 2 (10.100.5.200)
[10] Update DNS Search Domain (qa.arista.com)
[11] Update Admin Password (*****)
[12] Update Controller IP (10.106.6.4)
[13] Update NTP Option (Use default NTP servers)
[1] > 1
```

- After first-time setup is complete, press **Enter** to continue.

```
[Stage 1] Initializing system
[Stage 2] Configuring local node
Waiting for network configuration IP address on bond0 is 10.8.39.200 Generating
cryptographic keys
[Stage 3] Configuring system time
Initializing the system time by polling the NTP servers:
0.bigswitch.pool.ntp.org
1.bigswitch.pool.ntp.org
2.bigswitch.pool.ntp.org
3.bigswitch.pool.ntp.org
[Stage 4] Configuring cluster Cluster configured successfully. Current node ID is 27521
All cluster nodes:
Node 27521: 10.8.39.200:6642
First-time setup is complete!
Press enter to continue >
DMF Service Node (dmf-8.0-service-node #1) Log in as 'admin' to configure
```

- Connect one or more of the 10 GbE SFPs on the appliance hardware to a DMF out-of-band switch. The DMF Controller automatically detects each connected DMF Service Node Appliance and integrates the service node into the monitoring fabric.



Note: The DMF Controller will not establish a connection to the DMF Service Node Application if none of the 10 GbE SFPs on the appliance hardware are connected to a DMF out-of-band switch.

- Once connected to the DMF Controller, it can take several minutes for Service Node to update the software.

Figure 7-4: Service Node software update

```
(13) Update NTP Server 2 (time.windows.com)
[1] > 1
[Stage 1] Initializing system
[Stage 2] Configuring local node
Waiting for network configuration
IP address on bond0 is 10.240.156.5
Generating cryptographic keys
[Stage 3] Configuring system time
Initializing the system time by polling the NTP servers:
time.google.com
time.windows.com
[Stage 4] Configuring cluster
Cluster configured successfully.
Current node ID is 18414
All cluster nodes:
Node 18414: 10.240.156.5:6642

Checking for updates...
This may take a few minutes
The system may reboot while upgrading
```

- Login to the DMF Controller.
- Add the following service node configuration with the Service Node management interface's MAC address.

```
controller-1> enable
controller-1# config
controller-1(config)# service-node device-name
```

```
controller-1(config-service-node) # mac service-node-management-int-mac-address
controller-1(config-service-node) #
```

Note: Obtain the MAC address of the Service Node by logging into the service node via SSH and running the following command taking note of the **bond0** MAC address:



```
SN-1(config) # show local-node interfaces bond0
~~~~~ Interfaces ~~~~~
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond
mode Bond role
-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|
bond0 78:ac:44:05:65:b8 (Dell) up up active-backup

~~~~~ Addresses of Interfaces ~~~~~
# Interface Ip cidr
-|-----|-----|-----|-----|-----|-----|-----|-----|
1 bond0 10.240.130.26/25
2 bond0 fe80:0:0:0:7aac:44ff:fe05:65b8/64
```

14. Verify that the DMF Service Node is connected by entering the following command:

```
controller-1(config-service-node) # show managed-service-device
```

Enter this command in any CLI mode.

7.4 Creating Support Bundle on Service Node

The support bundle for the DANZ Monitoring Fabric (DMF) Service Node should be created from the DMF Controller ([Creating a Support Bundle](#)). But, if the DMF Service Node loses connectivity to the DMF Controller, a support bundle can be created by logging into the DMF Service Node.

The DMF Service Node CLI provides commands to automate the collecting, archiving, and uploading of critical data.

The following are the commands to configure Support Bundle auto-upload:

```
LG-SN(config) # service
LG-SN(config-service) # support auto-upload
<cr>
LG-SN(config-service) # support auto-upload
Enabled diagnostic data bundle upload
Use "diagnose upload support" to verify upload server connectivity
LG-SN(config-service) #
```

To check if auto-upload is enabled or not:

```
LG-SN(config-service) # show run service
! service
service
support auto-upload
LG-SN(config-service) #
```

The following is the command to generate the Support Bundle. After generating the support bundle, it uploads automatically. Please provide the support bundle ID to support personnel.

```
LG-SN(config-service) # support
Generating diagnostic data bundle for technical support. This may take several minutes...
```

```

Support Bundle ID: SMFUG-BS5S2
Local cli collection completed after 32.9s. Collected 33 commands (0.14 MB)
Local rest collection completed after 0.0s. Collected 3 endpoints (0.17 MB)
Local bash collection completed after 93.3s. Collected 133 commands (6.73 MB)
Local file collection completed after 8.4s. Collected 42 paths (1851.13 MB)
Collection completed. Signing and compressing bundle...
Support bundle created successfully
00:03:16: Completed
Generated Support Bundle Information:
Name : anet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS5S2.tar.gz
Size : 490MB
File System Path : /var/lib/floodlight/support/anet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS5S2.tar.gz
Url : https://10.240.130.8:8443/api/v1/support/anet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS5S2.tar.gz
Bundle id : SMFUG-BS5S2
Auto-uploading support anet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS5S2.tar.gz
Transfer complete, finalizing upload
Please provide the bundle ID SMFUG-BS5S2 to your support representative.
00:01:03: Completed
LG-SN(config-service)#

```

The **show support** command shows the status of the automatic upload.

```

LG-SN(config-service)# show support
# Bundle Bundle id Size Last modified Upload status
-----|-----|-----|-----|-----|
1 anet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS5S2.tar.gz SMFUG-BS5S2 490MB 2022-04-13 07:49:20.157000 UTC upload-completed
2 anet-support--LG-SN--2022-04-13--07-19-08Z--SI51T-BVJJB.tar.gz SI51T-BVJJB 488MB 2022-04-13 07:22:44.927000 UTC
3 anet-support-component--LG-SN--2021-05-19--22-47-17Z_0kc7zxw.tar.gz 462MB 2021-05-19 22:47:17.452000 UTC
LG-SN(config-service)#

```



Tip: Use the **diagnose upload support** command to check the reachability and health of the server before uploading the support bundle.

```

LG-SN(config-service)# diagnose upload support
Upload server version: diagus-master-76
Upload diagnostics completed successfully
00:00:04: Completed
Check : Resolving upload server hostname
Outcome : ok
Check : Communicating with upload server diagnostics endpoint
Outcome : ok
Check : Upload server healthcheck status
Outcome : ok
Check : Upload server trusts authority key
Outcome : ok
Check : Signature verification test
Outcome : ok
Check : Resolving objectstore-accelerate hostname
Outcome : ok
Check : Resolving objectstore-direct hostname
Outcome : ok
Check : Communicating with objectstore-accelerate
Outcome : ok
Check : Communicating with objectstore-direct
Outcome : ok
LG-SN(config-service)#

```

7.5 Upgrading Service Node Software From Release 7.x.x

Upgrading the DANZ Monitoring Fabric (DMF) Service Node, deployed in L2ZTN from **DMF-7.0.0** to a later version, will be automatically completed through a zero-touch upgrade if you upgrade a **DMF Release 7.0.x** Controller.

Upgrade of DMF Service Node deployed in L3ZTN from **DMF- 7.1.0** to a later version will be automatically completed through zero-touch when you upgrade a **DMF Release 7.1.x** Controller.

To verify that the Service Node is ready for the zero-touch upgrade, enter the following command from the CLI prompt on the active DMF Controller.

```
controller-1> show service-node sn-name zerotouch
```

Zerotouch status should be OK.

Installing and Configuring the DMF Recorder Node

This chapter describes the installation, initial configuration, and upgrade of the DMF Recorder Node.

8.1 Overview

The DANZ Monitoring Fabric (DMF) Recorder Node (RN) is a traffic recording appliance with Arista Networks software running on Dell, Inc. servers.

The RN records packets from the network to disk and retrieves specific packets from disk quickly, efficiently, and at scale. It integrates with DMF for a single-pane-of-glass solution. A single DMF Controller can manage multiple RNs, delivering packets for recording through out-of-band policies. The Controller also provides centralized APIs for interacting with RNs to perform packet queries.

A DMF out-of-band policy directs the recording of matching packets to one or more RNs. The out-of-band policy references the switch and port where the RN is attached to the fabric. The policy treats these as “dynamic” delivery interfaces identified by unique names. The DMF Controller also provides commands for viewing errors, warnings, statistics, and the status of connected RNs.

The RN provides an agent that collects statistics and health information from the Controller. The agent also allows the Controller to configure the RN, eliminating the need to administer any RN directly during normal operation separately. To the DMF Controller, the agent causes the RN to appear as a special type of switch. Use the REST API to query the RN directly.

The DMF Recorder Node appliance is based on Dell server hardware supporting the interfaces mentioned in [Figure 8-3: DMF Recorder Node \(HWA\) Rear Panel](#). For a detailed hardware specification, refer to the *DMF Hardware Guide*.



Note: Arista recommends using an iDRAC connection to the DMF Controller, DMF Service Node, Arista Analytics Node, and DMF Recorder Node appliances. This connection helps with troubleshooting issues. For more details, refer to the chapter on *Using iDRAC* later in this guide.

The RN's storage capacity and data interfaces are:

- 192 TB packet storage
- 25 GbE recording interface
- 25 GbE auxiliary interface

Connect the RN's data interface (shown in [Figure 8-3: DMF Recorder Node \(HWA\) Rear Panel](#)) to a DMF delivery switch, where traffic is delivered after it is processed.

The following figure illustrates the bezel on the larger (HWA) DMF Recorder Node.

Figure 8-1: DMF Recorder Node (HWA) Front Panel

Recorder-Bezel

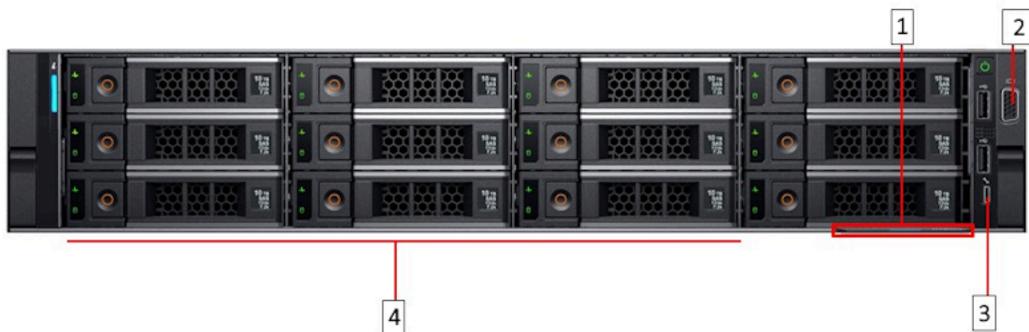


- | | | | |
|---|--|---|-----------------------------------|
| 1 | System identification button/indicator | 4 | LCD panel |
| 2 | Recorder Node Security Bezel | 5 | Power-on indicator / Power button |
| 3 | LCD menu buttons | 6 | USB ports |

The following figure illustrates the front panel of the DMF Recorder Node.

Figure 8-2: DMF Recorder Node (HWA) Front Panel

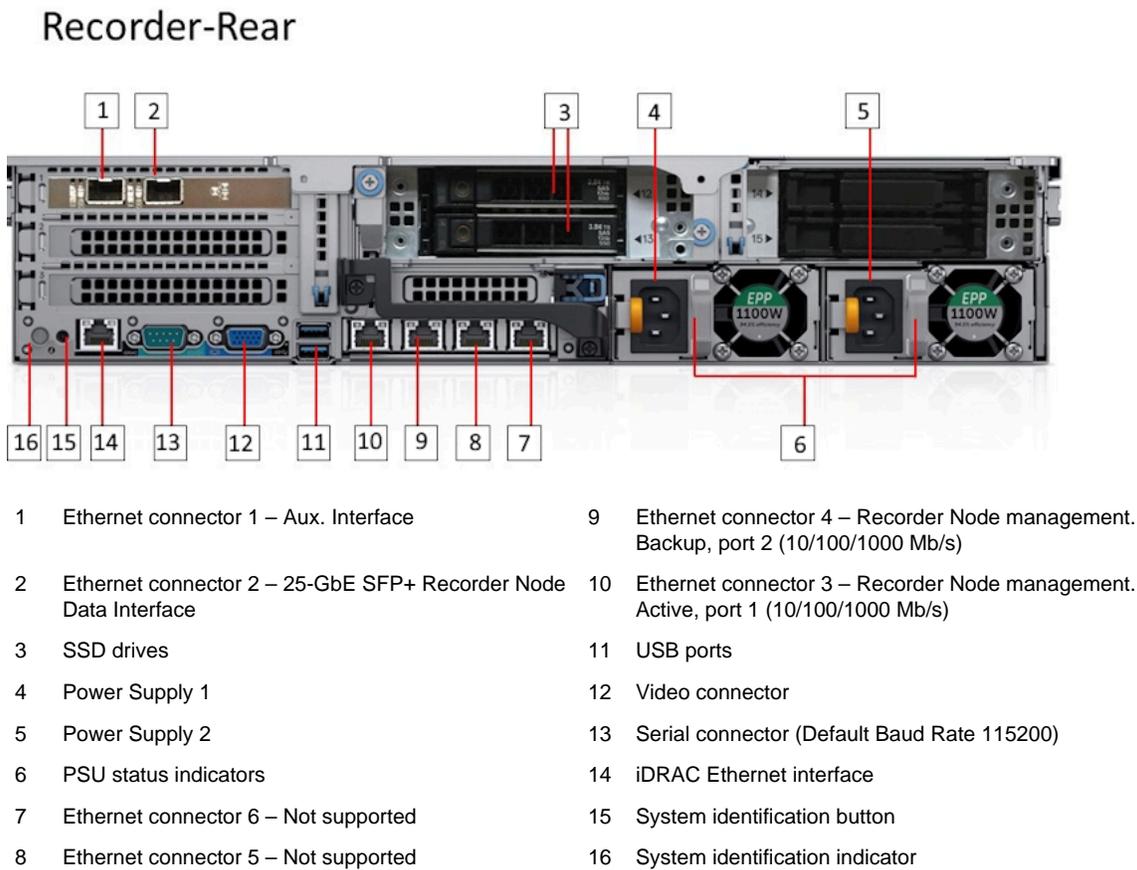
Recorder-Front



- | | |
|---|---------------------------|
| 1 | Information Tag |
| 2 | Video connector |
| 3 | Micro USB (not supported) |
| 4 | Hard drives |

The following figure illustrates the rear panel of the DMF Recorder Node.

Figure 8-3: DMF Recorder Node (HWA) Rear Panel



8.2 DMF Recorder Installation Procedure

Prerequisite: To install the Recorder Node (RN) software on a Dell server, complete the following steps.



Note: The appliance interfaces are on the back of the device.

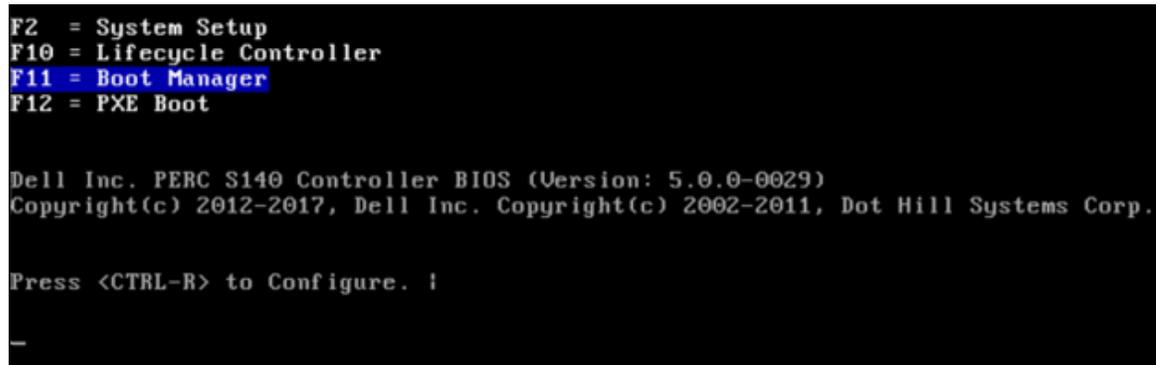


Note: System boot mode must be set to UEFI (default). Boot mode BIOS is not supported.

1. Rack the RN Appliance.
2. Connect the RN management interface port 1 to the management network.
3. Log in via the serial port or SSH using the admin account name. The baud rate is **115200**.
4. Insert a bootable USB drive in the RN USB port.
Refer to Appendix [Creating a USB Boot Image](#) to make a bootable USB drive.

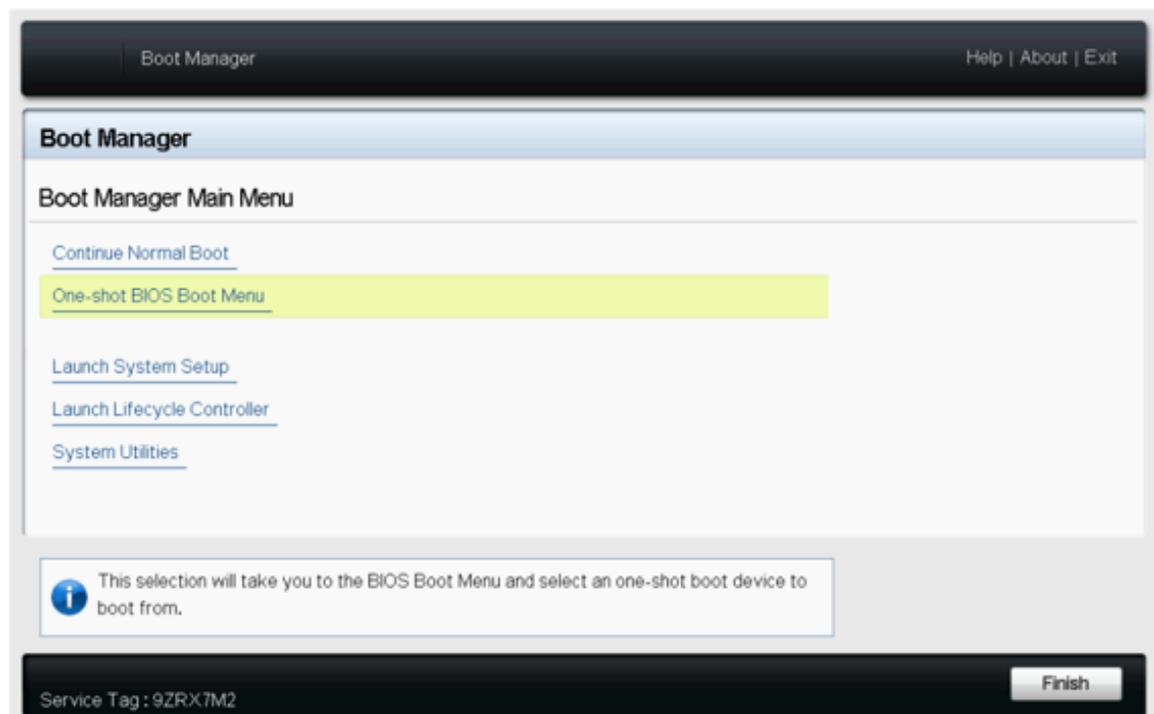
5. Power cycle the appliance.
6. Press **F11** to select the Boot Manager to allow booting from USB.

Figure 8-4: System Boot Manager Screen



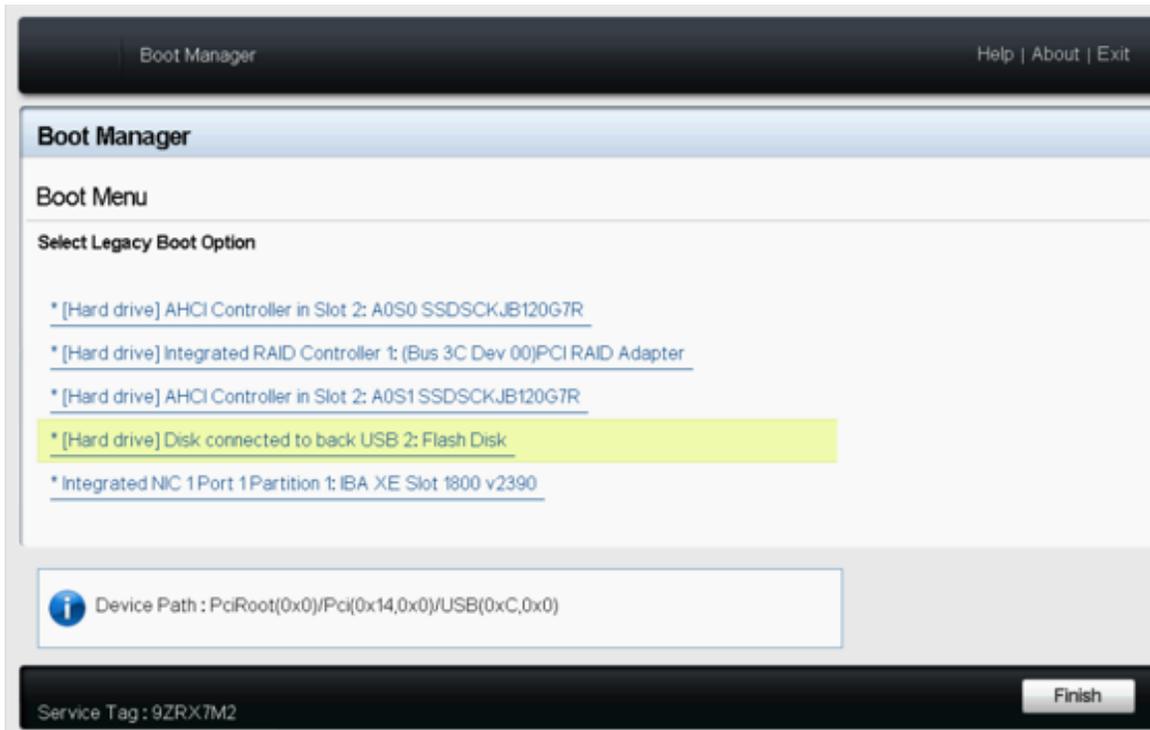
7. Select One-shot BIOS Boot Menu.
The Boot Manager screen is displayed in the following figure.

Figure 8-5: Boot Manager Main Menu



8. Select the USB drive.

Figure 8-6: Boot Menu



9. Respond to the system prompt to login in using the admin account:

```
recorder-node login: admin
(Press Control-C at any time to cancel and start over)
This product is governed by an End User License Agreement (EULA).
You must accept this EULA to continue using this product.
You can view this EULA from our website at:
https://www.arista.com/en/eula
Do you accept the EULA for this product? (Yes/No) [Yes] >
```

10. Type **yes** to accept the EULA, which is required to use the product. To view the EULA, type **view**, or refer to <https://www.arista.com/en/eula>.

The system displays the following messages.

```
Running system pre-check
Finished system pre-check
Starting first-time setup
```

11. Configure the recovery password.

```
Emergency recovery user password >
Emergency recovery user password (retype to confirm) >
Hostname > dmf-pr-740
```

12. Configure IP addresses for the management network and DNS servers.

```
[1] IPv4 only
[2] IPv6 only
[3] IPv4 and IPv6
> 1
IPv4 address [0.0.0.0/0] > 10.9.32.21/24
IPv4 gateway (Optional) > 10.9.32.1
DNS server 1 (Optional) > 10.3.0.4
DNS server 2 (Optional) >
DNS search domain (Optional) > qa.arista.com
Administrator password >
```

```
Administrator password (retype to confirm) >
Controller address if deployment mode is preconfigured (L3 ZTN) (Optional) > 10.111.35.101
```

13. If the RN is connected to the DMF Controller by a Layer 3 device (such as a router) in preconfigured (L3 ZTN) mode, enter the active DMF Controller's IP address.
14. Configure the administrator password.

```
Administrator password >
Administrator password (retype to confirm) >
```

15. Configure the NTP servers.

```
-----
Default NTP servers:
- 0.bigswitch.pool.ntp.org
- 1.bigswitch.pool.ntp.org
- 2.bigswitch.pool.ntp.org
- 3.bigswitch.pool.ntp.org
NTP server options:
[1] Use default NTP servers
[2] Use custom NTP servers
[1] > 1
```

16. Confirm the settings.

```
Please choose an option:
[ 1] Apply settings
[ 2] Reset and start over
[ 3] Update Recovery Password (*****)
[ 4] Update Hostname (dmf-pr-740)
[ 5] Update IP Option (IPv4 only)
[ 6] Update IPv4 Address (10.9.32.21/24)
[ 7] Update IPv4 Gateway (10.9.32.1)
[ 8] Update DNS Server 1 (10.3.0.4)
[ 9] Update DNS Server 2 (<none>)
[10] Update DNS Search Domain (qa.arista.com)
[11] Update Admin Password (*****)
[12] Update NTP Option (Use default NTP servers)
[1] >
The system displays the following messages.
[Stage 1] Initializing system
[Stage 2] Configuring local node
Waiting for network configuration
IP address on bond0 is 10.9.32.21
Generating cryptographic keys
[Stage 3] Configuring system time
Initializing the system time by polling the NTP servers:
0.bigswitch.pool.ntp.org
1.bigswitch.pool.ntp.org
2.bigswitch.pool.ntp.org
3.bigswitch.pool.ntp.org
[Stage 4] Configuring cluster
Cluster is already configured
First-time setup is complete!
```

17. Press **Enter** to complete the configuration.

8.3 Initial Configuration - GUI

After completing the installation, refer to the *DANZ Monitoring Fabric User Guide* to configure and operate the Recorder.

GUI Procedure

Complete the following steps to use the DANZ Monitoring Fabric (DMF) GUI to configure the Recorder Node (RN).

1. Select **Monitoring > Recorder Nodes** from the main menu and click the **Provision control (+)** icon.

Figure 8-7: Add Recorder Node



2. Enter the required details by specifying a **Name** and identifying the **MAC address** of the RN appliance NIC connected to DMF.



Tip: Choose the MAC address from the selection list if it has been discovered.

Figure 8-8: Provision Recorder Node

1. Info ✓ Name *

2. Indexing ✓

3. Network ✓ MAC Address
No Unconfigured MACs Found
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.

4. Storage ✓

Recording
Off On

Disk Full Policy
Rolling FIFO Stop and Wait ?

Max Packet Age
 minutes
The maximum age of a recorded packet in minutes. Packets older than this age will be deleted automatically.

Max Disk Utilization
 %
The maximum allowed utilization on the index and packet disks. The packet removal policy will be enforced at this limit.

Backup Disk Policy *

Pre Buffer
 minutes
Duration to record into a pre-buffer until an event occurs

Parse MetaWatch Trailer *
 ?

BACK NEXT RESET CANCEL SAVE

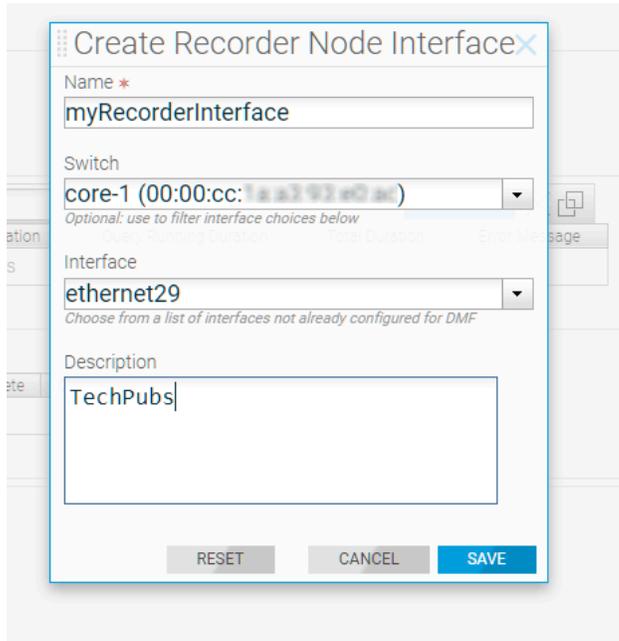
3. Click **Save**.

- Click the **Provision control (+)** at the top of the Interfaces section and enter the required information.

Figure 8-9: Add Interface



Figure 8-10: Provision Recorder Node



- Type an identifying **Name** (required) for the RN interface.
- Select the **Switch** and **Interface** to use to record the received traffic.
- Click **Save**.

8.4 Initial Configuration - CLI

CLI Procedure

To use the DMF CLI to perform the basic Recorder Node (RN) configuration, complete the following steps.

- Assign a name to the RN.

```
(config)# recorder-node device bt-recorder3
```

- Set the MAC address of the RN.

```
controller-1(config-recorder-node)# mac 18:66:da:fb:6d:b4
```

If the management MAC is unknown, you can determine it from the chassis ID of connected devices using the **show connected-devices** command.



Note: The following output is truncated and edited for documentation purposes.

```
controller-1> show connected-devices
# Switch      IF Name      SPAN? Device Name  Device Description  Chassis ID
-----|-----|-----|-----|-----|-----|
1 filter-1    ethernet1    False localhost    Arista Networks EOS 2c:dd:e9:37:bf:47
2 delivery-2  ethernet1    False localhost    Arista Networks EOS 2c:dd:e9:37:bf:47
3 delivery-2  ethernet43   False leaf1a     5c:16:c7:00:00:01   70:72:cf:c6:fe:f1
4 delivery-2  ethernet48   False qa-ibm-1   IBM NOS             74:99:75:69:f7:00
5 delivery-1  ethernet1    False leaf1a     5c:16:c7:00:00:01   70:72:cf:c6:fe:f1
6 delivery-1  ethernet2    False leaf1a     5c:16:c7:00:00:01   70:72:cf:c6:fe:f1
7 delivery-1  ethernet3    False leaf2a     5c:16:c7:00:00:01   70:72:cf:b5:e4:c0
8 delivery-1  ethernet4    False leaf2a     5c:16:c7:00:00:01   70:72:cf:b5:e4:c0
9 delivery-1  ethernet25   False dmf-rn-1    dmf-recorder-node   18:66:da:fb:6d:b4
```

3. The RN is *enabled by default*. If you ever need to enable it, you can use the following command.

```
controller-1(config-recorder-node)# record
```

4. Define the RN interface name.

```
controller-1(config)# recorder-fabric interface pr-intf-1
controller-1(config-recorder-fabric-interface)#
```

Assign any alphanumeric identifier for the name of the RN interface, which changes the submode to *config-recorder-fabric-interface*.

5. Assign a switch and interface and optionally provide a text description.

```
controller-1(config-recorder-fabric-interface)# recorder-interface switch delivery-1 ethernet25
controller-1(config-recorder-fabric-interface)# description 'Delivery point for recorder node'
controller-1(config-recorder-fabric-interface)#
```

6. Identify the RN interface by name in an out-of-band policy.

```
controller-1(config)# policy pkt-rec
controller-1(config-policy)# use-recorder-node pr-intf-1
```

7. Configure the DMF policy to identify the traffic to send to the RN.

```
controller-1(config-policy)# 1 match any
controller-1(config-policy)# filter-interface sw1-fill
controller-1(config-policy)# action forward
```

The following example forwards all traffic received in the monitoring fabric on filter interface *sw1-fill1* to the RN interface.

```
recorder-node device bt-recorder3
  mac 18:66:da:fb:6d:b4
recorder-fabric interface pr-intf-1
  description 'Delivery point for recorder node'
  recorder-interface switch delivery-1 ethernet25
policy pkt-rec
  action forward
  filter-interface sw1-fill1
  use-recorder-node pr-intf-1
  1 match any
```

8.5 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. DMF supports overriding the default configuration for an RN from the `config-recorder-node` submode for any RN.



Note: Currently, these options are available only from the CLI and not in the DMF 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 for the Recorder Node. This alias is associated with the MAC hardware address using the `mac` command.

Use any of the following commands from `config-recorder-node` submode to override the default configuration for the associated Recorder node.

- **banner:** Set recorder-node pre-login banner message
- **mac:** Configure MAC address for recorder-node name

Additionally, the following configurations can be overridden to use values specific to the recorder node or used in `merge` mode along with the configuration inherited from the Controller.

- **ntp:** Configure packet-recorder to override default timezone and NTP parameters
- **snmp-server:** Configure packet-recorder SNMP parameters and traps
- **logging:** Enable packet-recorder logging to the Controller
- **tacacs:** Set TACACS defaults, server IP address(es), timeouts and keys

To configure the recorder node to override the configuration inherited from the Controller, execute the following commands at the `config-recorder-node` submode:

- **ntp override-global:** Override global time config with packet-recorder time config
- **snmp-server override-global:** Override global SNMP config with packet-recorder SNMP config
- **snmp-server trap override-global:** Override global SNMP trap config with packet-recorder SNMP trap config
- **logging override-global:** Override global logging config with packet-recorder logging config
- **tacacs override-global:** Override global TACACS+ config with packet-recorder TACACS+ config

To configure the recorder node to work in a merge mode by merging its specific configuration with that of the Controller, execute the following commands at the `config-recorder-node` submode:

- **ntp merge-global:** Merge global time config with packet-recorder time config
- **snmp-server merge-global:** Merge global SNMP config with packet-recorder SNMP config
- **snmp-server trap merge-global:** Merge global SNMP trap config with packet-recorder SNMP trap config
- **logging merge-global:** Merge global logging config with -packet-recorder logging config

The TACACS+ configuration does not provide a command usable with the *merge* option: it can be inherited from the Controller or overridden to use only the recorder node-specific configuration.

DHCPv4-Based First-boot for DMF Controller and Managed Appliances

This chapter outlines a solution for provisioning DANZ Monitoring Fabric (DMF) appliances via PXE and automating the configuration of the first boot parameters using Ansible.

9.1 Introduction

Typically, the deployment of DANZ Monitoring Fabric (DMF) on supported hardware appliances involves two steps:

- Installing an appropriate image.
- Configuration of firstboot parameters such as IP address (DHCP/Static), DNS and NTP server address, admin password(s), cluster information etc.

In this context, **Pave** refers to the automation of the first-time installation of a DMF hardware appliance. This involves installing a DMF image on a hardware appliance and completing the first-boot configuration. First-boot configuration uses an Ansible playbook as an automation tool. In contrast, **Repave** refers to automating the **re-installation** of DMF images on supported DMF appliances. This involves the automated process of re-installing a DMF image on a DMF hardware appliance and completing the first-boot configuration.

Several prerequisites are required. The production/lab environment should have:

- A DHCP server that supports the configuration of DHCP options 66 and 67.
- A TFTP server that can serve the *bootloader* and the corresponding configuration.
- An NFS server to serve the net-bootable appliance image.
- A server with ansible-playbook and Arista-supported playbook modules installed.
- Supported DMF hardware appliances with preset boot order settings and a PXE-enabled management port.



Note: In the case of Repave action, this new feature introduces a new command (`boot pxe`) to change the boot order to PXE boot. On reboot, the appliance will automatically perform image re-installation.

9.2 Prepare Services (TFTP/NFS) for PAVE/REPAVE Operation

The following steps must be completed before deploying the DANZ Monitoring Fabric (DMF) ISO image on a TFTP server.

Create a directory by name images on the TFTP server and copy the ISO image to the TFTP server.

1. SSH to the server with `sudo` privileges and create a temporary directory using the command below.

```
$ mktemp -d /tmp/tmp.2syaj0amL7
```

2. Mount the ISO image to the directory created above.

```
$ mount /images/*.iso /tmp/tmp.2syaj0amL7
```

3. Copy the following files to the root TFTPboot directory.

```
$ cp /usr/lib/PXELINUX/pxelinux.0 /var/lib/tftpboot
$ cp /usr/lib/syslinux/modules/bios/ldlinux.c32 /var/lib/tftpboot
```

If the above files are unavailable on the system, obtain them via an `apt-get` or `yum` command.

```
$ apt-get install pxelinux
```

4. Update the `TFTP_DIRECTORY` variable with the parent directory created to store `bootloader` files.

```
$ sed -i "/^TFTP_DIRECTORY=/c\TFTP_DIRECTORY=/var/lib/tftpboot" /etc/default/tftpd-hpa
```

5. Create an appropriate folder under the TFTP root directory for each DMF appliance type.

```
$ mkdir /var/lib/tftpboot/dmf-controller/
$ mkdir /var/lib/tftpboot/dmf-service-node/
$ mkdir /var/lib/tftpboot/dmf-analytics-node/
$ mkdir /var/lib/tftpboot/dmf-recorder-node/
```

6. Create an appropriate folder on the NFS root directory for each DMF appliance type.

```
$ mkdir path_to_NFS_root_directory/dmf-controller/
$ mkdir path_to_NFS_root_directory/dmf-service-node/
$ mkdir path_to_NFS_root_directory/dmf-analytics-node/
$ mkdir path_to_NFS_root_directory/dmf-recorder-node/
```

7. Copy the files from the folder mounted earlier in Step 3.

```
$ cp "/tmp/tmp.2syaj0amL7/casper/vmlinuz" "/var/lib/tftpboot/dmf-controller"
$ cp "/tmp/tmp.2syaj0amL7/casper/initrd.lz" "/var/lib/tftpboot/dmf-controller"
$ cp -r "$/tmp/tmp.2syaj0amL7/." "path_to_NFS_root_directory/dmf-controller"
```

8. Update the ownership of the `tftp` parent directory with TFTP user account and restart the `tftp` service.

```
$ chown -R tftp:tftp /var/lib/tftpboot
$ systemctl restart tftpd-hpa
```

9. Copy kernel configurations that assist the DMF appliance booting with UEFI mode to locate the `bootloader` and `vmlinuz` files from the TFTP and NFS server. Arista recommends defining a name for menu entry so that it is easily remembered and prefixed with appliance type for differentiation, i.e., `DCA-DM-CDL-dmf-8.4.0`.

Note: In the currently supported UEFI boot mode, Arista recommends creating a PXE configuration file for each hardware MAC address, i.e., instead of the default `grub.cfg` use `grub.cfg-01-xx-xx-xx-xx-xx-xx`. Specify the MAC address in all lower-case with `:` (colon) replaced by `-` (dash).



To perform the action detailed in the note above, obtain the MAC / Hardware address of the management interfaces. Obtaining the MAC address can be done in two ways, as described below:

- a. If this is the first time the appliance is being installed, either:

- Use the Integrated Dell Remote Access Controller (iDRAC) web console to determine the MAC address of the management interface.

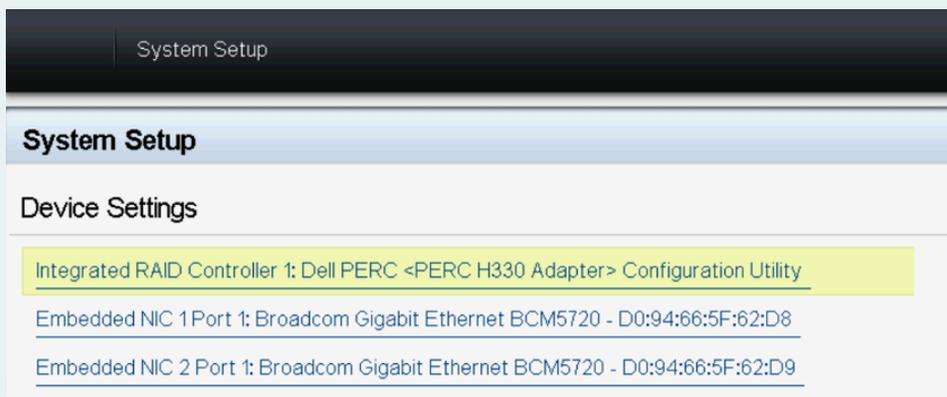
Figure 9-1: Using the iDRAC menu to obtain the Management Interface MAC Address

The screenshot shows the iDRAC web console interface. At the top, there is a navigation bar with 'Network Devices' selected. Below it, there are tabs for 'Summary', 'Embedded NIC 1', 'NIC Slot 2', and 'NIC Slot 3'. The main content area displays information for 'Embedded NIC 1: Broadcom Gigabit Ethernet BCM5720 - D0:94:66:5F:62:D8'. Under 'Port Properties', it lists 'Product Name' as 'Broadcom Gigabit Ethernet BCM5720 - D0:94:66:5F:62:D8', 'Vendor Name' as 'Broadcom Corp', and 'Number of Ports' as '2'. At the bottom, there is a table titled 'Ports and Partitioned Ports' with columns for 'Link Status', 'Port', 'Partition', 'Protocol', and 'Switch Connection ID'.

Link Status	Port	Partition	Protocol	Switch Connection ID
+ Up	1	Not Capable	NIC	Not Supported
+ Up	2	Not Capable	NIC	Not Supported

- Enter the device BIOS menu during boot-up and determine the MAC address of the management interface.

Figure 9-2: Using the BIOS menu to obtain the Management Interface MAC Address



- b. When performing a re-installation of DMF appliances, obtain the MAC address of the appliance via the CLI command.
- For DMF Controllers with SKU DCA-DM-C450 or DMF Analytics Nodes with SKU DCA-DM-AN450, execute the following commands on the Controller to obtain the relevant MAC address.

```
DCA-DM-C450# show local-node interfaces eno8303
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|-----|
eno8303 bond0 d0:8e:79:d4:1e:56 (Dell) d0:8e:79:d4:1e:56 (Dell) up up active
~ Address ~
None.
DCA-DM-C450# show local-node interfaces eno8403
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstat Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|-----|
eno8403 d0:8e:79:d4:1e:57 (Dell) d0:8e:79:d4:1e:57 (Dell) down down
~ Address ~
None.
DCA-DM-C450#
```

- For DMF Controllers with SKU DCA-DM-CDL, including all DMF Recorder Nodes, execute the following commands on the device to obtain the relevant MAC address.

```
DCA-DM-CDL# show local-node interfaces eno1
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|-----|
```

```

eno1      bond0  d0:94:66:21:b9:45 (Dell)  d0:94:66:21:b9:45 (Dell)  up      up      active
~ Address ~
None.
DCA-DM-CDL# show local-node interfaces eno2
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|-----|
eno2      bond0  d0:94:66:21:b9:45 (Dell)  d0:94:66:21:b9:46 (Dell)  down    down    down    backup
~ Address ~
None.
DCA-DM-CDL#

```

- For all DMF Service Nodes, execute the following commands on the device to obtain the relevant MAC address.

```

dmf-service-node# show local-node interfaces eno3
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|
eno3      bond0  78:ac:44:8a:59:52 (Dell)  78:ac:44:8a:59:52 (Dell)  up      up      up      active
~ Address ~
None.
dmf-service-node# show local-node interfaces eno4
----- Interfaces -----
Interface Master Hardware address Permanent hardware address Operstate Carrier Bond mode Bond role
-----|-----|-----|-----|-----|-----|
eno4      bond0  78:ac:44:8a:59:53 (Dell)  78:ac:44:8a:59:53 (Dell)  down    down    down
~ Address ~
None.
dmf-service-node#

```

- If both the management ports of the appliance are connected to the top of the rack management switch, Arista recommends generating a separate PXE configuration file for both management ports.

```

cat << EOF | sudo tee /var/lib/tftboot/boot/grub/grub.cfg-01-0a-0b-0c-0d-0e-0f
set default="0"
loadfont unicode
set gfxmode=auto
insmod all_video
insmod gfxterm
serial --unit=0 --speed=115200
serial --unit=1 --speed=115200
terminal_input console
terminal_input --append serial_com0
terminal_input --append serial_com1
terminal_output gfxterm
terminal_output --append serial_com0
terminal_output --append serial_com1
set timeout=5
menueentry "Install <INSERT_USER_FRIENDLY_NAME>" {
linux dmfc-controller/vmlinuz boot=casper netboot=nfs nfsroot=<ip-address of PXE server>
: /srv/install/dmf-controller toram noprompt ip=dhcp -- unattended_installation autofirstboot
}
via_ssh
initrd dmfc-controller/initrd.lz
}
EOF

```

9.3 DHCPv4-Based First-boot

- DHCPv4-based first-boot for DMF Controller and Managed Appliances

There are two main steps involved in accomplishing the automatic installation of an image on a supported DANZ Monitoring Fabric (DMF) hardware appliance and the completion of the first-boot configuration.

- [Auto-installation of images using PXE](#)
- [Auto-configuration of firstboot parameters to complete DMF HW appliance installation](#)

9.3.1 Auto Installation of Images

Auto-installation of a DMF image uses well-known services like DHCP, TFTP, NFS, and PXE. DMF images are PXE bootable, and using the services above, perform the auto-installation of images on the DMF appliances.

A high-level procedure for auto-installation is described below.

- Configure a DHCP server to provide a DHCP IP address.
- Enter the `next-server` IP address (This is the TFTP server IP address specified by **Option 66** configuration on the DHCP server).

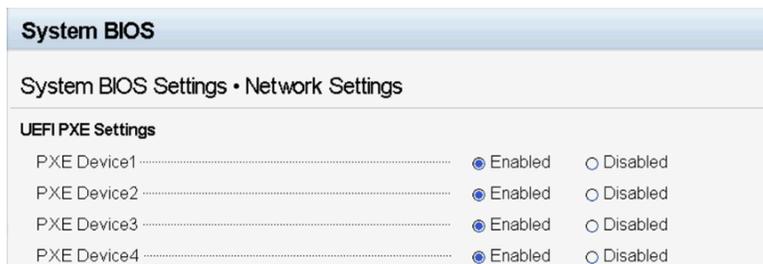
```
next-server <TFTP_SERVER_IP>;
class "pxeclient" {
match if substring (option vendor-class-identifier, 0, 9) = "PXEClient";
filename "boot/grub/x86_64-efi/core.efi"; # x86 EFI
}
```

- Arista recommends binding a static IP to the hardware MAC address of the DMF appliance. The following is a sample configuration using a Linux-based DHCP service.

```
group {
host <DMF_APPLIANCE_HOSTNAME> {
hardware ethernet <DMF_APPLIANCE_MAC>;
fixed-address <DESIRED_IP_ADDRESS>;
}
```

- On a DMF HW appliance using PXE boot, **enable** PXE on the management interface NIC card.

Figure 9-3: Enable interfaces in BIOS to be PXE bootable



- For the initial PAVE action, press the **F12** key during the initial boot process to manually trigger the PXE boot. This step is typically not required when reinstalling (REPAVE) the DMF appliance images.

Figure 9-4: Manual PXE Boot by pressing F12



- Power cycle the DMF HW appliance.

- The DHCP client on the NIC card sends a DHCP discovery request and gets an IP address and next-server (TFTP server) IP address.
- The management NIC gets the bootloader via PXE boot using a TFTP server. A bootloader config file with a filename based on the MAC address of the appliance is configured and saved on the DHCP server.
- The appliance obtains configuration parameters like the NFS mount, where appliance ISO images are stored.
- The appliance boots from the boot loader and obtains the DMF appliance ISO image from the NFS server.
- The appliance boots the ISO image and starts the installation process without user input.
- On reboot, the appliance again acquires a DHCP-based management IP address, sets up initial default user credentials, and waits for auto first-boot configuration via Ansible.



Note: For repave, the procedure is the same except that the DMF appliance is already running a DMF image and needs to boot from PXE again for re-installation.

9.3.2 Auto Configuration of First-boot

Auto-configuring of first-boot parameters uses Ansible and is accomplished by interacting with `auto-firstboot-cloud-plugin`, which is available in hardware appliances.

- Begin the initial configuration playbook in Ansible. The following is a sample YAML-based Ansible playbook file.

```
- name: Test Autofirstboot Properties Provider
gather_facts: false
hosts: Controllers
connection: "local"
run_once: true
tasks:
- name: Provide Autofirstboot Properties to Cluster
  arista.dmf.provisioner:
  config_json: "<Pave-Repave-config.json>"
  timeout: 3600
  log_dir: "logs"
```

- The configuration playbook obtains the JSON-based configuration file from the server executing the playbook. The JSON file contains specific sections for each DMF appliance that is being installed/re-installed. The following is an example:

```
{
  "<ACTIVE-CONTROLLER-IP>": {
    "initial-admin-password": "<ACTIVE-CONTROLLER-CLEAR-TEXT-PASSWD>",
    "initial-config-password": "pxe_temp_password",
    "dhcp-ip": "10.240.156.82",
    "appliance-type": "BMF",
    "firstboot-properties": {
      "admin-password": "<ACTIVE-CONTROLLER-CLEAR-TEXT-PASSWD>",
      "recovery-password": "<ACTIVE-CONTROLLER-CLEAR-TEXT-RECOVERY-PASSWD>",
      "hostname": "<ACTIVE-CONTROLLER-HOSTNAME>",
      "cluster-name": "<CLUSTER-NAME>",
      "cluster-description": "<CLUSTER-DESCRIPTION>",
      "ip-stack": "ipv4",
      "ntp-servers": ["<CLUSTER-NTP-SERVER-1>", "<CLUSTER-NTP-SERVER-2>"],
      "dns-servers": ["<CLUSTER-DNS-SERVER-1>", "<CLUSTER-DNS-SERVER-2>"]
    }
  },
  "<STANDBY-CONTROLLER-IP>": {
    "initial-config-password": "pxe_temp_password",
    "dhcp-ip": "<STANDBY-CONTROLLER-DHCP-IP>",
    "appliance-type": "BMF",
    "firstboot-properties": {
      "admin-password": "<STANDBY-CONTROLLER-CLEAR-TEXT-PASSWD>",
      "recovery-password": "<STANDBY-CONTROLLER-CLEAR-TEXT-RECOVERY-PASSWD>",
      "hostname": "<STANDBY-CONTROLLER-HOSTNAME>",

```

```

"cluster-name": "<CLUSTER-NAME>",
"cluster-to-join": "<ACTIVE-CONTROLLER-IP>",
"ip-stack": "ipv4"
},
"<MANAGED-APPLIANCE-IP>": {
"initial-config-password": "pxe temp password",
"dhcp-ip": "<STANDBY-CONTROLLER-DHCP-IP>",
"appliance-type": "BMFSN",
"firstboot-properties": {
"admin-password": "<MANAGED-APPLIANCE-CLEAR-TEXT-PASSWD>",
"recovery-password": "<MANAGED-APPLIANCE-CLEAR-TEXT-RECOVERY-PASSWD>",
"hostname": "<MANAGED-APPLIANCE-HOSTNAME>",
"ip-stack": "ipv4",
"controller_ip": "<ACTIVE-CONTROLLER-IP>",
"ntp-servers": ["<NTP-SERVER-1>","<NTP-SERVER-2>"],
"dns-servers": ["<DNS-SERVER-1>","<DNS-SERVER-2>"]
}
}

```



Note: The ansible script uses the *initial-config-password* to access the DMF appliance via SSH for the first time.

- Ansible pushes initial configuration data to the appliance via SSH using the default credential configured in the JSON file in the steps above.
- The appliance completes the initial configuration (first-boot).
- The appliance confirms to Ansible that the initial configuration is complete.

The following is an example run of the Ansible playbook:

```

ansible-playbook sample_playbook.yml --limit="10.240.156.82,10.240.156.84" -v
[DEPRECATION WARNING]: Ansible will require Python 3.8 or newer on the controller starting
with Ansible 2.12. Current version: 2.7.18 (default, Jul 1 2022, 12:27:04) [GCC 9.4.0].
This feature will be removed from ansible-core in version 2.12. Deprecation warnings can be
disabled by setting deprecation_warnings=False in ansible.cfg.
/root/.cache/pypoetry/virtualenvs/arista-dmf-vQFV2Q4_py2.7/lib/python2.7/site-packages/
ansible/parsing/vault/__init__.py:44: CryptographyDeprecationWarning: Python 2 is no longer
supported by the Python core team. Support for it is now deprecated in cryptography, and
will be removed in the next release.
from cryptography.exceptions import InvalidSignature
Using /etc/ansible/ansible.cfg as config file
PLAY [Test Autofirstboot Properties Provider]
*****TASK [Provide Autofirstboot Properties to Cluster]
*****Targeting all hosts specified in config [u'10.240.156.84', u'10.240.156.82']
[10.240.156.82 BMF (active)]: Waiting for SSH connection.
[10.240.156.84 BMF (standby)]: Waiting for SSH connection.
[10.240.156.84 BMF (standby)]: Established SSH connection.
[WARNING]: It is NOT recommended to disable SSL verification. Please upload a valid SSL
certificate to the https://10.240.156.84:8443 API server.
[10.240.156.82 BMF (active)]: Established SSH connection.
[WARNING]: It is NOT recommended to disable SSL verification. Please upload a valid SSL
certificate to the https://10.240.156.82:8443 API server.
[WARNING]: [10.240.156.82 BMF (active)]: Detected BIOS firmware. The BIOS PXE boot trigger has
been found to be unreliable. Upgrade to UEFI is recommended.
[WARNING]: [10.240.156.84 BMF (standby)]: Appliance already configured. Attempting PXE boot to
repave.
[WARNING]: [10.240.156.82 BMF (active)]: Appliance already configured. Attempting PXE boot to
repave.
[10.240.156.84 BMF (standby)]: Expecting appliance to reboot with IP = 10.240.156.84.
[10.240.156.82 BMF (active)]: Expecting appliance to reboot with IP = 10.240.156.82.
[10.240.156.84 BMF (standby)]: Verified repave succeeded.
[10.240.156.84 BMF (standby)]: Waiting for active to be configured.
[10.240.156.82 BMF (active)]: Verified repave succeeded.
[10.240.156.82 BMF (active)]: Successfully provided config to appliance.
[10.240.156.82 BMF (active)]: Waiting for appliance to apply config. Expecting final IP = 10.
240.156.82.
[10.240.156.82 BMF (active)]: Successfully retrieved firstboot logs.
[10.240.156.82 BMF (active)]: Config applied.
[10.240.156.82 BMF (active)]: Attempting to reset recovery password.
[10.240.156.82 BMF (active)]: Successfully reset recovery password.
[10.240.156.84 BMF (standby)]: Active configured. Proceeding with configuration.
[10.240.156.84 BMF (standby)]: Successfully provided config to appliance.
[10.240.156.84 BMF (standby)]: Waiting for appliance to apply config. Expecting final IP = 10.

```

```
240.156.84.
[10.240.156.84 BMF (standby)]: Successfully retrieved firstboot logs.
[10.240.156.84 BMF (standby)]: Config applied.
[10.240.156.84 BMF (standby)]: Attempting to reset recovery password.
[10.240.156.84 BMF (standby)]: Successfully reset recovery password.
changed: [10.240.156.82] => {"changed": true, "success": true, "summary": {"10.240.156.82": {
"appliance-type": "BMF", "expected-dhcp-ip": "10.240.156.82", "role": "active"}, "10.240.156.
84": {"appliance-type": "BMF", "expected-dhcp-ip": "10.240.156.84", "role": "standby"}}}
PLAY RECAP
*****56.82 : ok=1 changed=1 unreachable=0 failed=0 skipped=0 rescued=0 ignored=0
```

9.4 Assumption and Trust Model

PXE boot is fundamentally incompatible with zero-trust environments, and as a result, assumptions are made to establish the trust required to authenticate the appliances. This section provides a summary of the assumptions that underpin the security of this design.

The management network is trusted, such that:

- DHCP is secured.
 - The DHCP server (or DHCP relaying) is secure.
 - Rogue DHCP packets are dropped/blocked.
- Impersonation by MAC or IP spoofing is not possible due to either:
 - Assumption: Machines on the same L2 network are **trusted not to impersonate** other machines by presenting false identities (MAC addresses or IP addresses).
 - Guarantee: Machines in the same L2 network cannot impersonate other machines by presenting themselves with false identities as **enforced by network admins** who may, for example:
 - Pin a MAC address to a specific switch and switch port.
 - Pin an IP address to a MAC address statically (i.e., static ARP entry).
- Routers between the PXE TFTP/HTTP server and the target machine are secure/trusted to forward packets to the rightful owners (e.g., having correct routing tables and MAC address tables).
- The PXE (TFTP/NFS) server is secure and cannot be compromised, resulting in an attacker providing a malicious image.

Managing SNMP

This chapter describes how to manage SNMP services on a DANZ Monitoring Fabric (DMF) Controller.

10.1 SNMP Overview

SNMP provides a method for communication between an NMS or other client and agents (servers) on network devices, which send reports, called traps, regarding their operation and configuration. An SNMP agent manages and organizes the information as a collection of objects called MIBs.

In SNMPv3, an engineID identifies the agent (SNMP server), which helps prevent unauthorized SNMPv3 messages, such as traps, from being accepted or intercepted by unauthorized receivers. The engineID of the SNMP agent is required when configuring an SNMPv3 trap receiver to receive messages from an agent, including a DMF Controller or fabric switch.

In DMF, the engineID is auto-generated for the Controller and fabric switches. The engineID of the DMF Controller is configured for the local node. This configuration must be entered separately on the active and standby Controllers. The acceptable practice recommends configuring a different engineID for each Controller.

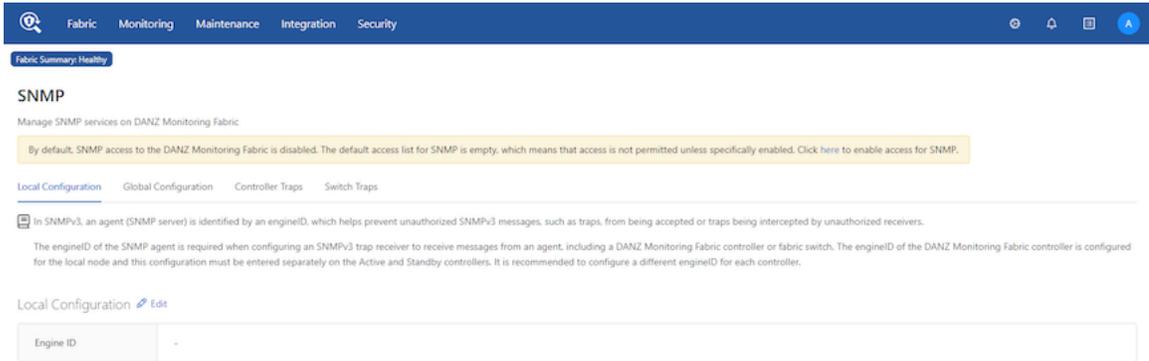
10.2 Using the DMF GUI to Configure SNMP

Complete the following steps to manage or view the DANZ Monitoring fabric (DMF) Controller SNMP configuration. SNMP configuration tabs and features include:

- Local Configuration
 - [Edit Local Configuration](#)
- Global Configuration
 - [Edit Global Configuration](#)
 - [Communities](#)
 - [Users](#)
- Controller Traps
 - [Edit Controller Traps](#)
- Switch Traps
 - [Edit Switch Traps](#)

1. Select **Maintenance > SNMP** from the DMF main menu.

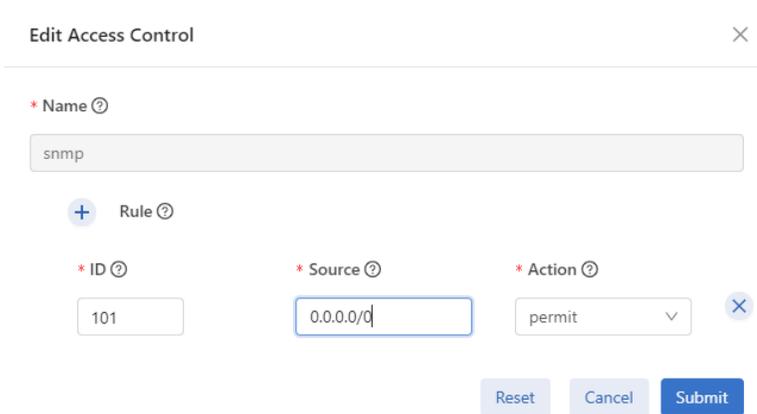
Figure 10-1: Configuring SNMP



 **Note:** By default, SNMP access is disabled.

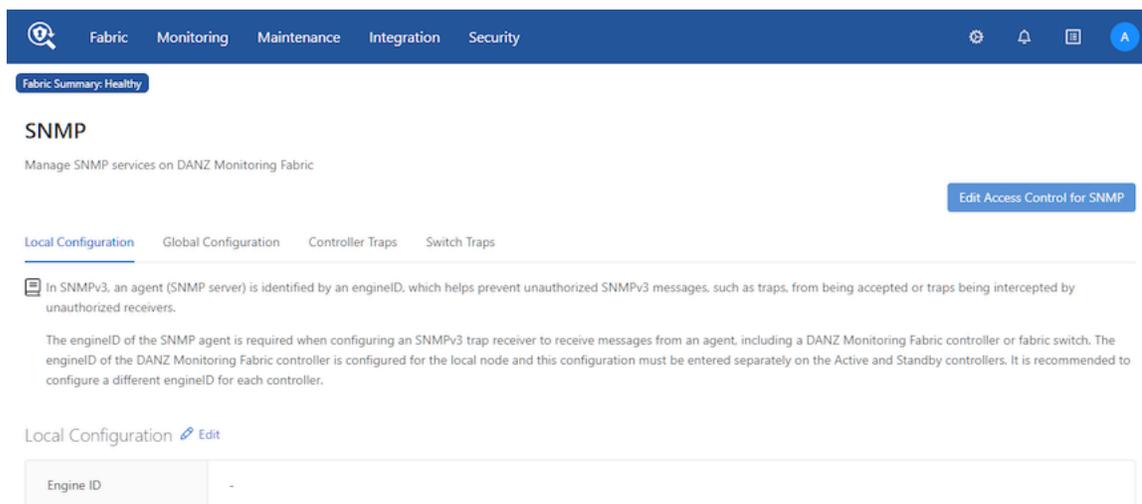
2. To enable access to SNMP for the Controller, click the link and enter the required fields, **ID**, **Source**, and **Action** (permit), in the **Edit Access Control** section.

Figure 10-2: Edit Access Control



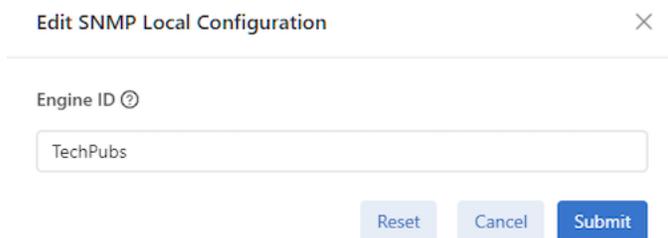
3. Click **Submit** to continue.

Figure 10-3: SNMP Enabled



4. Under **Local Configuration**, click **Edit** and enter an **Engine ID** value, as required.

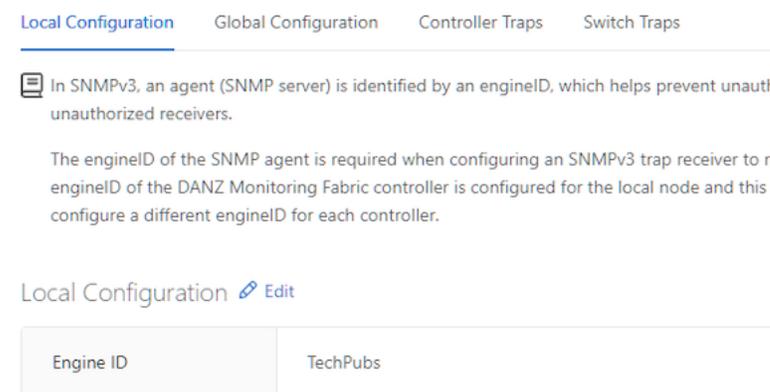
Figure 10-4: Edit SNMP Local Configuration



5. Click **Submit** to continue.

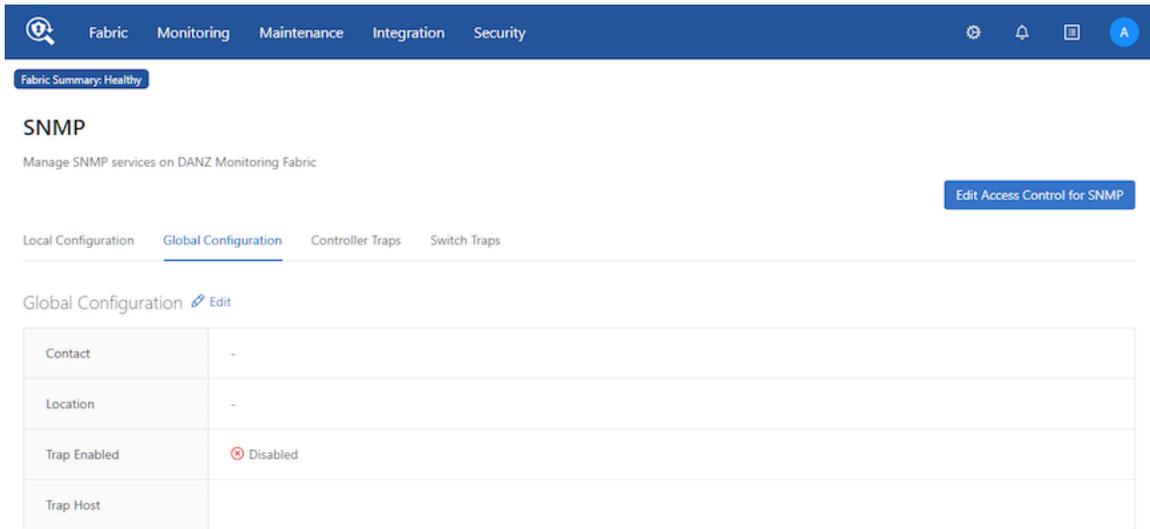
 **Tip:** Use the **Reset** button to clear the Engine ID value, if required.

Figure 10-5: Local Configuration



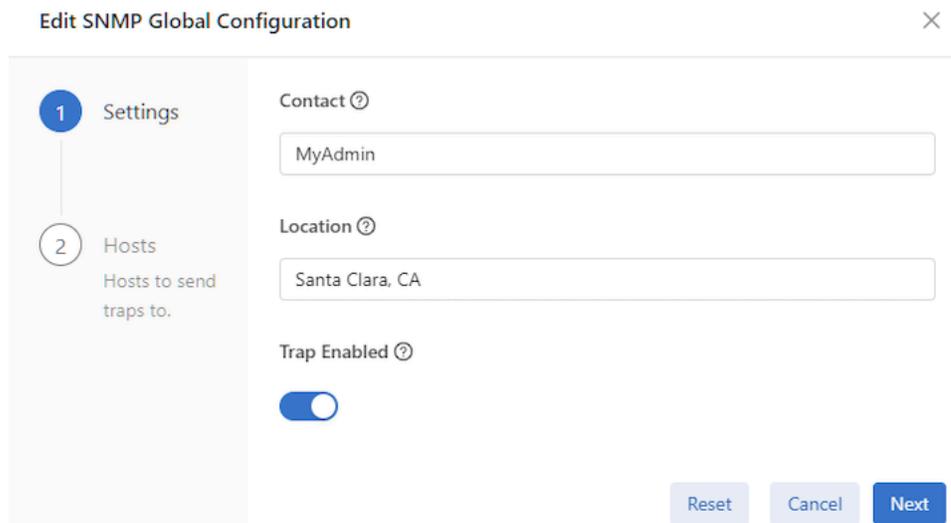
- To enable SNMP traps, select **Global Configuration**.

Figure 10-6: Global Configuration



- Click **Edit** and enter the **Contact** and **Location** details. Enable **Trap Enabled** by moving the selector switch to the right.

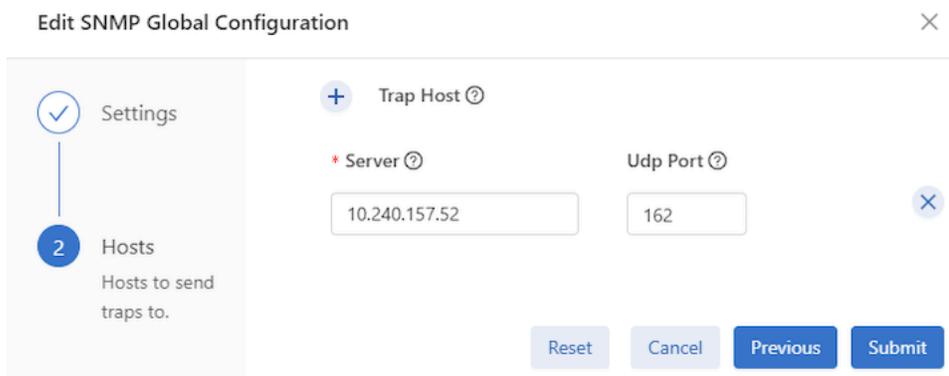
Figure 10-7: Edit SNMP Global Configuration



- Click **Next** to continue.

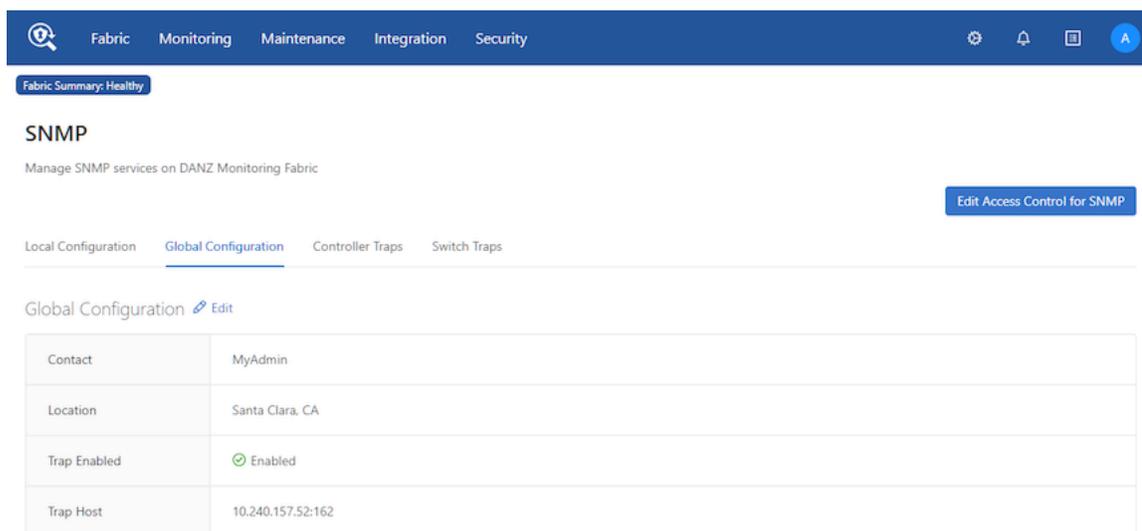
- Enter the **Trap Host** details for the **Server** and **UDP Port** (162 by default) using the **Provision control (+)** button.

Figure 10-8: SNMP Trap Host



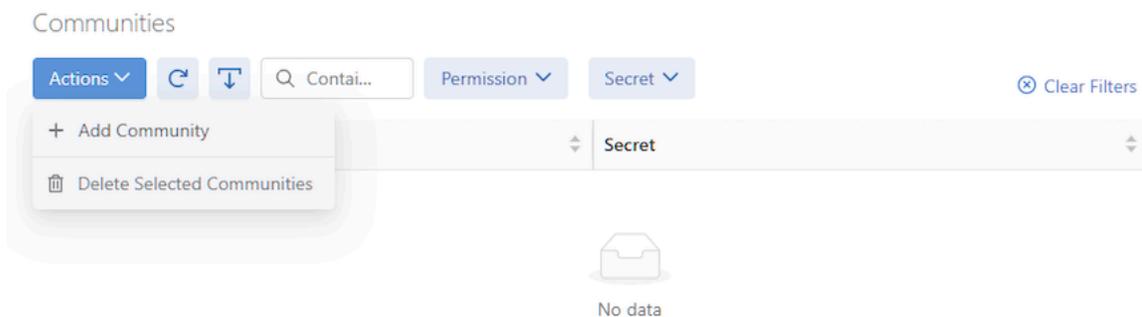
- Click **Submit**. The dashboard displays the information and confirms **Trap Enabled**.

Figure 10-9: Global Configuration Trap Enabled



- To create a new **Community**, select the **Actions** button under **Communities** and click **+ Add Community**.

Figure 10-10: Add Community



12. Select the **Permission** type (read-only) from the drop-down and enter the **Secret**.

Figure 10-11: Add Community Details

Add Community ✕

* Permission ⓘ

read-only

* Secret ⓘ

MySecret ✕ 👁

Requires at least 8 characters: 8/8

The Secret value will be encoded during submit.

Reset Cancel Submit

13. Click **Submit**—the dashboard updates with the Community details.

Figure 10-12: Communities

Communities

Actions 🔄 📄 Permission Secret

<input type="checkbox"/>	Actions	Permission	Secret
<input type="checkbox"/>	☰	read-only	013e1f375e08140a35

14. To create an SNMPv3 user, select the **Actions** button under **Users** and click **+ Add User**.

Figure 10-13: Add Users

Users

Actions 🔄 📄 Add Filters Name Clear Filters

+ Add User

🗑 Delete Selected Users

	Name	Auth Passphrase	Priv Passphrase	Priv Protocol
No data				

15. Enter the required information, such as the **Name** of the user, the **Authentication Passphrase** for the user, and the **Privacy Passphrase**. Use the **Privacy Protocol** drop-down to select Advanced Encryption

Standard (**AES**) or Data Encryption Standard (**DES**) encryption to encrypt the SNMP messages between the SNMP agent and the manager.

Figure 10-14: User Details and Encryption

The screenshot shows a form titled "Add User" with a close button (X) in the top right corner. The form contains the following fields:

- Name**: A text input field containing "MySNMPv3User".
- Auth Passphrase**: A text input field containing "aristarocks". Below the field, it says "Requires at least 8 characters: 11/8" and "The Auth Passphrase value will be encoded during submit." There are icons for clearing and toggling visibility.
- Priv Passphrase**: A text input field containing "SNMPv3rocks". Below the field, it says "Requires at least 8 characters: 11/8" and "The Priv Passphrase value will be encoded during submit." There are icons for clearing and toggling visibility.
- Priv Protocol**: A dropdown menu with "AES" selected.

At the bottom of the form are three buttons: "Reset", "Cancel", and "Submit".

- Click **Submit** to continue—the dashboard updates with the User details.

Figure 10-15: Users

The screenshot shows a dashboard titled "Users" with a table of users. The table has columns for Actions, Name, Auth Passphrase, Priv Passphrase, and Priv Protocol. There is one user entry: "MySNMPv3User".

Actions	Name	Auth Passphrase	Priv Passphrase	Priv Protocol
<input type="checkbox"/>	MySNMPv3User	71053822373c32272d100c05	033775263619725e410a1216	AES

At the bottom right of the table, it says "1-1 of 1 items" and "10 / page".

10.2.1 Configuring SNMP Traps

Complete the following steps to configure the SNMP traps sent to the trap host. SNMP Traps configurations include:

- [Controller Traps](#)
- [Switch Traps](#)

1. Select **Controller Traps** on the SNMP landing page.

Figure 10-16: Controller Traps

The screenshot shows the top navigation bar with 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security' tabs. Below the navigation bar is a 'Fabric Summary: Healthy' status indicator. The main heading is 'SNMP' with the subtext 'Manage SNMP services on DANZ Monitoring Fabric'. A blue button labeled 'Edit Access Control for SNMP' is on the right. Below this are tabs for 'Local Configuration', 'Global Configuration', 'Controller Traps' (which is selected), and 'Switch Traps'. A help icon and text 'Configure the traps sent to the trap host.' are present. Underneath, there is a section for 'Controller Traps' with an 'Edit' link. A table shows a single entry: 'Disk Percent' with a value of '-'.

2. Click **Edit** and enter the **Disk Percent** value.

Figure 10-17: Edit SNMP Controller Traps

The dialog box is titled 'Edit SNMP Controller Traps' with a close button (X) in the top right. It contains a label 'Disk Percent' with a help icon. Below the label is a text input field containing '60' and a percentage sign, with up and down arrow buttons. At the bottom of the dialog are three buttons: 'Reset', 'Cancel', and 'Submit'.

3. Click **Submit**. The dashboard displays the **Disk Percent** value.

Figure 10-18: Disk Percent

This screenshot is identical to Figure 10-16, but the 'Disk Percent' value in the table has been updated from '-' to '60%'.

4. Select **Switch Traps** on the SNMP landing page.

Figure 10-19: Switch Traps

The screenshot shows the SNMP configuration page for DANZ Monitoring Fabric. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', and 'Security'. Below the navigation bar, there is a 'Fabric Summary: Healthy' indicator. The main heading is 'SNMP' with the subtitle 'Manage SNMP services on DANZ Monitoring Fabric'. A button 'Edit Access Control for SNMP' is visible on the right. The page has tabs for 'Local Configuration', 'Global Configuration', 'Controller Traps', and 'Switch Traps'. A description states 'Configure the switch traps sent to the trap host.' Below this, there is a 'Switch Traps' section with an 'Edit' link. A table lists various trap types and their current status.

Trap Type	Status
Power Supply Unit (PSU) Status Change	-
Fan Status Change	-
Link Status Change	-
1-Minute CPU Load Threshold	-
5-Minute CPU Load Threshold	-
15-Minute CPU Load Threshold	-
Percent Idle	-
Percent Utilization	-
Memory free	-

5. Click **Edit** and enter the **Events** values:

- **PSU Status Change** in **seconds(s)**.
- **Fan Status Change** in **seconds(s)**.
- **Link Status Change** in **seconds(s)**.

Enable **Authentication Failure** by moving the selector switch to the right.

Figure 10-20: Edit SNMP Switch Traps - Events

The screenshot shows the 'Edit SNMP Switch Traps' configuration page with the 'Events' tab selected. On the left, a vertical navigation pane shows 'Events' (1), 'Thresholds' (2), and 'Thermal' (3). The main content area contains three sections: 'Power Supply Unit (PSU) Status Change' with a value of 10 seconds, 'Fan Status Change' with a value of 10 seconds, and 'Link Status Change' with a value of 1 second. Below these is the 'Authentication Failure' section, which has a toggle switch turned on. At the bottom right, there are 'Reset', 'Cancel', and 'Next' buttons.

6. Click **Next** to continue.
7. Enter the **Thresholds** values:
 - **1-Minute CPU Load Threshold** in percentage.
 - **5-Minute CPU Load Threshold** in percentage.
 - **15-Minute CPU Load Threshold** in percentage.
 - **Percent Idle** in percentage.
 - **Percent Utilization** in percentage.
 - **Memory Free** in bytes.
 - **Full-Match Flow Table** in percentage.

Figure 10-21: Edit SNMP Switch Traps - Thresholds

The screenshot shows the 'Edit SNMP Switch Traps' configuration page with the 'Thresholds' tab selected. On the left, the navigation pane shows 'Events' (1), 'Thresholds' (2), and 'Thermal' (3). The main content area contains seven sections: '1-Minute CPU Load Threshold' (80%), '5-Minute CPU Load Threshold' (80%), '15-Minute CPU Load Threshold' (80%), 'Percent Idle' (20%), 'Percent Utilization' (90%), 'Memory free' (4094 bytes), and 'Full-Match flow table util.' (50%). At the bottom right, there are 'Reset', 'Cancel', 'Previous', and 'Next' buttons.

8. Click **Next** to continue.
9. Enter the **Thermal** values:
 - **Min** in degrees Celsius.
 - **Max** in degrees Celsius.

- **Interval** in seconds(s). It must be equal to or greater than 10.
- **Status** from the drop-down list (None, All, Failed, Good, Missing).

Figure 10-22: Switch Traps - Thermal

Edit SNMP Switch Traps ×

✓ Events

✓ Thresholds

3 Thermal

Thermal ⓘ

Temperature ⓘ

Min ⓘ °C

Max ⓘ °C

Interval ⓘ second(s)

Status ⓘ

Reset Cancel Previous Submit

10. Click **Submit**. The dashboard displays the **Switch Traps** values.

Figure 10-23: Switch Traps Values

Local Configuration Global Configuration Controller Traps Switch Traps

Configure the switch traps sent to the trap host.

Switch Traps [Edit](#)

Power Supply Unit (PSU) Status Change	10 second(s)
Fan Status Change	10 second(s)
Link Status Change	1 second(s)
1-Minute CPU Load Threshold	80%
5-Minute CPU Load Threshold	80%
15-Minute CPU Load Threshold	80%
Percent Idle	20%
Percent Utilization	90%
Memory free	4 KB
Full-Match flow table util.	50%
Min	20 °C
Max	70 °C
Interval	10 second(s)
Status	All
Authentication Failure	☺

10.2.2 Configuring the System Name

Complete the following steps to configure the SNMP system name in the local configuration to a desired value such as a *fully qualified domain name (FQDN)*.

1. Select **Local Configuration** on the SNMP landing page.

Figure 10-24: Local Configuration

The screenshot shows the SNMP configuration page. At the top, there is a search icon and a menu icon. Below that, the title 'SNMP' is displayed, followed by the subtitle 'Manage SNMP services on DANZ Monitoring Fabric'. A button labeled 'Edit Access Control for SNMP' is visible. The 'Local Configuration' tab is active, with other tabs for 'Global Configuration', 'Controller Traps', and 'Switch Traps'. A help icon and text explain that in SNMPv3, an agent (SNMP server) is identified by an engineID. Below this, a paragraph states that the engineID of the SNMP agent is required when configuring an SNMPv3 trap receiver. At the bottom, the 'Local Configuration' section is shown with an 'Edit' link and a table:

Engine ID	-
System Name	DMF-MACSEC-C1.systest.bsn.sjc.aristanetworks.com

2. Click **Edit** and enter the chosen **System Name** string.
3. Click **Submit**. The dashboard displays the new **System Name** string.

10.3 Using the CLI to Configure SNMP

This section describes using the CLI to configure and manage SNMP settings for the DMF Controller cluster.



Note: To configure a separate SNMP server for switches or Service Nodes, configure an access list to permit access from required clients.

10.3.1 Configuring SNMP Access to the Controller

By default, SNMP access to the Controller is disabled. The default SNMP access list is empty, meaning access is not permitted unless specifically enabled.

The following commands enable access to the Controller by remote SNMP clients on the specified subnet:

```
controller-1 (config) # controller
controller-1 (config-controller) # access-control
controller-1 (config-controller-access) # access-list snmp
controller-1 (config-controller-access-list) # 10 permit from 10.8.67.0/24/0
```



Note: The `permit` command enables access to the Controller from an SNMP client in the subnetwork **10.8.67.0**.

To enable access from any subnet, use the access list entry **0.0.0.0/0** (IP v4) and **::/0**(IPv6), as in the following example:

```
controller-1(config)# controller
controller-1(config-controller)# access-control
controller-1(config-controller-access)# access-list snmp
controller-1(config-controller-access-list)# 10 permit from 0.0.0 .0/0
controller-1(config-controller-access-list)# 20 permit from ::/0
```

10.3.2 Configuring SNMP Access to the Analytics Node

To allow SNMP walk to the Analytics Node, use the steps highlighted in the section [Configuring SNMP Access to the Controller](#) above.

10.3.3 Identifying the SNMP Trap Receiver

To identify a host to receive SNMP traps while in the config mode, enter the `snmp-server host` command, which has the following syntax:

```
controller-1(config)# snmp-server host <ipaddress> [udp-port <udp-port>]
```

Replace **ipaddress** with the IP address of the host. Replace **udp-port** with the port number used by the SNMP traps. For example, the following command identifies a management system at **192.168.17.150** using **UDP port 162**.

```
controller-1(config)# snmp-server host 192.168.17.150 udp-port 162
```

UDP port 162 is the default for SNMP trap messages; **UDP port 161** is the default port for general SNMP messages.

The following are the SNMP traps generated by the Controller running on a VM or the hardware appliance:

```
Name OID Trap generation
-----
cpuload .1.3.6.2.4.1.2021.10.1.5.1 when load (average over 1 minute) > %90
memtotalfree .1.3.6.2.4.1.2021.4.11.0 when freemem (of entire Linux OS) < 50K
```

The following are the SNMP traps generated only by the hardware appliance:

```
cputemp .1.3.6.2.2.1.99.1.1.1.4.1001 when CPU core temp > vendor
specified threshold value
ambienttemp .1.3.6.2.2.1.99.1.1.1.4.2001 when chassis inlet temp >
vendor specified threshold value
powersupply .1.3.6.2.2.1.99.1.1.1.4.3001 when power consumption >
vendor specified threshold value
fan**speed .1.3.6.2.2.1.99.1.1.1.4.40** when fan speed < vendor
specified threshold
```

Configuring **disk-percent** trap will monitor the root partition and the `/var/log` partition. To configure the trap:

```
controller-1(config)# snmp-server trap
disk-percent set logging partition space use percentage at which to send trap
<disk-percent> Percent disk utilization (1..100)
controller-1(config)# snmp-server trap disk-percent 75
```

The following is the entry created in the `/etc/snmp/snmpd.conf` file when you configure the trap on the DMF controller:

```
monitor -r 30 -I dskPercent .1.3.6.2.4.1.2021.9.1.9.1 > 75
```

10.3.4 Configuring SNMP Settings

To set the SNMP community string, which is a password used by a management application for accessing SNMP information, enter the **snmp-server community** command from config mode, as in the following example:



Note: Even though the CLI has options for **ro** or **read-only** and **rw** or **read-write** types of community strings, DANZ Monitoring Fabric supports only the **ro** option.

```
controller-1(config)# snmp-server community ro <string>
```

This command sets the community string for read-only access to the SNMP trap server.



Note: To push the SNMP trap host configuration to the monitoring switches, configure the community string to access the MIBs on the controller and switches. The SNMP trap server uses the same community string to receive and process the traps.

To set the SNMP location, enter the **snmp-server location** command from config mode, as in the following example:

```
controller-1(config)# snmp-server location <location>
```

To set the SNMP contact, enter the **snmp-server contact** command from config mode, as in the following example:

```
controller-1(config)# snmp-server contact <contact>
```

To view the current SNMP configuration, enter the **show running-config snmp** command.



Note: The community string appears as a Type 7 encoded value in the running-config.

To monitor the Controller's `/var/log` and `root` partitions, configure the following trap:

- **disk-percent percent:** Replace **percent** with the percentage that triggers a trap when exceeded.



Note: Configuring the **disk-percent** trap on the Analytics Node will monitor the `/var/lib/analytics/data` folder, the `/var/log` folder, and the `root` partition.

To set the SNMP system name string to a desired value such as a *fully qualified domain name (FQDN)*, the Controller configuration must be updated in the `local` node mode and hence must be done individually for each Controller node in the cluster.



Note: The same configuration applies also to Service Nodes, Recorder Nodes, and Analytics Nodes.

You can enter a chosen string as shown in the example below:

```
controller-2# conf
controller-2(config)# local node
controller-2(config-local)# snmp-server
engine-id          Value for the SNMP engine ID, a text string up to 27 characters long
system-name        SNMP system name to expose (sysName)
DMF-MACSEC-2(config-local)# snmp-server system-name
<System-name>     <String>
controller-2(config-local)# snmp-server system-name DMF-C2.aristanetworks.com
controller-2(config-local)# end
controller-2# show run local
```

```
! local
local node
  hostname controller-2
  snmp-server system-name DMF-C2.aristanetworks.com
  interface management
  !
  ipv4
    ip 10.240.189.233/27 gateway 10.240.189.225
    method manual
    dns search qa.bsn.sjc.aristanetworks.com
    dns server 10.240.48.6
  !
  ipv6
    method manual
```

```
host ~ % snmpwalk -v2c -c bigswitch 10.240.189.233 sysName.0
SNMPv2-MIB::sysName.0 = STRING: DMF-C2.aristanetworks.com
```

Without the above configuration, an `snmpwalk` command would return the *hostname* in the `sysName` parameter, as shown in the following:

```
controller-2# conf
controller-2(config)# local node
controller-2(config-local)# no snmp-server system-name DMF-C2.aristanetworks.com
controller-2(config-local)# end
controller-2#
```

```
host ~ % snmpwalk -v2c -c bigswitch 10.240.189.233 sysName.0
SNMPv2-MIB::sysName.0 = STRING: controller-2
```

10.3.5 Configuring SNMP Switch Trap Thresholds

To configure the thresholds for the SNMP traps generated by fabric switches, use the following command:

```
[no] snmp-server switch trap {cpu-load <cpu-load> | cpu-load 5min <cpu-load5> | cpu-load 15min <cpu-load15> | fm-flow-table-util <util> | mem-free <mem-free> | percent-idle <percent> | percent-
```

utilization <percent> | psu-status <psu-status> | fan-status <fan-status> | link-status <link-status> | auth-fail | thermal [all | failed | good | missing | <interval> <min-temp> <max-temp>]

Use the following keywords with the `snmp-server switch trap` command as required.

- **auth-fail**: Sends a trap when an authentication attempt fails.
- **cpu-load *cpu-load***: Replace ***cpu-load*** with the threshold for CPU utilization.
- **fan-status**: Sends a trap when the fan status changes. Set the interval for monitoring between **10** and **100,000** seconds.
- **fm-flow-table-util *util***: Replace ***util*** with the percentage that triggers a trap when exceeded.
- **link-status**: Sends a trap when the status of a link changes. Set the interval for monitoring between **1** and **100,000** seconds.
- **mem-free *mem-free***: Replace ***mem-free*** with the threshold (in bytes) for memory utilization.
- **percent-idle *percent***: Replace ***percent*** with the percentage of CPU idle utilization that triggers a trap when exceeded.
- **percent-utilization *percent***: Replace ***percent*** with the with the percentage of CPU utilization that triggers a trap when exceeded.
- **psu-status**: Generate a trap when PSU status changes. Set the interval for monitoring between **10** and **100,000** seconds.
- **thermal**: Sends a trap when the thermal sensor status changes as specified using the following options.
 - **all**: Includes failed, good, and missing.
 - **failed**: Sends a trap when the thermal sensor fails.
 - **good**: Sends a trap when the thermal environment is normal.
 - **missing**: Sends a trip when the thermal sensor is not present.
 - **interval**: Sends the trip after the expiry of the specified interval. The range is **10** to **100,000** seconds.
 - [**min-temp** | **max-temp**]: A trap is generated when the temperature in degrees Celsius is less than **min-temp** or greater than **max-temp**.



Note: It is highly recommended to use **percent-idle** or **percent-utilization** instead of **cpu-load** trap.

10.3.6 SNMP Traps for DMF Service Node Appliance

The following are the SNMP traps supported by the DANZ Monitoring Fabric (DMF) Service Node appliance.

- PSU failed/recovered
- Fan failed/recovered
- Temp exceeded some threshold or came back to normal
- Interfaces up/down
- SN inaccessible by the Controller
- SN NetFlow GW is inaccessible
- Percent (%) packet drop exceeded some threshold

10.3.7 Managing the SNMPv3 Engine ID for Trap Receivers

SNMPv3 adds authentication and encryption to the features provided by earlier versions of SNMP (v1 and v2). DANZ Monitoring Fabric (DMF) supports the SNMPv3 user-based security model (USM) for message security through authentication and encryption.

In SNMPv3, an engineID identifies the agent (SNMP server), which helps prevent unauthorized SNMPv3 messages, such as traps, from being accepted or intercepted by unauthorized receivers. The engineID of the SNMP agent is required when configuring an SNMPv3 trap receiver to receive messages from an agent, including a DMF Controller or fabric switch.

In DMF, the engineID is auto-generated for the fabric switches. To view the engineID for a specific fabric switch, enter the following command:

```
controller-1> show switch <switch-name> running-config
```

For the DMF Controller, specify an **engine-ID** keyword that is used to generate the Controller engine-ID. The engine-ID keyword is a text string, up to 27 characters. To configure the engine-id, use the **snmp-server engine-id string** command from the **config-local-node** submode, as in the following example:

```
controller-1(config)# local node
controller-1(config-local)# snmp-server engine-id controller-1_EngineID
```

The engineID of the DMF Controller is configured for the local node. This configuration must be entered separately on the active and standby Controllers. The acceptable practice recommends configuring a different engineID for each Controller.



Note: The engine-id configuration is not included when applying a saved running-config to the Controller. The engine-id configuration must be reapplied using **snmp-server engine-id** command.

The **snmp-server engine-id** command sets the engine-ID for the Controller using the following format:

```
0x80001f8804 + <hex string>
```

where **hex string** is the ASCII hex version of the user-supplied string, which can be found using a tool like **xxd**:

```
$ echo "abcdef--g" | xxd -ps
6162636465662d2d670a
```

This command lets you calculate the engine ID, as in the following example.

```
snmp-server engine-id Controller2_Engine_ID
workstation~$ echo "Controller2_Engine_ID" | xxd -ps
436f6e74726f6c6c6572325f456e67696e655f49440a workstation$
```

The following is the output from the above with the trailing **0a** removed.

```
0x80001f8804
workstation~$ sudo cat /var/lib/snmp/snmpd.conf | grep old
oldEngineID 0x80001f8804436f6e74726f6c6c6572325f456e67696e655f4944 <-----
```

10.3.8 Configuring SNMPv3 Users

Use the `snmp-server user` command in config mode to create a user account for SNMP v3 access. When running an `snmpwalk` (`snmpget`, `snmpgetnext`, `snmpbulkget`) from a shell, passphrases should be enclosed in single quotes. Entering the passphrase with double quotes (" "), may result in an error. This command has the following syntax:

```
[no] snmp-server user <name> {auth [0] <cleartext passphrase> | 7 <auth-passphrase>} [priv {aes | des}]{[0] <cleartext passphrase> | 7 <priv-passphrase>}}
```

The following is the meaning of each keyword:

- **auth | auth 0 | auth 7:** Use a plaintext passphrase or a type 7 encoded passphrase.
- **cleartext-passphrase:** A cleartext passphrase from 8 to 64 alphanumeric characters including dash ("-") and space). A dash or whitespace is not allowed at the beginning or end of the passphrase. Other special characters are not allowed.
- **private-passphrase:** A type 0 encoded passphrase from 8 to 64 alphanumeric characters including dash ("-") and space. A dash or whitespace is not allowed at the beginning or end of the passphrase. Other special characters are not allowed.
- **type-7-passphrase:** A type 7 encoded passphrase from 8 to 128 alphanumeric characters including dash ("-") and space. The maximum text string length that can be used with a Type 7 encoder, which can be found online, is 64. A dash or whitespace is not allowed at the beginning or end of the passphrase. Other special characters are not allowed.
- **priv {aes | des}:** Optional keyword to perform Advanced Encryption Standard (AES) or Data Encryption Standard (DES) encryption of the following passphrase, which is used as an encryption key to encrypt the SNMP messages between the SNMP agent and the manager.
- **user username:** Up to 32 alphanumeric characters including dash ("-") and underscore ("_") Spaces are not permitted. After you configure the username with a plaintext passphrase, the output from the `show snmp-server` command displays the passphrases in Type7 encoded strings. The Controller's configuration gets pushed through zero touch networking (ZTN) to the connected fabric switches.



Note: DANZ Monitoring Fabric (DMF) only supports the `ro` or `read-only` type of community string option.

10.3.9 SNMPv3 Command Examples

Example 1. The `snmp_1` user is configured for authentication (authNoPriv) with the plaintext password `authauth1`.

```
controller-1(config)# snmp-server user snmp_1 auth authauth1
```

Example 2. The `snmp-2` user is configured for authentication (authNoPriv) with the plaintext password `authauth1`.

```
controller-1(config)# snmp-server user snmp-2 auth 0 authauth2
```

Example 3. The `snmp11` user is configured for authentication and DES encryption (authpriv) with the auth password `authauth11` and the encryption key `privpriv11`.

```
controller-1(config)# snmp-server user snmp11 auth 0 authauth11 priv des 0 privpriv11
```

Example 4. The *snmp21* user is configured for authentication and AES encryption (authpriv) with the auth password *authauth21* and the encryption key *privpriv21*.

```
controller-1(config)# snmp-server user snmp21 auth 0 authauth21 priv aes 0 privpriv21
```

The following are examples of Type7 encoded passphrases:

```
controller-1(config)# snmp-server user snmp1 auth 7 0207114f03071a35441f
controller-1(config)# snmp-server user snmp20 auth 7 0207114f03071a35441c59 priv des 7
021616521d161d285a1c59
controller-1(config)# snmp-server user snmp30 auth 7 0207114f03071a35441d59 priv aes 7
021616521d161d285a1d59
```

10.4 Configuring SNMP on a Specific Switch

Configuring SNMP for a specific switch does not affect the Controller or other switches. Otherwise, the configuration is similar to configuring SNMP at the Controller level, using the **Maintenance > SNMP** option.



Note: Before configuring SNMP for a specific switch, **enable** SNMP access to the Controller.

10.4.1 Using the GUI to Configure SNMP on a Specific Switch

To use the GUI to merge/override the default SNMP configuration with switch-specific SNMP configuration, complete the following steps:

1. Select **Fabric > Switches** and click the link for a specific switch.

Figure 10-25: Fabric Switches

Name	Connected	Quarantine Reason	Admin Status	ZeroTouch State	Management Interface	Domain	Priority	RDP Source IPv4 Address	Source IPv4 Address	Connected Since	Connection Time	Management IP Address	Allocated IP Address	Port	Interface Count	Flow Count
clt-core-1	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:54:24pm UTC	1M 1w 5d	1980	10.240.142.11	42239	33	na
clt-core-2	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:54:01pm UTC	1M 1w 5d	1980	10.207.118.28	33476	33	na
clt-delivr-1	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:53:57pm UTC	1M 1w 5d	1980	10.240.153.132	35002	56	na
clt-delivr-2	✓	--	✓ Up	Ok	--	--	--	--	--	Jan 31, 2024, 6:53:39pm UTC	1M 1w 5d	1980	10.240.153.131	42436	56	na
clt-fiber-1	✗	--	✓ Up	Reloading	Prefer Dedicated Management Interface	--	--	--	--	--	--	--	--	--	--	na
clt-fiber-2	✗	--	✓ Up	--	Prefer Dedicated Management Interface	--	--	--	--	--	--	--	--	--	--	na

2. On the Switches page, click the **Actions** control followed by **Configure Switch**.

Figure 10-26: Configure Switch Dialog

Configure Switch

1. Info ✓

2. PTP ✓

3. Clock ✓

4. SNMP ✓

5. SNMP Traps ✓

6. Logging ✓

7. TACACS ✓

8. sFlow ✓

9. LAG Enhanced Hash ✓

Name * cliff-core-1

Description

MAC Address

CC:1 aa:33:93:e0:ac

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 * Down Up

Management Interface * Prefer Dedicated Management Interface

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

IP Assignment Type

IP Address - IPv4 Address -

Tunnel Encapsulation Type * GRE

Change Password No Yes

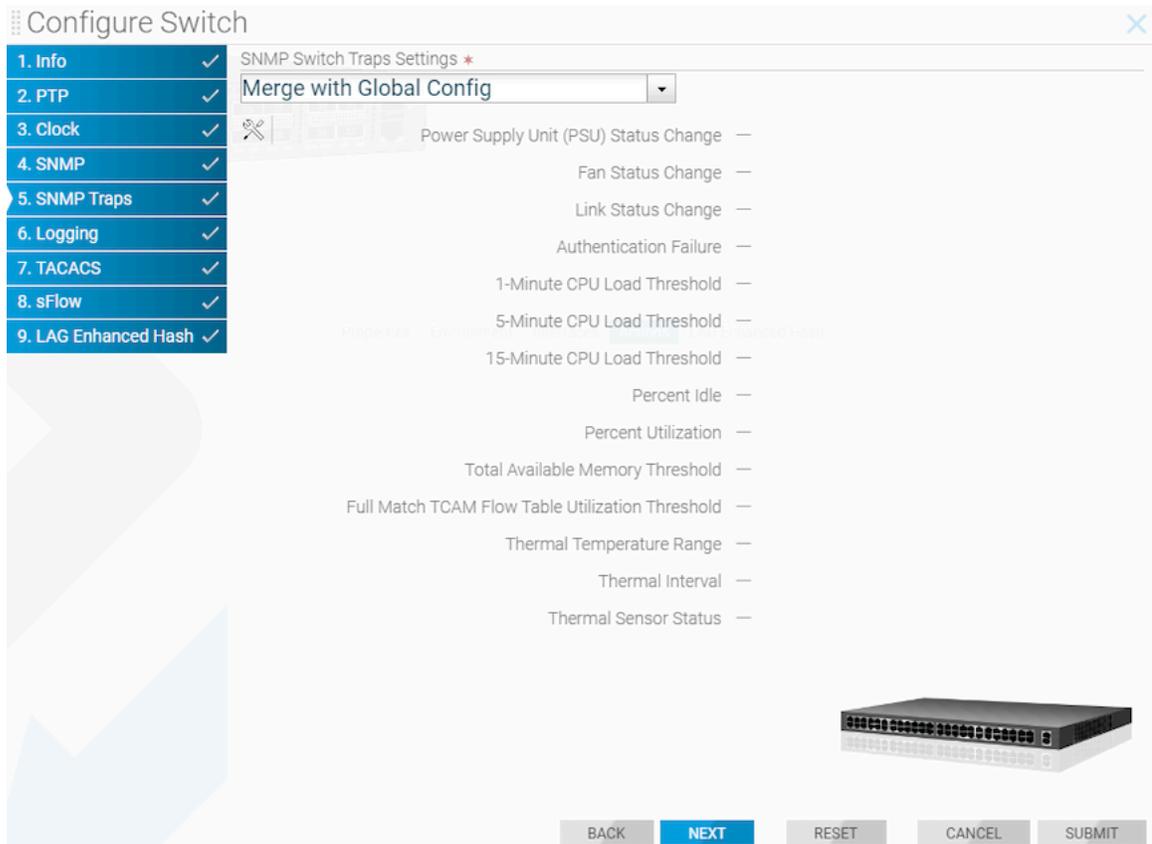
BACK NEXT RESET CANCEL SUBMIT

This page allows merging and overriding the default configuration pushed from the DANZ Monitoring Fabric (DMF) Controller with switch-specific SNMP configuration.

3. To merge or override the SNMP configuration, click the **SNMP** link. Choose from the **SNMP Settings** drop-down to **Merge with Global Config** or **Override Global Config**.
4. Make any changes required to the specific switch configuration, click **Next** to customize the SNMP traps, or click **Submit**.

- To merge or override the configuration for SNMP traps, click the **SNMP Traps** link and choose from the **SNMP Switch Trap Settings** drop-down to either **Merge with Global Config** or **Override Global Config**.

Figure 10-27: SNMP Traps



- Make any changes required to the specific switch configuration and click **Submit**.

10.4.2 Using the CLI to Configure SNMP on a Specific Switch



Note: Before entering SNMP commands from the config-switch submode, **enable** SNMP access to the Controller.

- When using the config-switch submode for a specific switch, configuration changes, including SNMP, do not affect the Controller or other switches. Otherwise, the configuration is similar to configuring SNMP in config mode at the Controller level.
- Entering the `snmp-server enable traps` command in config mode pushes snmp-server enable configuration to each connected fabric switch. Verify the switch configuration by entering the `show effective-config switch switch-name snmp` from the CLI, as in the following example.

```
controller-1(config)# snmp-server enable traps
```

- From the switch CLI:

```
controller-1(config)# show effective-config switch switch-btsw-1 snmp
! switch
switch switch-btsw-1
snmp-server enable traps
```

Like the GUI, use the CLI to merge or override the default SNMP configuration with switch-specific SNMP configuration. To do so, complete the following steps:

1. Add the SNMP configuration at the Controller. This is the default SNMP configuration pushed to all the switches. The following is an example configuration:

```
controller-1(config)# show running-config snmp
! snmp-server
snmp-server host 10.1.1.1
snmp-server enable traps
snmp-server community ro 7 02161159070f0c
snmp-server contact Alice
snmp-server location 'San Francisco'
snmp-server user user1 auth 7 0217135e191216344541
```

2. Configure switch-specific-parameters at the config-switch submode.

```
controller-1(config)# switch-btsw-1
controller-1(config-switch)# snmp-server host 10.1.1.2
controller-1(config-switch)# snmp-server contact Bob
controller-1(config-switch)# snmp-server location 'San Jose'
controller-1(config-switch)# snmp-server user user2 auth 0 qwertyuiop
```

3. In the config-switch submode, type either `merge-global` to merge the global config with the switch-specific config or `override-global` to override the global config with the switch config. When choosing neither, the switch inherits the global config, and any configuration added under the config-switch submode will be redundant.

```
controller-1(config-switch)# snmp-server merge-global
```

4. Check the SNMP configuration running on the switch using the CLI command `show effective-config switch switch-name snmp`:

```
controller-1(config-switch)# show effective-config switch switch-btsw-1 snmp
! switch
switch switch-btsw-1
snmp-server host 10.1.1.1
snmp-server host 10.1.1.2
snmp-server enable traps
snmp-server community ro 7 02161159070f0c
snmp-server contact Bob
snmp-server location 'San Jose'
snmp-server user user1 auth 7 0217135e191216344541
snmp-server user user2 auth 7 0207175f0d01072b4742
```

When using `merge-global`, the effective configuration on the switch is a merge of the global configuration and the switch-specific configuration.



Note: SNMP community, user, and host are of list-type. In merge-mode these list-type configurations append to potentially existing global config.

Below is an example with `override-global`:

```
controller-1(config-switch)# snmp-server override-global
controller-1(config-switch)# show effective-config switch switch-btsw-1 snmp
! switch
switch switch-btsw-1
snmp-server host 10.1.1.2
snmp-server contact Bob
snmp-server location 'San Jose'
snmp-server user user2 auth 7 0207175f0d01072b4742
```

When using `override-global`, the effective configuration on the switch is only the switch-specific configuration and completely overrides the default configuration inherited from the Controller.

5. Configuring SNMP traps using merge and override global commands is similar. See examples below:

```
controller-1(config)# snmp-server switch trap thermal all
controller-1(config)# snmp-server switch trap link-status 5
controller-1(config)# snmp-server switch trap percent-utilization 80
controller-1(config)# switch-btsw-1
controller-1(config-switch)# snmp-server switch trap thermal failed
controller-1(config-switch)# snmp-server switch trap link-status 1
controller-1(config-switch)# snmp-server switch trap percent-utilization 90
```

Example 1. merge-global

```
controller-1(config-switch)# snmp-server trap merge-global
controller-1(config-switch)# show effective-config switch switch-btsw-1 snmp-trap
! switch
switch switch-btsw-1
snmp-server switch trap thermal failed
snmp-server switch trap link-status 1
snmp-server switch trap percent-utilization 90
```

Example 2. override-global

```
controller-1(config-switch)# snmp-server trap override-global
controller-1(config-switch)# show effective-config switch switch-btsw-1 snmp-trap
! switch
switch switch-btsw-1
snmp-server switch trap thermal failed
snmp-server switch trap link-status 1
snmp-server switch trap percent-utilization 90
```

To limit SNMP access to clients in specific IP subnetworks, enter the **snmp-server community** command from the config-switch submode on the DMF Controller. This command has the following syntax:

```
snmp-server community {rw | ro} {<cleartext secret> |
0 <cleartext secret> | 7 <obfuscated secret>}
```

When using the **merge-global** and **override-global** commands at the config-switch submode, the SNMP community for the switch can be changed as shown in the following example:

SNMP configuration on the controller:

```
controller-1(config)# show running-config snmp
! snmp-server
snmp-server host 10.1.1.1
snmp-server community ro 7 02161159070f0c
snmp-server contact Alice
snmp-server location 'San Francisco'
snmp-server user user1 auth 7 0217135e191216344541
```

SNMP configuration on the switch:

```
controller-1(config-switch)# show run switch switch-btsw-1
! switch
switch switch-btsw-1
snmp-server override-global
snmp-server enable traps
snmp-server host 10.1.1.2
snmp-server community ro 7 021616521d071b24
snmp-server contact Bob
snmp-server location 'San Jose'
snmp-server user user2 auth 7 0207175f0d01072b4742
```

10.5 SNMP Clear Trap

SNMP trap messages are sent whenever a threshold is reached, or an HW failure happens, like PSU failure/removal. An SNMP clear trap message is sent whenever a threshold is less than the specified range or the HW failure is fixed, such as when the PSU starts working.

There is no command to enable this feature. This feature is automatically enabled when configuring the SNMP trap on the Controller.

SNMP traps that do not have associated clear traps have other ways of notifying state change. For example, link up and link down traps are sent when the link goes up and down. The `/etc/snmp/snmpd.conf` file lists all SNMP traps and clear trap settings.



Note: SNMP clear traps will be sent without any prior associated SNMP traps when the system comes up or there is any SNMP configuration change. Ignore these SNMP clear traps.

SNMP clear trap messages are not supported on DMF switches running EOS.

The following are switch traps for which clear traps will be sent:

- switch trap cpu-load
- switch trap fm-flow-table-util
- switch trap mem-free
- switch trap percent-idle
- switch trap percent-utilization

These are the appliance (Controller, Service Node, Recorder Node, Analytic Node) traps for which clear traps will be sent.

 **Note:** Upgrade the appliance iDRAC Firmware to the recommended version of **5.10.50.00** or later.

• disk-percent	
• memtotalfree	
• lowmemavailable	
• cpuload	
• cputemp	
• cpu1temp	
• ambienttemp	
• exhausttemp	
• powersupply	
• fanspeed	The number of fans on an appliance varies. Depending on the number of fans on the appliance, fanspeed clear traps are sent. Fan speed traps are named <i>fan1Aspeed</i> , <i>fan1Bspeed</i> , etc.
• psuCount	
• fanCount	

10.6 SNMP Trap Generation for Packet Drops and Link Saturation

Users wishing to be notified about packet drops or high link saturation in the DANZ Monitoring Fabric (DMF) fabric can receive SNMP traps for these events.

Specifically, when trap generation is enabled, the following events will send an SNMP trap to the configured trap collector:

- Transmit packet drop counter increase at a switch interface managed by DMF.
- Traffic saturation levels above the high water mark of 90% on a link managed by DMF.
- A drop in traffic saturation levels below the low water mark of 70% if the link was previously saturated.



Note: This feature is compatible with all hardware platforms with DMF 8.4 version or higher. However, this does not include the DMF Managed Analytics, Recorder Node, and Service Node appliances.

SNMP Trap Generation for Packet Drops and Link Saturation introduces the following OIDs:

OID	Type	Description
.1.3.6.1.4.1.37538.68.77.70	none	Root of the tree for custom DMF extensions.
.1.3.6.1.4.1.37538.68.77.70.1	none	Subtree root for link saturation warnings based on data from the endpoint applications/dmf/info/warnings/link-saturation.
.1.3.6.1.4.1.37538.68.77.70.1.1.0	integer	The number of new link saturation warnings that have been detected since the last polling interval (15 seconds).
.1.3.6.1.4.1.37538.68.77.70.1.2.0	integer	The number of link saturation warnings that have cleared since the last polling interval (15 seconds).
.1.3.6.1.4.1.37538.68.77.70.1.3.0	oidref	The OID to the start of the table containing all known link saturation warnings. Always holds the value: .1.3.6.1.4.1.37538.68.77.70.1.5 This object is sent as a varbind in the trap newLinkSaturationWarnings to indicate to the trap collector that this OID should be used in a walk for the updated list of warnings.
.1.3.6.1.4.1.37538.68.77.70.1.4.0	oidref	The OID to the start of the table containing all cleared link saturation warnings since the last polling interval. Always holds the value: .1.3.6.1.4.1.37538.68.77.70.1.6 This object is sent as a varbind in the trap linkSaturationWarningsCleared to indicate to the trap collector that this OID should be used in a walk for the updated list of cleared warnings.
.1.3.6.1.4.1.37538.68.77.70.1.5	string	Start/root of the table of link saturation warnings. Holds the string: "Table of current warnings"

SNMP Trap Generation for Packet Drops and Link Saturation OIDs (cont'd):

OID	Type	Description
.1.3.6.1.4.1.37538.68.77.70.1.5.1	string	[Table row 1] A string summarizing a link saturation warning as a dot-separated tuple of switch DPID, interface name, and traffic direction (either 'RX' or 'TX'). Example: "00:00:52:54:00:a2:d8:9d.ethernet2.TX" Indicates that port ethernet2 on the switch identified by DPID 00:00:52:54:00:a2:d8:9d has experienced transmit link saturation levels above 90%.
.1.3.6.1.4.1.37538.68.77.70.1.6	string	Start/root of the table of cleared link saturation warnings. Holds the string: "Table of cleared warnings"
.1.3.6.1.4.1.37538.68.77.70.1.6.1	string	[Table row 1] A string summarizing a recently-cleared link saturation warning as a dot-separated tuple of switch DPID, interface name, and traffic direction (either 'RX' or 'TX'). Example: "00:00:52:54:00:a2:d8:9d.ethernet2.TX" Indicates that the transmit link saturation warning for port ethernet2 on the switch identified by DPID 00:00:52:54:00:a2:d8:9d has cleared (dropped below the low watermark of 70%).
.1.3.6.1.4.1.37538.68.77.70.2	None	Subtree root for packet drop warnings which are issued when the TX drop counter of an interface increases. Based on the data from the endpoint <code>core/switch/dmf-stats/interface</code>
.1.3.6.1.4.1.37538.68.77.70.2.1.0	integer	The number of interfaces that have experienced an increase in the transmit drop counter since the last polling interval (15 seconds).
.1.3.6.1.4.1.37538.68.77.70.2.2.0	oidref	The OID to the start of the table containing all of the interfaces that have seen its transmit drop counter increase since the last polling interval. Always holds the value: .1.3.6.1.4.1.37538.68.77.70.2.3 This object is sent as a varbind in the trap <code>newPacketDropsWarnings</code> to indicate to the trap collector that this OID should be used in a walk for the updated list of interfaces with increased counters.
.1.3.6.1.4.1.37538.68.77.70.2.3	string	Start/root of the table of interfaces. Holds the string: "Table of interfaces with incremented drop counters"
.1.3.6.1.4.1.37538.68.77.70.2.3.1	string	[Table row 1] A string representing an interface that has experienced a transmit drop counter increase. Example: "00:00:52:54:00:a2:d8:9d.ethernet2" represents interface ethernet2 on the switch identified by DPID 00:00:52:54:00:a2:d8:9d.
.1.3.6.1.4.1.37538.68.77.70.2.3.2	counter64	[Table row 2] An unsigned integer holding the current counter value.

This feature introduces the following traps:

- newLinkSaturationWarnings
 - When new link saturation warnings are generated, in addition to its name, the trap contains the following varbinds:
 - .1.3.6.1.4.1.37538.68.77.70.1.1.0
 - .1.3.6.1.4.1.37538.68.77.70.1.3.0
- linkSaturationWarningsCleared
 - When one or more existing link saturation warnings are cleared, in addition to its name, the trap contains the following varbinds:
 - .1.3.6.1.4.1.37538.68.77.70.1.2.0
 - .1.3.6.1.4.1.37538.68.77.70.1.4.0
- newPacketDropsWarnings
 - When packet drop counters are seen incrementing on one or more managed interfaces, in addition to its name, the trap contains the following varbinds:
 - .1.3.6.1.4.1.37538.68.77.70.2.1.0
 - .1.3.6.1.4.1.37538.68.77.70.2.2.0

When traps are enabled, the Controller can be queried for the monitored data to send the traps. The MIB for this feature is rooted at the OID .1.3.6.1.4.1.37538.68.77.70. For a Controller at 192.0.2.1 with an SNMP community named **example**, a query using `snmpwalk` and the corresponding response may resemble the following:

```
$ snmpwalk -v2c -c example 192.0.2.1 .1.3.6.1.4.1.37538.68.77.70
iso.3.6.1.4.1.37538.68.77.70.1.1.0 = INTEGER: 0
iso.3.6.1.4.1.37538.68.77.70.1.2.0 = INTEGER: 0
iso.3.6.1.4.1.37538.68.77.70.1.3.0 = OID: iso.3.6.1.4.1.37538.68.77.70.1.5
iso.3.6.1.4.1.37538.68.77.70.1.4.0 = OID: iso.3.6.1.4.1.37538.68.77.70.1.6
iso.3.6.1.4.1.37538.68.77.70.1.5 = STRING: "Table of current warnings"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.1 = STRING: "00:00:52:54:00:a2:d8:9d.ethernet2.TX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.2 = STRING: "00:00:5c:54:00:f2:81:03.ethernet0.RX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.3 = STRING: "00:00:12:b0:03:22:07:a2.ethernet1.RX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.4 = STRING: "00:00:50:ef:0c:a2:16:f6.ethernet3.TX"
iso.3.6.1.4.1.37538.68.77.70.1.6 = STRING: "Table of cleared warnings"
iso.3.6.1.4.1.37538.68.77.70.2.1.0 = INTEGER: 2
iso.3.6.1.4.1.37538.68.77.70.2.2.0 = OID: iso.3.6.1.4.1.37538.68.77.70.2.3
iso.3.6.1.4.1.37538.68.77.70.2.3 = STRING: "Table of interfaces with incremented drop counters"
iso.3.6.1.4.1.37538.68.77.70.2.3.1.1 = STRING: "00:00:52:54:00:a2:d8:9d.ethernet2"
iso.3.6.1.4.1.37538.68.77.70.2.3.1.2 = STRING: "00:00:52:a4:31:f2:81:56.ethernet0"
iso.3.6.1.4.1.37538.68.77.70.2.3.2.1 = Counter32: 11
iso.3.6.1.4.1.37538.68.77.70.2.3.2.2 = Counter32: 1
```

The tables in the response are empty when there are no warnings:

```
iso.3.6.1.4.1.37538.68.77.70.1.1.0 = INTEGER: 0
iso.3.6.1.4.1.37538.68.77.70.1.2.0 = INTEGER: 0
iso.3.6.1.4.1.37538.68.77.70.1.3.0 = OID: iso.3.6.1.4.1.37538.68.77.70.1.5
iso.3.6.1.4.1.37538.68.77.70.1.4.0 = OID: iso.3.6.1.4.1.37538.68.77.70.1.6
iso.3.6.1.4.1.37538.68.77.70.1.5 = STRING: "Table of current warnings"
iso.3.6.1.4.1.37538.68.77.70.1.6 = STRING: "Table of cleared warnings"
iso.3.6.1.4.1.37538.68.77.70.2.1.0 = INTEGER: 0
iso.3.6.1.4.1.37538.68.77.70.2.2.0 = OID: iso.3.6.1.4.1.37538.68.77.70.2.3
iso.3.6.1.4.1.37538.68.77.70.2.3 = STRING: "Table of interfaces with incremented drop counters"
```

10.6.1 Using the CLI to Configure the SNMP Traps

This feature is enabled by enabling trap generation, configuring an SNMP community or USM user for sending the traps, and specifying a trap receiver host, as shown below.

```
(config)# snmp-server enable traps
(config)# snmp-server community ro example
(config)# snmp-server host 192.0.2.10
```

Traps are disabled when trap generation is disabled using the following command:

```
(config)# no snmp-server enable traps
```

CLI Show Commands

SNMP configurations, including traps, communities, or users, appear in the running-config using the **show** command.

```
# show running-config snmp

! snmp-server
snmp-server host 192.0.2.10
snmp-server enable traps
snmp-server community ro 7 02031c5a06160324
```

The configuration entry **snmp-server enable traps** indicates that trap support, including this feature, is enabled.

Syslog Messages

```
SDCACHE7001: Error while refreshing cache for Query{<query>,
includedStateTypes=[LOCAL_CONFIG, GLOBAL_CONFIG, OPERATIONAL]}
```

- <query> is one of:
 - basePath=/applications/dmf/info/warnings/link-saturation
 - basePath=/core/switch, selectedPaths=[dmf-stats/interface]
- Generated when the statistics collected and monitored for this feature could not be cached. The message is logged along with a stack trace containing further details.

```
SDCACHE7002: Query for data failed for Query{<query>,
includedStateTypes=[LOCAL_CONFIG, GLOBAL_CONFIG, OPERATIONAL]}
```

- <query> is one of:
 - basePath=/applications/dmf/info/warnings/link-saturation
 - basePath=/core/switch, selectedPaths=[dmf-stats/interface]
- Generated when the periodic query to the Controller for collecting the statistics monitored when evaluating if traps should be sent failed. The message is logged along with a stack trace containing further details.

```
SNMPEXT4001: Could not disable configuration <configuration>
```

- <configuration> is one of:

- `dmf_trap_monitors`
- `dmf_warning_extensions`
- Generated when the SNMP configurations associated with a trap, named [configuration], could not be disabled. The message is logged along with a stack trace containing further details.

SNMPEXT4002: Could not enable configuration <configuration>

- <configuration> is one of:
 - `dmf_trap_monitors`
 - `dmf_warning_extensions`
- Generated when the SNMP configurations associated with a trap, named [configuration], could not be enabled. The message is logged along with a stack trace containing further details.

Troubleshooting

If the CLI show command indicates:

- Traps are enabled.
- A community or user is configured, and
- A trap host is configured.

The trap host may not be reachable from the Controller, or statistics collection and monitoring for packet drops or link saturation may need to be enabled correctly. Check the logs for any messages described in the previous section using the following command:

```
> show logging controller | grep 'SDCACHE\|SNMPEXT'
```

Contact [Arista Technical Support](#) if any of these logs appear and:

- Traps are not sent from a system with known link saturation or packet drop warnings, e.g., via the CLI or GUI.
- A query, e.g., `snmpwalk` for OIDs rooted under `1.3.6.1.4.1.37538.68.77.70` does not result in responses or return errors.

Considerations

- The user cannot specify the link saturation trap thresholds. Traps are sent when link saturation levels cross the 90% threshold and are cleared once levels drop below 70%.
- The tree associated with this feature is only available through polling with `snmpwalk` when traps are enabled.
- The OIDs do not have resolvable names.
- Polling of link and interface states happens at a fixed interval; there may be a several-second delay between the occurrence of a trap-triggering event and sending the trap.

Using Authentication, Authorization, and Accounting

This chapter describes how to manage administrative access to the DANZ Monitoring Fabric (DMF) Controller using local groups and users (RBAC) and AAA servers in general, and it also provides specific configuration for TACACS+ or RADIUS servers.

11.1 Overview

Access privileges to the DANZ Monitoring Fabric (DMF) Controller are associated with groups. Each user assigned to the group obtains the access permissions associated with the group. The current version of DMF supports the following two groups:

- **Admin Group:** Root privileges with access to all modes, including debug modes.
- **Read-Only:** Read-only administrative access.

DMF also supports communication with a remote AAA server. All authentication, authorization, and accounting functions are set to local by default. The general options that control where these functions occur are:

- **Local only:** The accounts on the remote AAA server are ignored, using only the local database accounts.
- **Local primary and remote backup:** The accounts on the local database are used first. If the account is not found locally, DMF uses the database on the remote server.
- **Remote only:** Authentication and authorization do not fall back to local for users other than the admin and recovery accounts, no matter what happens.
- **Remote primary and local backup:** Uses the accounts on the remote AAA database first. If the remote server is unavailable, DMF uses the local database accounts.



Note: The admin and recovery user accounts are special accounts that cannot be authenticated remotely using RADIUS or TACACS+. These accounts are always authenticated locally to prevent administrative access from being lost in case a remote AAA server is unavailable. The following list summarizes the options available for each function:

The following list summarizes the options available for each function:

- **Accounting:** local, local and remote, or remote.
- **Authentication:** local, local then remote, remote then local, or remote.
- **Authorization:** local, local then remote, remote then local, or remote.



Note: Authorization falls back to local only if remote authorization fails because the remote AAA server is unreachable.

For information about using AAA with remote servers to manage administrative access to the DMF Controller, refer to the [Configuring AAA to Manage DMF Controller Access](#) section.

11.2 Using Local Groups and Users to Manage DMF Controller Access

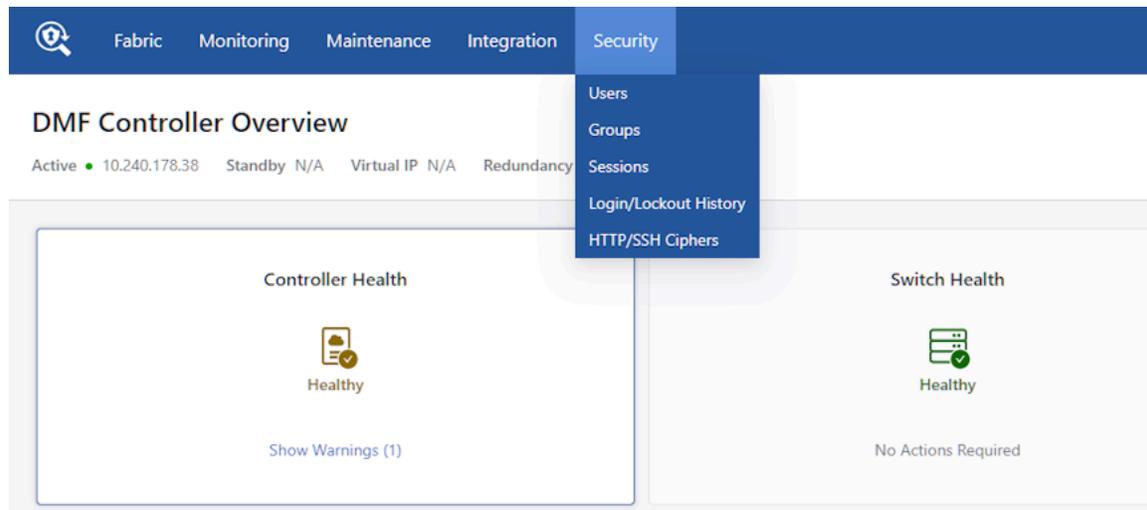
Administrative access to the monitoring fabric is managed by assigning interfaces to groups and then assigning user accounts to the group, which can then view and use the assigned interfaces.

11.2.1 Viewing Existing Groups and Users

To view the current allocation and assignment of groups and resources, enter the following command from any CLI mode on the active DANZ Monitoring Fabric (DMF) Controller.

To use the GUI to manage groups and users, select **Security** from the DMF GUI main menu.

Figure 11-1: Security



The system displays the Security page menu options.

- Users
- Groups
- Sessions
- Login/Lockout History
- HTTP/SSH Ciphers

Groups

Displays the following preconfigured default groups on the DMF Controller:

- **admin group:** Provides full administrative access to all interfaces. Global configuration options, such as assigning a role to a switch or interface, must be performed by the admin user.
- **read-only group:** Provides access to view and monitor fabric activity and switch configuration for all interfaces but does not allow changing configuration or clearing statistics.

All switch interfaces are assigned to both default groups. The admin user has read, use, and configure access to all the interfaces. Any user added to the read-only group has read-only access to all the interfaces.

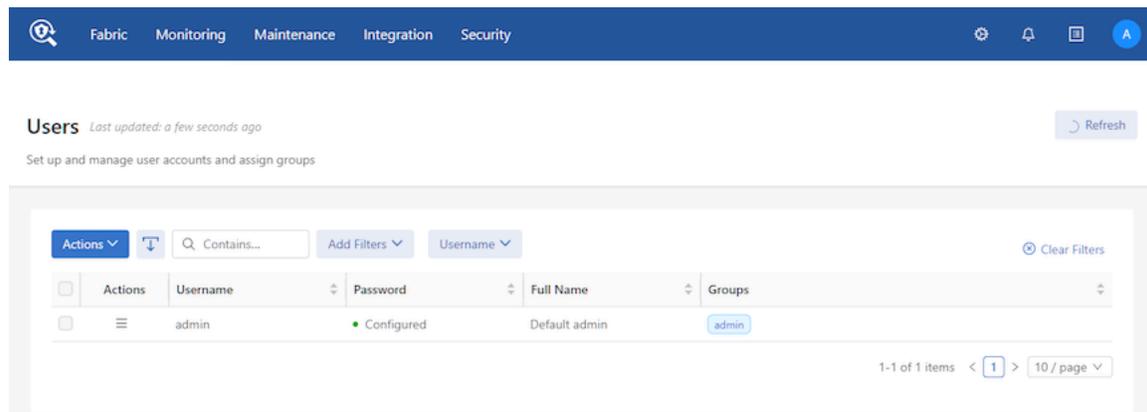
You can define a new group, in which case a user added to the group will be a restricted user with privileges to view and use only the interfaces assigned to the group.

11.2.2 Using the GUI to Manage User Accounts & Security Groups

To add a user to a group or to add a group, complete the following steps.

Add users from the main Security page or the following page when selecting **Security > Users** from the DMF GUI main menu.

Figure 11-2: Security > Users



1. Select **Actions > + Add User** in the Users section. The system displays the Add User dialog.

Figure 11-3: Add User

The 'Add User' dialog box is shown. It has a title bar with 'Add User' and a close button. The form contains the following fields:

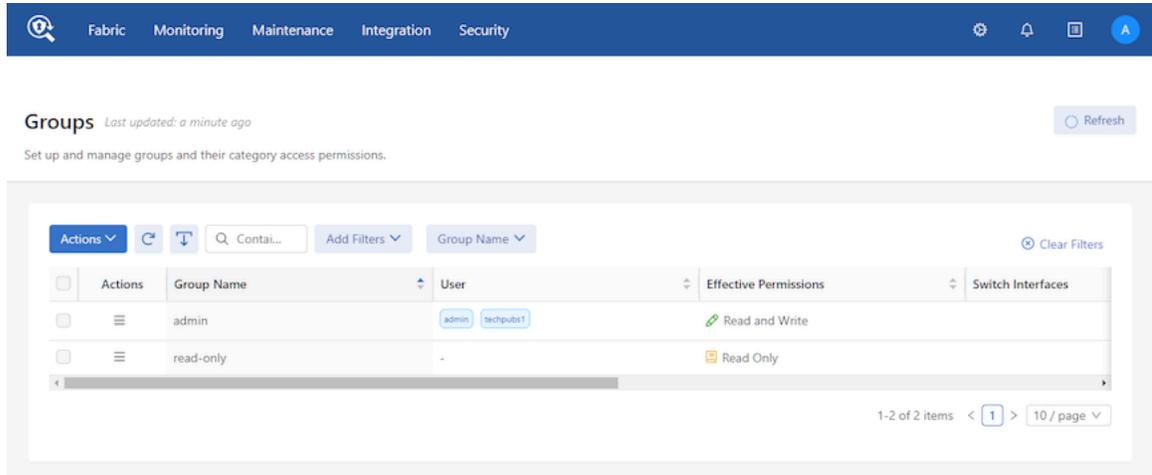
- Username:** techpubs1
- Password:** masked with dots and a toggle icon
- Full Name:** Technical Publications Writer One
- Group:** admin (selected from a dropdown menu)

At the bottom, there are 'Cancel' and 'Submit' buttons.

2. Click **Submit**.

3. To edit a user, click the **Actions > Edit** icon for the requisite **Username**.
4. To delete a user, select **Actions > Delete Selected Users**.
5. To add a **Security Group**, select **Security > Groups** in the Groups section.

Figure 11-4: Security Groups



6. Select **Actions > + Create Group**.

The system displays the Create Group menu.

Figure 11-5: Create Group

Create Group

This group will manage the categories based on permissions.

* Group Name

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

Telemetry

Mixed Permissions

Does Not Elevate

Inherit

Inherit

Inherit

Inherit

Inherit

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

7. Enter a name for the group and all other required group settings, including:
 - Category Access
 - **Default** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - **AAA** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - **App Log** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - **Sysops** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - **Time** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - **Telemetry** - Inherit, Does Not Elevate, Read Only, Read Write (drop-down)
 - Role Based Access Control (RBAC)

- **Policy** - Policies Configured, + **Add Policy**
 - **Service** - Service Configured, + **Add Service**
 - **Switch Interface** - Switch Interface Configured, + **Add Switch Interface**
 - **Recorder Node** - Recorder Nodes Configured, + **Add Recorder Node**
8. When complete, click **Create**. The table displays the new group.
 9. To edit a security group, click the **Actions > Edit** icon for the requisite **Group Name**.
 10. To delete a security group, select **Actions > Delete Selected Groups**.

11.2.3 Using the CLI to Manage Groups and User Accounts

To add a group, enter the `group` command from config mode, as in the following example:

```
controller-1(config)# group dc-group
```

To associate users with a group, enter the `associate` command from `config-group` mode, as in the following example:

```
controller-1(config-group)# associate user bob
controller-1(config-group)# associate user susan
```

To associate interfaces with the group, enter the `associate` command for each interface, as in the following example:

```
controller-1(config-group)# associate switch DMF-FILTER-SWITCH-1 interface TAP1-ethernet1
controller-1(config-group)# associate switch DMF-FILTER-SWITCH-1 interface TAP1-ethernet1
```

To associate other resources with the group, use the following options with the `associate` command:

```
[no]associate object{policy|service}<object-name>
[created-by<object-creator>[created-on<date-and-time>]]
```

Enter the `show group` command to display the currently configured groups and users.

To display the group configuration in the current running-config, enter the `show running-config group` command, as in the following example:

```
controller-1# show running-config group
! group
group admin
associate user admin
group dmf-qa
associate switch DMF-FILTER-SWITCH-1 interface ethernet1
associate switch DMF-FILTER-SWITCH-1 interface ethernet2
associate switch DMF-FILTER-SWITCH-1 interface ethernet25
associate switch DMF-FILTER-SWITCH-1 interface ethernet26
associate user test1
group read-only
```

To use the CLI, enter the following commands from config mode.

```
controller-1(config)# user bob
controller-1(config-local-user)# full-name Robert Smith
controller-1(config-local-user)# password
```

```

Password:
Re-enter:
controller-1(config-local-user)# group admin
controller-1(config-group)# associate user bob
controller-1(config-group)# show user
# User name Full name Groups
- |-----|-----|-----|
1 admin Default admin admin
2 bob Robert Smith admin

```

Enter the **no user** command from config mode to delete a user, as in the following example.

```

node1(config)# no user bob
controller-1(config)# show user
# User name Full name Groups
- |-----|-----|-----|
1 admin Default admin admin

```

To add a read-only user, run the following commands from config mode.

```

controller-1(config)# user john
controller-1(config-local-user)# full-name John Smith
controller-1(config-local-user)# password
Password:
Re-enter:
controller-1(config-local-user)# group read-only
controller-1(config-group)# associate user john
controller-1(config-group)# show user
# User name Full name Groups
- |-----|-----|-----|
1 admin Default admin admin
2 bob Robert Smith admin
3 john John Smith read-only

```

11.2.4 Changing User Passwords

To change the password for the admin user or other user accounts, select the **Actions > Edit** icon for the requisite **Username** from the **Security > Users** page.

Figure 11-6: Security Users

The screenshot shows the 'Users' page in a web application. The page title is 'Users' with a subtitle 'Set up and manage user accounts and assign groups'. There is a 'Refresh' button in the top right. Below the title is a table with the following columns: Actions, Username, Password, Full Name, and Groups. The table contains one row for the 'admin' user, with a status indicator 'Configured' and a group 'admin'. The interface also includes a search bar, 'Add Filters', and 'Clear Filters' options.

Actions	Username	Password	Full Name	Groups
	admin	Configured	Default admin	admin

1-1 of 1 items < 1 > 10 / page



Note: Entering a new password will overwrite the existing one. Leaving it blank retains the original password.

Figure 11-7: Change Password

The screenshot shows a web form titled "Edit User" with a close button (X) in the top right corner. The form contains the following fields:

- Username:** A text input field containing the text "admin".
- Password:** A text input field that is currently empty. Below the field is a small eye icon and a note: "Typing a new Password will overwrite the existing one. Leaving it blank will keep the original Password."
- Full Name:** A text input field containing the text "Default admin".
- Group:** A dropdown menu with "admin" selected and a downward arrow on the right.

At the bottom of the form are two buttons: "Cancel" (light blue) and "Submit" (dark blue).

Type the new password, confirm it, and click **Submit**.

11.2.5 Password Reset

Resetting the Password for the Recovery User

To reset the password for *the recovery* user, please follow one of the following procedures. Perform these steps on each cluster Controller as resetting the password of the *recovery* user on one Controller won't change it for the *recovery* user on the other Controller.

1. Using the Controller's bash:
 - a. Go to the Controller bash by executing `debug bash`.
 - b. Execute `sudo passwd recovery`.

```
admin@controller-1:~$ sudo passwd recovery
New password:
Retype new password:
passwd: password updated successfully
admin@controller-1:~$
```

2. From the recovery login account login:

For this to work, the current password for the *recovery* user is required.

```
recovery@controller-1:~$ passwd recovery
Changing password for recovery.
Current password:
New password:
Retype new password:
```

```
passwd: password updated successfully
recovery@controller-1:~$
```

3. Using the API `/api/v1/rpc/controller/os/action/system-user/reset-password`:

The API call below will reset the recovery user's password to *AdminAdmin*. The example below uses curl initiated from a Linux host, but any RESTful API client can call the API.

```
curl -g -H "Cookie: session_cookie=<session_cookie>"
'https://<controller IP>:8443/api/v1/rpc/controller/os/action/system-user/reset-password'
-d '{"user-name" : "recovery","password" : "AdminAdmin"}' -X POST
```

Resetting Password for admin and Other Local Users

To reset the password for *admin* and other local users, log in to the Controller using *recovery* user credentials. Use the `floodlight-reset-password` command to reset the user's password.

The example below example resets the *admin* user's password.

```
recovery@controller-1:~$ floodlight-reset-password --user admin
Enter new admin password:
Re-enter new admin password:
Password updated for user admin
recovery@controller-1:~$
```

The example below resets the password for the *guest* account, a *read-only* group user.

```
recovery@controller-1:~$ floodlight-reset-password --user guest
Enter new guest password:
Re-enter new guest password:
Password updated for user guest
recovery@controller-1:~$
```

11.2.6 Authentication with a User Token and REST API

Use the `access-token` command to create a long-lived token for authentication with external scripting, such as RESTful API. The token can be deleted (repudiated) at any time. DMF preserves a hashed version of the token in the running-config.

The following example shows an access token with the key *sam* assigned to the user *fred*.

```
controller-1(config)# user fred
controller-1(config-user)# access-token sam
access-token : Y9yVwLawjJ031SthnBKVh3XepIaJ6sSE
```

The system replies with the session cookie for use in the REST API.

To display the hashed version of the session cookie, enter the `show this` command or `show running-config fred`.

```
controller-1(config-user)# show this
! user
user fred
access-token sam 459d1d5e0bc42c5bbbbee091a984e2807592d11d8a5d4cbac110d32da3f0e436c 2018-01-08T15:44:02.353Z
hashed-password
method=PBKDF2WithHmacSHA512,salt=N9E9jaPIFTc_ZO9Oq0gd4A,rounds=25000,ph=true,
7eCf1PGAYqUw53vJ2bsGpSVEU5-D6ix5sHufFU13gFr3AHB4Jqj2eLNZzoo66y_qIRFOOOL8nc5oG5ilgJ1FxA
controller-1(config-user)#
```

The session cookie does not appear in the running-config; the hashed session cookie instead appears. Because the cookie is in the running-config, it persists over upgrade. To remove the token, enter the `no access-token` command, as in the following example:

```
controller-1(config-user)# no access-token sam
```

11.3 Configuring AAA to Manage DMF Controller Access

Authentication, Authorization, and Accounting (AAA) is a general standard for controlling and auditing access to network resources. Various implementations are possible using technologies such as LDAP, RADIUS, Kerberos, TACACS+, or a local database. The current version of DANZ Monitoring Fabric (DMF) supports AAA using a local database on the Controller node or with a remote TACACS+ or RADIUS server. All authentication, authorization, and accounting functions are set to local by default.

Authentication and authorization occur only once, and whatever privileges are associated with the first account found are used. Configuring options separately for authentication and authorization is possible, but Arista Networks recommends using the same settings for simplicity.

When one or more remote AAA servers are available, the best practice is to enable remote logging as the primary authorization method. The local database is only used if no remote server responds to the authorization request before the timeout. By default, the timeout is five seconds per server, with a maximum total timeout of **25** seconds for up to five remote servers.



Note: The admin and recovery user accounts are special accounts that cannot be authenticated remotely using RADIUS or TACACS+. These accounts are always authenticated locally to prevent administrative access from being lost in case a remote AAA server is unavailable.



Note: If the user fails to be authorized by any AAA server, the user account can be authorized as a member of the default group and chosen from the Default Group selection list. The options are admin or read-only when configuring the latter group.

DMF supports AAA services using RADIUS and TACACS+ servers using up to four servers of each type. The default timeout for each AAA server is **five** seconds, which was previously **30** seconds in earlier releases. The total aggregate timeout for all configured servers is **20** seconds.

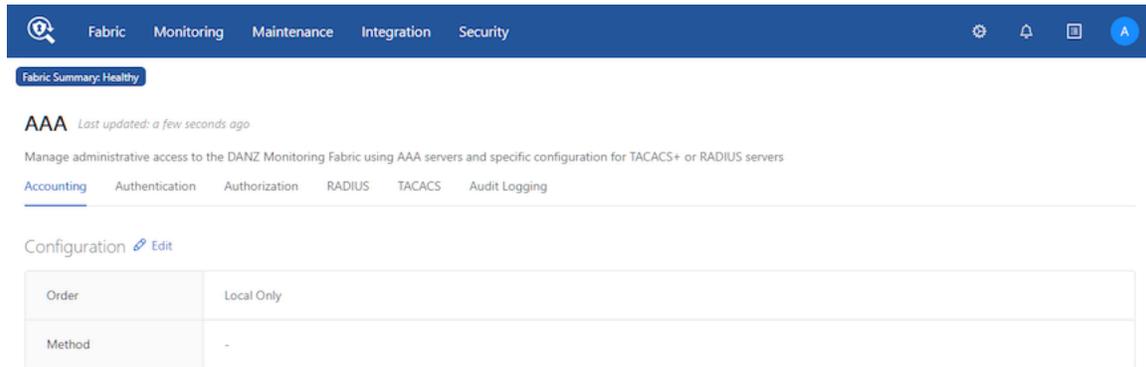
By default, the Controller waits **five** seconds before trying the next server or falling back to the local services provided by the Controller database if that is the configuration. You can change the default timeout, but the maximum aggregated time for all the servers is **25** seconds. Five servers are supported only with a default timeout of **five** seconds or less.

11.3.1 Using the GUI - Enabling Remote AAA Services

GUI Procedure

To manage the configuration of AAA settings on the DMF Controller, select **Maintenance > AAA**.

Figure 11-8: AAA Settings Page



AAA configuration tabs and features include:

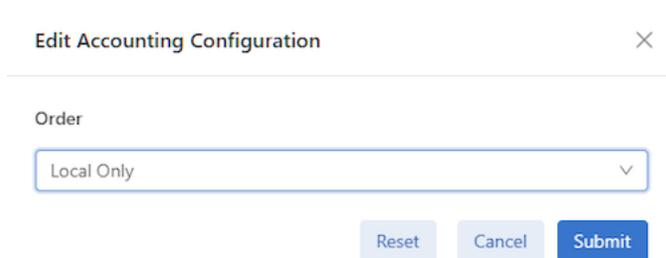
- Accounting
- Authentication
- Authorization
- RADIUS
- TACACS+
- Audit Logging

This page supports identifying up to five destination TACACS+ or RADIUS servers, configuring the server secret, and specifying a timeout.

Accounting

To log accounting log messages to a remote server, click **Accounting > Edit** in the Configuration section.

Figure 11-9: Accounting Settings



By default, DMF saves accounting logs on the local Controller. If local logging is disabled, these logs will not be available for analysis locally. To send the accounting logs to a remote server, select the **Order** preference:

- **Local Only** - Default.
- **Remote Only** - Enter a Remote Method, **TACACS** or **RADIUS** from the drop-down list.
- **Local Before Remote** - Enter a Remote Method, **TACACS** or **RADIUS** from the drop-down list.
- **Remote before Local** - Enter a Remote Method, **TACACS**, or **RADIUS** from the drop-down list.

Click **Submit**.



Tip: Use the **Reset** button to return to the default setting, if required.

Authentication

To configure the authentication method, click **Authentication > Edit** in the Configuration section.

Figure 11-10: Authentication Settings

Edit Authentication Configuration ×

1 Authentication

2 Password Requirements

Order
Local Only

Lockout Policy ⓘ
Mode
No Lockout

Reset Cancel Next

Select the **Authentication Order** preference:

- **Local Only** - Default.
- **Remote Only** - Enter a Remote Method, **TACACS** or **RADIUS** from the drop-down list.
- **Local Before Remote** - Enter a Remote Method, **TACACS** or **RADIUS** from the drop-down list.
- **Remote before Local** - Enter a Remote Method, **TACACS** or **RADIUS** from the drop-down list.

Lockout Policy Mode

Configuration for the time-based user lockout policy. In no-lockout mode, the configuration of this container is ignored, and lockout is disabled. In lockout-remote mode, this container's values (or, in their absence, the defaults) dictate the lockout behavior.

Select the desired **Lockout Policy Mode**:

- **No Lockout**
- **Lockout Remote**

Click **Next** and set the **Password Compliance Requirements** from the drop-down list:

- **None**
- **NIST-800-63B**
- **Custom-Check** - Enter the requirements for the detailed password (min/max, etc.).
- **No-Check**

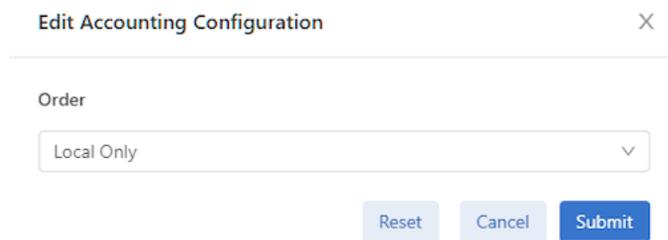
Click **Submit**.

Authorization

To use a remote server for authorization, click **Authorization** > **Edit** in the Configuration section.

Select the **Authorization** preference:

Figure 11-11: Authorization Settings



Edit Accounting Configuration X

Order

Local Only v

Reset Cancel Submit

Select the **Authorization Order** preference:

- **Local Only** - Default.
- **Remote Only** - From the drop-down list.
- **Local Before Remote** - From the drop-down list.
- **Remote before Local** - From the drop-down list.

Click **Submit**.

RADIUS

To configure RADIUS refer to the [Using RADIUS for Managing Access to the DMF Controller](#) section.

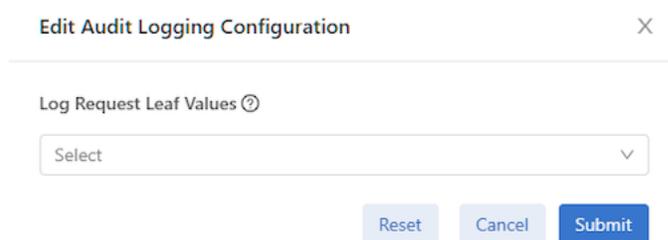
TACACS+

To configure TACACS+ refer to the [Using TACACS+ to Control Access to the DMF Controller](#) section.

Audit Logging

To configure the audit logging method, click **Audit Logging** > **Edit** in the Configuration section.

Figure 11-12: Audit Logging



Edit Audit Logging Configuration X

Log Request Leaf Values ⓘ

Select v

Reset Cancel Submit

Set the type of audit logging of request leaf values; if no value is specified, request leaf values will not be audit logged.

Select the **Log Request Leaf Values** from the drop-down list:

- **None** - Default.
- **Record All Request Values**

Click **Submit**.

11.3.2 Using the CLI - Enabling Remote AAA Services

CLI Procedure

Use the `aaa authorization role default` command to assign a default role to the current account. DMF uses the default group if authentication on a remote server is successful but no role is specifically assigned. This command does not apply to local authorization or authentication. If a local user is not associated with a group on the Controller, login is not allowed.



Note: Use the `authorization role default admin` command carefully because the effect is to provide every user account that authenticates successfully on a remote server with admin-level privileges unless the user account is specifically assigned to a different group.

11.4 Time-based User Lockout

Starting in the **DMF 8.0** release, DANZ Monitoring Fabric supports time-based user lockout functionality. Users are locked out for '**t2**' time when attempting with '**n**' incorrect passwords within '**t1**' time.

Locked-out users must either be cleared of lockout or wait for the lockout period to expire before attempting another login with a correct password. The feature is disabled by default.

To enable, use the following command:

```
aaa authentication policy lockout failure <number of failed attempts> window <within t1 time>
duration <lockout for t2 time>
```

- The value range for 'failure' is **1** to **255**.
- The value range for 'window' and 'duration' is **1** to **4294967295** seconds ($2^{32}-1$).

The example below locks out any user out for **15** minutes when attempting three incorrect logins within **3** minutes.

```
controller-1(config)# aaa authentication policy lockout failure 3 window 180 duration 900
```



Note:

- This feature affects only remote logins such as SSH/GUI/REST API using username/password. Console-based login and password-less authentications such as SSH keys, Single Sign-on, and access tokens are unaffected. Locked-out users can still access the Controller via console/password-less authentication.
- The feature is node-specific concerning the functionality, i.e., if user1 is locked out of accessing the active Controller in the cluster, they can log in to the standby Controller with the correct

password and vice versa. Lockout user information is not persistent across Controller reboot/fail-over.

To view if a user is locked out, admin-group users can run the following command:

```
show aaa authentication lockout
```

```
controller-1# show aaa authentication lockout
User name Host          Failed Logins  Lockout Date          Lockout Expiration
-----|-----|-----|-----|-----|
admin    10.240.88.193      1             2020-09-08 16:07:36.283000 PDT 2156-10-15 22:35:51.283000 PDT
```

To clear the lockout for a user, admin-group users can run the following command:

```
clear aaa authentication lockout user <username>
```

To clear all the locked-out users use the following command:

```
clear aaa authentication lockout
```

The following example shows how to clear a locked-out **admin** user:

```
controller-1# clear aaa authentication lockout user admin
controller-1# show aaa authentication lockout
None.
```

A **Recovery** user will also be locked out if attempting with incorrect passwords.

To check if the user is locked out, use the '**pam_tally2**' tool:

```
admin@controller-1:~$ sudo pam_tally2 -u recovery
Login Failures Latest failure From
recovery 9 09/08/20 16:16:04 10.95.66.44
```

To reset the lockout for the user, use the following command:

```
admin@controller-1:~$ sudo pam_tally2 --reset --user recovery
Login Failures Latest failure From
recovery 9 09/08/20 16:16:04 10.95.66.44
admin@controller-1:~$ sudo pam_tally2 -u recovery
Login Failures Latest failure From
recovery
```



Note: The **window** parameter does not apply to the **Recovery** user login as the '**pam_tally2**' tool does not support it.

11.5 Using TACACS+ to Control Access to the DMF Controller

Use remote Authentication, Authorization, and Accounting (AAA) services employing a TACACS+ server to control administrative access to the switch CLI. The following table lists the accepted Attribute-Value (AV) pairs:

Attributes	Values
BSN-User-Role	admin
	read-only
	bigtap-admin
	bigtap-read-only



Note: The remotely authenticated admin and bigtap-admin users and the read-only and bigtap-read-only users have the same privileges. The bigtap-admin and bigtap-read-only values are supported to allow the creation of DMF-specific entries without affecting the admin and read-only TACACS+ server entries.

A remotely authenticated admin user has full administrative privileges. Read-only users on the switch must be remotely authenticated. Read-only access is not configurable for locally authenticated user accounts.

Read-only users can only access login mode, from where they can view most show commands, with some limitations, including the following:

- TACACS, SNMP, and user configuration are not visible to the read-only user in the output from the **show running-config** command.
- **show snmp**, **show user**, and **show support** commands are disabled for the read-only user.



Note: Local authentication and authorization take precedence over remote authentication and authorization.

Configure privileges on the remote TACACS+ server using the following attribute-value pairs:

- Supported attribute name: BSN-User-Role
- Supported attribute values: admin, read-only,

Use the TACACS+ server to maintain administrative access control instead of using the Controller's local database. However, it is best practice to maintain the local database as the secondary authentication and authorization method in case the remote server becomes unavailable.

DANZ Monitoring Fabric requires the following configuration on TACACS+ servers in addition to the configuration required on the Controller.

Authentication Method

- Configure the TACACS+ server to accept ASCII authentication packets. Do not use the single connect-only protocol feature.
- The DANZ Monitoring Fabric TACACS+ client uses the ASCII authentication method. It does not use PAP.

Device Administration

- Configure the TACACS+ server to connect to the device administration login service.
- Do not use a network access connection method, such as PPP.

Group Memberships

- Create a bigtap-admin group. Make all DANZ Monitoring Fabric users part of this group.

- TACACS+ group membership is specified using the BSN-User-Role AVPair as part of TACACS+ session authorization.
- Configure the TACACS+ server for session authorization, not for command authorization.



Note: Specify the `BSN-User-Role` attribute as *Optional* in the `tac_plus.conf` file to use the same user credentials to access ANET and non-ANET devices.

11.5.1 Using the GUI to Add a TACACS+ Server

To identify a TACACS+ server to provide remote AAA services, complete the following steps. Repeat this procedure to identify up to five servers.

Procedure



Note: To set the secret and timeout values for all the TACACS+ servers, click the controls under the **Default Settings** section. Otherwise, set these values individually for each server.

Figure 11-13: Maintenance AAA

The screenshot shows the AAA Maintenance GUI. The top navigation bar includes Fabric, Monitoring, Maintenance, Integration, and Security. The main content area is titled 'AAA' and has tabs for Accounting, Authentication, Authorization, RADIUS, and TACACS. Under the TACACS tab, there is a 'Default Settings' section with fields for 'Timeout' (5 seconds) and 'Secret' (empty). Below this is a 'TACACS Servers' table with columns for Actions, Server Address, State, and Last Updated. A table with one row is visible, showing a server at 192.168.1.10 in a 'Configured' state, last updated on Apr 12, 2022. There are also filter and pagination controls at the bottom of the table.

1. Click the **Add TACACS Server** from the **Actions** menu in the TACACS Servers table.

Figure 11-14: Create TACACS+ Server Dialog

The dialog box is titled 'Add TACACS Server' and has a close button (X) in the top right corner. It contains two input fields: 'Server Address' with a question mark icon and a value of '192.168.1.20', and 'Secret' with a question mark icon and an empty field. At the bottom, there are two buttons: 'Cancel' and 'Submit'.



Note: Do not use the pound character (#) in the TACACS secret, as it will be interpreted as the start of a comment in the PAM config file.

2. Enter the IP address of the TACACS+ server.
3. Type the password required to access the server in the Secret field.



Note: Click the lock icon to encrypt the password if plain-text passwords are not used in the AAA environment.

4. Click **Submit**.

11.5.2 Using the CLI to Enable Remote Authentication and Authorization on the DMF Fabric Controller

CLI Procedure

Use the following commands to configure remote 'login' authentication and authorization. The examples use the SSH default for connection type.

```
controller-1(config)# tacacs server host 10.2.3.201
controller-1(config)# aaa authentication login default group tacacs+ local
controller-1(config)# aaa authorization exec default group tacacs+ local
```

As a result, all users in the bigtap-admin group on TACACS+ server **10.2.3.201** have full access to the DANZ Monitoring Fabric Controller.

11.5.3 Using the CLI to Add a TACACS+ Server

To view the current TACACS+ configuration, enter the `show running-config tacacs` command.

Complete the following steps to configure the DMF Controller with TACACS+ to control administrative access to the switch.

1. Identify the IP address of the TACACS+ server and any key required for access using the `tacacs server` command with the following syntax:

```
tacacs server host <server> key {<plaintext-key> | 0 <plaintext-key> | 7 <encrypted-key>}
```

Enable up to four AAA servers by repeating this command for each server. For example, using a plaintext key, the following command enables TACACS+ with the server running at **10.2.3.4**.

```
controller-1(config)# tacacs server 10.2.3.4 key 0 secret
```

If the key is omitted, an empty key is used.



Note: Do not use the pound character (#) in the TACACS+ secret, as it will be interpreted as the start of a comment in the PAM config file.

2. Encrypt each TACACS+ server connection using a pre-shared key. To specify a key for a specific host, use one of the following:

```
controller-1(config)# tacacs server host <ip-address> key <plaintextkey>
controller-1(config)# tacacs server host <ip-address> key 0 <plaintextkey>
controller-1(config)# tacacs server host <ip-address> key 7 <plaintextkey>
```

Replace *plaintextkey* with a password up to 63 characters in length. This key can be specified either globally, or for each individual host. The first two forms accept a plaintext (literal) key, and the last form accepts a pseudo-encrypted key, such as that displayed with `show running-config`.

The following is an example using the key 7 option followed by the encrypted string:

```
controller-1(config)# tacacs server 10.2.3.4 key 7 0832494d1b1c11
```

To configure a global key, use the following command:

```
controller-1(config)# tacacs server key 0 secret
```

The global key value is used if no key is specified for a given host. If no key is specified globally and no key is specified for a given host, then an empty key is assumed.



Note: Be careful configuring TACACS+ to avoid disabling access to the DANZ Monitoring Fabric (DMF) Controller.

3. Configuring device-specific TACACS+ server parameters overriding that of the global TACACS+ servers is possible. This applies to switches, service nodes, and recorder nodes. The configuration must be made from the config-device submenu. An example to configure a switch-specific TACACS+ server is described below:

```
controller-1(config)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch)# tacacs override-global
controller-1(config-switch)# tacacs server host 1.1.1.1 key 7 020700560208
```

Similarly, TACACS+ server keys and timeout values can also be overridden:

```
controller-1(config)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch)# tacacs override-global
controller-1(config-switch)# tacacs server timeout 8
controller-1(config-switch)# tacacs server key 0 qwerty
```

To move back to using the globally defined TACACS+ servers, run the `no tacacs override-global` command at the *config-switch* submenu.

4. To view the TACACS+ configuration on a specific switch, use the `show effective-config switch switch-name tacacs` command:

```
controller-1(config-switch)# show effective-config switch DMF-DELIVERY-SWITCH-1 tacacs
! switch
switch DMF-DELIVERY-SWITCH-1
tacacs server host 1.1.1.1 key 7 020700560208
tacacs server timeout 8
```

The TACACS+ key value displays as a `type7` secret instead of plaintext.

11.6 Setting up a Tac_Plus Server

After installing the `tac_plus` server, complete the following steps to set up authentication and authorization for the DMF Controller with the TACACS server:

1. Configure users and groups.

2. In the `/etc/tacacs/tac_plus.conf` file, specify the user credentials and group association.

```
# user details
user = user1 {
member = anet-vsa-admin
login = des a9qtD2JXeK0Sk
}
```

3. Configure the groups to use one of the AV pairs supported by the DMF Controller (for example, `BSN-User-Role="admin"` for admin users).

```
# group details
# ANET admin group
group = anet-vsa-admin {
service = exec {
BSN-User-Role="admin"
}
}
# ANET read-only group
group = anet-vsa-read-only {
service = exec {
BSN-User-Role="read-only"
}
}
```



Note: Different TACACS servers need different ways to define the attributes. The following is an example of configuring an Aruba Clearpass server.

```
<TacacsServiceDictionaries>
<TacacsServiceDictionary dispName="Big Switch Networks" name="shell:ip">
<ServiceAttribute dataType="String" dispName="BSN User Role" name="BSN-User-Role"/>
</TacacsServiceDictionary>
```

4. Configure the TACACS+ server and AAA on the DMF Controller.

```
tacacs server host <IP address> key server's secret>
aaa authentication login default group tacacs+ local
aaa authorization exec default group tacacs+ local
aaa accounting exec default start-stop locals group tacacs+
```

This configuration allows authentication and authorization to connect to the TACACS+ server to verify user credentials and privileges. Checking the user account locally only occurs when the remote server is unreachable. In this example, accounting is set to store audit logs locally and send them to the remote server.

Refer to your AAA server documentation for further details or instructions for setting up other servers.

11.6.1 Using the Same Credentials for DMF and Other Devices

To use the same user credentials to access DMF and a non-DMF device, the `BSN-User-Role` attribute must be specified as *Optional* in the `tac_plus.conf` file, as shown in the following example.

```
group = group-admin {
default service = permit
service = exec {
optional BSN-User-Role = "admin"
}
}
```

11.6.2 RBAC-Based Configuration for Non-Default Group User

To create an RBAC configuration for a user in a non-default group, complete the following steps:

1. Create a group AD1.

```
group AD1
```

Do not associate any local users.

2. Use the same group name on the TACACS+ server and associate a user to this group.



Note: The attribute should be **BSN-User-Role**, and the value should be the **group-name**.

The following is an example from the open TACACS+ server configuration.

```
group = AD1 {  
  service = exec {  
    BSN-User-Role="AD1"  
  }  
}
```

3. After you create the group, associate a user to the group.

```
user = user3 {  
  member = AD1  
  login = cleartext user3  
}
```

11.7 Using RADIUS for Managing Access to the DMF Controller

By default, the Authentication and Authorization functions are set to local while the Accounting function is disabled. The only supported privilege levels are as follows:

- admin: Administrator access, including all CLI modes and debug options.
- read-only: Login access, including most show commands.



Note: RADIUS does not separate authentication and authorization, so be careful when authorizing a user account using a remote RADIUS server to use the correct password configured for the user on the remote server.

The admin group provides full access to all network resources, while the read-only group provides read-only access to all network resources.



Note: The admin and recovery user accounts cannot be authenticated remotely using RADIUS. These accounts are always authenticated locally to prevent administrative access from being lost in case a remote AAA server is unavailable.

DANZ Monitoring Fabric also supports remote AAA server (RADIUS) communication. The following summarizes the options available for each function:

- Accounting: local, local and remote, or remote.
- Authentication: local, local then remote, remote then local, or remote.

- Authorization: local, local then remote, remote then local, or remote.



Note: Fallback to local authentication occurs only when the remote server is unavailable, not when authentication fails.

Configure privileges on the remote RADIUS server using the attribute-value pairs shown in the following table:

Supported attribute names	Supported attribute values
BSN-User-Role	admin bigtap-admin read-only bigtap-read-only

The BSN-AV-Pair attribute sends CLI command activity accounting to the RADIUS server.

11.7.1 Using the GUI to Add a RADIUS Server

To identify a RADIUS server to provide remote AAA services, complete the following steps. Repeat this procedure to identify up to five servers.

GUI Procedure

1. Select **Maintenance > AAA > RADIUS** from the GUI.

Figure 11-15: Maintenance AAA: Radius Section

The screenshot displays the AAA configuration interface. At the top, there are navigation tabs: Fabric, Monitoring, Maintenance, Integration, and Security. The 'Maintenance' tab is active, and the 'RADIUS' sub-tab is selected. Below the navigation, the 'AAA' section is shown, with 'Last updated: a minute ago'. Underneath, there are sub-tabs for Accounting, Authentication, Authorization, RADIUS, and TACACS. The 'Default Settings' section includes a table with the following data:

Parameter	Value
Timeout	5 seconds
Secret	-
Auth Port	1812
Acct Port	1813

Below the default settings is the 'RADIUS Servers' section. It features a search bar, 'Add Filters', and a 'Server Address' filter. A table lists the configured RADIUS servers:

Actions	Server Address	State	Last Updated
<input type="checkbox"/>	10.240.130.59	Configured	Apr 12, 2022 09:21:32:613pm

At the bottom right, there is a pagination control showing '1-1 of 1 items' and '10 / page'.

2. Click the **Add RADIUS Server** from the **Actions** menu in the RADIUS Servers table.

Figure 11-16: Create RADIUS Server Dialog

Add RADIUS Server X

* Server Address ⓘ

192.168.1.10

Secret ⓘ

Cancel Submit

3. Enter the IP address of the RADIUS server.
4. Type the password required to access the server in the Secret field.



Note: Click the lock icon to encrypt the password if plain-text passwords are not used in your AAA environment.

5. Click **Submit**.

11.7.2 Using the CLI to Add a RADIUS Server

Use the following command to identify the remote RADIUS server:

```
radius server host <server-address> [timeout <{timeout}>][key <{<plaintext>} | 0 <{plaintext}> | 7  
{  
<secret>}]
```

For example, the following command identifies the RADIUS server at the IP address **192.168.17.101**.

```
controller-1(config)# radius server host 192.168.17.101 key admin
```

You can enter this command up to five times to identify multiple RADIUS servers. The Controller tries to connect to each server in the order they are configured.

11.7.3 Setting up a FreeRADIUS Server

After installing the FreeRADIUS server, complete the following steps to set up authentication and authorization for the DMF Controller with the RADIUS server.

Procedure

1. Create the BSN dictionary and add it to the list of used dictionaries.

```
create dictionary /usr/share/freeradius/dictionary.arista with the contents below:  
VENDOR Big-Switch-Networks 37538  
BEGIN-VENDOR Arista-Networks  
ATTRIBUTE BSN-User-Role 1 string
```

```
ATTRIBUTE BSN-AVPair 2
string
END-VENDOR Arista-Networks
```



Note: Make sure that the `BSN-User-Role` attribute is in the first position in the Radius attribute dictionary, followed by `BSN-AVPair`.

2. Include the arista dictionary in the radius dictionary file: `/usr/share/freeradius/dictionary`

```
$INCLUDE dictionary.arista
```

3. Configure a sample user with admin and read-only privileges.

The following is an example that defines and configures a user, opens the user file `/etc/freeradius/users`, and inserts the following entries:



Note: This example shows how the VSA is associated with the user and its privileges. In an actual deployment, a database and encrypted password are necessary.

```
"user1" Cleartext-Password := "passwd"
BSN-User-Role := "read-only",
```

The following example authorizes `user2` for RBAC group `AD1`:

```
"user2" Cleartext-Password := "passwd"
BSN-User-Role := "AD1",
```

The following example authorizes `user3` for RBAC group `admin`:

```
"user3" Cleartext-Password := "passwd"
BSN-User-Role := "admin",
```

4. Configure the RADIUS server and AAA on the DMF Controller.

```
radius server host <IP address> key server's secret>
aaa authentication login default group radius local
aaa authorization exec default group radius local
aaa accounting exec default start-stop group radius local
```

This configuration allows authentication and authorization to connect to the RADIUS server to verify user credentials and privileges. AAA fallback to local occurs only when the remote server is unreachable. In this example, accounting is set to store audit logs locally and send them to the remote server.

5. Add the DMF Controller subnet to the allowed subnets (`clients.conf`) on the RADIUS server.

This entry is required if access to the RADIUS server is limited to allowed clients/subnets. The following is an example of the `clients.conf` file:

```
client anet {
ipaddr = 10.0.0.0/8
secret = <server's secret>
}
```

6. Restart the FreeRADIUS service on the server to enable the configuration.

The following is an example accounting record sent from the Controller to the RADIUS server after adding the `BSN-AVPair` attribute to the `/usr/share/freeradius/dictionary.arista` file.

```
root@radius-bsn1:/var/log/freeradius/radacct/10.8.41.11# tail -f detail-20180123
```

```
Tue Jan 23 17:48:22 2018
Acct-Session-Id = "Session@9c7e1872"
Acct-Status-Type = Interim-Update
BSN-AVPair = "auth_description=session/
9c7e18722d6b9334ff6cac2924ae06baa4cdc6b53756e93120145cfd7712b466"
User-Name = "admin"
BSN-AVPair = "remote_address=10.1.2.86"
BSN-AVPair = "cmd_args=show version"
NAS-IP-Address = 10.8.41.11
Acct-Unique-Session-Id = "8d80262884a1abde"
Timestamp = 1516729702
```

11.8 Custom admin role

Use the custom admin roles feature, which allows the definition of a group that grants the privilege to read or write like a built-in admin group, but only for specific use cases/categories. Configure the required privileges for each predefined set of categories.

This configuration allows a user to define a more flexible group role, such as a group that will enable its users to read and write most of the general configuration like admin, but not the AAA-related parts of the configuration.

This approach helps to create user profiles on DANZ Monitoring Fabric with limited admin access restricted only for the required features.

11.8.1 Categories

Users can be created and associated to groups with the following categories:

category:AAA

Configuration and state related to AAA, including but not limited to local user management, group management *AAA group creation and user association*.

```
Controller-1> enable
Controller-1# configure
Controller-1(config)# group aaa-mgmt-group
Controller-1(config-group)# permission category:AAA privilege read-write
Controller-1(config)# user aaa-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group aaa-mgmt-group
Controller-1(config-group)# associate user aaa-user
AAA group user sample configuration
AAA group user sample configuration
Controller-1(config)# aaa authentication login default local
Controller-1(config)# tacacs server host <REMOTE_TACACS_SERVER_IP>
```

```
Controller-1(config)# show group aaa-mgmt-group configured-permission
# Permission Privilege
1 category:AAA read-write
Controller-1(config)# show group aaa-mgmt-group effective-permission
# Effective-permission Inferred privilege
- |-----|-----|
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA read-write
4 category:AAA/SENSITIVE does-not-elevate
5 category:APPLLOG inferred-does-not-elevate
6 category:SYSOPS inferred-does-not-elevate
```

```

7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate

```

When AAA group user tries to configure sensitive information

category:AAA/SENSITIVE:

Sensitive data included in configuration and state related to AAA, including but not limited to local user management, group management.

AAA/SENSITIVE group creation and user association

```

Controller-1> enable
Controller-1# configure
Controller-1(config)#
Controller-1(config)# group aaa-sen-mgmt-group
Controller-1(config-group)# permission category:AAA/SENSITIVE privilege read-write
Controller-1(config)# user aaa-sen-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group aaa-mgmt-group
Controller-1(config-group)# associate user aaa-sen-user
Controller-1(config)# show group aaa-sen-mgmt-group configured-permission
# Permission Privilege
1 category:AAA/SENSITIVE read-write
Controller-1(config)# show group aaa-sen-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA read-write
4 category:AAA/SENSITIVE inferred-read-write
5 category:APPLLOG inferred-does-not-elevate
6 category:SYSOPS inferred-does-not-elevate
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate
Controller-1(config)#

```

AAA group user sample configuration

```

Controller-1(config)# tacacs server host 10.240.176.159 key 12323243
Controller-1(config)#

```

category:TIME:

Configuration and state related to system time and NTP. *TIME* group creation and user association

```

Controller-1> enable
Controller-1# configure
Controller-1(config)#
Controller-1(config)# group time-mgmt-group
Controller-1(config-group)# permission category:TIME privilege read-write
Controller-1(config)# user time-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group time-mgmt-group
Controller-1(config-group)# associate user time-user

```

```

Controller-1(config)# show group time-mgmt-group configured-permission
# Permission Privilege
1 category:TIME read-write
Controller-1(config)# show group time-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate

```



```

Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group sysops-mgmt-group
Controller-1(config-group)# associate user sysops-user

```

```

Controller-1(config)# show group sysops-mgmt-group configured-permission
# Permission Privilege
1 category:SYSOPS read-write
Controller-1(config)# show group sysops-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA inferred-does-not-elevate
4 category:AAA/SENSITIVE inferred-does-not-elevate
5 category:APPLLOG inferred-does-not-elevate
6 category:SYSOPS read-write
7 category:SYSOPS/SENSITIVE does-not-elevate
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate

```

SYSOPS group user sample configuration

```

Controller-1(config)# boot partition alternate
Controller-1(config)#

```

category:SYSOPS/SENSITIVE:

Sensitive data included in Configuration and state related to system time and NTP.

SYSOPS group creation and user association

```

Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group sysops-mgmt-group
Controller-1(config-group)# permission category:sysops/SENSITIVE privilege read-write
Controller-1(config)# user sysops-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group sysops-mgmt-group
Controller-1(config-group)# associate user sysops-user

```

```

Controller-1(config)# show group sysops-mgmt-group configured-permission
# Permission Privilege
1 category:SYSOPS/SENSITIVE read-write
Controller-1(config)# show group sysops-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA inferred-does-not-elevate
4 category:AAA/SENSITIVE inferred-does-not-elevate
5 category:APPLLOG inferred-does-not-elevate
6 category:SYSOPS read-write
7 category:SYSOPS/SENSITIVE read-write
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate

```

SYSOPS group user sample configuration

```

Controller-1(config)# connect switch dell-5048-253-ru30
Controller-1(config)#

```

category:APPLLOG:

Configuration and state related to appliance log level.

APPLOG group creation and user association

```
Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group APPLOG-mgmt-group
Controller-1(config-group)# permission category:APPLOG privilege read-only
Controller-1(config)# user APPLOG-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group APPLOG-mgmt-group
Controller-1(config-group)# associate user APPLOG-user
```

```
Controller-1(config)# show group APPLOG-mgmt-group configured-permission
# Permission Privilege
1 category:APPLOG read-write
Controller-1(config)# show group APPLOG-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA inferred-does-not-elevate
4 category:AAA/SENSITIVE inferred-does-not-elevate
5 category:APPLOG read-write
6 category:SYSOPS inferred-does-not-elevate
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate
```

APPLOG group user sample configuration

```
Controller-1(config)# show logging controller
Controller-1(config)#
```

category:DEFAULT

General configuration and states

DEFAULT group creation and user association

```
Controller-1> enable
Controller-1# configure
Controller-1(config)# group DEFAULT-mgmt-group
Controller-1(config-group)# permission category:DEFAULT privilege read-write
Controller-1(config)# user DEFAULT-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group DEFAULT-mgmt-group
Controller-1(config-group)# associate user DEFAULT-user
```

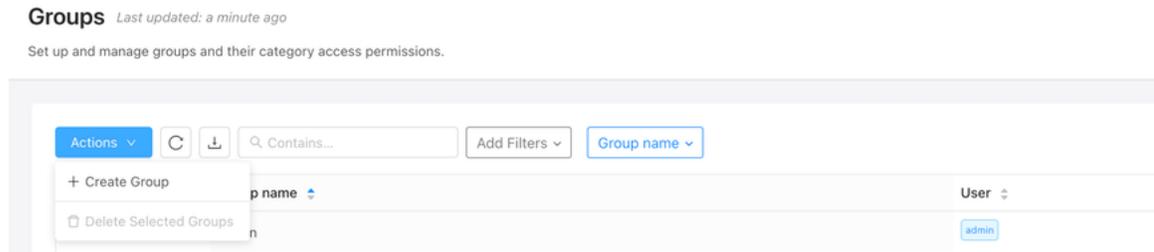
```
Controller-1(config)# show group DEFAULT-mgmt-group configured-permission
# Permission Privilege
1 category:DEFAULT read-write
Controller-1(config)# show group DEFAULT-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT read-write
2 category:DEFAULT/SENSITIVE inferred-read-write
3 category:AAA inferred-read-write
4 category:AAA/SENSITIVE inferred-read-write
5 category:APPLOG inferred-read-write
6 category:SYSOPS inferred-read-write
7 category:SYSOPS/SENSITIVE inferred-read-write
8 category:TIME inferred-read-write
9 category:TIME/SENSITIVE inferred-read-write
```

GUI

Navigate to the following page to create the custom admin groups from GUI.

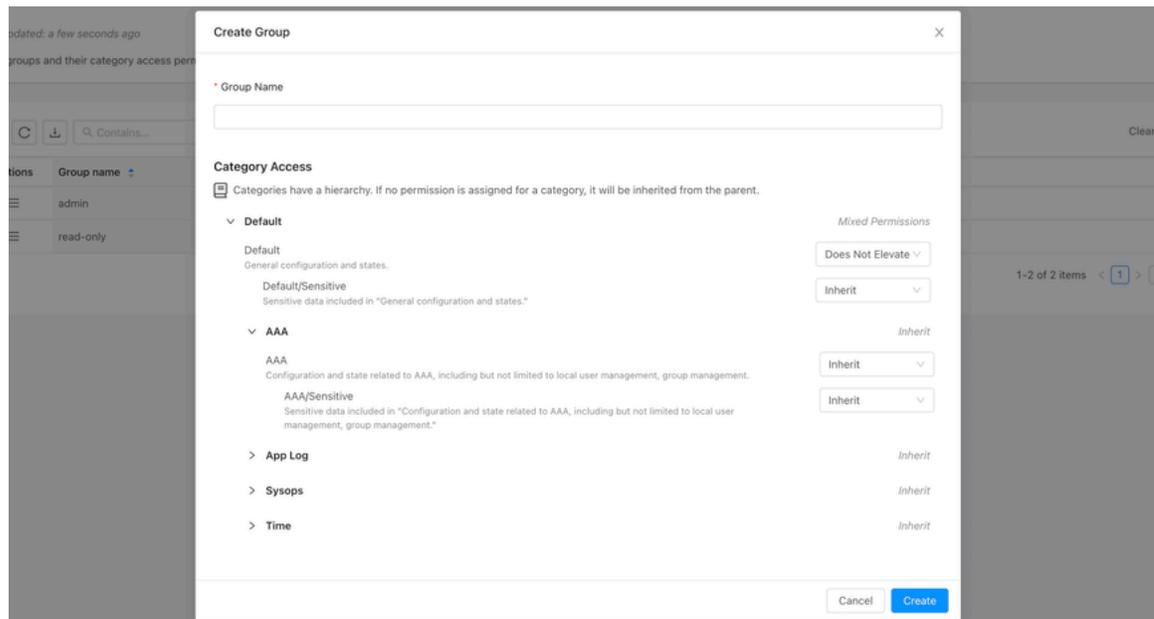
1. Login to DMF.

Figure 11-17: Group Creation



2. Proceed to **Security > Groups** and add the required settings.

Figure 11-18: Custom Group



11.8.2 Permissions & Privileges

Privileges are the permissions set of each category.

- **Read Only:** Grants the user privileges to read only the configurations for the enabled category.
- **Read-Write:** Grants the user read and write privileges for configurations for the enabled category.
- **Does Not Elevate:** The user can neither read nor write configurations for that category. Any category, such as AAA, TIME, etc., configured with **Does Not Elevate** will not inherit from the parent category.
- **Inherit:** Categories will inherit their permission from their parent category. The DEFAULT category is considered the root of all categories.

For example, if **TIME/SENSITIVE** is set to **inherit** and **TIME** has **read-only**, then users associated with this group will inherit from the parent category, and it will effectively have the privilege of **read-only** for **TIME/SENSITIVE**.

Similarly, if **AAA** is set to **inherit** and **DEFAULT** has **read-only**, then users associated with this group can read all configurations related to AAA.

11.8.3 Category Identification

DMF helps identify which feature comes under what category, achieved using the following command.

Syntax: :: **category <feature_name>**

For example, the AAA feature.

```
Controller-1# configure
Controller-1(config)#
Controller-1(config)# category aaa authentication login default local
syntax of variant:
aaa authentication login default {local | group {tacacs+ | radius} | local group {tacacs+ |
radius}
| group {tacacs+ | radius} local}
no aaa authentication login default [local | group {tacacs+ | radius} | local group {tacacs+ |
radius} | group {tacacs+ | radius} local]
Configure authentication parameters
AAA "aaa authentication login default local"
Although the category can be configured, allow-all read access is enabled for the command
Exact matching commands: 1, similar commands (same prefix): 1
```

In the above output, the AAA "aaa authentication login default local" line signifies that the above feature comes under the AAA category.

11.8.4 Group Management

View Group permissions:

```
NS-178-37(config-group)# group aaa-user
Controller-1(config)# permission category:AAA privilege read-write
Controller-1(config)# show group aaa-user configured-permission
# Permission Privilege
1 category:AAA read-write
Controller-1(config)# show group aaa-user effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT inferred-does-not-elevate
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate
3 category:AAA read-write
4 category:AAA/SENSITIVE inferred-read-write
5 category:APPLLOG inferred-does-not-elevate
6 category:SYSOPS inferred-does-not-elevate
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME inferred-does-not-elevate
9 category:TIME/SENSITIVE inferred-does-not-elevate
Controller-1(config)#
```

Summary:

Category: AAA is set as **read-write** privilege as per the user configuration Category: AAA/SENSITIVE is set as **read-write** since it inherits from the immediate parent, i.e., Category: AAA which is "read-write".

Rest of the categories are set as **does-not-elevate** since they inherit from Category: DEFAULT

Example to configure a user part of all categories except AAA:

```

Controller-1(config)# group all-except-AAA-group
Controller-1(config-group)# permission category:DEFAULT privilege read-write
Controller-1(config-group)# permission category:AAA privilege does-not-elevate
Controller-1(config-group)# show group all-except-AAA-group configured-permission
# Permission Privilege
1 category:AAA does-not-elevate
2 category:DEFAULT read-write
Controller-1(config-group)# show group all-except-AAA-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT read-write
2 category:DEFAULT/SENSITIVE inferred-read-write
3 category:AAA does-not-elevate
4 category:AAA/SENSITIVE inferred-does-not-elevate
5 category:APLOG inferred-read-write
6 category:SYSOPS inferred-read-write
7 category:SYSOPS/SENSITIVE inferred-read-write
8 category:TIME inferred-read-write
9 category:TIME/SENSITIVE inferred-read-write
Controller-1(config-group)#

```

Summary:

Category: DEFAULT is set as “read-write” privilege as per the user configuration. Category: AAA is set as “does-not-elevate” privilege as per the user configuration. Category: AAA/SENSITIVE is set as “does-not-elevate” since it inherits from the immediate parent, i.e., Category: AAA which is “does-not-elevate”. The rest of the categories are set as “read-write” since they inherit from Category: DEFAULT.

11.8.5 Remote Users

With Custom Admin, associate users with the groups with required privileges using remote TACACS+ and RADIUS servers. To do that, follow these steps:

1. Create a user **bns-user1** and a desired password.
2. Create a group “**BSN-User-Role**” with the required categories and privileges and associate the above user to the group.
3. Configure the username and password, and the group created steps 1 and 2 on the TACACS/ RADIUS servers.
4. Login to the DANZ Monitoring Fabric (DMF) Controller using the above credential.

11.8.6 Commonly used profiles

User Administration

Users belonging to this group can only manage users and groups.

```

Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group user-administration
Controller-1(config-group)# permission category:AAA privilege read-write
Controller-1(config-group)# permission category:DEFAULT: does-not-elevate
Controller-1(config)# user time-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group time-mgmt-group

```

```
Controller-1 (config-group) # associate user time-user
```

```
Controller-1 (config) # show group user-administration configured-permission  
# Permission Privilege  
1 category:AAA read-write  
2 category:DEFAULT does-not-elevate
```

```
Controller-1 (config) # show group user-administration effective-permission  
# Effective-permission Inferred privilege  
1 category:DEFAULT does-not-elevate  
2 category:DEFAULT/SENSITIVE inferred-does-not-elevate  
3 category:AAA read-write  
4 category:AAA/SENSITIVE inferred-read-write  
5 category:APPLOG inferred-does-not-elevate  
6 category:SYSOPS inferred-does-not-elevate  
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate  
8 category:TIME inferred-does-not-elevate  
9 category:TIME/SENSITIVE inferred-does-not-elevate
```

Network-admin

Users belonging to this group will have admin without the privilege to manage AAA.

```
Controller-1> enable  
Controller-1# configure terminal  
Controller-1 (config) #  
Controller-1 (config) # group network-admin  
Controller-1 (config-group) # permission category:AAA privilege does-not-elevate  
Controller-1 (config-group) # permission category:DEFAULT: read-write  
Controller-1 (config) # user time-user  
Controller-1 (config-user) # password <user_configured_password>  
Controller-1 (config-user) # group time-mgmt-group  
Controller-1 (config-group) # associate user time-user
```

```
Controller-1 (config) # show group network-admin configured-permission  
# Permission Privilege  
1 category:AAA does-not-elevate  
2 category:DEFAULT read-write  
Controller-1 (config) # show group network-admin effective-permission  
# Effective-permission Inferred privilege  
1 category:DEFAULT read-write  
2 category:DEFAULT/SENSITIVE inferred-read-write  
3 category:AAA does-not-elevate  
4 category:AAA/SENSITIVE inferred-does-not-elevate  
5 category:APPLOG inferred-read-write  
6 category:SYSOPS inferred-read-write  
7 category:SYSOPS/SENSITIVE inferred-read-write  
8 category:TIME inferred-read-write  
9 category:TIME/SENSITIVE inferred-read-write
```

Auditor

Users belonging to this group can read everything for auditing.

```
Controller-1> enable  
Controller-1# configure terminal  
Controller-1 (config) #  
Controller-1 (config) # group auditor  
Controller-1 (config-group) # permission category:DEFAULT: read-only  
Controller-1 (config) # user time-user  
Controller-1 (config-user) # password <user_configured_password>  
Controller-1 (config-user) # group time-mgmt-group  
Controller-1 (config-group) # associate user time-user
```

```
Controller-1 (config) # show group auditor configured-permission
```

```

# Permission Privilege
1 category:DEFAULT read-only
Controller-1(config)# show group auditor effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT read-only
2 category:DEFAULT/SENSITIVE inferred-read-only
3 category:AAA inferred-read-only
4 category:AAA/SENSITIVE inferred-read-only
5 category:APPLLOG inferred-read-only
6 category:SYSOPS inferred-read-only
7 category:SYSOPS/SENSITIVE inferred-read-only
8 category:TIME inferred-read-only
9 category:TIME/SENSITIVE inferred-read-only

```

Filtered-auditor

Users belonging to this group can audit everything but cannot see sensitive data.

```

Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group filtered-auditor
Controller-1(config-group)# permission category:DEFAULT: read-only
Controller-1(config-group)# permission category:AAA/SENSITIVE privilege does-not-elevate
Controller-1(config-group)# permission category:DEFAULT/SENSITIVE privilege does-not-elevate
Controller-1(config-group)# permission category:TIME/SENSITIVE privilege does-not-elevate
Controller-1(config-group)# permission category:SYSOPS/SENSITIVE privilege does-not-elevate
Controller-1(config)# user time-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group time-mgmt-group
Controller-1(config-group)# associate user time-user

```

```

Controller-1(config)# show group filtered-auditor configured-permission
# Permission Privilege
1 category:DEFAULT read-only
Controller-1(config)# show group filtered-auditor effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT read-only
2 category:DEFAULT/SENSITIVE does-not-elevate
3 category:AAA inferred-read-only
4 category:AAA/SENSITIVE does-not-elevate
5 category:APPLLOG inferred-read-only
6 category:SYSOPS inferred-read-only
7 category:SYSOPS/SENSITIVE does-not-elevate
8 category:TIME inferred-read-only
9 category:TIME/SENSITIVE does-not-elevate

```

11.8.7 Commonly used Category-feature Matrix

Table 3: Software Requirements for DANZ Monitoring Fabric

CATEGORY	FEATURE
AAA	TACACS
AAA/SENSITIVE	TACAS-PASSWORD
AAA	RADIUS
AAA/SENSITIVE	RADIUS-PASSWORD
AAA	ACCCOUNTING
AAA	AUTHENTICATION
AAA	AUTHORIZATION
AAA	USERS
AAA	CLEAR AAA parameters
AAA	CLEAR SESSION parameters
TIME	TIME-Time Zone
TIME	TIME-NTP-servers
TIME/SENSITIVE	TIME-NTP-KEYS
APPLOG	Show logging commands
SYSOPS/SENSITIVE	connect
SYSOPS	boot
SYSOPS	clear async
SYSOPS	Delete dump files
SYSOPS	Reload Controller
DEFAULT	auto-vlan-mode
DEFAULT	auto-vlan-range
DEFAULT	Clear debug counters
DEFAULT	Clear statistics
DEFAULT	Policy
DEFAULT	System restart local-node
DEFAULT	filter-interface-group

Known Limitations

- Viewing audit logs is only possible through the built-in admin user.
- There must be a built-in admin user to upgrade the DMF Controller and managed appliances.
- RBAC has a higher preference over custom admin groups.
- The custom group feature doesn't apply to DMF Managed Appliances.

Management and Control Plane Security

This chapter describes options for increasing the security of management access to the DMF Controller node.

12.1 Management Plane Security

The management plane network is used by the administrator, whether locally or remotely, to reach the Controller management interfaces. DANZ Monitoring Fabric (DMF) uses standard, well-known cryptographic technology, such as RSA and AES. Still, system administrators must choose strong passwords and change them frequently, according to well-established security best practices.

All services the Controller uses are enabled by default except for SNMP, which is disabled by default. Refer to the [Protocol Access Required to the DMF Controller](#) section to block or permit specific protocols to the management interface.

For example, the control plane is the network between the Controllers and the switches to carry OpenFlow control traffic. The following are general requirements and recommendations for deployment:

- The Controller must be on the same Layer 2 network as the switches—physically isolated data, control, and management plane networks.
- The only devices on the control plane are switches and Controllers.
- Make the control plane network not routed or minimally IP access restricted via its egress router.
- Physically secure the management and data plane networks (for example, locks on the cage doors).

Many of the Zero-Touch Networking (ZTF) protocols (DHCP, such as ONIE, Controller discovery, and image download) and the OpenFlow protocol are not authenticated. They are subject to spoofing in an untrusted network. The following are best practices regarding securing the control plane within the switched fabric.

- The control plane network is “Layer 2 trusted,” meaning the attacker cannot spoof Layer 2 messages on the control network. In practice, this means the control plane network should be an isolated VLAN, ideally containing only the Controller and switches.
- Harden the switch management interface against Layer 3 attacks (all services are authenticated, unnecessary services are turned off, and so forth).
- The network should not be reachable by Layer 3 protocols. If Layer 3 access is required, the administrator should maintain a Layer 3 allowlist of hosts that can access the control network, for example, using an ACL on the edge router.

12.2 Importing the Controller Private Key and Certificate

This section describes how to import a private key and a certificate to the Controller after copying it to the Controller using the `copy` command.

To import a private key to the Controller, enter the `private-key` command in the `config-controller` submode:

```
[no] private-key <controller-key-name>
```

Replace `controller-key-name` with the name of the private key. Use the `no` version of the command to remove the `private-key`.

To import the Controller certificate, use the `certificate` command in `config-controller` submode.

```
[no] certificate <name>
```

Replace the `name` with the name assigned to the Controller certificate. Use the `no` version of the command to remove the `certificate`.

Import the private key and certificate to the Controller using the `copy` command.

12.3 Using Certificates Signed by a CA for GUI Access to the Controller

By default, SSL is enabled on the Controller using a self-signed certificate. Complete the following steps to install a certificate signed by a public or private CA.

Procedure

1. Generate the Certificate Signing Request (CSR) and the private key for the Controller.

Perform this operation on any workstation that supports OpenSSL. The following example shows the operation performed on a Linux workstation.

```
root@Ubuntu-12:~/openssl-ca/admin# openssl req -newkey rsa:2048 -nodes -keyout controller.  
key -new -out controller.csr  
Generating a 2048 bit RSA private key  
.....+++  
.....+++  
writing new private key to 'controller.key'  
-----  
You are about to be asked to enter information that will be incorporated into your  
certificate request.  
What you are about to enter is what is called a Distinguished Name or a DN. There are  
quite a few fields but you can leave some blank  
For some fields there will be a default value,  
If you enter '.', the field will be left blank.  
-----  
Country Name (2 letter code) [AU]:US  
State or Province Name (full name) [Some-State]:California  
Locality Name (eg, city) []:Santa Clara  
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Arista Networks Organizational  
Unit Name (eg, section) []:Engineering  
Common Name (e.g. server FQDN or YOUR name) []:DMF Secure Certificate  
Email Address []:admin@arista.com  
Please enter the following 'extra' attributes  
to be sent with your certificate request  
A challenge password []:anet1234  
An optional company name []:Arista  
root@Ubuntu-12:~/openssl-ca/admin#
```

```

root@Ubuntu-12:~/openssl-ca/admin# ls -ltr
total 8
-rw-r--r-- 1 root root 1708 Feb 7 15:39 controller.key
-rw-r--r-- 1 root root 1184 Feb 7 15:39 controller.csr
root@Ubuntu-12:~/openssl-ca/admin#

```

2. Submit the CSR to the CA and get the certificate signed.

Submit the CSR to the trusted CA for browsers used to access the DMF GUI. For organizations using GUI-based CAs, copy the contents of the CSR to the CA for signature.

The following example shows the operation performed on a Linux workstation.

```

root@Ubuntu-12:~/openssl-ca# openssl ca -config openssl-ca.cnf -policy signing_policy -
extensions signing_req -out admin/controller.pem -infile admin/controller.csr
Using configuration from openssl-ca.cnf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName :Printable:'US'
stateOrProvinceName :ASN.1 12:'California'
localityName :ASN.1 12:'Santa Clara'
organizationName :ASN.1 12:'Arista Networks'
organizationalUnitName:ASN.1 12:'Engineering'
commonName :ASN.1 12:'DMF Secure Certificate'
Certificate is to be certified until Nov 3 23:41:17 2020 GMT (1000 days) Sign the
certificate? [y/n]:y
1 out of 1 certificate requests certified, commit? [y/n]y Write out database with 1 new
entries
Data Base Updated
root@Ubuntu-12:~/openssl-ca#
root@Ubuntu-12:~/openssl-ca/admin# ls -ltr
total 16
-rw-r--r-- 1 root root 1708 Feb 7 15:39 controller.key
-rw-r--r-- 1 root root 1184 Feb 7 15:39 controller.csr
-rw-r--r-- 1 root root 5882 Feb 7 15:41 controller.pem
root@Ubuntu-12:~/openssl-ca/admin#

```

3. Copy the signed certificate to the Controller:

```

controller-1# copy scp://root@10.8.67.3:/root/openssl-ca/admin/controller.pem cert://
root@10.8.67.3's password:
controller.pem
5.74KB - 00:00
controller-1# copy scp://root@10.8.67.3:/root/openssl-ca/admin/controller.key private-key://
/controller- private.key
root@10.8.67.3's password:
controller.key
1.67KB - 00:00
controller-1#

```

4. Verify that the certificate was copied correctly:

```

controller-1# show secure
<SNIP>
----- Cert -----
# Name
-|-----|
1 DMF Secure Certificate 2 QA CA
3 ovsclient
----- Csr -----
# Name
-|-----|
1 12358.controller.cluster
2 32591.controller.cluster
----- Private Keys-----
Name Algorithm Value
-----|-----|-----|
controller-private.key sha256 DB:6D:C1:01:E2:CD:71:C4:AA:54:FA:6F:3F:80:4E:C7:25:4C:A9:2A:CA:7F:F5:44:CF:37:3C:C7:67:93:19:BB
ovsclient sha256 EB:88:0C:9D:EE:37:AA:BA:1A:6E:7B:F9:6E:7F:89:45:69:C4:7F:58:D3:18:D2:DC:49:16:2E:1D:2A:2B:94:89
controller-1#

```

5. Apply the certificate and private key.

```

controller-1(config-controller)# certificate DMF\Secure\Certificate
controller-1(config-controller)# private-key controller-private.key

```

6. Display the Controller security configuration.

```

controller-1(config-controller)# show this

```

```

! controller
controller
certificate 'DMF Secure Certificate'
cluster-name DMF_Cluster
private-key controller-private.key
access-control
!
access-list api
1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list gui
1 permit from ::/0
2 permit from 0.0.0.0/0
!
access-list ssh
1 permit from ::/0
2 permit from 0.0.0.0/0
controller-1(config-controller)#

```

7. Access the DMF GUI using a browser and display the certificate.
8. After connecting to the Controller, click the padlock icon to the left of the location field to display information about the certificate.

12.3.1 Replacing the Certificate

Please use the following steps to replace the Controller's certificate.

Scenario 1: Using the same CSR used to sign the current certificate.

Obtain a newly signed certificate from CA using the same CSR and copy it to the Controller using the following command:

```
# copy new certificate from the source cert://
```

For example:

```
# copy scp://root@10.240.88.130/root/openssl-ca/certificate.pem cert:// root@10.240.8
8.130 password certificate.pem
6.49KB - 00:00
#
```

No other action is needed as the current certificate will be overwritten when copying the new one.

Scenario 2: Does not have the same CSR for the current certificate.

1. Generate a new CSR and the private key.
2. Sign the CSR to get the new certificate.
3. Import/copy the certificate to the Controller. The current certificate will be overwritten if the Common Name matches the new one.
4. Import/copy the new private key to the Controller. The private key will be overwritten if the file name is the same as the old one. In that case, there is no need for any config changes.

Assuming the Common Name and the private key dest file names are different from the original ones, remove the old certificate and private key and install a new certificate and private key.

To remove the old certificate and private key, use the following commands:

```
C1(config)# controller
```

```
C1(config-controller)# no certificate certificate name
C1(config-controller)# no private-key private-key name
C1(config-controller)#
```

To configure the new certificate and private key, use the following commands:

```
C1(config)# controller
C1(config-controller)# certificate new certificate name
C1(config-controller)# private-key new private-key name
C1(config-controller)#
```

12.4 Managing the Controller HTTP and SSH Ciphers, Protocols, and Data Integrity Algorithms

Use the `crypto` command to enter the config-crypto submode to configure settings for HTTP and SSH. Use the `http` and `ssh` commands in the config-crypto submode to configure the ciphers and protocols. Configure the list of enabled ciphers, protocols, or algorithms by appending to the list.

Use the `no` version of this command with any keyword to remove the specific cipher, protocol, or algorithm. Use the `no` version of the command without a keyword to restore the list to the default value. Use the CLI help feature to identify the supported ciphers, protocols, or data integrity (MAC) algorithms.

12.4.1 Configuring HTTP Ciphers

Enter the following commands to configure the ciphers for HTTP:

```
controller-1(config)# crypto
controller-1(config-crypto)# http
controller-1(config-crypto-http)# cipher <index> <cipher-name>
```



Note: When you configure a set of ciphers instead of using the default set, please make sure to configure at least one or all the three ciphers mentioned below in addition to your choice. ECDHE-RSA-CHACHA20-POLY1305, ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384.

Example:

```
controller-1(config)# crypto
controller-1(config-crypto)# http
controller-1(config-crypto-http)# cipher 1 <your choice of cipher-name>
controller-1(config-crypto-http)# cipher 2 <your choice of cipher-name>
controller-1(config-crypto-http)# cipher 3 <your choice of cipher-name>
controller-1(config-crypto-http)# cipher 21 ECDHE-RSA-CHACHA20-POLY1305
controller-1(config-crypto-http)# cipher 22 ECDHE-RSA-AES128-GCM-SHA256
controller-1(config-crypto-http)# cipher 23 ECDHE-RSA-AES256-GCM-SHA384
```

12.4.2 Configuring HTTP Protocols

Starting in the DANZ Monitoring Fabric 8.4 release, the TLSv1.3 HTTPS protocol is supported. DMF supports TLSv1.3 and TLSv1.2 by default, with the TLSv1.3 protocol preferred for TLS connections.

Enter the following commands to configure the protocols for HTTP:

```
controller-1(config)# crypto
controller-1(config-crypto)# http
controller-1(config-crypto-http)# protocol <index> <protocol-name>
```

12.4.3 Configuring SSH Ciphers

Configured SSH ciphers and MAC algorithms on the Controller are pushed to the switches running Switch Light OS via ZTN. With this enhancement, users can also restrict the SSH ciphers and MAC algorithms on the switches.

Enter the following commands to configure the ciphers for SSH:

```
controller-1(config)# crypto
controller-1(config-crypto)# ssh
controller-1(config-crypto-ssh)# cipher <index> <cipher-name>
```

12.4.4 Configuring SSH Data Integrity Algorithms

Enter the following command to configure data integrity (MAC) algorithms for SSH:

```
controller-1(config)# crypto
controller-1(config-crypto)# ssh
controller-1(config-crypto-ssh)# mac <index> <mac-name>
```

12.4.5 Changes to Supported MACs/Ciphers/SSH Keys

Please note the below changes to supported MACs/Ciphers and SSH Keys in DANZ Monitoring Fabric starting with the **DMF 8.0** release due to the upgrade of Ubuntu OS from **16.04** to **20.04**. Arista Networks refined the list of supported ciphers to represent better what can be configured and those default enabled when TLSv1.3 protocol is enabled.

- **The default list of SSH MACs has changed:** hmac-ripemd160-etm@openssh.com and hmac-ripemd160 have been removed from the default list of SSH MACs.
- The new default list of SSH MACs is:

```
hmac-sha2-512-etm@openssh.com
hmac-sha2-256-etm@openssh.com
umac-128-etm@openssh.com
hmac-sha1-etm@openssh.com
hmac-sha2-512
hmac-sha2-256
umac-128@openssh.com
hmac-sha1
```

- The following SSH MACs are obsolete and no longer supported:

```
hmac-ripemd160
hmac-ripemd160@openssh.com
hmac-ripemd160-etm@openssh.com
```

- The following SSH ciphers are obsolete and no longer supported:

```
arcfour
arcfour128
arcfour256
```

```
blowfish-cbc
cast128-cbc
```

- **Changes in SSH host keys:**

```
ssh_host_dsa_key is obsolete and no longer supported.
ssh_host_ed25519_key is now available, along with ssh_host_ecdsa_key
and ssh_host_rsa_key that have been available since past releases.
```

- **Removed SSL ciphers:**

```
CAMELLIA128-SHA;
CAMELLIA256-SHA;
DES-CBC3-SHA;
DH-DSS-AES128-GCM-SHA256;
DH-DSS-AES128-SHA;
DH-DSS-AES128-SHA256;
DH-DSS-AES256-GCM-SHA384;
DH-DSS-AES256-SHA;
DH-DSS-AES256-SHA256;
DH-DSS-CAMELLIA128-SHA;
DH-DSS-CAMELLIA256-SHA;
DH-DSS-DES-CBC3-SHA;
DH-DSS-SEED-SHA;
DH-RSA-AES128-GCM-SHA256;
DH-RSA-AES128-SHA;
DH-RSA-AES128-SHA256;
DH-RSA-AES256-GCM-SHA384;
DH-RSA-AES256-SHA;
DH-RSA-AES256-SHA256;
DH-RSA-CAMELLIA128-SHA;
DH-RSA-CAMELLIA256-SHA;
DH-RSA-DES-CBC3-SHA;
DH-RSA-SEED-SHA;
DHE-DSS-AES128-GCM-SHA256;
DHE-DSS-AES128-SHA;
DHE-DSS-AES128-SHA256;
DHE-DSS-AES256-GCM-SHA384;
DHE-DSS-AES256-SHA;
DHE-DSS-AES256-SHA256;
DHE-DSS-CAMELLIA128-SHA;
DHE-DSS-CAMELLIA256-SHA;
DHE-DSS-SEED-SHA;
DHE-RSA-CAMELLIA128-SHA;
DHE-RSA-CAMELLIA256-SHA;
DHE-RSA-SEED-SHA;
ECDH-ECDSA-AES128-GCM-SHA256;
ECDH-ECDSA-AES128-SHA;
ECDH-ECDSA-AES128-SHA256;
ECDH-ECDSA-AES256-GCM-SHA384;
ECDH-ECDSA-AES256-SHA;
ECDH-ECDSA-AES256-SHA384;
ECDH-ECDSA-DES-CBC3-SHA;
ECDH-ECDSA-RC4-SHA;
ECDH-RSA-AES128-GCM-SHA256;
ECDH-RSA-AES128-SHA;
ECDH-RSA-AES128-SHA256;
ECDH-RSA-AES256-GCM-SHA384;
ECDH-RSA-AES256-SHA;
ECDH-RSA-AES256-SHA384;
ECDH-RSA-DES-CBC3-SHA;
ECDH-RSA-RC4-SHA;
ECDHE-ECDSA-DES-CBC3-SHA;
ECDHE-ECDSA-RC4-SHA;
ECDHE-RSA-DES-CBC3-SHA;
ECDHE-RSA-RC4-SHA;
EDH-DSS-DES-CBC3-SHA;
EDH-RSA-DES-CBC3-SHA;
PSK-3DES-EDE-CBC-SHA;
PSK-RC4-SHA;
RC4-MD5;
RC4-SHA;
SEED-SHA;
SRP-3DES-EDE-CBC-SHA;
SRP-DSS-3DES-EDE-CBC-SHA;
SRP-DSS-AES-128-CBC-SHA;
SRP-DSS-AES-256-CBC-SHA;
SRP-RSA-3DES-EDE-CBC-SHA;
```

```
DHE-PSK-AES128-CBC-SHA;
DHE-PSK-AES128-CBC-SHA256;
DHE-PSK-AES128-GCM-SHA256;
DHE-PSK-AES256-CBC-SHA;
DHE-PSK-AES256-CBC-SHA384;
DHE-PSK-AES256-GCM-SHA384;
DHE-PSK-CHACHA20-POLY1305;
DHE-RSA-CHACHA20-POLY1305;
ECDHE-PSK-AES128-CBC-SHA;
ECDHE-PSK-AES128-CBC-SHA256;
ECDHE-PSK-AES256-CBC-SHA;
ECDHE-PSK-AES256-CBC-SHA384;
ECDHE-PSK-CHACHA20-POLY1305;
ECDHE-RSA-CHACHA20-POLY1305;
PSK-AES128-CBC-SHA256;
PSK-AES128-GCM-SHA256;
PSK-AES256-CBC-SHA384;
PSK-AES256-GCM-SHA384;
PSK-CHACHA20-POLY1305;
RSA-PSK-AES128-CBC-SHA;
RSA-PSK-AES128-CBC-SHA256;
RSA-PSK-AES128-GCM-SHA256;
RSA-PSK-AES256-CBC-SHA;
RSA-PSK-AES256-CBC-SHA384;
RSA-PSK-AES256-GCM-SHA384;
RSA-PSK-CHACHA20-POLY1305
```

- Added SSL ciphers:

```
ECDHE-ECDSA-CHACHA20-POLY1305;
```

- Conditionally enabled ciphers:



Note: Enabled by default for TLSv1.3, the SSL ciphers below cannot be configured using `crypto/http/ciphers` configuration.

```
TLS_AES_256_GCM_SHA384;
TLS_CHACHA20_POLY1305_SHA256;
TLS_AES_128_GCM_SHA256
```

12.5 Inherit MAC and Cipher Configuration

This feature provides the ability to mirror the SSH/HTTPS cryptographic configuration of the DMF Controller to the managed appliances (i.e., service nodes and recorder nodes) and the SSH cryptographic configuration of the Controller to the EOS switches.

12.5.1 Using the CLI to Configure SSH and HTTPS

The configuration that a managed appliance or EOS switch receives is intended for the Controller itself. Configuring a cipher or message authentication code (MAC) on the Controller will automatically be reflected onto a managed appliance or EOS switch.

SSH and HTTPS Cryptographic Configuration Syntax

```
(config)# crypto
(config-crypto)# ssh
(config-crypto-ssh)# cipher number algorithm
(config-crypto-ssh)# mac number algorithm
(config-crypto-ssh)# http
```

```
(config-crypto-http)# cipher number algorithm
(config-crypto-http)# protocol number algorithm
```

The following is a configuration example using common algorithms.

```
(config)# crypto
(config-crypto)# ssh
(config-crypto-ssh)# cipher 1 3des-cbc
(config-crypto-ssh)# mac 1 hmac-md5
(config-crypto-ssh)# http
(config-crypto-http)# cipher 1 AES128-GCM-SHA256
(config-crypto-http)# cipher 2 ECDHE-RSA-CHACHA20-POLY1305
(config-crypto-http)# protocol 2 SSLv2
```

12.5.2 Verify the Cryptographic Configuration

Check the cryptographic configuration of the Controller using the `show running-config` command, as shown in the following example, and verify the settings in the crypto section.

```
# show running-config
.
.
.
! crypto
crypto
!
ssh
cipher 1 3des-cbc
mac 1 hmac-md5
.
.
.
```

All ciphers/protocols/MACs of the HTTPS/SSH cryptographic configuration supported on the Controller are supported on the managed appliances, with one caveat listed in the Limitations section below. Check the HTTPS/SSH cryptographic configuration by reviewing the running-config of a managed appliance, as shown below for a Recorder Node.

```
# show recorder-node device rn1 running-config
.
.
.
! crypto
crypto
!
ssh
cipher 1 3des-cbc
mac 1 hmac-md5
.
.
.
```



Note: EOS does not support all SSH ciphers and MACs that the Controller does.

The following SSH MAC algorithms supported by the Controller are **not** supported by EOS:

- hmac-md5-96-etm@openssh.com (HMAC-MD5 in “encrypt-then-mac” mode)

- hmac-md5-etm@openssh.com (HMAC-MD5 in “encrypt-then-mac” mode)
- hmac-sha1-96-etm@openssh.com (HMAC-SHA1 in “encrypt-then-mac” mode)
- umac-64-etm@openssh.com (message authentication code based on universal hashing (UMAC) in “encrypt-then-mac” mode)
- umac-64@openssh.com (UMAC)

The following SSH cipher algorithm supported by the Controller is not supported by EOS:

- rijndael-cbc@lysator.liu.se (Rijndael in CBC mode)

This difference can be seen when reviewing the running-config of the Controller and the ZTN-generated running-config of an EOS switch:

```
# show running-config
.
.
.
! crypto
crypto
!
ssh
cipher 1 rijndael-cbc@lysator.liu.se
cipher 2 3des-cbc
mac 1 hmac-md5-etm@openssh.com
mac 2 hmac-sha2-512-etm@openssh.com
.
.
.
```

```
# show switch switch-name running-config
.
.
.
cipher 3des-cbc
mac hmac-sha2-512-etm@openssh.com
.
.
.
```

Only the supported ciphers/MACs get added to the `running-config` of the EOS switch. To review the disallowed MACs/ciphers when generating the running-config of the switch, use the following `show` command:

```
# show fabric warnings feature-unsupported-on-device
# Name Warning
-|-----|-----|-----|-----|-----|-----|
1 core1 rijndael-cbc@lysator.liu.se is not a supported cipher on EOS switches
2 core1 hmac-md5-etm@openssh.com is not a supported mac on EOS switches
```

Syslog Messages

No syslog messages are generated when the DMF Controller’s cryptographic configuration is mirrored to the managed appliances and EOS switches.

12.5.3 Limitations

1. There are limitations to the HTTPS configuration (some options may cause ZTN protocol/communication and controller-to-controller communication failures). The following HTTPS protocol versions are required to be used to avoid communication failures:
 - a. TLSv1.2
 - b. TLSv1.3
2. It is not apparent when a cipher/MAC is not reflected onto an EOS switch (due to it being unsupported). To ascertain this scenario, check the Controller's running-config and the switch's ZTN-generated running-config and compare them (alternatively, check the "show fabric warnings" command output to review any generated warnings).
3. An ECDSA-based cryptographic cipher configuration inherited by the managed appliances will cause a failure in communication with the Controller.

12.6 Protocol Access Required to the DMF Controller

12.6.1 Management Plane Access

You can control access to the DMF Controller for specific protocols, and in the case of SSH, you can allow access only from specific IP addresses or subnetworks. The following table summarizes the TCP/UDP protocol ports that DMF uses. The CLI `access-list` option column shows the options for the ports that can be enabled or disabled using the CLI `access-list` command (**`config-controller-access`** submode). The ports listed are open by default on the Controller, except for SNMP, which is disabled by default.

These ports must also be open on any device, such as a router or firewall, that connects the management console or application to the DMF Controller.

Protocol	Port	Application	CLI access-list option	Match criteria
HTTP	TCP 80	GUI auto-redirect		
HTTPS	TCP 443	GUI remote access	gui, applicable to Controller, Service Node, Recorder Node, Analytics Node	Default any, configurable
HTTPS	TCP 8443	REST API	api, applicable to Controller, Service Node, Recorder Node, Analytics Node	Default any, configurable
ICMP/ ICMPv6	ICMP/ ICMPv6			selected ICMP types
ICMP/ ICMPv6	ICMP/ ICMPv6			selected ICMP types
SNMP	UDP 161, 162	SNMP, applicable to Controller, Service Node, Analytics Node	snmp	Default none, configurable
SSH	TCP 22	CLI remote access	ssh, applicable to Controller, Service Node, Recorder Node, Analytics Node	Default any, configurable
syslog	UDP 514			
vce-api	UDP 7443	vCenter integration	vce-api	Enabled by default

 **Note:** Be careful when configuring firewall rules for the SSH protocol, which is permitted from all subnetworks by default. The option exists to restrict SSH access to one or more specific subnetworks. However, this denies access from all other subnetworks. If no connectivity from a specified subnetwork is available, accessing the Controller is only through the local console.

12.6.2 Control Plane Access for DMF Controller

The following ports must be open between the DMF Controller and any connected devices. No further configuration is required if all devices are in the same Layer 2 network as the DMF Controller. However, if any DMF nodes or fabric switches connect over a Layer 3 network, these ports must be open on any firewalls or routers that connect the devices to the DMF Controller.

Table 4: DMF Controller

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	22	Both Directions	SSH	Customer	Switches, managed appliances
TCP	49	Out	TACACS+		Customer TACACS+ server
TCP	53	Out	DNS		Customer DNS server
UDP	53	Out	DNS		Customer DNS server
UDP	67	Out	DHCP		Customer DHCP server
UDP	68	In	DHCP	Customer DHCP server	
TCP	80	In	ZTN ONIE	Switches	
UDP	123	Both Directions	NTP	Switches, Service Node, Recorder Node, Analytics Node	Customer NTP server
UDP	161	In	SNMP	Customer	
UDP	162	Out	SNMP Traps		Customer
TCP	443	In	GUI	Customer	
UDP	514	Out	Syslog		Customer Syslog server
UDP	1813	Out	RADIUS		Default RADIUS accounting port
UDP	5353	In	ZTN MDNS	Switches, Service Node, Recorder Node	
TCP	6379	Out	Controller Stats		Analytics Nodes
TCP	6642	Both Directions	Cluster Sync	Controller HA	Controller HA
TCP	6653	In	OpenFlow	Switches, Recorder Node, Service Node	
TCP	7443	In	VCE API	vCenter API	
TCP	8443	Both Directions	Floodlight REST API	Customer, Recorder Node	Recorder Node, Service Node
TCP	8443	Out	Controller-to-switch traffic		Fabric switches using the Switch Light OS
TCP	8843	In	ZTN	Switches, Service Node, Recorder Node	

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	9379	Out	Analytics Node Replicated Redis		Analytics Node



Note: Starting from DMF release 8.5.0, Controllers use port 8443 to apply configurations to fabric switches using the Switch Light OS. In earlier releases it was port 443. If using Switch Light OS, network administrators must update the firewall configuration to permit traffic on port 8443 between Controllers and fabric switches.

To enable SNMP access to the Controller or to restrict access to the Controller for the REST API, web-based GUI access, or SSH applications, complete the following steps.

Procedure

1. Enter the **controller** command from config mode to enter config-controller submode.

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

2. Enter the **access-control** command from config-controller submode.

```
controller-1(config-controller)# access-control
controller-1(config-controller-access)#
```

3. Enter the **access-list** command from config-controller-access submode followed by the protocol for which you want to configure a rule.

The protocols for which you can configure rules include the following:

- **api:** Enter the config-controller-access-list submode for REST/API access to the Controller.
- **gui:** Enter the config-controller-access-list submode for web-based GUI access to the Controller.
- **ns-api:** Enter the config-controller-access-list submode to manage NS-API access to the Controller, for example, from OpenStack or vCenter.
- **ssh:** Enter the config-controller-access-list submode for SSH access to the Controller.
- **snmp:** Enter the config-controller-access-list submode for SNMP access to the Controller.

By default, the access list for all services except for SNMP is **0.0.0.0/0**, which allows access from any IPv4 subnetwork and **::/0**, which allows access from any IPv6 subnetwork. For SNMP, the access list is empty, meaning access is not permitted unless specifically enabled. SNMP UDP **port 161** is blocked on Controllers and fabric switches by default. You must configure an SNMP access list using the Controller's CLI to communicate using SNMP.

For example, the following command enters **config-controller-access-list** submode for the SSH protocol:

```
controller-1(config-controller-access)# access-list ssh
controller-1(config-controller-access-list)#
```

When configuring the following access list on the Controller, the ACL is pushed to all connected switches, and the SNMP ACL is applied to each switch *ma1* management interface.

```
controller-1(config-controller-access)# access-list snmp
controller-1(config-controller-access-list)# 1 permit from ::/0
controller-1(config-controller-access-list)# 2 permit from 10.0.0.0/8
```

- Specify the subnetworks from which access is permitted for the specified protocol. Specify the subnetwork followed by a slash and the number of bits in the subnet mask.

For example, the following commands allow access to the SSH protocol only from the subnetwork **192.168.1.0**:

```
controller-1(config-controller-access)# access-list
ssh controller-1(config-controller-access-list)# 10 permit from 192.168.1.0/24
```

- To view the current firewall configuration for the Controller, enter the **show running-config** command and see the access-control section.

12.6.3 Protocol Access Required to the DMF Controller - Sync

Sync has been added to the access list. All traffic is permitted for IPv4 (**0.0.0.0/0**) and IPv6 (**:::0**) by default. If the active standby Controller High Availability (HA) pair is deployed in different L3 subnets, permitting all traffic for sync services can be a security risk. For a secure connection between the active and standby Controller, the access list for syncing should only permit the active and standby Controller's management IP addresses.

Procedure

- Add a rule to permit active Controller management IP address. Do not overwrite the existing default rule for sync. Use a different rule number when adding a new rule for the *sync* access list. The example below shows how to configure Rule ID 3 (3 permit from 10.240.130.17/32).

```
DMF-CTL2(config)# show controller access-control access-list
# Access-list Rule Action Source
-- |-----|----|-----|-----|
1 api 1 permit ::/0
2 api 2 permit 0.0.0.0/0
3 gui 1 permit ::/0
4 gui 2 permit 0.0.0.0/0
5 ntp 1 permit ::/0
6 ntp 2 permit 0.0.0.0/0
7 snmp 1 permit 0.0.0.0/0
8 ssh 1 permit ::/0
9 ssh 2 permit 0.0.0.0/0
10 sync 1 permit ::/0
11 sync 2 permit 0.0.0.0/0
12 vce-api 1 permit ::/0
13 vce-api 2 permit 0.0.0.0/0
DMF-CTL2(config-controller-access-list)# controller
DMF-CTL2(config-controller)# access-control
DMF-CTL2(config-controller-access)# access-list sync
DMF-CTL2(config-controller-access-list)# 3 permit from 10.240.130.17/32
DMF-CTL2(config-controller-access-list)#
```

- Add another rule to permit standby Controller management IP address. Do not overwrite the existing default rule for *sync*. Use a different rule number when adding a new rule for *sync* access list. The example below shows how to configure Rule ID 4 (4 permit from 10.240.130.16/32).

```
DMF-CTL2(config-controller-access-list)# 4 permit from 10.240.130.16/32
DMF-CTL2(config-controller-access-list)# show controller access-control access-list
# Access-list Rule Action Source
-- |-----|----|-----|-----|
1 api 1 permit ::/0
2 api 2 permit 0.0.0.0/0
3 gui 1 permit ::/0
4 gui 2 permit 0.0.0.0/0
5 ntp 1 permit ::/0
6 ntp 2 permit 0.0.0.0/0
7 snmp 1 permit 0.0.0.0/0
8 ssh 1 permit ::/0
9 ssh 2 permit 0.0.0.0/0
10 sync 1 permit ::/0
11 sync 2 permit 0.0.0.0/0
```


Table 5: DMF Switch

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	22	In	SSH	Customer, DMF Controller	
TCP/UDP	53	Out	DNS		Customer DNS Server
UDP	67	Out	DHCP		Customer DHCP Server
UDP	68	In	DHCP	Customer DHCP Server	
TCP	80	In	cAPI ¹	DMF Controller	
UDP	123	Out	NTP		Customer NTP Server
UDP	161	In	SNMP	Customer	
UDP	162	Out	SNMP Trap		Customer
UDP	514	Out	Syslog		Customer Syslog Server
UDP	5353	Out	ZTN MDNS		DMF Controller
TCP	6030	In	Telemetry ¹		
UDP	6343	Out	sFlow ^{®2}		Analytics Node
UDP	6380	Out	Control Packets		Analytics Nodes
TCP	6653	Out	OpenFlow		DMF Controller
TCP	8443	In	BigDB	DMF Controller	
TCP	8843	Out	ZTN		DMF Controller

¹ Port to be opened on firewalls for Arista switches running EOS to communicate with the DMF Controller.

² sFlow[®] is a registered trademark of Inmon Corp.

12.6.5 Control Plane Access for DMF Service Node

The following ports must be open for the DMF Service Node to communicate with the DMF Controller, Analytics Node, and other services (e.g., NTP, DHCP, etc.). No further configuration is required if all devices are in the same Layer 2 network as the DMF Controller. However, if the DMF Controller and Service Nodes connect over a Layer 3 network, these ports must be open on any firewalls or routers.

Table 6: DMF Service Node

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	22	In	SSH	Customer, DMF Controller	
TCP	49	Out	TACACS+		Customer TACACS+ Server
TCP/UDP	53	Out	DNS		Customer DNS Server
UDP	67	Out	DHCP		Customer DHCP Server
UDP	68	In	DHCP	Customer DHCP	
UDP	123	Out	NTP		Customer NTP Server
UDP	161	In	SNMP	Customer	
UDP	162	Out	SNMP Trap		Customer SNMP Trap Server
UDP	514	Out	Syslog		Customer Syslog Server
UDP	1812	Out	Default RADIUS Authentication port		Customer RADIUS Server
UDP	1813	Out	Default RADIUS Accounting port		Customer RADIUS Server
UDP	5353	Out	ZTN MDNS		DMF Controller
TCP	6030	In	Telemetry		
TCP	6634	In	Telemetry		
TCP	6653	Out	OpenFlow		DMF Controller
TCP	8443	Both Direction	Floodlight REST API	DMF Controller	DMF Controller
TCP	8843	Out	ZTN		DMF Controller

12.6.6 Control Plane Access for DMF Recorder Node

The following ports must be open between the DMF Recorder Node and any connected devices. No further configuration is required if all devices are in the same Layer 2 network as the DMF Recorder Node. However, if the DMF Controller, Analytics Node, or fabric switches connect over a Layer 3 network, these ports must be open on any firewalls or routers that connect the devices to the DMF Recorder Node.

Table 7: DMF Recorder Node

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	22	In	SSH	Customer, DMF Controller	
TCP	49	Out	TACACS+		Customer TACACS+ Server
TCP/UDP	53	Out	DNS		Customer DNS Server
UDP	67	Out	DHCP		Customer DHCP Server
UDP	68	In	DHCP	Customer DHCP Server	
UDP	123	Out	NTP		Customer NTP Server
UDP	161	In	SNMP	Customer	
UDP	162	Out	SNMP Trap		Customer SNMP Trap Server
TCP	443	In	Stenographer Query API	Customer, DMF Controller	
UDP	514	Out	Syslog		Customer Syslog Server
UDP	1812	Out	Default RADIUS Authentication port		Customer RADIUS Server
UDP	1813	Out	Default RADIUS Accounting port		Customer RADIUS Server
TCP	2049	Both	NFS		Customer NSF Server
UDP	2049	Both	NFS		Customer NFS Server
UDP	5353	Out	ZTN MDNS		DMF Controller
TCP	6653	Out	OpenFlow		DMF Controller
TCP	8443	Both Direction	Floodlight REST API	DMF Controller	DMF Controller
TCP	8843	Out	ZTN		DMF Controller

12.6.7 Control Plane Access for Analytics Node

The following ports must be open between the Analytics Node and any connected devices. No further configuration is required if all devices are in the same Layer 2 network as the Analytics Node. However, if the Analytics Node connects over a Layer 3 network, these ports must be open on any firewall or router.

Table 8: DMF Analytics Node

Protocol	Port	Direction	Application	In Flows	Out Flows
TCP	22	In	SSH	Customer	
TCP	25		SMTP		Analytics Nodes to Mail Server
TCP	49	Out	TACACS+		Customer TACACS+ Server
TCP/UDP	53	Out	DNS		Customer DNS Server
UDP	67	Out	DHCP		Customer DHCP Server
UDP	68	In	DHCP	Customer DHCP Server	
UDP	123	Out	NTP		Customer NTP Server
UDP	161	In	SNMP	Customer	
UDP	161	In	SNMP	from Analytics Nodes to DMF switch	
UDP	162	Out	SNMP Trap		Customer
UDP	162	Out	SNMP Trap		from Analytics Nodes to DMF switch
TCP	443	In	GUI	Customer	
TCP	467		SMTP		Analytics to Mail Server
UDP	514	Out	Syslog		Customer Syslog Server
UDP	1812	Out	Default RADIUS Authentication port		Customer RADIUS Server
UDP	1813	Out	Default RADIUS Accounting port		Customer RADIUS Server
UDP	2055	In	NetFlow v5	DMF Service Nodes and Switches	

Protocol	Port	Direction	Application	In Flows	Out Flows
UDP	4739	In	IPFIX & NetFlow v9	DMF Service Nodes and Switches	
TCP	5043	Both Direction	Active Directory	Customer Active Directory Server	Customer Active Directory Server
UDP	6343	In	sFlow ^{®1}	DMF Switches	
TCP	6379	Both Direction	Controller Stats	Controller to Analytics VIP	
UDP	6380	In	Control Packets	DMF Switches	
TCP	6642	Both Direction	Analytics Cluster sync	HA controller	HA controller
TCP	8443	Both Direction	Floodlight REST API	Customer	Managed Appliances
TCP	9379	Both Direction	Replicated Redis	DMF Controller to Analytics Node VIP	
TCP	9379	Out	Analytics Node Replicated Re- dis Server (for dpid.port -> Filter Name)		Analytics Node

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

Using iDRAC with a Dell R430, R630, or R730 Server

This chapter provides step-by-step instructions for using the integrated Dell Remote Access Controller (iDRAC) Enterprise Version to install the DMF Controller image on Dell R430 servers and the DMF Service Node Appliance image on Dell R630 or R730 servers.

This chapter was validated using *iDRAC Enterprise Version 2.10.10* and *DMF Release 6.0.1*. The instructions were tested with *DMF Controller Release 6.0.1* on Dell R430 and *DMF Service Node Appliance Release 6.0.1* on Dell R630 or R730.

The instructions in this chapter tested using macOS are similar to those tested using Windows OS. Use Internet Explorer, Chrome, or Safari browsers to access the iDRAC web interface.



Note: The iDRAC features require a separate iDRAC Enterprise license on the Dell R430/R630/R730 hardware appliances.

To purchase an iDRAC Enterprise license, contact Dell Sales at <https://www.dell.com/learn/us/en/19/campaigns/contact-us-phone-number>.

Alternatively, purchase an iDRAC Enterprise license from Dell Digital Locker (DDL). <https://www.dell.com/support/kbdoc/en-us/000130349/how-to-obtain-idrac-enterprise-licenses-from-dell-digital-locker-ddl>.

13.1 Setting Up and Configuring iDRAC

Complete the following steps to set up iDRAC on a Dell R430, R630, or R730 server.

Procedure

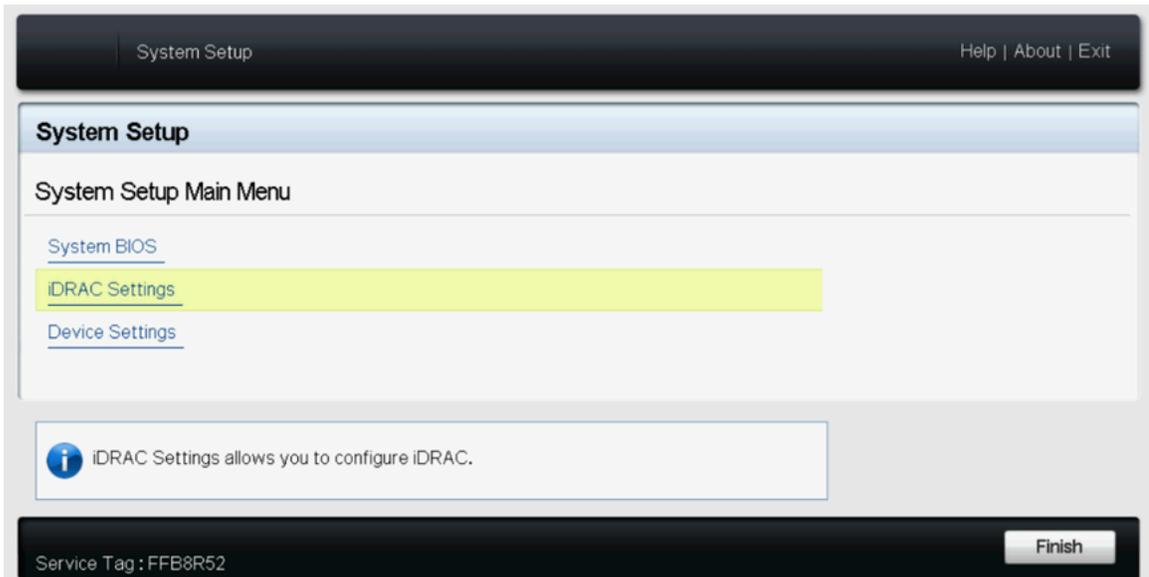
1. Turn on the Dell PowerEdge Server.
2. Connect a monitor to the VGA port and the keyboard to the USB port.



Note: When using the serial port to connect to the R430, R630, or R730, set the baud rate to **115200**.

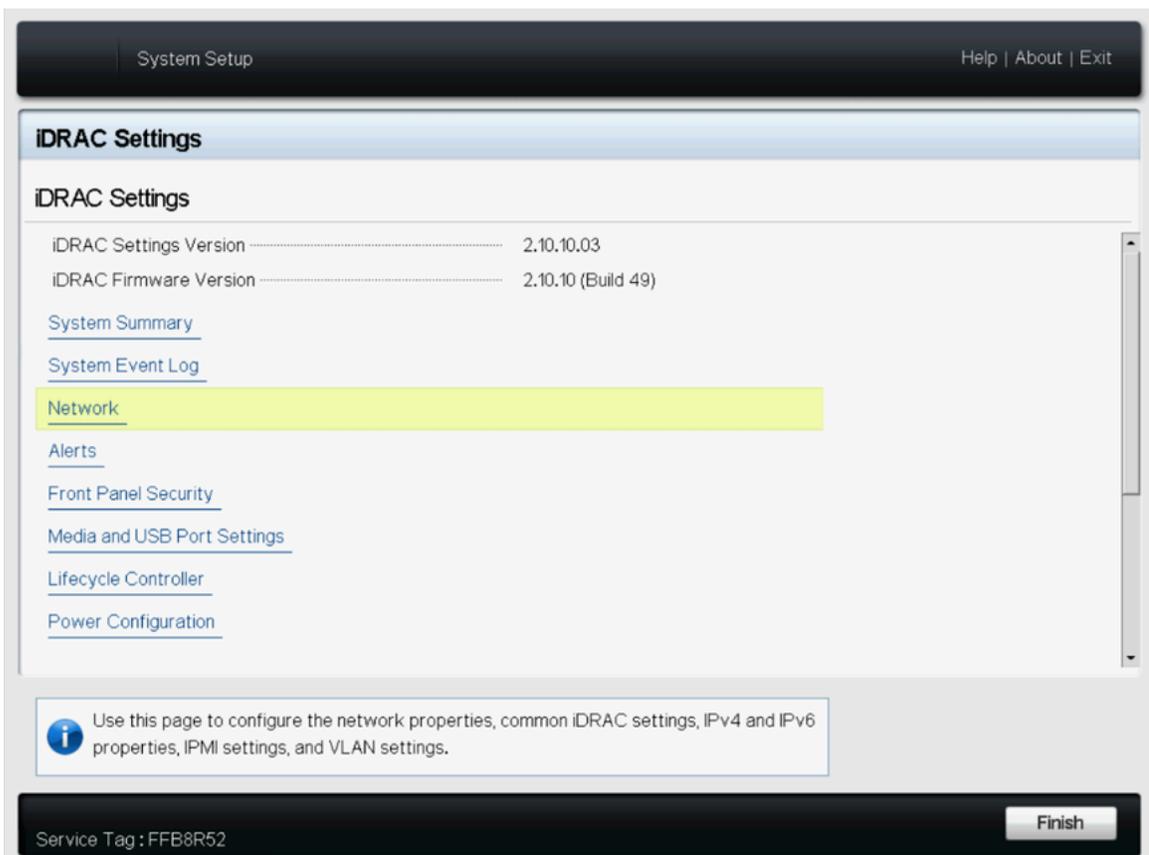
3. Press **<F2>** to enter the **System Setup Main Menu** screen.

Figure 13-1: System Setup Main Menu



4. Select **iDRAC Settings** from the **System Setup Main Menu** (**F2 > iDRAC Settings**).

Figure 13-2: iDRAC Settings



- Use the arrow keys to select **Network**.

Figure 13-3: iDRAC Settings > Network

System Setup Help | About | Exit

iDRAC Settings

iDRAC Settings • Network

NETWORK SETTINGS

Enable NIC Disabled Enabled

NIC Selection LOM1

Failover Network None

MAC Address 44:A8:42:35:29:FB

Auto Negotiation Off On

Auto Dedicated NIC Disabled Enabled

Network Speed 10 Mbps 100 Mbps 1000 Mbps

Active NIC Interface LOM1

Duplex Mode Half Duplex Full Duplex

COMMON SETTINGS

Register DRAC on DNS Disabled Enabled

Select Enabled to enable NIC. When NIC is enabled, it activates the remaining controls in this group. When a NIC is disabled, all communication to and ... (Press <F1> for more help)

Service Tag: FFB8R52 Back

- Select an option from the **NIC Selection** list for use with iDRAC. In this example, LOM1. The options provided by both Dell R430 and R630 or R730 servers include **LOM1**, **LOM2**, **LOM3**, and **LOM4**. In addition to the LOM options, the Dell R630 or R730 provides a dedicated NIC for use with iDRAC. To use the dedicated NIC with a Dell R630 or R730, select Dedicated from the NIC Selection list.

7. Configure the iDRAC IPv4 address, netmask, gateway, and DNS addresses, as shown on the following page.

Figure 13-4: iDRAC Settings > Network (Completed Entries)

System Setup Help | About | Exit

iDRAC Settings

iDRAC Settings • Network

IPv4 SETTINGS

Enable IPv4 Disabled Enabled

Enable DHCP Disabled Enabled

Static IP Address

Static Gateway

Static Subnet Mask

Use DHCP to obtain DNS server addresses Disabled Enabled

Static Preferred DNS Server

Static Alternate DNS Server

IPv6 SETTINGS

Enable IPv6 Disabled Enabled

Enable Auto-configuration Disabled Enabled

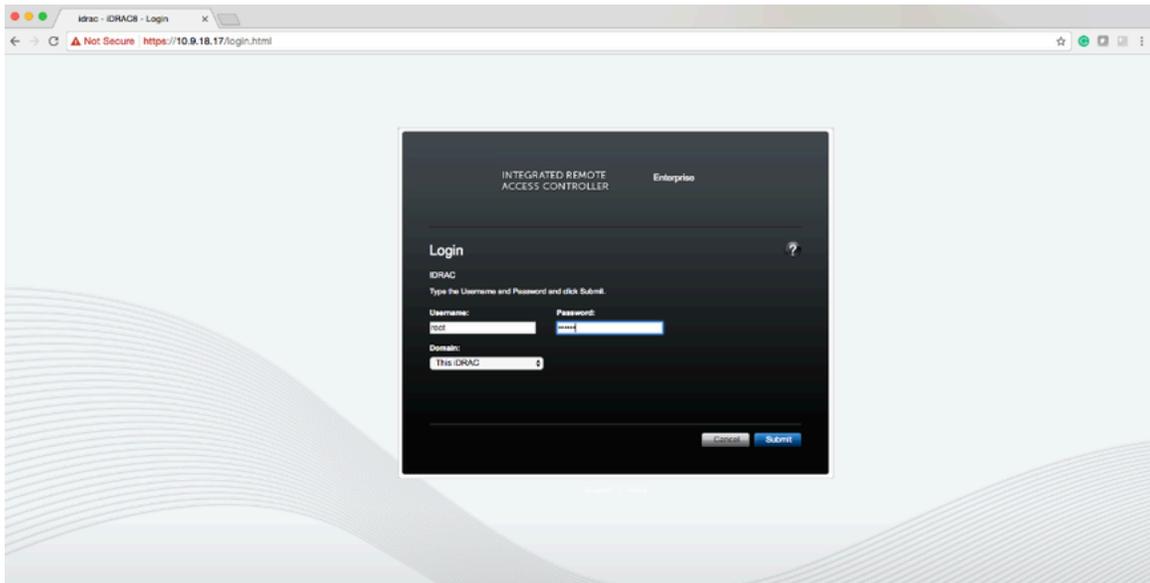
i Select Enabled to enable NIC. When NIC is enabled, it activates the remaining controls in this group. When a NIC is disabled, all communication to and ... (Press <F1> for more help)

Service Tag : FFB8R52 Back

8. After completing the iDRAC configuration, press the **Esc** key to display the Exit menu.
9. Select **Save Changes and Exit** and then press **Enter** to keep the changes.
10. From a web browser, type **DRAC-IPv4-address** in the browser address bar and press **Enter**.

The system displays the iDRAC web interface.

Figure 13-5: iDRAC Web Interface Login Window



11. Enter the username and password.

Default Password

By default, all PowerEdge servers ship with a unique iDRAC password, randomly generated at the factory, that is located on the pull-out Information Tag located on the front of the chassis, near the server asset tag. Customers who choose this option need to note this unique random password and use it to log in to iDRAC for the first time. For security purposes, Dell strongly recommends changing the default password.

Legacy Password

Customers who prefer to use the known legacy password *calvin* should choose this option. One reason to select this option would be to ensure conformance with any existing scripts. For security purposes, Dell strongly recommends changing the legacy password.

Refer to <https://www.dell.com/support/kbdoc/en-us/000133536/dell-poweredge-what-is-the-default-username-and-password-for-idrac> for the official Dell information.

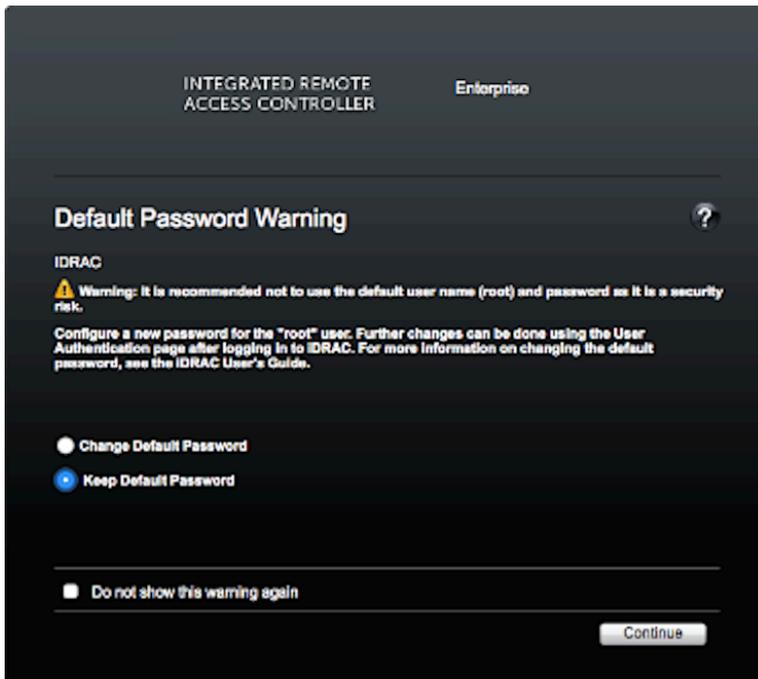
The pull-out information card is shown in the figure below:



Check out the bottom label on the pull-out information card to find the default password, which in this example is **94AXYXFERKNW**:

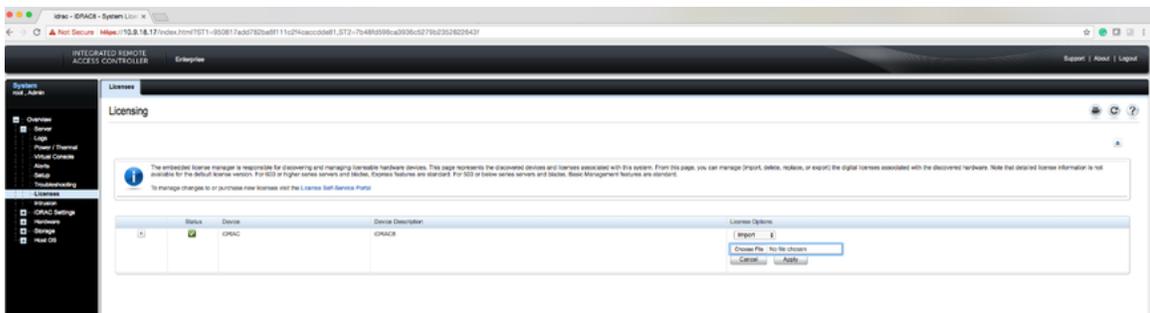


Figure 13-6: Change the Default Password



12. To change the default password, type the new password, when prompted, and type it again to confirm. In this example, the password stayed the same.

Figure 13-7: Licensing Page



13. To apply the iDRAC Enterprise Version license select the license option and then select import.
14. Select the license file and click **Apply**.
15. After applying the Enterprise Version license, log out for the license to take effect.

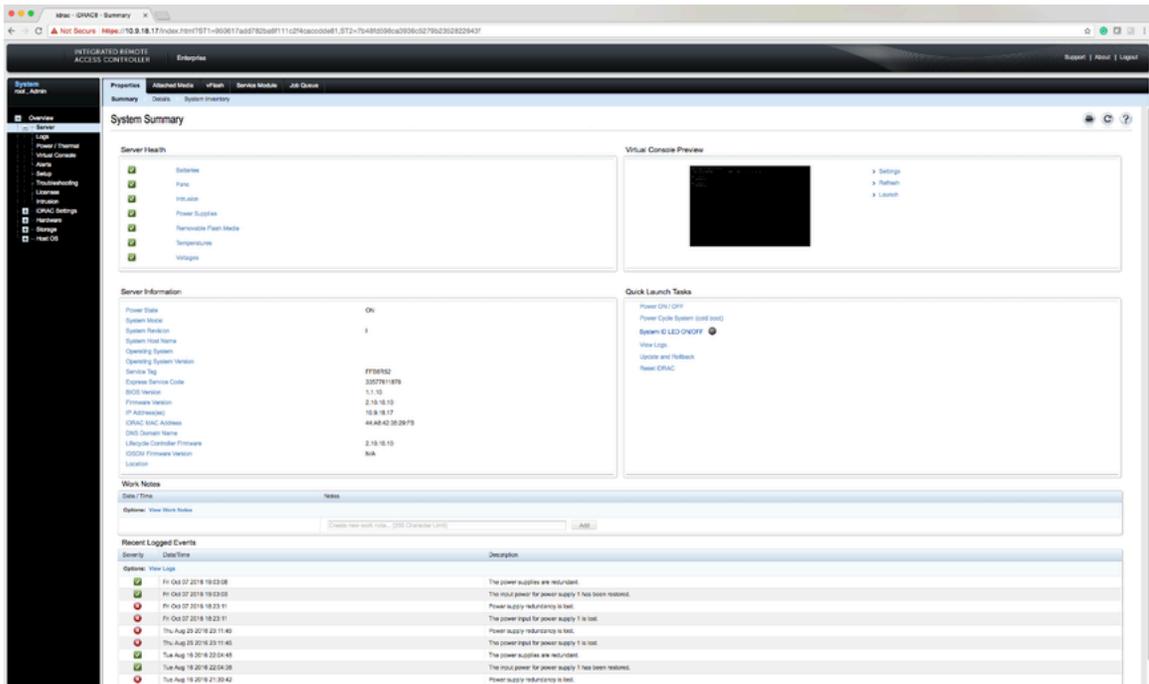


Note: You need to log out from the iDRAC web interface, but there is no need to close the web browser.

16. Log in to the iDRAC web interface using the new credentials (if changed). After applying the license, logging out and logging in enables the Enterprise Version license.

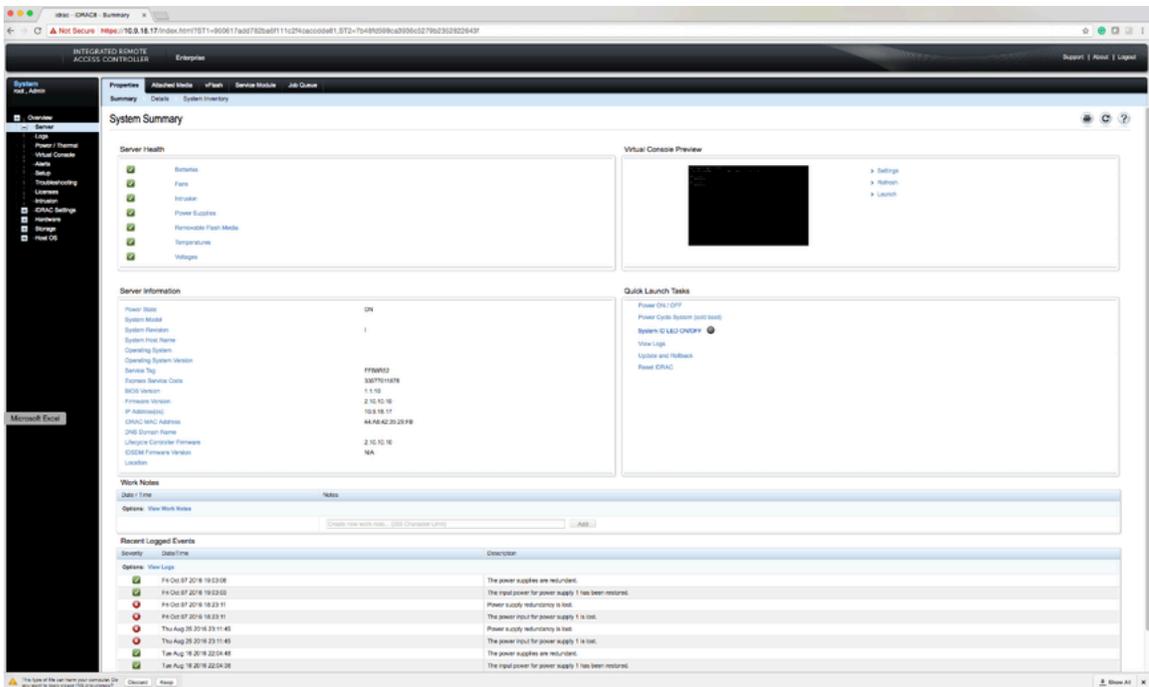
17. Using the Virtual Console Preview option, click **Launch**.

Figure 13-8: System Summary Page



18. When prompted, click **Keep** to confirm the operation.

Figure 13-9: System Summary Page



19. Click the `viewer.jnlp` link, as shown in the following example.

Figure 13-10: System Summary Page

The screenshot shows the iDRAC System Summary page. The page is divided into several sections:

- Server Health:** A list of health indicators with green checkmarks, including Battery, Fans, Inlet Air, Power Supplies, Removable Flash Media, Temperatures, and Voltages.
- Virtual Console Preview:** A black rectangular area representing the virtual console, with links for Settings, Refresh, and Launch.
- Server Information:** A table of system details:

Power State	ON
System Model	
System Revision	1
System Host Name	
Operating System	
Operating System Version	
Service Tag	FF8W102
Express Service Code	130791879
BIOS Version	1.14.0
Firmware Version	2.10.10.10
IP Address(es)	10.8.18.17
iDRAC MAC Address	4A:6A:4D:30:29:F9
OME Device Name	
Lifecycle Controller Firmware	2.10.10.10
EGS/OM Firmware Version	N/A
Location	
- Quick Launch Tasks:** A list of tasks including Power On (ON), Power Cycle System (power boot), System ID LED ON/OFF, View Logs, Update and Refresh, and Power iDRAC.
- Work Notes:** A section for notes with a "New" button and a "View Work Notes" link.
- Recent Logged Events:** A table of system events:

Severity	Date/Time	Description
Info	Fri Oct 07 2016 18:03:08	The power supplies are redundant.
Info	Fri Oct 07 2016 18:03:08	The input power for power supply 1 has been restored.
Info	Fri Oct 07 2016 18:03:11	Power supply redundancy is lost.
Info	Fri Oct 07 2016 18:03:11	The power input for power supply 1 is lost.
Info	Thu Aug 20 2016 23:11:45	Power supply redundancy is lost.
Info	Thu Aug 20 2016 23:11:45	The power input for power supply 1 is lost.
Info	Tue Aug 16 2016 22:04:48	The power supplies are redundant.
Info	Tue Aug 16 2016 22:04:38	The input power for power supply 1 has been restored.

20. When prompted, open the `viewer.jnlp` file.

Figure 13-11: System Summary Page

This screenshot is identical to Figure 13-10, but with a Java security warning dialog box overlaid in the center. The dialog box contains the following text:

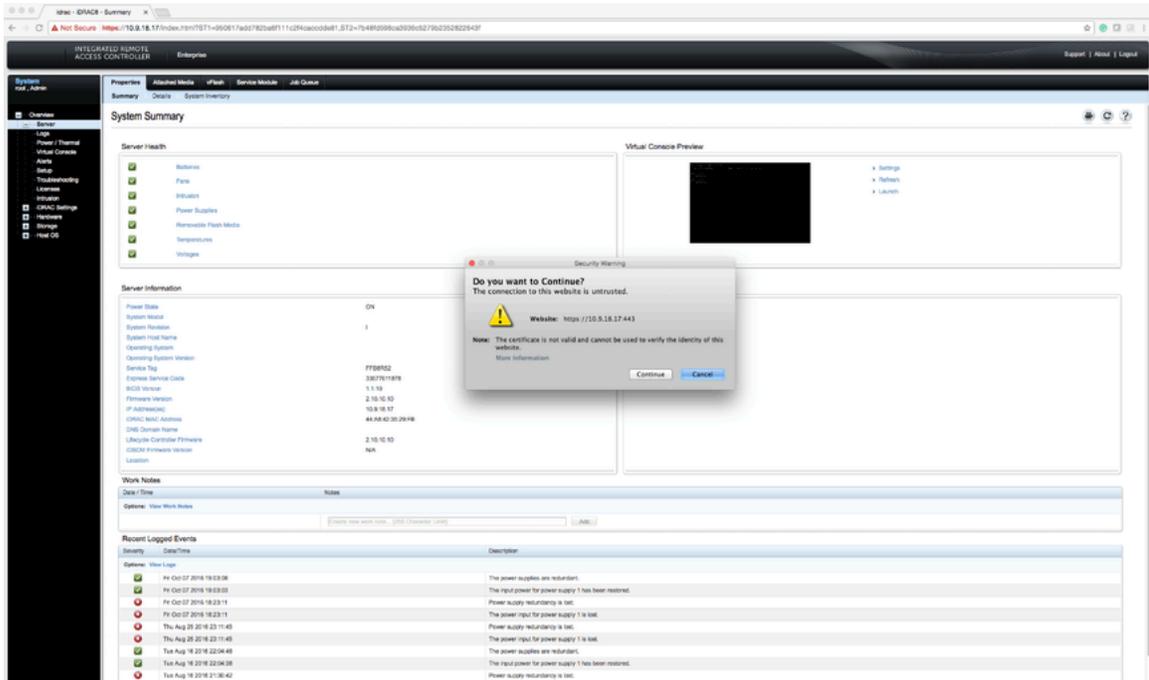
"viewer (135.jnlp)" is an application downloaded from the Internet. Are you sure you want to open it?

Options: Cancel, Show Web Page, Open



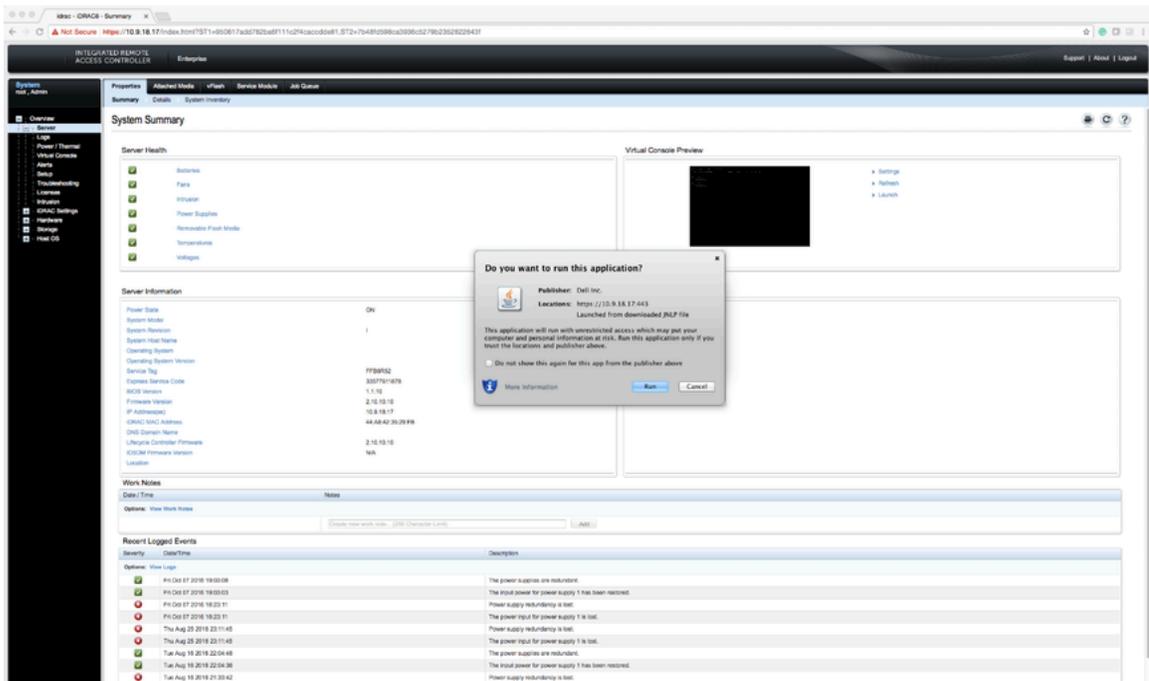
Note: JAVA is required, and the correct version downloads automatically.

Figure 13-12: System Summary: Continue Prompt



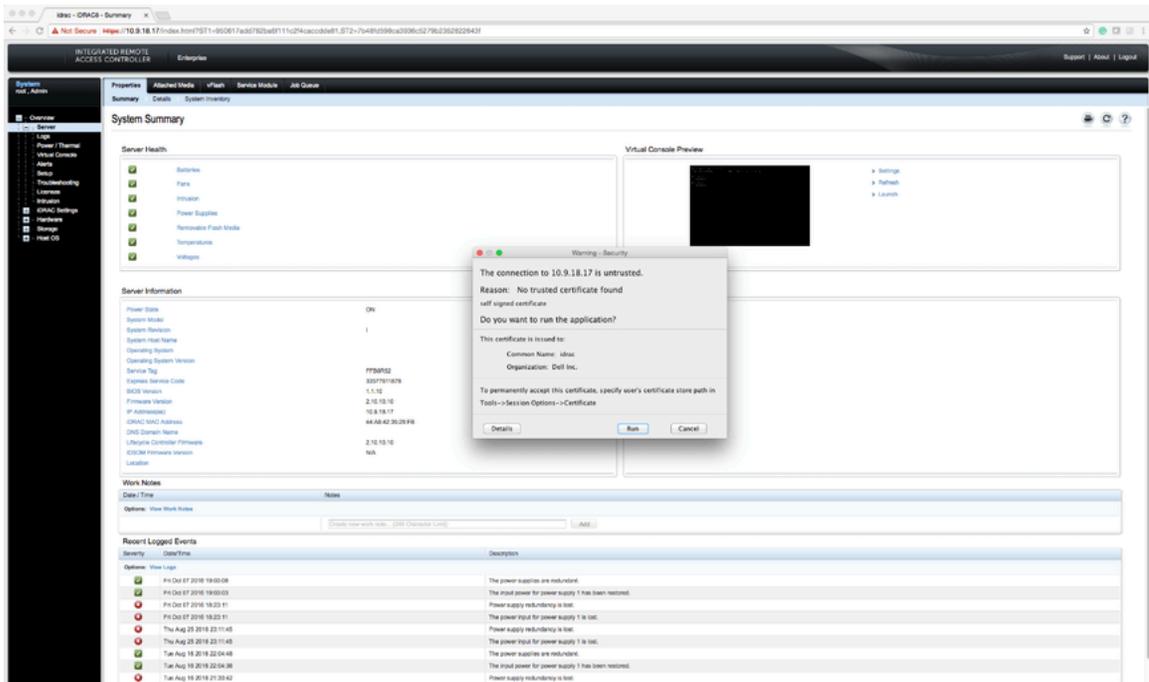
21. When prompted, click **Continue**.

Figure 13-13: System Summary: Run Prompt



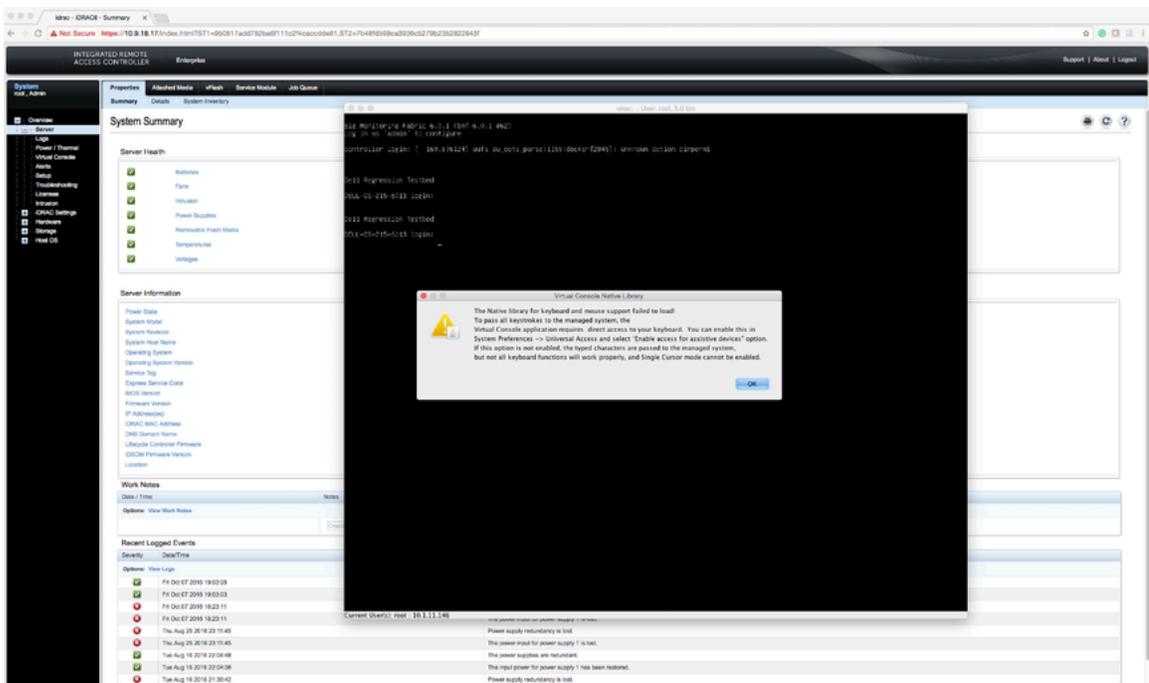
22. When prompted, click **Run**.

Figure 13-14: Confirm Connection to Untrusted Network



23. When prompted to continue with an untrusted certificate, click **Run**.

Figure 13-15: System Summary: Confirmation Prompt



24. When prompted, click **OK**.

iDRAC launches the virtual console window.

13.2 Using iDRAC to Install the DMF Controller or DMF Service Node Image

Before beginning, set up iDRAC as described in the [Setting Up and Configuring iDRAC](#) section, and then complete the instructions in this section.

After installing the Controller software image using iDRAC, follow the instructions in [Installing and Configuring the DMF Controller](#) to complete the DMF Controller installation and initial configuration.

For **DMF Release 6.3.2** onwards, the procedure for using iDRAC to install the DMF Controller or DMF Service node image is the same.

Complete the following steps using iDRAC to install the DMF Controller or Service Node image on a Dell R630/R730.

Procedure

1. Direct your browser to the iDRAC web interface and log in to launch the iDRAC virtual console.

Figure 13-16: iDRAC Virtual Console Window



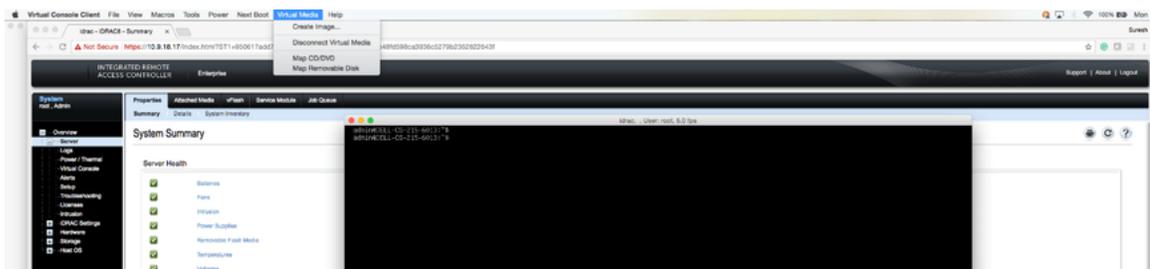
2. Select **Virtual Media > Connect Virtual Media**.

Figure 13-17: Virtual Media > Connect Virtual Media Option



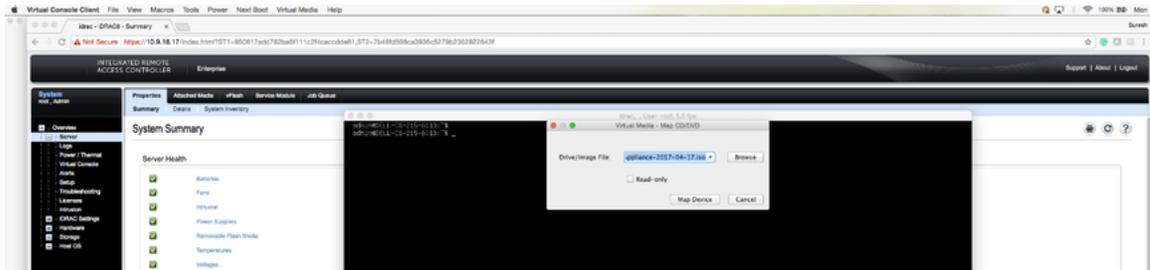
3. Click **Virtual Media** again and select **Map CD/DVD** this time.

Figure 13-18: Virtual Media > Map CD/DVD Option



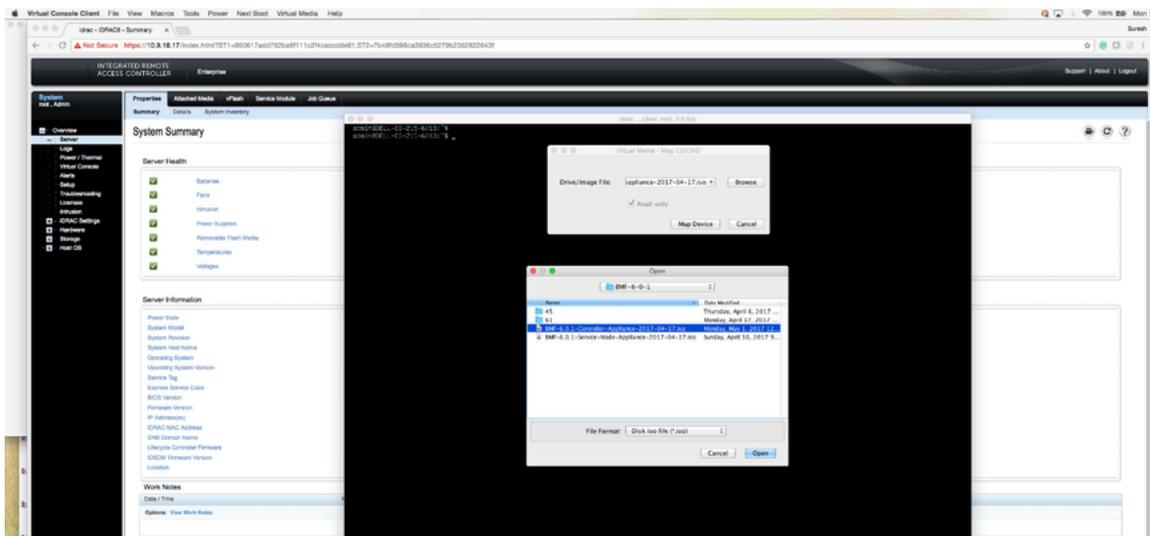
- Click **Browse** to choose the DMF Controller ISO file.

Figure 13-19: Virtual Media > Map CD/DVD Browse Option



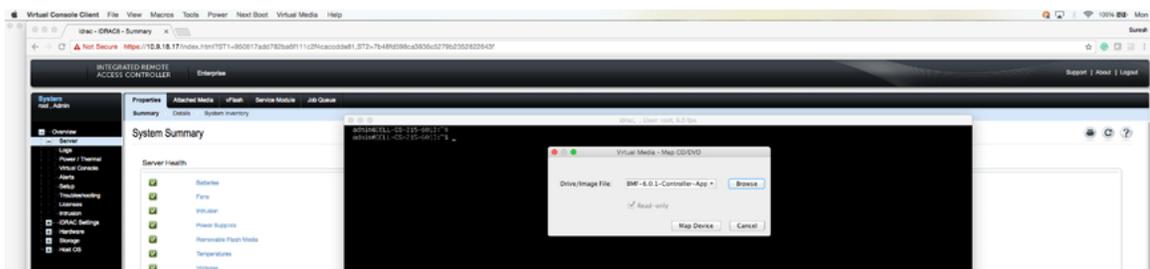
- Select the **DMF Controller ISO** image and click **Open**.

Figure 13-20: Open DMF Controller ISO Image



- Click **Map Device**.

Figure 13-21: Map Device



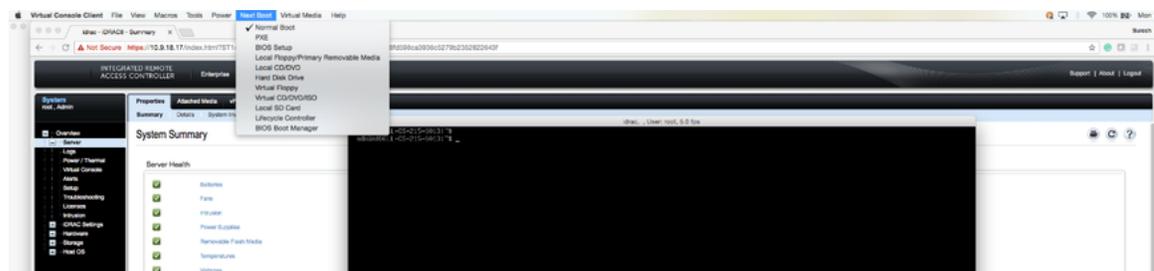
This action maps the DMF Controller ISO file to a Virtual CD/DVD on the Virtual Media menu.

Figure 13-22: Virtual Media > DMF Controller ISO Mapped to a Virtual CD/DVD



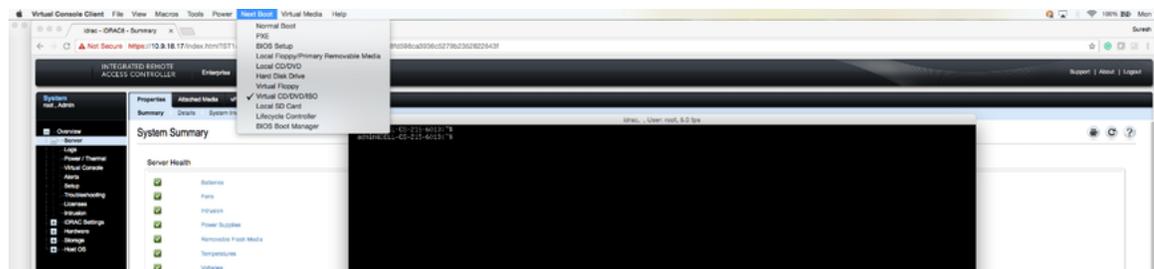
7. Click **Next Boot**.

Figure 13-23: Next Boot Menu



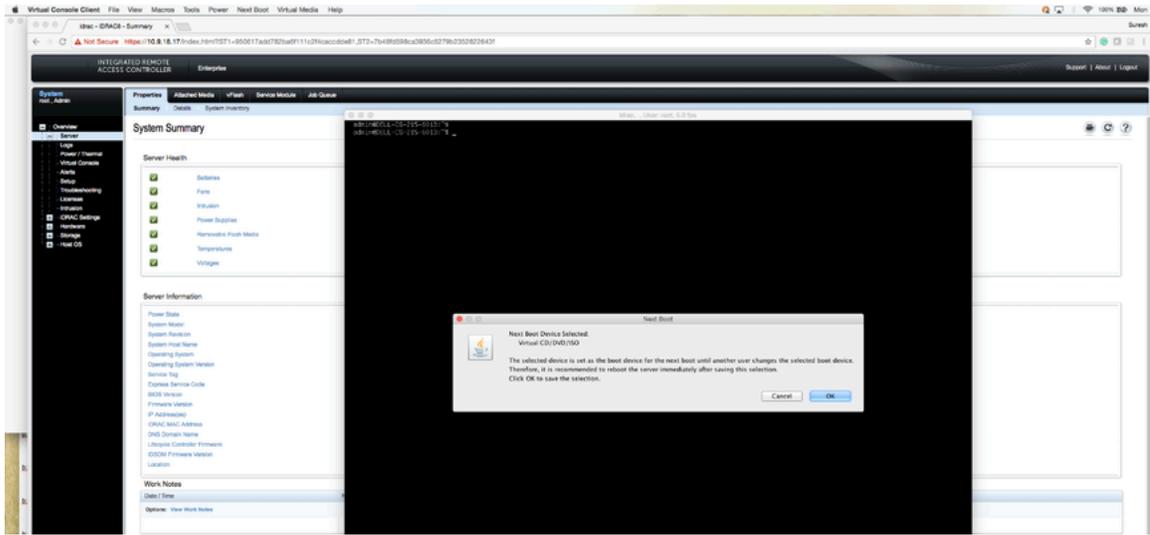
8. Select **Virtual CD/DVD** from the **Next Boot** menu.

Figure 13-24: Next Boot > Virtual CD/DVD/ISO Option



- When prompted, click **OK**.

Figure 13-25: Next Boot Confirmation Prompt



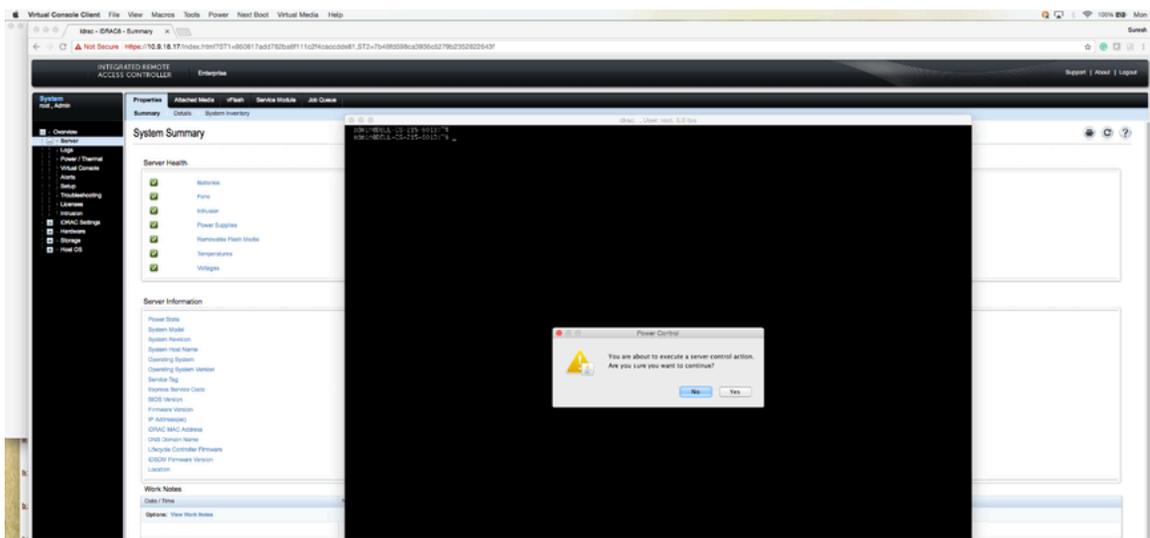
- Click **Power** and select **Restart System (warm boot)**.

Figure 13-26: Power Restart System (warm boot)



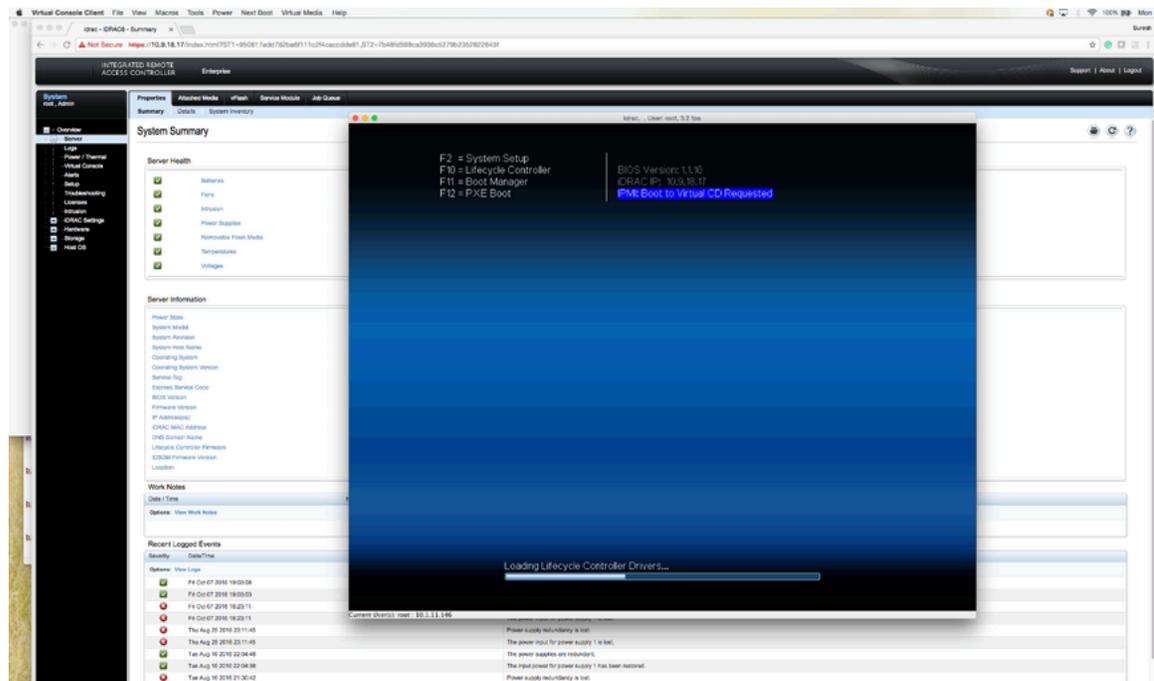
- When prompted, click **OK**.

Figure 13-27: Power Control Confirmation Prompt



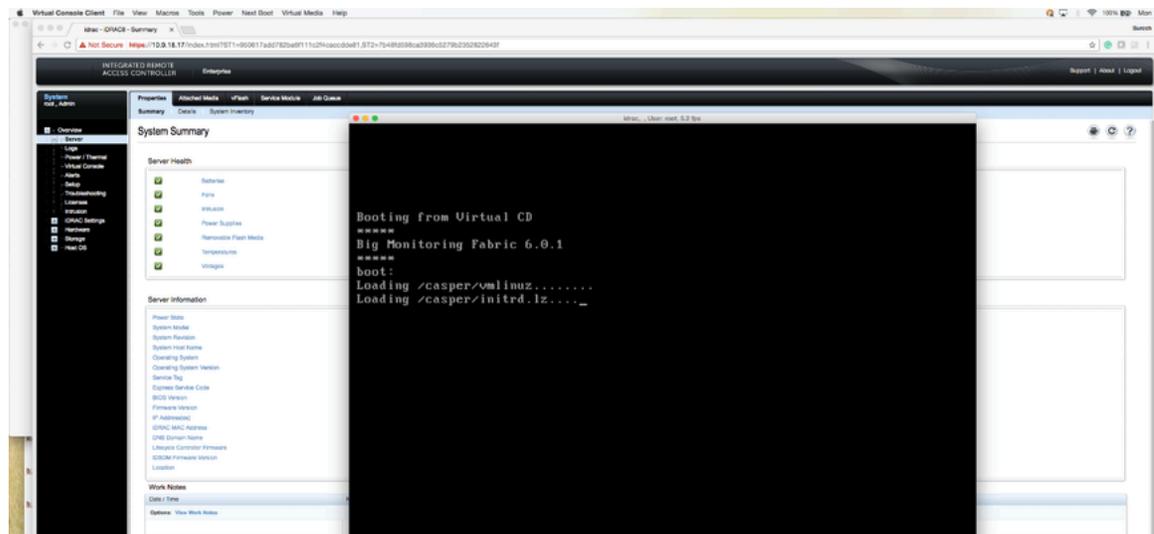
When the server boots up, it selects the Virtual CD/DVD as the boot device.

Figure 13-28: Booting from Virtual CD/DVD



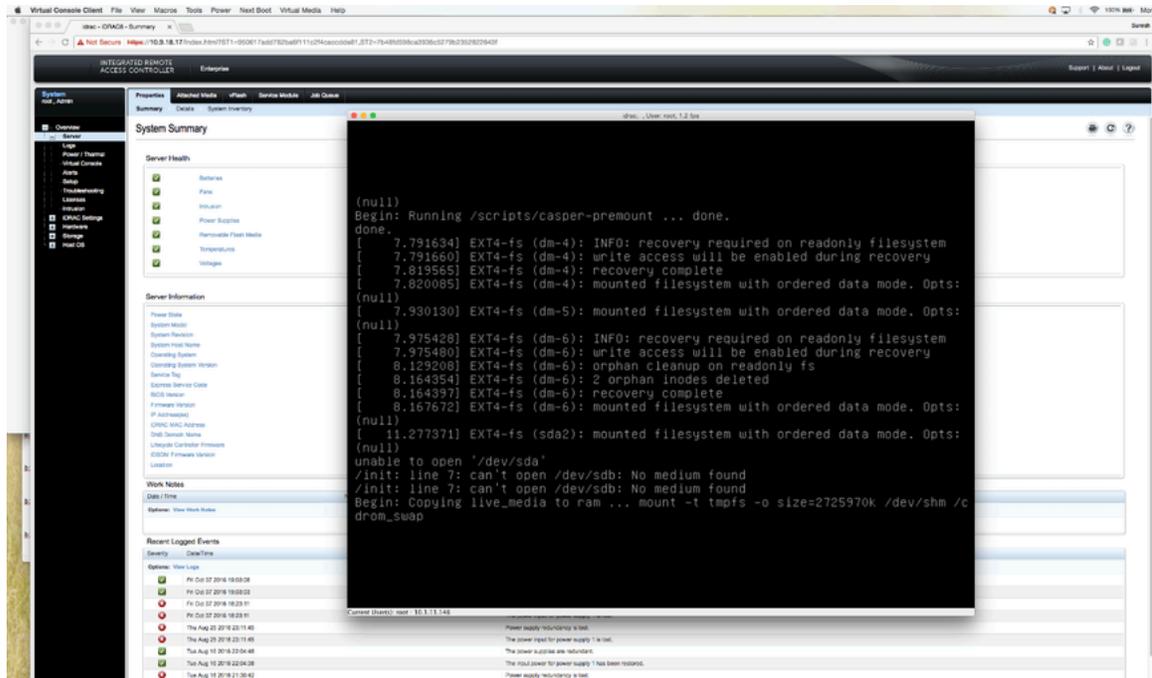
The server displays its status as it boots from Virtual CD/DVD.

Figure 13-29: Server Boot Status



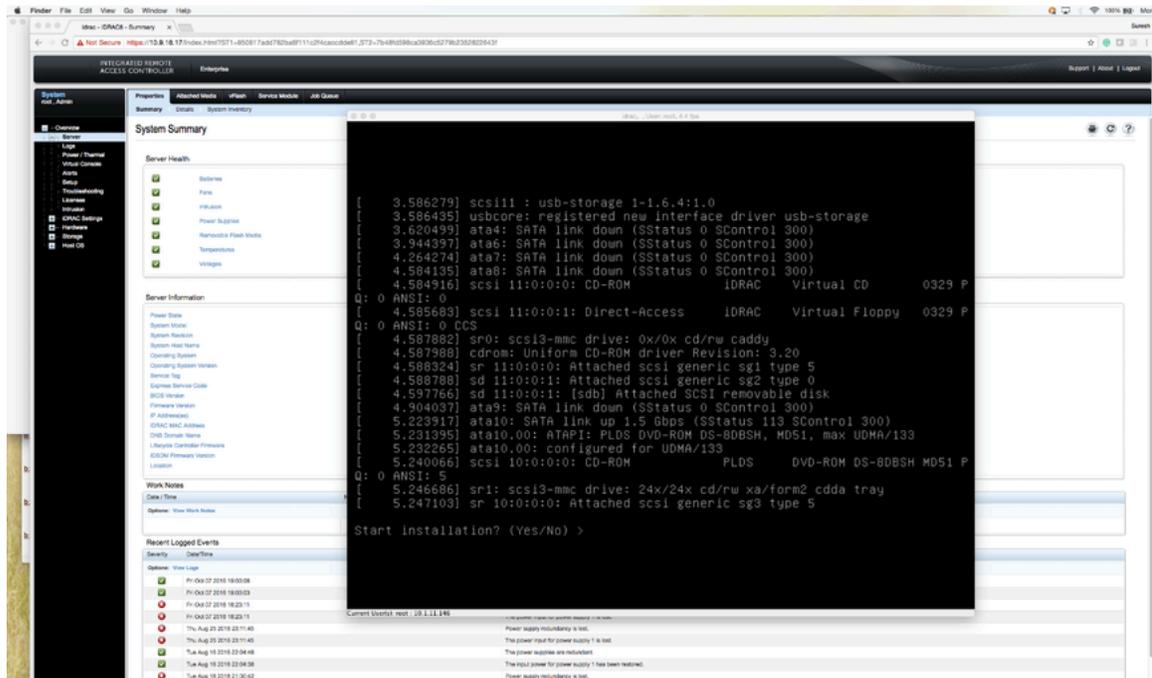
Note: Depending on the network speed, it may take a while to download the ISO image to the server.

Figure 13-30: Server Boot Status



12. When prompted, type **yes** to install the DMF Controller image on the server.

Figure 13-31: Installation Prompt



To complete the installation and initial configuration, refer to the instructions in this document.

Using iDRAC with a Dell R440 or R740 Server

This Chapter provides step-by-step instructions for using the integrated Dell Remote Access Controller (iDRAC) Enterprise Version to install DMF software images on Dell server hardware.

This Chapter was validated using iDRAC Enterprise. The instructions are similar for the following installation options.

- *DMF Controller Release 6.3.1* on the Dell R440 Server
- *Analytics 6.3.1* on the Dell R440 Server
- *DMF Recorder* on the Dell R740 Server

The instructions in this document tested using macOS are similar to those tested using Windows OS. You can use Internet Explorer, Chrome, or Safari browsers to access the iDRAC web interface.



Note: DMF Controller Nodes, Arista Analytics Nodes, DMF Recorder Nodes, and DMF Service Nodes using Dell R440/R640/R740 server hardware include an iDRAC Enterprise license.

14.1 Setting Up and Configuring iDRAC

To set up iDRAC on a Dell R440 or R740 server, complete the following steps.

Procedure

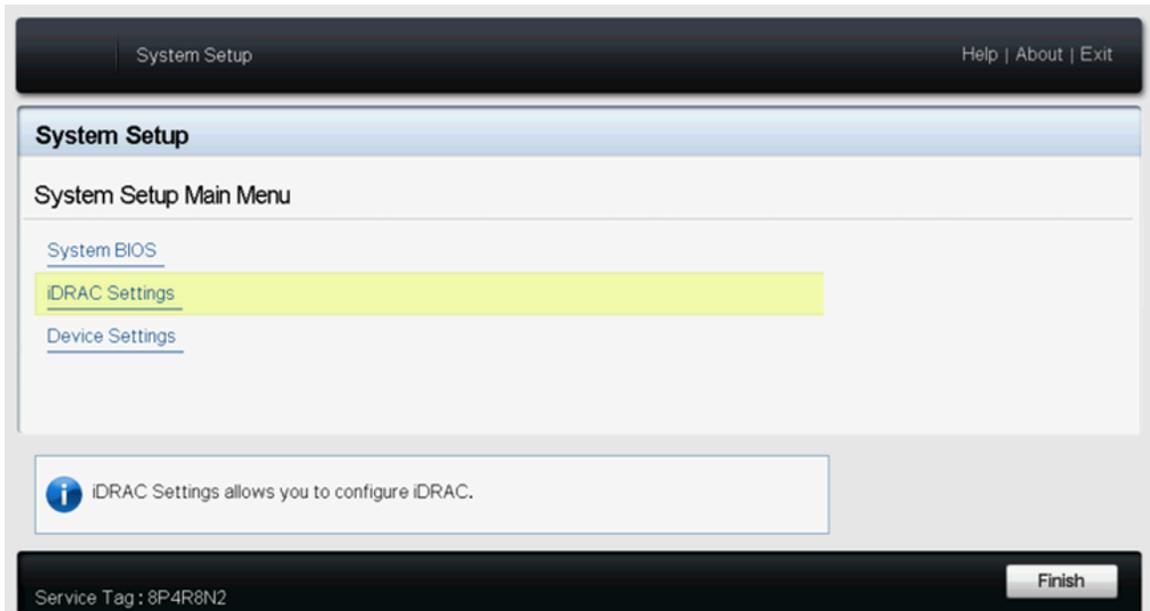
1. Turn on the Dell PowerEdge Server.
2. Connect a monitor to the VGA port and the keyboard to the USB port.



Note: When using the serial port to connect to the R440 and the R740, set the baud rate to **115200**.

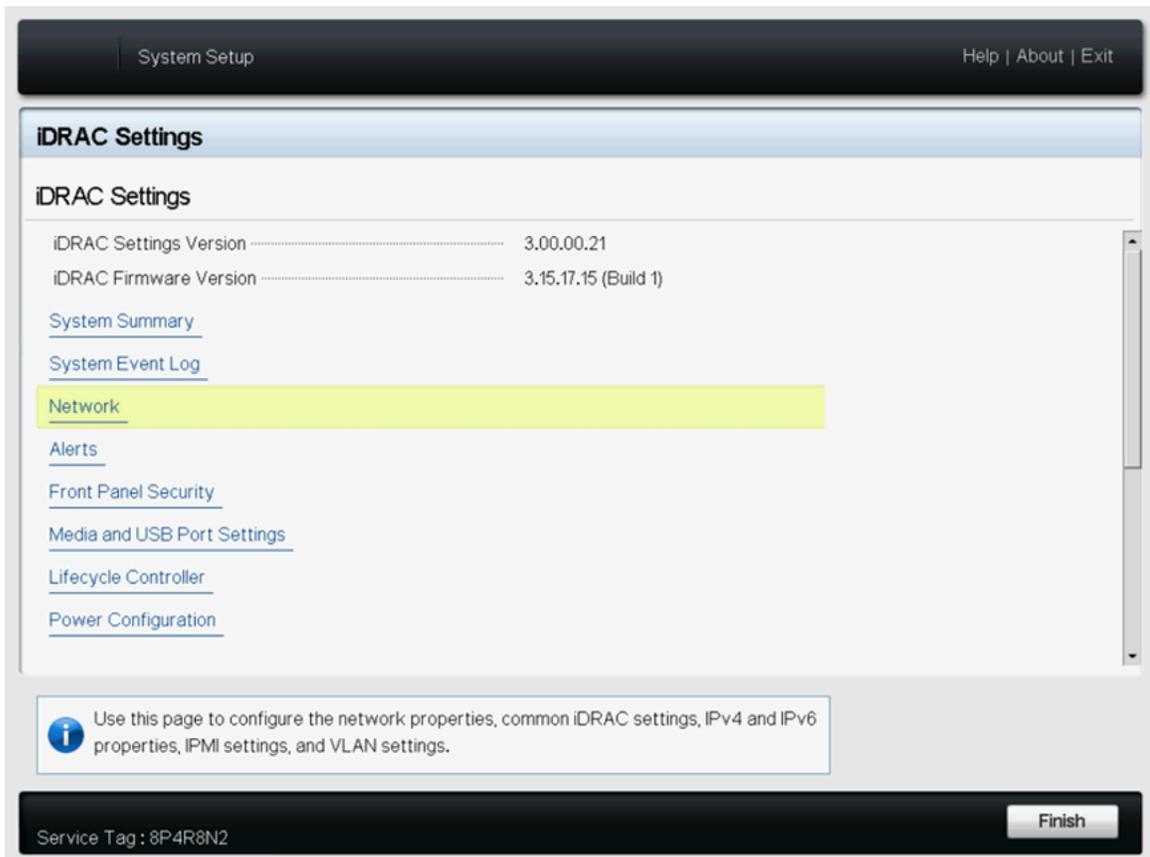
3. Press **F2** to enter the **System Setup Main Menu** screen.

Figure 14-1: System Setup Main Menu



4. Select **iDRAC Settings** from the **System Setup** Main Menu (**F2 > iDRAC Settings**).

Figure 14-2: iDRAC Settings



5. Use the arrow keys to select **Network**.

Figure 14-3: iDRAC Settings Network

System Setup Help | About | Exit

iDRAC Settings

iDRAC Settings • Network

NETWORK SETTINGS

Enable NIC	<input type="radio"/> Disabled	<input checked="" type="radio"/> Enabled	
NIC Selection	<input checked="" type="radio"/> Dedicated	<input type="radio"/> LOM1	<input type="radio"/> LOM2
Failover Network	<input checked="" type="radio"/> None		
MAC Address	D0:94:66:3A:1D:FA		
Auto Negotiation	<input type="radio"/> Off	<input checked="" type="radio"/> On	
Auto Dedicated NIC	<input checked="" type="radio"/> Disabled	<input type="radio"/> Enabled	
Network Speed	<input type="radio"/> 10 Mbps	<input type="radio"/> 100 Mbps	<input checked="" type="radio"/> 1000 Mbps
Active NIC Interface	Dedicated		
Duplex Mode	<input type="radio"/> Half Duplex	<input checked="" type="radio"/> Full Duplex	

COMMON SETTINGS

Register DRAC on DNS	<input checked="" type="radio"/> Disabled	<input type="radio"/> Enabled
----------------------	---	-------------------------------

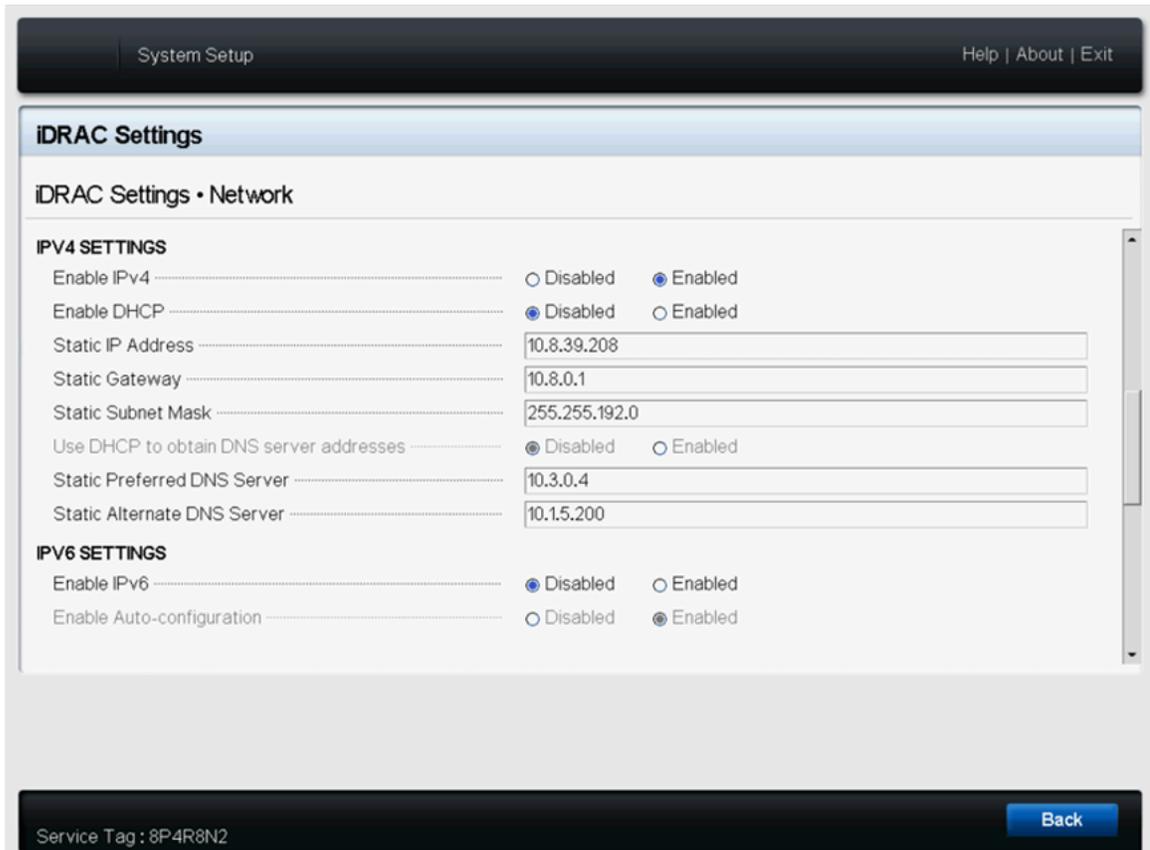
Help Message: Select Enabled to enable NIC. When NIC is enabled, it activates the remaining controls in this group. When a NIC is disabled, all communication to and ... (Press <F1> for more help)

Service Tag : 8P4R8N2 Back

6. Select the **Dedicated** option from the NIC selection list for use with iDRAC. A dedicated iDRAC port is available on R440 and R740 Dell servers.

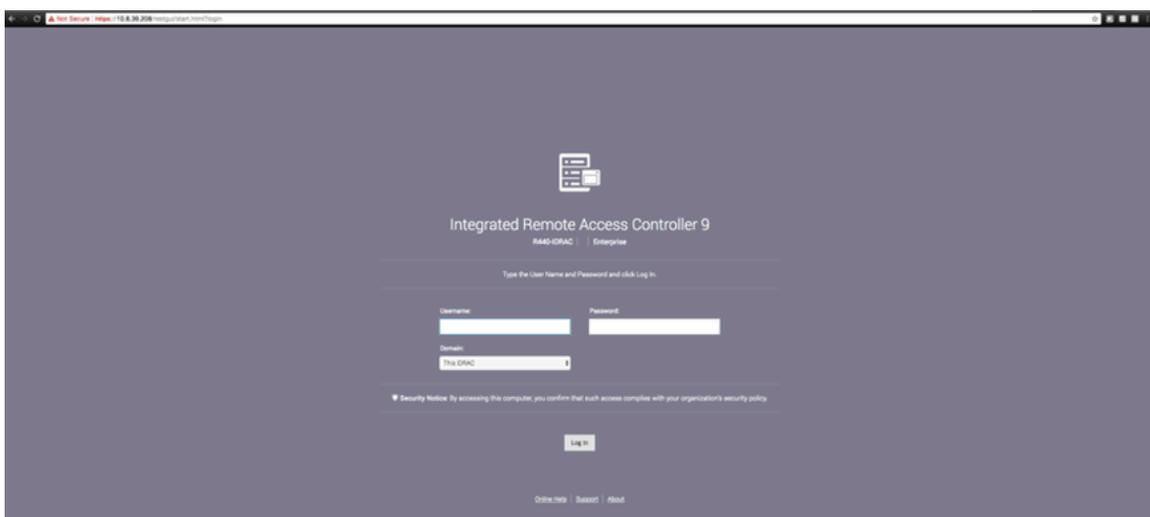
7. Configure the iDRAC IPv4 address, netmask, gateway, and DNS addresses.

Figure 14-4: iDRAC Settings Network (Completed Entries)



8. After completing the iDRAC configuration, press the **Esc** key to display the Exit menu.
9. Select **Save Changes and Exit** and then press **Enter** to keep the changes.
10. From a web browser, type **DRAC-IPv4-address** in the browser address bar and press **Enter**. The system displays the iDRAC web interface.

Figure 14-5: iDRAC Web Interface Login Window



11. Enter the username and password.

Default Password

By default, all PowerEdge servers ship with a unique iDRAC password, randomly generated at the factory, that is located on the pull-out Information Tag located on the front of the chassis, near the server asset tag. Customers who choose this option need to note this unique random password and use it to log in to iDRAC for the first time. For security purposes, Dell strongly recommends changing the default password.

Legacy Password

Customers who prefer to use the known legacy password *calvin* should choose this option. One reason to select this option would be to ensure conformance with any existing scripts. For security purposes, Dell strongly recommends changing the legacy password.

Refer to <https://www.dell.com/support/kbdoc/en-us/000133536/dell-poweredge-what-is-the-default-username-and-password-for-idrac> for the official Dell information.

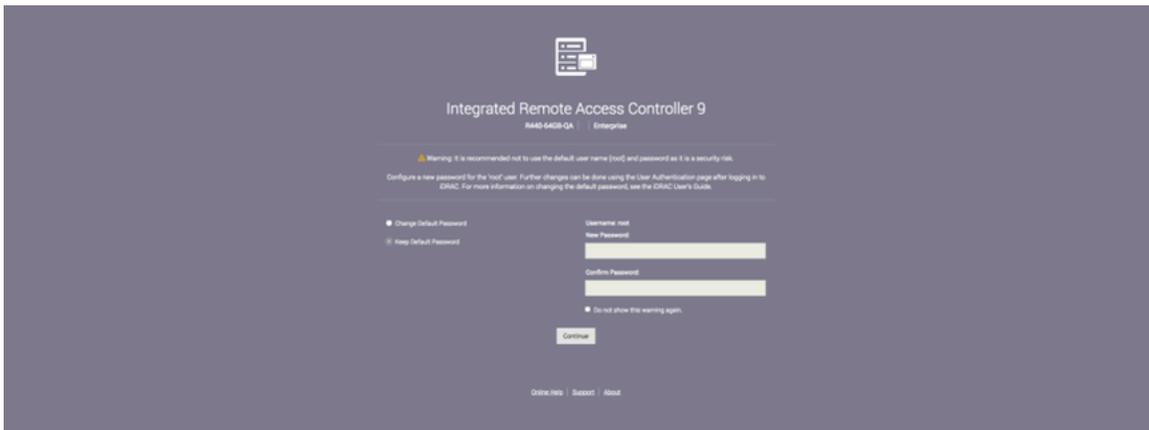
The pull-out information card is shown in the figure below:



Check out the bottom label on the pull-out information card to find the default password, which in this example is *94AXYXFERKNW*:



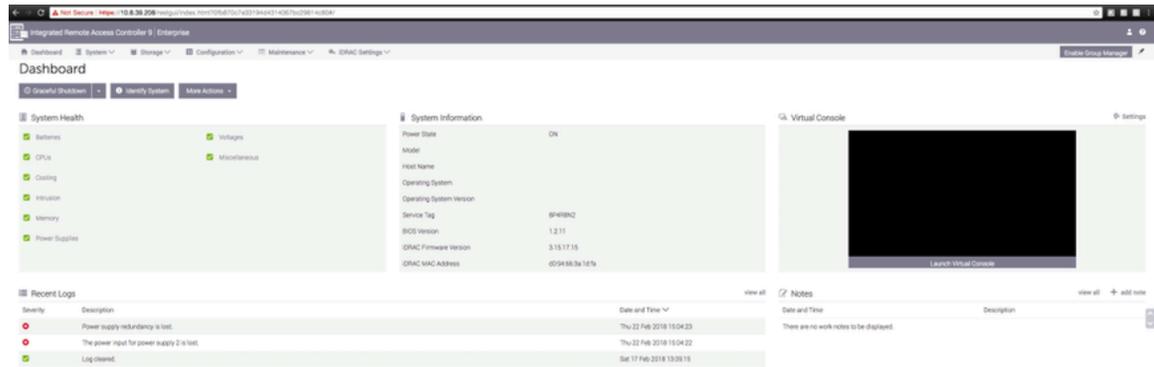
Figure 14-6: Change the Default Password



12. To change the default password, type the new password, when prompted, and type it again to confirm. In this example, the password stayed the same.
13. Log in to the iDRAC web interface using the new credentials (if changed).

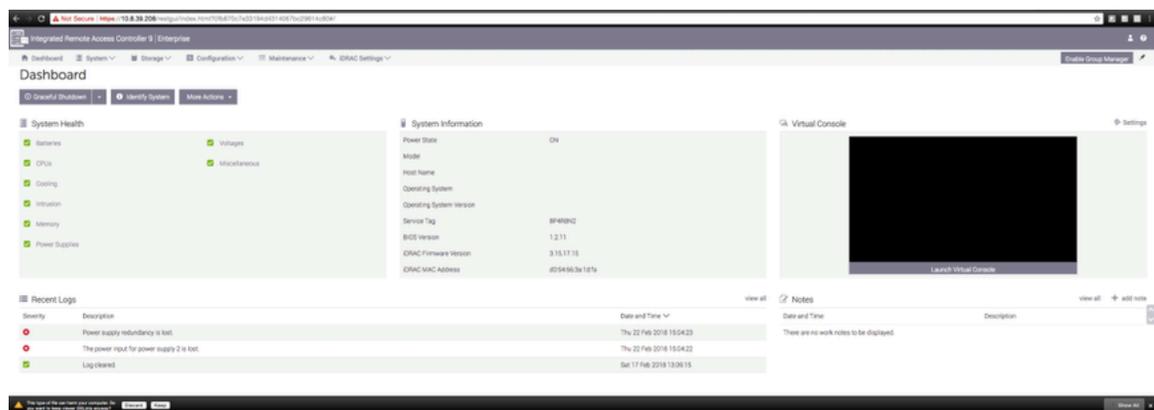
- Using the Virtual Console window option, click **Launch Virtual Console**.

Figure 14-7: System Summary Page



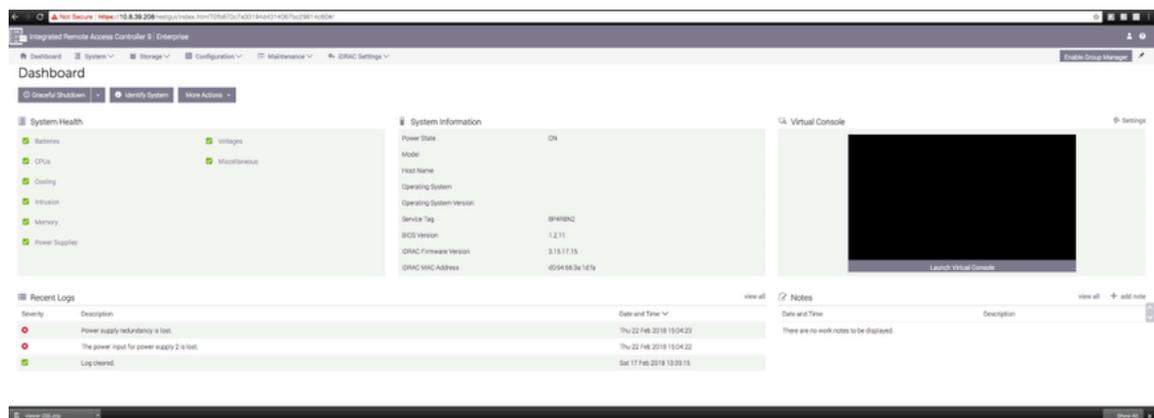
- When prompted at the bottom of the screen, click **Keep** to confirm the operation.

Figure 14-8: System Summary Page



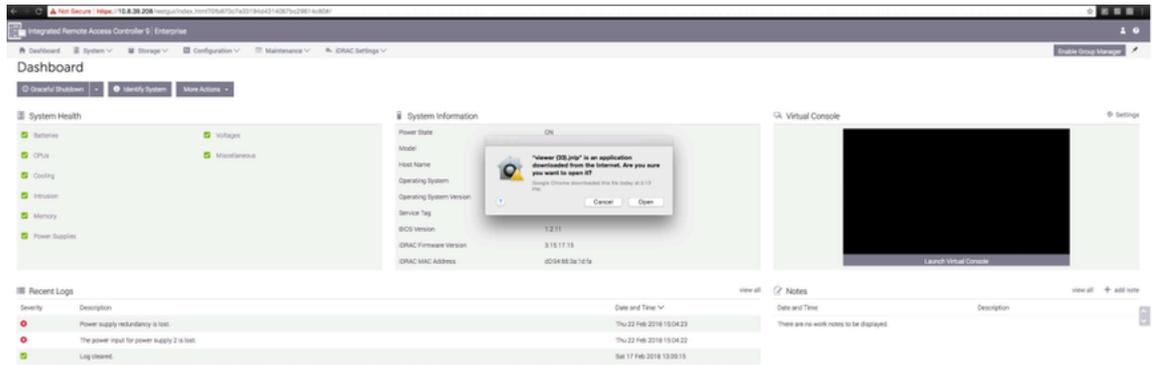
- Click the `viewer.jnlp` link, as shown in the following example.

Figure 14-9: System Summary Page



- When prompted, open the `viewer.jnlp` file.

Figure 14-10: System Summary Page



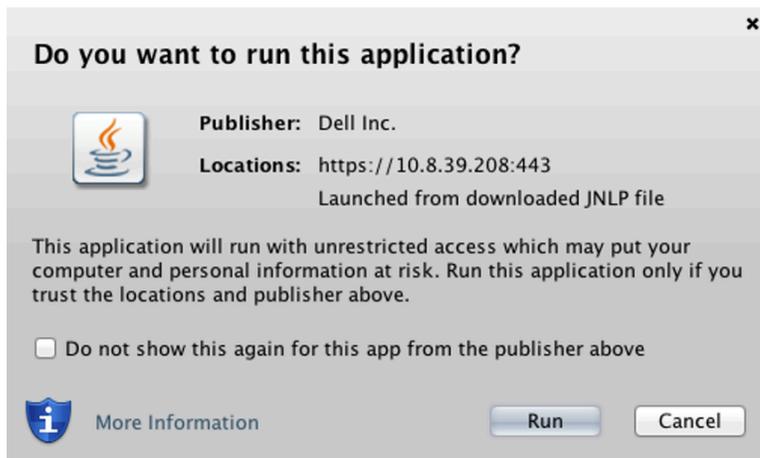
Note: Java is required, and the correct version downloads automatically.

Figure 14-11: System Summary: Continue Prompt



- When prompted, click **Continue**.

Figure 14-12: System Summary: Run Prompt



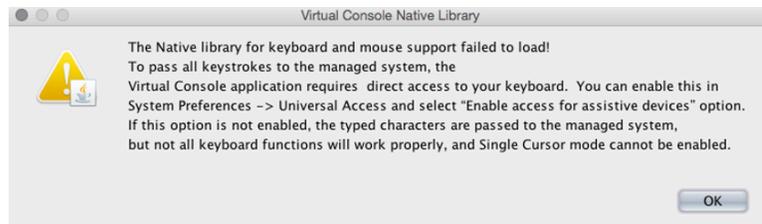
19. When prompted, click **Run**.

Figure 14-13: Confirm Connection to Untrusted Network



20. When prompted to continue with an untrusted certificate, click **Run**.

Figure 14-14: System Summary: Confirmation Prompt



21. When prompted, click **OK**.
iDRAC launches the virtual console window.

14.2 Using iDRAC to Install a DMF Controller, Analytics, or Recorder Software Image

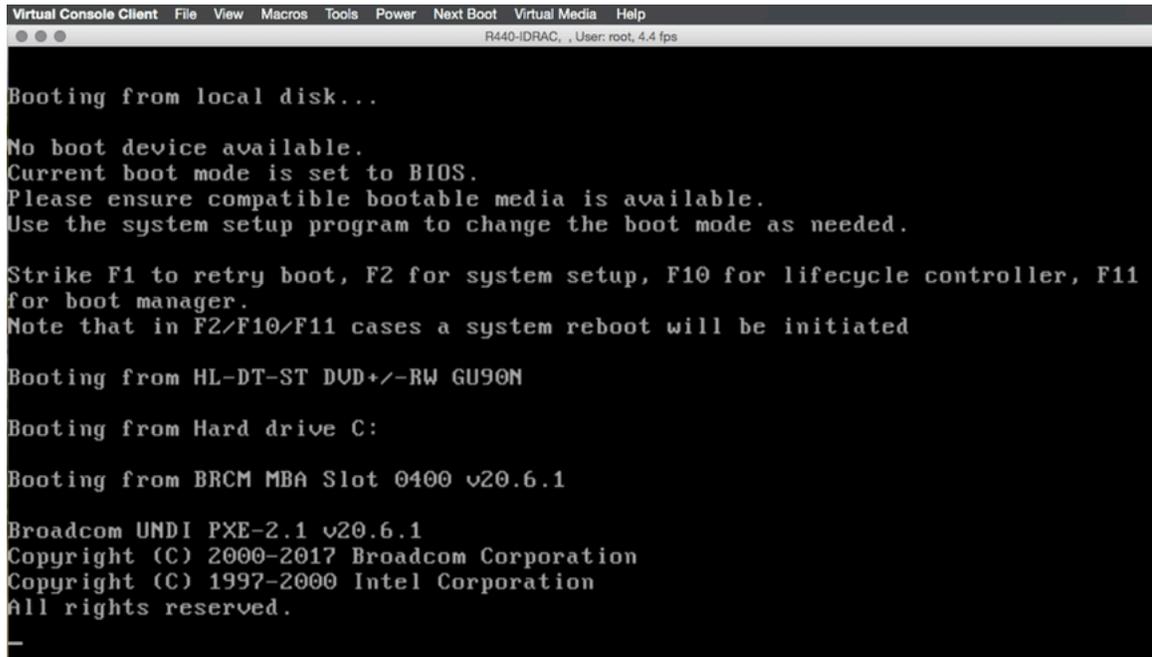
Before beginning, set up iDRAC as described in the [Setting Up and Configuring iDRAC](#) section, and then complete the instructions in this section. The procedure is similar for installing DMF Controller, Analytics, and Recorder software images.

After installing the software image using iDRAC, complete the installation and initial configuration described in this document's previous chapters.

To use iDRAC to install the software image, complete the following steps:

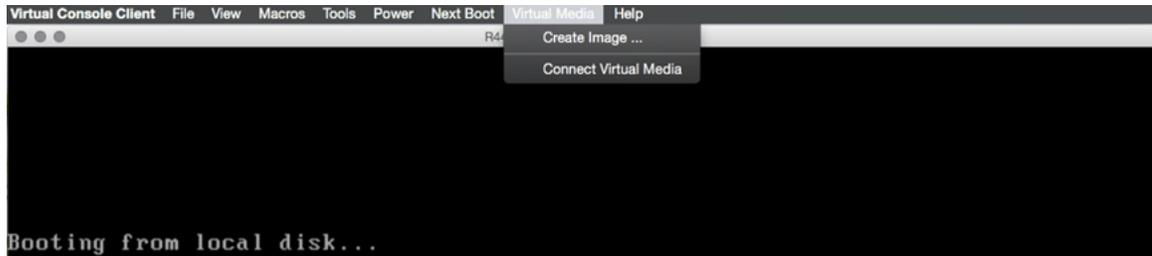
1. Direct your browser to the iDRAC web interface and log in to launch the iDRAC virtual console.

Figure 14-15: iDRAC Virtual Console Window



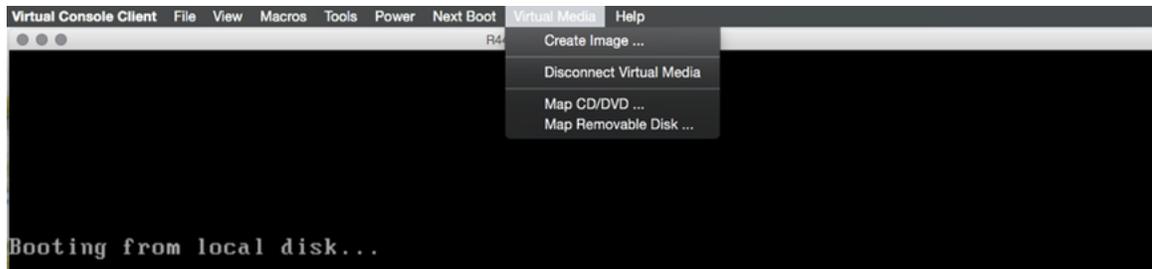
2. Select **Virtual Media > Connect Virtual Media**.

Figure 14-16: Virtual Media Connect Virtual Media Option



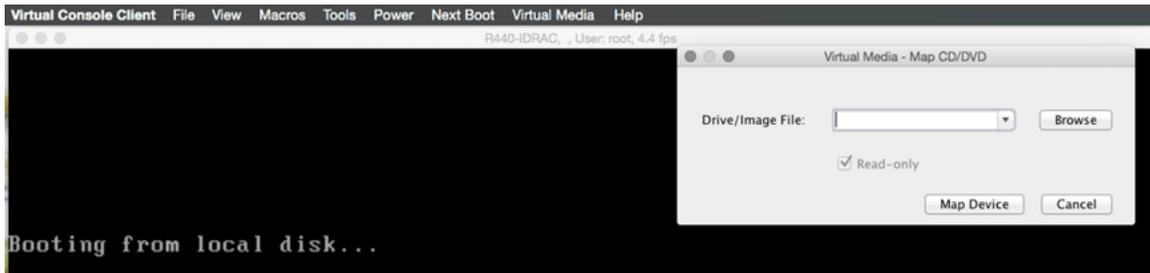
3. Click **Virtual Media** again and select **Map CD/DVD** this time.

Figure 14-17: Virtual Media Map CD/DVD Option



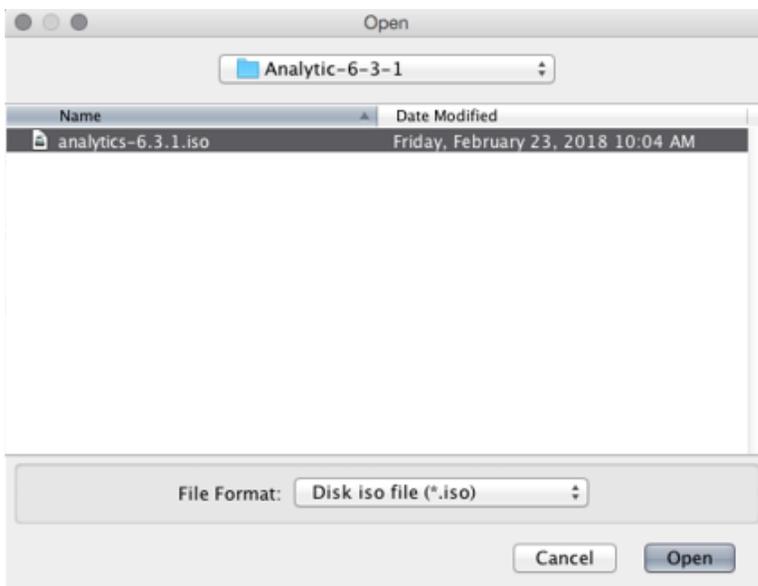
- Click **Browse** to select the ISO file you want to install.

Figure 14-18: Virtual Media Map CD/DVD Browse Option



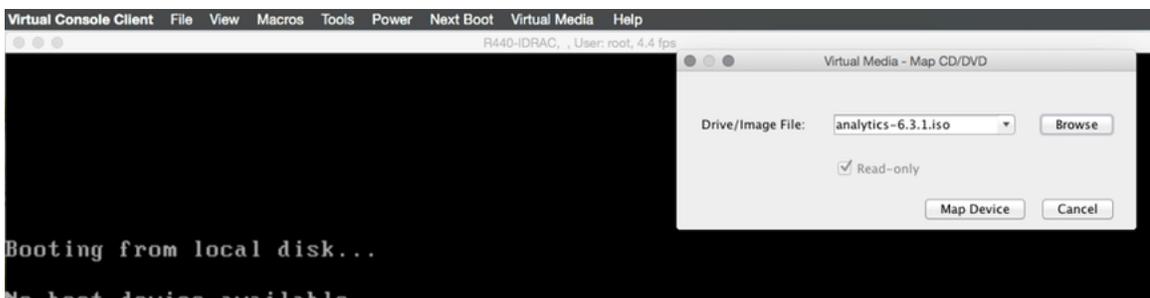
- Select the Analytics, DMF Controller, or Recorder software image to install and click **Open**. The following example shows selecting the Arista Analytics image.

Figure 14-19: Open DMF Controller ISO Image



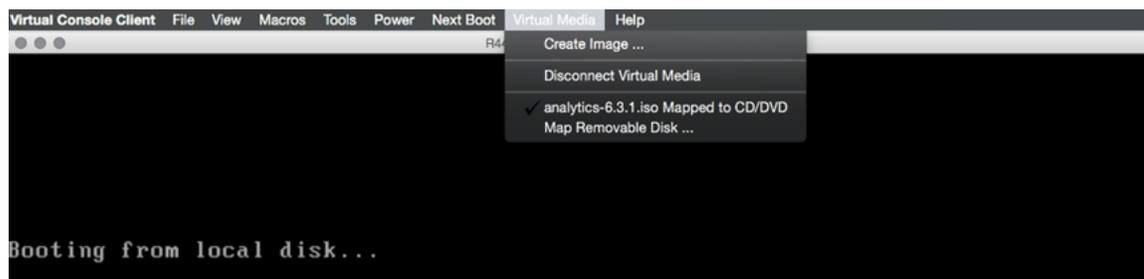
- Click **Map Device**.

Figure 14-20: Map Device



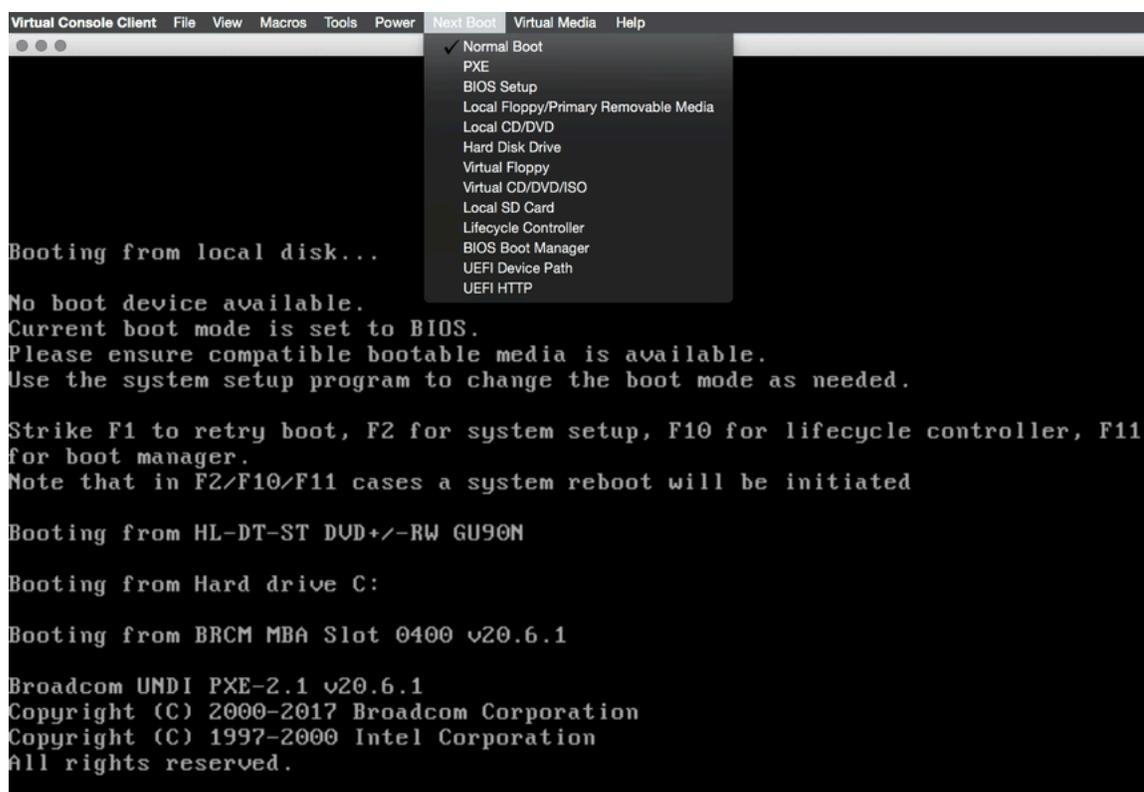
This action maps the DMF Controller ISO file to a Virtual CD/DVD on the Virtual Media menu.

Figure 14-21: Virtual Media DMF Controller ISO Mapped to a Virtual CD/DVD



7. Click **Next Boot**.

Figure 14-22: Next Boot Menu



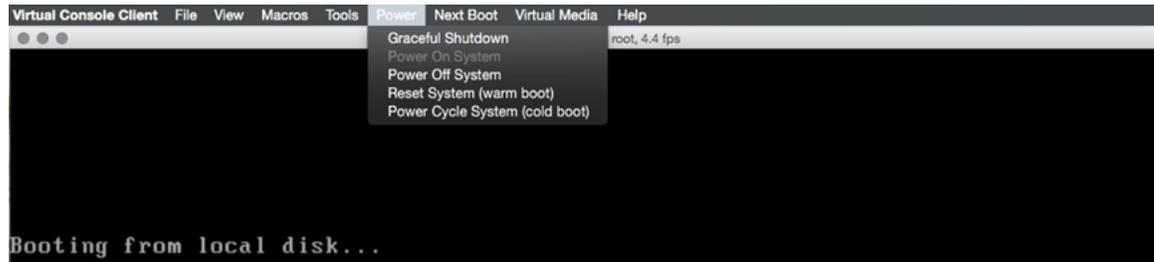
8. Select **Virtual CD/DVD** from the **Next Boot** menu.
9. When prompted, click **OK**.

Figure 14-23: Next Boot Confirmation Prompt



- Click **Power** and select **Restart System (warm boot)**.

Figure 14-24: Power Restart System (warm boot)



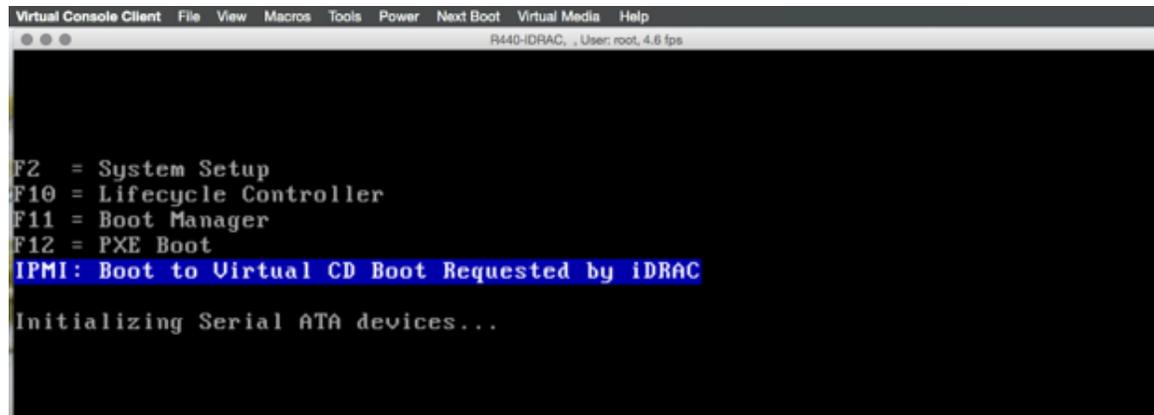
- When prompted, click **OK**.

Figure 14-25: Power Control Confirmation Prompt



When the server boots up, it selects the Virtual CD/DVD as the boot device.

Figure 14-26: Booting from Virtual CD/DVD



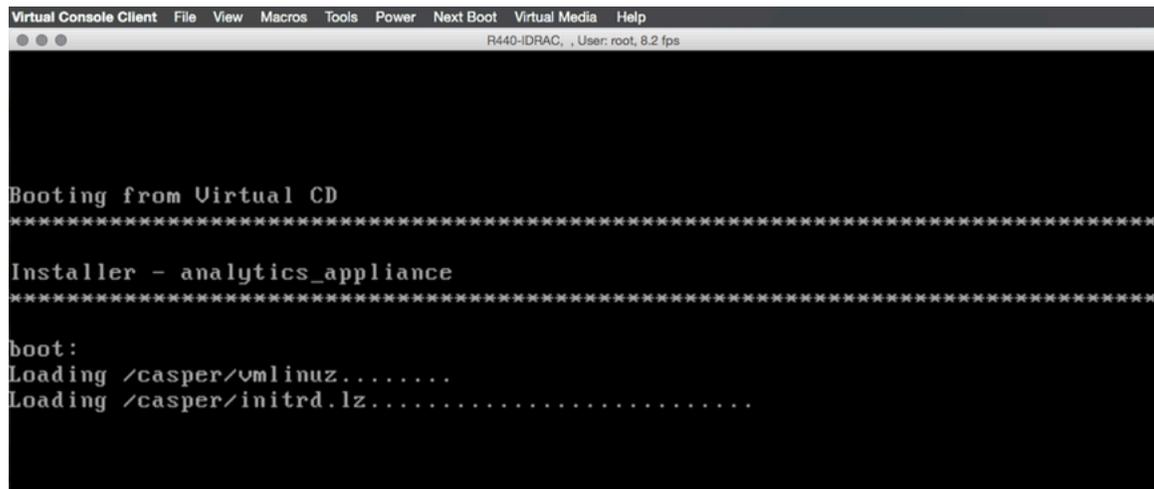
The server displays its status as it boots from Virtual CD/DVD.

Figure 14-27: Server Boot Status



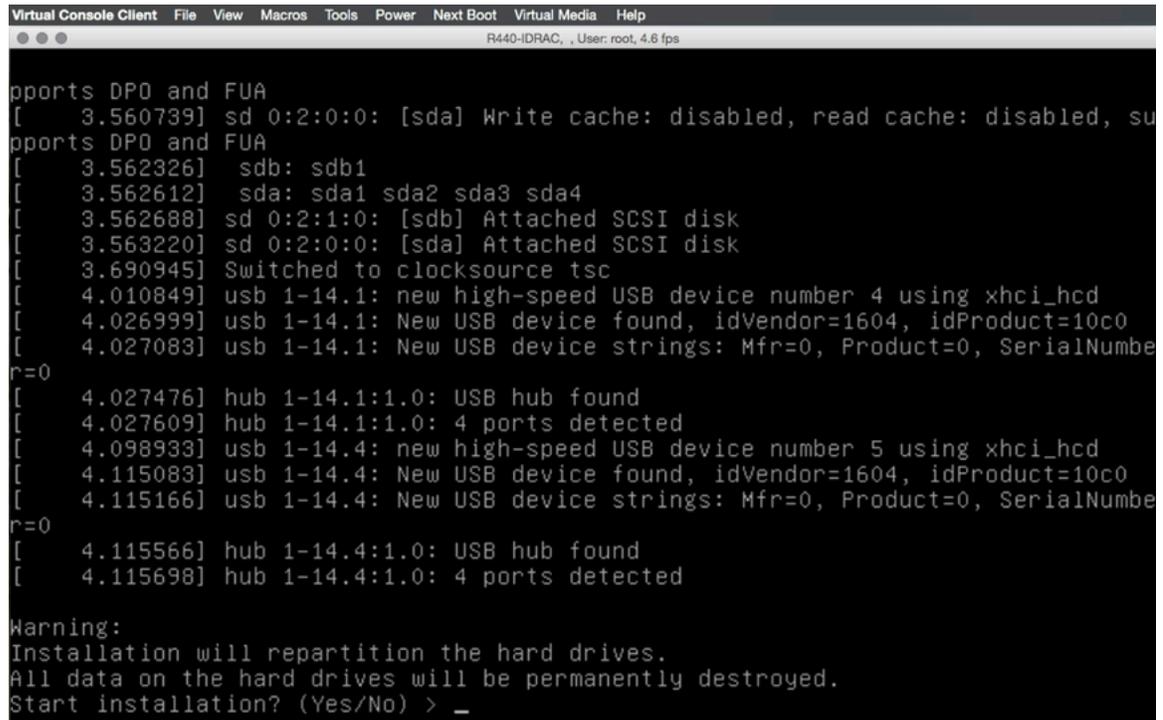
Note: Depending on the network speed, it may take a while to download the ISO image to the server.

Figure 14-28: Server Boot Status



- When prompted, type **yes** to install the DMF Controller image on the server.

Figure 14-29: Installation Prompt



```
Virtual Console Client  File  View  Macros  Tools  Power  Next Boot  Virtual Media  Help
R440-iDRAC, User: root, 4.6 fps

pports DP0 and FUA
[ 3.560739] sd 0:2:0:0: [sda] Write cache: disabled, read cache: disabled, su
pports DP0 and FUA
[ 3.562326] sdb: sdb1
[ 3.562612] sda: sda1 sda2 sda3 sda4
[ 3.562688] sd 0:2:1:0: [sdb] Attached SCSI disk
[ 3.563220] sd 0:2:0:0: [sda] Attached SCSI disk
[ 3.690945] Switched to clocksource tsc
[ 4.010849] usb 1-14.1: new high-speed USB device number 4 using xhci_hcd
[ 4.026999] usb 1-14.1: New USB device found, idVendor=1604, idProduct=10c0
[ 4.027083] usb 1-14.1: New USB device strings: Mfr=0, Product=0, SerialNumbe
r=0
[ 4.027476] hub 1-14.1:1.0: USB hub found
[ 4.027609] hub 1-14.1:1.0: 4 ports detected
[ 4.098933] usb 1-14.4: new high-speed USB device number 5 using xhci_hcd
[ 4.115083] usb 1-14.4: New USB device found, idVendor=1604, idProduct=10c0
[ 4.115166] usb 1-14.4: New USB device strings: Mfr=0, Product=0, SerialNumbe
r=0
[ 4.115566] hub 1-14.4:1.0: USB hub found
[ 4.115698] hub 1-14.4:1.0: 4 ports detected

Warning:
Installation will repartition the hard drives.
All data on the hard drives will be permanently destroyed.
Start installation? (Yes/No) > _
```

Switch CPLD Upgrade Procedure

For detailed instructions on upgrading the CPLD on the Dell switches, please refer to the *Firmware Upgrade for Dell Switches* document at <https://www.arista.com/assets/data/pdf/Dell-Switches-Firmware-Upgrade-Manual.pdf>.

Switch ONIE Upgrade Procedure

For detailed instructions on updating ONIE (Open Network Install Environment) on the Dell switches, please refer to the *Firmware Upgrade for Dell Switches* document at <https://www.arista.com/assets/data/pdf/Dell-Switches-Firmware-Upgrade-Manual.pdf>.

ONIE is a small operating system typically pre-installed as firmware on bare metal network switches. It provides an environment for automated provisioning.

Removing an Existing OS from a Switch

This appendix provides information on how to uninstall the OS from a switch.

A.1 Reverting from DMF (Switch Light OS) to EOS - 7050X3 and 7260CX3

This appendix details the procedure to revert the Arista 7050X and 7260X Series switch from Switch Light OS to EOS.

Procedure

To revert from SWLOS to EOS, complete the following steps:

1. Connect to the switch via a serial connection and confirm that a SWLOS is installed. The serial console output of the switch should approximate the following example.

```
Connected to 10.240.130.2.
Escape character is '^]'.
Switch Light OS SWL-OS-DMF-8.2.0(0), 2022-03-25.05:22-fdf3fa6
Site1-F1 login:
```

2. Log in to the switch via a serial connection and reboot the switch.
3. Enter Aboot by interrupting the boot process by pressing **Control-C**.

```
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0)
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
```

4. Once in Aboot, change the directory to **/mnt/flash**.

```
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
Aboot# cd /mnt/flash
Aboot# pwd
/mnt/flash Aboot#
```

5. List the files under the **/mnt/flash** directory and check the **boot-config** file for current swi.

```
Aboot# pwd
/mnt/flash
Aboot# ls
AsuFastPktTransmit.log SsuRestoreLegacy.log lost+found
EOS-4.23.3M.swi aboot-chainloader.swi onl
Fossil boot-config persist
SWL-INSTALLER.swi debug schedule
SsuRestore.log fastpkttx.backup startup-config
Aboot#
Aboot# cat boot-config
SWI=flash:aboot-chainloader.swi Aboot#
```

6. Edit the `boot-config` file to point to the existing EOS swi under the `/mnt/flash` directory.

```

About# vi boot-config
About#
SWI=flash:aboot-chainloader.swi
~
SWI=flash:/EOS-4.23.3M.swi
~
About# cat boot-config
SWI=flash:/EOS-4.23.0F.swi

```

7. Reboot the switch. The switch should boot up with the EOS image.

```

About# reboot
Requesting system reboot
Press Control-C now to enter About shell
Booting flash:/EOS-4.23.0F.swi
[ 13.231125] Starting new kernel
starting version 219
Failed to apply ACL on /dev/kvm: Operation not supported
Welcome to Arista Networks EOS 4.23.0F

```

A.2 Reverting from DMF (EOS) to EOS - 7280R

This appendix details the procedure to revert the Arista 7280R Series switch from DMF EOS to UCN EOS.

Procedure

To revert from DMF (EOS) to UCN EOS, complete the following steps:

1. Log in to the switch via a serial connection and reboot the switch.

```

Connected to 10.240.130.2.
Escape character is '^]'.
SAND-3 login: admin
Output to this terminal is being recorded for diagnostic purposes.
Note that only output that is visible on the console is recorded.
SAND-3>en
SAND-3#reload
! Signing certificate used to sign SWI is not signed by root certificate.
Proceed with reload? [confirm]

```

2. Enter About and interrupt the boot process by pressing **Control-C**.

```

agesawrapper_amdinitearly() returned AGESA_SUCCESS
Watchdog enabled, will fire in 2 mins
CBFS: 'Master Header Locator' located CBFS at [200:ffffc0)
CBFS: Locating 'normal/romstage'
CBFS: Found @ offset 5b3d40 size 7b7c
About 9.0.3-4core-14223577
Press Control-C now to enter About shell
^CWelcome to About.
About#

```

3. Once in About, change the directory to `/mnt/flash`.

```

Press Control-C now to enter About shell
^CWelcome to About.
About#
About# cd /mnt/flash
About# pwd
/mnt/flash About#

```

4. List the files under the `/mnt/flash` directory and check the `boot-config` file for current swi.

```

About# pwd
/mnt/flash

```

```

Aboot# ls
AsuFastPktTransmit.log
EOS-4.25.2F.swi
EOS-4.27.2F-26021868.uppsaladmfre1-x86_64.swi
EOS-4.27.2F-26021868.uppsaladmfre1-x86_64.swi.tmp
Fossil
SsuRestore.log
SsuRestoreLegacy.log
boot-config
debug
fastpkttx.backup
x86_64
lost+found
persist
schedule
startup-config
zerotouch-config
ztn-boot-info
Aboot#
Aboot# cat boot-config
SWI=flash:/EOS-4.27.2F-26021868.uppsaladmfre1-x86_64.swi Aboot#

```

5. Remove the **startup-config** file and edit the **boot-config** file to point to the existing EOS swi under the **/mnt/flash** directory.

```

Aboot# rm -rf startup-config
Aboot# vi boot-config
Aboot#
SWI=flash:aboot-chainloader.swi
~
SWI=flash:/EOS-4.25.2F.swi
~
Aboot# cat boot-config
SWI=flash:/EOS-4.25.2F.swi

```

6. Reboot the switch. The switch should boot up with the EOS image.

```

Aboot# reboot
Requesting system reboot
Press Control-C now to enter Aboot shell
Booting flash:/EOS-4.25.2F.swi
[ 13.231125] Starting new kernel
starting version 219
Failed to apply ACL on /dev/kvm: Operation not supported
Welcome to Arista Networks EOS 4.25.2F

```

A.3 Removing the existing OS from a Switch

Some switch platforms may have a preexisting operating system (OS) installed. When installing the Switch Light OS on top of an existing OS, there is a chance of failure. To avoid this issue, first uninstall any existing OS on the switch.

For example, to use a Dell switch with Force 10 OS (FTOS) pre-installed, remove FTOS before installing Switch Light OS. If FTOS is not first deleted, Switch Light installation may fail.

When you boot the switch, if only ONIE options are listed in the switch GNU GRUB boot menu, the switch does not have an existing OS installed. The following example shows the prompts that indicate no OS is installed on the switch.

Example 4: ONIE Prompts for a switch without an OS installed.

```

GNU GRUB version 2.02~beta2+e4a1fe391
+-----+
|*ONIE: Install OS |

```

```
| ONIE: Rescue |
| ONIE: Uninstall OS |
| ONIE: Update ONIE |
| ONIE: Embed ONIE |
| |
+-----+
```

If the switch prompt looks something like this, skip this section and proceed directly to the following section to install Switch Light OS.

Procedure

To delete FTOS from a Dell switch, complete the following steps:

1. Confirm an OS is installed on the switch.
2. Another OS exists if other options besides ONIE appear in the boot menu. The following example shows the options provided by FTOS installed on a Dell switch.

```
GNU GRUB version 2.02~beta2+e4a1fe391
+-----+
|*FTOS |
| FTOS-Boot Line Interface |
| DIAG-OS |
| ONIE |
| |
+-----+
```

3. After the switch has booted and the prompt for FTOS is displayed, change to **enable** mode.

```
DellEMC>enable
The SupportAssist EULA acceptance option has not been selected. SupportAssist can be
enabled once the SupportAssist EULA has been accepted. Use the: 'support-assist activate
' command to accept EULA and enable SupportAssist.
DellEMC#Feb 13 22:36:44 %STKUNIT1-M:CP %SEC-4-ENABLE_PASSW_NOT_CONFIGURED: Enable password
is required for authentication but
not configured (by default from console)
Feb 13 22:36:44 %STKUNIT1-M:CP %SEC-5-AUTHENTICATION_ENABLE_SUCCESS: Enable authentication
success on console DellEMC#
```

4. Reload the switch, do not save the configuration, and confirm the operation when prompted.

```
DellEMC# reload
System configuration has been modified. Save? [yes/no]: no
Proceed with reload [confirm yes/no]: yes
```

The following messages are displayed.

```
Feb 13 22:37:17 %STKUNIT1-M:CP %CHMGR-5-RELOAD: User request to reload the chassis syncing
disks... done
unmounting file systems...
```

5. To uninstall the OS, choose **ONIE** from the GNU GRUB boot menu, as shown in the following example.

```
GNU GRUB version 2.02~beta2+e4a1fe391
+-----+
|*FTOS |
| FTOS-Boot Line Interface |
| DIAG-OS |
|*ONIE |
| |
+-----+
```

The ONIE submenu is displayed, as shown in the following example.

```
GNU GRUB version 2.02~beta2+e4a1fe391
```

```
+-----+
| ONIE: Install OS |
| ONIE: Rescue |
| *ONIE: Uninstall OS |
| ONIE: Update ONIE |
| ONIE: Embed ONIE |
| EDA-DIAG |
| |
+-----+
```

6. From the ONIE submenu, choose ONIE Uninstall OS.

The uninstall process can take up to 15 minutes. After completion, the switch will automatically reboot.

The OS Uninstall log is displayed, starting with information about the existing OS, as shown in the following example.

```
ONIE: OS Uninstall Mode ...
Version : 3.27.1.1
Build Date: 2016-09-07T16:44-0700
Info: Mounting kernel filesystems... done.
Info: Mounting ONIE-BOOT on /mnt/onie-boot ...
<SNIP>
```

The following messages are displayed when the process is complete and the switch reboots.

```
Requesting system reboot
sd 4:0:0:0: [sda] Synchronizing SCSI cache
reboot: Restarting system
reboot: machine restart
```



Note: FTOS was successfully uninstalled if only the ONIE options are shown after the switch reboots, as in Example 1 at the beginning of this section. When you see only ONIE options, proceed to the next section to install Switch Light OS.

Creating a USB Boot Image

This appendix provides details on how to create a bootable USB with Switch Light OS.

B.1 Creating the USB Boot Drive with MacOS X

Complete the following steps to create a bootable USB drive on MacOS X.

Procedure

1. Insert the USB drive into a USB port on the Mac computer.
Inserting the USB drive mounts the USB drive, but it must be unmounted to create a bootable disk.
2. Open a Mac OS terminal window.
3. Enter the `diskutil` command to list all the mounted disks, as in the following example:

```
diskutil list
```



Note: You can also use the MacOS Disk Utility GUI application (applications/utilities) to identify the mounted disks and unmount the USB drive.

4. Identify the `/dev/disk<x>` label for the inserted USB drive.
5. Unmount the USB drive (this is different than ejecting) using the following command.

```
diskutil unmountdisk /dev/disk<x>
```



Note: Replace `<x>` with the unique numeric identifier assigned by the system.

6. Enter the `sudo dd` command in the terminal window to make the USB drive bootable.

```
sudo dd if=<path to iso image> of=/dev/rdisk<x> bs=1024m
```



Warning: Using the `dd` command with the wrong disk name can erase the installed OS or other vital information.

Copy the Service Node appliance ISO image to the USB drive using this command. Using `/dev/rdisk` makes the copying faster (`rdisk` stands for a raw disk).

Replace `<x>` with the drive identifier for the USB drive and replace `<path to iso image>` with the filename and path to the location where you downloaded the Service Node ISO image.

For example, the following command copies the file `dmf-service-node.iso` to **disk2**:

```
sudo dd if= dmf-service-node.iso of=/dev/rdisk2 bs=1024m
```

Copying the image to the USB drive can take up to ten minutes.

To monitor the progress of the write operation, enter the following command in a separate terminal window.

```
$ while sudo killall -INFO dd; do sleep 5; done
```

```
disk util eject
```

Alternatively, select Eject from the File menu.

B.2 Creating the USB Boot Image with Linux

Complete the following steps to create a bootable USB drive using Linux.

Procedure

1. Insert the USB drive into a USB port on the Linux workstation.
2. Enter the following command in a Linux terminal window to identify the USB drive.

```
disk -lu
```

On Linux, the USB drive is typically `/dev/sdb`.

3. Verify that the USB drive is not currently mounted, or unmount it if it is. Use the `mount` command to list the currently mounted devices.
4. Use the `sudo dd` command to make the USB drive bootable by copying the Service Node ISO image.

```
# sudo dd if=<path to iso image> of=/dev/sdb bs=4096
```



Warning: Using the `dd` command with the wrong disk name can erase the installed OS or other vital information.

Replace **<path to iso image>** with the filename and path to the location where you downloaded the Service Node ISO image. For example, the following command copies `dmf-service-node.iso` to the USB drive:

```
# sudo dd if=dmf-service-node.iso of=/dev/sdb bs=4096
```

Copying the image to the USB drive can take up to ten minutes.

5. Eject the USB drive from your Linux workstation.

B.3 Creating a USB Boot Image Using Windows

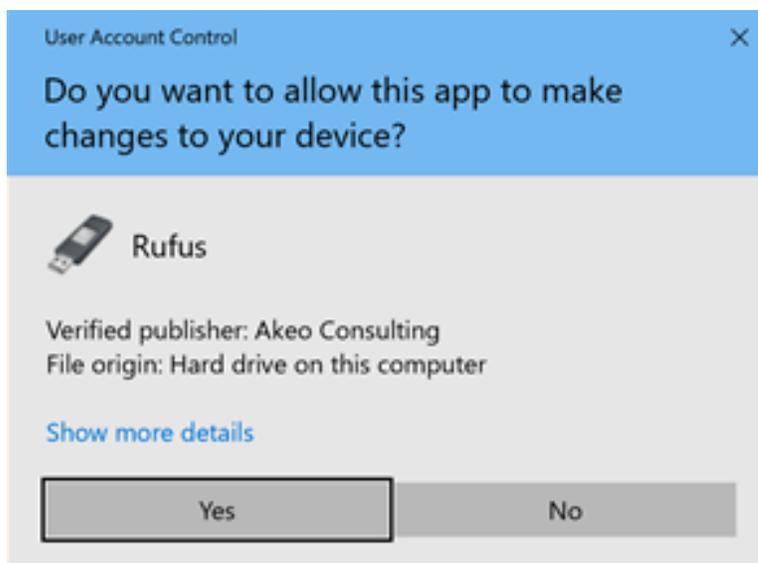
Several Windows utilities are available for building a USB boot image from an ISO image. The following procedure uses the Rufus bootable image program.

To build a USB boot image using Windows, complete the following steps.

Procedure

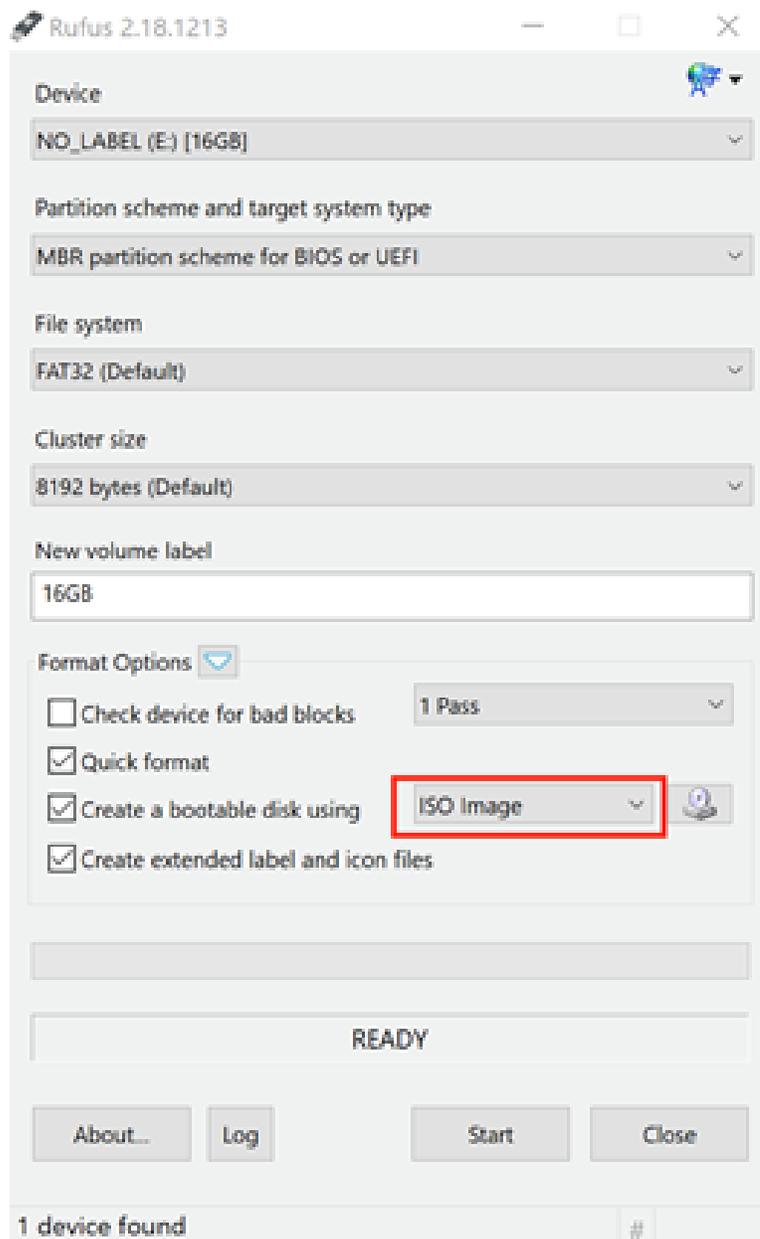
1. Download the Rufus utility from <https://rufus.akeo.ie/>.
2. After downloading the utility, double-click the **rufus.exe** file.

Figure B-1: User Account Control



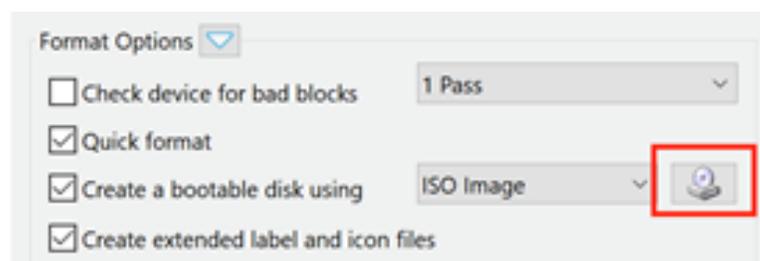
3. Click **Yes** to allow the changes required for installation.

Figure B-2: Rufus: Create an ISO Image Option



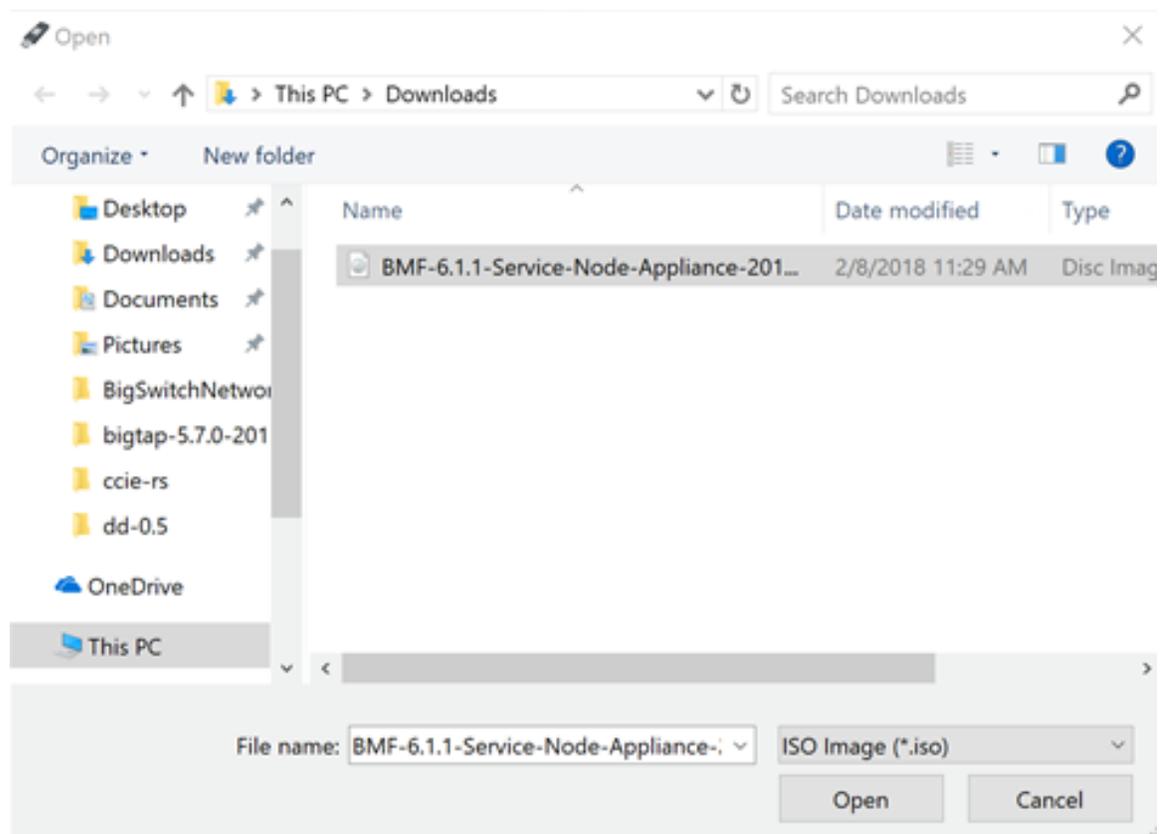
4. To create a bootable disk, select **ISO Image**.

Figure B-3: Rufus: Select ISO Image



5. Click the **CD-ROM** icon.

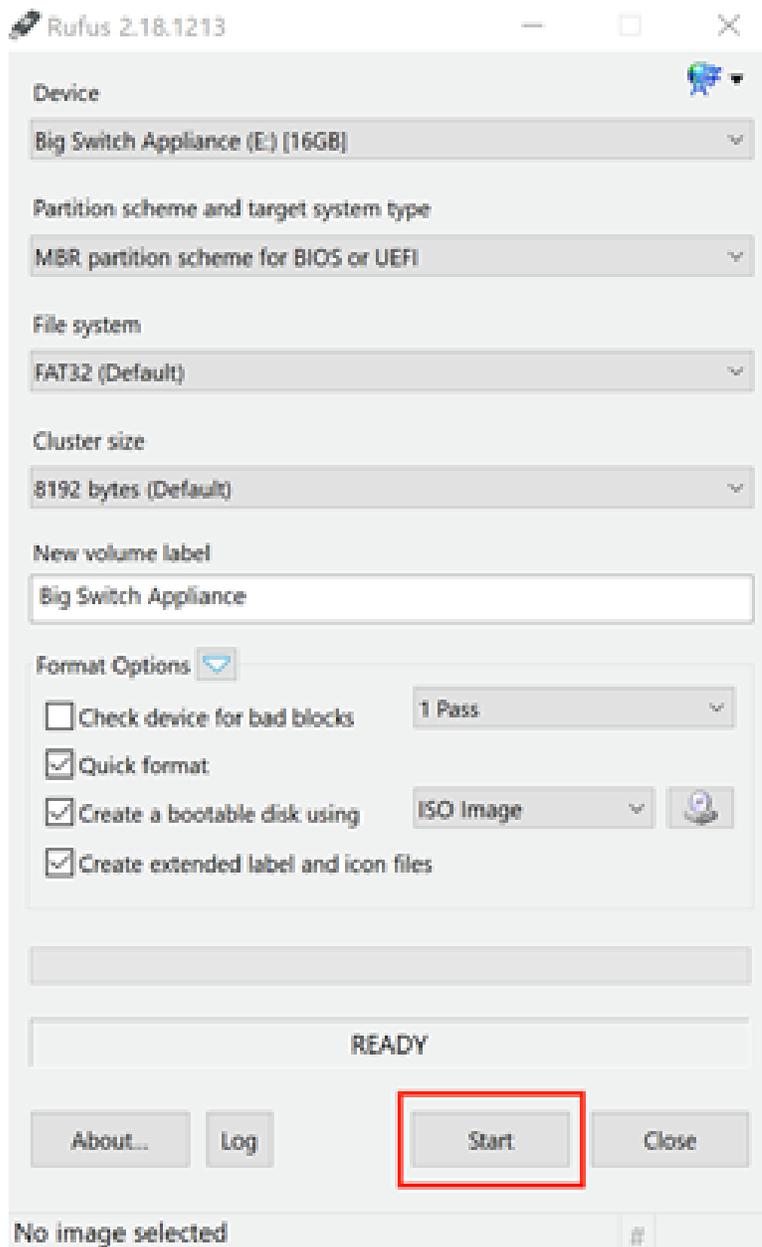
Figure B-4: Open ISO Image File



6. Select the file to use and click **Open**.

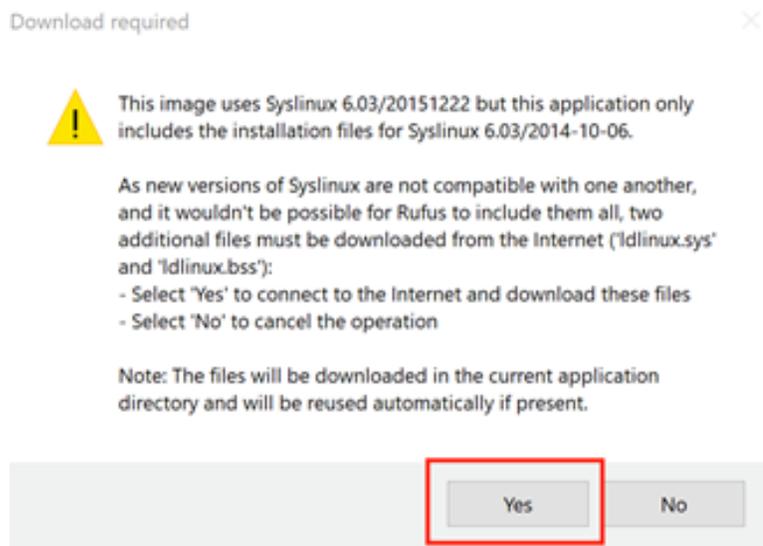
7. Click **Start** to burn the ISO image to USB.

Figure B-5: Rufus: Start



If an upgrade to syslinux is required, the system displays the following dialog box.

Figure B-6: User Account Control

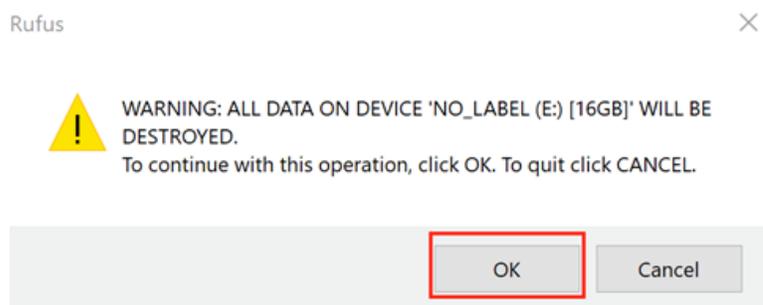


8. If this prompt appears, click **Yes** to continue.

9. When prompted to use DD mode or ISO mode, choose **ISO**.

The system displays a warning that the data on the USB drive will be destroyed, and a new image will be installed.

Figure B-7: Erasing Data Warning



10. Click **OK** to confirm the operation.

Installing a Controller VM

This appendix describes how to install and configure a VM for the DANZ Monitoring Fabric Controller.

C.1 General Requirements

The minimum hardware required for installing the VM software image on a supported Azure, KVM or ESXi version is as follows:

- 8 vCPUs with a minimum scheduling of 2 GHz.
- 32 GB of virtual memory.
- 400 GB hard disk.
- 1 virtual network interface reachable from the physical switches.



Note: The DANZ Monitoring Fabric Hardware Appliance is recommended for production deployment because VM performance depends on many factors in the hypervisor setup.

C.2 Installing on VMware ESXi/vSphere

C.2.1 Prerequisites

- A DMF Controller image in OVA format (.ova file).
- An ESXi host can be part of vSphere cluster. Refer to *the DMF Hardware Compatibility List* for supported versions of ESXi and vCenter.
- A virtual network has already been created.
- A machine with an installed vSphereClient or webClient.



Note: Do not use *ESXi 5.1 GA*. A known issue exists where installing a large VM causes the ESXi host to crash. Check the VMware website for more information.

C.2.2 VM Installation

To install the VM, complete the following steps:

1. Log in to vCenter or ESXi host with vSphere Client. The following uses vCenter 6.7.
2. Right-click on the host where the Controller VM will be deployed and select **Deploy OVF Template**.

3. Browse or enter the path to the OVA file and click **Next**.
4. Enter the name of the VM and click **Next**.
5. Select **Compute Resource** and click **Next**.
6. Leave the provision format at the default, select the datastore, and click **Next**.
7. Select **Network Mapping**. Map to the network created for the DMF Controller and switches. Click **Next**.
8. Click **Next** and click **Finish**.
9. Start the VM and complete the steps in Chapter 2, [Installing and Configuring the DMF Controller](#), to set up the Controller.

C.2.3 vMotion support for Virtual Controller

VMware vMotion is supported for Controller pairs. This support is offered only for **vCenter 7.0.2**.

The following are some additional points to remember when performing vMotion of the virtual Controllers:

- vMotion can be performed on each of the Controllers in the HA pair individually at separate times.
- Performing vMotion of the standby Controller first is recommended.
- Always back up the Controller configuration before performing vMotion.

C.3 Installing on Ubuntu KVM

C.3.1 Prerequisites

- DMF Controller virtual disk in **qcow2** format (**.qcow2** file).
- Ubuntu host with Virtual Machine Manager installed.
- Ubuntu host is connected to the management network via a bridge **br0**.

C.3.2 VM Installation

To install the VM, complete the following steps:

1. Copy **DMF-Controller-VM.qcow2** to **/var/lib/libvirt/images/**.
2. Start **Virtual Machine Manager**, and choose **Create a new virtual machine**.
3. Provide a name for the new VM and click **Import existing disk image** options. Click **Forward** to continue.
4. Set the existing storage path to point to the provided DMF Controller image. Press **Forward** to continue.
5. Set the **Memory (RAM)** and **CPU**. Allocate at least 4G RAM and 2 CPU instances for the image. Click **Forward** to continue.
6. Select the checkbox **Customize configuration before install**. Expand **Advanced options**, change to **Specify shared device name**.
7. Enter the **Bridge name: br0**, to bind the Controller virtual machine to the br0 bridge interface created previously. Click **Finish** to continue.

8. Under the **Processor** section, expand **Configuration**. Select the **Copy host CPU** configuration option. This may improve performance dramatically, depending on your VM host. Click **Apply** to save the changes.
9. Under the **Disk 1** section, expand **Advanced options and Performance options**. Set the options on this page as follows:
 - Disk bus: VirtIO
 - Storage format: qcow2
 - Cache mode: Writeback
 - IO mode: default
10. Click **Apply** to save the changes.
11. Under the **NIC** section, set the **Device model** to **virtio**. Click **Apply** to save the changes.
12. Select **Begin installation** to create the virtual machine.
13. Follow the steps in [Installing and Configuring the DMF Controller](#) to set up the Controller.

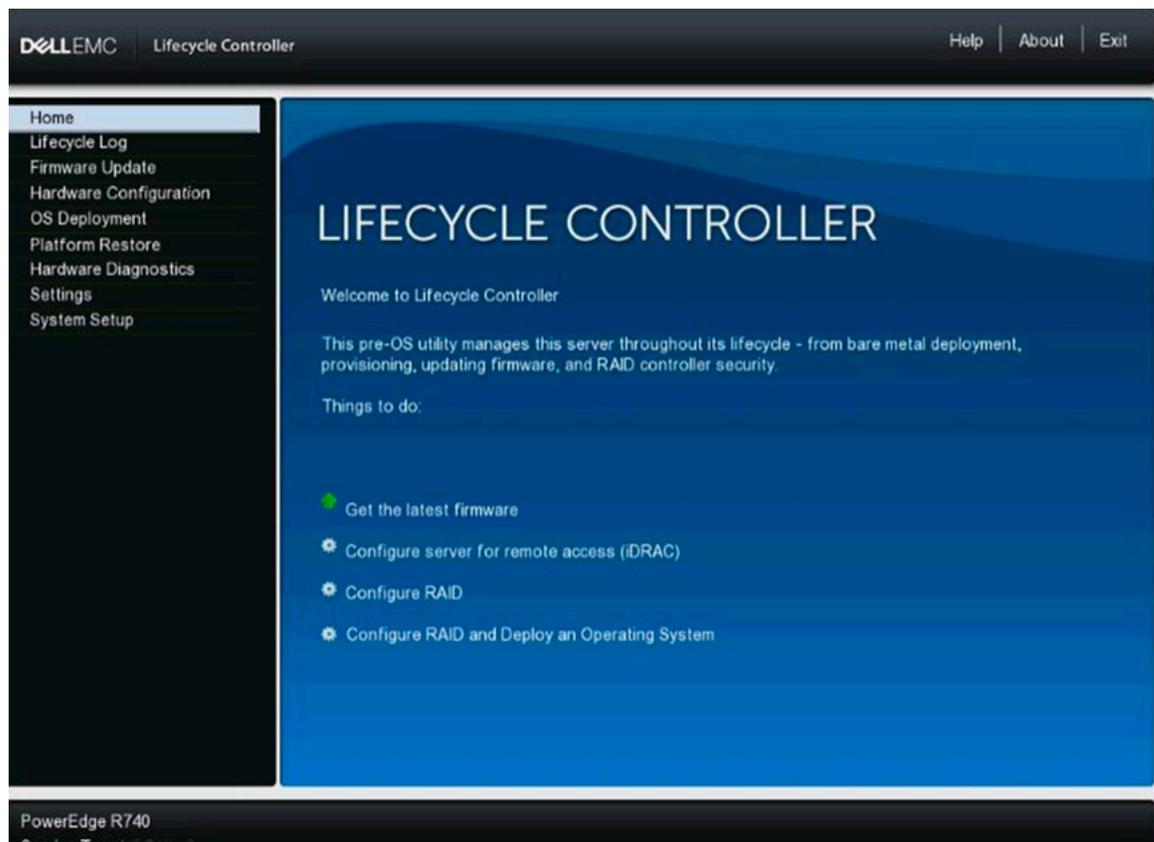
Erasing DMF Appliance

Overwrite data to erase data stored on the DMF appliance securely. This section describes how to do this using the Dell LifeCycle Controller. There is an erase API for the DMF Recorder Node, but it does not securely remove the data from its disk. Instead, it unlinks files in the Index and Packet partitions so the file system can reclaim the space for future packets and indices. This approach was a design decision because unlinking files is faster than overwriting data. To securely erase data and to prevent anyone from accessing the data, use the following procedure.

D.1 Using the Dell Lifecycle Controller

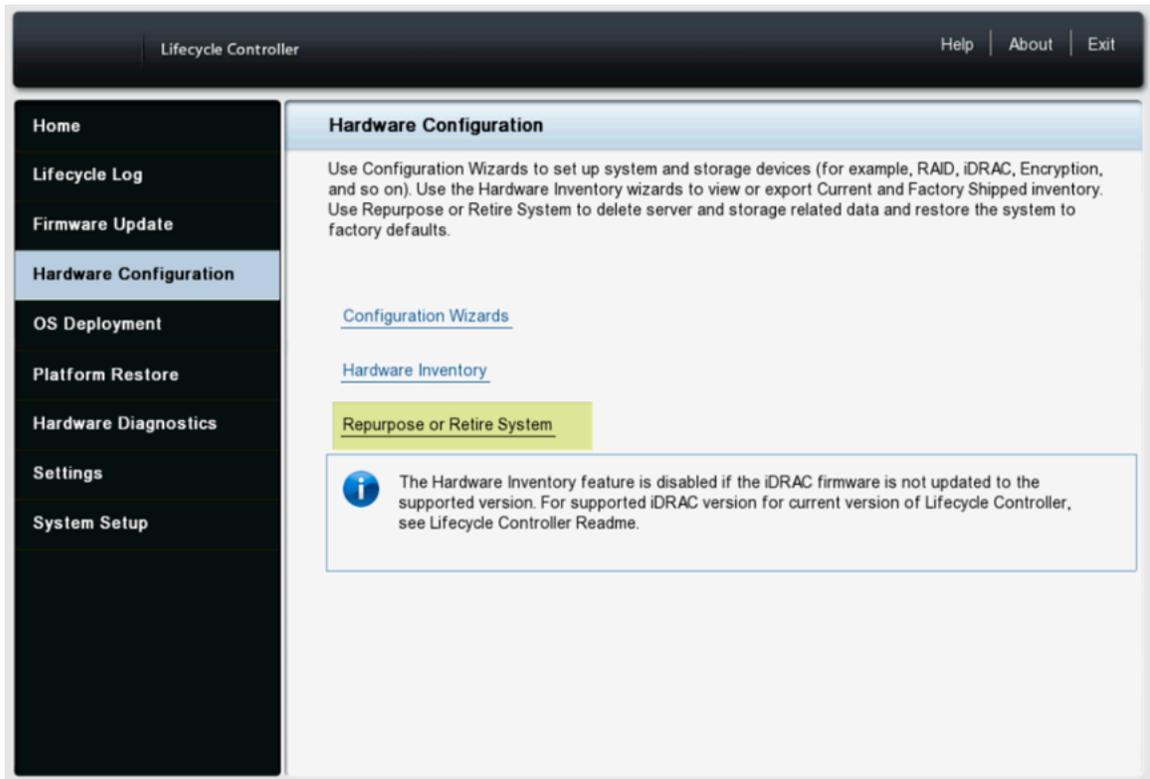
1. Restart the Recorder Node.
2. During POST, press **F10** to enter the Lifecycle Controller GUI.

Figure D-1: Dell Lifecycle Controller



3. Select **Hardware Configuration**. Click on **Repurpose or Retire System**.

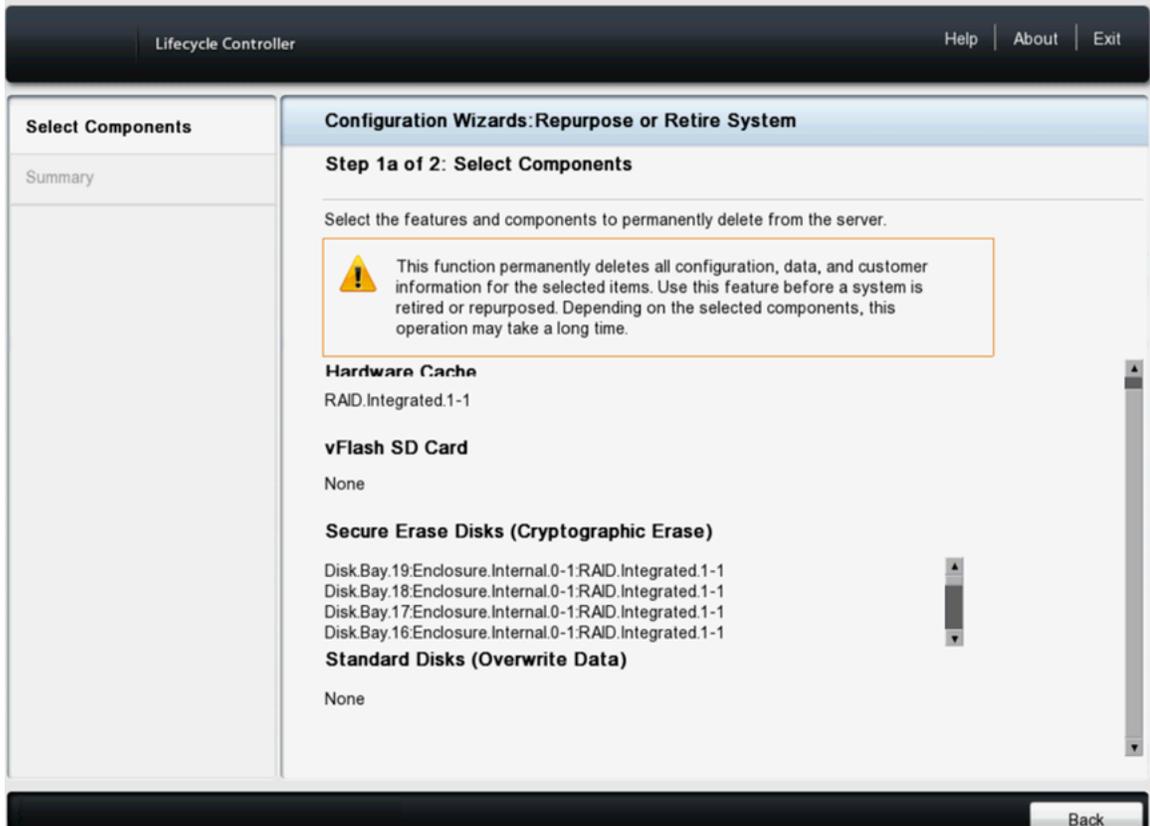
Figure D-2: Dell Lifecycle Controller



The Retire or Repurpose System function enables the removal of data from the server by erasing server non-volatile stores and data stored on Hard Disk Drives (HDDs), Self-Encrypting Drive (SED), Instant Secure Erase (ISE), and Non-Volatile Memory drives (NVMe).

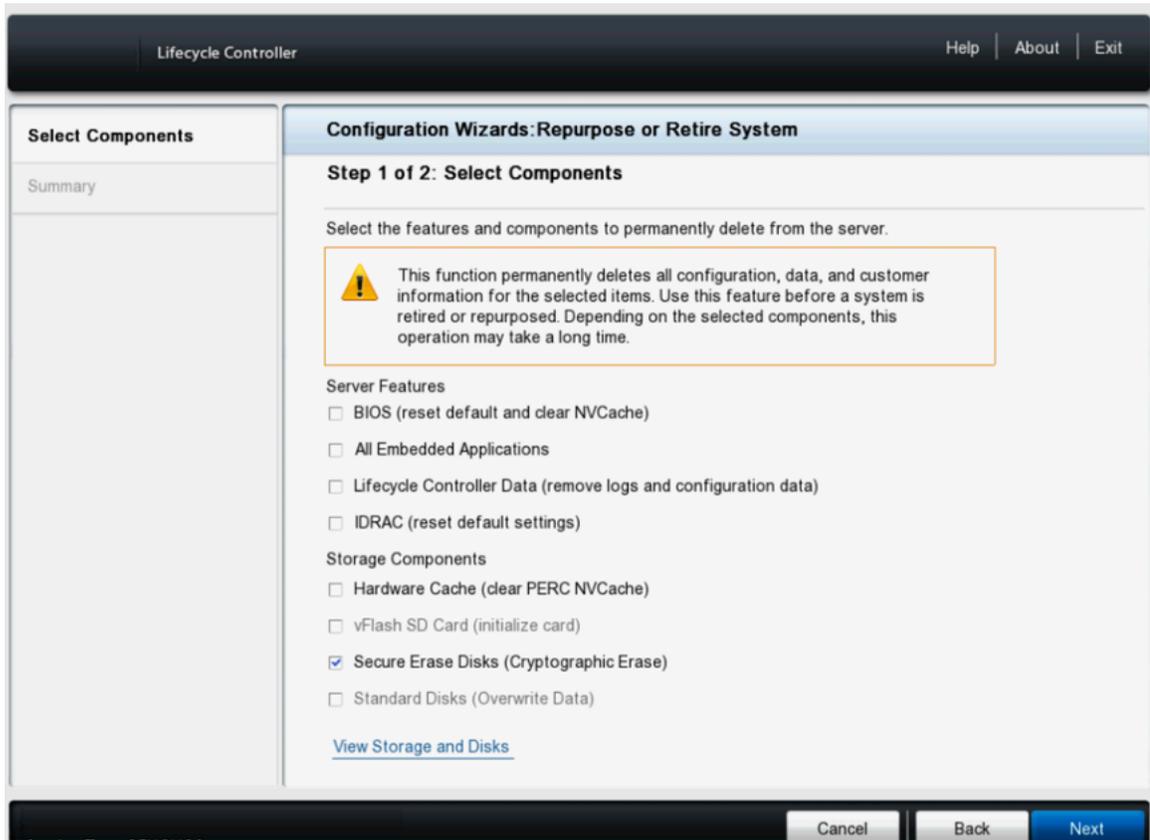
4. Click on **View Storage and Disks** to display all the drives attached to the server supported for erasure. Only drives that can be erased and detected are displayed.

Figure D-3: Dell Lifecycle Controller Hardware Configuration



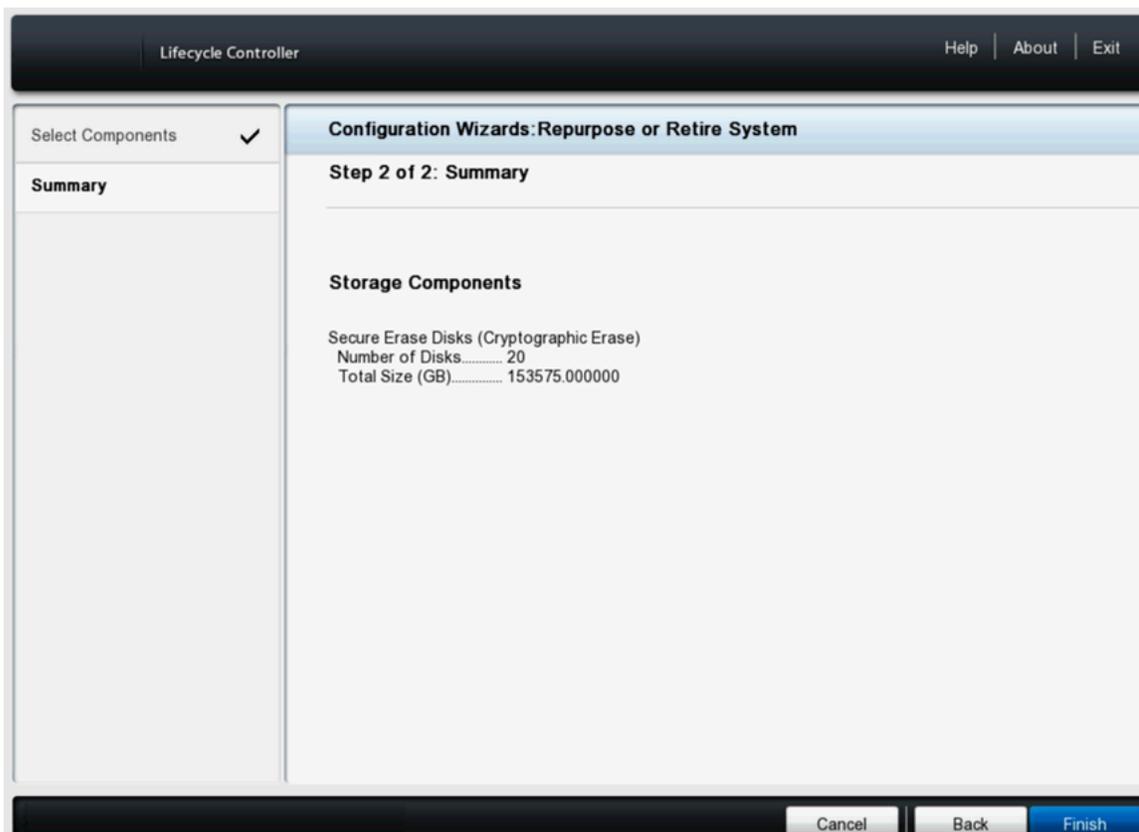
5. Click **Back** and select to return to **Step 1** (or **1a**). **Select Secure Erase Disk (Cryptographic Erase)**. Select **Secure Erase Disk (Cryptographic Erase)** and **Standard Disks (Overwrite Data)** if the system detects both types.

Figure D-4: Dell Lifecycle Controller > Repurpose/Retire a system



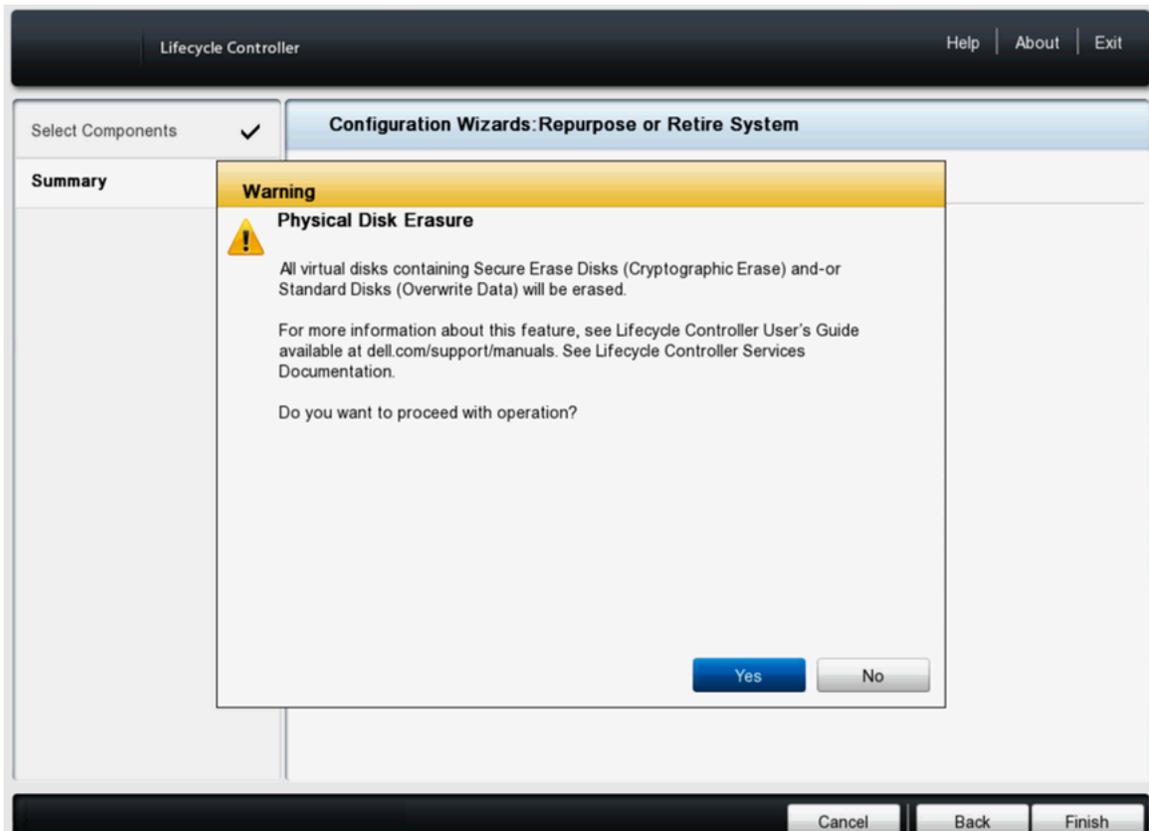
- Click **Next** to view the **Step 2** Summary page, which displays the drives that will be erased. Click **Finish**.

Figure D-5: Dell Lifecycle Controller Repurpose/Retire a system



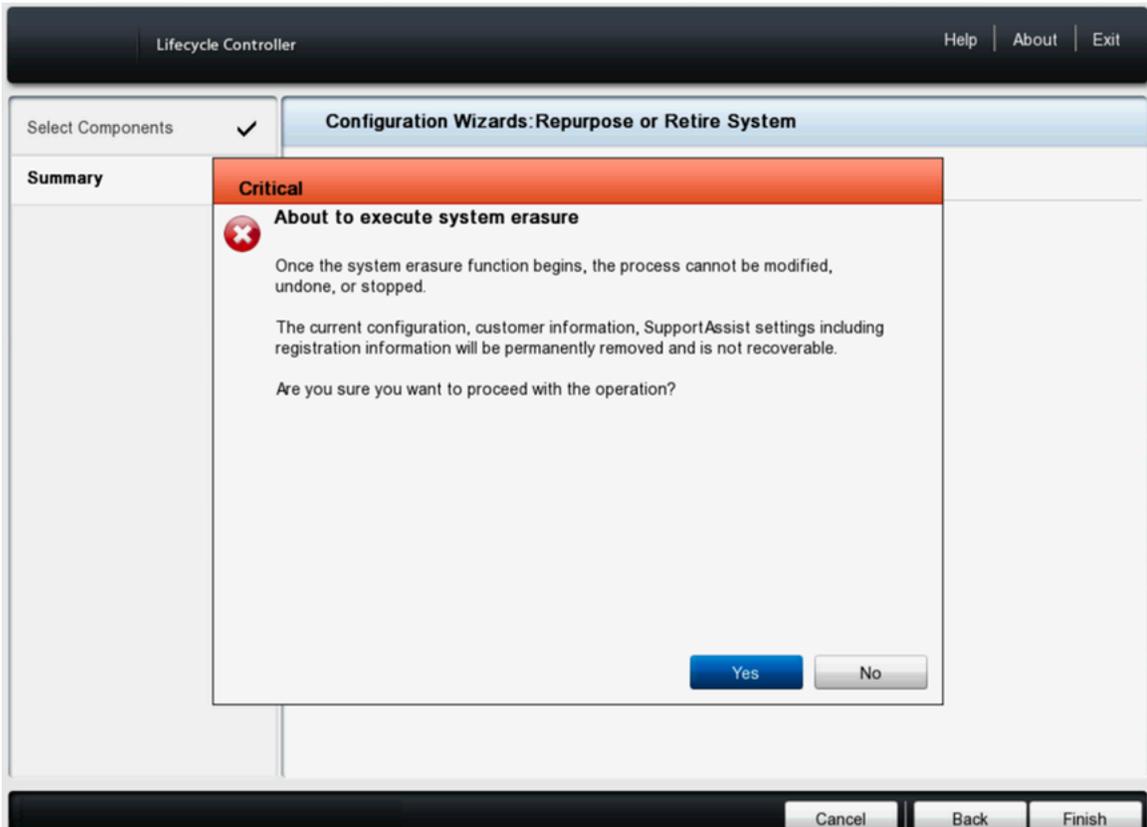
- Warning Physical Disk Erasure message appears about the erasure of the disks. Click **Yes**.

Figure D-6: Dell Lifecycle Controller Repurpose/Retire a system



8. A **Critical** message indicates that the erasure process cannot be stopped once started. Click **Yes**.

Figure D-7: Dell Lifecycle Controller Repurpose/Retire a system



-
9. The DMF appliance will reboot, and during POST, the display will show the Automated Task Application to erase the disks.

Figure D-8: DMF Appliance



10. The Automated Task Application dialog box displays the task of erasing drives along with a progress bar. After the task finishes, the DMF Recorder Node turns off.



Note: Depending on the amount of data on the DMF appliance, the disk erasure process can take some time.

Figure D-9: Dell Lifecycle Controller Repurpose/Retire a system

The screenshot displays the 'Automated Task Application' interface. The main title is 'Erase Cryptographic Disks (JID_046622862127)'. The current status is 'Task in Progress'. The task time limit is 1918:54:00, and the elapsed time is 00:42. The task progress is shown as 1 of 3 tasks completed, with a total elapsed time of 00:00:42. A message box indicates 'Tasks are running normally.' and a warning box states 'Do not restart, press CTRL+ALT+DEL, or turn off the server. The system will restart automatically when complete.'

Property	Value
Current Status	Task in Progress
Task Time Limit	1918:54:00
Elapsed Time	00:42
Task	1 of 3
Total Elapsed Time	00:00:42

Tasks are running normally.

Do not restart, press CTRL+ALT+DEL, or turn off the server. The system will restart automatically when complete.

Reforming Controller HA Cluster

Removing default IPv4 or IPv6 permit entry before adding a specific permit rule in the sync access list will permanently break communication between active and standby Controllers. Follow the procedure outlined in this appendix to recover from a split Controller HA cluster.

E.1 Controller Cluster Recovery

Controller-1 (IP: **192.168.55.11**, node-id:**23955**) is active and **Controller-2** (IP:**192.168.39.44**, node-id:**1618**) is standby. Retrieve the Node-id using the `show controller details` command in the CLI.

```
DMF-CTL2(config-controller-access-list)# show controller details
Cluster Name : DMF-7050
Cluster UID : a5de38214971de42aa7b51b96ac7345f4f228b20
Cluster Virtual IP : 10.240.130.18
Redundancy Status : redundant
Redundancy Description : Cluster is Redundant
Last Role Change Time : 2022-11-05 00:56:04.862000 UTC
Cluster Uptime : 2 months, 1 week
# IP          Hostname @ Node Id Domain Id State   Status   Uptime
-----|-----|-----|-----|-----|-----|-----|
1 192.168.39.44 DMF-CTL2 * 1618   1       active  connected 2 weeks, 2 days
2 192.168.55.11 DMF-CTL1  23955  1       standby connected 2 weeks, 2 days
~~~~~ Failover History ~~~~~
# New Active Time completed Node Reason Description
-----|-----|-----|-----|-----|-----|-----|
1 22049      2022-11-05 00:55:35.994000 UTC 22049 cluster-config-change Changed connection
state: cluster configuration changed
```

Procedure

1. [**Controller-1**] Add the sync 2 permit from the **0.0.0.0/0** rule.
2. **Controller-2** remains a standby so the user cannot add default rule to access-list sync until it transitions to active.
3. [**Controller-2**] Run this command on **Controller-2**, `system reset-connection switch all` changing **Controller-2** to active.
4. [**Controller-2**] On **Controller-2**, add the default rule to access-list sync 2 permit from the **0.0.0.0/0** rule.
5. [**Controller-2**] On **controller-2**, use the debug bash then run the following command:

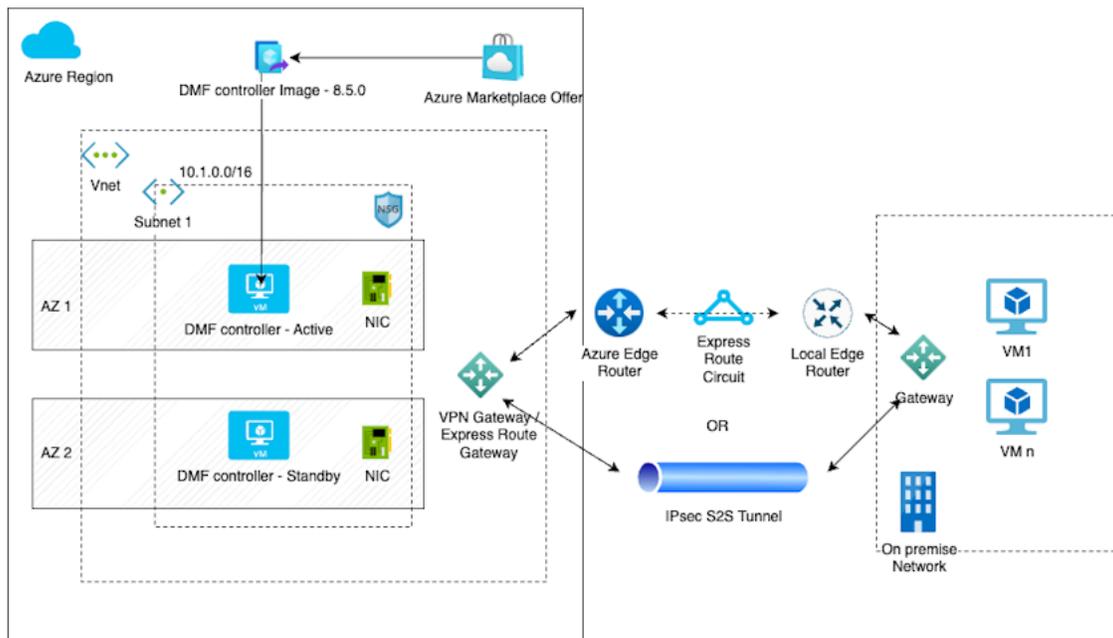
```
sudo bootstrapool -ks /etc/floodlight/auth_credentials.jceks --set 23955,192.168.55.11,
6642
Node id 23955 is the old active Controller-1 node-id and ip address is old active
Controller-1 ip address.
```

6. Wait for the cluster to reform.
7. **Controller-1** and **Controller-2** may change their role after this recovery procedure, that is **Controller-2** may become active.

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.

Figure F-1: Customer Azure Infrastructure



Customer Azure Infrastructure

The DMF Controller in Azure feature enables the registration of VM deployments in Azure and supports **auto-firstboot** using Azure **userData** or **customData**.

F.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 CDRROM.

`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 the DMF Controllers.

The following table lists details of the first boot parameters for the `auto-firstboot` configuration.

Firstboot Parameters - Required Parameters

Key	Description	Valid Values
<code>admin_password</code>	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
<code>recovery_password</code>	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 style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>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.</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 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 host names or FQDN strings	
ntp_servers	The NTP servers for the cluster to use.	no	List of the host names of FQDN strings <code>0.bigswitch.pool.ntp.org</code> <code>1.bigswitch.pool.ntp.org</code> <code>2.bigswitch.pool.ntp.org</code> <code>3.bigswitch.pool.ntp.org</code>	

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"
  ]
}
```

F.2 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.

F.3 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 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>

F.4 Syslog Messages

- Azure DMF Controller VMs can be accessed via an `ssh` login.
- `systemctl` should be in a **running** state without any failed units for the Controllers to be in a healthy state as shown in the following example:

```
dmf-controller-0-vm> debug bash
admin@dmf-controller-0-vm:~$ sudo systemctl status
dmf-controller-0-vm
State: running
Jobs: 0 queued
Failed: 0 units
```

F.5 Troubleshooting

- 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.
- Azure DMF Controller VMs can be accessed via `ssh`. Firstboot logs can be accessed using `debug bash` as shown below:

```
dmf-controller-0-vm> debug bash
admin@dmf-controller-0-vm:~$ less /var/log/floodlight/firstboot/firstboot.log
```

- For debugging parameter validation errors, access the parameter validation results using `debug bash` as shown below:

```
dmf-controller-0-vm> debug bash
admin@dmf-controller-0-vm:~$ less /var/lib/floodlight/firstboot/validation-results.json
```

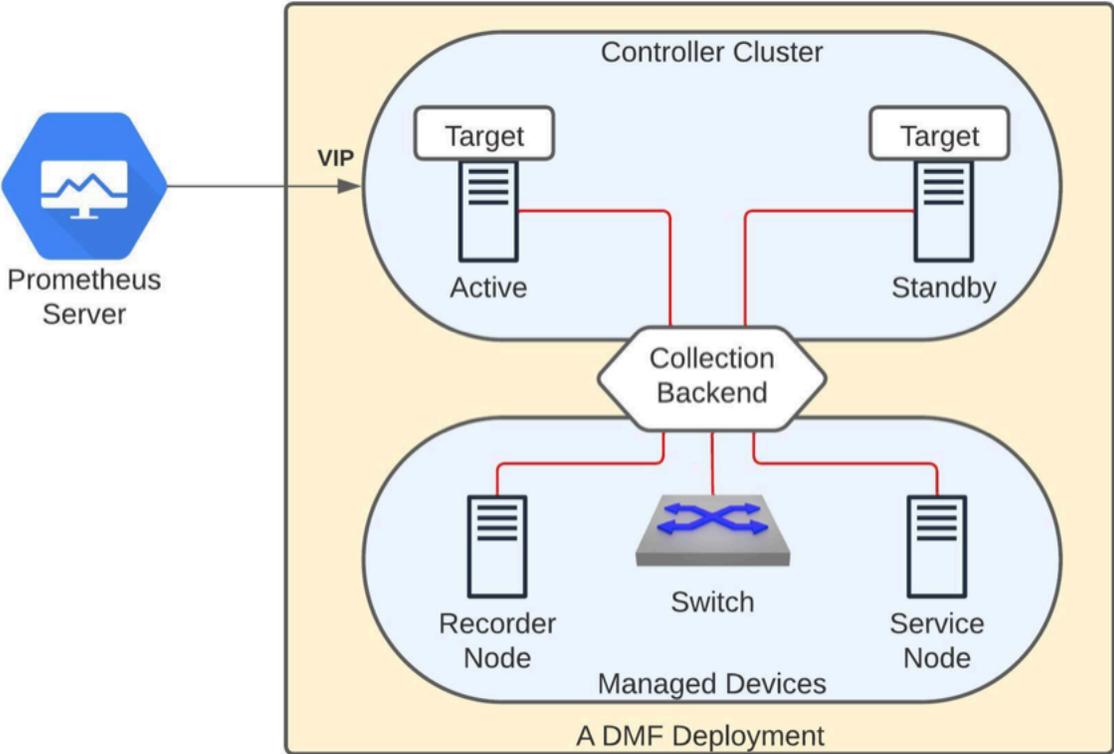
Prometheus Endpoint Support for Infrastructure Metrics

Prometheus is an open source monitoring and alerting toolkit. It collects and stores metrics from different sources in a time-series database. Prometheus offers a powerful query language, which allows users to analyze and visualize the collected data in real-time. With its robust alerting system, Prometheus can also notify users of potential issues which helps with their timely resolution.

Starting with the DMF 8.5.0 release, the Prometheus server can scrape metrics from a DMF (DANZ Monitoring Fabric) deployment. The DMF Controller exposes interface counters, CPU usage, memory usage, sensor states, and disk usage statistics from all the devices including the Controllers from a single Prometheus endpoint.

Deployment

Figure G-1: DMF Deployment with Prometheus Server



The aforementioned diagram shows a DMF deployment with an active/standby Controller cluster. In this environment, each Controller collects metrics from all the devices it manages as well as from both the Controller nodes. It then exposes them via the telemetry endpoint `/api/v1/metrics/prometheus`. This is an authenticated endpoint that listens on port 8443 and supports both Prometheus and OpenMetrics exposition formats.

Even though each Controller is capable of serving the telemetry information, it is recommended to use cluster's virtual IP (VIP) in the Prometheus configuration in order to achieve seamless continuity in the event of a Controller failover.

Configuration

No additional configuration is necessary on the DMF Controller to enable the metric collection. However, to allow the Prometheus service to access the new HTTP telemetry endpoint, a user access-token needs to be created. An admin user can choose an existing user or create a dedicated user with correct privileges and generate a token in order for Prometheus to fetch metrics from a fabric. The following sections describe the necessary configurations.

Permission

The group that the user belongs to needs to have sufficient permission to query the Prometheus endpoint. The following table summarizes the behavior.

Group	Behavior
admin	An access-token generated for a user in the <code>admin</code> group will have access to the telemetry endpoint.
read-only	An access-token for a user in the <code>read-only</code> group will not have access to the telemetry endpoint.
Any custom group	An access-token for a user in a group with <code>TELEMETRY</code> or <code>DEFAULT</code> permission but not with <code>DEFAULT/SENSITIVE</code> permission will have access.

To set telemetry permission for a custom 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
```

Access Token

Generate an access token for the user using the following command.

```
dmf-controller(config)# user username
dmf-controller(config-user)# access-token descriptive name for the access-token
access-token : ZxyHXL0QyOhDUogT8wjZj7ouSiVtWNB3
```

Prometheus Service

The following configuration needs to be added to the Prometheus server to fetch metrics from the Controller's `/api/v1/metrics/prometheus` endpoint periodically.

```
scrape_configs:
- job_name: <job_name>
  scheme: https
  authorization:
  type: Bearer
  credentials: <access-token>
  metrics_path: /api/v1/metrics/prometheus
  scrape_interval: <interval>
  static_configs:
  - targets:
  - <vip>:8443
  tls_config:
  insecure_skip_verify: true
```

The table below depicts the recommended configurations for this feature.

Configuration	Value
Credential	The corresponding access token created on the Controller
Scrape Interval	The minimum supported interval is 10s
Target	Use the VIP of the DMF Controller cluster
TLS	If a self-signed certificate is used on the Controller, add <code>insecure_skip_verify: true</code>

Please refer to the configuration guidelines of the specific Prometheus version you are using in the production.

Limitations

- Every device does not support every metric. Check the [Metrics Summary](#) section for more details.
- 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 have any effect on the counter values reported by the telemetry. The value is monotonically increasing and corresponds to the total count since the device was last powered up. This value only gets reset when the device is rebooted.

Notes

- The configured name of a managed device (e.g., switch, recorder node etc.) on the Controller is used as the value of the `device_name` label for all the metrics corresponding to it. In the case of a Controller, the configured hostname is used in the `device_name` label. Thus, these names are expected to be unique in a specific DMF deployment.
- It is possible that no metrics are collected from a device for a short time period. This may happen when the device is rebooting or when the Controllers experience a failover event.
- Prometheus will add additional metrics, e.g., `scrape_duration_seconds`, `scrape_samples_post_metric_relabeling`, `scrape_samples_scraped`, and `scrape_series_added`. These metrics do not collect any data from the DMF fabric and can be ignored.
- A metric representing a status that can take a fixed set of values is represented as a `StateSet` metric. Each possible state is reported as a separate metric. The current state is reported with value 1 and the other states with value 0. The state itself is reported as a label with the same name as that of the metric. For example, `device_psu_oper_status` will display multiple metrics for the operational status of a PSU (Power Supply Unit). If a PSU, `psu1` is in the failed state, then the metric `device_psu_oper_status{name="psu1", device_psu_oper_status="failed"}` will report value 1. At the same time, we will see the metric `device_psu_oper_status{name="psu1", device_psu_oper_status="ok"}` with value 0 for another state `ok`.
- Internally, all the metrics are fetched at 10 sec frequency except the ones associated with the sensors. These are currently collected at every minute.

Troubleshooting

If no metric is collected or no change in the metrics is visible on Prometheus for a few minutes, please follow the following troubleshooting steps:

- The Controller cluster is reachable over its VIP and there is an elected active Controller.

- You can retrieve the telemetry states in Prometheus exposition format using the token you created. Use command `curl -k -H "Authorization: Bearer <token>" https://<vip>:8443/api/v1/metrics/prometheus`
- The telemetry connection to the devices is active by running `show telemetry connection` command. Check the *Telemetry Collector* section of the *DANZ Monitoring Fabric User Guide* for more details.

Resources

- [Prometheus - Monitoring system & time series database](#)
- [OpenMetrics Specification](#)

Appendix

Metrics Summary

The following table shows the details of each metric exposed by the DMF fabric. The supported metric column shows that if a metric is generally supported on the device type. However, some specific platforms or hardware might not report a specific metric.

Metric	Description	Supported device type				
		Ctrl	SWL	EOS	SN	RN
device_cgroup_cpu_percentage	The normalized CPU utilization percentage by a control group	Y	Y	N	Y	Y
device_cgroup_memory_bytes	The memory used by a control group in bytes	Y	Y	N	Y	Y
device_config_info	The informational metrics of a managed device	N	Y	Y	Y	Y
device_cpu_utilization_percentage	The percentage utilization of a CPU on a device of the fabric	Y	Y	Y	Y	Y
device_fan_oper_status	The current operational status of a fan	Y	Y	Y	Y	Y
device_fan_rpm	The current rate of rotation of a fan	Y	Y	Y	Y	Y
device_fan_speed_percentage	The percentage of the max rotation capacity that a fan is currently rotating	N	Y	N	N	N
device_interface_in_broadcast_packets_total	The total number of broadcast packets received on the interface of a device	Y	Y	Y	Y	Y
device_interface_in_discards_packets_total	The number of discarded inbound packets by the interface of a device even though no errors had been detected	Y	Y	Y	Y	Y
device_interface_in_errors_packets_total	The number of inbound packets discarded at the interface of a device for errors	Y	Y	Y	Y	Y
device_interface_in_fcs_errors_packets_total	The number of received packets on the interface of a device with erroneous frame check sequence (FCS)	Y	Y	Y	Y	Y
device_interface_in_multicast_packets_total	The total number of multicast packets transmitted from the interface of a device	Y	Y	Y	Y	Y
device_interface_in_octets_total	The total number of octets received on the interface of a device	Y	Y	Y	Y	Y
device_interface_in_packets_total	The total number of packets received on the interface of a device	Y	Y	Y	Y	Y
device_interface_in_unicast_packets_total	The total number of unicast packets received on the interface of a device	Y	Y	Y	Y	Y
device_interface_operational_status	The operational state of an interface of a device	Y	Y	Y	Y	Y
device_interface_out_broadcast_packets_total	The total number of broadcast packets transmitted from the interface of a device	Y	Y	Y	Y	Y

device_interface_out_discards_packets_total	The number of discarded outbound packets at the interface of a device even though no errors had been detected.	Y	Y	Y	Y	Y
device_interface_out_errors_packets_total	The the number of outbound packets that could not be transmitted by the interface of a device because of errors	Y	Y	Y	Y	Y
device_interface_out_multicast_packets_total	The total number of multicast packets transmitted from the interface of a device	Y	Y	Y	Y	Y
device_interface_out_octets_total	The total number of octets transmitted from the interface of a device	Y	Y	Y	Y	Y
device_interface_out_packets_total	The total number of packets transmitted from the interface of a device	Y	Y	Y	Y	Y
device_interface_out_unicast_packets_total	The total number of unicast packets transmitted from the interface of a device	Y	Y	Y	Y	Y
device_memory_available_bytes	The current available memory at a device of the fabric	Y	Y	Y	Y	Y
device_memory_total_bytes	The total memory of this device	Y	Y	N	Y	Y
device_memory_utilized_bytes	The current memory utilization of a device of the fabric	Y	Y	Y	Y	Y
device_mount_point_space_available_megabytes	The amount of unused space on the filesystem	Y	Y	N	Y	Y
device_mount_point_space_usage_percentage	The used space in percentage	Y	Y	N	Y	Y
device_mount_point_space_utilized_megabytes	The amount of space currently in use on the filesystem	Y	Y	N	Y	Y
device_mount_point_total_space_megabytes	The total size of the initialized filesystem.	Y	Y	N	Y	Y
device_psu_capacity_watts	The maximum power capacity of a power supply	N	N	Y	N	N
device_psu_input_current_amps	The input current drawn by a power supply	Y	Y	Y	Y	Y
device_psu_input_power_watts	The input power to a power supply	Y	Y	N	Y	Y
device_psu_input_voltage_volts	The input voltage to a power supply	Y	Y	Y	Y	Y
device_psu_oper_status	The current operational status of a power supply unit	Y	Y	Y	Y	Y
device_psu_output_current_amps	The output current supplied by a power supply	N	Y	Y	N	N
device_psu_output_power_watts	The output power of a power supply	N	Y	Y	N	N
device_psu_output_voltage_volts	The output voltage supplied by a power supply	N	Y	Y	N	N

device_thermal_oper_status	The current operational status of a thermal sensor	Y	Y	N	Y	Y
device_thermal_temperature_celsius	The current temperature reported by a thermal sensor	Y	Y	Y	Y	Y

* Ctrl = Controller, SWL = A switch running Switch Light

EOS = A switch running Arista EOS, SN = Service Node, RN = Recorder Node.

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 SNMP MIB Reference Guide*