



DANZ Monitoring Fabric Deployment Guide

Arista Networks

www.arista.com

DANZ Monitoring Fabric Deployment Guide

DOC-06656-01

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 2023 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos, and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks is subject to Arista Network 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.....	2
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.....	5
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.....	6
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.....	15
2.5 Managing the DMF Controller Cluster.....	16
2.5.1 Verifying Cluster Configuration.....	16
2.5.2 Configuring the Cluster Virtual IP.....	17
2.5.3 Setting the Time Zone.....	18
2.5.4 Viewing Controller and Cluster Status.....	19
2.5.5 Saving and Restoring Controller Configuration.....	20
2.6 Copying Files Between a Workstation and a DANZ Monitoring Fabric Controller.....	21
2.6.1 Snapshot File Management Using REST API.....	22
2.6.2 Convert a Text-based running-config into a Snapshot.....	23
2.7 Managing DMF Sessions.....	28
2.8 Managing and Viewing Logs.....	29
2.8.1 Sending Logs to a Remote Syslog Server.....	30
2.8.2 Using the GUI to Configure Remote Logging.....	30
2.8.3 Using the CLI to Configure Remote Logging.....	31
2.8.4 Viewing Log Files on the Controller.....	31
2.8.5 Administrative Activity Logs.....	31
2.9 REST API Logging.....	32
2.9.1 Restricting Size of REST API Body.....	33
2.10 Syslog Over TLS.....	34
2.10.1 Overview.....	34
2.10.2 Configuration.....	35
2.11 Creating Support Bundle.....	36
2.12 NIST 800-63b Password Compliance.....	38
2.12.1 Configuration.....	38
2.13 Custom Password Compliance.....	39
2.14 Switch Management Interfaces not Mirroring Controller Management Interface ACLs.....	40
2.14.1 Configuration using the CLI.....	41
2.14.2 CLI Show Commands.....	42

2.14.3	Limitations.....	43
2.15	Recovery Procedure.....	44
2.15.1	Recovery from a Single Controller Failure.....	44
2.15.2	Recovery from a Dual Controller Failure.....	44
Chapter 3: DANZ Monitoring Fabric Deployment Topologies.....		46
3.1	DANZ Monitoring Fabric Topologies.....	46
3.1.1	Single-Switch Topology.....	46
3.1.2	Two-Tier Topology.....	46
3.1.3	Three-Tier Any-Tap-to-Any-Tool Topology.....	47
Chapter 4: Installing DMF Switches.....		49
4.1	HTTPS Support for Controller Hosted URLs using ZTP.....	49
4.2	Zero Touch Fabric Provisioning Modes.....	49
4.3	Using L2 ZTF (Auto-Discovery) Provisioning Mode.....	50
4.3.1	Requirements.....	50
4.3.2	Switch Installation Procedure.....	51
4.3.3	Arista Switch Installation Procedure for 7050X Series and 7260X Series.....	54
4.3.4	Allocating IPv4 Addresses to Fabric Switches.....	57
4.3.5	Using the GUI to Allocate IPv4 Addresses.....	57
4.3.6	Using the CLI to Allocate IPv4 Addresses.....	57
4.4	Using L3 ZTN (Pre-Configured) Switch Provisioning Mode.....	58
4.4.1	Installing a Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	58
4.4.2	Installing Arista 7050X and 7260X Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	62
4.4.3	Installing Arista 7280R Series Switch Using L3 ZTF (Preconfigured) Provisioning Mode.....	65
4.4.4	Configuring the Switch Static IP and Controller IP in Interactive ZTF Mode.....	71
4.4.5	Installing Arista 7050X and 7260X Series using dhcp with bootfile-name option.....	72
4.4.6	Installing Arista 7280R Series using dhcp with bootfile-name option.....	73
4.4.7	Using DHCP with Default URL for Switch Installation in Preconfigured Provisioning Mode.....	75
4.5	Registering a Switch After Initial Deployment.....	76
4.5.1	Using the GUI to Register a Switch.....	76
4.5.2	Using the CLI to Register a Switch.....	77
4.6	Changing the ZTF Mode After Deployment.....	78
4.6.1	Changing to Layer 3 (Pre-Configured) Switch Provisioning Mode.....	78
4.6.2	Using the GUI to Change the Switch Provisioning Mode.....	78
4.6.3	Using the CLI to Change the Switch Provisioning Mode.....	79
4.6.4	Changing to Layer 3 ZTF (Preconfigured) Mode.....	79
4.6.5	Changing to Layer 2 ZTF (Auto-Discovery) Mode.....	80
4.6.6	System Reinstall for an EOS Switch.....	81
Chapter 5: Managing Switches and Interfaces.....		83
5.1	Configuring Link Aggregation.....	83
5.1.1	Using the GUI to Configure Link Aggregation Groups.....	83
5.1.2	Using the CLI to Configure Link Aggregation Groups.....	85
5.2	Connecting Directly to a Switch.....	86
5.2.1	Manually Configuring Enhanced Hashing for Load Distribution.....	86
5.2.2	Configuring Enhanced Hashing.....	87
5.2.3	Symmetric Load Balancing.....	87
5.2.4	GTP Hashing.....	87
5.3	Overriding the Default Switch Configuration.....	88

5.4 Configuring Switch Interfaces.....	89
5.4.1 Forward Error Correction.....	91
5.4.2 Autonegotiation.....	92
5.4.3 Manually Setting the Interface Speed.....	93
5.4.4 Using Breakout Cables.....	93
5.4.5 Verifying Switch Configuration.....	94
Chapter 6: DMF Upgrade Procedures.....	96
6.1 Upgrading the Controller.....	96
6.1.1 Upgrade Procedure Summary.....	96
6.1.2 Upgrade Options.....	97
6.2 Copying the ISO Image File to the Controller.....	97
6.3 Staging the Upgrade.....	98
6.4 Launching the Upgrade.....	99
6.5 DMF Switch Automatic Upgrade.....	99
6.6 Verifying the Upgrade.....	100
6.7 Verify Persistent IPAM Assigned IP Addresses after Upgrade.....	100
6.8 Rolling Back an Upgrade.....	102
Chapter 7: Installing and Upgrading the DMF Service Node.....	104
7.1 Overview.....	104
7.2 Connecting the Service Node.....	105
7.3 Service Node Setup and Initial Configuration.....	106
7.4 Creating Support Bundle on Service Node.....	109
7.5 Upgrading Service Node Software From Release 7.x.x.....	110
Chapter 8: Installing and Configuring the DMF Recorder Node.....	111
8.1 Overview.....	111
8.2 DMF Recorder Installation Procedure.....	113
8.3 Initial Configuration - GUI.....	117
8.4 Initial Configuration - CLI.....	118
8.5 Changing the Recorder Node Default Configuration.....	119
Chapter 9: DHCPv4 Based Firstboot for DMF Controller and Managed Appliances.....	121
9.1 Introduction.....	121
9.2 Steps to Prepare your Services (TFTP/NFS) for PAVE/REPAVE Operation.....	121
9.3 DHCPv4 based firstboot for DMF Controller and Managed Appliances.....	124
9.4 Assumption and Trust Model.....	128
Chapter 10: Managing SNMP.....	129
10.1 SNMP Overview.....	129
10.2 Using the DMF GUI to Configure SNMP.....	129
10.2.1 Configuring SNMP Traps.....	133
10.3 Using the CLI to Configure SNMP.....	135
10.3.1 Configuring SNMP Access to the Controller.....	135
10.3.2 Configuring SNMP Access to the Analytics Node.....	136
10.3.3 Identifying the SNMP Trap Receiver.....	136
10.3.4 Configuring SNMP Settings.....	137
10.3.5 Configuring SNMP Switch Trap Thresholds.....	137

10.3.6	SNMP Traps for DMF Service Node Appliance.....	138
10.3.7	Managing the SNMPv3 Engine ID for Trap Receivers.....	138
10.3.8	Configuring SNMPv3 Users.....	139
10.3.9	SNMPv3 Command Examples.....	140
10.4	Configuring SNMP on a Specific Switch.....	140
10.4.1	Using the GUI to Configure SNMP on a Specific Switch.....	140
10.4.2	Using the CLI to Configure SNMP on a Specific Switch.....	141
10.5	SNMP Clear Trap.....	144
10.6	SNMP Trap Generation for Packet Drops and Link Saturation.....	145
10.6.1	Using the CLI to Configure the SNMP Traps.....	149

Chapter 11: Using Authentication, Authorization, and Accounting..... 151

11.1	Overview.....	151
11.2	Using Local Groups and Users to Manage DMF Controller Access.....	151
11.2.1	Viewing Existing Groups and Users.....	152
11.2.2	Using the GUI to Manage User Accounts.....	152
11.2.3	Using the CLI to Manage Groups and User Accounts.....	153
11.2.4	Changing User Passwords.....	155
11.2.5	Password Reset.....	155
11.2.6	Managing Groups.....	156
11.2.7	Authentication with a User Token and REST API.....	157
11.3	Configuring AAA to Manage DMF Controller Access.....	157
11.3.1	Enabling Remote AAA Services.....	158
11.4	Time-based User Lockout.....	159
11.5	Using TACACS+ to Control Access to the DMF Controller.....	161
11.5.1	Using the GUI to Add a TACACS+ Server.....	162
11.5.2	Using the CLI to Enable Remote Authentication and Authorization on the DMF Fabric Controller.....	163
11.5.3	Using the CLI to Add a TACACS+ Server.....	163
11.6	Setting up a tac_plus Server.....	164
11.6.1	Using the Same Credentials for DMF and Other Devices.....	165
11.6.2	RBAC-Based Configuration for Non Default Group User.....	165
11.7	Using RADIUS for Managing Access to the DMF Controller.....	166
11.7.1	Using the GUI to Add a RADIUS Server.....	166
11.7.2	Using the CLI to Add a RADIUS Server.....	167
11.7.3	Setting up a freeradius Server.....	168
11.8	Custom admin role.....	169
11.8.1	Categories.....	169
11.8.2	PERMISSIONS & PRIVILEGE:.....	175
11.8.3	Category Identification:.....	175
11.8.4	Group managements:.....	175
11.8.5	Remote Users.....	176
11.8.6	Commonly used profiles.....	176
11.8.7	Commonly used Category-feature Matrix.....	179

Chapter 12: Management and Control Plane Security..... 180

12.1	Management Plane Security.....	180
12.2	Importing the Controller Private Key and Certificate.....	180
12.3	Using Certificates Signed by a CA for GUI Access to the Controller.....	181
12.3.1	Replacing the Certificate.....	183
12.4	Managing the Controller HTTP and SSH Ciphers, Protocols, and Data Integrity Algorithms.....	184
12.4.1	Configuring HTTP Ciphers.....	184
12.4.2	Configuring HTTP Protocols.....	184
12.4.3	Configuring SSH Ciphers.....	185

12.4.4	Configuring SSH Data Integrity Algorithms.....	185
12.4.5	Changes to Supported MACs/Ciphers/SSH Keys.....	185
12.5	Inherit MAC and Cipher Configuration.....	187
12.5.1	Using the CLI to Configure SSH and HTTPS.....	187
12.5.2	Verify the Cryptographic Configuration.....	188
12.5.3	Limitations.....	190
12.6	Protocol Access Required to the DMF Controller.....	190
12.6.1	Management Plane Access.....	190
12.6.2	Control Plane Access for DMF Controller.....	191
12.6.3	Protocol Access Required to the DMF Controller - Sync.....	193
12.6.4	Control Plane Access for DMF Switches.....	195
12.6.5	Control Plane Access for DMF Service Node.....	196
12.6.6	Control Plane Access for DMF Recorder Node.....	198
12.6.7	Control Plane Access for Analytics Node.....	199
12.7	Enabling Secure Control Plane.....	201
12.7.1	Certificate Requirements and Recommendations.....	202
12.7.2	Installing X.509 Digital Certificates on the Active and Standby Controllers.....	203
12.8	Using the GUI to Enable Secure Control Plane.....	203
12.9	Using the CLI to Enable Secure Control Plane.....	208
12.9.1	Verification and Troubleshooting.....	212
12.9.2	Installing X.509 Digital Certificates on the Switches.....	212
12.9.3	Securing the Control Channel Between the Controller and the Switches.....	213
12.9.4	Adding Switches After Enabling Secure Control Plane.....	214
12.9.5	Replacing a Controller Appliance After Control Plane Security is Enabled.....	215
12.10	Replacing Certificates of Controllers and Switches.....	215
12.10.1	Replacing to-be-expired certificates of controllers and switches.....	215
12.10.2	Replacing expired certificates of switches.....	216
12.10.3	Replacing expired certificates of controller and switches.....	216
 Chapter 13: Using iDRAC with a Dell R430, R630, or R730 Server.....		218
13.1	Setting Up and Configuring iDRAC.....	218
13.2	Using iDRAC to Install the DMF Controller or DMF Service Node Image.....	228
 Chapter 14: Using iDRAC with a Dell R440 or R740 Server.....		234
14.1	Setting Up and Configuring iDRAC.....	234
14.2	Using iDRAC to Install a DMF Controller, Analytics, or Recorder Software Image.....	241
 Chapter 15: Switch CPLD Upgrade Procedure.....		248
 Chapter 16: Switch ONIE Upgrade Procedure.....		249
 Appendix A: Removing existing OS from switch.....		250
A.1	Reverting from DMF (Switch Light OS) to EOS - 7050X3 and 7260CX3.....	250
A.2	Reverting from DMF (EOS) to EOS - 7280R.....	251
A.3	Removing the existing OS from a Switch.....	252
 Appendix B: Creating a USB Boot Image.....		255
B.1	Creating the USB Boot Drive with MacOS X.....	255
B.2	Creating the USB Boot Image with Linux.....	256
B.3	Creating a USB Boot Image Using Windows.....	256

Appendix C: Installing a Controller VM.....	262
C.1 General Requirements.....	262
C.2 Installing on VMware ESXi/vSphere.....	262
C.2.1 Prerequisites.....	262
C.2.2 VM Installation.....	262
C.2.3 vMotion support for Virtual Controller.....	263
C.3 Installing on Ubuntu KVM.....	263
C.3.1 Prerequisites.....	263
C.3.2 VM Installation.....	263
 Appendix D: Erasing DMF Appliance.....	 264
D.1 Using the Dell LifeCycle Controller.....	264
 Appendix E: Reforming Controller HA Cluster.....	 273
E.1 Controller Cluster Recovery.....	273
 Appendix F: References.....	 274
F.1 Related Documents.....	274

Getting Started

This chapter describes the software and hardware requirements and provides a summary of the steps for deploying the DANZ Monitoring Fabric (DMF) controller.

1.1 DMF Installation Prerequisites

Follow the prerequisite steps below:

Table 1: DMF Prerequisites

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

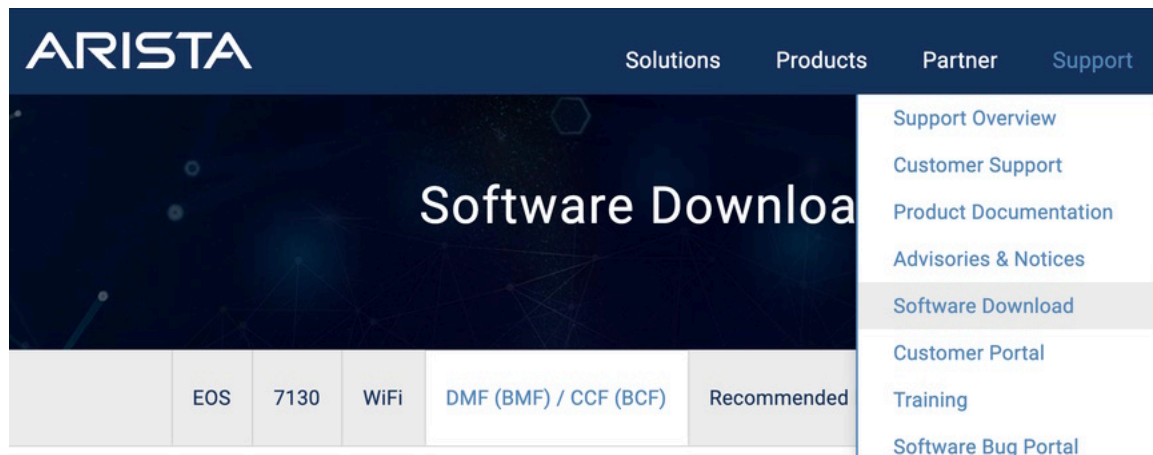
1.2 Downloading Software

The Arista Support Portal provides links to software packages and documentation for 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.

You can download documents of the appropriate release from the **DMF - BMF Product Documentation** folder tree. To download software, expand the **DMF - BMF Software Downloads** folder tree and you can find 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.

Task	Document
1. Identify new features and any software or upgrade issues for the current release.	<i>Release Notes</i>
2. Verify 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 and 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.

- **CLI Reference Guide** - Describes the command syntax and function of each command and keyword provided by the DMF Controller CLI.
- **Deployment Guide** - Procedures for installation, upgrade, and initial DMF controller configuration that is common to all DMF features and use cases.
- **Hardware Compatibility List** - Lists software and hardware that has been 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 List of REST API calls.
- **SNMP MIB 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 controller:
 - Controller Software.
 - Controller Hardware Appliance or KVM/ESX Virtual Machine (For VM requirements, refer to [Installing a Controller VM](#)).
- Switch Requirements
 - For the supported switch hardware refer to the **DANZ Monitoring Fabric 8.4 Hardware Compatibility List**.
 - 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 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 DMF controllers.	Complete the First boot script. <code>show controller</code>	Controller View	Installing and Configuring the DMF Controller
Monitoring Fabric Switches <ul style="list-style-type: none"> Name and MAC address. 	<code>switch switch-1 mac <mac-address></code>	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. 	<code>service-node <name>mac <service-node-eth0-mac></code> Complete the First boot script for initial configuration. <code>managed-service netflow 1 netflow <L3-Delivery interface> service-interface switchapp-ft-as6700-1 port-channel1</code>	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. 	<code>packet-recorder device <recorder-name>mac <mac-address></code> <code>packet-recorder device <name> packet-recorder interface <name></code> <code>packet-recorder-interface switch <switch> <interface use=packet-recorder <interface></code>	Monitoring > Packet Recorders 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. 	<code>switch switch-1 interface ethernet1role filter interface-name span-tap-1</code> <code>switch switch-1 interface ethernet2 role delivery interface-name tool-1</code> <code>policy p1 10 match tcp action forward</code> <code>filter-interface span-tap-1delivery-interface tool-1 tunneling</code>	Monitoring > Interfaces Monitoring > Policies	<i>DANZ Monitoring Fabric 8.4 User Guide</i>
Arista Analytics	Complete the First boot script and perform initial configuration	Settings > Analytics Configuration	<i>Arista Analytics 8.4 User Guide</i>

Installing and Configuring the DMF Controller

This chapter describes the basic procedure for installing the DANZ Monitoring Fabric controller software.

2.1 Connecting to the Controller

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

The initial setup of the controller can be completed in any of the following ways:

- Attach a monitor and keyboard to the appliance console port directly.
- Use a terminal server to connect to the appliance console port and telnet to the terminal server.
- If the controller is deployed as a VM, connect to the VM console.



Note: It is recommended to have an iDRAC connection to the Controller, Service Node, Analytics Node and Recorder Node appliances. This 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 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 at 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. Make sure that the port baud rate of **115200** is in effect 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 telnet or SSH to the hardware appliance serial console through the terminal server.

2.2 Configuring the Active Controller Using the First Boot Script

Once you connect 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 DMF controller:

- IP address for the active and standby controller
- 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 prior to **6.1** was **9600**. If you upgrade an existing controller appliance having a serial interface set to **9600** baud, you will need to change the terminal setting to **115200** after performing an upgrade.

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

1. Insert the bootable USB drive into the controller appliance USB port.

Refer to Appendix B: Creating a USB Boot Image to make a bootable USB drive.



Note: When you power on the hardware appliance for the Active controller, you are prompted to press **Enter** to begin the installation.

2. 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
-----
```

3. To accept the agreement, type the command:**Yes**.

Refer to Appendix B: Creating a USB Boot Image to make a bootable USB drive.



Note: If you type the command: **No**, the system prompts you to Log out. You must accept the EULA to install the product.

After you type 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 >
```

4. 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
```

5. Choose the IP forwarding mode when prompted.

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

6. Choose the method for assigning IPv4 addresses.

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



Note: In case the user decides to select **[2] Automatic via DHCP**, it is recommended that the DHCP address be reserved for the MAC address of the Controller's management port.

7. 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.

In case the user selects **[1]** in **Step 6** above, the system will then prompts the user to enter the IPv4 address and related configuration 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.

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



Note: If you enter 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 if you want 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
```

9. Type **1** when you are configuring the Active controller.

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

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

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

The system then prompts you for the name of each NTP server to use for setting the system time, which is 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.

11. Enter the IP address or fully-qualified domain name of the NTP server for your network.

You can complete the initial configuration without specifying the NTP server, but a functioning NTP configuration is required for deploying the DMF controller in a production environment.

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] >
```

12. To apply the configuration, type 1.



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

After you enter the option to apply the settings, the system applies your 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 you set for the cluster during installation. You can also login to the Active controller using SSH with the IP address that you assigned, 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


For operational resilience, it is recommended to deploy the DANZ Monitoring Fabric in a two-node cluster. First, set up the first (active) cluster node as described above. Then, 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 subnet, which changes how controllers communicate with each other. Active and standby controllers communicate using controller management **IPv4** address. This enables provisioning 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 controller should be less than 300ms.

When you power on the hardware appliance or VM with the pre-installed DMF software, you are prompted to login as `admin` for the first-time configuration.

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

 **Note:** In addition, ports on controller-facing management-switch port must come up within 5 seconds.

When you power on the hardware appliance for the Active controller, you are prompted 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.

 **Note:** *DANZ Monitoring Fabric 8.4.0 (dmf-8.4.0 #11)*

Log in as `'admin'` to configure

```
controller login: admin
```

Use the default account (`admin`) without a password when the system is booting from the factory default image.

Use the default account (**admin**) without a password when the system is booting from the factory default image. The system displays the following prompt, which asks you 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: If you type **No**, the system prompts you to Log out. You must accept the EULA to install the product.

After you type **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**.

After you type **Yes** to accept the EULA, the system prompts for a password to be used for emergency recovery of the controller node, as shown below:

```
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: In case the user decides to select [2] **Automatic via DHCP**, it is recommended that the DHCP address be reserved 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) In case the user had selected [1] **Manual** in Step 7, the system will now prompt the user to enter the IPv4 address and other configuration for the Standby controller node.

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



Note: If the user enters 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 will not prompt 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 will now prompt the user for the IPv4 address of the Active controller for the cluster the standby will join.

```
Existing Controller IP > 10.106.8.2/24
```

Next the user is prompted to enter the IP address of the Active controller for the cluster.

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

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

13. The system displays the First Boot Menu. All the details provided by the user, thus far, will be displayed for a final verification. The user will now need to press **1** to apply the settings and start the process of the Standby Controller to join 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 user will next be advised that the current DMF cluster does not have **secure control** enabled. This is an optional feature that can be enabled later. At this stage, it is recommended to press **1** to and resume the process of 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: You can connect to the DMF controller at the IP address assigned to either the Active or Standby controller, but 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 a existing Standby controller to a different subnet, standby controller management IPv4 address has to be changed before moving the appliance to a different subnet. Before making any changes make sure that all the underlying network configuration is setup properly for connectivity.

1. Changing the management IP address on the standby controller. There are 2 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 controller 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: Using REST API to replace controller management ip address instead of removing and adding new management ip address. If this method is used than there is no need to use controller IDRAC or console access to change the management ip address. Login as admin user and use the below REST API to replace controller management ip address from bash.

```
Controller-2# deb bash
admin@DMF-Controller-2:~$ curl -g -H "Cookie: session_cookie=$FL_SESSION_COOKIE" 'http://localhost:8080/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 a existing standby controller to a different controller cluster.

1. Remove the standby controller from 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 as specified in section [Joining Standby Controller to Existing Cluster](#) ** if joining a 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 how to connect to the DMF controller.

To access the Active DMF controller, use the IP address of the Active controller. You can also use the Virtual IP (VIP) of the cluster, after one has been configured, as described in the [Configuring the Cluster Virtual IP](#) section.

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

2.4.1 Using the DANZ Monitoring Fabric CLI

Once the DMF controllers are up and running, log in to the DMF controller using the VM local console or by using SSH.

CLI commands are divided 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 have a more significant effect on system configuration and that can only be entered after entering the configure command. You also access sub-modes from this mode.

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

When you login, 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 on-line help for the command, enter the **help** command followed by the command.

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

You can pipe commands through other commands like **grep** for analysis and troubleshooting DANZ Monitoring Fabric operations. You can see 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
```

You can 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 GUI lets you perform similar operations to the CLI using a graphic user interface instead of text commands and options. You can use the DANZ Monitoring Fabric GUI with recent versions of any of the following supported browsers:

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

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

Figure 2-1: Accessing the DANZ Monitoring Fabric GUI



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

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

After accepting the prompts, the system displays the login prompt.

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

The main menu for the DANZ Monitoring Fabric GUI is displayed as in the following figure.

Figure 2-2: DANZ Monitoring Fabric GUI Main Menu



When you log in to the DANZ Monitoring Fabric GUI you see a landing page, the Controller Dashboard, and a menu bar at the top with sub-menus containing options for setting up DANZ Monitoring Fabric and for monitoring network activity. The menu bar includes the following sub-menus:

- **Fabric:** Manage DANZ Monitoring Fabric switches and interfaces.
- **DANZ Tap:** Manage DANZ Tap policies, services, and interfaces.
- **Maintenance:** Configure fabric wide settings (clock, SNMP, AAA, sFlow, Logging, Analytics Configuration).
- **Integration:** Manage integration of vCenter or NSX-V instances to allow monitoring traffic using DMF.
- **Security:** Manage administrative access.
- **Profile:** Display or change user preferences, change the password, or sign out.

2.5 Managing the DMF Controller Cluster

This section describes how to configure settings that apply to a DANZ Monitoring Fabric cluster or to both Active and Standby controllers. Most configuration is completed on the Active controller only, and the configuration is synchronized with the Standby controller.

2.5.1 Verifying Cluster Configuration

A DMF Out-of-Band HA cluster consists of two controller nodes, master and one slave. Keep the following conditions in mind when configuring the cluster:

- Both master and slave must be in the same IP subnet.

- Firewall configurations are separate for master and slave, so you must manually keep the configuration consistent between the two nodes.
- NTP service is required to establish a cluster. Starting with **DMF 7.0.0**, the master controller provides the NTP service for the cluster and connected switches.
- 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)#

```

- To display the current controller roles from either the Master or Slave node, enter the **show controller** command, as in the below 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 that lets you connect 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. Active and standby controller have to be in 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 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 > Clock** from the main menu.

Figure 2-3: Maintenance Clock



This displays the time on the system you are using to access the DMF GUI and the controller, information about the currently configured NTP server, and provides an option for forcing immediate NTP synchronization with the NTP server, which in **DMF 8.0.0** and later, runs on the DMF Master controller.

CLI Procedure

To set the time zone for the controller, from config-controller submode, use the **set timezone** command, which has the following syntax:

```
[no] clock timezone <time-zone>
```

Replace **time-zone** with the specific string for the time zone you want to use. To see a list of supported values, press **Tab** after the **clock set** command. You will notice that certain values, such as **America/**, are followed by a slash (/). These values must be followed by a keyword for the specific time zone. The supported time zones can be seen 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: Starting in *DANZ Monitoring Fabric Release 8.1.0*, 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 clock timezone
```

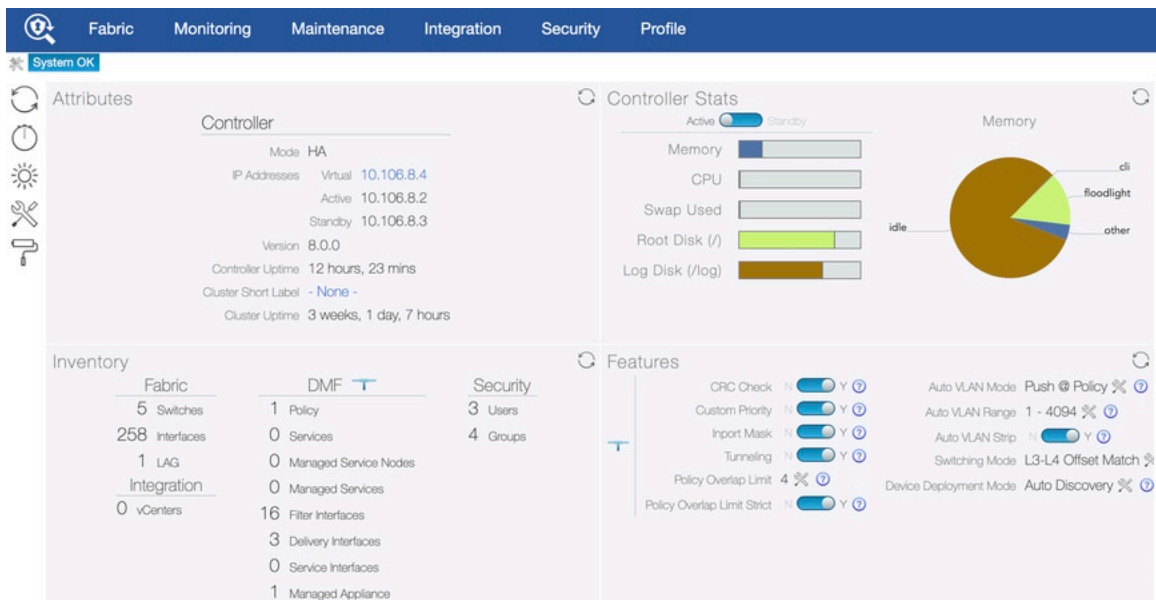


Warning: If you manually set the clock for the DMF controller using the clock set command, this may affect database reconciliation between the nodes in a controller cluster. Because time stamps are used to identify records that should be updated, it is recommended to adjust any time skew using the ntpdate 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 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 2-4: Cluster and Controller Status



The number of warnings and errors are listed and highlighted in yellow in the upper left corner of the page. This page also provides configuration options that let you view and change fabric-wide settings.

CLI Procedure

View the overall controller status by entering the **show controller** command, as in the following example:

```
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
----- Failover History -----
# 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-control**: 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

If the controller is set up properly, no regular maintenance is required. However, it is recommended to periodically save a copy of the controller running-config to a location outside the controller. You can do this using the copy command from the controller to an FTP server, as in the following example.

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

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

To restore controller configuration from a remote server, first 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 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:
```

4. Run first-time setup again to reconfigure the controller.
5. 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
```

6. Restore the cluster.

Remove the standby controller; then boot factory default on the Active node to start a new cluster.



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

2.6 Copying Files Between a Workstation and a DANZ Monitoring Fabric Controller

You can use the `SCP` command followed by the keywords shown below on a workstation connected to the 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 <filename> admin@<controller-ip>://<keyword>
```

For example, the following command copies the controller ISO image file from a workstation to the image file partition on the controller, from where it can be installed on the controller.

```
scp 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**.

You can use any user account belonging to the admin group. Replace the **admin** username in the above example 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. You can use **PSCP** on a Windows workstation to copy the file to the controller, using the following command:

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

You can use **SCP** 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 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 Snapshot File Management Using REST API

Starting in the **DMF 7.2 release**, snapshots can be managed using REST API. The API actions and their CLI equivalents are as follows:

- **Take** : copy running-config snapshot://my-snapshot
- **Fetch** : copy http[s]://snapshot-location snapshot://my-snapshot
- **Read** : show snapshot my-snapshot
- **Apply** : copy snapshot://my-snapshot running-config
- **List** : show snapshot
- **Delete** : delete snapshot my-snapshot

Take:

Below example shows how to save the current running-config to **snapshot://<file name>**.

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

Fetch:

Retrieve a snapshot found at “**source-url**” and save it locally as “**testsnap**”.

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

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 but the API will give 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"]'
```

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

Introduced in DMF 8.4, a new 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

DMF-8.4 adds a choice to 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 configuration 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 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 is updated immediately (for example, the hostname).

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
```

The transaction can be deleted using the following command:

```
# delete transaction
```



Note: While the system configuration is not updated, the snapshot is 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_SK1aj
niQ,rounds=25000,ph=true,0vtPyup3h5JThGFLff-1zw-42-BV7tG7Sm99RÖT1OmZCZj1zcWLJj
9Lc28mxkQI1-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 dmf204-sn
mac e4:43:4b:48:58:ac
! switch
switch gt160
mac c4:ca:2b:47:97:bf
admin hashed-password $6$ppXOyA92$0hibVW63R0t1T3f7NRUFxPEWUb4b6414dTGEayrr
Xcw5or/ZDxm/ZNvotQ7AQfVMo7OZ1I.yDLwrnlVXrZkV3.
!
interface Ethernet1
speed 10G
!
interface Ethernet5
speed 25G
switch hs160
mac c4:ca:2b:b7:44:83

```

```

admin hashed-password $6$McvJd94$vrXDNkr2OSz3kiZSYPCfuIbcaBuIcoC7ywlVeFF
d7oAgLnleVcV6NyEFZnykje4ILUjmJPWdWeu3LaF4sWzd/
!
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/RmErwpDXvJE20UnEXK
oLQffodjsyIlnZ1nG54X5Cq5qgb6uTGXs1TMYkqBWurLh1
!
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-sha1
! 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)# 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

To view and configure settings for remote sessions established to the DANZ Monitoring Fabric controller, switches and managed appliances complete the following steps:

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

Figure 2-5: Security Sessions

<input type="checkbox"/>	This Session	ID	Username ▲	Full Name	Groups	Last IP Address	Created	Last Used
<input checked="" type="checkbox"/>		7a1d7d40...	admin	Default admin	admin	10.95.66.14	Today, 8:56:18pm Pacific Standard Time	Today, 10:39:13pm Pacific Standard Time
<input type="checkbox"/>		b797ae98...	admin	Default admin	admin	10.106.8.10	Today, 12:22:40pm Pacific Standard Time	Today, 12:22:40pm Pacific Standard Time
<input type="checkbox"/>		bfa39bea...	admin	Default admin	admin	10.95.64.210	Today, 10:20:07am Pacific Standard Time	Today, 10:38:38pm Pacific Standard Time
<input type="checkbox"/>		c49ef680...	admin	Default admin	admin	10.95.66.14	Today, 9:41:50pm Pacific Standard Time	Today, 9:42:10pm Pacific Standard Time
<input type="checkbox"/>		ca40e639...	admin	Default admin	admin	127.0.0.1	Today, 10:11:27am Pacific Standard Time	Today, 10:11:27am Pacific Standard Time
<input type="checkbox"/>		daeb9538...	admin	Default admin	admin	127.0.0.1	Today, 10:16:44am Pacific Standard Time	Today, 10:16:44am Pacific Standard Time

Nov 20, 2020, 6:39:13am UTC Show: 10 25 100 All (1 - 6 / 6)

This page displays the sessions currently established. To forcibly end a session, select **Delete** from the menu control for the session.

- To configure the default settings for remote connections to the controller, click **Settings** control at the top of the page.

Figure 2-6: Security Sessions Configurations

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

- 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 causes the sixth session connection attempt by the same user account to fail. If more than five sessions are already open, the excess oldest sessions are closed.

This limit applies to sessions established through the GUI, CLI, or REST API, whether directed to the DANZ Monitoring Fabric fabric switches, managed appliances, Active or Standby controller and/or to the cluster virtual IP address.



Note: 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.

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 aaa session expiration to **10** minutes:

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

2.8 Managing and Viewing Logs

By default, the DANZ Monitoring Fabric fabric switches send syslog messages to both the Active and Standby controllers. Syslog messages are disk persistent and are removed only based on time and rotation policy.

When you configure an external syslog server, the controllers send the syslog messages to the external server and keep a local copy of the syslogs on the controller. If multiple external syslog servers are configured, the syslog messages are sent 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 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.

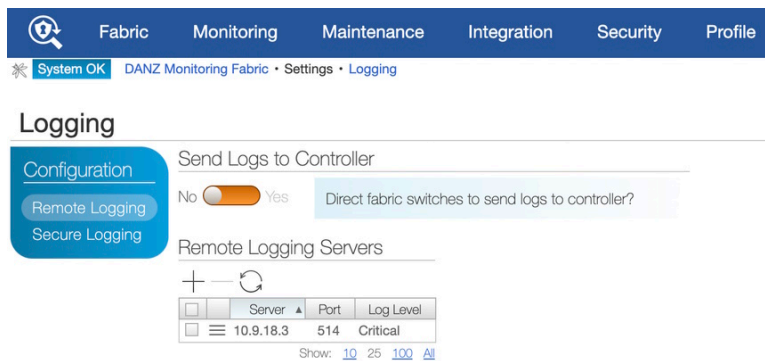
The controllers do not share their own syslog with the other controller in the cluster. In other words, the Active controller does not send its syslogs to the Standby controller, and the Standby controller does not share its syslogs with the Active controller. You access the individual controller logs from either the Active and 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 2-7: Maintenance Logging

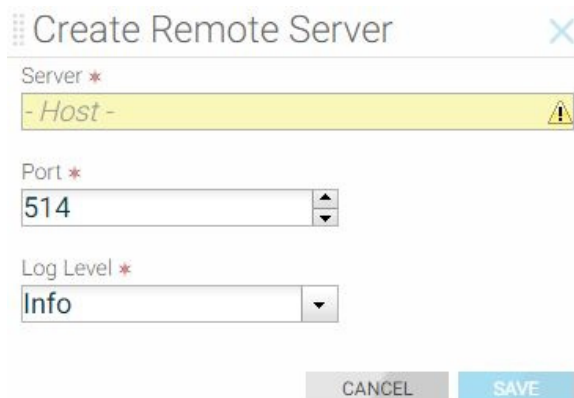


2. To enable remote logging, identify a remote syslog server, by clicking the **Provision control (+)** at the top of the Remote Servers table.



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 2-8: Create 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 is added to 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, which has 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 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_i
drac._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_i
drac._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_i
drac._tcp.local.
...
controller-1>
```

2.8.5 Administrative Activity Logs

User activities are logged to the floodlight log file on the DMF controller. The following events are logged:

- CLI commands entered
- Login/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 in *DANZ Monitoring Fabric 8.4* release, REST API body can be logged to `audit.log`. Prior to this release, REST API calls from GUI or REST-clients are logged in `audit.log` without the data (or) body of the request. In this release, data/body of the rest 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-re
quest-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/9d0a66315f7d9e0df8f2478fe7c0b3d77cec25e865e4e135f0f
9e28237570b70"
user="admin" remote-address="fd7a:629f:52a4:20d0:1ca8:28ed:6f59:cd47" session-
id=
"9d0a66315f7d9e0df8f2478fe7c0b3d77cec25e865e4e135f0f9e28237570b70"
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\", \"us
er[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. With no configuration the **max-body-size** is set to **9223372036854775807** bytes.

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

Below example configuration restricts the body of the rest-api to 16K bytes. If the body size is greater than 16K bytes, the rest call is rejected. This can be useful in DOS attack cases where some intruder tries to send very large configuration repeatedly to make the controller busy in parsing it for applying. This configuration can be enabled to quickly workaround the attack 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 and the configuration required to implement this feature.

2.10.1 Overview

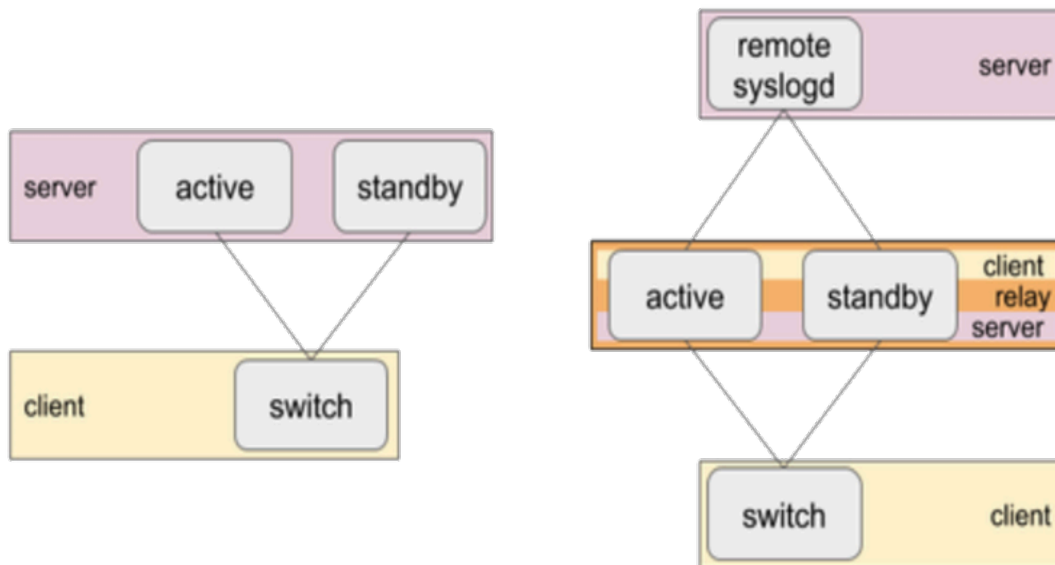
Transport Layer Security (TLS) is used in DANZ Monitoring Fabric to secure syslog messages, which may originate or traverse a non-secure network in transit to the syslog server. This helps mitigate the following primary threats:

- 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 may examine the contents of a syslog messages, 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 counters disclosure of message contents.

Starting with **DANZ Monitoring Fabric 7.3**, syslog over TLS is automatically enabled between controllers and switches when CPSEC is configured on the fabric. Switches send secure syslog messages to active and standby controllers and both controllers relay those messages to remote syslog server.

The following diagram illustrates how TLS over syslog works when messages are sent to one or more remote servers.

Figure 2-9: Syslog Over TLS Between DMF Controller and Remote Syslog Server



In this scenario, when you enable TLS on the controller, this enables a client relay server on the controller. A client relay server includes a server and a client. The server receives messages and hands them off to the local client on the controller. The client on the controller then sends the message to the remote syslog server.

The client relay server simplifies certificate management because you need to manage only the certificates for the controllers and the remote syslog server, rather than for every switch. Secure syslog cannot be enabled between switches and remote syslog server directly. To enable end-to-end (**Switches >**

Controller(s) > Remote server) secure syslog, user has to configure syslog over TLS between controller and remote server.

To complete the deployment of syslog over TLS, the administrator of the remote syslog server must generate certificates for the remote server and the controller using the same public or private CA. The root certificate of the trusted CA must then be copied and installed on the controller.




Note: When upgrading from a BMF version prior to **7.3** deployment with logging remote configured, the syslog port is automatically set to **514** after the upgrade is complete, even though this port was not explicitly configured. Before configuring syslog over TLS, remove the remote server configuration using the `no logging remote <server-ip>` command. Then reconfigure the remote server and the correct port for syslog over TLS (**6514**) will be assigned.

After implementation, when the controller connects to the remote syslog server, the syslog server replies with its public key to allow the controller to encrypt the syslog messages.

There is no restriction requiring all syslog messages to be sent only using TLS if the controller is configured to send messages to an unauthenticated syslog server. To ensure that all syslog messages are encrypted, the DMF controller administrator must remove any unauthenticated syslog servers from the controller configuration.

The following are the high-level steps required to deploy Syslog over TLS:

1. Initialize CPSEC on the DMF Active controller.
This generates the certificates and keys for the connection between the controller and the switches. Refer to [Enabling Secure Control Plane](#) for steps to configure CPSEC on the fabric.
2. Configure syslog on the DMF Active controller.
 **Note:** The syslog subsystem listens for CPSEC-related TLS Events to configure itself with keys/certs.
3. Externally generate keys and signed certificates for the remote syslog servers and the DMF controllers.
4. Import keys/certs to the Active and Standby controllers, and to the remote syslog servers.
5. Enable remote logging on the Active controller.
6. Enable TLS on the connection between the controllers and the remote servers.

2.10.2 Configuration

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

1. Generate the certificates for the remote syslog servers and the DMF controller using the same trusted public or private CA.
2. Copy the trusted CA root certificate, the controller certificate, and the controller 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**, this is cacert.pem.

5. Identify a certificate for mutual authentication.

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



Note: In the example in **Step 2**, this 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**, this 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, dmesg, and journal by default.

2.11 Creating Support Bundle

Collection of running configuration and log files are critical to understanding the fabric behavior when the fabric is in a fault state.

DMF CLI provides the commands to automate the process of collecting, archiving, and uploading the critical data. These commands cover all devices which are part of DMF fabric such as controllers, switches, DMF Service Node, DMF Recorder Node, and Arista Analytics Node.

The following are the commands to configure 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 or not:

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

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

```
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-Contro
ller--2020-11-24--18-31-39Z--
SGPVW-BZ3MM.tar.gz
Url : https://10.2.1.103:8443/api/v1/support/anet-support--DMF-Controll
er--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 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

```

The **diagnose upload support** command is used to check the reachability/health of the server to which the support will be uploaded. Below is an example output of checks that are performed when running the command. When support bundle upload fails for some reason, the command can be used 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#

```

Once the issues are fixed, upload can be retried using the `upload support support bundle file name` command. The same command can be used if auto-upload was not configured prior to taking the support bundle.

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

2.12 NIST 800-63b Password Compliance

This feature activates password compliance for local accounts on 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 going forward will enforce the rules, it does not have an effect on existing passwords. Updating all passwords after the NIST-800-63b compliance has been enabled 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 so that password compliance can be enforced for the entire 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 needs to be set separately on the controller and the Arista Analytics Node so that password compliance can be enforced 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, it is recommended to update local account passwords for DMF devices. Refer to below commands 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>
```

Packet Recorder admin password update

```
Controller(config)# packet-recorder <packet-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 **BMF-7.2.1** and above, existing user credentials continue to work. In case of a FIPS image, password compliance is enabled by default and this doesn't affect existing user accounts. It is recommended to change recovery / local account passwords after upgrade from previous FIPS image. In the case of first time installation of a FIPS image, 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 DMF devices (controller, switches, DMF Service Node, Arista Analytics Node, and DMF Recorder Node).

Starting in **DMF 8.2 release**, custom password requirements are supported for local users.

To enable custom password compliance method, do the following:

```
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 password. It is recommended to reset/reconfigure the user password to adhere to the password requirements. Without configuring the custom password compliance, configuring the requirements is not effective. So user must configure the *custom* compliance method and the password requirements to use this feature.

Once the custom password compliance is enabled, customers can 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
```

So when a user tries to configure/reset password that does not meet any one of the requirements, error will be displayed:

```
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 **show run aaa authentication**:

```
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 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 per-managed-device basis or globally for all managed devices on the CLI. The controller and analytics node are disincluded from receiving this config type (when performed globally).

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

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

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 to the user, outlining the IP address of the management interface 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 intra-fabric-only-access (IFOA) enforced on them. The list of services and their corresponding ports are shown in the table below. An # means enforcing IFOA on that port on that managed device type is impossible.

Conversely, an # means enforcing IFOA on that port on that managed device type is possible. There may be some information beside the # indicating what runs on that port on the managed device.

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

If you want 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
```



Note: It is also possible to push a config that does not override the entire config 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.

Since we can merge the global/default access-control configuration with 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 ma1 acl subnet 10.243.254.20/32 proto tcp port 8443 accept
interface ma1 acl subnet fe80::5054:ff:fe8:b844/128 proto tcp port 8443 accept
interface ma1 acl subnet 0.0.0.0/0 proto tcp port 8443 drop
interface ma1 acl subnet ::/0 proto tcp port 8443 drop
interface ma1 acl subnet ::/0 proto udp port 161 accept
interface ma1 acl subnet 0.0.0.0/0 proto udp port 161 accept
interface ma1 acl subnet 0.0.0.0/0 proto udp port 161 drop
interface ma1 acl subnet ::/0 proto udp port 161 drop
interface ma1 acl subnet 10.243.254.20/32 proto tcp port 22 accept
interface ma1 acl subnet fe80::5054:ff:fe8:b844/128 proto tcp port 22 accept
interface ma1 acl subnet ::/0 proto tcp port 22 accept
interface ma1 acl subnet 0.0.0.0/0 proto tcp port 22 accept
interface ma1 acl subnet 0.0.0.0/0 proto tcp port 22 drop
interface ma1 acl subnet ::/0 proto tcp port 22 drop
interface ma1 acl default accept
.
.
.
```

Note that 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 config 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 doesn't 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 remaining 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 (config will be synced 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 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/control
ler.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 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 is synced 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 so they get the 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.

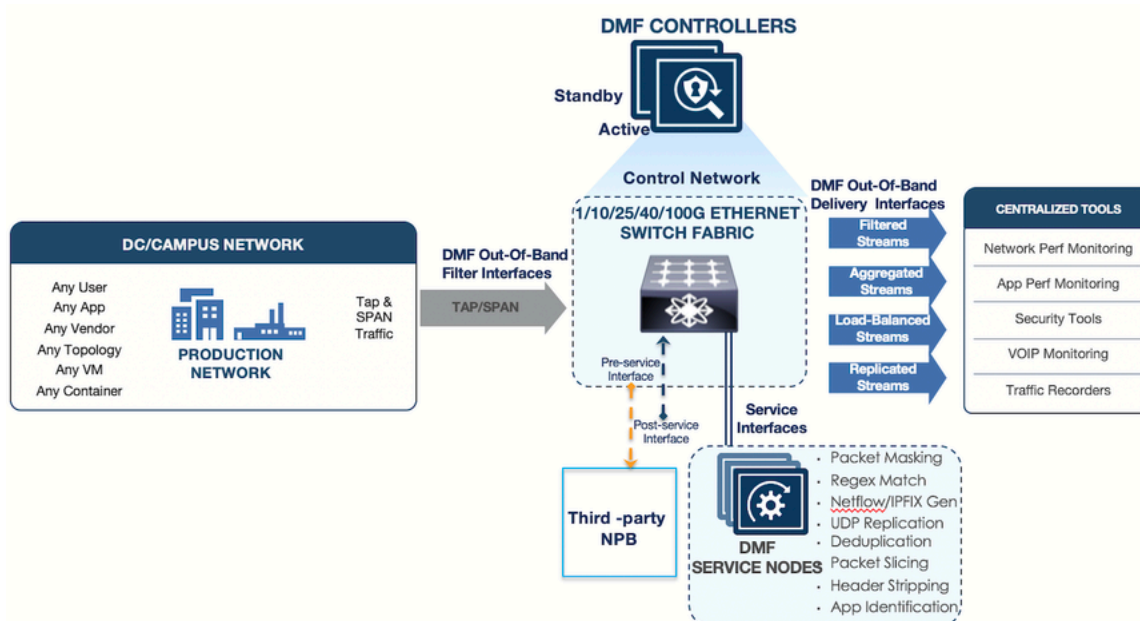
3.1 DANZ Monitoring Fabric Topologies

You can implement out-of-band monitoring using DANZ Monitoring Fabric, with a single switch or with a large number of switches for a scalable, high-availability topology that supports thousands of filter and delivery ports. This section provides a summary of the recommended deployments.

3.1.1 Single-Switch Topology

This 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 a variety of packet manipulation operations, such as time stamping and packet slicing.

Figure 3-1: Single Switch Topology



This design option is most useful in the following scenarios:

- The environment does not yet need to scale beyond the interfaces provided by a single DMF Out-of-Band switch.
- When filter and delivery ports are physically dispersed throughout the data center, a single-switch topology improves cable management.

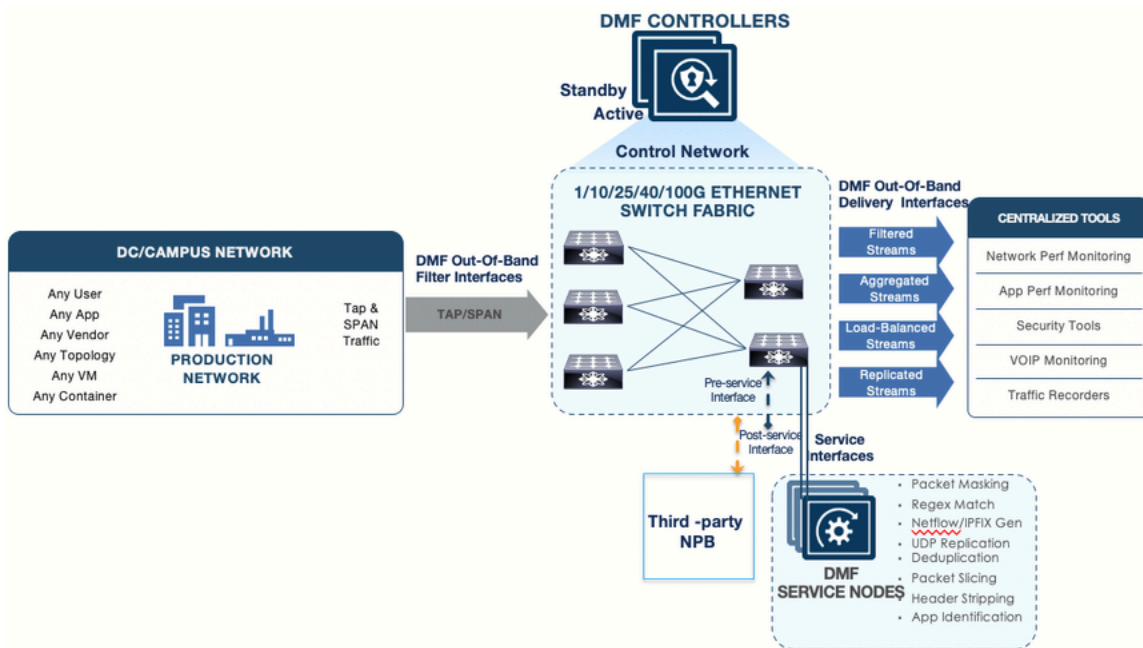
3.1.2 Two-Tier Topology

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

- Medium-to-high port scalability requirement.

- Production network TAPs are dispersed across the datacenter and require aggregation.
- Tools are physically consolidated and the traffic needs to be aggregated.

Figure 3-2: Two-Tier Topology



When deploying this topology, ensure the following requirements:

- Core links should be used only between monitoring switches in different tiers. These links can be either 10, 25, 40, or 100G depending on bandwidth requirements and port availability.
- 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 oversubscription design.
- Service nodes should be connected to the delivery switch so that the aggregated traffic can be sent to the service nodes (NPBs).

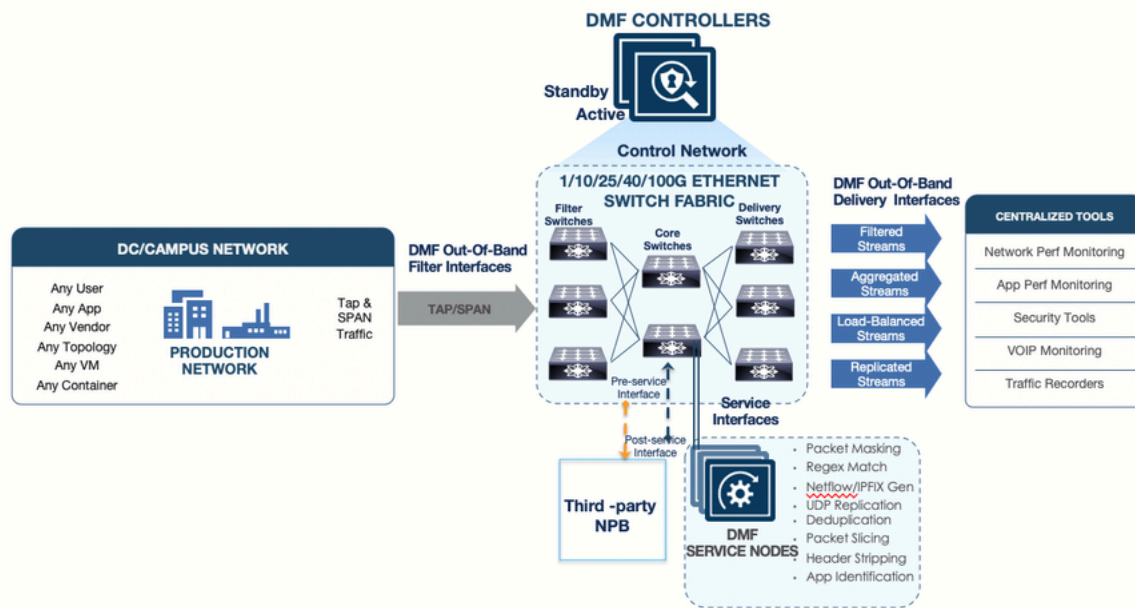
3.1.3 Three-Tier Any-Tap-to-Any-Tool Topology

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

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

- 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: Three-Tier (Any-TAP-to-Any-Tool) Topology



With this topology, ensure the following requirements:

- 40 or 100G links are recommended to the core switches.
- Connect each filter switch to at least two core tier switches for redundancy. More ports can be connected between tiers depending on oversubscription design and port availability.
- 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.

Installing DMF Switches

This chapter describes how to install DANZ Monitoring Fabric switches and perform initial setup and configuration.

DMF release 8.4 supports secure HTTPS connectivity for controller-hosted URLs using ZTP.

Prior to DMF version 8.4 the controller used HTTP to access ZTP install scripts and software images. HTTP does not provide the security required in today's network environments, so the need for HTTPS support arose in those customer environments where all port 80 traffic (HTTP) is blocked. Blocking HTTP access makes the DHCP-based installation of Switchlight and other required software impossible. This new feature allows access to ZTP install scripts and software images via secure HTTPS.

4.1 HTTPS Support for Controller Hosted URLs using ZTP

DMF release 8.4 supports secure HTTPS connectivity for controller-hosted URLs using ZTP.

Prior to DMF version 8.4 the controller used HTTP to access ZTP install scripts and software images. HTTP does not provide the security required in today's network environments, so the need for HTTPS support arose in those customer environments where all port 80 traffic (HTTP) is blocked. Blocking HTTP access makes the DHCP-based installation of Switchlight and other required software impossible. This new feature allows access to ZTP install scripts and software images via secure HTTPS.

This feature does not require any special configuration.

Use the CLI `show switch-image url` command to display the URLs for the ZTP install script and images.

The output contains HTTP and HTTPS URLs for the script and each available image, as shown in the following example.

```
C1> show switch-image url
# File                               Url
      Alternative Url
-----|-----
1 arista-ztp-install-script http://<controller IP>/switchlight/arista-ztp-
install-script
2 arista-ztp-install-script https://<controller IP>/switchlight/arista-ztp-
install-script
3 install-amd64                    http://<controller IP>/switchlight/install-amd64
4 install-amd64                    https://<controller IP>/switchlight/install-amd64
5 update-amd64                     http://<controller IP>/switchlight/amd64
6 update-amd64                     https://<controller IP>/switchlight/amd64
7 update-aristaeos                 http://<controller IP>/eos/x86_64
8 update-aristaeos                 https://<controller IP>/eos/x86_64
```

4.2 Zero Touch Fabric Provisioning Modes

Fabric switch installation can be completed in one of the following two modes:

- **Layer 2 Zero Touch Fabric (L2ZTF, Auto-discovery switch provisioning mode):** In this mode, which is the default, Switch ONIE software automatically discovers the controller via IPv6 local link addresses and downloads and installs the appropriate Switch Light OS image from the controller. This method of

installation requires that all the fabric switches and the DMF controller are in the same Layer 2 network (IP subnet). Also, if the fabric switches require IPv4 addresses to communicate with SNMP or other external services, you must configure IPAM, which provides the controller with a range of IPv4 addresses to allocate to the fabric switches.

- **Layer 3 Zero Touch Fabric (L3ZTF, Preconfigured switch provisioning mode):** When fabric switches are in a different Layer 2 network from the controller, log in to each switch individually to configure network information and download ZTF installer. Subsequently, the switch automatically downloads Switch Light OS from the controller. This mode requires that communication between the controller and the fabric switches occurs 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

All the fabric switches in a single fabric must be installed using the same mode. If you have any fabric switches in a different IP subnet than the controller, you must use Layer 3 mode for installing all the switches, even those in the same Layer 2 network as the controller. Installing switches in mixed mode, with some switches using ZTF in the same Layer 2 network as the controller, while other switches in a different subnet are installed manually or using dhcp is unsupported.

4.3 Using L2 ZTF (Auto-Discovery) Provisioning Mode

ZTF is used to provision and install DMF switches that are in same Layer 2 management network as the DMF controllers. ZTF uses the Open Network Install Environment (ONIE) boot loader for automating switch installation and configuration. Supported fabric switches are shipped with an ONIE network-enabled boot image. For a list of supported monitoring fabric switches, refer to the ***DANZ Monitoring Fabric 8.4 Hardware Compatibility List***. During switch boot up, each switch gets the Switch Light OS software from the DMF controller.



Note: Switches from different vendors can be deployed in the same fabric. However, use of cables to connect switches from different vendors 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 requirements for using ZTF to install fabric switches:

- The supported fabric switches are listed in the ***DANZ Monitoring Fabric 8.4 Hardware Compatibility List*** for your controller. version.
- Connect the management Ethernet interface of each physical switches to the management network and power it up.
- Connect the DMF controller appliance management interface to the same Layer2 management network as the management Ethernet interface of every physical switch.
- When upgrading switches from a previous deployment, make sure you use the Switch Light OS image compatible with your controller version.
- Assign a range of IPv4 addresses to be assigned using IPAM, if switches must communicate with SNMP, NTP, syslog, or other IPv4 services.



Note: DANZ Monitoring Fabric ZTF is implemented using the IPv6 link local address, which is autogenerated.



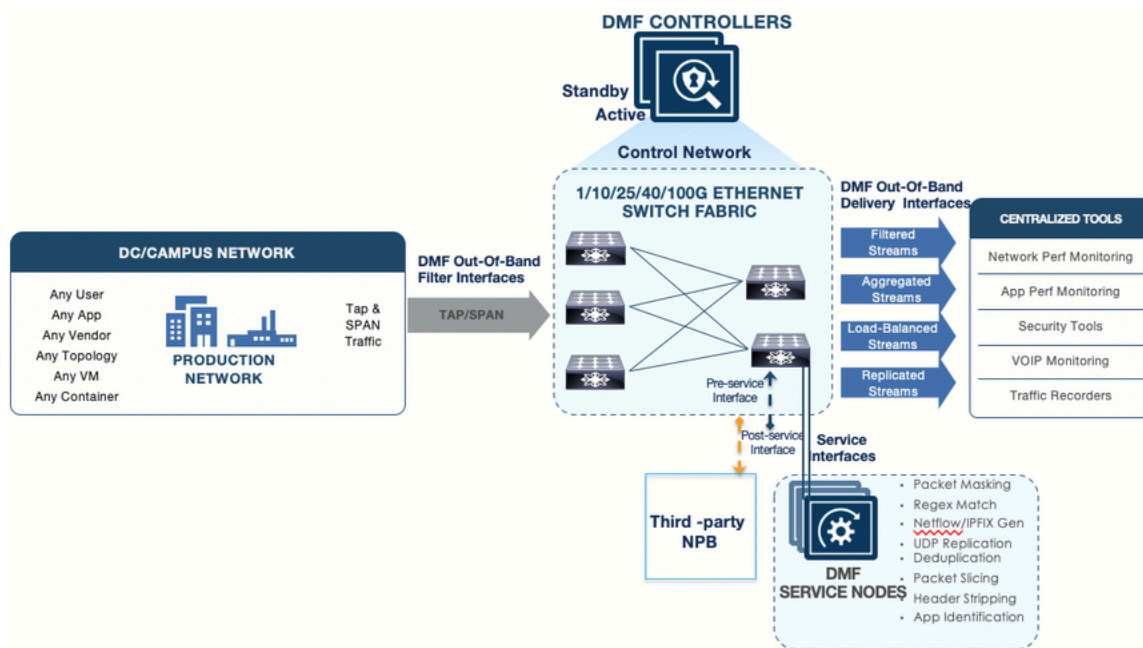
Warning: If a switch has an operating system already installed, following the instructions in [Removing the existing OS from a Switch](#) before connecting the switch to a DMF controller.

4.3.2 Switch Installation Procedure

This section describes how to use ZTF to perform a fresh installation of one or more supported DMF switches, using the Switch Light OS included with the controller software.

The figure below 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 you cannot view the switch label (it will not be visible if you have already racked the switches), you can get 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, 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 mode is the default mode. When you enter the `show running-config` command, you should see `deployment-mode auto-discovery`.

3. Power 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 takes you to 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.
```

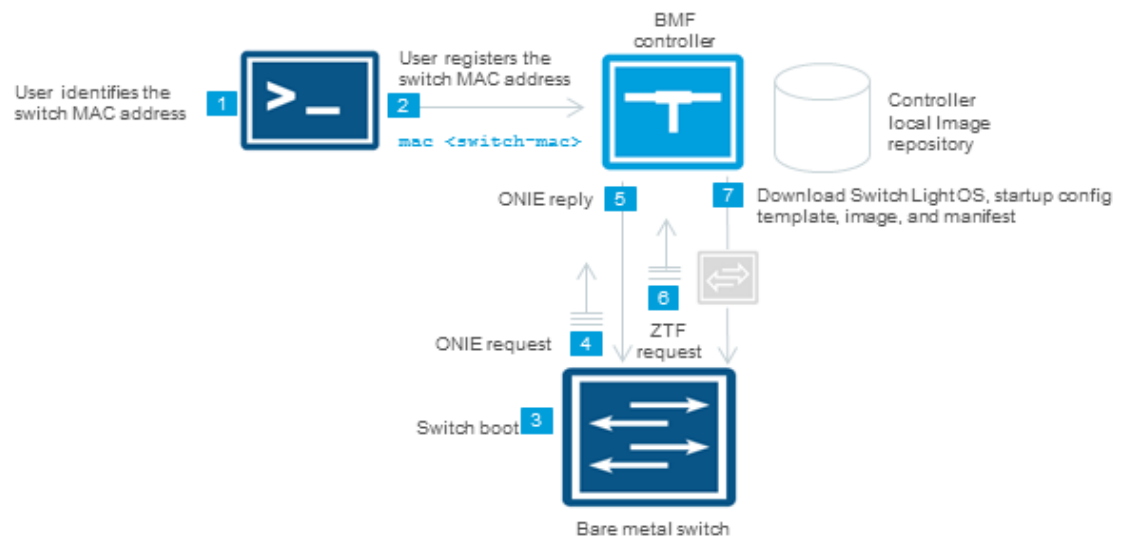
- b. Select **ONIE: Install OS**.

This puts the switch into the installer mode, and the rest of the process is performed 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 that occurs after registering a connected fabric switch on the controller.

Figure 4-2: Switch Boot and ZTF Configuration



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

The rest of the steps in the installation and configuration process happen automatically after you power on the registered and physically connected switch. These automated steps are shown in Figure 20 8 and are summarized below.



Note: Steps 5 through 7 below are performed automatically during the ZTF process and no user intervention should be required. 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 to 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.

```
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 DMF. The initial boot process cannot be performed automatically in the same L2 domain like existing DMF supported switches that are running ONIE.

Initial installation of Switch Light OS on the Arista platforms is accomplished by dropping to the Aboot shell interface at boot and telling it to boot the Switch Light switch image. This operation will cause Switch Light to be installed 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 console cable to the Arista switch. Then power on or reboot.
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. The gateway IP address can be added by using either `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
```

4. Identify the MAC address of the ma1 management interface of the switch. Use the `ifconfig -a` command. The HWAddr is the MAC address. The MAC address is also printed on a label that can be found on the rear of the switch.



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, please take care to press **Control-C** exactly at the prompt and not later. In case 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.

```

About# 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

```

5. On another 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 **#ter-1**.

```

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

```

6. On the DMF controller, issue the command **show switch-image url** to obtain the URL for the Switch Light boot image. The URL that is needed is **Update-amd64**. Make note of this URL as it is needed in the next steps.

```

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

```

7. Back on to the switch About shell, boot with the URL of the SwitchLight image. The command is **boot url**. This is the full console log of a successful Switch Light OS installation in L2 ZTN mode.

```

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
4444.69: Running SwitchLight install...
...
About 9.0.3-4core-14223577
Press Control-C now to enter About 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_s
ize=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-7050s
x3-48yc8-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:

```

8. Once the switch has booted up successfully, the Arista switch will show in connected state when viewed from the DMF controller.

```

DMF-CTRL(config)# show switch DMF-F1
# Switch Name IP Address State Pipeline Mode
- |-----|-----|-----|-----|
1 DMF-F1 fe80::7272:cfff:febd:dcbc%9 connected 13-14-match-push-vlan
DMF-CTRL(config)#

```


4.3.4 Allocating IPv4 Addresses to Fabric Switches

When using L2ZTF, the DMF controllers and fabric switches use Link-local IPv6 for communication. In order for switches to be able to communicate with external (IPv4) services, you must configure IP address management (IPAM), which assigns an IPv4 address to the switch from a configured pool of addresses. You may want to do this to enable a fabric switch in L2ZTN mode to communicate with external services, such as NTP, SNMP, and syslog. Note that 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. If you try to configure IPAM when the provisioning mode is set to Layer 3 (Preconfigured), an error message is displayed.

4.3.5 Using the GUI to Allocate IPv4 Addresses

To allocate a pool of IPv4 addresses, complete the following steps:

1. Select **Fabric > Switches** from the main menu.

Figure 4-3: Allocating IPv4 Addresses to a Switch

Starting IP	Ending IP	Subnet Mask Length	Addresses Allocated	Addresses Used	Utilization
10.240.130.11	10.240.130.29	25	19	9	47.4%

This page lists the switches connected to the DMF controller and provides an option for configuration a pool of IPv4 addresses to use for IPAM assignment to the fabric switches.

2. Click the **Settings** control at the top of the IP Address Allocation section.

Figure 4-4: Configure Switch IP Allocation

Starting IP	Ending IP	Subnet Mask Length	Addresses Allocated
10.240.130.11	10.240.130.29	25	19

3. Enable the checkbox for an existing range, or configure a new range.
To configure a new range, click the **Provision control (+)** at the top of the **IP Ranges** section, and enter the range of IP addresses on the dialog that is displayed.
4. When finished selecting IP address ranges to add to the pool, click **Submit**.

4.3.6 Using the CLI to Allocate IPv4 Addresses

To allocate a pool of IPv4 addresses and configure the DNS server and default gateway, complete the following steps:

1. Enter config-ipam-switch submode by entering the `ipam switch` command.

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

2. Identify the DNS server to be used by the fabric switches by entering the `dns-server` command.

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

3. Identify the default gateway server to be used by the fabric switches by entering the `dns-server` command.

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

4. Identify the range of addresses to be used by the fabric switches by entering the `ip-range` command.

```
controller-1(config-ipam-switch)# ip-range 192.168.1.100 192.168.1.200
subnet-mask-length 24
```



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

5. Enable IPAM IPv4 address allocation to the fabric switches by entering the `allocate` command.

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

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

When a switch is in a different subnet than the controller, you need to configure the network information in switches and download the ZTF installer from the controller downloads the Switch Light OS image.



Note: Switches from different vendors can be deployed in the same fabric. However, use of cables to connect switches from different vendors 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.

Make sure you enter the `deployment-mode pre-configured` command 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 DMF provisioning mode set to pre-configured, complete the following steps:

1. Confirm that the switch has ONIE installed.



Note: Note: To do this, power on the switch, 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, which are listed in the **DANZ Monitoring Fabric 8.4 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 IP address in **Step 2a** is not assigned then follow **Steps 2b** and **2c**. If IP address is already assigned in **Step 2a**, then skip to **Step 3**.

- a. Check IP addressing on 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 assigning IP address to management port.

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

- c. Adding 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 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`.

4. For amd-64 platform boxes (for example, Dell S6100):
 - a. Reboot and select ONIE from GNU/GRUB menu.
 - b. 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 takes you to the ONIE install mode: `run onie_bootcmd`
 - c. Type `onie-discovery-stop` at the installer mode in ONIE.

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

5. Verify that switch can ping to 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
```

6. Use the image-url path from the controller for the platform amd-64.

```
ONIE:/ # install_url http://10.8.25.31/switchlight/install-amd64
discover: installer mode detected.
Stopping: discover... done.
```



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.
```

7. 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#
```

8. Enter **zcs** to drop to **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)
```

9. 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: If you try to configure IPAM when the provisioning mode is set to Layer 3 (Preconfigured), an error message is displayed.

10. Configure the IP address of the DMF controller (ZTN server).

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

11. Set the DNS information.

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

12. 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)
```

13. 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 you see the final messages, the installation is completed.

14. 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 : CNOWKIFYN7793164P0004
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)#
```

15. You can log in to the switch and enter the **show** command to display the switch status, when you see the following prompt.

```
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

The initial installation of Switch Light OS on the Arista platforms must be performed manually in the DMF similar to what is required on with ONIE enabled switches. Use this guide for DMF controllers configured for L3 ZTN mode.



Note: L3 ZTN means the controller is setup for deployment-mode pre-configured.

Initial installation of Switch Light OS on the Arista switch is accomplished by dropping to the Aboot shell interface at boot and telling it to boot the Switch Light switch image. This operation will cause Switch Light to be installed 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 currently running EOS.

Procedure

1. Attach console cable to the Arista switch. Then power on or reboot.
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. The MAC address is also printed on a label that can be found on the rear of the switch.

```
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 another 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 URL for the Switch Light boot image. The URL that is needed is **Update-amd64**. Make note of this URL that will be needed in the next steps. 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 SwitchLight 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 SwitchLight 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. Static IP is used in the example here.

```
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 manual ZTN configuration. Type **Yes** to complete the SwitchLight 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 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-Abboot-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:

```

17. Once the switch has booted up successfully, the Arista switch will show in connected state when viewed from the DMF controller.

```

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

DMF 8.1.0 is the first release to support the Arista 7280R Series switches. Instead of running Switch Light OS, the 7280R switches will be the first set of switches to run on EOS operating system while deployed in the **DANZ Monitoring Fabric**. In **DMF 8.1.0**, the supported models are the 7280SR, 7280SR2, and 7280CR3.

The controller must be configured for L3 ZTN mode in order for the 7280R switches to connect. The 7280R switches will not be able to connect to the DMF controller while in L2 ZTN mode (L2 ZTN is the default setting).



Note: L3 ZTN means the controller is setup for deployment-mode pre-configured.

Also in **DMF 8.1.0**, the DMF fabric does not support mixing 7280R switches running EOS with any other switches running the Switch Light OS. All switches on DMF fabric must be all Arista 7280R models.

This is a one-time setup needed to load the DMF compatible EOS image. Once you have this setup, the next controller upgrade will automatically upgrade the switches as well.

Perform these steps on the 7280R 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 update-aristaeos. Make note of this URL which will be needed later when copying the image to the switch. Do not use the other URLs.

```
DMF(config)# show switch-image url
File           Url
-----
install-amd64  http://172.24.210.21/switchlight/install-amd64
update-amd64   http://172.24.210.21/switchlight/amd64
update-aristaeos http://172.24.210.21/eos/i686
```

2. Attach a console connection to the Arista 7280R Series switch. Then power on the switch or reboot it.
3. Interrupt the boot process with **Control-C** to drop into the About 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
About 9.0.3-4core-14223577
Press Control-C now to enter About shell
^CWelcome to About.
About#
```

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**.

```
About# 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.

```
About# ifconfig ma1 172.24.210.61 netmask 255.255.252.0
About# 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.

```
About# cd /mnt/flash
About#
About# 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 beginning step (**Step 1**). The command is **wget url**.

```
About# wget http://172.24.210.21/eos/i686
Connecting to 172.24.210.21 (172.24.210.21:80)
i686 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:/i686
```

10. Verify the boot-config was saved.

```
About# cat /mnt/flash/boot-config
SWI=flash:/i686
```

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

```
About# reboot
About# [ 1096.250482] sysrq: SysRq : Remount R/O
Requesting system 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
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
PMxC0 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
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
```

```

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
Booting flash:/i686
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. Login as admin. From 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-Abboot-norcal9-9.0.3-4core-14223577 Wed Nov 13
 22:08:55 UTC 2019
bootblock starting...
Family_Model: 00660f01
PMxC0 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
PMxC0 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
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
Warning - AGESA callout: platform_PcieSlotResetControl not supported
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
Abboot 9.0.3-4core-14223577
Press Control-C now to enter Abboot shell
Booting flash:/i686
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 in show version. Type: **show version**. In this example, the System **MAC = d4af.f754.195b**. This MAC address can also be obtained by looking at 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: i686
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 format of 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-Ma1)# ip address 172.24.210.89/22
```

16. 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 1 controller. ZTN configuration download will begin after configuring this part.

```
localhost(config-if-Ma1)# management dmf
localhost(config-mgmt-dmf)# controller address 172.24.210.21
localhost(config-mgmt-dmf)# no disabled
```

17. 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
```

18. On the DMF controller, the switch will show connected as well.

```
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 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 zcsh CLI, complete the following steps:

1. From the Switch Light OS prompt for the switch you are configuring, 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 **zcsh** to drop to **ztn-config** mode.

```
loader# zcsh
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 you have 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.
```

4.4.5 Installing Arista 7050X and 7260X Series using dhcp with bootfile-name option

Starting with **DMF 8.2**, the installation of Switch Light OS on Arista switches can be automated by using an Arista ZTP boot script that is 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 by 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 is 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 URL of the Arista ZTP boot script. 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. Example here is the `/etc/dhcp/dhcpd.conf` file from an ISC DHCP server. The DMF switches can be brought up using A) **vendor-class-identifier** or B) **switch hardware address**.



Note: The edits required depend on your network environment and the type of switches.

- a. For a large DMF deployment one can use vendor-class-identifier as shown below. **DMF switches should use 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 UCN switches, then 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 (dhcpd) after editing `/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.

```
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)#
```

4.4.6 Installing Arista 7280R Series using dhcp with bootfile-name option

The installation of EOS on Arista SAND platforms (7280x) can be automated by using an Arista ZTP boot script that is available on the DMF controller. This procedure is applicable to all Arista 7280R Series platforms that are supported in DMF.

The Arista ZTP boot script is served to the Arista switch by 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 and from the DMF controller and configures the appropriate boot settings on the Arista switch.



Note: Controller must be setup for deployment-mode pre-configured.

Procedure

1. Connect to the Arista switch to get the system MAC access. The system MAC address is the HWaddr for interface ma1 in About. The MAC is this example is **D4:AF:F7:F9:EE:38**.

```
About# ifconfig ma1
ma1 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 **filter-1**.

```
switch DMF-F1
mac d4:af:f7:54:19:5a
```

3. From the DMF controller, obtain the URL of the Arista ZTP boot script. 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. Example here is the `/etc/dhcp/dhcpd.conf` file from an ISC DHCP server. The DMF switches can be brought up using A) **vendor-class-identifier** or B) **switch hardware address**.



Note: The edits required depend on your network environment and the type of switches.

- a. For a large DMF deployment one can use vendor-class-identifier as shown below. DMF switches should use 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 UCN switches, then 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 (dhcpd) after editing `/etc/dhcp/dhcpd.conf` file.
6. The Arista ZTP boot script will be downloaded and executed at the ZTP phase.
7. On DMF controller, verify the switch is connected.

```
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)#
```

4.4.7 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 when using fabric switches of the same architectural type, in this case, using 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 how a dhcp configuration file looks with both 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, then 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 case of DMF, this is optional as the `default-url` provides for SWI and configurations as well.
3. Restart the dhcp server process (`dhcpd`) 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, simply register the name and MAC address of the switch with the Active DMF controller. The switch downloads a compatible Switch Light OS image and configuration from the controller and the registered switch name is used to refer to the switch in the CLI output, as well as GUI displays.

4.5.1 Using the GUI to Register a Switch

Procedure

1. Select **Fabric > Switches** from the main menu.

Figure 4-5: Fabric Switches Option

The screenshot shows the 'Fabric Switches' configuration page. The top navigation bar includes 'Fabric', 'Monitoring', 'Maintenance', 'Integration', 'Security', and 'Profile'. The main content area is divided into two sections: 'IP Address Allocation' and 'Switches'.

IP Address Allocation:

- Status: Enabled
- DNS Server: 10.240.48.6
- Gateway: 10.240.130.1
- Total Allocated Addresses: 19

Switches Table:

Name	Connected	Admin Status	Management Interface	Connected Since	Connection Time	Management IP Address	Allocated IP Address	Port
DMF-D1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 23, 2020, 6:44:47pm UTC	4d 7h 12m	fe80::7272:cfff:fe1:7c7b99	10.240.130.18/25	49998
DMF-F1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 23, 2020, 6:45:13pm UTC	4d 7h 12m	fe80::7272:cfff:febd:dc9c99	10.240.130.19/25	49132
DMF-F2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:34:16pm UTC	5d 7h 26m	fe80::7272:cfff:fe1:7ecb99	10.240.130.11/25	60877
Site2-Core1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:34:20pm UTC	5d 7h 26m	fe80::7272:cfff:febd:477299	10.240.130.12/25	54374
Site2-Core2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:34:26pm UTC	5d 7h 26m	fe80::7272:cfff:febd:5b4899	10.240.130.13/25	51960
Site2-D1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:34:21pm UTC	5d 7h 26m	fe80::7272:cfff:febd:df5699	10.240.130.14/25	49479
Site2-F1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:34:58pm UTC	5d 7h 26m	fe80::7272:cfff:fe7:c5ef99	10.240.130.15/25	42305
Site2-F2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> Up	--	Nov 22, 2020, 6:33:50pm UTC	5d 7h 26m	fe80::ce37:abff:fe2c:97ea99	10.240.130.16/25	38918

This page lists the switches connected to the 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 to the left of the switch and select Configure from the pull-down menu that appears.

Figure 4-6: Configure Switch (Page 1)

This dialog lets you 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.

This dialog also provides access to a series of dialogs that you can use to override the default configuration that is pushed from the DMF controller to the switch. To get to another page, click the numbered link for the page or click **Next**.

- Type the alias for the switch in the **Name** field.
- Type the MAC address and click **Submit**.

4.5.2 Using the CLI to Register a Switch

Enter the `switch switch-name` command to enter the config-switch submode, from where you 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, 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-1314-push-vlan
2 bigtap-switch-2  fe80::ce37:abff:fe60:cf8a%2  connected  bigtap-1314-push-vlan
3 bigtap-switch-3  fe80::ce37:abff:fea0:9071%2  connected  bigtap-1314-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
2 DMF-DELIVERY-SWITCH-1 cc:37:ab:60:d4:74 (Edgecore) 00:00:cc:37:ab:60:d4:74 connected
3 DMF-FILTER-SWITCH-1 8c:ea:1b:26:73:6f (Edgecore) 00:00:8c:ea:1b:26:73:6f connected
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

You cannot use ZTF to install the switches when the switch management network is connected to the DMF controllers through a Layer 3 network. However, when a switch is in a different subnet than the controller, you can 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 switches and controllers are in the same L2 broadcast domain, the default switch deployment-mode (auto-discovery) can be used for L2-ZTF deployment. If the switches and controllers are not in the same L2 broadcast domain, the provisioning mode should be changed to pre-configured to enable L3-ZTF deployment. The entire fabric must be in a single provisioning mode; using the auto-discovery switch provisioning mode is only supported if all the switches are in the same Layer 2 domain.

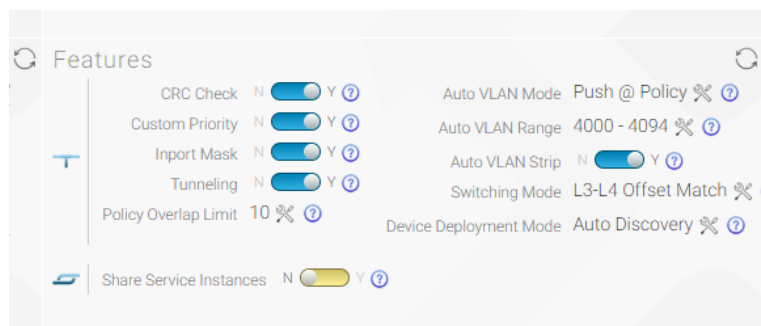
4.6.2 Using the GUI to Change the Switch Provisioning Mode

To use the GUI to change the switch provisioning mode, complete the following steps.

Procedure

1. Click the **DMF logo** in the upper left corner of the GUI Main menu to display the controller landing page.
2. On the controller landing page, click the **Settings** control in the Features section.

Figure 4-7: Changing the Switch Provisioning Mode



Note: For information about the other options shown in this section, refer to the DANZ Monitoring Fabric Out-of-Band User Guide.

- Click the Settings controller to the right of the **Device Provisioning Mode** option.

Figure 4-8: Configure Device Provisioning Mode



- Move the slider to **Pre-Configured**, click **Submit**, 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

To use the CLI to change from Auto-discovery (L2-ZTF) Mode to preconfigured (L3-ZTF) Mode, complete the following steps.

Procedure

- Disable IPAM by entering the following commands.

IPAM is supported only in L2-ZTF mode. You must disable it to move to L3-ZTF mode. Before you can disable IPAM, you must remove all the IPAM configuration in **conf#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 DMF controller.

```
controller-1 (config) # deployment-mode pre-configured
```

- Configure the switches using dhcp or static IP addresses.



Note: If not using dhcp, you must assign a static IP using the switch CLI (PCLI) on each fabric switch.

- Configure the Active and Standby DMF controller IP address on each fabric switch.
- Reboot each switch.

If switches and controllers are in the same L2 broadcast domain, the default switch deployment-mode (auto-discovery) can be used for L2-ZTF deployment. If the switches and controllers are not in the same L2 broadcast domain, the provisioning mode should be changed to pre-configured to enable L3-ZTF deployment. The entire fabric must be in a single provisioning mode; using the auto-discovery switch provisioning mode is only supported if all the switches are in the same Layer 2 domain.

4.6.4 Changing to Layer 3 ZTF (Preconfigured) Mode

To change to Preconfigured (L3-ZTF) Mode from Auto Discovery (L2-ZTF) mode, complete the following steps.

Procedure

- 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 mal ip-address dhcp
interface mal ip-address <ip-address>/<prefix> gateway <gateway-address>
A) Setting mal interface of switch to use DHCP.
(ztn-config) interface mal ip-address dhcp
(ztn-config) show
IP Option: DHCP
Controllers: fe80::250:56ff:fea2:df9b
(ztn-config)
B) Setting mal interface of switch for static IP
(ztn-config) interface mal 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 DMF controller to enable Layer 3 mode switch installation, whether using the manual method or dhcp with 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

To change to Auto-Discovery (L2-ZTF) Mode from Preconfigured (L3-ZTF) Mode, complete the following steps.

Procedure

1. If required, move the switches or the controllers so that they are in same L2 broadcast domain.



Note: L2-ZTF requires that all the switches and controllers are located in same broadcast domain.

2. Change the switch provisioning mode by entering the following command on the Active 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 mal 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

A system reinstall is accomplished by removing the local startup-config/zerotouch-config on the switch so the 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: 13-ztn currently configured
system switch reinstall: 13-ztn implies switches are remote
system switch reinstall: 13-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 once 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 not present. The DMF ZTN process on the switch needs to be kick-started 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 the state is stuck in reloading (and there's a `Zerotouch state error hint`/`Zerotouch state error msg` output), please contact [Arista Support](#).

Managing Switches and Interfaces

This chapter describes how to manage switches and interfaces after installing the monitoring fabric switches. It includes the following sections.

5.1 Configuring Link Aggregation

Link Aggregation lets you combine multiple LAN links and/or cables in parallel. Link aggregation provides a high level of redundancy and higher transmission speed.



Note: When connecting a LAG to a DMF Service Node appliance, member links can be connected to multiple DMF Service Node appliances with data ports of same speed.

DMF provides a configurable method of hashing for load distribution among the members of a LAG. The enhanced hashing algorithm, automatically assigns the best hashing type for the switch and traffic. This setting also lets user manually select the packet types and fields used for load distribution among the members of a port-channel interface. For the supported switch platforms, enhanced mode and symmetric hashing are enabled by default. With symmetric hashing, bidirectional traffic between two hosts going out on a port channel is distributed on the same member port.

The default hashing option use the best available packet header field that applies to each packet and that is supported by the switch. These fields can include the following:

- IPv4
- IPv6
- MPLS (disabled by default)
- L2GRE packet

If none of these headers can be used, Layer 2 header fields (source MAC address, destination MAC address, VLAN-ID, and ethertype) are used to distribute traffic among the LAG member interfaces. Hashing on the following packet header fields is enabled by default:



Note: VN-tagged packets and QinQ packets are treated as L2 packets and Layer 2 headers are used to distribute traffic among LAG member interfaces for these packets:

- hash l2 dst-mac eth-type src-mac vlan-id
- hash ipv4 dst-ip src-ip
- hash ipv6 dst-ip src-ip
- hash l2gre inner-l3 dst-ip src-ip
- hash symmetric

5.1.1 Using the GUI to Configure Link Aggregation Groups

To view, manage, or create Link Aggregation Groups (LAGs) in the monitoring fabric, complete the following steps.

Procedure

1. Select **Fabric > LAGs** from the main menu.

Figure 5-1: Link Aggregation Groups

The screenshot shows a navigation bar with 'Fabric', 'Monitoring', 'Maintenance', 'Integration', 'Security', and 'Profile'. Below it, a breadcrumb trail reads 'System OK DANZ Monitoring Fabric > LAGs'. The main heading is 'Link Aggregation Groups (LAGs)'. Below the heading is a table with a search bar 'Filter table rows' and a 'FILTER' button. The table has the following columns: Switch, Name, Hash Type, State Flags, Members, Tx Packet Count, Tx Byte Count, Tx Packet Rate, Tx Bit Rate, Tx Peak Bit Rate, Tx Peak Packet Rate, Tx Peak Byte Time, Tx Peak Packet Time, and Rx Packet Count. The table is currently empty, and the text 'No lags' is displayed at the bottom right of the table area.

2. To create a new LAG, click the **Provision control (+)** at the upper left corner of the table.

Figure 5-2: Create LAG

The 'Create LAG' dialog box contains the following fields:

- Name ***: A text input field containing 'pg-1'.
- Hash Type**: A dropdown menu with 'Auto' selected. Below it is the text 'Best hash based on hardware compatibility'.
- Switch ***: A dropdown menu with '- Select Switch -' selected.

The bottom of the dialog box shows two buttons: 'Cancel' and 'Save'.

3. Enter a name for the LAG.

- Select the switch where you want to create the LAG.

Figure 5-3: Create LAG (populated)

- Select the interfaces to include in the LAG and click **Submit**.

5.1.2 Using the CLI to Configure Link Aggregation Groups

Use the `lag-interface` command to enter the *config-switch-lag-if* submode, where you can define the LAG member interfaces and specify the type of load distribution (hashing) to use for the LAG.

Use the `member` command to add an interface to a LAG. Enter this command for each interface you want to add to the LAG. To remove an interface, use the `no member` version of the command.

For example, the following commands add two interfaces to a LAG named *my-lag*.

```
controller-1(config)# switch DMF-FILTER-SWITCH-1
controller-1(config-switch)# lag-interface mylag
controller-1(config-switch-lag-if)# member ethernet13
controller-1(config-switch-lag-if)# member ethernet14
```

To configure multiple delivery interfaces as a LAG, complete the following steps.

- Assign a name to the LAG and enter `config-switch-lag-if` submode.

```
controller-1(config)# switch DMF-DELIVERY-SWITCH-1
controller-1(config-switch)# lag-interface lag1
controller-1(config-switch-lag-if)#
```

2. Assign members to the LAG.

```
controller-1(config-switch-lag-if)# member ethernet39
controller-1(config-switch-lag-if)# member ethernet40
```

3. To view the LAGs configured, enter the show lag command, as in the following example.

```
controller-1> show lag
```

5.2 Connecting Directly to a Switch

For installing the Switch Light OS individually on a switch in a different Layer 2 domain, or for troubleshooting the switch, you can telnet or SSH to the switch.

To allow SSH to a switch, if using ZTF for installing switches from the DMF controller in the same Layer 2 domain, you must configure an IPAM IP address pool. Alternatively, you can 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:e
bff: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, type debug admin command at the ztn-config prompt that is displayed.

```
(ztn-config) debug admin
DMF-FILTER-SWITCH-1>
```

This provides access to the switch CLI prompt, which you can use for directly managing or troubleshooting the switch. The following commands are available from 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

5.2.1 Manually Configuring Enhanced Hashing for Load Distribution

In some scenarios, you may want to manually select the bytes in each packet that are used for load distribution among the members in a LAG.

To manually configure enhanced hashing, enter the `hash-type enhanced` command from the `config-switch-lag-if` submode. Then enter the `lag-enhanced-hash` command to enter `config-switch-hash` submode, where you 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 8.4 Hardware Compatibility List*.

Changes in hash configuration do not affect the LAG configuration, so you do not have to reconfigure LAGs after changing the hash type.

5.2.2 Configuring Enhanced Hashing

To configure enhanced hashing, use the `lag-enhanced-hash` command to enter the ***config-switch-hash*** submode. Then use the `hash` command to identify the hash type and the specific fields to use 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}
```

5.2.3 Symmetric Load Balancing

Enhanced hashing supports symmetric load balancing (enabled by default) for switch platforms that support this feature.



Note: Symmetric hashing is supported on specific switches and for IP and 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. The recommended configuration for optimal symmetric behavior is to enable hashing on source IP address and destination IP address.



Note: In some scenarios, Layer 4 protocol ports can be used to improve load balancing efficiency, but these fields are not used by default because they cannot be used if packet fragmentation is likely to occur.

5.2.4 GTP Hashing

Generic Tunneling Protocol (GTP) hashing provides more even distribution of GTP-encapsulated packets among the members of a port-group. When GTP hashing is enabled, 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, the TEID in GTP header is used for hashing GTP traffic instead of using L4 ports. You must configure both `I4-dst-port` and `I4-src-port` in 'hash ipv4' or 'hash ipv6' for proper operation.

5.3 Overriding the Default Switch Configuration

After completing switch installation, further switch configuration, including software upgrades, is managed from the DMF controller.



CAUTION: All configuration changes related to fabric switches must be made either through the controller CLI or the controller GUI, which provides *DANZ Monitoring Fabric 8.4 Deployment Guide* configuration options in the config-switch submenu 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, which eliminates the need for box-by-box configuration. However it is possible merge/override the default configuration pushed from the DMF controller with switch specific configuration for some parameters. These parameters are as follows:

- Clock
- SNMP and SNMP Traps
- Logging
- TACACS

We currently support 2 types of overriding mechanisms: * **override-global:** Only the switch-specific config is applied * **merge-global:** The global config and switch-specific config 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 configs; take it as-is in the resultant config, else: * If the key exists in both global and switch configs: 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 configs; take that list as-is in the result config, else: * If it exists in both global configuration as well as per-switch configuration then merge with this rule:

- If the global and switch-specific config has an entry with the same key, the switch-specific list entry completely replaces the entry from the **global-config**, else:
- All entries from the switch-specific config 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.
2. To override any of the default switch configuration settings, click the **Menu** control to the left of the switch and select **Configure** from the pull-down menu that appears.

Figure 5-4: Configure Switch (Page 1)

This dialog provides access to a series of dialogs that you can use to override the default configuration that is pushed from the DMF controller to the switch.

3. To get to another page, click the numbered link for the page or click **Next**.
4. After making any changes required, click **Submit**.

CLI Procedure

If you need to override the default configuration for a specific switch, enter the **config-switch** submode for the specific switch and use the commands available, which you can see 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) #
```

You can 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.

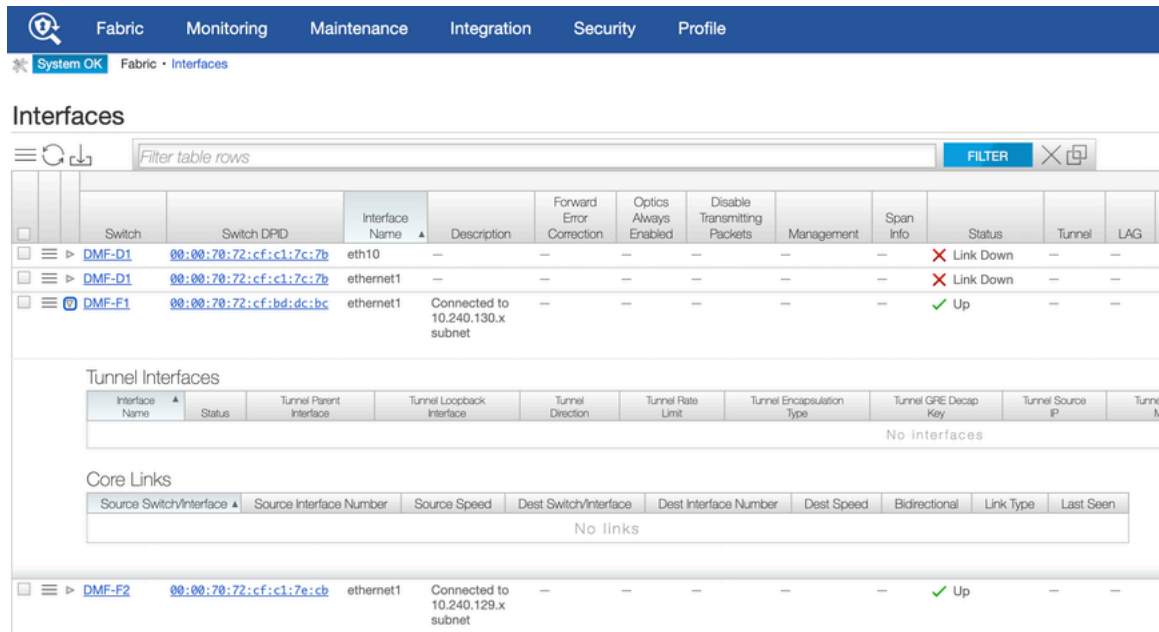
5.4 Configuring Switch Interfaces

To use the GUI to configure switch interfaces, complete the following steps.

Procedure

1. Select **Fabric > Interfaces** from the GUI Main menu.

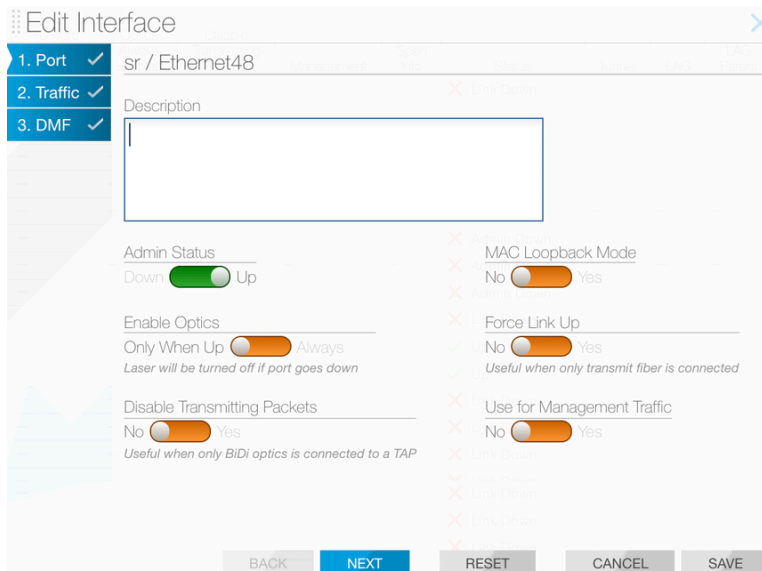
Figure 5-5: Fabric Interfaces Option



This page lets you monitor and configure the interfaces on switches connected to the 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 that are currently using the interface.

To configure an interface, click the Menu control to the left of the interface and select Configure from the pull-down menu that is displayed.

Figure 5-6: Configuring Interface Settings - Page 1



This dialog provides access to three pages. Page 3. DMF page lets you configure the interfaces for use in policies. For further information refer to the **DANZ Monitoring Fabric 8.4 User Guide**.

Page 1. 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.

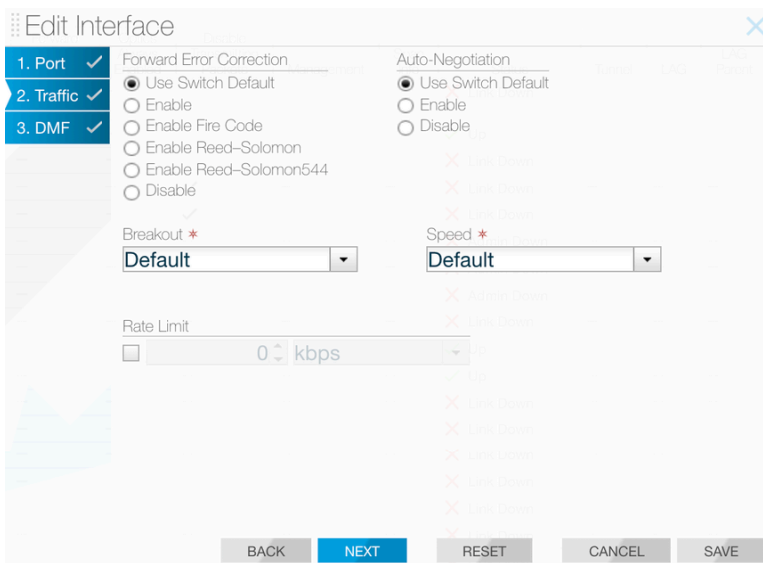
- **MAC Loopback Mode:** Returns traffic to the originating interface.
- **Force Link Up:** This is useful to enable when only the transmit fiber is connected.
- **Description:** Assign a description for the interface.



Note:

- Ideally force link up configuration should be applied only on Delivery interface.
 - This configuration allows L1 on port to stay UP, even when optical fiber cable is connected only in Tx direction.
 - This feature helps to black hole traffic if applied on links between switches.
2. If you are finished, click **Save**. To configure traffic options, click **Next**.
If you click Next, the system displays the following page.

Figure 5-7: Configuring Interface Settings - Page 2



This page provides the following options.

- Forward Error Correction
 - Auto-Negotiation
 - Breakout
 - Speed
 - Rate Limit
3. After making any changes required, click **Save**.

5.4.1 Forward Error Correction

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, using enter following command 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 current port context.
- **enabled** – Enable if possible in 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.

5.4.2 Autonegotiation

The following summarizes the effect of each option:

- autoneg enabled – Enable if possible in current port context.
- autoneg disabled – Disable if possible in current port context.

Autonegotiation can be enabled or disabled for the following interface configurations:

- 100G DAC in 100G mode
- 1G-BASE-SX
- 1G-BASE-LX

GUI Procedure

Use **Page 2** of the **Edit Interface** dialog to enable or disable Auto-Negotiation, or to restore the default.

CLI Procedure

To use the CLI to explicitly enable or disable auto-negotiation, 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.

5.4.3 Manually Setting the Interface Speed

To manually set the interface speed for an interface, from **con#g-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

5.4.4 Using Breakout Cables

Breakout (splitter) cables are supported by DANZ Monitoring Fabric to allow splitting a single 40G, 100G, 200G, 400G port into individual sub-interfaces.

For a list of the switches, ports, and breakout cables that are supported, refer to the **DANZ Monitoring Fabric 8.4 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 find out breakout capable ports from the controller, enter the following command and you will see

```
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 `conf# mode`, as in the following example.

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

3. Enter the `conf#-switch` submode and then the `conf#-switch-if` submode, where you can 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, 1g, 10g, 25g, 40g, 100g
55 ethernet54/2 5c:16:c7:13:d5:d6 (Big Switch) up      up      40g fiber, 10g fiber, 1g, 10g, 25g, 40g, 100g
56 ethernet54/3 5c:16:c7:13:d5:d7 (Big Switch) up      up      40g fiber, 10g fiber, 1g, 10g, 25g, 40g, 100g
57 ethernet54/4 5c:16:c7:13:d5:d8 (Big Switch) up      up      40g fiber, 10g fiber, 1g, 10g, 25g, 40g, 100g
```

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`).

5.4.5 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 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 count and down count, which indicates 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.

The following is an 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
    
```

DMF Upgrade Procedures


This chapter describes how to upgrade the DMF controller and fabric switches after initial installation. It includes the following sections.

6.1 Upgrading the Controller

The default serial console baud rate on **BMF 6.1** and later hardware appliances is **115200**. It is not recommended to use the serial interface to perform an upgrade. However, if you upgrade an existing controller appliance having a serial interface set to **9600** baud, change the terminal setting to **115200** after performing the upgrade to **BMF 6.1.0** or later.


 **Note:** The upgrade process checks for supported versions when upgrading. Refer to the most recent **DANZ Monitoring Fabric 8.4.0 Release Notes** for a list of supported versions for upgrade.

The DMF controller platform supports two partitions for controller software images. The Active partition contains the Active image that is 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:** When you copy an upgrade bundle to the controller, it overwrites any older image currently in the images repository. If you try to download a bundle that is identical to the version in the images repository, the operation fails when the checksum is calculated after copying is complete.

After the upgrade image is copied to the image repository, you can use the **upgrade stage** command to copy the upgrade image to the Alternate partition and prepare the controller for the upgrade.

After the ISO image file is copied to the Alternate partition, the boot image remains in the Active partition. When you enter the **upgrade launch** command, the upgrade image is changed to the boot image and the controller reboots.

 **Note:** 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. When the upgrade process is completed, the Alternate and Active partitions have been 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, you can use the **boot partition alternate** command to restore the controller to its previous software and configuration.

Log files are written in a separate partition which is 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 cluster option, upgrade command is executed cluster-wide. To run cluster option user must be logged into the active controller.

launch: Complete the upgrade process by rebooting the controller from the alternate partition and by 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.

The following options can be used with the **launch** keyword:

- **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 upgrade.
- **pre-launch-check:** Identifies the status of the controller in regard to 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 you to create space on the local file system. After the image is copied, if the integrity of the image cannot be verified, the system displays a message and the copy is not completed. After the image is copied, a warning message is also displayed 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 DANZ Monitoring Fabric Controller](#) section.

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 remote server to all the nodes in 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 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 controller. This command prepares the platform for upgrade by populating the Alternate partition with the contents of the new image. This step is separated from the launch step, where the upgraded image becomes the Active image so the platform can be prepared 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 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 in the process, 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, along with 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 will be 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 “upgrade stage” is interrupted, it is required to stage the image again in order to launch 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, then applied to the new cluster. The new cluster requests switches from the old cluster, adjusts its local configuration to describe this new ownership, and begins rebooting the switches.

The `upgrade cluster launch` command manages the transition from the current version to the next version for both active and standby controller. This includes various steps, including rebooting all the controllers in the cluster, rebooting all of 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 Active controller via controller management IP or virtual IP and enter the `upgrade cluster launch` command:

```
controller-1# upgrade cluster launch
```

The system prompts you to proceed with the reboot. Enter **yes**. Standby controller reboots and displays the following messages, while the Active controller waits for 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
```

If you are using an SSH terminal to connect to the controller, your session is terminated. You can 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 the one 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 versions of the DANZ Monitoring Fabric and Switch Light OS.

To verify the upgrade of the controllers, login 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** 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 should be persistent across DMF upgrades in single-node mode as they are in HA mode.

Configuration

No additional configuration changes are required and keeps IPAM assigned IP configuration persistent across cluster upgrades. However, it assumes that IPAM has been previously configured.

CLI Show Commands

Use the following **show** commands to display switch configurations.

To display the switch configuration use the following command:

```
> show running-config switch
! switch
switch core1
mac 52:54:00:48:07:9f

switch delivery1
mac 52:54:00:5a:19:4d

switch filter1
mac 52:54:00:f4:15:2e
```

To display the IPAM configuration use the following command:

```
> show running-config ipam switch
! ipam
ipam switch
dns-server 172.22.22.40
gateway 10.243.248.1
ip-range 10.243.248.66 10.243.248.70 subnet-mask-length 21
allocate
```

To display the IPAM IP allocation for switches use the following command:

```
> show ipam switch
Start : 10.243.248.66
End   : 10.243.248.70
Count : 5
Used  : 3
~ Allocated IP Addresses ~
# IpSwitch
-|-----|-----|
1 10.243.248.66 core1
2 10.243.248.67 delivery1
3 10.243.248.68 filter1
```

To display the switch status use the following command:

```
> show switch
# Switch Name IP Address      State      Pipeline Mode
-|-----|-----|-----|-----|
1 core1       fe80::5054:ff:fead:7fae%6 connected 13-14-match-push-vlan
2 delivery1   fe80::5054:ff:fe15:49%6  connected 13-14-match-push-vlan
3 filter1     fe80::5054:ff:feaa:f643%6 connected 13-14-match-push-vlan
```



Note: This `show` command doesn't list IPv4 addresses but only link-local IPv6 addresses for the switches. Check the configuration of a given switch using the command below.

To display the switch running-config to show the IP address, the subnet mask length, gateway, and DNS server use the following command:

```
> show switch core1 running-config
datapath id 00:00:52:54:00:ad:7f:ae
interface ma1 ip-address 10.243.248.66/21
ip default-gateway 10.243.248.1
dns-server 172.22.22.40
```

Syslog Messages

Use the following command `show logging controller | grep MGMTIPMGR` to examine the IPAM IP allocation steps in the syslog INFO messages; they all have the log-id prefix "MGMTIPMGR":

```
> show logging controller | grep MGMTIPMGR
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1008: Reconciling [core1, delivery1, filter1]
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1006: Setting feature enabled to true
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1001: Adding IP Address range IPAddressRangeKey
[start=10.243.248.66, end=10.243.248.70, subnetMask=21]
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1009: Successfully allocated 10.243.248.68 to filter1
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1009: Successfully allocated 10.243.248.67 to delivery1
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1009: Successfully allocated 10.243.248.66 to core1
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1010: Attempting to allocate addresses for switches
[core1, delivery1, filter1]
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1008: Reconciling [core1, delivery1, filter1]
```

```
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1006: Setting feature enabled to true
INFO [org.projectfloodlight.core.ipalloc.ManagementIPManager:EventPipe
lineStage:ipam] MGMTIPMGR1001: Adding IP Address range IPAddressRangeKey
[start=10.243.248.66, end=10.243.248.70, subnetMask=21]
```

Troubleshooting

- Use the **show** commands above to understand the configuration.
- Change the corresponding configuration, such as adding or deleting an IPAM IP range, and examine the log outputs.
- Locate relevant log output in the floodlight syslog at `/var/log/floodlight/floodlight.log`
- Enable extended tracing logs to obtain more information on related events. (See additional details below.)

Limitations

- The switch IP addresses persist across upgrades. However, in case of a controller failure or simultaneous reboot (single or HA controllers), switch IP addresses are reassigned from the configured IPAM IP address range(s), which may differ from the original assignments.
- The **show switch** command doesn't display IPv4 addresses. This behavior has stayed the same but can confuse the end user. One may examine the running config of the switch to cross-check the assigned IPv4 number.

6.8 Rolling Back an Upgrade

If you decide 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 make sure that 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

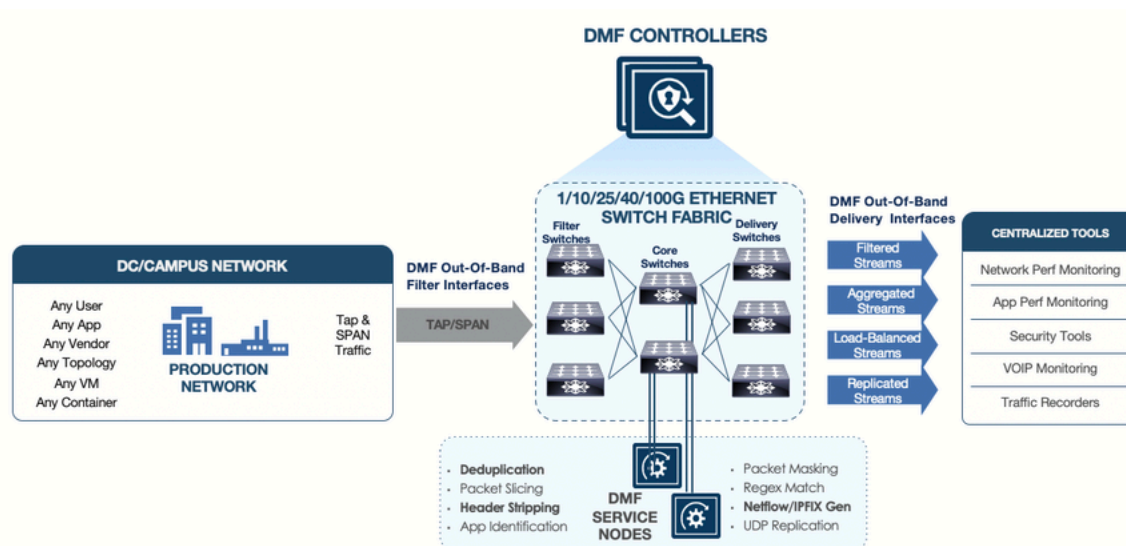
This chapter describes how to install the DANZ Monitoring Fabric DMF Service Node and includes the following sections.

7.1 Overview

The DMF Service Node provides advanced packet matching and modification capabilities that can be applied to monitored traffic. The DMF Service Node is an optional component in DANZ Monitoring Fabric when you need the advanced features it provides. The DMF Service Node provides 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



The DMF controller automatically detects each connected DMF Service Node after basic installation is complete 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


You 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 can be connected in two ways:

- Management interface (1 GbE), which is connected to the management switch. Using the management interface limits throughput to 1 GbE and shares the bandwidth with management traffic to the DMF controller.
- 10 GbE Service Node interfaces (SNI), which are connected to DMF switches. Connecting to this interface supports up to 10 GbE throughput.

Note: It is recommended to have an iDRAC connection to the DMF Controller, DMF Service Node, Arista Analytics Node and DMF Recorder Node appliances. This helps in easy troubleshooting of issues. Please refer to the chapter on “Using iDRAC” later in this guide for more details.


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 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 DMF Service Node appliance.



Note: Hyperthreading should be disabled on the hardware appliance before installing the Service Node software to avoid performance issues. If hyperthreading is disabled after the Service Node software is installed, performance will be affected and reinstalling the software will be required to resolve the issue.

To run the first boot setup, which performs the initial configuration required for a new DMF Service Node, complete the following steps.

Procedure

1. Rack the DMF Service Node Appliance.

The appliance interfaces are on the rear 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 Ethernet2) to the network management switch.
- One serial interface (db9).
- Four 10-GbE SFP ports with the R640; 16 10-GbE ports with the R740: You can connect these ports to a DMF switch.

2. Power on the DMF Service Node server appliance.

3. Log in via the serial port using the admin account name. The baud rate is **115200**. If you are using a terminal server to connect, make sure 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. http://
www.arista.com/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 an example, which you should 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 DMF Controller by a Layer 3 device (such as a router) in preconfigured (L3 ZTN) mode, enter the Active DMF controller IP address.



Note: Starting with **BMF 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
```

8. 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
```

9. 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
```

10. Connect one or more of the 10G 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 10G SFPs on the appliance hardware are connected to a DMF Out-of-Band switch.

11. 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

```
([1] 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
```

12. Login to the DMF controller.
13. Add following service node configuration with service node management interface 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: MAC address of the Service Node can be obtained by logging into service node via SSH and running the following command:

```
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

```

and note the bond0 MAC address.

14. Verify that the DMF Service Node is connected by entering the following command.

```
controller-1(config-service-node)# show managed-service-device
```

This command can be entered from any CLI mode.

7.4 Creating Support Bundle on Service Node

Support bundle for DMF Service Node should be created from DMF controller ([Creating Support Bundle](#)). But, in case of DMF Service Node losing connectivity to DMF controller, support bundle can be created on DMF Service Node by logging into DMF Service Node.

DMF Service Node CLI provides commands to automate the process of collecting, archiving, and uploading the 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)#

```

Following is the command to generate Support Bundle. Once support bundle is generated it will be automatically uploaded. Please provide support bundle id to support personal.

```

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--202
2-04-13--07-46-01Z--

```

```

SMFUG-BS5S2.tar.gz
Url : https://10.240.130.8:8443/api/v1/support/onet-support--LG-SN--2022-04-
13--07-46-01Z--SMFUG-BS5S2.tar.gz
Bundle id : SMFUG-BS5S2
Auto-uploading support onet-support--LG-SN--2022-04-13--07-46-01Z--SMFUG-BS
5S2.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 automatic upload.

```

LG-SN(config-service)# show support
# Bundle
-----|-----|-----|-----|-----|
1 onet-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 onet-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 onet-support-component--LG-SN--2021-05-19--22-47-17Z_0kc7zxc.tar.gz | | 462MB | 2021-05-19 22:47:17.452000 UTC |
LG-SN(config-service)#

```

The diagnose upload support can be used to check the reachability/health of the server to which support bundle will be up- loaded to.

```

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 DMF Service Node, deployed in L2ZTN from **BMF-7.0.0** to a later version will be automatically completed through zerotouch upgrade, if you upgrade a **BMF Release 7.0.x** controller. Upgrade of DMF Service Node deployed in L3ZTN from **BMF- 7.1.0** to later version will be automatically completed through zerotouch when you upgrade a **BMF Release 7.1.x** controller. To verify that the Service Node is ready for the zerotouch 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 how to perform the installation, initial configuration, and upgrade of the DMF Recorder Node. It includes the following sections.

8.1 Overview

The DMF Recorder is a traffic recording appliance consisting of software provided by Arista Networks, running on servers provided by Dell, Inc.

The Recorder records packets from the network to disk and recalls specific packets from disk quickly, efficiently, and at scale. The Recorder is integrated with DMF for a single-pane-of-glass. A single DMF controller can manage multiple Recorders, delivering packets for recording through Out-of-Band policies. The controller also provides central APIs for interacting with Recorders to perform packet queries.

An DMF Out-of-Band policy directs matching packets to be recorded to one or more Recorders. The out-of-band policy defines the switch and port used to attach the Recorder to the fabric. The policy treats these as “dynamic” delivery interfaces and adds them to the policy with unique names. The DMF controller also provides commands for viewing errors, warnings, statistics, and the status of connected Recorders.

The Recorder provides an OpenFlow agent that collects statistics and health information from the controller. The OpenFlow agent also allows the controller to configure the Recorder, eliminating the need for to separately administer any Recorder directly during normal operation. To the DMF controller, the OpenFlow agent causes the Recorder to appear as a special type of switch. You can use the REST API to directly query the Recorder.

The DMF Recorder Node is based on the Dell 740 server hardware and is available with the following interfaces.

- Two management interfaces (10/100/1000 Mb/s)
- One serial interface (db9)
- One VGA interface
- Two USB ports
- One dedicated iDRAC port



Note: It is recommended to have an iDRAC connection to the DMF Controller, DMF Service Node, Arista Analytics Node and DMF Recorder Node appliances. This helps in easy troubleshooting of issues. Please refer to the chapter on “Using iDRAC” later in this guide for more details.

The DMF Recorder Nodes storage capacity and the bandwidth provided by the data interfaces. DMF Recorder Node

- 192 TB packet storage capacity
- Two 25-Gb SFP ports
- Two 10Gb Copper ports

The following figure illustrates the bezel on the larger (HWA) DMF Recorder Node.

Figure 8-1: DMF Packet Recorder Node (HWA) Front Panel

Recorder-Bezel




- | | |
|--|-------------------------------------|
| 1 System identification button / indicator | 4 LCD panel |
| 2 Packet 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

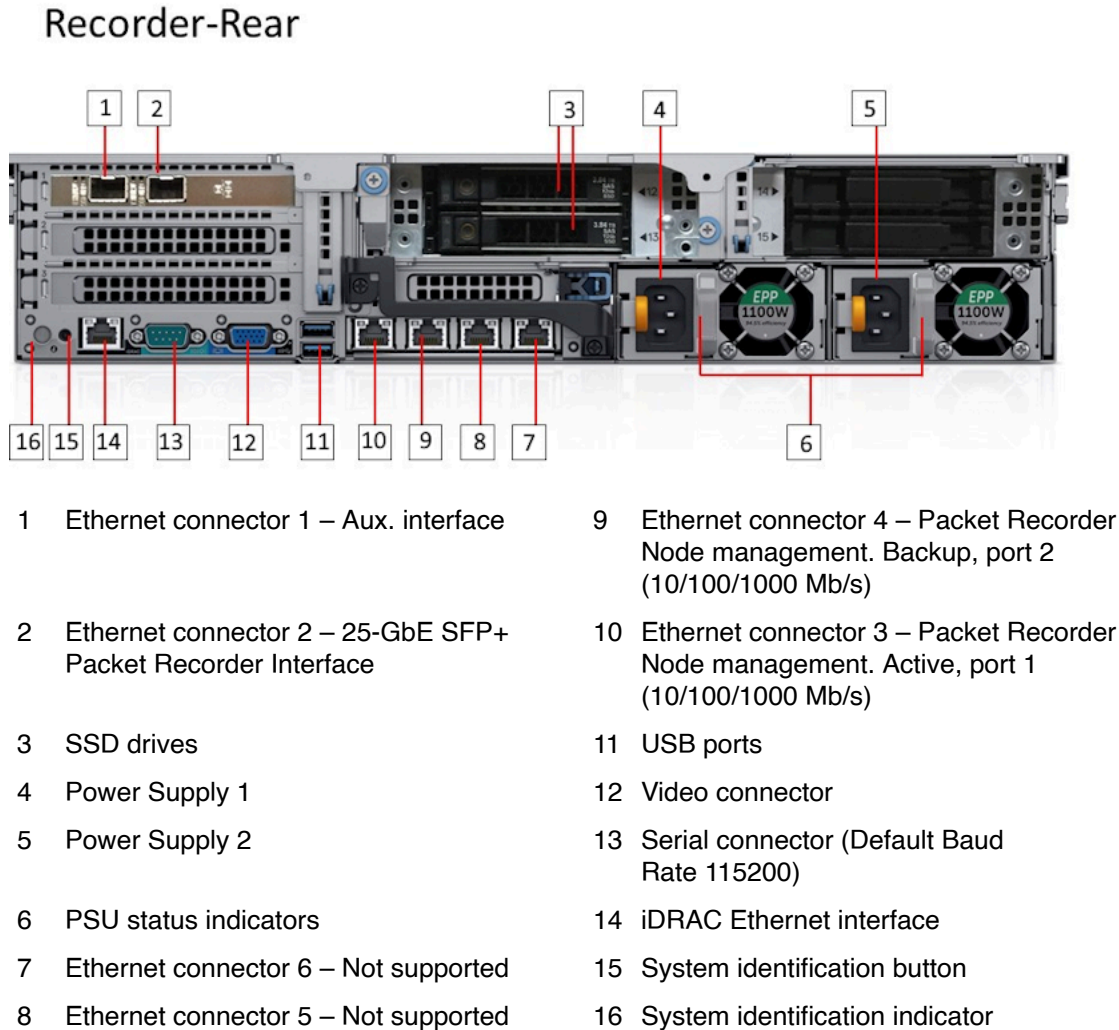
Recorder-Front



- 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

Pre-Requirement: A fresh installation of the DMF Recorder Node is required to upgrade to Release 7.1.x from earlier releases. To install the Recorder software on a Dell server, complete the following steps:

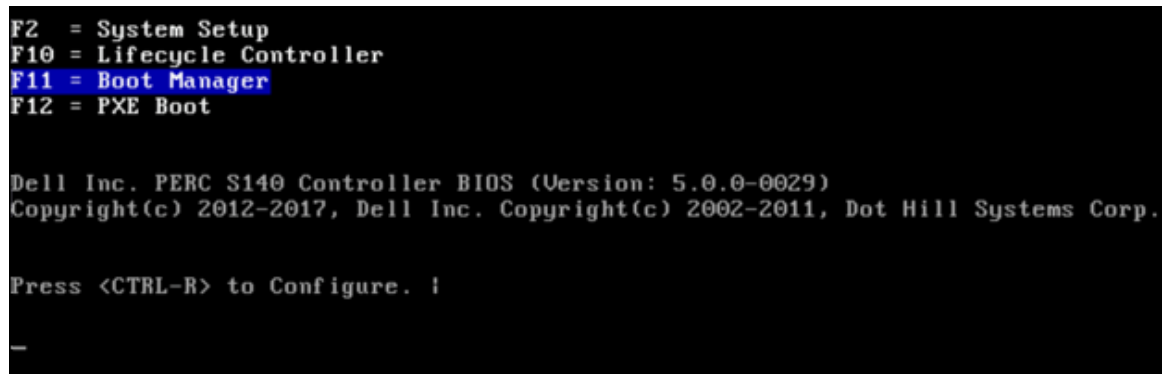
1. Rack the DMF Recorder Node Appliance.

The appliance interfaces are on the rear of the device, where the power cord is connected.
2. Connect the bottom leftmost Recorder management interface (**Gb 10**) 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 DMF Recorder Node USB port.

Refer to Appendix [Creating a USB Boot Image](#) to make a bootable USB drive.
5. Power cycle the appliance.

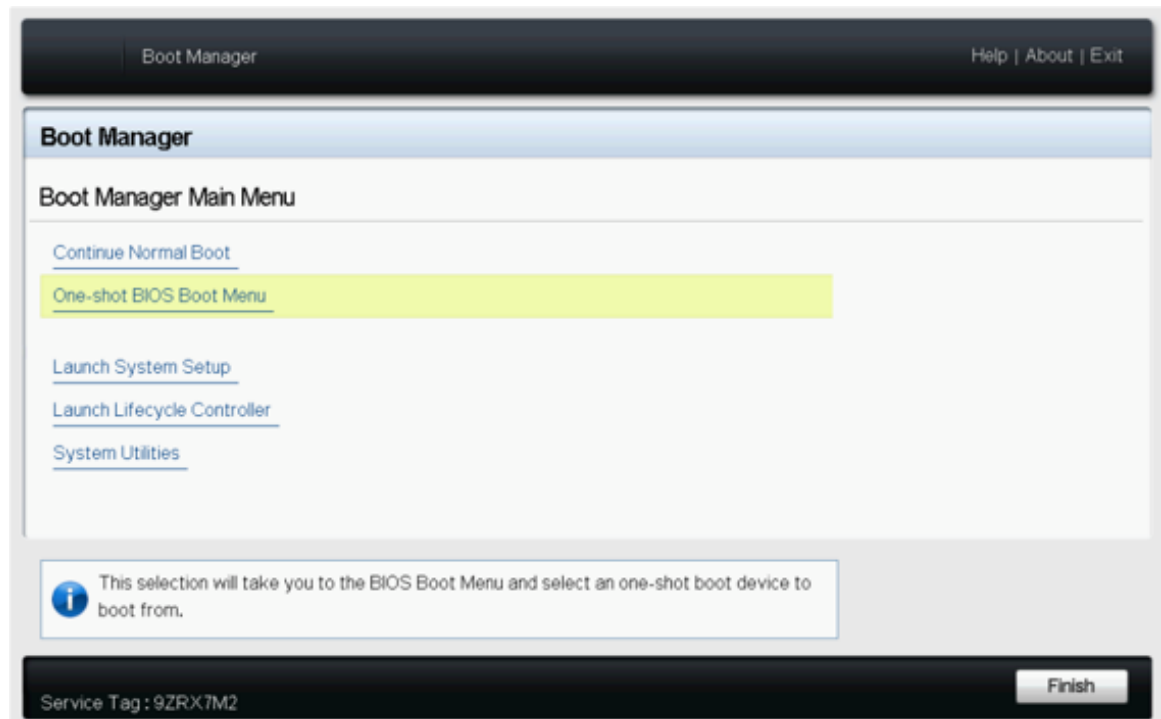
6. Press **F11** to select the Boot Manager to allow booting from USB.

Figure 8-4: System Boot Manager Screen



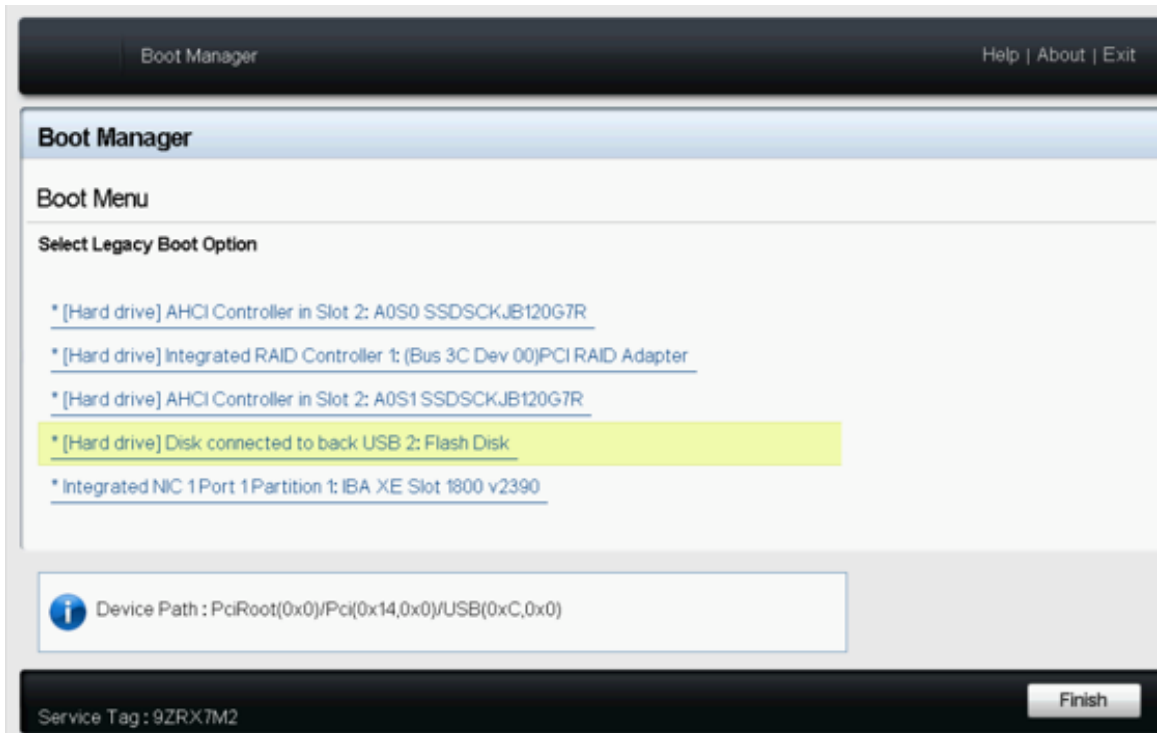
7. Select One-shot BIOS Boot Menu.
The Boot Manager screen is displayed as shown in the following figure.

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 <http://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 DMF Recorder Node is connected to DMF Controller by a Layer 3 device (such as a router) in preconfigured (L3 ZTN) mode, enter the Active DMF controller 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. To complete the configuration, press **Enter**.

8.3 Initial Configuration - GUI

After completing the installation, refer to the *DANZ Monitoring Fabric 8.4 User Guide* to configure and operate the Recorder.

GUI Procedure

To use the DMF GUI to configure the Recorder, complete the following steps.

1. Select **Monitoring > Packet Recorders** from the main menubar.

Figure 8-7: Provision Packet Recorder

The screenshot shows a web-based configuration window titled "Provision Packet Recorder". It contains the following fields and controls:

- Name ***: A text input field containing "myrecorder".
- MAC Address**: A dropdown menu showing "No Unconfigured MACs Found". Below it is a note: "Drop-down includes connected switches without a fabric role and addresses from failed ZTN requests. Choose from the drop-down or enter a new value expected to connect in the future. When a switch with the entered MAC connects, this configuration will be applied to it."
- Recording**: A toggle switch currently set to "Off".
- Disk Full Policy**: A dropdown menu currently set to "Rolling FIFO". A "Stop and Wait" option with a question mark icon is also visible.
- Pre Buffer**: A spinner control set to "minutes". Below it is a note: "Duration to record into a pre-buffer until an event occurs".
- At the bottom, there are three buttons: "RESET", "CANCEL", and "SAVE".

2. Type a name for the Recorder
3. Identify the MAC address of the Recorder appliance NIC connected to DMF.
If the MAC address has been discovered, you can choose it from the selection list.
4. Click **Save**.

- Click the **Provision control (+)** at the top of the Interfaces section.

Figure 8-8: Provision Packet Recorder

- Type an identifying name for the Recorder interface.
- Select the switch and interface to use for receiving traffic to record.

8.4 Initial Configuration - CLI

CLI Procedure

To use the DMF CLI to perform the basic Recorder configuration, complete the following steps.

- Assign a name to the Recorder:

```
(config) # packet-recorder device bt-recorder3
```

- Set the MAC address of the Recorder.

```
controller-1(config-packet-recorder) # mac 18:66:da:fb:6d:b4
```

If the management MAC is unknown, you can determine it from the chassis ID of connected devices.

```
controller-1> show connected-devices
> Switch  IF Name  DMF name SPAN?  Device Name  Device Description  Chassis ID  Port ID  Port Description  Management Address  Protocol
-----
1 bt-lb9-1  ethernet      50 False      bt-recorder3  dmf-recorder-node,SNHLZYHH2  18:66:da:fb:6d:b4  3c:fd:fe:1f:0f:82  enp                130s0f110.4.100.200  LLDP
```

- Enable the Recorder.

```
controller-1(config-packet-recorder) # record
```

4. Define the Recorder interface name.

```
controller-1(config)# packet-recorder device pr-intf-1
controller-1(config-packet-recorder)#
```

You can assign any alphanumeric identifier for the name of the Recorder interface, which changes the submode to **config-bigtap-pkt-rec**.

5. Assign a switch and interface and optionally provide a text description.

```
controller-1(config-packet-recorder)# description 'Delivery point for
packet-recorder'
controller-1(config-packet-recorder)# packet-recorder-interface switch
00:00:70:72:cf:c7:cd:7d ethernet37
```

6. Identify the Recorder interface by name in an Out-of-Band policy:

```
controller-1(config)# policy pkt-rec
controller-1(config-policy)# use-packet-recorder pr-intf-1
```

7. Configure the DMF policy to Identify the traffic to send to the Recorder.

```
controller-1(config-policy)# 1 match any
controller-1(config-policy)# filter-interface sw1-fill
```

The following example forwards all traffic received in the monitoring fabric on filter-interface **sw1-#1** to the Recorder interface. Below is an example of the packet-recorder configuration:

```
packet-recorder pr-intf-1
description 'Delivery point for packet-recorder'
packet-recorder-interface switch 00:00:70:72:cf:c7:cd:7d ethernet37

policy pkt-rec
action forward
filter-interface sw1-fill
use-packet-recorder pr-intf-1
1 match any
```

8.5 Changing the Recorder Node Default Configuration

Configuration settings are automatically downloaded to the Recorder node from the DMF controller, which eliminates the need for box-by-box configuration. However, you can override the default configuration for a Recorder node from the config-packet-recorder submode for any Recorder node.



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

To change the CLI mode to **con#g-packet-recorder**, enter the following command from config mode on the Active DMF controller.

```
controller-1(config)# packet-recorder device <instance>
```

Replace **instance** with the alias you want to use for the Recorder Node. This alias is associated with the MAC hardware address, using the **mac** command.

Use any of the following commands from **con#g-packet-recorder** 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 below configurations can be overridden to use values specific to the recorder node or can also be used in a merge-mode along with the configuration inherited from the 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 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 **conf#g-packet-recorder** 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 **conf#g-packet-recorder** 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

Tacacs configuration does not have a merge option. It can either be inherited completely from the controller or overridden to use only the recorder node specific configuration.

DHCPv4 Based Firstboot for DMF Controller and Managed Appliances

This chapter outlines a solution for provisioning DANZ Monitoring Fabric appliances via PXE and automating the configuration of the first boot parameters using Ansible.

9.1 Introduction

Typically the deployment of DANZ Monitoring Fabric 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 the completion of the firstboot configuration. Firstboot configuration uses Ansible playbook as a automation tool. In contrast, **Repave** refers to the automation of 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 the completion of the firstboot configuration.

To accomplish the above tasks there are a couple of prerequisites that need to be met. The production / lab environment should have:

- A DHCP server that supports the configuration of DHCP option 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 with Arista supported playbook modules installed.
- Supported DMF hardware appliances with preset boot settings order and with PXE enabled management port



Note: In 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 reinstallation.

9.2 Steps to Prepare your Services (TFTP/NFS) for PAVE/REPAVE Operation

The following steps need to be completed before the DMF ISO image can be deployed 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 not available on your system, you can obtain them via an **apt-get** or **yum** command.

```
$ apt-get install pxelinux
```

4. Update the TFTP_DIRECTORY variable with parent directory created to store boot loader files.

```
$ sed -i "/^TFTP_DIRECTORY=/c\TFTP_DIRECTORY=/var/lib/tftpboot" /etc/default/tftpd-hpa
```

5. Create an appropriate folder under your 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 your 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 files from the folder that was 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 ownership of tftp parent directory with TFTP user account and restart tftp service.

```
$ chown -R tftp:tftp /var/lib/tftpboot
$ systemctl restart tftpd-hpa
```

9. Copy kernel configurations that assist the DMF appliance that is booting with UEFI mode, to locate **bootloader** and **vmlinuz** files from TFTP and NFS server. It is recommended to define a name for menu entry so that it easily remembered and prefixed with appliance type for differentiation i.e., *DCA-DM-CDL-dmf-8.4.0*.



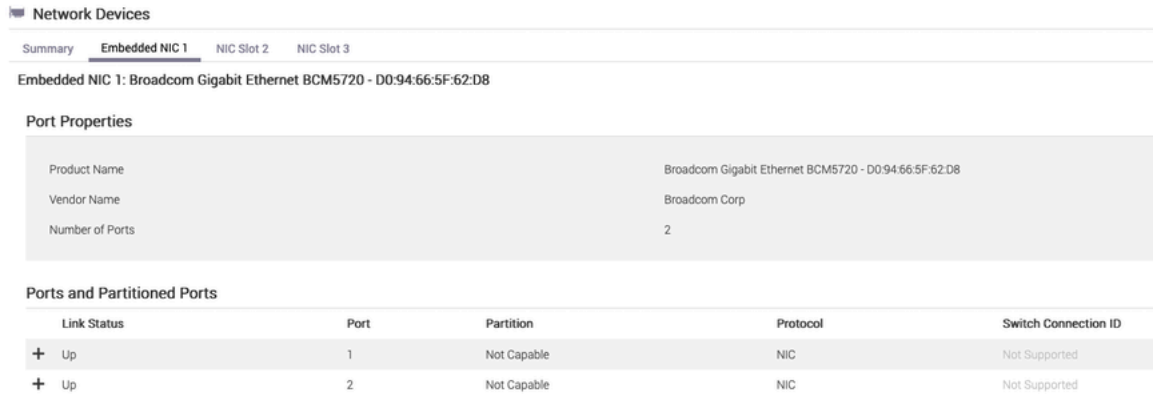
Note: In the currently supported UEFI boot mode, it is recommended to create a PXE configuration file for each hardware mac address i.e., instead of the default **grub.cfg** use **grub.cfg-01-xx-xx-xx-xx-xx-xx**. The MAC address should be specified in all lower-case with : replaced by -.

To do the aforementioned, the user will need to obtain the MAC / Hardware address of the management interfaces. This can be done in two ways as described below:

- a. If this is the first time the appliance is being installed, the user can either:

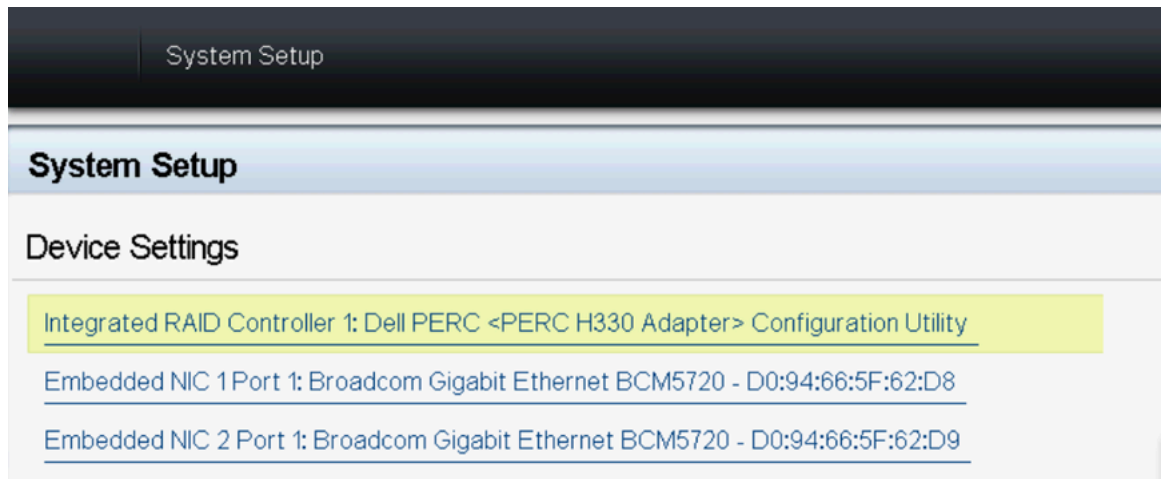
- Use the Integrated Dell Remote Access Controller (iDRAC) web console to determine the MAC address of the management interface.

Figure 9-1: Using the iDRAC menu to obtain the Management Interface MAC Address



- 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. If the user intends to do re-installation of DMF appliances, the MAC address of the appliance can be obtained via the CLI command.
 - In case of DMF Controllers with SKU DCA-DM-C450 or DMF Analytics Nodes with SKU DCA-DM-AN450 execute the following two commands, on the controller, to obtain the relevant MAC address.

```
DCA-DM-C450# show local-node interfaces eno8303
-----
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
-----
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#
```

- In case of DMF Controllers with SKU DCA-DM-CDL and all DMF Recorder Nodes execute the following two commands, on the device, to obtain the relevant MAC address.

```
DCA-DM-CDL# show local-node interfaces eno1
-----
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       backup
~ Address ~
None.
DCA-DM-CDL#

```

- For all DMF Service Nodes, execute the following two 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         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
~ Address ~
None.
dmf-service-node#

```

- If both the management ports of the appliance are connected to the top of the rack management switch, the user is recommended to generate a separate PXE configuration file for both the management ports.

```

cat << EOF | sudo tee /var/lib/tftpboot/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
menuentry "Install <INSERT_USER_FRIENDLY_NAME>" {
linux dmf-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 dmf-controller/initrd.lz
}
EOF

```

9.3 DHCPv4 based firstboot for DMF Controller and Managed Appliances

There are two main steps involved in accomplishing automatic installation of an image on a supported DMF hardware appliance and the completion of the firstboot configuration. Auto installation of images using PXE and auto configuration of firstboot parameters to complete DMF HW appliance installation.

- **Auto-installation of Images** - Auto installation of DMF image uses well known services like DHCP, TFTP, NFS and PXE. DMF images are PXE bootable and using the aforementioned services, we can accomplish auto-installation of images on the DMF appliances.

A high level procedure for auto-installation is given below.

- User configures DHCP server to provide DHCP IP address:

- Next-server IP address (This is the TFTP server IP address specified by **Option 66** configuration on DHCP server).

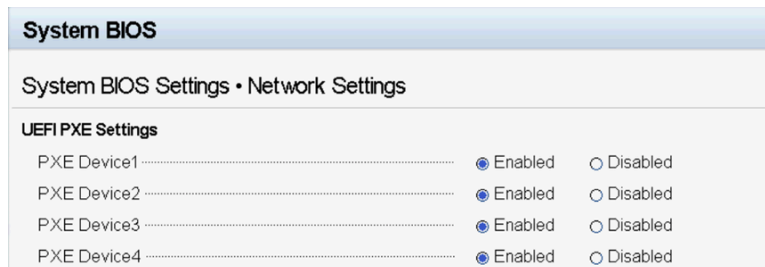
```
next-server <TFTP_SERVER_IP>;
class "pxeclient" {
  match if substring (option vendor-class-identifier, 0, 9) =
    "PXEClient";
  filename "boot/grub/x86_64-efi/core.efi"; # x86 EFI
}
```

- It is recommended to bind a static IP to the hardware MAC address of the DMF appliance. Sample configuration using linux based DHCP service is shown below.

```
group {
  host <DMF_APPLIANCE_HOSTNAME> {
    hardware ethernet <DMF_APPLIANCE_MAC>;
    fixed-address <DESIRED_IP_ADDRESS>;
  }
}
```

- On a DMF HW appliance where you want 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 PXE boot. This is not required when the user intends to reinstall (REPAVE) the DMF appliance images.

Figure 9-4: Manual PXE Boot by pressing F12



- Power cycle the DMF HW appliance.
- DHCP client on NIC card sends DHCP discover and request and gets IP address and next-server (TFTP server) IP address.
- Then 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 also gets other configuration parameters like NFS mount where appliance ISO images are stored.
- It boots from the boot loader and then gets the DMF appliance ISO image from the NFS server.
- The appliance then boots the ISO image and starts the installation process without needing user input.

- On reboot, the appliance again acquires a DHCP based management IP address, sets up initial default user credentials and waits for auto firstboot 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.

- **Auto Configuration of Firstboot** - Auto configuring of firstboot parameters uses Ansible. This is accomplished by interacting with auto-firstboot-cloud-plugin available in hardware appliances.
- Customer initiates initial configuration playbook in Ansible. A sample YAML based anisble playbook file is shown below:

```
- 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 will get the JSON based configuration file from the server where the playbook is being executed. The JSON file should have specific sections for each DMF appliance that is being installed/re-installed. An example is shown below.

```
{
  "<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 initial-config-password is used by the ansible script to access the DMF appliance via SSH for the first time.

- Ansible pushes initial configuration data to appliance via SSH, using default credential configured in the JSON file in **Step 1** above.
- Appliance completes initial configuration (firstboot).
- Appliance confirms to Ansible that initial configuration is complete. An example run of the ansible-playbook is shown below.

An example run of the ansible-playbook is shown below.

```
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 must be made in order to establish the trust required to authenticate the appliances. This section provides a summary of the assumptions which underpin the security this design provides.

- 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 DMF controller. It includes the following sections.

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. The information managed by an SNMP agent is organized as a collection of objects called MIBs.

In SNMPv3, an agent (SNMP server) is identified by an engineID, which helps prevent unauthorized SNMPv3 messages, such as traps, from being accepted or traps being 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.

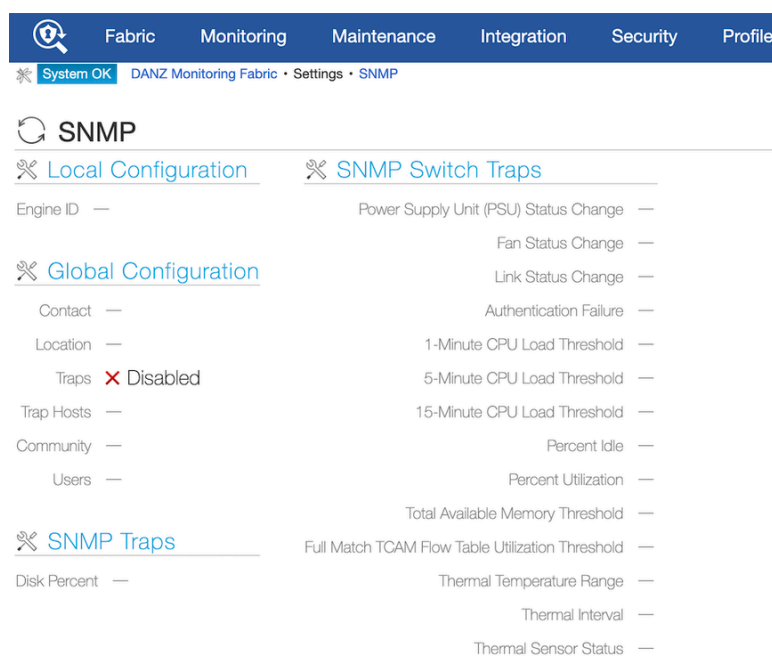
In DMF, the engineID is autogenerated for the controller and fabric switches. The engineID of the DMF controller is configured for the local node and this configuration must be entered separately on the Active and Standby controllers. It is recommended to configure a different engineID for each controller.

10.2 Using the DMF GUI to Configure SNMP

To manage or view the controller SNMP configuration, complete the following steps:

1. Select **Maintenance > SNMP** from the main menu.

Figure 10-1: Configuring SNMP



By default, the SNMP server is disabled.

- To enable access to SNMP for the controller, click the **Settings** control in the Configuration section.

Figure 10-2: Configuring SNMP Settings (Page 1 of 4)

Configure SNMP

- 1. Settings ✓ Contact
- 2. Communities ✓ myadmin
- 3. Hosts ✓ Location
- 4. Users ✓ 000 mystreet, anywhere

Traps
Disabled Enabled

BACK NEXT RESET CANCEL SUBMIT

- To enable SNMP traps, move the slider to **Enabled** in the **Traps** section.
- Type the appropriate values for the **Contact**, **Location**, and **Community** that helps identify the controller during SNMP communications.
- Click **Next**.

Figure 10-3: Configuring SNMP Settings

Configure SNMP

- 1. Settings ✓
- 2. Communities ✓
- 3. Hosts ✓
- 4. Users ✓

<input type="checkbox"/>	Secret	Permission
<input type="checkbox"/>	02161159070f0c	RO

Show: 10 25 100 All

BACK NEXT RESET CANCEL SUBMIT

- To apply one or more SNMP community strings, enable the checkbox for the entry.

- To define a new community string, click the **Provision control (+)** at the top of the table.

Figure 10-4: SNMP Community

Configure SNMP Community

Secret *

mysecret

Secret will be encoded

Permission *

Read-Only Read-Write

CANCEL APPEND

- Type the **community string** in the Secret field.
- Set the permission to **Read-Write** if required by sliding the Permission slider to the Right.
- Click **Append**.
- Enable the checkbox of the newly added community string and click **Next**.

Figure 10-5: Identifying the SNMP Trap Receiver

Configure SNMP

1. Settings ✓

2. Communities ✓

3. Hosts ✓

4. Users ✓

Server UDP Port

No SNMP trap hosts

BACK NEXT RESET CANCEL SUBMIT

- To add a trap receiver, enable the checkbox for an existing entry on the table.

- To add a new entry, click the **Provision control (+)**.

Figure 10-6: Configuring SNMP Trap Host

Configure SNMP Trap Host

Server *
my-trap-receiver

UDP Port *
162

RESET CANCEL APPEND

- Type the IP address and port number for the NMS or other SNMP client to which the controller should send SNMP trap messages and click **Append**.



Note: *UDP port 162* is the default for SNMP trap messages; *UDP port 161* is the default port for general SNMP messages.

- Enable the checkbox for the newly added trap receiver and click **Submit**.
- To create an SNMPv3 user, click **Next** to display page 4 of the dialog.

Figure 10-7: Configuring SNMPv3 User

Configure SNMP

1. Settings ✓
2. Communities ✓
3. Hosts ✓
4. Users ✓

<input type="checkbox"/>	Name ▲	Auth Passphrase	Private Passphrase	Private Protocol
No users				

BACK NEXT RESET CANCEL SUBMIT

- Enable the **checkbox** to add an existing user.

- To define a new SNMPv3 user, click **Provision** control at the top of the table.

Figure 10-8: Configuring SNMPv3 User

- Type the **name** of the user, the **authentication passphrase** for the user, and the **Private Passphrase** for encrypting messages.



Note: You can use the Private Protocol option to perform Advanced Encryption Standard (AES) or Data Encryption Standard (DES) encryption to encrypt the SNMP messages between the SNMP agent and the manager.

- Click **Append**.
- Enable the checkbox for each user and click **Submit**.

10.2.1 Configuring SNMP Traps

To configure the SNMP traps sent to the trap host, complete the following steps:

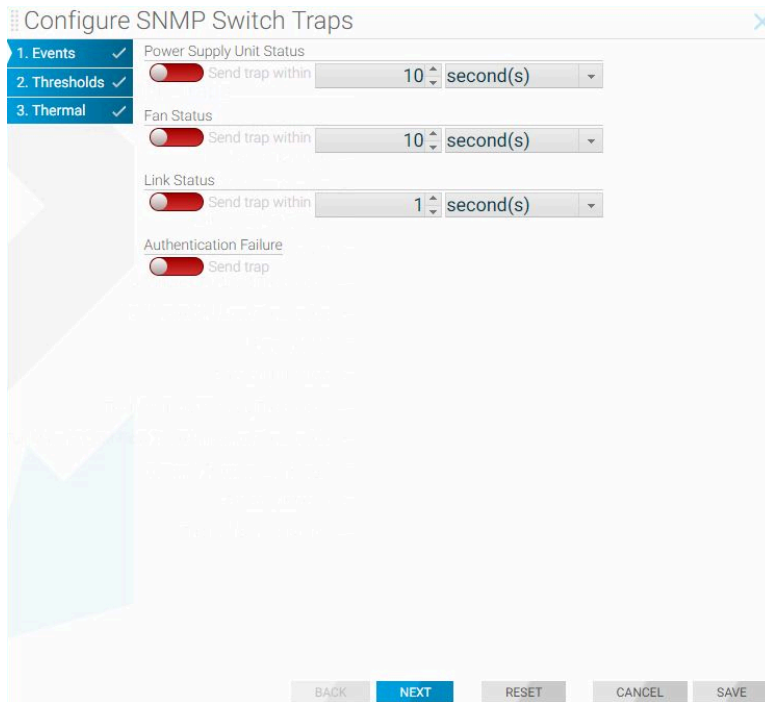
- Select **Maintenance > SNMP** and on the SNMP landing page, click the Settings control to the left of the SNMP Traps section.

Figure 10-9: Configure SNMP Traps

- Change the percentage of disk utilization or disable the trap by moving the slider to the left. Configuring the disk-percent trap enables monitoring the size of `/var/log` and the root partitions.
- Click **Submit**.

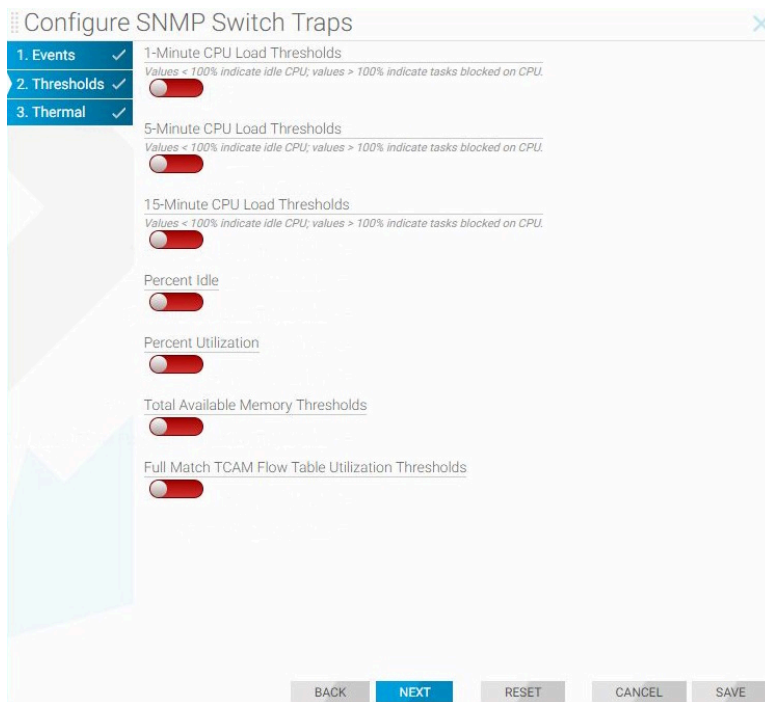
- To enable or disable specific traps, click the **Settings** control in the **Switch Traps** section.

Figure 10-10: Setting Threshold Traps: Events



- Enable or disable the event traps on this page, as required.
For the status traps, select the interval after which to send the trap.
- After making any changes required, click **Next**.

Figure 10-11: Setting Threshold Traps: Thresholds



This page lets you configure threshold-based traps.

7. Enable or disable any of these traps, and set the percentages, as required. For the Total Available Memory Threshold, select the units (bytes, KB, MB, or GB), and specify the number of units.
8. Click **Next**.

Figure 10-12: Thermal Traps

The screenshot shows a configuration window titled "Configure SNMP Switch Traps". On the left, there is a sidebar with three tabs: "1. Events", "2. Threshholds", and "3. Thermal", each with a checkmark. The "Thermal" tab is active. The main area is titled "Threshholds and Interval" and contains a red toggle switch that is turned on. Below the toggle, there is an "Interval" field set to "10" with a unit dropdown menu set to "second(s)". There are also "Min" and "Max" fields, both with a unit dropdown menu set to "°C". Below these fields is a "Sensor Status" dropdown menu set to "None". At the bottom of the window, there are five buttons: "BACK", "NEXT", "RESET", "CANCEL", and "SAVE".

This dialog lets you enable and configure the minimum and maximum temperature for thermal traps, and the interval between sending traps. You can also enable a trap to be sent when the sensor status is failed, missing, good, or all.

9. After completing the SNMP trap configuration, click **Save**.

10.3 Using the CLI to Configure SNMP

This section describes how to use 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 access list for SNMP is empty, which means that access is not permitted unless specifically enabled.

The following commands enable access to the controller by remote SNMP clients on the specified subnetwork:

```
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, from 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
```

Starting in **DMF 7.3 release**, configuring **disk-percent** trap will monitor root partition in addition to **/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 `/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, currently DANZ Monitoring Fabric supports only the `ro` option.

```
controller-1(config)# snmp-server community ro <string>
```

This sets the community string for read-only access to the SNMP trap server.



Note: For SNMP trap host configuration to be pushed to the monitoring switches, the community string or the SNMPv3 user must be configured on the controller.

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 is displayed in the running-config as a Type 7 encoded value.

To monitor controller's `/var/log` and `root` partitions, configure the following trap:

- `disk-percent percent`: Replace **percent** with the percentage that triggers a trap when it is exceeded.



Note: Configuring the `disk-percent` trap on the Analytics Node will monitor the `/var/lib/analytics/data` folder in addition to the `/var/log` folder and the `root` partition.

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>]}
```

The following keywords can be used with the `snmp-server switch trap` command.

- `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 it is 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 it is exceeded.
- **percent-utilization percent**: Replace **percent** with the with the percentage of CPU utilization that triggers a trap when it is 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 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 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). **BMF Release 6.1.0** introduced support for the SNMPv3 user-based security model (USM) for message security through authentication and encryption.

In SNMPv3, an agent (SNMP server) is identified by an engineID, which helps prevent unauthorized SNMPv3 messages, such as traps, from being accepted or traps being 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 autogenerated 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 **conf#g-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 and this configuration must be entered separately on the Active and Standby controllers. It is recommended to configure 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
61626364656662d2d670a
```

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 configuration gets pushed through zero touch networking (ZTN) to the connected fabric switches.



Note: Currently DANZ Monitoring Fabric supports only the `ro` or `read-only` type of community string option.

10.3.9 SNMPv3 Command Examples

In the following example the `snmp_1` user is configured for authentication (authNoPriv) with the plaintext password `authauth1`.

```
controller-1(config)# snmp-server user snmp_1 auth authauth1
```

In the following example, 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
```

In the following example, 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
```

In the following example, 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 you can configure SNMP for a specific switch, you must 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.

- On the Switches page, click the **Settings** control for **Configure Switch**.

Figure 10-13: Configure Switch Dialog

The screenshot shows the 'Configure Switch' dialog for switch DMF-D1. The left sidebar lists configuration steps: 1. Info, 2. Clock, 3. SNMP, 4. SNMP Traps, 5. Logging, 6. TACACS, 7. sFlow, and 8. LAG Enhanced Hash. The main configuration area includes:

- Name ***: DMF-D1
- MAC Address**: 70:72:cf:c1:7c:7b (with a drop-down menu and explanatory text: '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.')
- Description**: (empty text box)
- Admin Status ***: Down (toggle switch)
- Change Password**: No (toggle switch)
- Management Interface ***: Prefer Dedicated Management Interface (dropdown menu with explanatory text: '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.')

 A bottom sidebar contains actions: Configure Switch, Clear Switch Config, Create Tunnel, Manage Firmware, Provision Interfaces, Reset Controller, Show Config, Network, and Add Backup.

This page lets you merge/override the default configuration pushed from the DMF controller with switch specific SNMP configuration.

- To merge/override the SNMP configuration, click the **3. SNMP** link. Choose from the **SNMP Settings** drop down to either **Merge with Global Config** or **Override Global Config**.
- Make any changes required to the specific switch configuration and click **Next** if you want to customize the SNMP traps, or click **Submit** if you are done.
- To merge/override the configuration for SNMP traps, click the **4. SNMP Traps** link and choose from the **SNMP Switch Trap Settings** drop down to either **Merge with Global Config** or **Override Global Config**.
- 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 you can enter SNMP commands from the config-switch submode, you must 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 very similar to configuring SNMP in config mode at the controller level.
- When you enter the `snmp-server enable traps` command in config mode, this pushes snmp-server enable configuration to each connected fabric switch. You can 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, CLI can also be used to merge/override the default SNMP configuration with switch specific SNMP configuration. To do so, complete the following steps:

1. Add SNMP configuration at the controller. This is the default SNMP configuration that is pushed to all the switches. An example configuration is shown below:

```
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. Steps shown below:

```
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 global config with switch specific config or **override-global** to override the global config with the switch config. If neither is chosen, 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 config 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
```

It is seen that with 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 configs 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
```

It is seen that with `override-global`, the effective configuration on the switch is only the switch specific configuration. The default configuration inherited from the controller is completely overridden.

5. Configuring SNMP traps using `merge` and `override global` commands are 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 with `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 with `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>}
```

- Using the `merge-global` and `override-global` commands at the `config-switch` submode, the SNMP community for the switch can be changed. An example configuration is shown below:

SNMP config at 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 config at 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 HW failure happens like PSU failure/removal. SNMP clear trap message is sent whenever threshold is less than user specified range or HW failure is fixed such as PSU starts working.

There is no command to enable this feature. This feature is automatically enabled when SNMP trap is configured in the controller.

SNMP traps which 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. All SNMP traps and clear trap settings are listed under `/etc/snmp/snmpd.conf` file.



Note: SNMP clear traps will be sent without any prior associated SNMP traps when system comes up or there is any SNMP config 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: Appliance IDARC Firmware version should be upgraded to recommended version **5.10.50.00** or later.

- disk-percent
- memtotalfree
- lowmemavailable
- cpuload
- cputemp
- cpu1temp
- ambienttemp
- exhausttemp
- powersupply
- fanspeed
- psuCount
- fanCount

Number of fans on appliance vary. Depending on number of fans on the appliance fanspeed clear traps are sent. Fan speed traps are named ***fan1Aspeed, fan1Bspeed*** etc.

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 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 watermark of 90% on a link managed by DMF.
- A drop in traffic saturation levels to below the low watermark of 70% if the link was previously saturated.



Note: This feature is compatible with all hardware platforms compatible with DMF 8.4. 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.e
thernet2.TX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.2 = STRING: "00:00:5c:54:00:f2:81:03.e
thernet0.RX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.3 = STRING: "00:00:12:b0:03:22:07:a2.e
thernet1.RX"
iso.3.6.1.4.1.37538.68.77.70.1.5.1.4 = STRING: "00:00:50:ef:0c:a2:16:f6.e
thernet3.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.e
thernet2"
iso.3.6.1.4.1.37538.68.77.70.2.3.1.2 = STRING: "00:00:52:a4:31:f2:81:56.e
thernet0"
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, are 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` do not result in responses or return errors.

Considerations

- The link saturation trap thresholds cannot be specified by the user. 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 on 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 DMF controller using local groups and users (RBAC), AAA servers in general, and it also provides specific configuration for TACACS+ or RADIUS servers.

11.1 Overview

Access privileges to the DMF controller are associated with groups. Each user assigned to the group obtains the access permissions associated with the group. In the current version of DMF, the following two groups are supported:

- **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. By default, all authentication, authorization, and accounting functions are set to local. The general options that control where these functions occur are:

- **Local only:** The accounts on the remote AAA server are ignored; only the local database accounts are used.
- **Local primary and remote backup:** The accounts on the local database are used first. If the account is not found locally, the database on the remote server is used.
- **Remote only:** Authentication and authorization does not fall back to local for users other than the admin and recovery accounts, no matter what happens.
- **Remote primary and local backup:** The accounts on the remote AAA database are used first. If the remote server is not available, the accounts in the local database are used.



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 not available.

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 as a result of the remote AAA server being 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

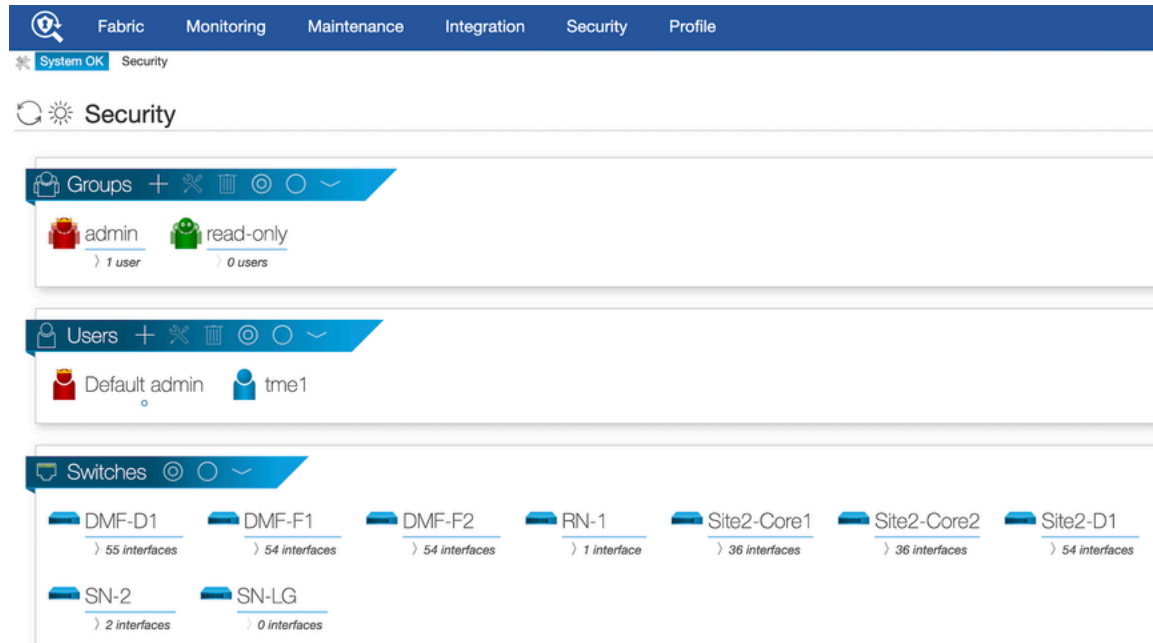
Administrative access to the monitoring fabric is managed by assigning interfaces to groups and then assigning user accounts to the group, which can then view and use the assigned interfaces.

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 DMF controller.

To use the GUI to manage groups and users, select **Security** from the DMF GUI main menu. The system displays the Security page.

Figure 11-1: Security



This page shows the following default groups that are preconfigured 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 for viewing and monitoring 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.

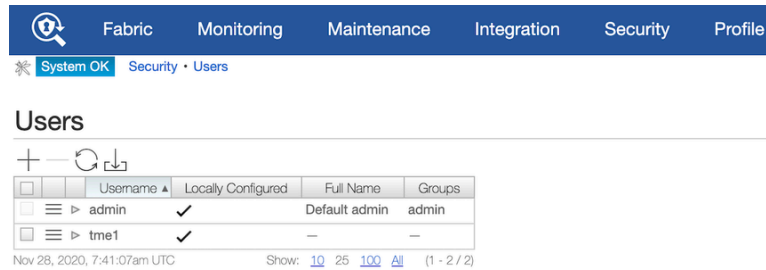
You can define a new group, in which case, a user added to the group will be a restricted user, with privileges to view and use only the interfaces assigned to the group.

11.2.2 Using the GUI to Manage User Accounts

To add a user to a group, complete the following steps.

You can add users from the main Security page, or from the following page that appears when you select **Security > Users** from the DMF GUI main menu.

Figure 11-2: Security > Users



1. Click the **Provision control (+)** at the top of the Users section, or at the top of the table displayed when you select **Security > Users** from the DMF GUI main menu.
The system displays the Add User dialog.
2. After completing this page, click **Next**.
The system displays Page 2 (Groups).
The Groups page provides a list of existing groups to which you can add the current user.
3. To add a group to this list, click the **Provision control (+)** in the upper left corner.
The system displays the Manage Groups dialog.
4. To create a new group, click the **Provision control (+)** at the upper left corner of the Manage Groups dialog.
5. Type a name for the group and click **Submit**.
The new group is added to the table.
6. Enable the checkbox for the new group and click **Append Selected**.
The group is added to the Create User dialog.
7. Enable the checkbox for the new group and click **Save**.

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 `conf#g-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-creato
r> [created-on<date-and-time>]]
```

To display the currently configured groups and users, enter the **show group** command.

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
```

To delete a user, enter the **no user** command from config mode, 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 **Change Password** from the Menu control to the left of the user on the **Security > Users** page.

Figure 11-3: Change Password

Type the new password and confirm it and click **Save**.

11.2.5 Password Reset

Resetting Password for recovery User

To reset password for *recovery* user, please follow one of the following procedures. The steps need to be performed on both the controllers of the cluster as resetting the password of *recovery* user on one controller won't change it for the *recovery* user on the other controller.

1. Using Controller's bash
 - a. Go to 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 recovery account login.

For this to work, customer needs to know the current password for *recovery* user.

```
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`:

Below API call will reset the *recovery* user's password to *AdminAdmin*. Example is given below is using `curl` initiated from a Linux host but any rest client can be used to 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 password for *admin* and other local users, login to controller using *recovery* user credentials. Use `floodlight-reset-password` to reset the user's password.

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:~$
```

Below example resets password for *guest* who is 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 Managing Groups

You can manage groups from the Security page or from the following page that is displayed when you select **Security > Groups** from the DMF GUI main menu.

Figure 11-4: Security Groups

The screenshot shows the DMF GUI interface for managing security groups. At the top, there is a navigation bar with tabs for Fabric, Monitoring, Maintenance, Integration, Security (selected), and Profile. Below the navigation bar, there is a status indicator 'System OK' and a breadcrumb 'Security > Groups'. The main content area is titled 'Groups' and contains a table with the following data:

	Name ▲	# Users	# Unique Users	# Shared Users	# Interfaces	# Unique Interfaces	# Shared Interfaces
<input type="checkbox"/>	admin	1	1	0	—	—	—
<input type="checkbox"/>	read-only	0	0	0	—	—	—

At the bottom of the screenshot, there is a timestamp 'Nov 28, 2020, 7:49:41am UTC' and a 'Show:' dropdown menu with options '10', '25', '100', and 'All', followed by '(1 - 2 / 2)'.

To view the interfaces and users assigned to a group, click the **Expansion** control to the left of the group. To add a user to the group, click the **Provision control (+)** at the upper left corner of the Users table.

To associate an interface with the group, complete the following steps.

1. Click the **Provision control (+)** at the upper left corner of the Switch Interfaces table.
The system displays the **Add Interfaces to Groups** page.
2. To add an interface to this table, click **Provision control (+)** at the upper left corner of the table.
The system displays the **Append Interfaces** page.
3. Enable the checkbox for each interface you want to add to the group and click **Append Selected**.
The interfaces are added to the **Add Interfaces to Group** page.
4. Enable the checkbox for each interface you want to add to the group and click **Save**.
The interfaces are now associated with the group and any users in the group will have privileges to view and use the interfaces.

11.2.7 Authentication with a User Token and REST API

You can use the `access-token` command to create a long-lived token that can be used for authentication with external scripting, such as using RESTful API. The token can be deleted (repudiated) at any time. A hashed version of the token is preserved 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 that can be used 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 459d1d5e0bc42c5bbbee091a984e2807592d11d8a5d4cbac110d
32da3f0e436c 2018-01-
08T15:44:02.353Z
hashed-password
method=PBKDF2WithHmacSHA512,salt=N9E9jaPIFTc_ZO9Oq0gd4A,rounds=25000,ph=true,
7eCf1PGAYqUw53vJ2bsGpSVEU5-D6ix5sHufFUi3gFr3AHB4Jqj2eLNZzoo66y_qIRFOOOL8nc5
oG5ilgJlFxA
controller-1(config-user)#
```

The session cookie does not appear in the running-config, instead the hashed session cookie 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, which can be implemented in different ways, such as with LDAP, RADIUS, Kerberos, TACACS+, or a local database. The current version of DMF supports AAA using a local database on the controller node, or with a remote TACACS+ or RADIUS server. By default, all authentication, authorization, and accounting functions are set to local.

Authentication and authorization occur only once, and whatever privileges are associated with the first account found are used. You can configure options separately for authentication and authorization, but it is recommended to use 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 method of authorization. The local database is then used only 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 not available.



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, which you can choose from the Default Group selection list. The options are admin or read-only, when the latter group has been configured.

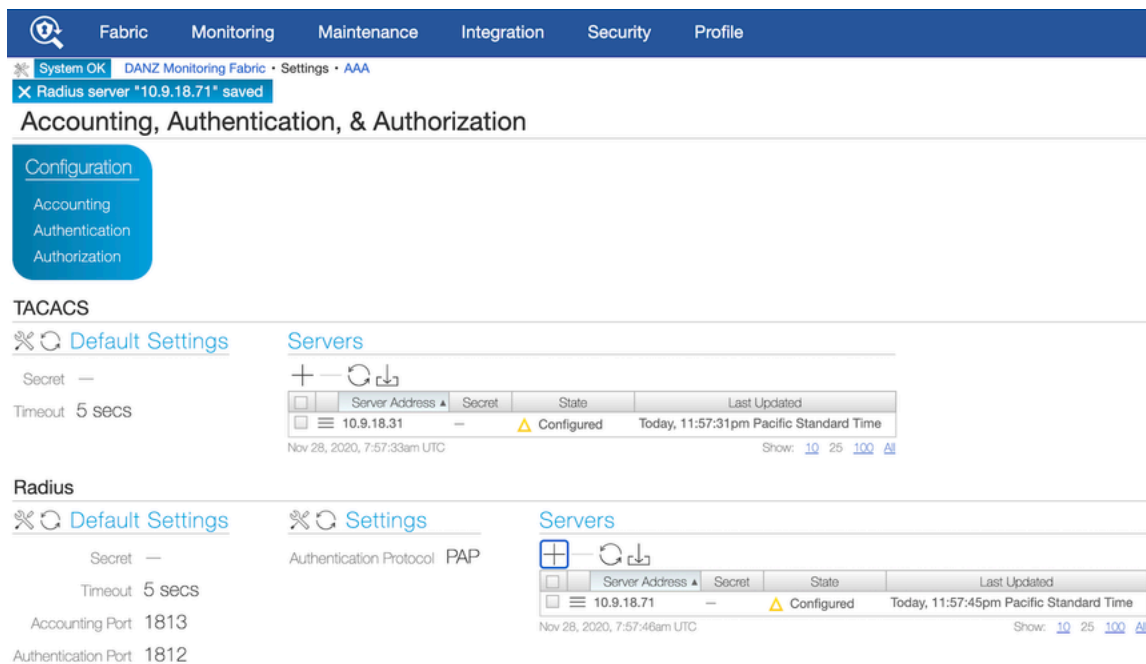
DMF supports AAA services using both RADIUS and TACACS+ servers. Up to four servers of each type can be used. The default timeout for each AAA server is **#ve** seconds, which has been changed from **30** seconds in previous releases. The total aggregate timeout for all configured servers is a maximum of **20** seconds.

By default, the controller waits **#ve** seconds before trying the next server, or before falling back to 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 **#ve** seconds or less.

GUI Procedure

To manage how AAA settings are configured on the DMF controller, select **Maintenance > AAA**. The system displays the page shown in the following dialog.

Figure 11-5: AAA Settings Page




This page lets you identify up to five destination TACACS or RADIUS servers, configure the server secret, and specify the timeout.

11.3.1 Enabling Remote AAA Services

GUI Procedure

To log accounting log messages to a remote server, click the **Accounting** link in the Configuration section.

Figure 11-6: Accounting Settings



By default, accounting logs are saved on the local controller, where they are available for searching and analysis using the **Visibility > Analytics** option. If local logging is disabled, these logs will not be available for analysis locally.

To send the accounting logs to a remote server move the slider to **No** and select **RADIUS** or **TACACS** from the selection list. To use a remote server for authentication, click the **Authentication** link in the Configuration section.

Figure 11-7: Authentication Settings

Configuration

Accounting

Authentication

Authorization

Authentication Sources

Set prioritized list of authentication stores. At least 1 source must be set.

Remote 1st

Local 2nd

Remote Type *

TACACS

Select the primary authentication method from the 1st selection list, and select the secondary authentication method from the 2nd selection list. Then select **RADIUS** or **TACACS** from **Remote Type** selection list and click **Save**.

To use a remote server for authorization, click the **Authorization** link in the Configuration section.

Figure 11-8: Authorization Settings

Configuration

Accounting

Authentication

Authorization

Authorization Sources

Set prioritized list of authorization stores. At least 1 source must be set.

Remote 1st

Local 2nd

Remote Type *

TACACS

Default Group

- Group -

CLI Procedure

Use the `aaa authorization role default` command to assign a default role to the current account. The default group is used 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 **DMF 8.0 release**, DANZ Monitoring Fabric supports time-based user lockout functionality. Users will be locked out of login for '**t2**' time when attempting with '**n**' incorrect passwords within '**t1**' time.

Locked out users have to be cleared of lockout or they have to wait for the lockout period to expire before attempting login with 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>
```

- Value range for 'failure' can be from **1** to **255**.
- Value range for 'window' and 'duration' can be from **1** to **4294967295** seconds ($2^{32}-1$).

Below example will lock 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, password-less authentications such as ssh-keys, Single Sign-on and access-token are not affected. Locked out users can still access controller via console/password-less authentication.
- The feature is node-specific with respect to the functionality. i.e.: if user1 is locked out accessing Active controller in the cluster, they would be able to login to Standby controller with 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,

```
clear aaa authentication lockout
```

The following example shows how to clear the "admin" user who got locked out.

```
controller-1# clear aaa authentication lockout user admin  
controller-1# show aaa authentication lockout  
None.
```

'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: 'window' parameter does not apply to 'recovery' user login as the '*pam_tally2*' tool does not support the same.

11.5 Using TACACS+ to Control Access to the DMF Controller

You can use remote Authentication, Authorization, and Accounting (AAA) services using 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 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 is 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 compared to remote authentication and authorization.

Privileges at the remote TACACS+ server must be configured using the following attribute-value pairs:

- Supported attribute name: BSN-User-Role
- Supported attribute values: admin, read-only,

You can use a TACACS+ server to maintain administrative access control instead of using the controller local database, although it is a best practice to maintain the local database as the secondary method of authentication and authorization 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: To use the same user credentials to access ANET and non-ANET devices, the BSN-User-Role attribute must be specified as “Optional,” in the `tac_plus.conf` file.

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, you can set these values individually for each server.

Figure 11-9: Maintenance AAA

The screenshot shows the Maintenance AAA configuration page. At the top, there are navigation tabs: Fabric, Monitoring, Maintenance, Integration, and Security. Below these, the 'AAA' section is active, with sub-tabs for Accounting, Authentication, Authorization, RADIUS, and TACACS. The 'Default Settings' section shows a Timeout of 5 seconds and a Secret field. The 'TACACS Servers' section features a table with columns for Actions, Server Address, State, and Last Updated. A single server is listed with the address 192.168.1.10, state 'Configured', and last updated on Apr 12, 2022 at 10:32:35:269pm. The table has 1 of 1 items and a page size of 10.

1. Click the **Add TACACS Server** from Actions menu at top of the TACACS Servers table.

Figure 11-10: Create TACACS+ Server Dialog

The dialog box is titled 'Add TACACS Server' and contains two input fields. The first field is labeled 'Server Address' and contains the IP address '192.168.1.20'. The second field is labeled 'Secret' and is currently empty. At the bottom of the dialog, there are two buttons: 'Cancel' and 'Submit'.



Note: Do not use the pound character (#) in the TACACS secret. 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 your 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
```

Now, all users in bigtap-admin group on TACACS+ server **10.2.3.201** have full access to 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

To configure the DMF controller with TACACS+ to control administrative access to the switch, complete the following steps.

1. Identify the IP address of the TACACS+ server and any key required for access, using the `tacacs server` command, which has the following syntax:

```
tacacs server host <server> key {<plaintext-key> | 0 <plaintext-key> | 7 <encrypted-key>}
```

You can enable up to four AAA servers by repeating this command for each server. For example, the following command enables TACACS+ with the server running at **10.2.3.4**, using a plaintext key.

```
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. It will be interpreted as the start of a comment in the PAM config file.

2. Each TACACS+ server connection can be encrypted 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 below command:

```
controller-1(config)# tacacs server key 0 secret
```

If no key is specified for a given host, then the global key value is used. If no key is specified globally and no key is specified for a given host, then an empty key is assumed.



Note: Be careful when you configure TACACS+ to avoid disabling access to the DANZ Monitoring Fabric controller.

3. It is possible to configure device specific TACACS server parameters that overrides that of the global TACACS servers. This applies to switches, service-nodes and recorder nodes. The configuration has to be made from the config-device submode. 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 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 `no tacacs override-global` at the `conf#-switch` submode.

4. To view the TACACS+ configuration on a specific switch, use `show effective-config switch switch-name tacacs` command as in the following example.

```
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 is displayed 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 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. An example on how this should be configured for Aruba Clearpass server is as below:

```
<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 sets authentication and authorization to first connect to the TACACS+ server to verify user credentials and privileges. The user account will be checked locally 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.

For further details, or for instructions for setting up other servers, refer to your AAA server documentation.

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 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, 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 that is 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 not available.

DANZ Monitoring Fabric also supports communication with a remote AAA server (RADIUS). 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.

Privileges at the remote RADIUS server must be configured 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 is used for sending 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 11-11: Maintenance AAA: Radius Section

The screenshot shows the AAA configuration page with the RADIUS tab selected. Under 'Default Settings', there is a table with the following values:

Timeout	5 seconds
Secret	-
Auth Port	1812
Acct Port	1813

Below the settings is a 'RADIUS Servers' table with one entry:

Actions	Server Address	State	Last Updated
<input type="checkbox"/>	10.240.130.59	Configured	Apr 12, 2022 09:21:32:613pm

2. Click the **Add RADIUS Server** from Actions menu at the top of the RADIUS Servers table.

Figure 11-12: Create RADIUS Server Dialog

The dialog box titled 'Add RADIUS Server' contains the following fields and buttons:

- Server Address**: Input field containing '192.168.1.10'
- Secret**: Empty input field
- Cancel**: Button
- Submit**: Button

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 the right of the Secret field 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 in which 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 DMF-User-Role 1 string
ATTRIBUTE DMF-AVPair 2
string
END-VENDOR Arista-Networks
```

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 is intended only to show how the VSA is associated with the user and its privileges. In an actual deployment, a database and encrypted password is 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 sets authentication and authorization to first 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 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>
}
```

- 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 which grants the privilege to read or write like a built-in admin group, but only for certain use-cases/categories. Privilege can be configured for each predefined set of categories.

This allows a user to define a more flexible group role such as a group which allows its users to read and write most of general configuration like admin, but not the AAA related parts of the configuration.

This helps to create user profiles on DANZ Monitoring Fabric with a 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
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 inferred-does-not-elevate
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME read-write
9 category:TIME/SENSITIVE does-not-elevate

```

TIME group user sample configuration

```

Controller-1(config)# ntp server <NTP_SERVER_IP>
Controller-1(config)#

```

category:TIME/SENSITIVE:

Sensitive data included in Configuration and state related to system time and NTP.

TIME/SENSITIVE group creation and user association

```

Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group time-mgmt-group
Controller-1(config-group)# permission category:TIME/SENSITIVE 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/SENSITIVE read-write
Controller-1(config)# show group time-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 inferred-does-not-elevate
7 category:SYSOPS/SENSITIVE inferred-does-not-elevate
8 category:TIME read-write
9 category:TIME/SENSITIVE read-write

```

TIME group user sample configuration

```

Controller-1(config)# ntp key 5 sha1 0 abcdfa5cdfabcdnfabcdfa3cdfabcdnfabcdfabcdnf
Controller-1(config)#

```

category:SYSOPS

Configuration and state related to system operation

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 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 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:APPROG 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:APPROG:

Configuration and state related to appliance log level.

APPROG group creation and user association

```

Controller-1> enable
Controller-1# configure terminal
Controller-1(config)#
Controller-1(config)# group APPROG-mgmt-group
Controller-1(config-group)# permission category:APPROG privilege read-only
Controller-1(config)# user APPROG-user
Controller-1(config-user)# password <user_configured_password>
Controller-1(config-user)# group APPROG-mgmt-group
Controller-1(config-group)# associate user APPROG-user

```

```

Controller-1(config)# show group APPROG-mgmt-group configured-permission
# Permission Privilege
1 category:APPROG read-write
Controller-1(config)# show group APPROG-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:APPROG 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

```

APPROG 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
```

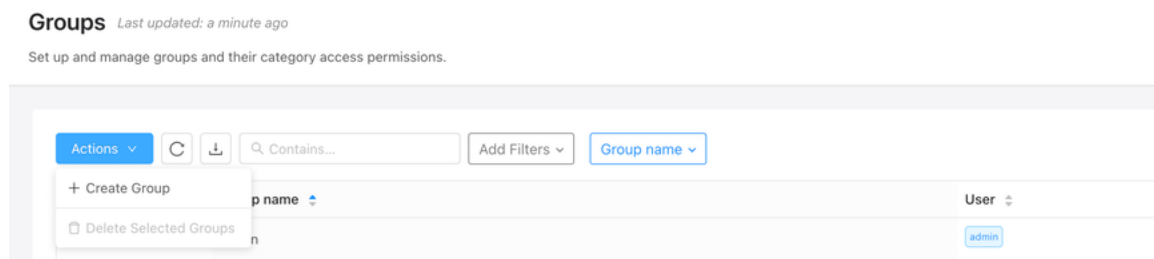
```
Controller-1(config)# show group DEFAULT-mgmt-group configured-permission
# Permission Privilege
1 category:DEFAULT read-write
Controller-1(config)# show group DEFAULT-mgmt-group effective-permission
# Effective-permission Inferred privilege
1 category:DEFAULT read-write
2 category:DEFAULT/SENSITIVE 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 into the following page to create the custom admin groups from GUI.

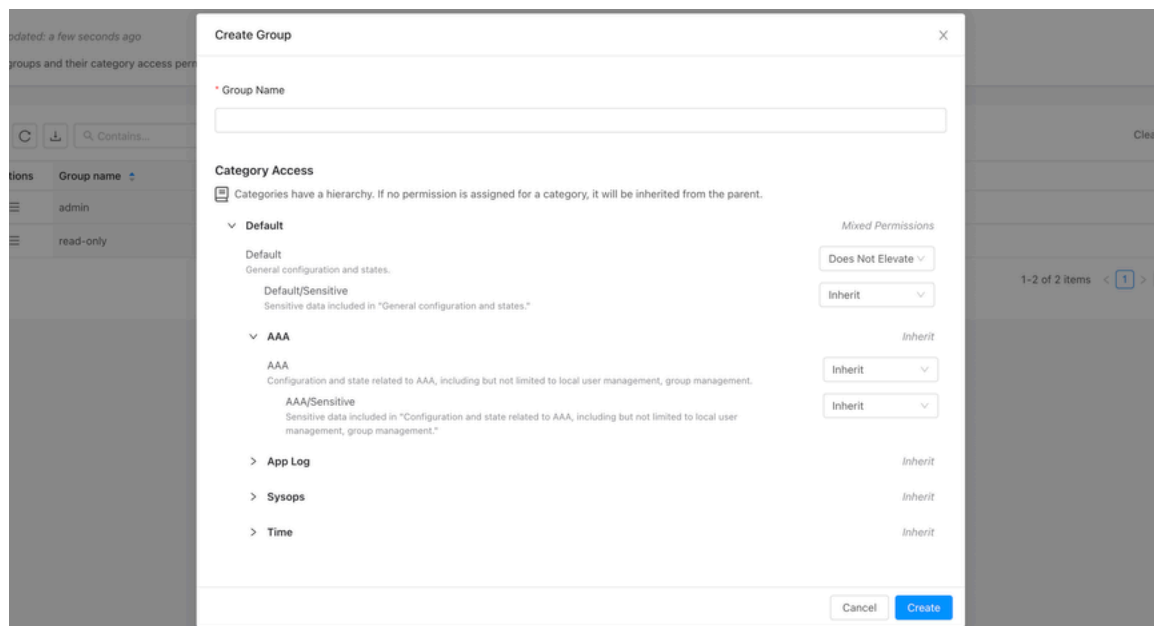
Step 1 Login to DMF.

Figure 11-13: group_creation



Step 2 Security > Groups.

Figure 11-14: custom_group



11.8.2 PERMISSIONS & PRIVILEGE:

Privileges are the permissions set of each category.

- **Read Only:** User will be granted to only read the configurations for the enabled category
- **Read-Write:** User will be granted to read and write configurations for the enabled category
- **Does Not Elevate:** User can neither read nor write any 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 its permission from its 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 privilege **read-only** for **TIME/SENSITIVE**.

Similarly, if **AAA** is set to **inherit** and **DEFAULT** has **read-only** then users associated with this group will be able to read all configurations related to AAA.

11.8.3 Category Identification:

DMF helps to identify which feature comes under what category. This can be achieved by using the following command.

Syntax: :: **category <feature_name>**

For example 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 line **AAA "aaa authentication login default local"** signifies that the above feature comes under the AAA category.

11.8.4 Group managements:

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 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-pe
rmission
# Permission Privilege
1 category:AAA does-not-elevate
2 category:DEFAULT read-write
Controller-1(config-group)# show group all-except-AAA-group effective-pe
rmission
# 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:APPLLOG 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 immediate parent i.e. Category: AAA which is “does-not-elevate” Rest of the categories are set as “read-write” since they inherit from Category: DEFAULT.

11.8.5 Remote Users

With Custom Admin, users can be associated with the groups with required privileges using remote TACACS and RADIUS servers. This can be established using the following steps:

1. Create a user **bn-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, password and the group created on step 1, 2 on the TACACS/RADIUS servers.
4. Login to the DANZ Monitoring Fabric controller using the above credential.

11.8.6 Commonly used profiles

User-administration

User 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:APPROG 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

User 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:APPROG 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

User 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:APPROG 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

User 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:APPROG 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	ACCOUTING
AAA	AUTHENTICATION
AAA	AUTHORISATION
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:

- Audit logs can be viewed only through built in admin user.
- Need to be an built-in admin user to upgrade DMF controller and managed appliances.
- RBAC has higher preference over custom admin groups.
- Custom group feature won't be applicable for 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 the network used by the administrator, whether working locally or remotely, to reach the controller management interfaces. DMF uses standard, well known cryptographic technology, such as RSA and AES, but it is the responsibility of system administrators to choose strong passwords and change them frequently, according to well- established security best practices.

All services used by the controller are enabled by default except for SNMP, which is disabled by default. To block or permit specific protocols to the management interface, refer to the [Protocol Access Required to the DMF Controller](#) section.

The control plane is the network between the controllers and the switches, for example, 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 dataplane 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) along with the OpenFlow protocol itself are not authenticated and are subject to spoofing in an untrusted network. The following are best practices in regard to securing the control plane within the switched fabric.

- The control plane network is “Layer 2 trusted,” meaning that the attacker cannot spoof Layer 2 messages on the control network. In practice, this means that the control plane network should be an isolated VLAN, which ideally contains only the controller and switches.
- The switch management interface should be hardened against Layer 3 attacks (all services are authenticated, unnecessary services 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 whitelist 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 it has been copied to the controller using the copy command.

To import a private key to the controller, enter the `private-key` command in `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 *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. To install a certificate signed by a public or private CA, complete the following steps.

Procedure

1. Generate the Certificate Signing Request (CSR) and the private key for the controller.

This operation can be performed 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.

The CSR should be submitted to the trusted CA for browsers used to access the DMF GUI. For organizations using GUI based CAs, simply 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/control
ler.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/control
ler.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/control
ler.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 that was used to sign the current certificate.

Obtain a new 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@10.240.88.130 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. Current certificate will be overwritten if the Common Name matches with the new one.
4. Import/copy the new private key to the controller. Private-key will be overwritten if the file name is same as the old one. In that case, no need for any config changes.

Assuming the CN and the private-key dest file names are different than the original ones, remove the old cert and private-key and install a new cert and private key.

To remove 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 where you can 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 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 DANZ Monitoring Fabric 8.4 release, TLSv1.3 HTTPS protocol is supported. Both TLSv1.3 and TLSv1.2 are supported by default and the former is the preferred one 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 restrict the SSH ciphers and MAC algorithms on the switches as well.

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 **DMF 8.0** release, due to upgrade of Ubuntu OS from **16.04** to **20.04**. Arista Networks refined the list of supported ciphers to better represent what can be configured and those default enabled when TLSv1.3 protocol is enabled.

- **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: The below set of SSL ciphers are enabled by default for TLSv1.3 and 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 example below 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:

1. hmac-md5-96-etm@openssh.com (HMAC-MD5 in “encrypt-then-mac” mode)
2. hmac-md5-etm@openssh.com (HMAC-MD5 in “encrypt-then-mac” mode)
3. hmac-sha1-96-etm@openssh.com (HMAC-SHA1 in “encrypt-then-mac” mode)
4. umac-64-etm@openssh.com (message authentication code based on universal hashing (UMAC) in “encrypt-then-mac” mode)
5. umac-64@openssh.com (UMAC)

The following SSH cipher algorithm supported by the controller are not supported by EOS:

1. 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 ciphers/MACs that are supported 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 TCP/UDP protocol ports that are used by DMF are summarized in the following table. The options for the ports that can be enabled or disabled using the CLI `access-list` command (***config-controller-access*** submode) are shown in the CLI access-list option column. 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 by default is permitted from all subnetworks. You can restrict SSH access to one or more specific subnetwork. However, access is then denied from all other subnetworks. If no connectivity from a specified subnetwork is available, the only method of accessing the controller is 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. If all devices are in the same Layer 2 network as the DMF controller, no further configuration is required. However, if any DMF nodes or fabric switches are connected 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 Direction	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
UDP	67	Out	DHCP		Customer DHCP server
UDP	68	In	DHCP	Customer DHCP server	
TCP	80	In	ZTN ONIE	DMF Switches	
UDP	123	Both Direction	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 Direction	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 Direction	Floodlight REST API	Customer, Recorder Node	Recorder Node, Service Node
TCP	8843	In	ZTN	Switches, Service Node, Recorder Node	
TCP	9379	Out	AN Replicated Redis		Analytics Node

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, which means that access is not permitted unless specifically enabled. By default, SNMP UDP **port 161** is blocked on controllers and fabric switches. You need to configure an SNMP access-list using the controller 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 the following access-list is configured 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
```

4. 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 for the SSH protocol only from 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
```

5. 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. By default all traffic is permitted for IPv4 (**0.0.0.0/0**) and IPv6 (**::/0**). If Active- Standby controller HA pair are deployed in different L3 subnets, permitting all traffic for sync service can be a security risk. For secure connection between active and standby controller, access-list for sync should only permit active and standby controller management IP address.

Procedure

1. First 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 sync access-list. In below example Rule id 3 is used.

```
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)#
```

2. Then 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. In below example Rule id 4 is used.

```
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
12 sync         3      permit  10.240.130.17/32
13 sync         4      permit  10.240.130.16/32
14 vce-api      1      permit  ::/0
15 vce-api      2      permit  0.0.0.0/0
DMF-CTL2(config-controller-access-list)#
```

3. Then remove the default permit entry.

```
DMF-CTL2(config-controller-access-list)# no 1 permit
DMF-CTL2(config-controller-access-list)# no 2 permit
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 3 permit 10.240.130.17/32
11 sync 4 permit 10.240.130.16/32
12 vce-api 1 permit ::/0
13 vce-api 2 permit 0.0.0.0/0
DMF-CTL2(config-controller-access-list)#

```

4. Verify cluster state using **show controller details** and make sure cluster state is redundant.

```

DMF-CTL2(config)# 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 10.240.130.17 DMF-CTL2 * 22049 1 active connected 2 weeks, 2 days
2 10.240.130.16 DMF-CTL1 27671 1 standby connected 2 weeks, 2 days
-----|-----|-----|-----|-----|-----|
# 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
DMF-CTL2(config)#

```

If default rules are removed before adding the new rules for sync access-list, controller cluster communication can break. To recover from this, please refer to [Controller Cluster Recovery](#).

12.6.4 Control Plane Access for DMF Switches

The following ports must be open for DMF Switches to communicate with DMF Controller, Analytics Node, and other services (eg. NTP, DHCP etc). If all devices are in the same Layer 2 network as the DMF controller, no further configuration is required. However, if any DMF Controller and fabric switches are connected over a Layer 3 network, these ports must be open on any firewalls or routers that connect the devices to the DMF controller.

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	
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
UDP	6343	Out	sFlow		Analytics Node
UDP	6380	Out	Control Packets		Analytics Nodes
TCP	6653	Out	OpenFlow		DMF Controller
TCP	8843	Out	ZTN		DMF Controller

12.6.5 Control Plane Access for DMF Service Node

The following ports must be open for DMF Service Node to communicate with DMF Controller, Analytics Node, and other services (eg. NTP, DHCP etc). If all devices are in the same Layer 2 network as the DMF controller, no further configuration is required. However, if

DMF Controller and Service Nodes are connected 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	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. If all devices are in the same Layer 2 network as the DMF Recorder Node, no further configuration is required. However, if DMF Controller, Analytics Node or fabric switches are connected 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 NFS 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. If all devices are in the same Layer 2 network as the Analytics Node, no further configuration is required. However, if Analytics Node is connected 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

Protocol	Port	Direction	Application	In Flows	Out Flows
UDP	2055	In	NetFlow v5	DMF Service Nodes and Switches	
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	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

12.7 Enabling Secure Control Plane

The Secure Control Plane feature encrypts communications between the DMF controller and the fabric switches using certificates issued by a third-party Certificate Authority (CA).

The first step is to download the CA root certificate and import it to the primary controller. Certificate Signing Requests (CSRs) are automatically generated by the controllers at first boot and by the fabric switches

when connected to a controller in secure control plane provisioning mode. The CSRs are copied from the local file repository on the Active controller to the CA. The signed certificates are then copied back to the Active controller.



Note: For CPSEC certificate management, one has to use their own private CA to sign controller and switch certificates. Public CA is not supported and for private CA, there is no support for Intermediate CA.

Advanced Encryption Standard (AES) is a standard that uses symmetric block ciphers to provide robust data encryption. AES-NI is a new hardware implementation of AES that accelerates encryption and decryption for supported Intel-based servers. The DMF controller hardware appliance supports AES-NI. However, this feature is not supported when upgrading the controller appliance from earlier versions of DMF and an error message is displayed.

When upgrading from an earlier DMF version, if you need to enable AES-NI hardware support on the DMF hardware appliance, contact Arista Networks technical support for help.

After the certificates are installed on the controller and the switches, the fabric can be locked down so that all further communications between the controller and the switches, including OpenFlow messages, occur over Transport Layer Security (TLS).

To configure a fabric with this feature, provisioning should occur within a trusted network with at least one controller that is isolated from external traffic. Once the Crypto-Secure Controller-Switch Interface feature is fully configured and the provisioning network is in Lockdown state, it can be safely connected to external traffic.

Before implementing this feature, it is assumed that both controllers are installed and joined in an Active-Standby cluster. Each controller uses its own certificate signed by a Certificate Authority (CA). The CA root certificate is synchronized between cluster members, and is pushed to the fabric switches. When failover occurs, switches will trust the new Active controller because it presents a certificate signed by the same CA as the original Active controller.



Note: If using IPv4 for connectivity between the controller and switches, IPAM must be enabled on the fabric switches for NTP service to work correctly. If the time between the switches and controllers is not synchronized, a certificate may appear invalid.

To enable secure control plane, complete the following steps on the primary controller.

12.7.1 Certificate Requirements and Recommendations

DMF supports certificates in PEM file format. Certificates contained in a PKCS#12 archive (.p12) can be converted to PEM file using the open source tool `openssl`, available on Linux and Mac OS. The command for conversion is as follows:

```
openssl pkcs12 -in arista.p12 -out certificate.pem
```

Certificate rollover is not supported, so it is recommended to issue certificates with a long life-time.

Certificates in DMF are used for mutual authentication, using ECDHE_RSA and RSA ciphers.

The common name issued must match the common names in the CSR generated by DMF, which are as follows:

- For controllers: `{{$(id).controller.cluster}}`
- For switches: `{{$(mac).switch.cluster}}`

Certificates must contain two X509 Key Usage attributes: Digital Signature and Key Encipherment, as shown below:

```
X509v3 Key Usage:  
Digital Signature, Key Encipherment
```

If present, the X509v3 Extended Key Usage Attribute must include TLS Web Client Authentication and TLS Web Server Authentication, as follows:

```
X509v3 Extended Key Usage:
  TLS Web Client Authentication, TLS Web Server Authentication
```

12.7.2 Installing X.509 Digital Certificates on the Active and Standby Controllers

The Active and Standby controllers are treated separately for the purpose of installing certificates. The process is the same and is repeated for both controllers, except you only have to import the CA root certificate to the Active controller and it is automatically synchronized with the Standby controller, as well as the fabric switches. To install the certificates on each controller, complete the following steps.

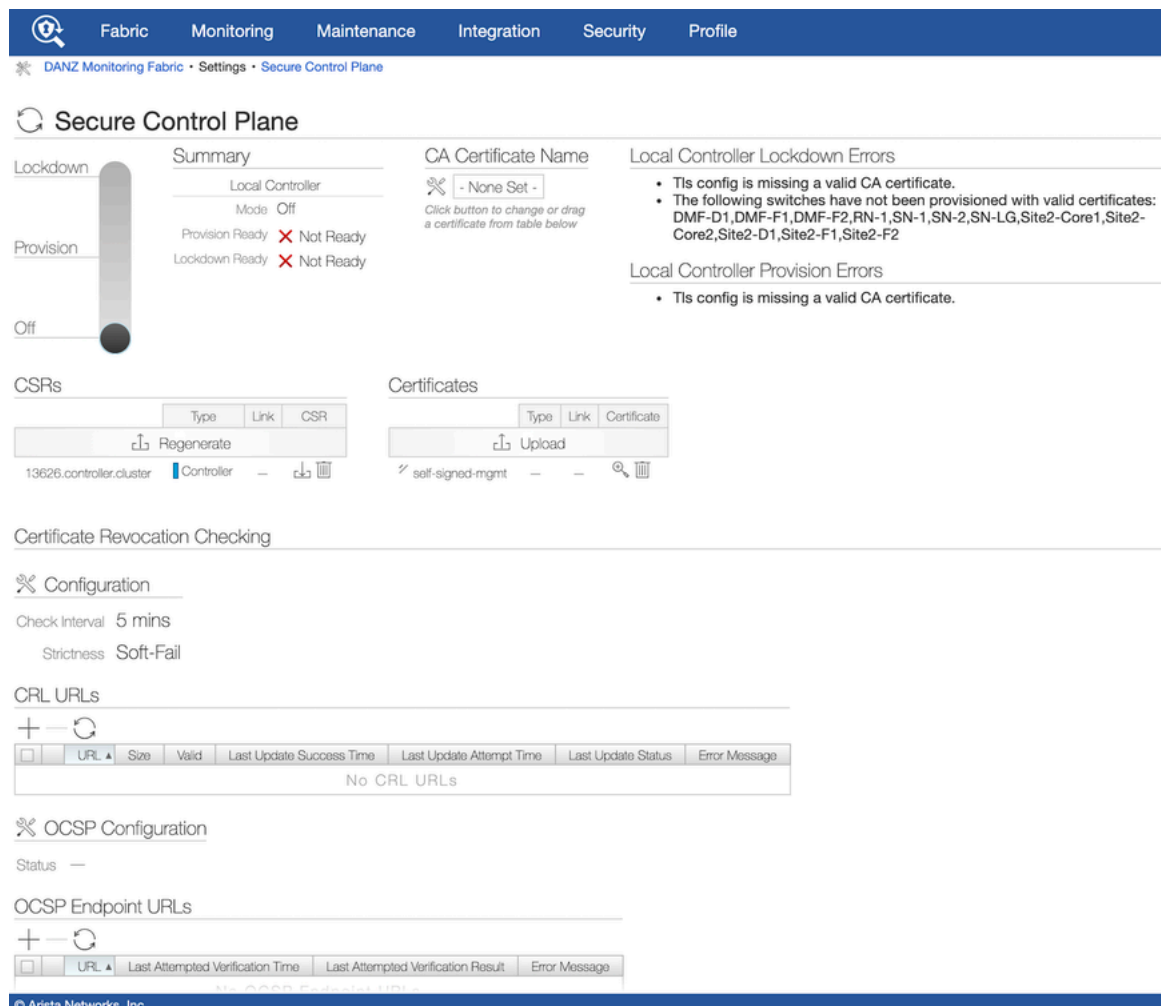
After each controller has been configured using the First Boot process, it has an auto-generated Certificate Signing Request (CSR). By default, the secure control plane is off.

12.8 Using the GUI to Enable Secure Control Plane

Perform the following steps:

1. Select the **Maintenance > Secure Control Plane** option.

Figure 12-1: Secure Control Plane Off (Default)



The figure above lists the CSRs for the current controller and connected fabric switches.



Note: The errors listed on the right side of the page must be cleared before the Secure Control Plane can enter Lockdown mode.

The slider in the upper left corner lets you change the controller secure plane mode from Off (the default) to Provision and Lockdown.

2. Copy the root certificate from the CA website to the controller, using the following CLI command.

```
copy scp://<CA-URL><root-cert>//<cert-id>
```


Replace **<CA-URL>** with the URL and credentials required to access the CA server. Replace **<root-cert>** with the file name for the root certificate, which is provided by the CA. Replace **<cert-id>** with a string that identifies the root certificate when it is copied to the controller local repository.

For example, the following command copies the CA root certificate and assigns cacert-example as the certificate ID on the controller.

```
controller-1 (config) # copy scp://root@10.8.67.3:/root/openssl-ca/cacert.pem  
cert://cacert-example
```

3. Click the **Upload** control in the Certificates section.

Figure 12-2: Upload Certificates



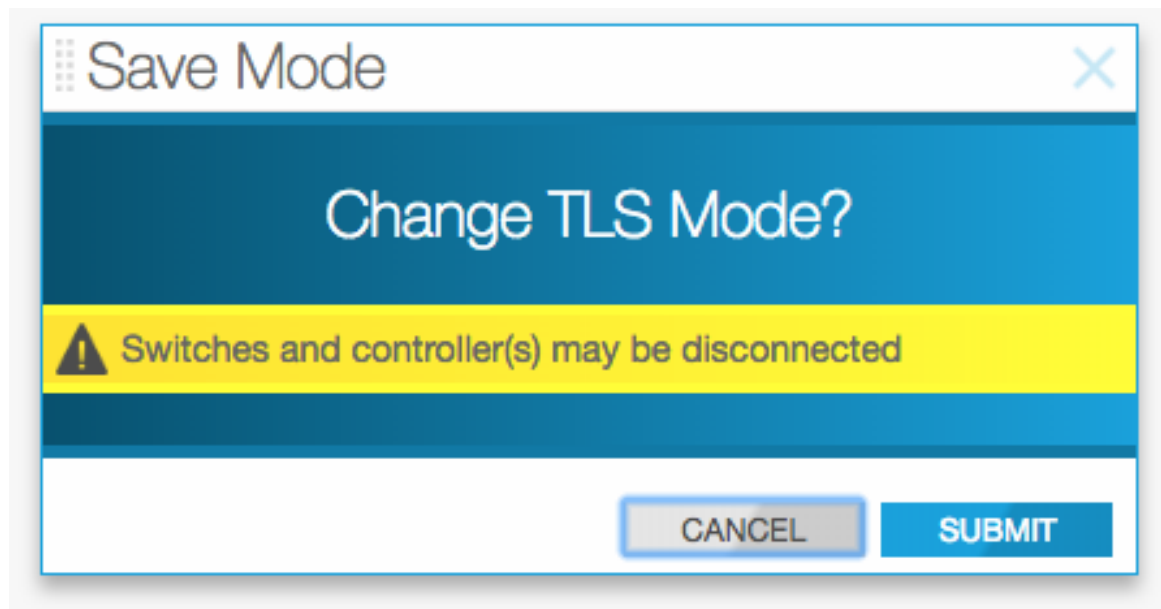
4. Drag and drop the CA root certificate to the CACERT/cacert field to import the CA root certificate, or click the CACERT/cacert
In the dialog that appears, type the certificate ID assigned to the CA root certificate and click **Submit**.
5. Click the link in the CSR column to download the Active and Standby controller CSRs and send it to the CA as directed by the CA.
6. Copy each signed controller certificate from the CA to the Active controller by entering the copy scp command.

```
copy scp://<CA-URL> cert://<controller-cert>
```

The CSRs are automatically removed after the signed certificates are copied from the CA.

7. Click the **Upload** link in the Certificates section and copy and paste the signed Active controller certificate to the dialog.
Copy from (including) "BEGIN CERTIFICATE" to (including "END-CERTIFICATE."
8. Click the slider from **No** to **Provision**.

Figure 12-3: Confirm Change of TLS Mode



9. Click **Submit** to confirm the operation.

After moving the controller to **Provision** mode, a CSR is generated for each connected fabric switch, as shown in the figure below.

Figure 12-4: Secure Control Plane—Switch CSRs (Provision Mode)

Secure Control Plane

Lockdown: Provision

CA Certificate Name: CACERT/cacert

Lockdown

- The following switches have not been provisioned with valid certificates: [leaf0-a, leaf0-b, leaf1-a, leaf1-b, spine0, spine1]

CSRs	Type	Link	CSR
ec-f4-bb-fb-de-41.switch.cluster	Switch	leaf1-b	↓
70-72-cf-bd-58-34.switch.cluster	Switch	spine1	↓
70-72-cf-ae-b6-34.switch.cluster	Switch	spine0	↓
70-72-cf-c7-c1-ed.switch.cluster	Switch	leaf1-a	↓
34-17-eb-f7-a0-c4.switch.cluster	Switch	leaf0-b	↓
70-72-cf-b5-f4-5c.switch.cluster	Switch	leaf0-a	↓

Certificates	Type	Link	Certificate
CACERT/cacert	—	—	📄
4529.controller.cluster	Controller	—	📄
24906.controller.cluster	Controller	—	📄

Switch IP Allocation

Status: Enabled

DNS Server: 10.3.0.4

Gateway: 10.8.67.1

Total Allocated Addresses: 46

Starting IP	Ending IP	Subnet Mask Length	Addresses Allocated	Addresses Used	Utilization
10.8.67.200	10.8.67.245	24	46	6	13.0%

Show: 10 25 100 All

- Click the download link for each switch CSR and send each switch CSR to the CA as directed by the CA. When the certificate is signed by the CA, the URL for obtaining the certificate is sent to the applicant.
- Enter the `copy scp` command to copy the signed certificate from the CA to the controller. This command has the following syntax.

```
copy scp://<CA-URL> <cert>://<switch-cert>
```

For example, the following command copies the signed switch certificate from the CA to the Active controller.

```
controller-1(config)# copy scp://root@10.8.67.3:/root/openssl-ca/70-72-cf-ae-b6-34.switch.cluster.pem cert://70-72-cf-ae-b6-34.switch.cluster
```

After each switch certificate is copied to the controller, it is automatically downloaded to the switch and synchronized with the Standby controller.

Figure 12-5: Upload Certificates


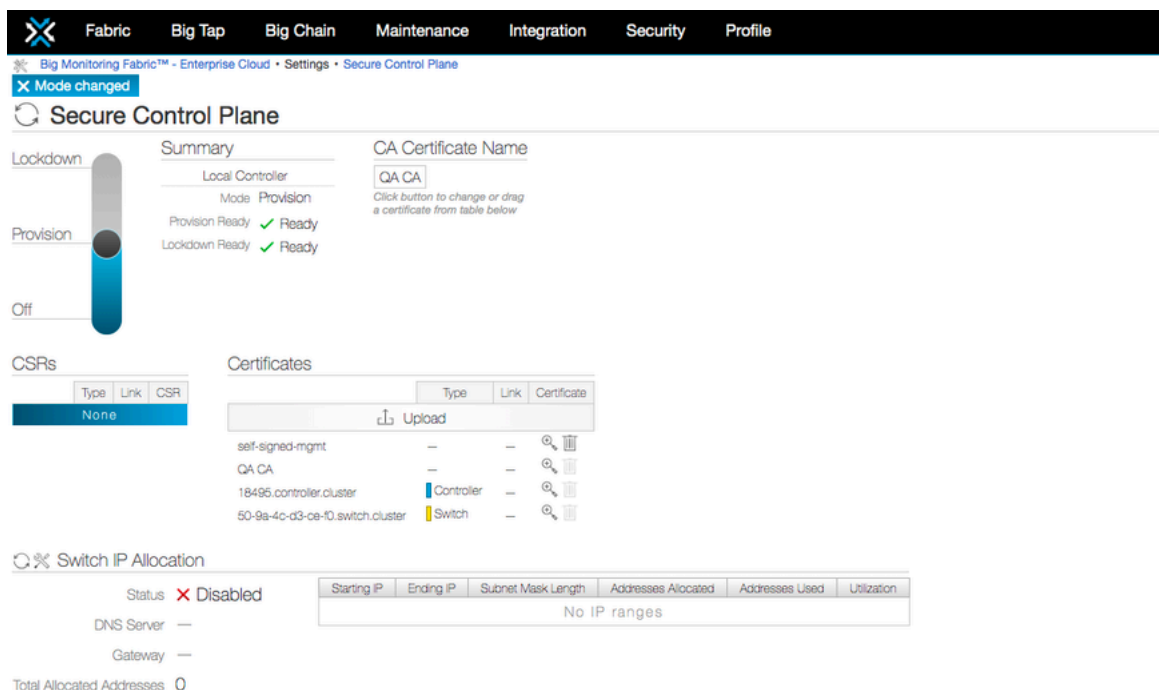


Figure 12-6: Secure Control Plane: Errors Cleared



The one error remaining stating that the switch is not provisioned with a valid certificate clears after the switches reboot into lockdown mode, assuming the switch certificate has been signed by the CA and copied to the Active controller.

After all the switch certificates are uploaded to the controller, click the slider from Provision to Lockdown, as shown in the following figure.

Figure 12-7: Secure Control Plane—Lockdown

Secure Control Plane

Lockdown Provision Off

Summary
Mode Lockdown
CA Certificate Name CACERT/cacert

CSRs

Type	Link	CSR
None		

Certificates

Type	Link	Certificate
CACERT/cacert	—	—
4529.controller.cluster	Controller	—
24906.controller.cluster	Controller	—
ec-14-bb-fb-de-41.switch.cluster	leaf1-b	—
70-72-cf-bd-58-34.switch.cluster	spine1	—
70-72-cf-ae-b6-34.switch.cluster	spine0	—
70-72-cf-c7-c1-ed.switch.cluster	leaf1-a	—
70-72-cf-b5-f4-5c.switch.cluster	leaf0-a	—
34-17-eb-f7-a0-c4.switch.cluster	leaf0-b	—

Switch IP Allocation

Status Enabled
DNS Server 10.3.0.4
Gateway 10.8.67.1
Total Allocated Addresses 46

Starting IP	Ending IP	Subnet Mask Length	Addresses Allocated	Addresses Used	Utilization
10.8.67.200	10.8.67.245	24	46	6	13.0%

Show: 10 25 100 All

12. Confirm the operation on the dialog that is displayed.

Once the controller is in Lockdown state, switches can only join the fabric if they have certificates signed by the trusted CA. New switch certificates can be installed by connecting the new switches to a controller in a trusted network in Provision mode. If security in the production network is not a concern, you can return the production controller to Provision mode until the new switches have certificates installed, and then move the production controller back to Lockdown state.

12.9 Using the CLI to Enable Secure Control Plane

The following CLI commands are used to implement and manage this feature:

- **secure control plane ca <cert>**: (Active controller only) to import the CA root certificate, which registers the CA as a trusted CA on the controller. This trust relationship is synchronized from the Active controller to the Standby controller and the connected fabric switches.
- **secure control plane off | provision | lockdown**: to change the controller control plane mode.
- **copy csr://<> scp://<>**: submit CSR to CA.
- **copy scp://<> cert://<>**: import certificate from CA.
- **delete cert <cert-name>**: to delete an existing certificate.
- **show running-config control plane**: view control plane security configuration

- **show secure control plane**: view the operational state of control plane security and to view CSRs and certificate names.
- **system regenerate csr**: create a new CSR, which is required when changing to a different CA.

Perform the following steps:

1. To identify the node ID that is automatically assigned to the Active and Standby controllers, enter the following command:

```
controller-1> show controller details
Cluster Description : T6
Cluster Virtual IP : 10.8.67.12
Redundancy Status : redundant
Redundancy Description : Cluster is Redundant
Last Role Change Time : 2020-10-19 15:34:34.656000 UTC
Cluster Uptime : 6 hours, 10 minutes
# IP @ Node Id Domain Id State Status Uptime
-|-----|-----|-----|-----|-----|-----|
1 10.8.67.11 18082 1 active connected 4 hours, 40 minutes
2 10.8.67.10 * 27141 1 standby connected 3 minutes
```

The node ID is used to automatically generate a certificate signing request (CSR) for each controller during First Boot.

2. To determine the current status of control place security and display the CSR for the Active and Standby controller, enter the **show secure control plane** command, as in the following example:

```
controller-1> show secure control plane
State Configured mode : off
State Current mode : off
Provision Issue : TLS must be fully configured before entering provision or
lock mode
Provision Issue : Every controller in the cluster must have a valid
certificate imported
before entering provision or lock mode
Lockdown Issue : The following switches have not been provisioned with valid
certificates:
[filter0-a, filter0-b, filter1-a, delivery0, delivery1]
~~~~~ Switches ~~~~~
# Switch State
-|-----|-----|
1 DMF-CORE-SWITCH-1 unprovisioned
2 DMF-DELIVERY-SWITCH-1 unprovisioned
3 DMF-FILTER-SWITCH-1 unprovisioned
~ Certs ~
# Name
-|-----|
None
~~~~~ Csrs ~~~~~
# Name
-|-----|
1 18082.controller.cluster
2 27141.controller.cluster
```

The CSR is automatically generated for the Active and Standby controllers during first boot and named by appending **controller.cluster** to the controller node ID.

3. (Active controller only) Copy the CA root certificate to the Active controller by entering the following command:

```
copy scp://<CA-URL><root-cert> cert://<cert-id>
```

Replace **<CA-URL>** with the URL and credentials required to access the CA server. Replace **<root-cert>** with the file name for the root certificate, which is provided by the CA. Replace **<cert-id>** with a string that identifies the root certificate when it is copied to the controller local repository.

For example, the following command copies the CA root certificate and assigns **cacert-example** as the certificate ID on the controller:

```
controller-1(config)# copy scp://root@10.8.67.3:/root/openssl-ca/cacert.pem
cert:
//cacert-example
```

4. To verify that the certificate is added, enter the following command:

```
controller-1> show secure control plane
State Configured mode : off
State Current mode : off
Provision Issue : TLS must be fully configured before entering provision or
lock
mode Provision Issue : Every controller in the cluster must have a valid
certificate
imported before entering provision or lock mode
Lockdown Issue : The following switches have not been provisioned with valid
certificates: [filter0-a, filter0-b, filter1-a, delivery0, delivery1]
~~~~~ Switches ~~~~~
# Switch      State
-|-----|-----|
1 filter0-a  unprovisioned
2 filter0-b  unprovisioned
3 filter1-a  unprovisioned
4 delivery0  unprovisioned
5 delivery1  unprovisioned
~ Certs ~
# Name
-|-----|
1 cacert-example <====Use this certificate ID to import the CA root
certificate to the Active controller
~~~~~ Csrs ~~~~~
# Name
-|-----|
1 18082.controller.cluster
2 27141.controller.cluster
```

5. Import the CA root certificate to the Active controller by entering the following command:

```
secure control plane ca <cert-id>
```

Replace **<cert-id>** with the certificate identifier you assigned to the CA root certificate when you copied it from the CA. For example, the following command imports the root certificate **cacert-example**.

```
controller-1(config)# secure control plane ca cacert-example
```

This causes the CA associated with the root certificate to become the trusted CA for the Active controller. This trust relationship is then synchronized with the Standby controller and connected fabric switches.



Note: Only one trusted CA is allowed per fabric. To generate a new CSR to send to a different CA, enter the **system regenerate csr** command from Enable mode.

6. Enter the **copy csr** command on the Active controller to send the CSR for the Active and Standby controllers to the third-party CA or other Public Key Infrastructure (PKI). This command has the following syntax:

```
copy csr://<controller-csr> scp://<CA-URL>
```

For example, the following command copies the CSR of the Active controller (18082.controller.cluster) to the CA server:

```
controller-1(config)# copy csr://18082.controller.cluster
scp://root@10.8.67.3:/root/openssl-ca/18082.controller.cluster.csr
```

The controller CSR is then signed by the PKI and a notification is sent to the applicant containing the URL for the signed certificates.

The command for copying the CSR of the Standby controller to the CA, in this example, is as follows:

```
controller-1(config)# copy csr://27141.controller.cluster
scp://root@10.8.67.3:/root/openssl-ca/27141.controller.cluster.csr
```

7. Copy each signed controller certificate from the CA to the Active controller by entering the **copy scp** command:

```
copy scp://<CA-URL> cert://<controller-cert>
```

For example, the following command copies the Active controller certificate from the CA server to the Active controller file repository:

```
controller-1(config)# copy scp://root@10.8.67.3:/root/openssl-ca/1
8082.controller.
cluster.pem cert://18082.controller.cluster
```

The command for copying the Standby controller certificate in this example is as follows:

```
controller-1(config)# copy scp://root@10.8.67.3:/root/openssl-ca/2
7141.controller.cluster.
pem cert://27141.controller.cluster
```

8. To verify that the controller certificates have been copied, enter the **show secure control plane** command:

```
...
~~~~~ Certs ~~~~~
# Name
-|-----|
1 18082.controller.cluster
2 27141.controller.cluster
3 cacert-example
~ Csr ~
None.
```



Note: The CSRs are automatically removed once the signed certificates are copied from the CA.

9. After both controllers have their certificates registered, enter the following command to begin provisioning:

```
controller-1(config)# secure control plane provision
```

After the certificates are installed on the Active and Standby controllers, each switch can run Zero Touch Fabric (ZTF) using HTTPS using the server-side certificate installed on the controller.

When the switch starts, it auto-generates a key-pair and a self-signed certificate. The switch uses its self-signed certificate to go through initial ZTF.

12.9.1 Verification and Troubleshooting

To enter provision mode, the following must be completed:

- The CA root certificate must be imported and configured on the controllers.
- The controller device certificates must be imported.
- IPAM must be enabled for the fabric switches, if using IPv4 for communication between the controller and switches.

To verify that the controller is in provision mode, enter the `show secure control plane` command:

```
DMF2(config)# show secure control
plane State Configured mode : provision
State Current mode : provision
...
~~~~~ Csrs ~~~~~
# Name
-|-----|
1 34-17-eb-f7-a0-c4.switch.cluster
2 70-72-cf-ae-b6-34.switch.cluster
3 70-72-cf-b5-f4-5c.switch.cluster
4 70-72-cf-bd-58-34.switch.cluster
5 70-72-cf-c7-c1-ed.switch.cluster
```

This example shows that the controller is in control plane provision mode and displays the CSRs that were generated for each connected fabric switch. The CSR names are generated by appending `switch.cluster` to the MAC address of each switch.

12.9.2 Installing X.509 Digital Certificates on the Switches

After the initial switch configuration is completed, the process of installing certificates on the switch begins automatically. The switch auto-generates a CSR and sends it periodically to the controller. The switch then periodically requests its signed certificate from the controller until it receives a reply.

To install the certificates on each switch, complete the following steps on the Active controller.

Procedure

1. Enter the `copy csr` command to copy the switch CSR from the controller to the PKI. This command has the following syntax:

```
copy csr://<controller-URL> scp://<CA-URL>
```

For example, the following command copies the CSR for a fabric switch with MAC address `70-72-cf-ae-b6-34` to the CA:

```
controller-1(config)# copy csr://70-72-cf-ae-b6-34.switch.cluster scp://
root@10.8.67.3:/
root/openssl-ca/70-72-cf-ae-b6-34.switch.cluster.csr
```



Note: The certificate is signed by the PKI and the URL for obtaining the certificate is sent to the applicant.

2. Enter the `copy scp` command to copy the signed certificate from the PKI to the controller. This command has the following syntax:

```
copy scp://<CA-URL> <cert>://<switch-cert>
```

For example, the following command copies the signed certificate from the CA to the Active controller:

```
controller-1(config)# copy scp://root@10.8.67.3:/root/openssl-ca/70-72-cf-ae-b6-34.switch.cluster.pem cert://70-72-cf-ae-b6-34.switch.cluster
```



Note: After each switch certificate is copied to the controller, it is automatically downloaded to the switch and synchronized with the Standby controller.

12.9.3 Securing the Control Channel Between the Controller and the Switches

After the certificates are installed on the switches and the controller, you can secure communications between the controller and the switches so that all messages are encrypted using TLS. After the switch receives its certificate, it automatically reboots and then periodically announces to the controller that it is provisioned and ready to enter lockdown state.

Procedure

1. To view the status of the switches at this stage, enter the `show secure control plane` command, as in the following example:

```
controller-1> show secure control plane
State Configured mode : provision
State Current mode : provision
~~~~ Switches ~~~~
# Switch State
-|-----|-----|
1 DMF-CORE-SWITCH-1 lockdown
2 DMF-DELIVERY-SWITCH-1 lockdown
3 DMF-FILTER-SWITCH-1 lockdown
~~~~~ Certs ~~~~~
# Name
-|-----|
1 18082.controller.cluster
2 27141.controller.cluster
3 34-17-eb-f7-a0-c4.switch.cluster
4 70-72-cf-ae-b6-34.switch.cluster
5 70-72-cf-b5-f4-5c.switch.cluster
6 70-72-cf-bd-58-34.switch.cluster
7 70-72-cf-c7-c1-ed.switch.cluster
8 cacert-example
~ Csr ~
None.
```

2. To view the secure control plane status of a specific switch, enter the following command:

```
show switch <switch> secure control plane
Replace <switch> with the name of a specific switch. For example, the following command
displays the secure control plane status of delivery0:
controller-1> show switch delivery0 secure control plane
Next tls mode : off
Certificate issuer : /countryName=US/stateOrProvinceName=CA/organizationName=Arista
Networks/localityName=Santa Clara/commonName=Engineering/organizationalUnitName=Switch
Light/emailAddress=support@arista.com
Certificate start : 2015-12-07 09:36:50 UTC
Tls mode : off
Certificate validity : False
Certificate end : 2025-12-04 09:36:50 UTC
Certificate modulus :
F4AFDA90D24F1860A49566C1F1B159F76E04ADE7D9918EF658382DB6C15CA545F63A6B9643BC9C01BE564048C4A59346DFBACF8E
A227A1DDC7D90618B563789E18D3BF3C63D9D02973B0799F222076288373F70C4CAD184E89985459DA4E6618A4B2074E647A8138
E6C45714DBA1EA1940653E9BC8E67A70EAD4AE45FABB7C55904C9232FFEA7D36044230864C534B3839D246936B6A0DEF7ED19D01
0C35F03C2095BA49E07D669A0CB88A040FE4F6DF9CCF7CEEBFE0A134A8A3AF41C818470761367F3DAFD2AA4A20E8C65EA5E92FEE
1C12224B0B145B7A93BBAF9A3B1CFB7B22E31E4849FFC0603746726925ABE0CD31B7A1327A6D284EF111439FF1142B
Key modulus :
F4AFDA90D24F1860A49566C1F1B159F76E04ADE7D9918EF658382DB6C15CA545F63A6B9643BC9C01BE564048C4A59346DFBACF8E
A227A1DDC7D90618B563789E18D3BF3C63D9D02973B0799F222076288373F70C4CAD184E89985459DA4E6618A4B2074E647A8138
E6C45714DBA1EA1940653E9BC8E67A70EAD4AE45FABB7C55904C9232FFEA7D36044230864C534B3839D246936B6A0DEF7ED19D01
0C35F03C2095BA49E07D669A0CB88A040FE4F6DF9CCF7CEEBFE0A134A8A3AF41C818470761367F3DAFD2AA4A20E8C65EA5E92FEE
```

```
1C12224B0B145B7A93BBAF9A3B1CFB7B22E31E4849FFC0603746726925ABE0CD31B7A1327A6D284EF111439F8FF1142B
Certificate subject : /countryName=US/stateOrProvinceName=CA/organizationName=Arista
Networks/localityName=Santa Clara/commonName=Engineering/organizationalUnitName=Switch
Light/emailAddress=support@arista.com
```

3. After all switches are provisioned with certificates, enter the **secure control plane lockdown** command to enter lockdown.

```
controller-1(config)# secure control plane lockdown
```

This command affects the Active and Standby controllers. After the control plane is secured, all communication between the controllers and the fabric switches occurs using TLS. After executing this command, switches automatically disconnect from the controllers and then reconnect using TLS.



Note: While a controller is in lockdown mode, a physical switch requires a certificate issued by a CA trusted by the controller to connect to the fabric.

4. To verify that the controllers are in lockdown mode, enter the **show secure control plane** command again, as in the following example:

```
controller-1> show secure control plane State Configured mode : lockdown
State Current mode : lockdown
~~~~ Switches ~~~~
# Switch                               State
-|-----|-----|
1 DMF-CORE-SWITCH-1                    lockdown
2 DMF-DELIVERY-SWITCH-1                lockdown
3 DMF-FILTER-SWITCH-1                  lockdown
~~~~~ Certs ~~~~~
# Name
-|-----|
1 18082.controller.cluster
2 27141.controller.cluster
3 34-17-eb-f7-a0-c4.switch.cluster
4 70-72-cf-ae-b6-34.switch.cluster
5 70-72-cf-b5-f4-5c.switch.cluster
6 70-72-cf-bd-58-34.switch.cluster
7 70-72-cf-c7-c1-ed.switch.cluster
8 cacert-example
~ Csr ~
None.
```

After the controllers have entered lockdown mode, no additional switches can join the fabric until the controller control plane is returned to provision mode.

To return the controller control plane to provision mode, enter the following command:

```
controller-1(config)# secure control plane provision
```

12.9.4 Adding Switches After Enabling Secure Control Plane

After control plane security is enabled on the production (untrusted) network and the controller enters lockdown mode, no switches can be added to the network unless they have certificates installed by a CA trusted by the controller.

If you have a separate provisioning network, with the controller in provisioning mode, follow the procedure for installing certificates on the new switches in the “Installing X.509 Digital Certificates on the Switches” section. Then just connect the switches after the certificates are installed to the production controller, which is still in lockdown mode.

The switches provisioned in a separate trusted network can communicate with the production controller because the newly provisioned fabric switches and the production controller have certificates issued by the same trusted CA.

To add more switches directly to the production network, complete the following steps:

Procedure

1. Return the production controller to provisioning mode.
2. Connect the new switches to the production controller.
3. Provision the new switches as described in the “Installing X.509 Digital Certificates on the Switches” section.
4. Return the production controller to lockdown mode.

12.9.5 Replacing a Controller Appliance After Control Plane Security is Enabled

To replace a DMF Hardware Appliance after control plane security has been enabled, complete the following steps:

Procedure

1. Remove the configuration of the hardware appliance by entering the following command:

```
controller-1(config)# system remove-node
```

2. Complete the installation on the new appliance and join the existing cluster. The cluster control plane encryption moves to a not-ready state, and switches cannot connect to the new controller.
3. Sign the CSR of the new controller and import it, as described in the [Installing X.509 Digital Certificates on the Active and Standby Controllers](#) section.

The cluster returns to lockdown state and switches can connect to the new controller.

12.10 Replacing Certificates of Controllers and Switches

12.10.1 Replacing to-be-expired certificates of controllers and switches



Note: The next three sections detail the procedure to replace the certificates and the assumption is that the old CSRs that were used for signing/setting up CPSEC the first time are not available for signing again. So the user would have to regenerate the CSRs when trying to replace the certificates.

Certificates expiration information is shown in `show secure` and in `show fabric warning`. Certificates that will expire in 184 days are shown as warnings.

```
controller-1# show fabric warning
~~~~~ CPSEC warnings ~~~~~
# Warning description
-|-----|
1 Node local certificate is expiring in 183.9 days.
2 Switch dmf-filter-sw-1 certificate is expiring in 183.9 days.
3 Switch dmf-filter-sw-2 certificate is expiring in 183.9 days.
4 Switch dmf-filter-sw-3 certificate is expiring in 183.9 days.
5 Switch dmf-delivery-sw-1 certificate is expiring in 183.9 days.
6 Switch dmf-delivery-sw-2 certificate is expiring in 183.9 days.
7 Switch dmf-core-sw-1 certificate is expiring in 183.9 days.
8 Switch dmf-core-sw-2 certificate is expiring in 183.9 days.
```



Warning: Replacing the certificates is HITFUL procedure as the switches will reboot to turn off and turn on CPSEC.




Note: It is better to replace the to-be-expired certificates before they expire as we need switch console access to turn off CPSEC if the certs are expired.

Steps:

1. Unconfigure lockdown mode by `no secure control plane lockdown` on active node. Switches will reboot to go from “tls strict” (CPSEC lockdown) to “tls off” (CPSEC off).
2. Delete all controller and switches certificates by `delete cert <cert name>` on active node.
3. Regenerate CSRs for both controllers by `system regenerate csr` on each node.
4. Sign the CSRs by CA and import them to active node.
5. Execute `system regenerate switch <switch name>` for each switch on the active node. This step is needed to remove old certificates on the switches and regenerate new CSRs. Deleting the switch certificates (in Step 2) on the controller does not remove them from the switches.
6. Configure `secure control plane provision` and wait for all switches to send their CSRs to controller.
7. Sign the switch CSRs by CA and import them to active node.
8. Active node pushes the certificates to switches. Switches will go for a reboot to get the new certificate installed and lock themselves down.
9. Configure `secure control plane lockdown` to get fabric into CPSEC lockdown mode.

12.10.2 Replacing expired certificates of switches

Steps:

1. Configure no secure control plane lockdown on the active controller. Switches will go for a reboot to turn “tls” off and connect to the controllers again.
 -  **Note:** In case a switch is not connected to the fabric (or) go into a loop due to tls is not turned off properly, do one of the following (both requires console connection to the switch):
 - a. “ONIE Uninstall OS” if the switch is in reboot loop (control+C to get out of the loop and issue ‘reload’ again)
 - b. On the switch CLI, configure `tls mode off` which will force the switch to go for a reboot.
2. Once switches connect back to the controller, delete all the switch certificates on the active controller. This will remove the certificates stored on the controller.
3. Execute `system regenerate switch <switch name>` on the active controller. This will remove / reset existing certificates on the switches and generate new CSR/key. Deleting the switch certificates on the controller (in **Step 2**) does not remove the expired certificates on the switches.
4. Configure `secure control plane provision` on the active controller so that switches will send their new CSR to controller.
5. Sign the CSRs by CA and import the certificates to active controller. Switches will go for a reboot and install the new certificates to lock themselves down.
6. Once switches connect back to controller and in lockdown mode, Configure secure control plane lockdown on the active controller.

12.10.3 Replacing expired certificates of controller and switches

Controller and switch certificates are expired in the below output. When this happens, active controller will be in ‘split-brain’ mode and standby will be in ‘inconsistent’ mode.

```

controller-1(config-switch)# show fabric warning
~~~~~ CPSEC warnings ~~~~~
# Warning description
-|-----|
1 Node local certificate has expired.
2 Node local certificate is expired.
3 Switch dmf-filter-sw-1 certificate is expired.
4 Switch dmf-filter-sw-2 certificate is expired.
5 Switch dmf-filter-sw-3 certificate is expired.
6 Switch dmf-delivery-sw-1 certificate is expired.
7 Switch dmf-delivery-sw-2 certificate is expired.
8 Switch dmf-core-sw-1 certificate is expired.

```



```
9 Switch dmf-core-sw-2 certificate is expired.
```



Note: When both controllers and switch certificates are expired, we need switch console access to do **Step 2**.

Steps:

1. Configure **no secure control plane lockdown** on active node. This will force the active/standby selection to happen and floodlight process on one of the controllers will restart to join the fabric as standby. Also the switches will be disconnected from the fabric and run in headless mode.
2. Because switches are disconnected from the controllers, configuration done in step 1 is not pushed to switches. We need to login to switches via console (or if admin password is set and IPAM configured, user can “ssh” to switch) and configure following command on all switches. Switches will reboot to turn CPSEC off and join the controllers.

```
tls mode off
```

```
dmf-core-sw-1#config
dmf-core-sw-1(config)# tls mode off
System will be in 'off' mode after the next reboot.
```

3. Delete all controller and switch certificates.
4. Generate CSRs for controllers by doing **system regenerate csr** on both the nodes.
5. Sign the CSRs by CA and import the certificates to active controller.
6. Execute **system regenerate switch <switch name>** on the active controller. This will remove / reset existing certificates on the switches and generate new CSR/key. Deleting the switch certificates on the controller (in Step 3) does not remove the expired certificates on the switches.
7. Configure **secure control plane provision** on active node and switches will send the new CSRs to controller.
8. Sign the switch CSRs by CA and import the certificates to controller.
9. Switches will reboot to turn CPSEC on with new certificates and lock themselves down.
10. Once switches are in CPSEC lockdown mode, configure **secure control plane lockdown** on active node.

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 BMF 6.0.1 Controller image on Dell R430 servers and the DMF Service Node Appliance image on Dell R630 or R730 servers.

This chapter has been tested using **iDRAC Enterprise Version 2.10.10** and **BMF Release 6.0.1**. The instructions were tested with **BMF 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 have been tested using MAC OS but the steps are similar using Windows OS. You can 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: <https://www.dell.com/learn/us/en/19/campaigns/contact-us-phone-number>.

Alternatively, an iDRAC Enterprise license can be purchased 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

To set up iDRAC on a Dell R430, R630, or R730 server, complete the following steps.

Procedure

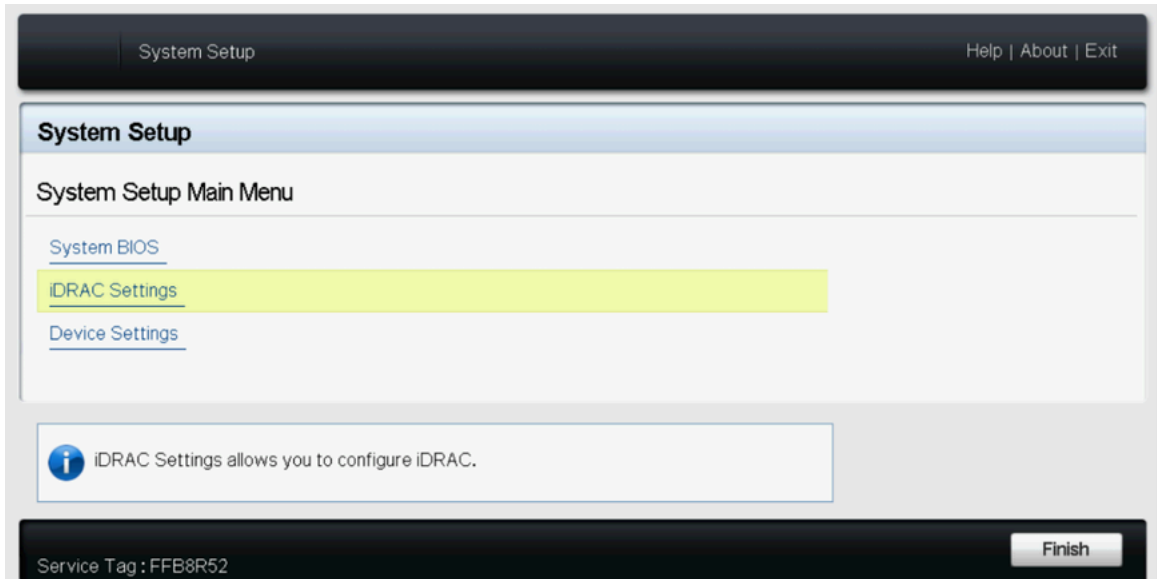
1. Power on the Dell PowerEdge Server.
2. Connect a monitor to the VGA port and connect the keyboard to the USB port.



Note: When using the serial port to connect to the R430, R630, or R730, the baud rate should be set to 115200.

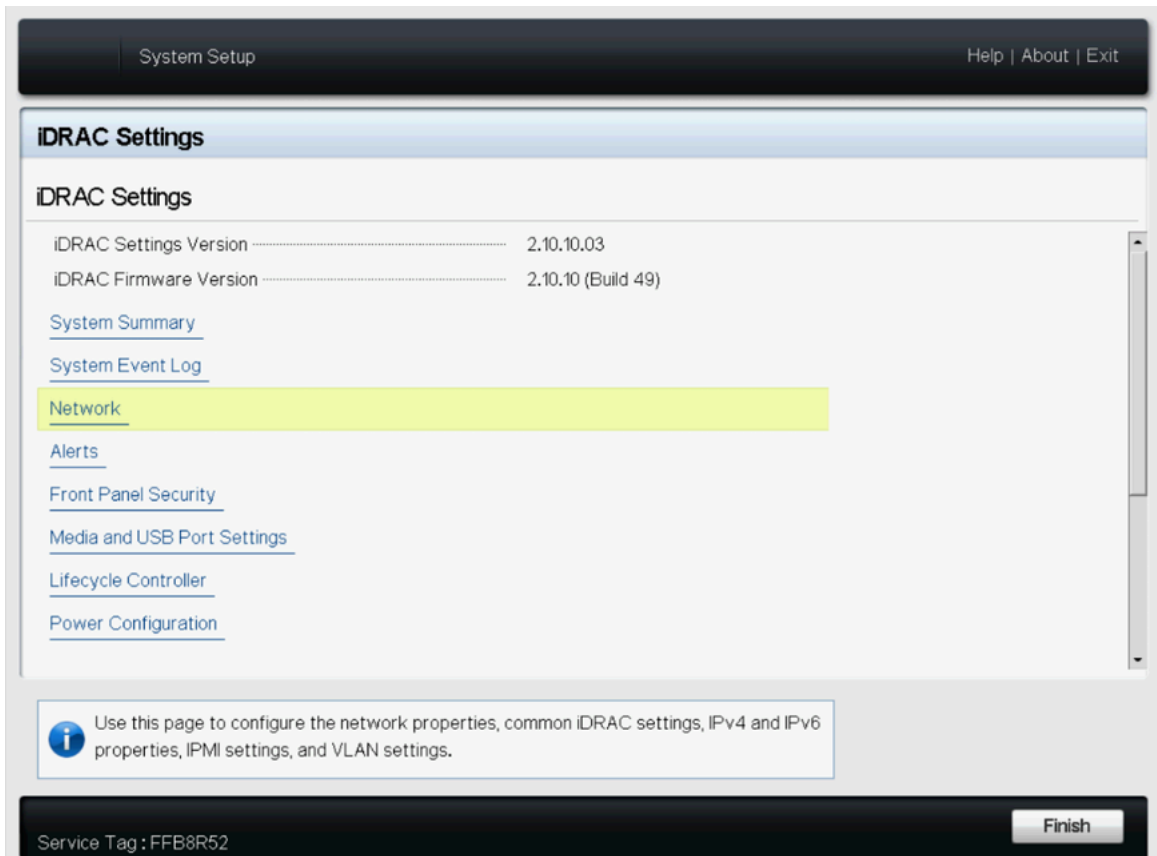
3. Press <F2> to enter the **System Setup Main Menu** screen.

Figure 13-1: System Setup Main Menu



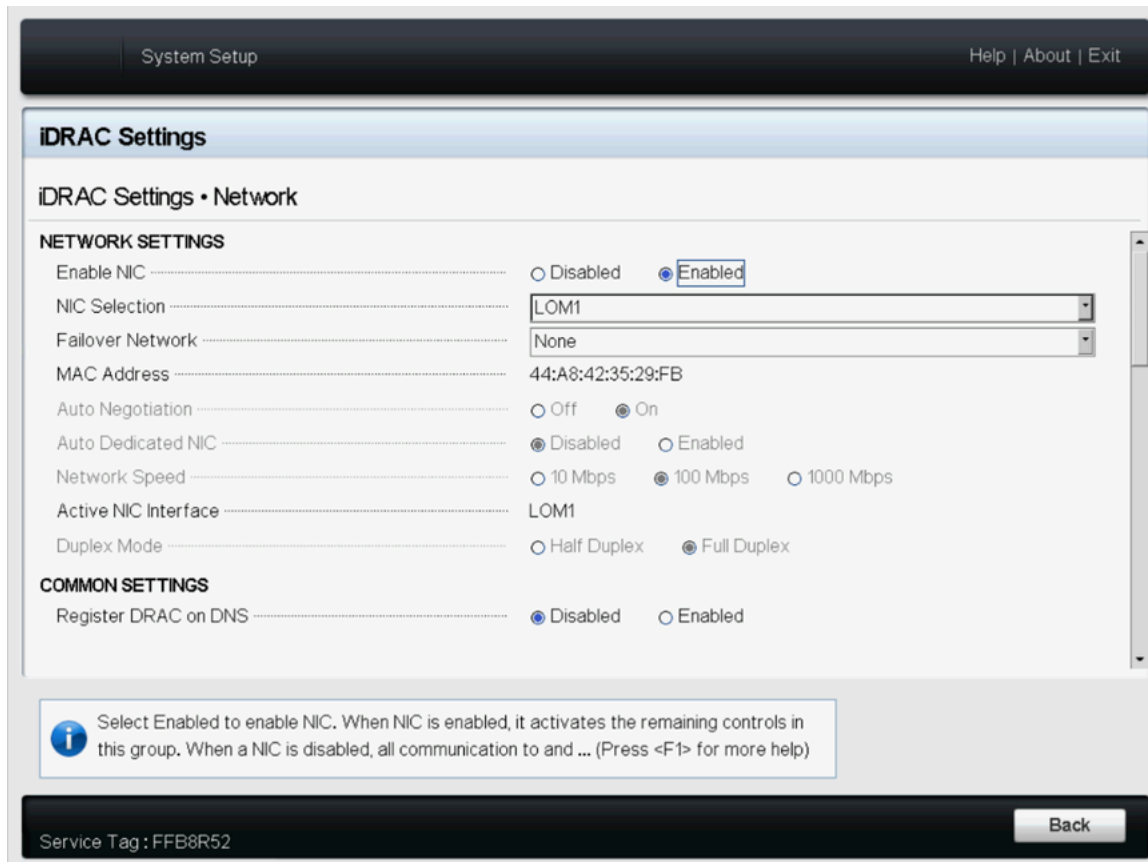
4. Select **iDRAC Settings** from the **System Setup Main Menu** (F2 > iDRAC Settings).

Figure 13-2: iDRAC Settings



5. Use the arrow keys to select **Network**.

Figure 13-3: iDRAC Settings > Network



6. Select an option from the **NIC Selection** list for use with iDRAC. In this example, LOM1 is selected. 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.

- Configure the iDRAC IPv4 address, netmask, gateway, and DNS addresses, as show 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

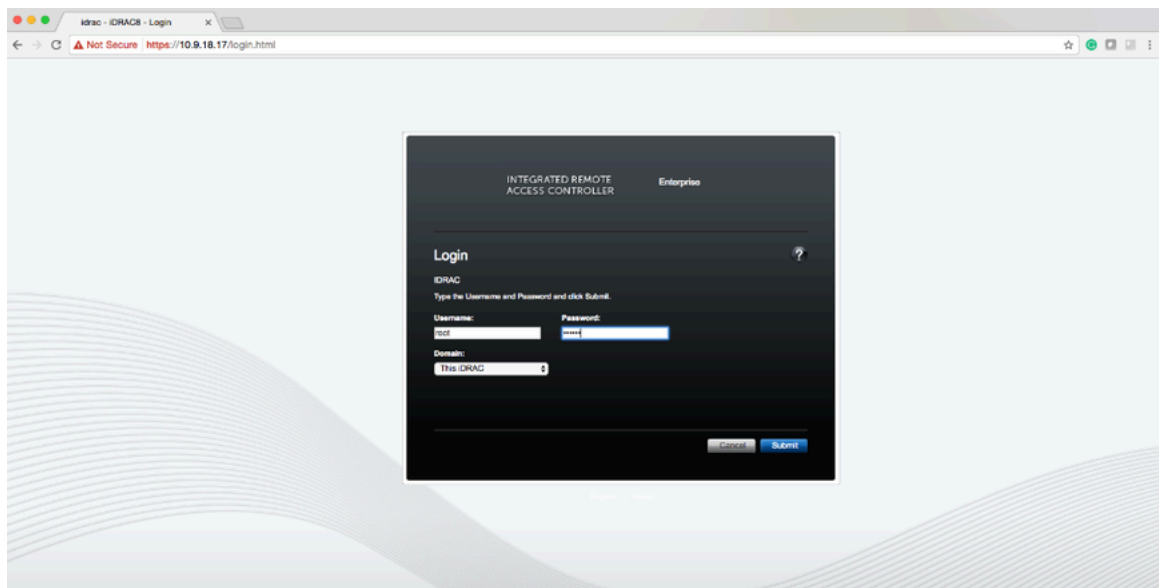
Info Select Enabled to enable NIC. When NIC is enabled, it activates the remaining controls in this group. When a NIC is disabled, all communication to and ... (Press <F1> for more help)

Service Tag : FFB8R52 Back

- After completing the iDRAC configuration, press the **Esc** key to display the Exit menu.
- Select **Save Changes and Exit** and then press **Enter** to keep the changes.
- From a web browser, type **DRAC-IPv4-address** in the browser address bar and press **Enter**.

The system displays the iDRAC web interface, as shown on the following page.

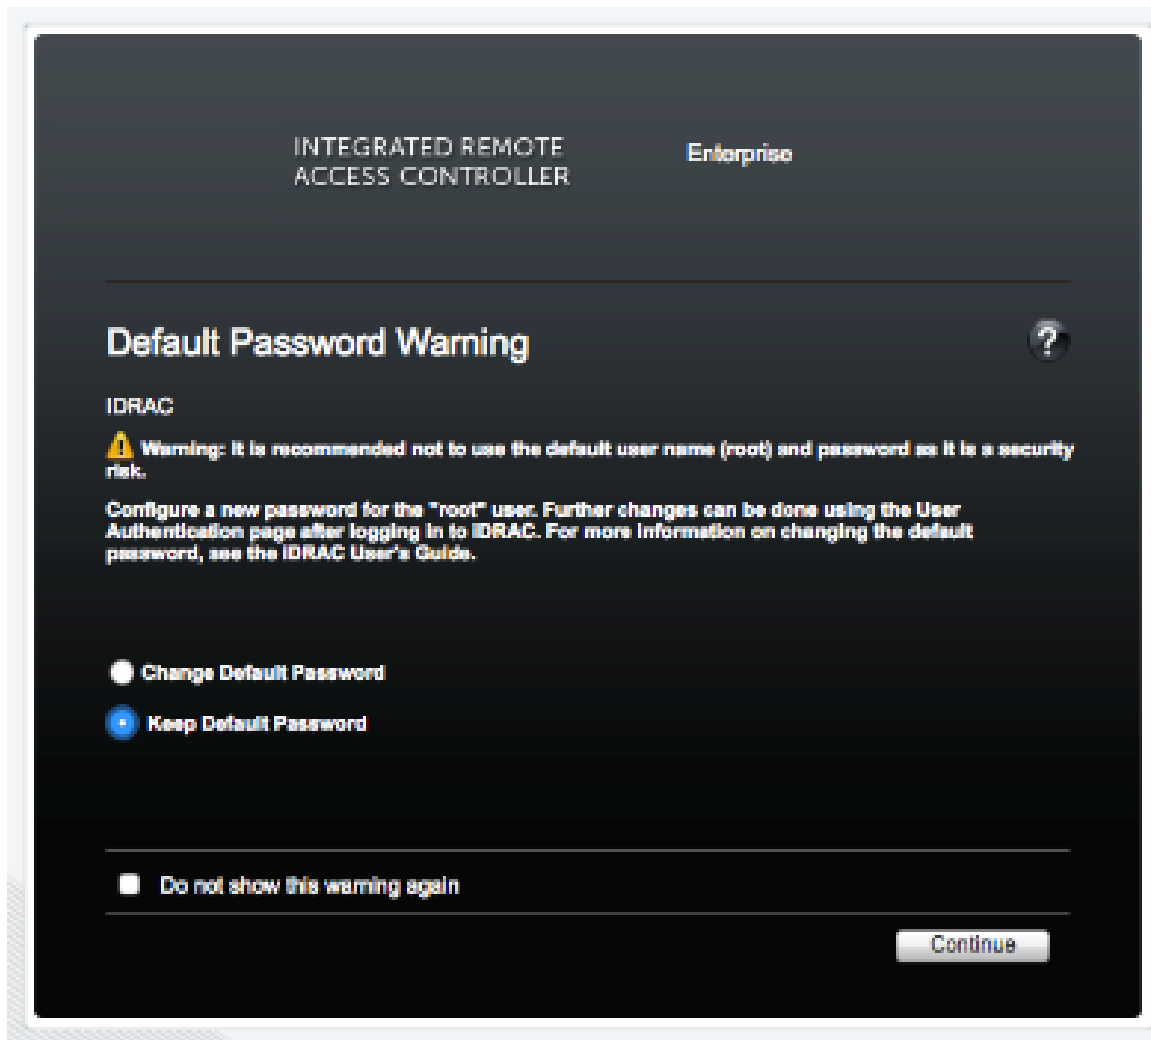
Figure 13-5: iDRAC Web Interface Login Window



11. Enter the user name and password.
The default login credentials are as follows:
 - Username: *root*

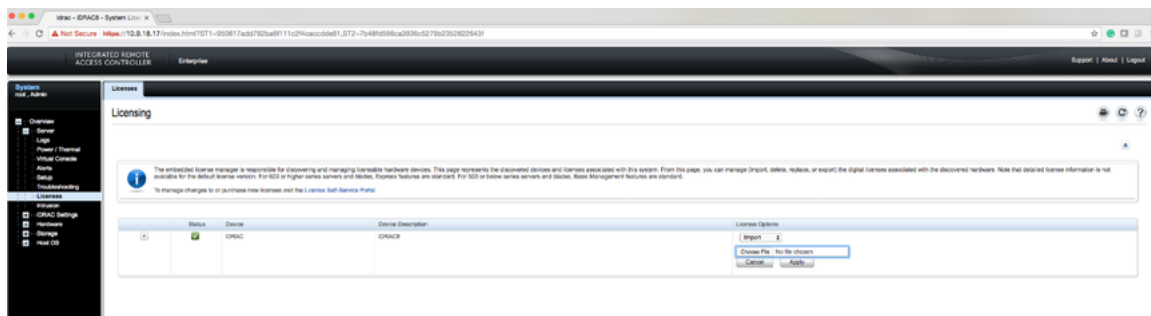
- Password: *calvin*

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 has not been changed.

Figure 13-7: Licensing Page



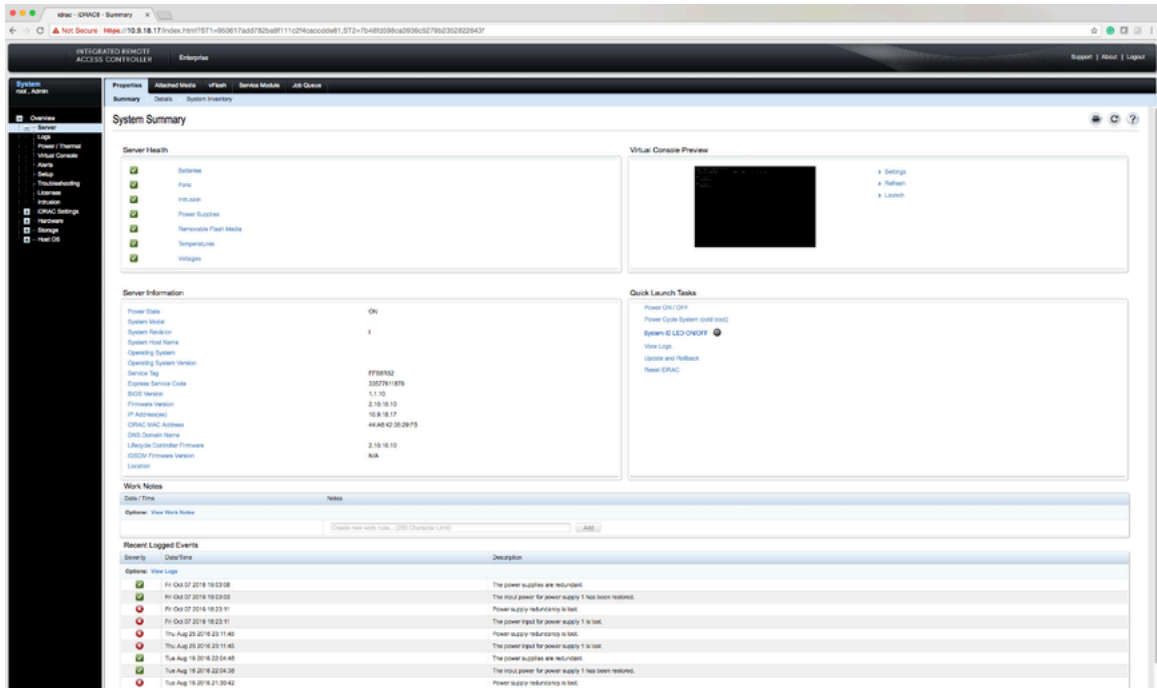
13. To apply the iDRAC Enterprise Version license select the license option and then select import.
14. Select the license file and click **Apply**.
15. After applying the Enterprise Version license, log out from the top right corner for the License to take effect.



Note: You need to log out from the iDRAC web interface, but there is no need to close the web browser.

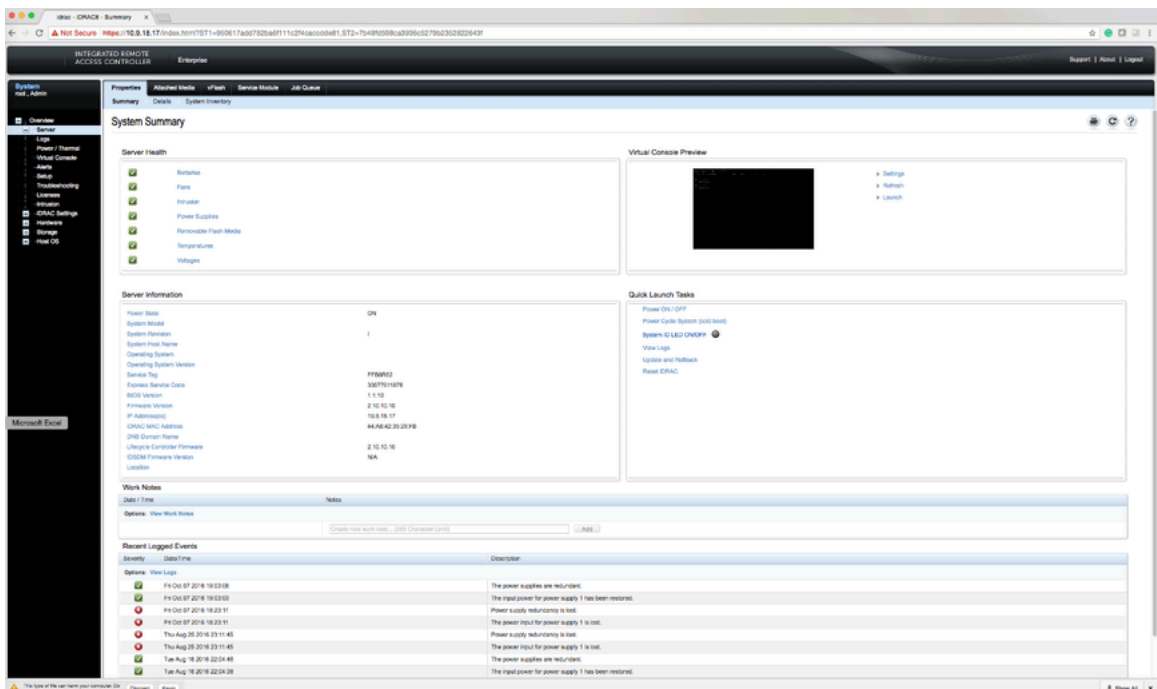
16. Log in to the iDRAC web interface using the new credentials if they were changed. Logging out and logging in after applying the license enables the Enterprise Version license.
17. On the right side of the page, at the **Virtual Console Preview** option, click **Launch** as shown on the following page.

Figure 13-8: System Summary Page



18. When prompted at the bottom of the screen, click **Keep** to confirm the operation.

Figure 13-9: System Summary Page



19. Click the `viewer.jnlp` link at the bottom of the page, 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 components with status indicators (green checkmarks for OK, red X for error). Components include: Batteries, Fans, Inlet Air, Power Supplies, Removable Flash Media, Temperature, and Voltage.
- Virtual Console Preview:** A large black 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	FF84952
Express Service Code	339791679
BIOS Version	1.14.0
Firmware Version	2.10.10.10
IP Address(es)	10.8.10.17
iDRAC MAC Address	4A:4B:42:30:2979
iDRAC Domain Name	
Lifecycle Controller Firmware	2.10.10.10
iDRAC Firmware Version	N/A
Location	
- Quick Launch Tasks:** A list of tasks including Power ON/Off, Power Cycle System (soft boot), System ID LED ON/Off, View Logs, Update and Rollback, and Reset iDRAC.
- Work Notes:** A section for adding notes, with a "Notes" header and a "Add" button.
- Recent Logged Events:** A table of system events:

Severity	Date/Time	Description
OK	Fri Oct 07 2016 18:03:08	The power supplies are redundant.
OK	Fri Oct 07 2016 18:03:08	The input power for power supply 1 has been restored.
OK	Fri Oct 07 2016 18:03:11	Power supply redundancy is lost.
OK	Fri Oct 07 2016 18:03:11	The power input for power supply 1 is lost.
OK	Thu Aug 20 2016 23:11:45	Power supply redundancy is lost.
OK	Thu Aug 20 2016 23:11:45	The power input for power supply 1 is lost.
OK	Tue Aug 16 2016 22:04:48	The power supplies are redundant.
OK	Tue Aug 16 2016 22:04:36	The input power for power supply 1 has been restored.

20. When prompted, open the `viewer.jnlp` file, as shown on the following page.

Figure 13-11: System Summary Page

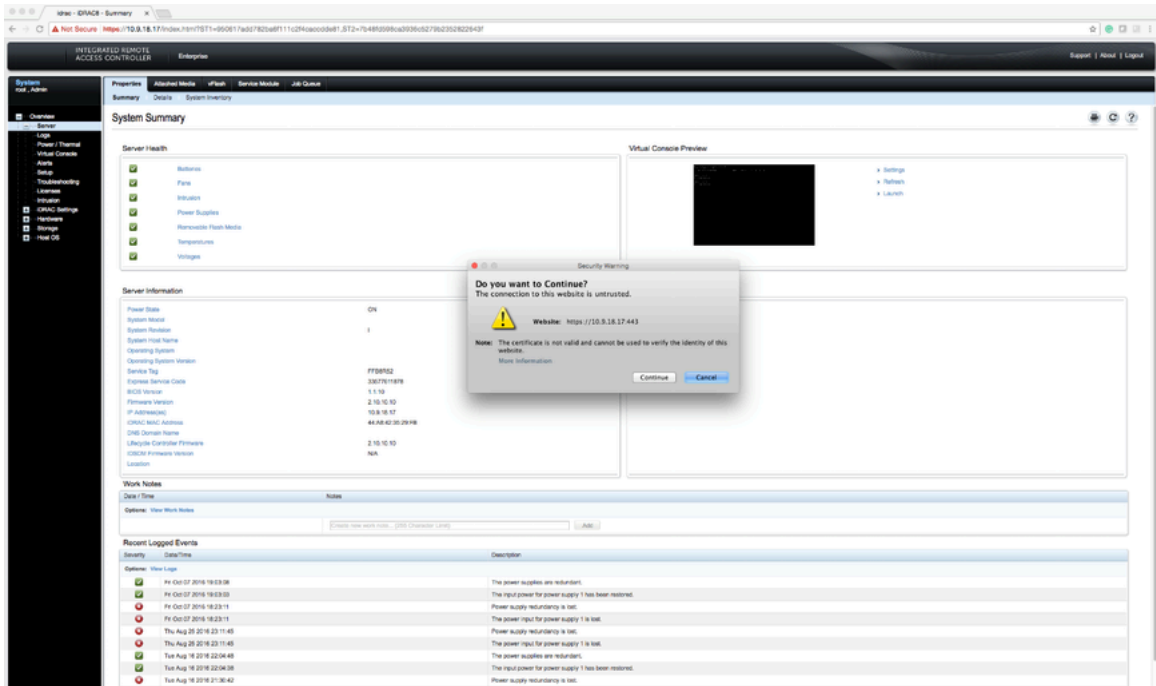
This screenshot is identical to Figure 13-10, but with a Java security warning dialog box overlaid in the center. The dialog box contains the following text:

"viewer.jnlp" is an application downloaded from the Internet. Are you sure you want to open it?
 Details: This file was downloaded from the Internet on 10/07/2016 18:03:11.
 Buttons: Cancel, Show Web Page, Open




Note: JAVA is required and the correct version is automatically downloaded.

Figure 13-12: System Summary: Continue Prompt



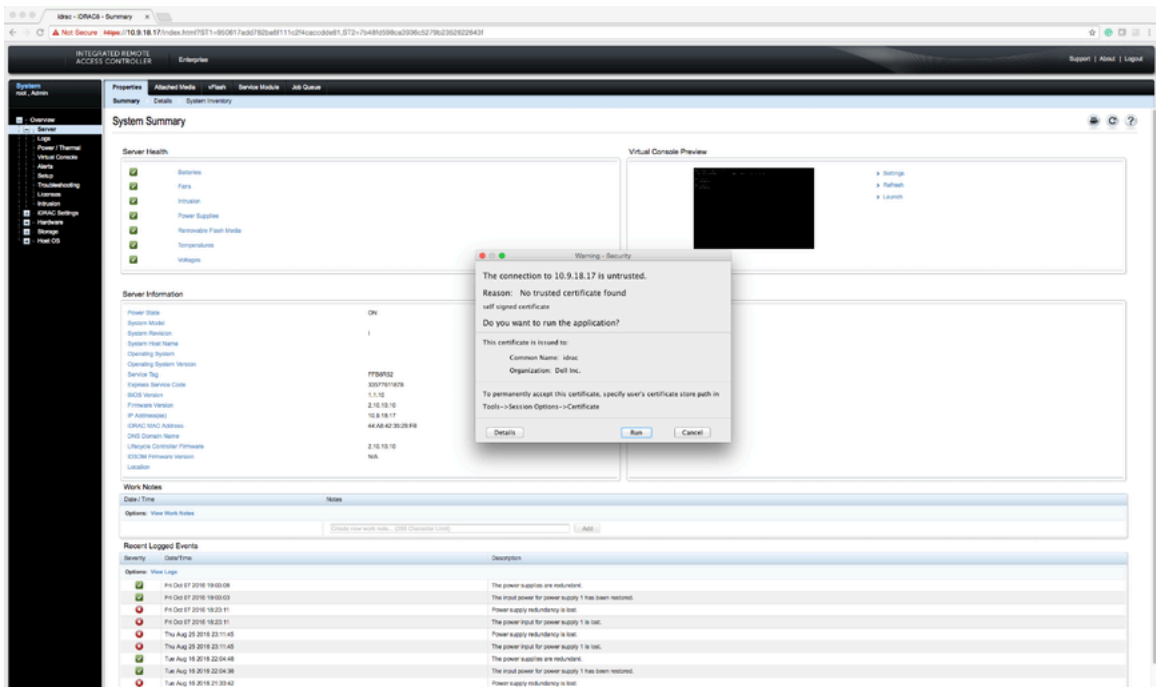
21. When prompted, click **Continue**.

Figure 13-13: System Summary: Run Prompt



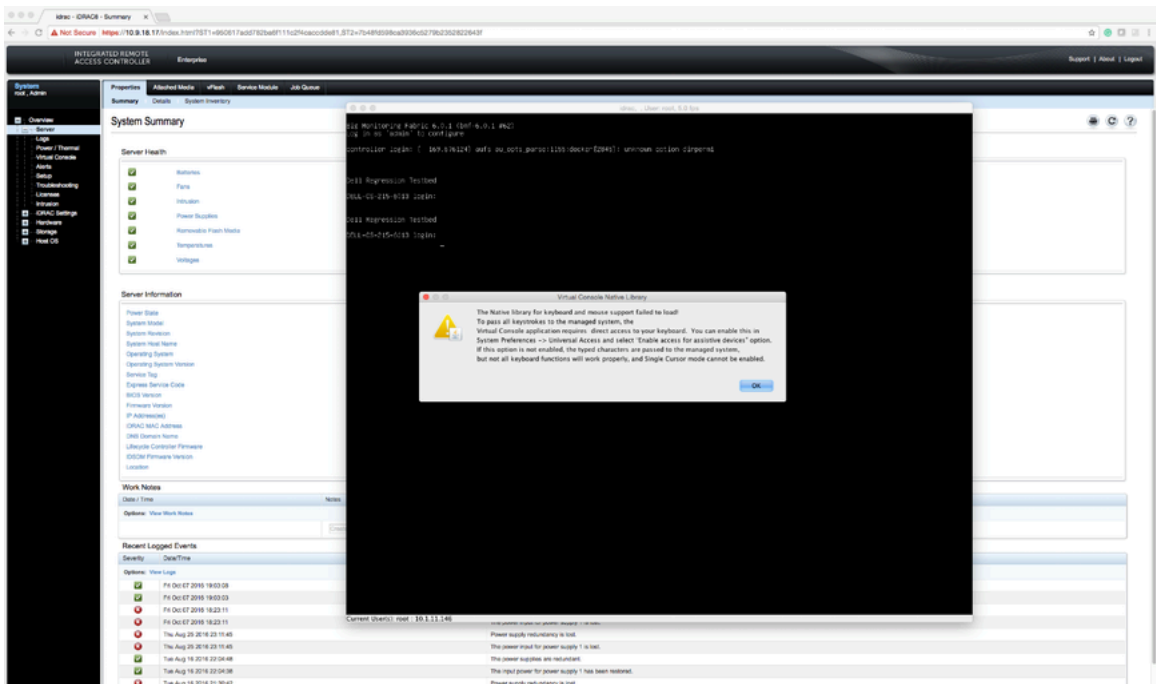
22. When prompted, click Run.

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 given in, [Installing and Configuring the DMF Controller](#) for completing the DMF controller installation and initial configuration.

For **DMF Release 6.3.2** onwards, the procedure for using iDRAC for installing the DMF Controller or DMF Service node image is the same.

To use iDRAC to install the DMF Controller or Service Node image on a Dell R630/R730, complete the following steps.

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** as shown on the following page.

Figure 13-17: Virtual Media > Connect Virtual Media Option




3. Click **Virtual Media** again and this time select **Map CD/DVD**, as shown on the following page.

Figure 13-18: Virtual Media > Map CD/DVD Option



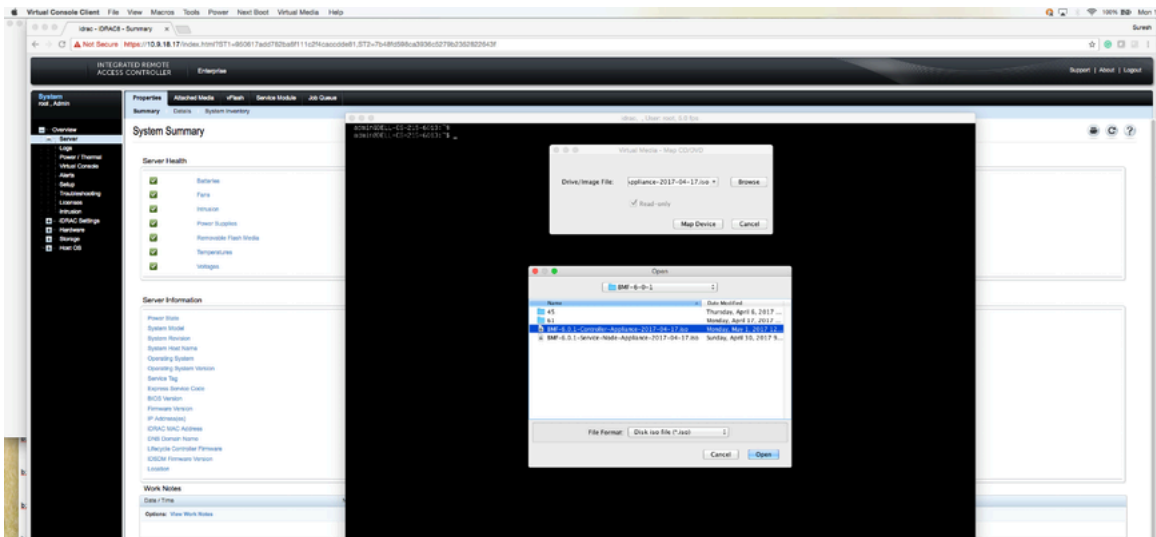
- Click **Browse** to select the DMF Controller ISO file.

Figure 13-19: Virtual Media > Map CD/DVD Browse Option




- Select the DMF Controller ISO image, and click **Open**, as shown on the following page.

Figure 13-20: Open DMF Controller ISO Image



- Click **Map Device**, as shown on the following page.

Figure 13-21: Map Device




This maps the DMF Controller ISO file to a Virtual CD/DVD on the Virtual Media menu, as shown on the following page.

Figure 13-22: Virtual Media > DMF Controller ISO Mapped to a Virtual CD/DVD



7. Click **Next Boot** as shown on the following page.

Figure 13-23: Next Boot Menu



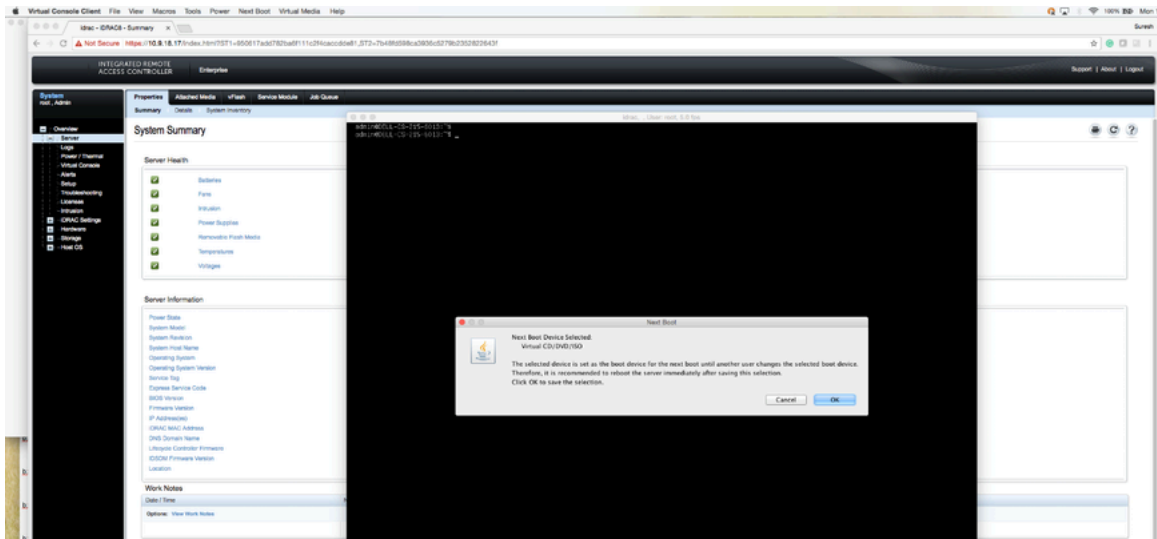
8. Select **Virtual CD/DVD** from the **Next Boot** menu, as shown on the following page.

Figure 13-24: Next Boot > Virtual CD/DVD/ISO Option



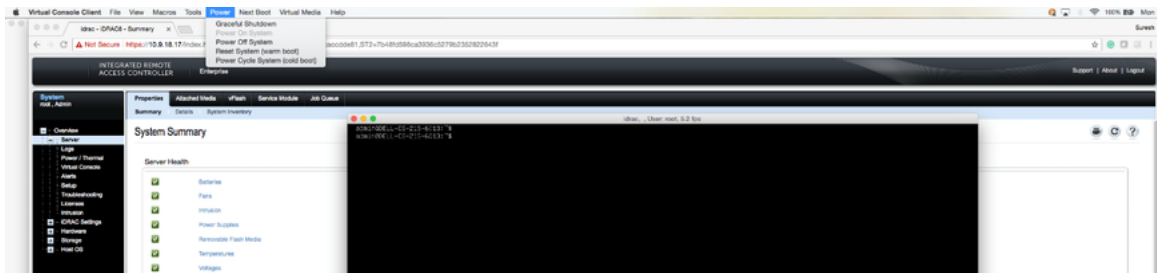
9. When prompted, click **OK**, as shown on the following page.

Figure 13-25: Next Boot Confirmation Prompt




10. Click **Power** and select **Restart System (warm boot)**, as shown on the following page.

Figure 13-26: Power Restart System (warm boot)



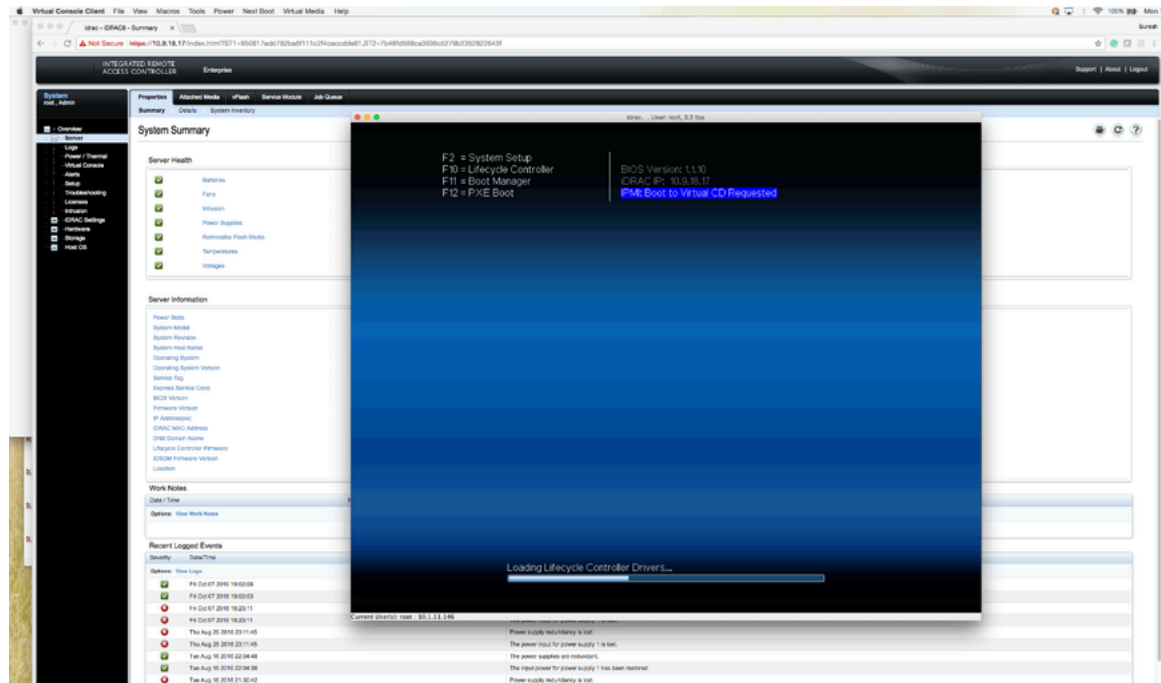
11. When prompted, click **OK**, as shown on the following page.

Figure 13-27: Power Control Confirmation Prompt



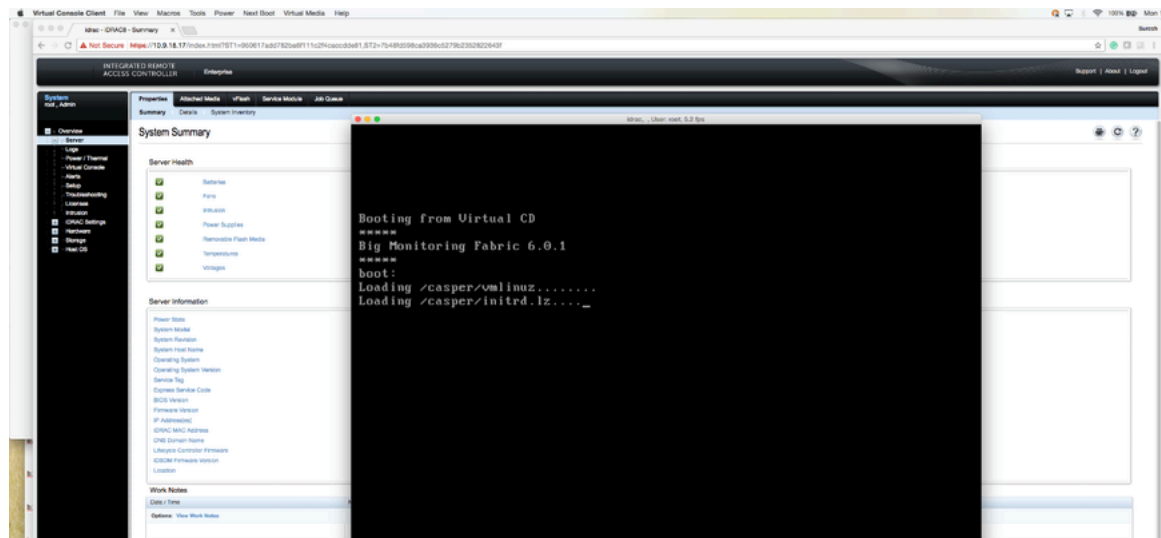
When the server boots up, it now selects the Virtual CD/DVD as the boot device, as show on the following page.

Figure 13-28: Booting from Virtual CD/DVD



The server displays its status as it boots from Virtual CD/DVD, as show on the following page.

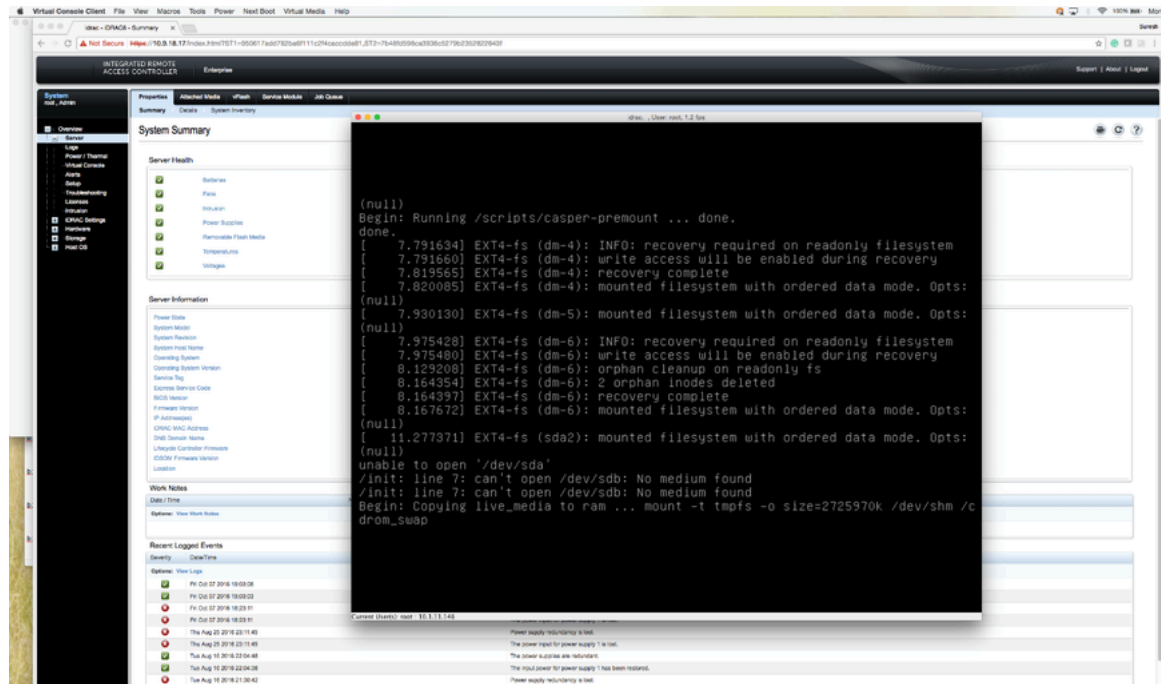
Figure 13-29: Server Boot Status





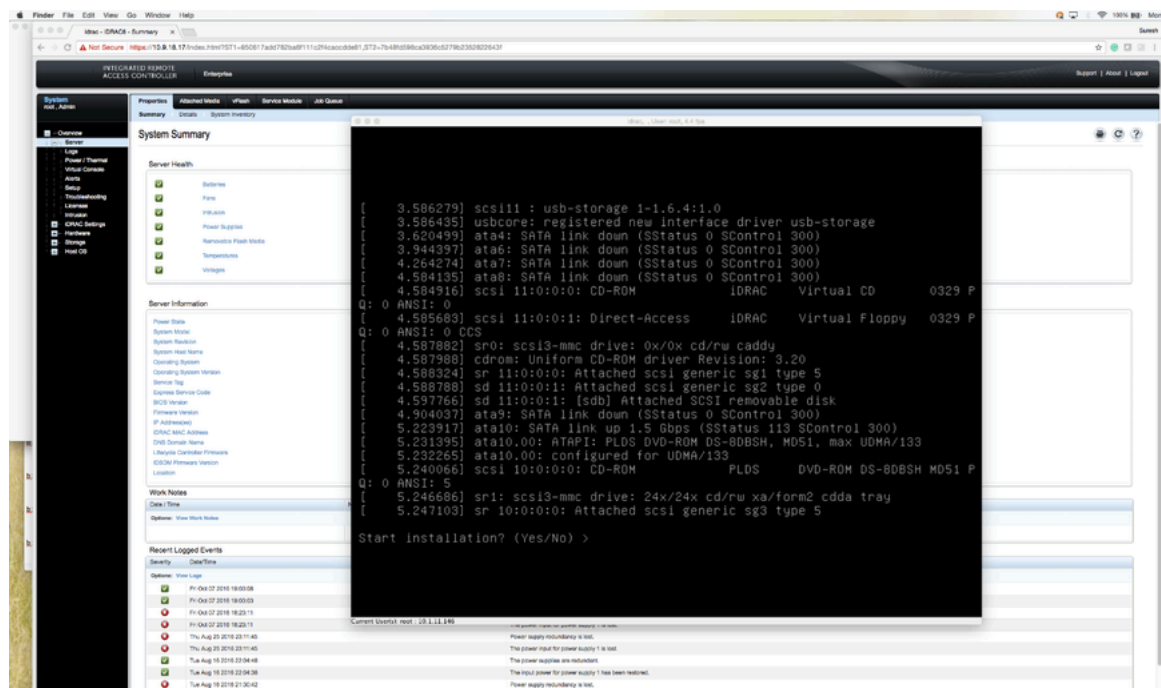
Note: Depending on the network speed, it may take a while to download the ISO image to the server.

Figure 13-30: Server Boot Status



- When prompted, type **yes** to start installing the DMF Controller image on the server, as shown on the following page.

Figure 13-31: Installation Prompt



To complete the installation and initial configuration, refer to the instructions given earlier 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 has been tested using iDRAC Enterprise. The instructions are similar for the following installation options.

- *BMF 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 have been tested using MAC OS but the steps are similar using Windows OS. You can use Internet Explorer, Chrome, or Safari browsers to access the iDRAC web interface.



Note: An iDRAC Enterprise license is included in all of the DMF Controller Nodes, Arista Analytics Nodes, DMF Recorder Nodes, and DMF Service Nodes that use the Dell R440/R640/R740 server hardware.

14.1 Setting Up and Configuring iDRAC

To set up iDRAC on a Dell R440 or R740 server, complete the following steps.

Procedure

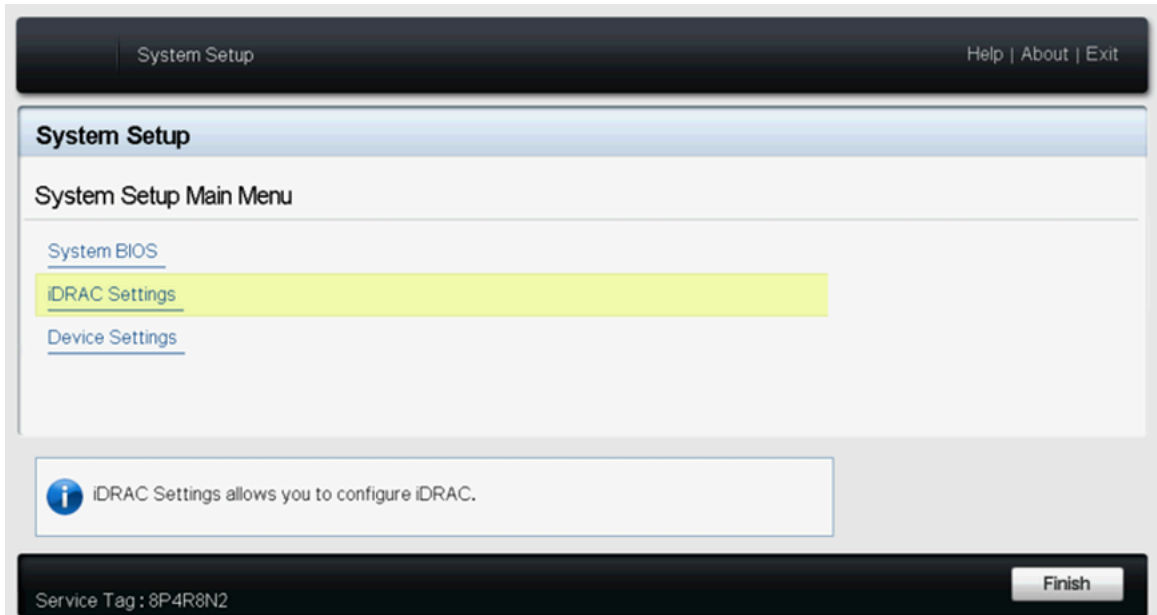
1. Power on the Dell PowerEdge Server.
2. Connect a monitor to the VGA port and connect the keyboard to the USB port.



Note: When using the serial port to connect to the R440 and the R740, the baud rate should be set to **115200**.

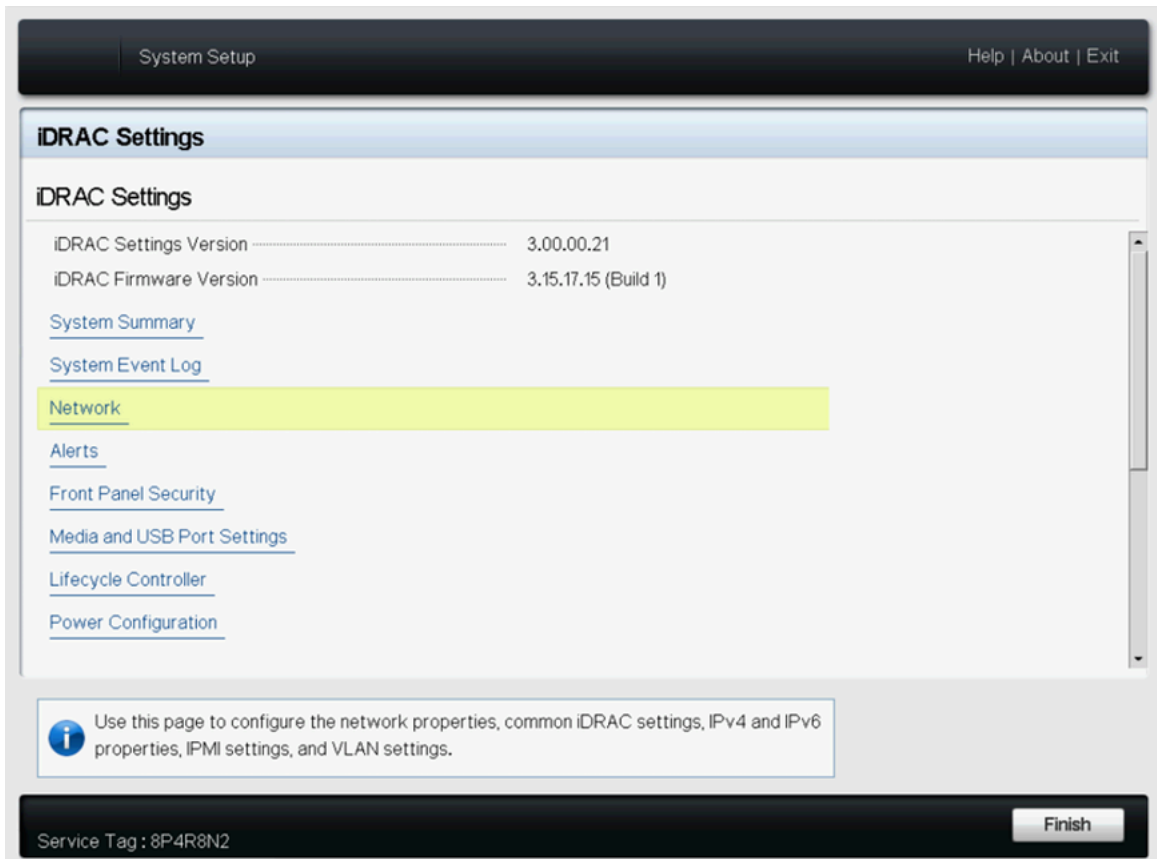
3. Press **F2** to enter the **System Setup Main Menu** screen.

Figure 14-1: System Setup Main Menu



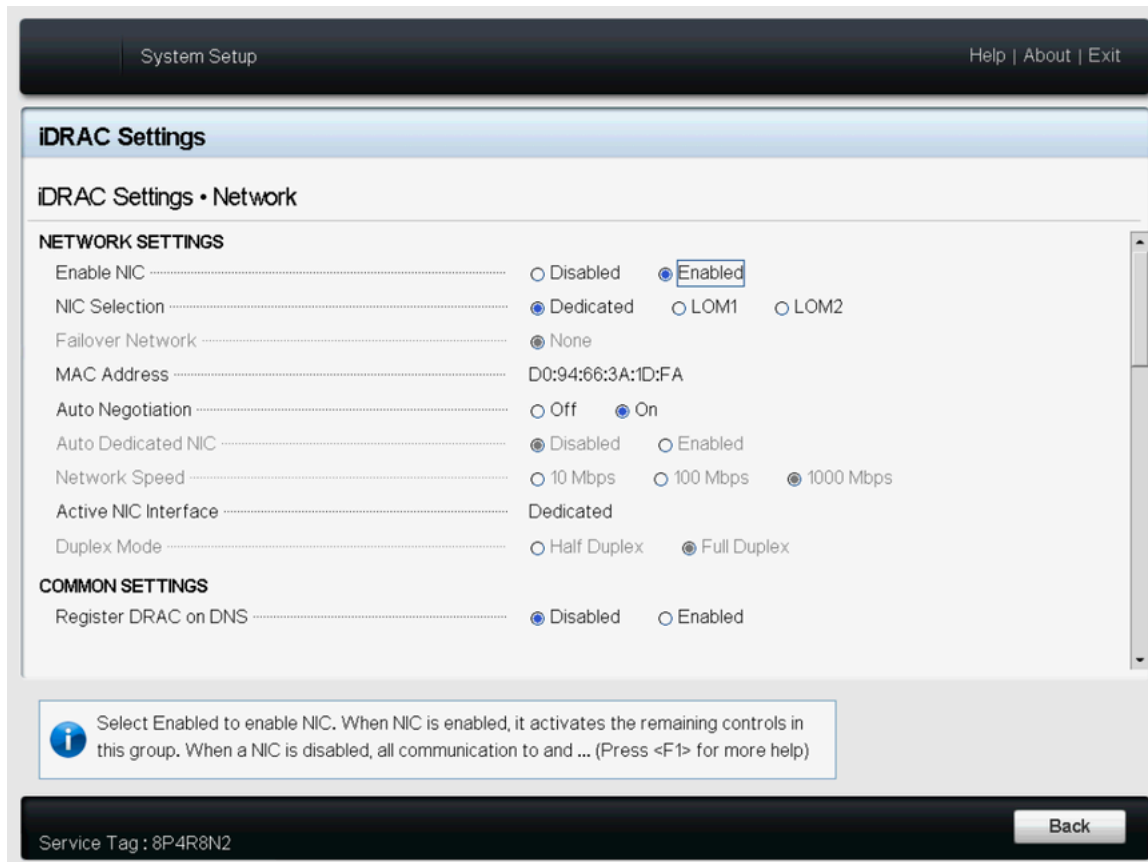
4. Select **iDRAC Settings** from the **System Setup Main Menu** (**F2 > iDRAC Settings**).

Figure 14-2: iDRAC Settings



5. Use the arrow keys to select **Network**.

Figure 14-3: iDRAC Settings Network



6. Select the **Dedicated** option from the NIC selection list for use with iDRAC. Dedicated iDRAC port available on R440 and R740 Dell servers.

- Configure the iDRAC IPv4 address, netmask, gateway, and DNS addresses, as shown on the following page.

Figure 14-4: 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

Service Tag : 8P4R8N2 Back

- After completing the iDRAC configuration, press the **Esc** key to display the Exit menu.
- Select **Save Changes and Exit** and then press **Enter** to keep the changes.
- From a web browser, type **DRAC-IPv4-address** in the browser address bar and press **Enter**. The system displays the iDRAC web interface, as shown on the following page.

Figure 14-5: iDRAC Web Interface Login Window

Integrated Remote Access Controller 9

RAMS-iDRAC | Enterprise

Type the User Name and Password and click Log In.

Username

Password

Domain

Security Notice: By accessing this computer, you confirm that such access complies with your organization's security policy.

Overview | Support | About

- Enter the user name and password.

Default iDRAC Login

The default iDRAC login credentials are as follows:



Note:

- Username: *root*
- Password: *calvin*

Secure Password

For iDRAC9, a new feature called secure password is available.

- The iDRAC secure password is available on the back of the system information tag (Service Tag) under “iDRAC Default Password” as shown in figure below.
- If you have not opted for secure default access to iDRAC, then the default password should be blank. In this case, the default username and password (*root/calvin*) apply.

Figure 14-6: Sticker with secure default password


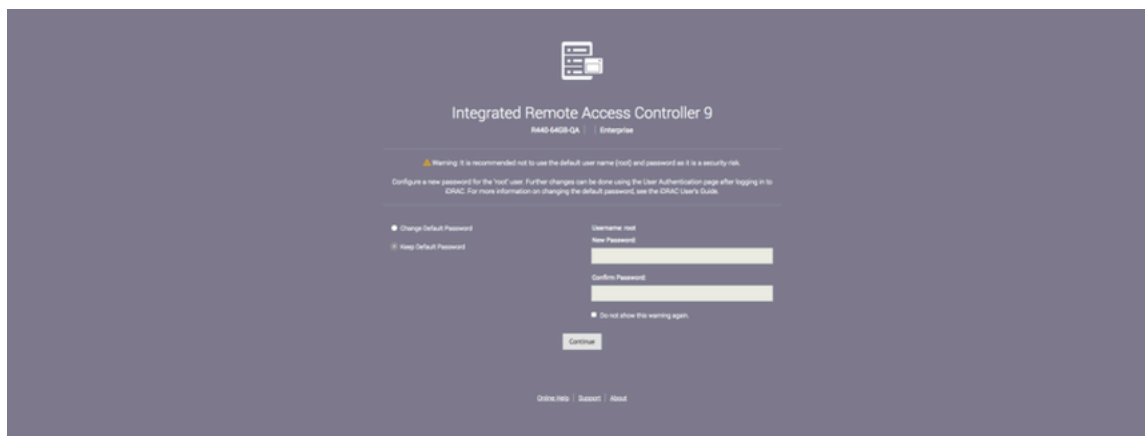


Table 9:

1	Information tag (Top view)	3	OpenManage Mobile (OMM) label
2	Information tag (Bottom view)	4	iDRAC MAC address and iDRAC secure password label

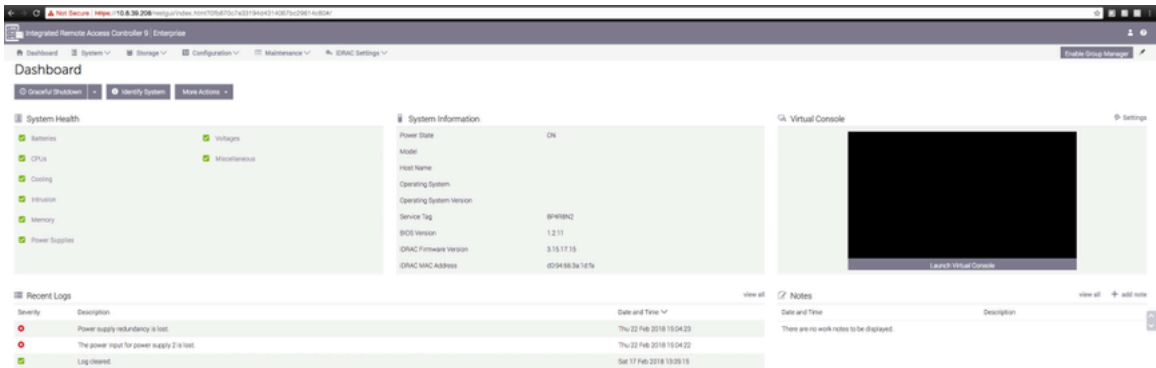
Figure 14-6: Change the Default Password



12. To change the default password, type the new password, when prompted, and type it again to confirm. In this example, the password has not been changed.

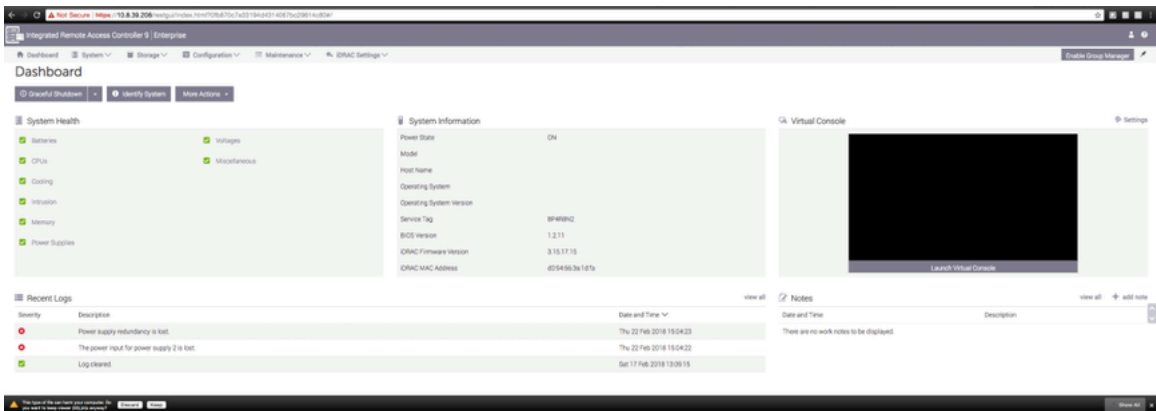
13. Log in to the iDRAC web interface using the new credentials if they were changed.
14. On the right side of the page, at the Virtual Console window option, click **Launch Virtual Console** as shown on the following page.

Figure 14-8: System Summary Page



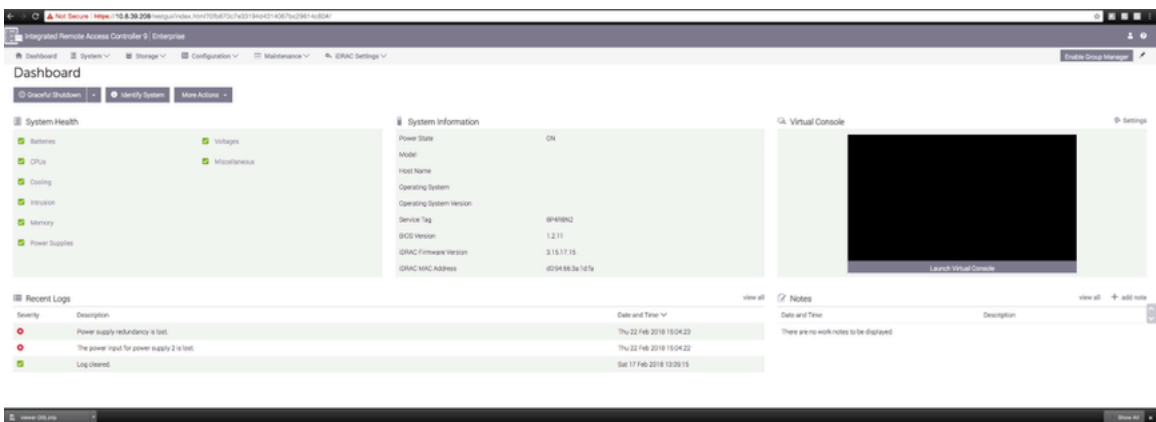
15. When prompted at the bottom of the screen, click **Keep** to confirm the operation.

Figure 14-9: System Summary Page



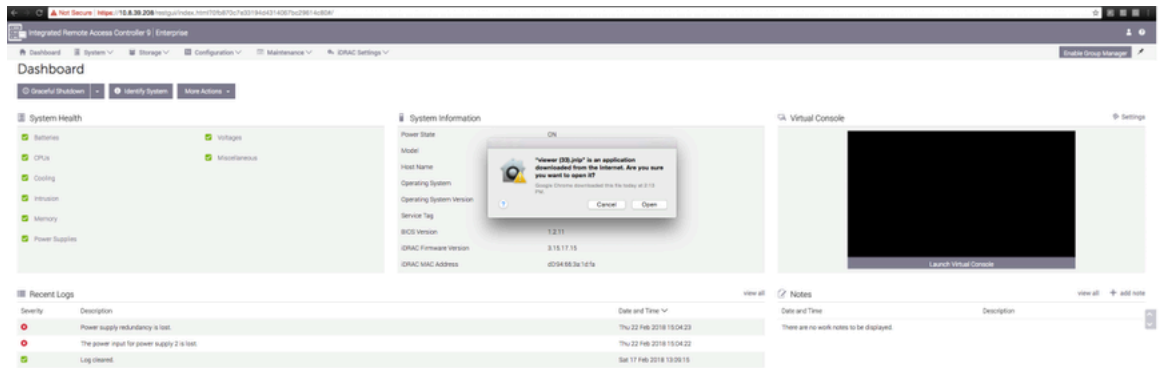
16. Click the `viewer.jnlp` link at the bottom of the page, as shown in the following example.

Figure 14-10: System Summary Page




17. When prompted, open the `viewer.jnlp` file, as shown on the following page.

Figure 14-11: System Summary Page



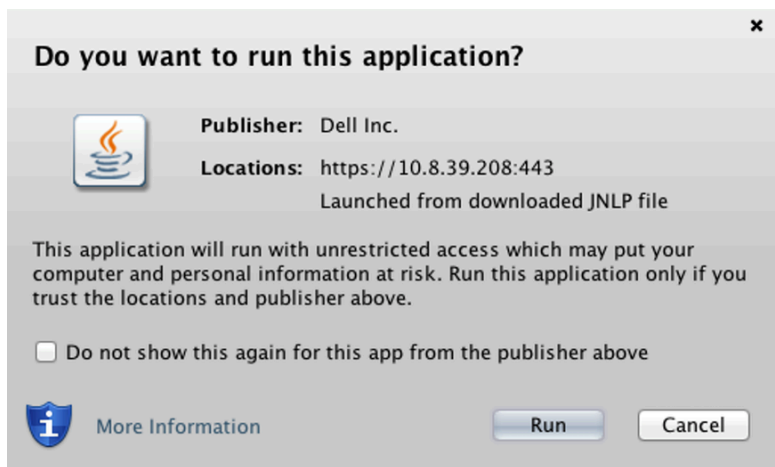
Note: Java is required and the correct version is automatically downloaded.

Figure 14-12: System Summary: Continue Prompt




18. When prompted, click **Continue**.

Figure 14-13: System Summary: Run Prompt




19. When prompted, click **Run**.

Figure 14-14: Confirm Connection to Untrusted Network



20. When prompted to continue with an untrusted certificate, click **Run**.

Figure 14-15: 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 as described in the previous chapters in this document.

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 14-16: iDRAC Virtual Console Window




2. Select **Virtual Media > Connect Virtual Media** as shown on the following page.

Figure 14-17: Virtual Media Connect Virtual Media Option



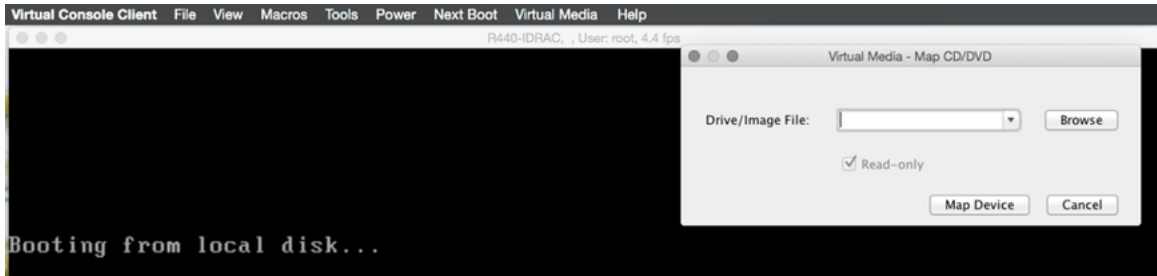
3. Click **Virtual Media** again and this time select **Map CD/DVD**, as shown on the following page.

Figure 14-18: Virtual Media Map CD/DVD Option




- Click **Browse** to select the ISO file you want to install.

Figure 14-19: Virtual Media Map CD/DVD Browse Option



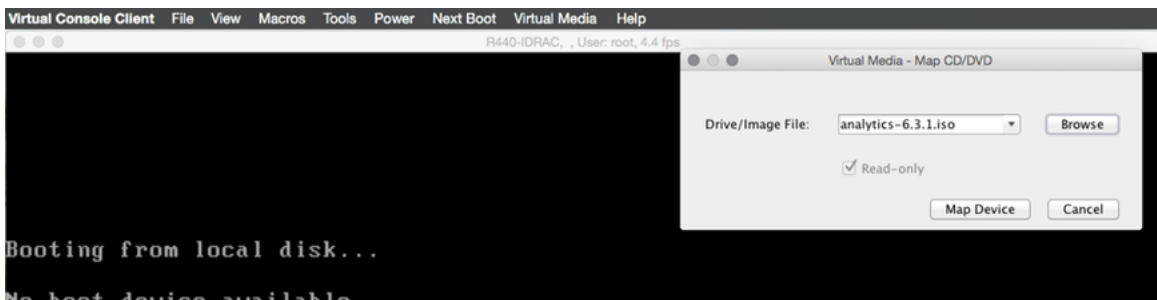
- Select the Analytics, DMF Controller, or Recorder software image you want to install and click **Open**. The following example shows selecting the Arista Analytics image.

Figure 14-20: Open DMF Controller ISO Image



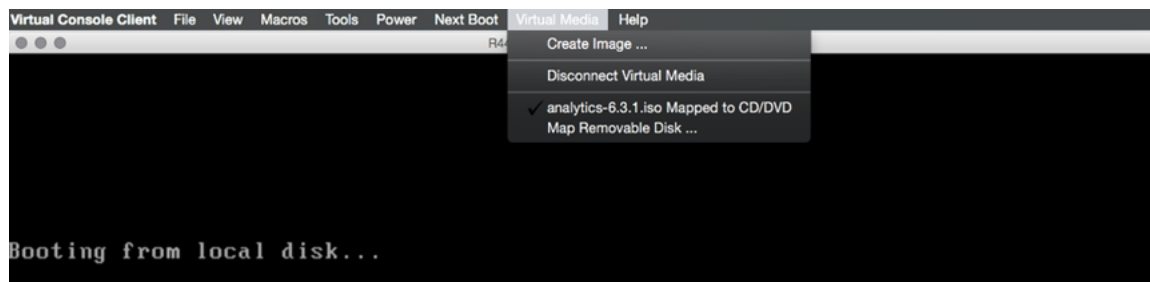
- Click **Map Device**, as shown on the following page.

Figure 14-21: Map Device




This maps the DMF Controller ISO file to a Virtual CD/DVD on the Virtual Media menu, as shown on the following page.

Figure 14-22: Virtual Media DMF Controller ISO Mapped to a Virtual CD/DVD



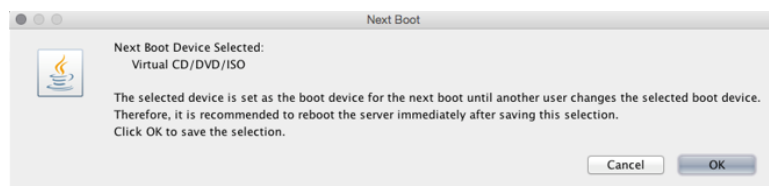
7. Click **Next Boot** as shown on the following page.

Figure 14-23: Next Boot Menu



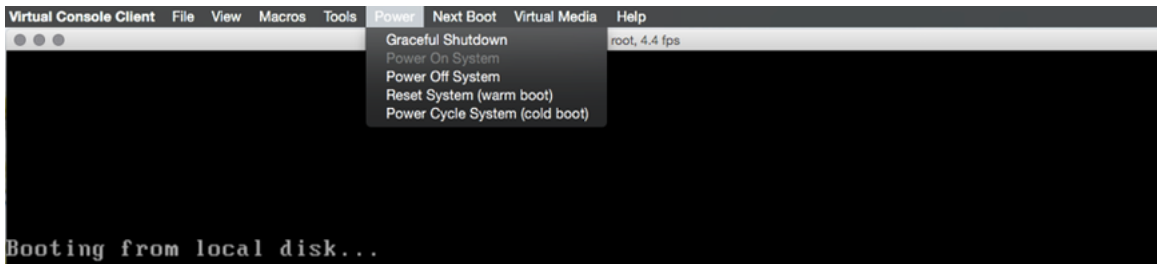
8. Select **Virtual CD/DVD** from the **Next Boot** menu.
9. When prompted, click **OK**, as shown on the following page.

Figure 14-24: Next Boot Confirmation Prompt



- Click **Power** and select **Restart System (warm boot)**, as shown on the following page.

Figure 14-25: Power Restart System (warm boot)



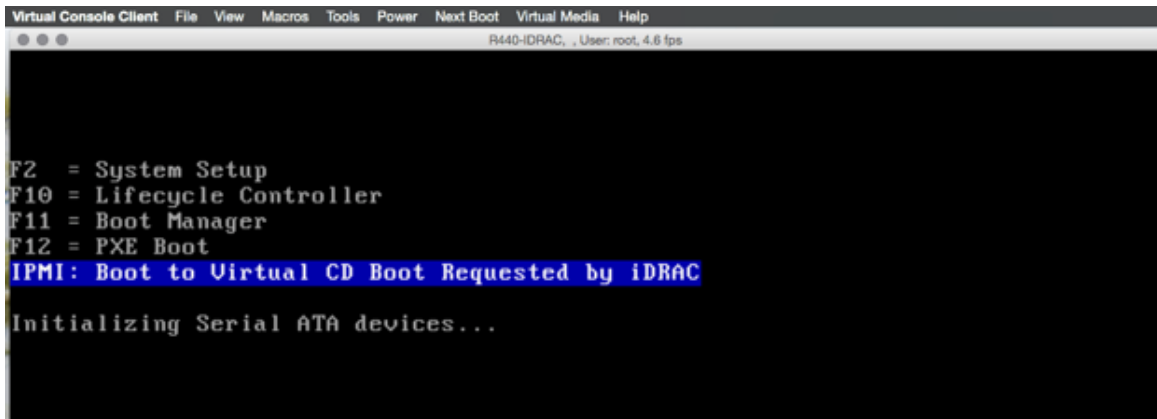
- When prompted, click **OK**, as shown on the following page.

Figure 14-26: Power Control Confirmation Prompt



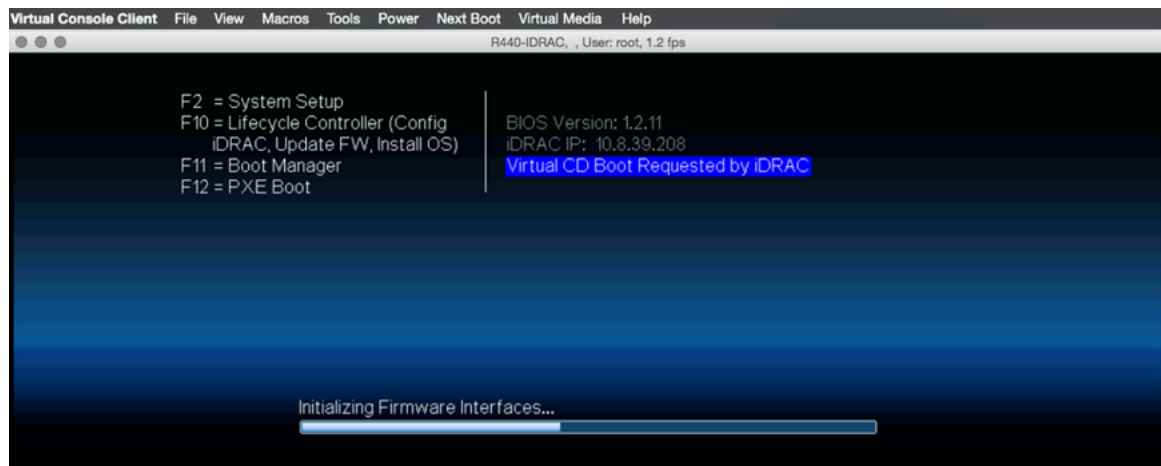
When the server boots up, it now selects the Virtual CD/DVD as the boot device, as shown on the following page.

Figure 14-27: Booting from Virtual CD/DVD



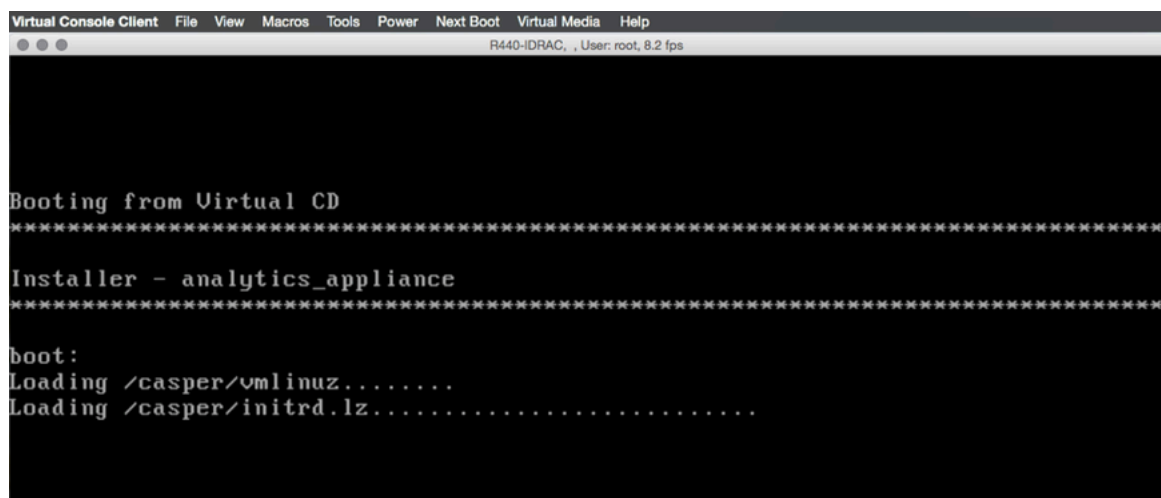
The server displays its status as it boots from Virtual CD/DVD, as show on the following page.

Figure 14-28: Server Boot Status



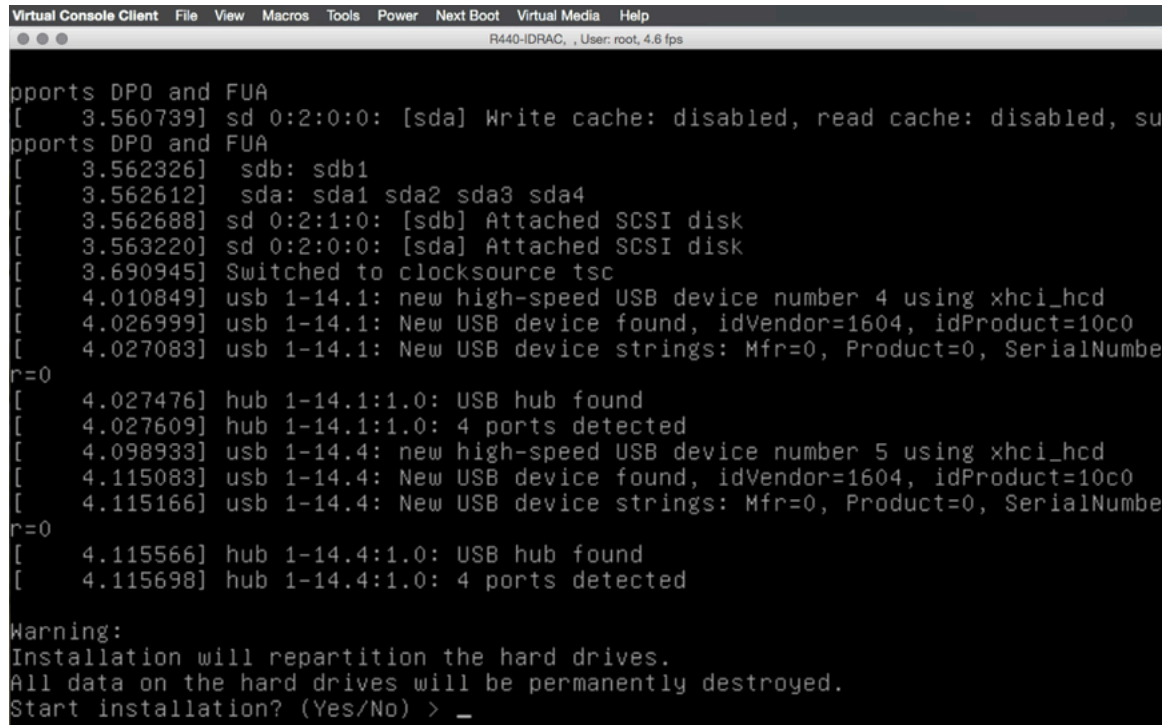
Note: Depending on the network speed, it may take a while to download the ISO image to the server.

Figure 14-29: Server Boot Status



- When prompted, type **yes** to start installing the DMF Controller image on the server, as shown on the following page.

Figure 14-30: 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 upgrade 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>.

Certain Dell switches support DMF controller assisted firmware update. Refer to HCL guide for identifying controller assisted firmware package for the specific platforms.

To install/upgrade the firmware, complete the following steps:

```
CTL (config) # system install switch <name> onie|cpld reboot
```


Switch ONIE Upgrade Procedure

For detailed instructions update 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 that is typically pre-installed as firmware on bare metal network switches. It provides an environment for automated provisioning.

Certain Dell switches support DMF controller assisted firmware update. Refer to HCL guide for identifying controller assisted firmware package for the specific platforms.

To install/upgrade the firmware, complete the following steps:

```
CTL(config) # system install switch <name> onie|cpld reboot
```

Removing existing OS from switch

This appendix provides information on how to uninstall OS from 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 switch via serial connection and confirm that an SWLOS is installed on the switch. Serial console output of the switch should look something like this.

```
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 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 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 **/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 existing EOS swi which should be under **/mnt/flash** directory.

```
Aboot# vi boot-config
Aboot#
SWI=flash:aboot-chainloader.swi
~
SWI=flash:/EOS-4.23.3M.swi
~
Aboot# cat boot-config
SWI=flash:/EOS-4.23.0F.swi
```

7. Reboot the switch. Switch should boot-up with EOS image.

```
Aboot# reboot
Requesting system reboot
Press Control-C now to enter Aboot 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 step:

1. Log in to the switch via 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 Aboot by interrupting 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
Aboot 9.0.3-4core-14223577
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
```

3. Once in Aboot, change directory to **/mnt/flash**.

```
Press Control-C now to enter Aboot shell
^CWelcome to Aboot.
Aboot#
Aboot# cd /mnt/flash
```

```
Abboot# pwd
/mnt/flash Abboot#
```

4. List the files under `/mnt/flash` directory and check the `boot-config` file for current swi.

```
Abboot# pwd
/mnt/flash
Abboot# ls
AsuFastPktTransmit.log
EOS-4.25.2F.swi
EOS-4.27.2F-26021868.uppsaladmfirel-i686.swi
EOS-4.27.2F-26021868.uppsaladmfirel-i686.swi.tmp
Fossil
SsuRestore.log
SsuRestoreLegacy.log
boot-config
debug
fastpkttx.backup
i686
lost+found
persist
schedule
startup-config
zerotouch-config
ztn-boot-info
Abboot#
Abboot# cat boot-config
SWI=flash:/EOS-4.27.2F-26021868.uppsaladmfirel-i686.swi Abboot#
```

5. Remove the `startup-config` file and edit the `boot-config` file to point to existing EOS swi which should be under `/mnt/flash` directory.

```
Abboot# rm -rf startup-config
Abboot# vi boot-config
Abboot#
SWI=flash:aboot-chainloader.swi
~
SWI=flash:/EOS-4.25.2F.swi
~
Abboot# cat boot-config
SWI=flash:/EOS-4.25.2F.swi
```

6. Reboot the switch. Switch should boot-up with EOS image.

```
Abboot# reboot
Requesting system reboot
Press Control-C now to enter Abboot 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 the Switch Light OS is installed 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, first 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 at the switch GNU GRUB boot menu, the switch does not have an existing OS installed. The following example shows an example of 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 section that follows to install Switch Light OS.

Procedure

To delete FTOS from a Dell switch, complete the following steps:

1. Confirm that an OS is already installed on the switch.
2. Another OS is installed if other options besides ONIE are displayed at 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**.

Uninstall 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>
```

When the process is complete the following messages are displayed 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

To create a bootable USB drive on MacOS X, complete the following steps.

Procedure

1. Insert the USB drive into a USB port on the Macintosh.
This automatically mounts the 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.

Use this command to copy the Service Node appliance ISO image to the USB drive. 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
```

It can take up to ten minutes to copy the image to the USB drive.

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

To create a bootable USB drive on using Linux, complete the following steps.

Procedure

1. Insert the USB drive into a USB port on the Linux workstation.
2. In a Linux terminal window, enter the following command 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
```

It can take up to ten minutes to copy the image to the USB drive.

5. Eject the USB drive from your Linux workstation.

B.3 Creating a USB Boot Image Using Windows

A number of 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/>.

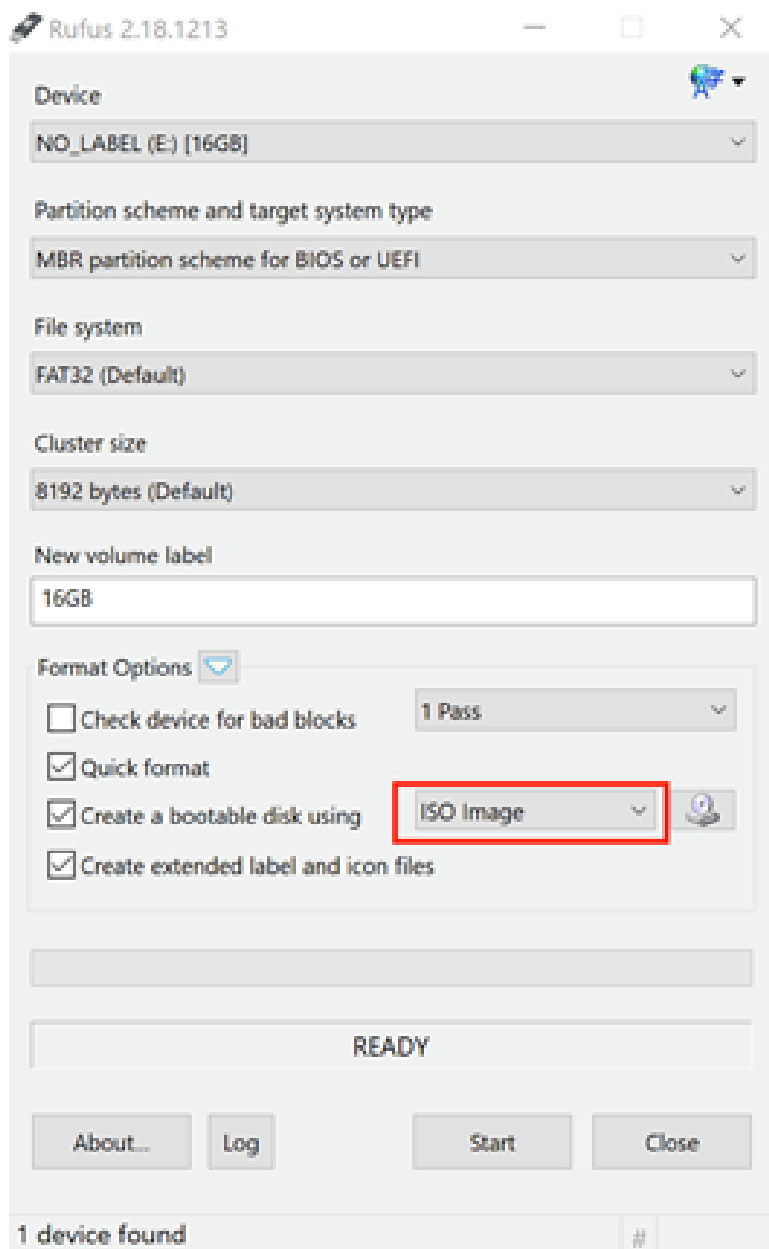
2. After downloading the utility, double-click the **rufus.exe** file.

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



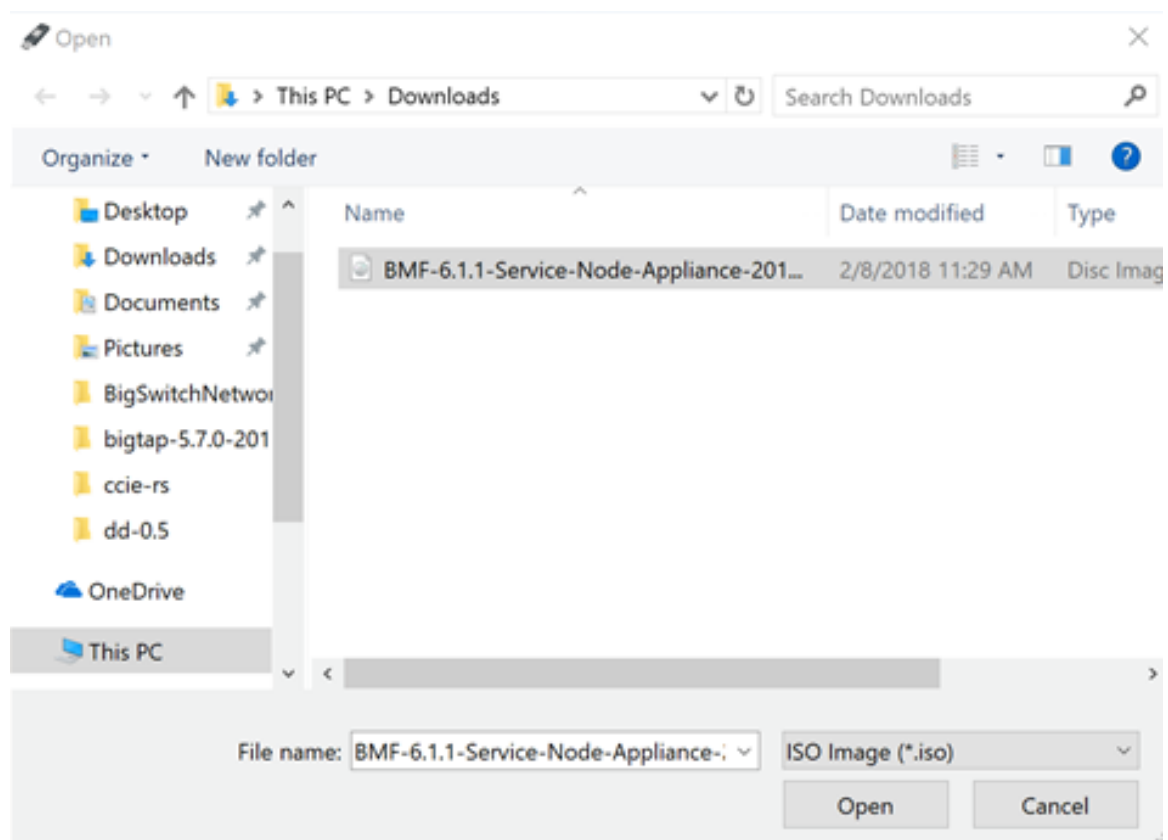
4. To create a bootable disk select **ISO Image**.

Figure B-3: Rufus: Select ISO Image



5. Click the **CD-ROM** icon.

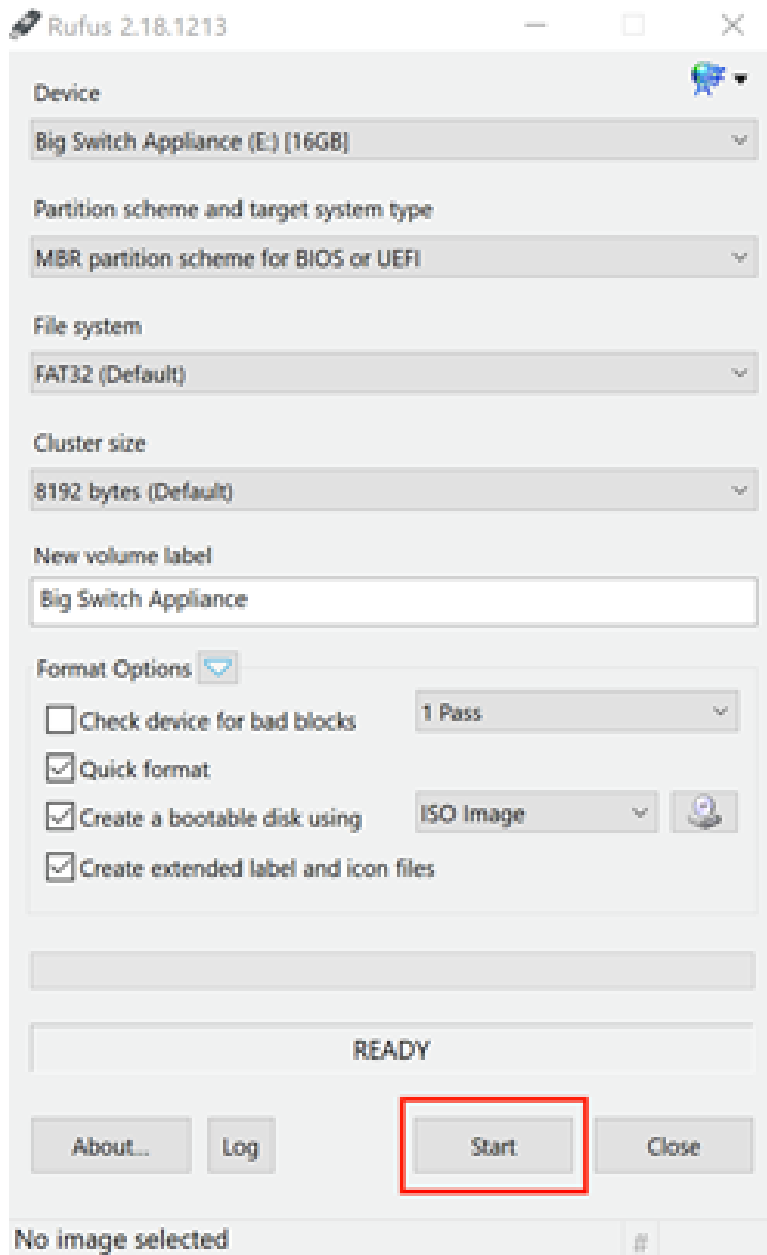
Figure B-4: Open ISO Image File



6. Select the file to use and click **Open**.

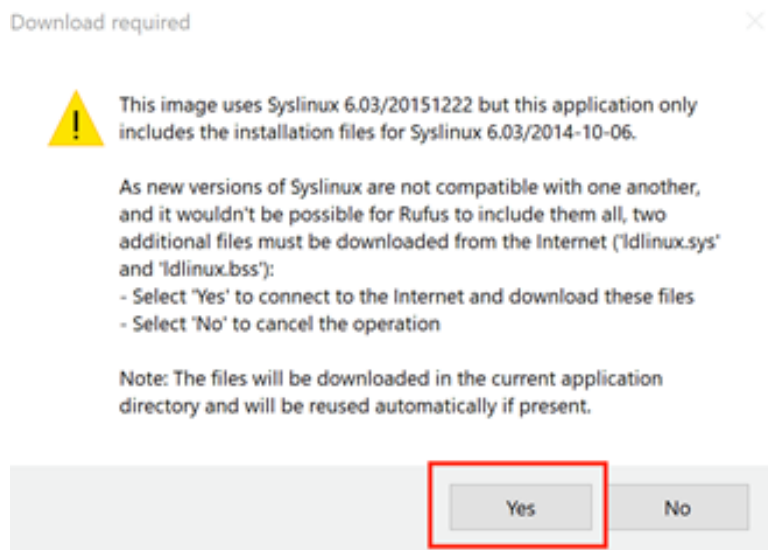
7. Click **Start** to burn the ISO image to USB.

Figure B-5: Rufus: Start



If 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 is going to be destroyed and a new image is going to 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 KVM or ESXi version are as follows:

- 4 vCPUs with a minimum scheduling of 1 GHz.
- 8 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

- DMF controller image in OVA format (.ova file).
- ESXi host, can be part of vSphere cluster. Refer to *DMF Hardware Compatibility List* for supported versions of ESXi and vCenter.
- A virtual network already created.
- A machine with vSphereClient or webClient.



Note: Do not use *ESXi 5.1 GA*. It has a known issue where installing a large VM causes ESXi host to crash. Check the VMware web site for more information.

C.2.2 VM Installation

To install the VM, complete the following steps:

1. Login to vCenter or ESXi host with vSphere Client. The following example is based on an vCenter 6.7.
2. Right click on host where Controller VM is to 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 provision format at the default, and select the datastore then click **Next**.
7. Select **Network Mapping**. You should map to the network created for DMF controller and switches. Click **Next**.
8. Click **Next** and click **Finish**.
9. Start the VM and complete the steps shown 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 the Controller pairs starting from **DMF 8.2**. 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 backup 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**, choose **Create a new virtual machine**.
3. Provide a name for the new VM and click **Import existing disk image** options. Then click **Forward** to continue.
4. Set the existing storage path to point to the provided DMF controller image. Then 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 can sometimes 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. Now, follow the steps in [Installing and Configuring the DMF Controller](#) to set up the controller.


Erasing DMF Appliance

To securely erase data stored on DMF appliance, it must be overwritten. This section describes how to do this using Dell LifeCycle Controller. For the DMF Recorder Node there is a erase API, but it does not securely remove the data from its disk. Instead, it unlinks files in the Index and Packet partitions, so the space can be reclaimed by the file system for future packets and indices. This was a design decision, because unlinking files is very fast compared to overwriting data. For data that needs to be erased securely, to prevent anyone from accessing it use the below procedure.

D.1 Using the Dell LifeCycle Controller

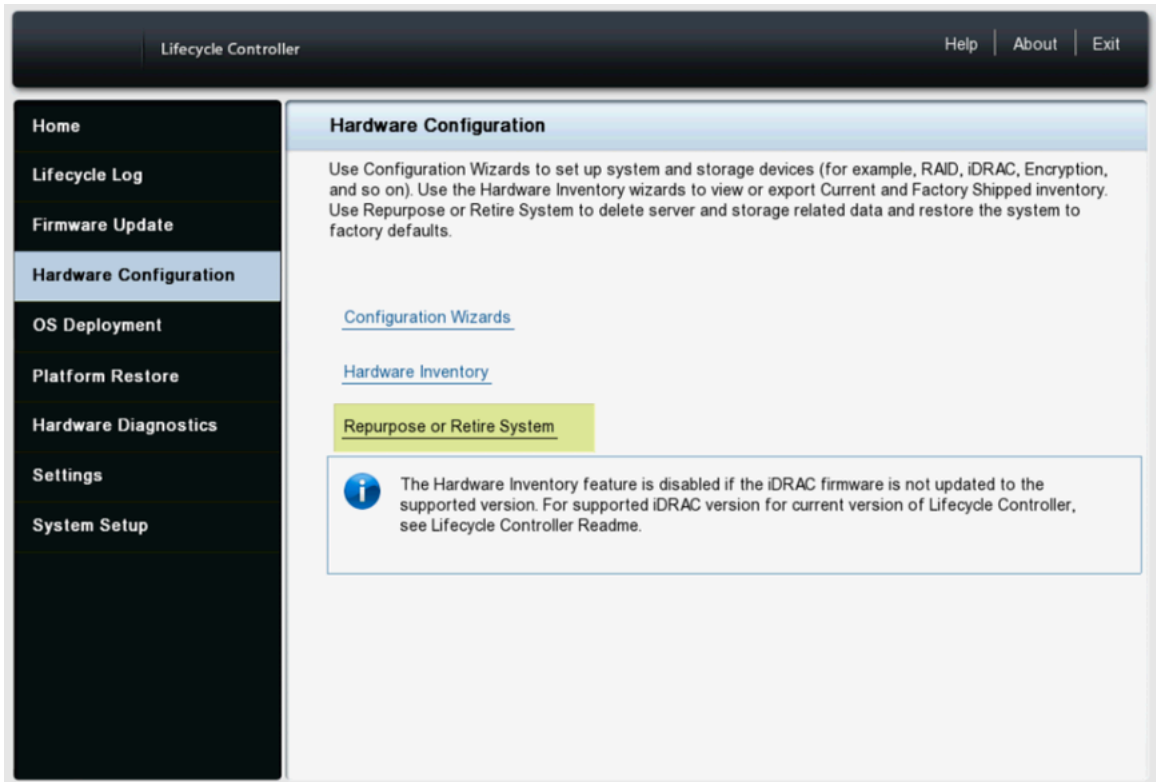
1. Restart the Recorder Node.
2. During POST, press **F10** to start LifeCycle Controller GUI.

Figure D-1: Dell Lifecycle Controller



3. From left pane select **Hardware Configuration**. Click on **Repurpose or Retire System**.

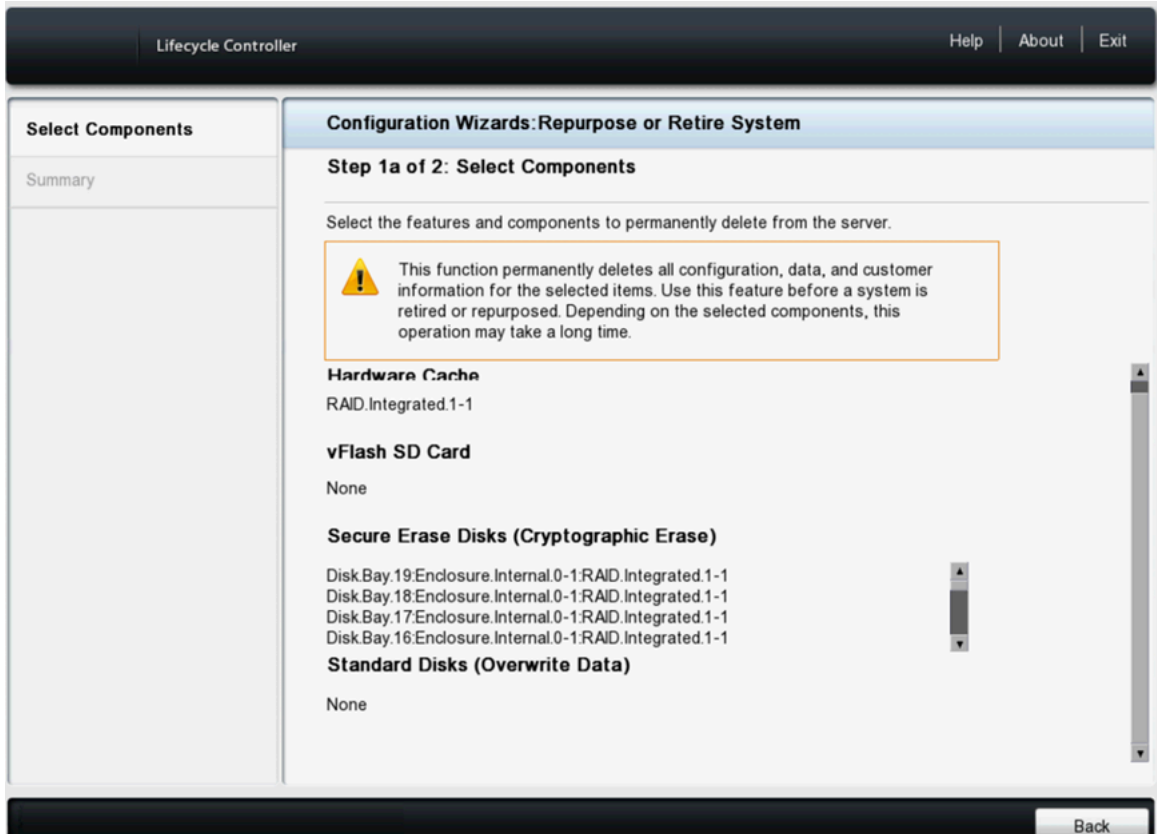
Figure D-2: Dell Lifecycle Controller



The Retire or Repurpose System function enables 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).

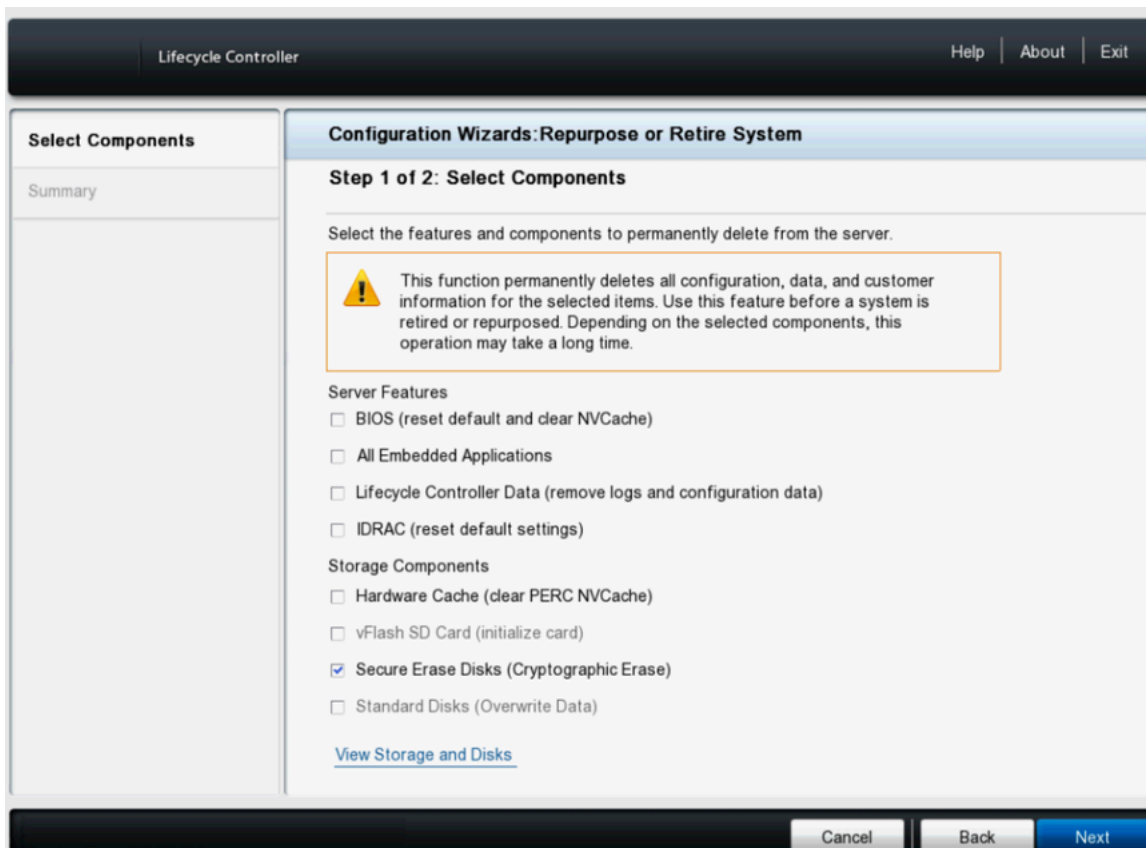
4. Click on **View Storage and Disks** to display all the drives attached to the server which are supported for erase. Only drives that can be erased and detected are displayed.

Figure D-3: Dell Lifecycle Controller Hardware Configuration




- Click **Back** and select to return to **Step 1** (or **1a**). **Select Secure Erase Disk (Cryptographic Erase)**. If both type of detected by the system, then select **Secure Erase Disk (Cryptographic Erase)** and **Standard Disks (Overwrite Data)**.

Figure D-4: Dell Lifecycle Controller > Repurpose/Retire a system

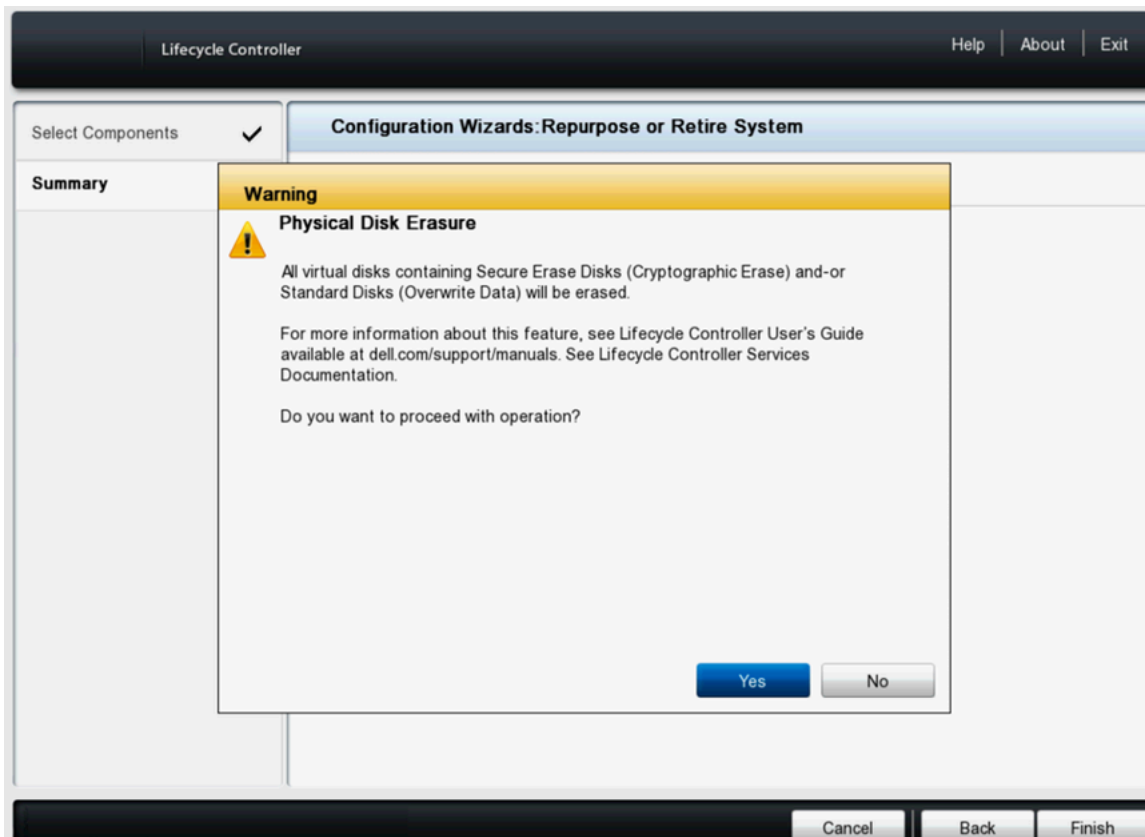


6. Click **Next** to view **Step 2 Summary** page. This page will show the drives that will be erased. Click **Finish**.
- Figure D-5: Dell Lifecycle Controller Repurpose/Retire a system**

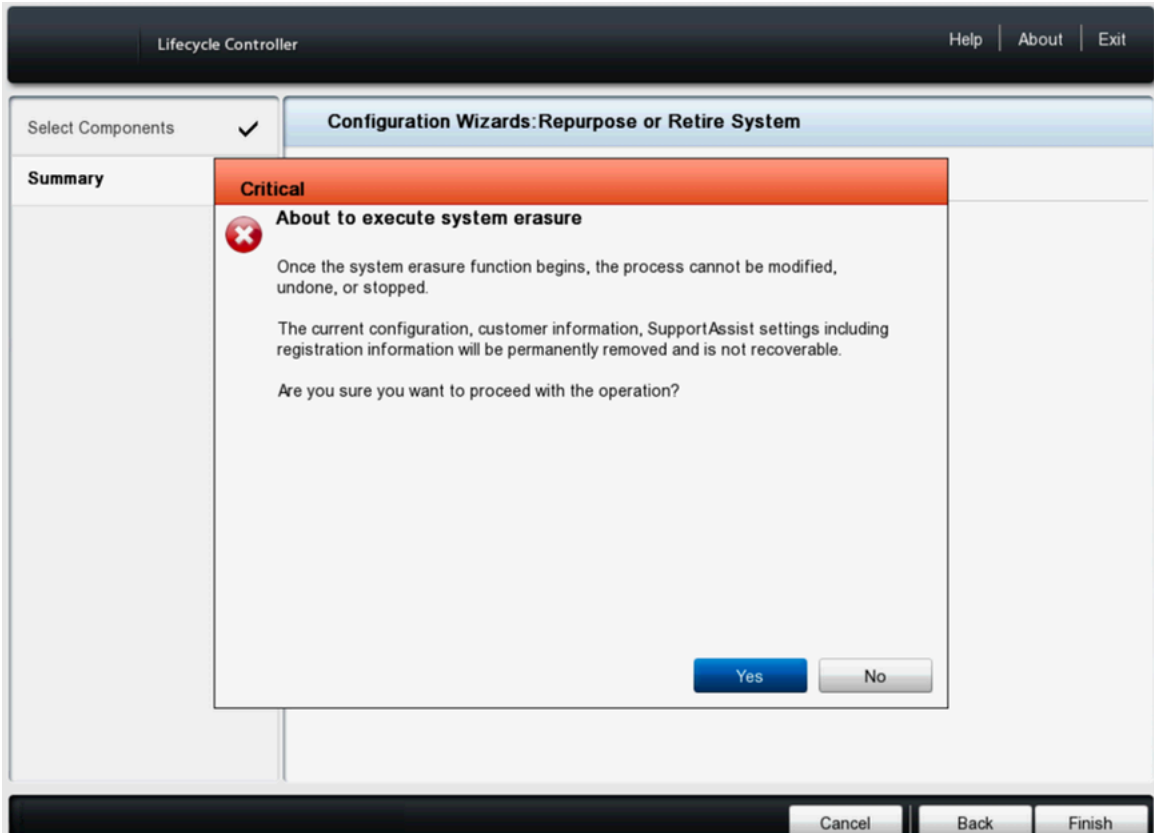


- Warning Physical Disk Erasure message appears about erasure of disks. Click **Yes**.

Figure D-6: Dell Lifecycle Controller Repurpose/Retire a system



8. Critical message appears indicating that erasure process cannot be stopped once started. Click **Yes**.
- Figure D-7: Dell Lifecycle Controller Repurpose/Retire a system**



9. DMF appliance will reboot and during POST, display will indicate entry to Automated Task Application to erase the disks.

Figure D-8: DMF Appliance



10. In Automated Task Application dialog box, the task of erasing drives is displayed along with progress bar. After the task has completed, DMF Recorder Node is turned off.



Note: Depending on the amount of data on DMF appliance, 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. On the left, a navigation pane lists tasks: 'Erase Cryptographic D...', 'Erase Cleanup', and 'Host Shutdown (JID_04...)', with 'Erase Cleanup' selected. The main area shows the details for the task '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. A progress bar is visible below the elapsed time. The task list shows 'Task 1 of 3' and 'Total Elapsed Time 00:00:42'. At the bottom, there are two informational boxes: a blue box stating 'Tasks are running normally.' and an orange box with a warning icon stating 'Do not restart, press CTRL+ALT+DEL, or turn off the server. The system will restart automatically when complete.'

Erase Cryptographic Disks (JID_046622862127)	
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 specific permit rule in 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. Node-id can be retrieved from show controller details 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 back the sync 2 permit from **0.0.0.0/0** rule.
2. **Controller-2** still stays as a standby so the user cannot add default rule to access-list sync until it transitions to active.
3. [**Controller-2**] Run this cmd on **Controller-2** “**system reset-connection switch all**”. This will change **Controller-2** to active.
4. [**Controller-2**] Then on **Controller-2** add back default rule to access-list sync 2 permit from **0.0.0.0/0** rule.
5. [**Controller-2**] On **controller-2** go to debug bash then run the following cmd.

```
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.

References

F.1 Related Documents

The following documentation is available for **DANZ Monitoring Fabric 8.4.0**:

- *DANZ Monitoring Fabric 8.4.0 Release Notes*
- *DANZ Monitoring Fabric 8.4 User Guide*
- *DANZ Monitoring Fabric 8.4 Deployment Guide*
- *DANZ Monitoring Fabric 8.4 Hardware Compatibility List*
- *DANZ Monitoring Fabric 8.4 Hardware Guide*
- *DANZ Monitoring Fabric 8.4 Verified Scale Guide*
- *DANZ Monitoring Fabric 8.4 SNMP MIB Reference Guide*