Dell Networking MXL 10/40GbE Switch IO Module

Getting Started Guide

March 2014

Notes, Cautions, and Warnings

NOTE: A NOTE indicates important information that helps you make better use of your device.

CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

———————————————

Regulatory Model DF10MXL

March 2014   P/N WMDP6   Rev. A02

Contents

Introduction

This document provides basic information about the Dell Networking MXL 10/40GbE Switch IO Module, including how to install the switch in the Dell PowerEdge M1000e Enclosure and perform the initial configuration.

For more detailed information about any of the basic installation steps, refer to the *Dell PowerEdge M1000e Enclosure Owner's Manual* on the Dell Support website at http://support.dell.com/manuals.

Product Description

The MXL 10/40GbE Switch is a layer 2/3 blade switch with two fixed 40GbE ports on the base module and support for two optional plug-in modules. The switch operates in a PowerEdge M1000e Enclosure, which can support up to 32 servers and six MXL 10/40GbE Switches.

The MXL 10/40GbE Switch runs the Dell Networking operating system (OS), providing switching, bridging, and routing functionality for transmitting data, storage, and server traffic as follows:

- GbE/10GbE connections through the midplane to up to 32 server ports in the chassis
- GbE, 10GbE, or 40GbE connections through uplink ports to top-of-rack (ToR) switches
- 40GbE stacking connections to other MXL 10/40GbE Switches in a switch stack

In a data center network, the MXL 10/40GbE Switch provides converged network support and interoperates with Dell and third-party ToR devices. The switch supports data center bridging (DCB) features, and optimizes connectivity between servers and storage devices over Fiber Channel over Ethernet (FCoE) and internet small computer system interface (iSCSI) links.

By providing increased 40GbE bandwidth for device interconnection in a shared network storage environment, with the possibility of splitting 40GbE quad small form-factor pluggable plus (QSFP+) uplinks into 10GbE small form-factor pluggable plus (SFP+) connections, the MXL 10/40GbE Switch is perfectly positioned to help transition a data center with multiple speed requirements.

Hardware Overview

This section contains information about device characteristics and modular hardware configurations for the MXL 10/40GbE Switch.

Internal Ports

The MXL 10/40GbE Switch provides thirty-two 1/10-Gigabit Ethernet internal ports. The internal ports are connected to server blades through the M1000e chassis midplane. Each port can operate in either 1GbE or 10GbE mode. Internal ports are numbered 1 to 32.

The MXL 10/40GbE Switch also provides an internal Ethernet interface—the out-of-band (OOB) interface—which is dedicated to switch management. The OOB interface is connected to the chassis management controller (CMC) through the chassis midplane. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

Front Panel

Figure 1-1 shows the MXL 10/40GbE Switch front panel:

Figure 1-1.   MXL 10/40GbE Switch: Front Panel I

Expansion Slot 1   Expansion Slot 0   USB Console Port

40GbE QSFP+ Ports

Base Module

The MXL 10/40GbE Switch provides two native 40-Gigabit Ethernet fixed ports on the base module for uplink connections. You can use these ports with 4x10G breakout cables to operate as 10GbE uplink ports.

In addition, you can configure the native 40GbE ports as stacking ports. You can connect up to six MXL 10/40GbE Switches (in the same or different chassis) in a single stack. For more information, refer to Assembling a Switch Stack.

FlexIO™ Plug-in Modules

The MXL 10/40GbE Switch (Figure 1-2) provides 10/40GbE external connectivity by using the following FlexIO plug-in modules in the two expansion slots:

- 4-Port 10-Gigabit Ethernet module using SFP+ optics (SR or LR) or direct-attach cables (1m, 3m, or 5m DAC)
- 4-Port 10GBASE-T module using RJ-45 connector (copper) cables

     NOTE: The 10Mb speed is not supported on the 4-Port 10GBASE-T module. Only 100Mb, 1GbE, and 10GbE speeds are supported.

- 2-Port 40-Gigabit Ethernet QSFP+ module which you can use for 40GbE connections or 10GbE SFP+ connections using 4x10G breakout cables.

    By default, the 40GbE ports on a 2-Port 40GbE QSFP+ module come up in 4x10GbE (quad) mode as eight 10GbE ports. To change a port from 4x10GbE to 40GbE mode of operation, enter the **no stack-unit port portmode quad** command.

    ```
    Dell(conf)# no stack-unit unit-number port port-number portmode quad
    ```

    stack-unit *unit-number*: Enter the number of the stack unit to be reset. The range is from 0 to 5. To display the stack-unit number, enter the **show system brief** command.

    port *port-number*: Enter the port number of the 40GbE QSFP+ port to be split. Valid values for slot 0: 41 or 45; for slot 1: 49 or 53 (refer to Port Numbering).

    portmode quad: Identifies the port as a split 10GbE SFP+ port.

Save the configuration and reload the switch.

```
Dell# write memory
Dell# reload
```

- Up to two FC Flex IO modules can be inserted on an I/O Aggregator, thereby enabling up to 32G of FC throughput per module or 64G throughput per switch. All the ports support 2/4/8G FC and SFP+ optics. The FC Flex IO modules provide flexibility for pairing with the other FlexIO modules. The two FC Flex IO modules that offer up to 8 8Gb Fibre Channel ports for uplink traffic in addition to the fixed two 40GbE ports on the MXL 10/40GbE Switch and I/O Aggregator.

Figure 1-2.   MXL 10/40GbE Switch: Plug-in Modules

2-Port 40GbE QSFP+ Module

4-Port 10GbE SFP+ Module

4-Port 10GBASE-T Module (100Mb/1GbE/10GbE)
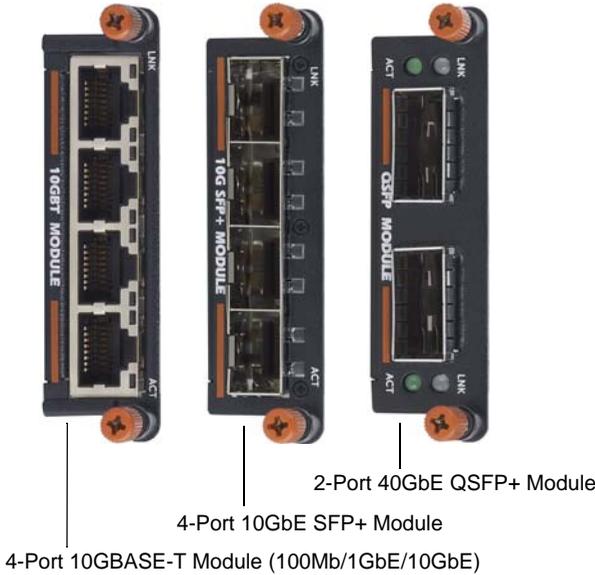
Figure 1-3.   FC FlexIO Modules

> 📝 **NOTE:** You can only hot-swap plug-in modules of the same type without requiring a reboot. For example, you can replace a 2-Port 40GbE QSFP+ module only with another 2-Port 40GbE QSFP+ module.

On the MXL 10/40GbE Switch, you can configure uplink ports of the same speed on different modules in the same link aggregation group (LAG). You can also use individual ports for uplink connections.

> 📝 **NOTE:** A maximum of 16 ports are supported in a LAG.

For switch stacking, you must use a 40GbE port on the base module or on a 2-Port 40GbE QSFP+ plug-in module. Stacking is not supported on 10GbE ports.

All FlexIO plug-in modules, transceivers, and attach cables are sold separately.

Port Numbering

When installed in a PowerEdge M1000e Enclosure, the MXL 10/40GbE Switch ports are numbered 33 to 56 from the bottom to the top of the switch:

- 40GbE base-module ports:
    - In 40GbE mode of operation, the ports are numbered 33 and 37.
    - In 4x10GbE mode of operation, the ports are numbered 33 to 36 and 37 to 40.

    For information about how to change a 40GbE port to 4x10GbE mode, refer to Splitting 40GbE QSFP+ Ports into 10GbE SFP+ Ports.

- 2-Port 40-GbE QSFP+ module operating in the default 4x10GbE mode:
    - In expansion slot 0, the ports are numbered 41 to 44 and 45 to 48.
    - In expansion slot 1, the ports are numbered 49 to 52 and 53 to 56.

    For information about how to change a port from 4x10GbE to 40GbE mode of operation, refer to FlexIOTM Plug-in Modules.

- 2-Port 40-GbE QSFP+ module operating in 40GbE mode:
    - In expansion slot 0, the ports are numbered 41 and 45.
    - In expansion slot 1, the ports are numbered 49 and 53.

- 4-Port 10-GbE SFP+ or 10GBASE-T module:
    - In expansion slot 0, the ports are numbered 41 to 44.
    - In expansion slot 1, the ports are numbered 49 to 52.

- The FC Flex IO modules ports operate:
    - In expansion slot 0, the ports are numbered 41 to 44.•
    - In expansion slot 1, the ports are numbered 49 to 52.

USB Ports

Figure 1-4.   USB Ports on Front Panel
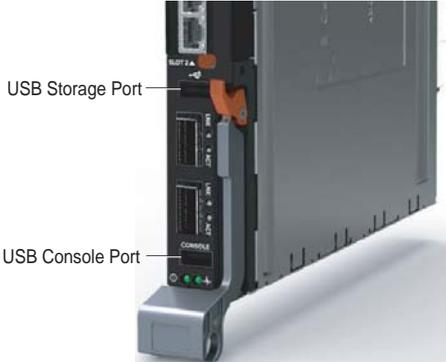
To configure the switch through an RS-232 serial interface, use the lower USB console port (Figure 1-4). This port provides a direct connection to the switch and allows you to access the command line interface (CLI) from a console terminal connected to the port through the provided serial cable (with USB type-A to female DB-9 connectors).

The console port supports asynchronous data of eight data bits, one stop bit, no parity bit, and no flow control. The default baud rate is 9600 bps.

The upper USB port functions as an external flash drive that you can use to store configuration files and scripts.

System and Port LEDs

The front panel of the MXL 10/40GbE Switch contains light emitting diodes (LEDs) that provide information about the status of the switch (Figure 1-5).

Figure 1-5.   System LEDs on Front Panel

Table 1-1 describes system LED conditions.

Table 1-1.   System LEDs

System LED	Color	Meaning
Power	Green	Power is being supplied to the switch.
	Off	The switch does not have power.
Status	Blue	The switch is operating normally as a standalone switch or as a stack master.
	Off	The switch is not the stack master.
	Amber	A fault has occurred or the switch is booting.

Each plug-in module also contains LEDs that provide information about the link status and traffic activity on a port (Figure 1-6).

Figure 1-6.   Port LEDs on Modules
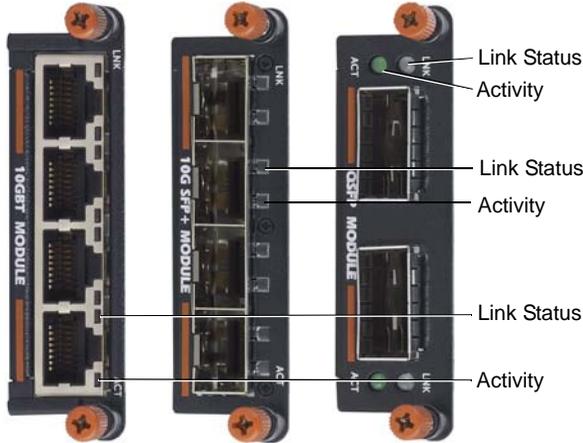
Table 1-2 describes the LED status of the 10GbE BASE-T, 10GbE SFP+, and 40GbE QSFP+ ports.

Table 1-2.   Port LED Status

Port LED	Color	Meaning
Link Status	Off	The port is down.
	Green	The port is up and can transmit traffic at maximum speed: A QSFP+ port can transmit at 40G. An SFP+ port can transmit at 10G. A BASE-T port can transmit at 10G.
	Yellow	The port is up and is transmitting traffic at lower than maximum speed: A 40GbE QSFP+ port is transmitting at 10G. A 10GbE SFP+ or 10GBASE-T port is transmitting at 1G or 100Mb.
Activity	Off	No traffic is being transmitted or received on the port.
	Blinking Green	Traffic is being transmitted or received on the port.

Table 1-3 describes the LED status of a 40GbE QSFP+ port that is split into four 10GbE SFP+ ports using a 4x10G breakout cable.

Table 1-3. LED Status of a 40GbE QSFP+ Port with Breakout Cable

Port LED	Color	Meaning
Link Status	Off	All four 10GbE ports on a breakout cable are down.
	Yellow	At least one of the four 10GbE ports on a breakout cable is up.
Activity	Off	No traffic is being transmitted on any 10GbE port on the breakout cable.
	Blinking Green	Traffic is being transmitted or received on at least one of the 10GbE ports on the breakout cable.

Installation

Site Preparation

Before installing the switch or switches, make sure that the chosen installation location meets the following site requirements:

- **Clearance** — There is adequate front and rear clearance for operator access. Allow clearance for cabling, power connections, and ventilation.

- **Cabling** — The cabling is routed to avoid sources of electrical noise such as radio transmitters, broadcast amplifiers, power lines, and fluorescent lighting fixtures.

- **Ambient Temperature** — The ambient switch operating temperature range is 10° to 35ºC (50° to 95ºF).

  NOTE: Decrease the maximum temperature by 1°C (1.8°F) per 300 m (985 ft.) above 900 m (2955 ft.).

- **Relative Humidity** — The operating relative humidity is 8% to 85% (non-condensing) with a maximum humidity gradation of 10% per hour.

Unpacking the Switch

Package Contents

When unpacking each switch, make sure that the following items are included:

- One Dell Networking MXL 10/40GbE Switch IO Module
- One USB type A-to-DB-9 female cable
- *Getting Started Guide*
- *Safety and Regulatory Information*
- *Warranty and Support Information*
- *Software License Agreement*

Unpacking Steps

NOTE: Before unpacking the switch, inspect the container and immediately report any evidence of damage.

1 Place the container on a clean, flat surface and cut all straps securing the container.

2 Open the container or remove the container top.

3 Carefully remove the switch from the container and place it on a secure and clean surface.

4 Remove all packing material.

5 Inspect the product and accessories for damage.

Installing and Configuring the Switch

After you unpack the MXL 10/40GbE Switch, refer to the flow chart in Figure 1-7 for an overview of the steps you must follow to install the blade and perform the initial configuration.

To see if a switch is running the latest Dell Networking OS version, use the **show version** command. To download an Dell Networking OS version, go to **http://support.dell.com**.

Figure 1-7.   MXL 10/40GbE Switch: Installation and Configuration Flow Chart

Insert Switch Blade and Power On

↓

Connect Console

↓

Press any key to enter X-Loader or U-Boot menu?

— Yes → X-Loader CLI

— No | Do not press a key ↓

U-Boot CLI

— Yes → Run Memory Tests/ Configure Switch Settings

X-Loader CLI → Run Memory Tests/ Configure Switch Settings → Reboot

— No | Do not press a key ↓

Load Image from Flash to RAM

↓

Enter Wizard?

— Yes → Easy Setup Wizard Configuration

— No ↓

Jumpstart waits for DHCP Offer with FTOS Image and Configuration

To exit Jumpstart:
From the console, enter
FTOS> enable
FTOS# stop jumpstart
FTOS# reload-type normal

↓

Manual Initial Configuration → Assemble a Switch Stack → Switch Configuration

Installing the Switch Blade in a PowerEdge M1000e

After you unpack the switch blade, slide it into one of the open I/O module slots in the back of a PowerEdge M1000e.

The M1000e is a 10U rack-mountable blade chassis that holds:

- Server blades: Eight full-height, 16 half-height, or 32 quarter-height blades
- Switch blades: Six I/O modules and two integrated chassis management controllers

Server blades are installed in the front of the chassis (Figure 1-8); switch blades are installed in the back of the chassis (Figure 1-9).

Figure 1-8.   PowerEdge M1000e: Front View with Server Blades

16 half-height server blades
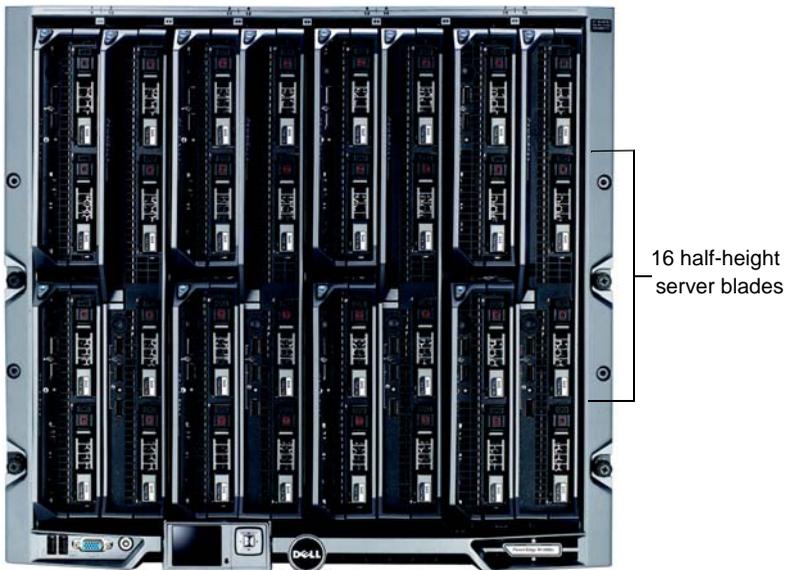
Figure 1-9.   PowerEdge M1000e: Back View with Six MXL 10/40GbE Switch Blades

After you slide the MXL 10/40GbE Switch in so that the connectors on the back of the blade touch the chassis midplane, the switch receives power from the chassis and automatically powers on. The chassis management controller (CMC) in the chassis validates that the switch blade is a supported I/O module before powering it on.

When the switch powers on, the Boot loader loads the image from the local flash. The image initializes the hardware and brings the switch up in operational mode.

Connecting a Console Terminal

After the MXL 10/40GbE Switch powers on, complete all external cabling connections and connect a terminal to the blade to configure the switch.

> **NOTE:** *If you are installing a stack of MXL 10/40GbE Switches*, connect the terminal to the console port on the Master Switch. If you connect the terminal to a member (non-Master) switch, you will not be able to use the CLI. For more information, refer to Assembling a Switch Stack.
> Read the Release Notes for this product before proceeding. You can download the release notes from the Dell Support website at support.dell.com/manuals.

To monitor and configure the switch via the serial console, use the USB console port on the front panel of the switch (Figure 1-4) to connect it to a VT100 terminal or to a computer running VT100 terminal emulation software. The console port is implemented as a data terminal equipment (DTE) connector.

The following equipment is required to use the console port:

- VT100-compatible terminal or a desktop or a portable computer with a serial port running VT100 terminal emulation software, such as Microsoft HyperTerminal.
- A serial cable (provided) with a USB type-A connector for the console port and DB-9 connector for the terminal.

To connect a terminal to the switch console port, perform the following tasks:

1 Connect the DB-9 connector on the serial cable to the terminal or computer running VT100 terminal emulation software.

2 Configure the terminal emulation software as follows:

   a Select the appropriate serial port (for example, COM 1) to connect to the console.

   b Set the data rate to 9600 baud.

   c Set the data format to 8 data bits, 1 stop bit, and **no** parity.

   d Set the flow control to **none**.

   e Set the terminal emulation mode to **VT100**.

   f Select **Terminal keys** for Function, Arrow, and Ctrl keys. Ensure that the setting is for Terminal keys (not Microsoft Windows keys).

Connect the USB connector on the cable directly to the switch console port. The console port on the MXL 10/40GbE Switch is located below the fixed 40GbE ports (Figure 1-4).

Invoking the X-Loader and U-Boot CLIs

During the boot process, you can perform various configuration tasks by accessing the X-Loader and U-Boot CLIs, such as running memory tests (X-Loader) and activating the backup image or recovering a password (U-Boot).

You are first prompted to enter the X-Loader CLI by pressing any key when the following message displays: `Hit any key to stop autoboot.` If you do not press a key, the boot process continues and you are prompted to enter the U-Boot CLI by pressing any key.

After performing any of the X-Loader or U-Boot tasks, the switch automatically reboots when you exit a CLI. To continue with the boot process without entering either CLI, do not press a key.

Performing the Initial Configuration

Prerequisites: Before you perform the initial switch configuration, make sure that:

- The MXL 10/40GbE Switch was never configured before and is in the same state as when you received it.
- The MXL 10/40GbE Switch booted successfully when it powered on.

Perform the initial switch configuration through the console port. After the initial configuration, you can manage the switch from the already connected console port or through a remote connection.

Before you start, to perform the initial switch configuration, you must obtain the following information from your network administrator:

- The IP address to be assigned to the out-of-band (OOB) interface for device management.
- The IP subnet mask for the OOB interface.
- The IP address of the OOB interface default gateway.

These settings are necessary to allow remote management of the switch through a Telnet (Telnet client) or HTTP (Web browser) connection.

> **NOTE:** The switch is configured with a default user name (root) and password (calvin).

Initial Configuration Procedure

Perform the initial configuration from the connected console by using the CLI.

Configuring the SNMP Server

To set up the SNMP account and provide read/write privileges, use the `snmp-server community` command.

1 Create an SNMP community string with read/write privileges by entering the **snmp-server community** *<community_name>* **rw** *<ACL_NAME>* command in CONFIGURATION mode.

2 Enter IP ACCESS LIST mode by naming a standard IP access list by entering the **ip access-list standard** *<ACL_NAME>* in CONFIGURATION mode.

3 Configure the management IP address to manage the SNMP query of the chassis by entering the **permit** *<Management_System_IP>* command in CONFIGURATION-IP ACCESS-LIST-STANDARD mode.
With this configuration, the SNMP query is accepted only from the configured IP.

> **NOTE:** If you want any SNMP management entity to access the chassis, ignore steps 2 and 3.

Configuring the Username and Password

To access the system remotely, configure a system username, password, and privilege level.

1 Configure a username, password, and privilege level by entering the **username** *username* **password** *<LINE> privilege level* command in CONFIGURATION mode.

Configuring the Management Port IP Address

Assign IP addresses to the management ports in order to access the system remotely.

1 Enter INTERFACE mode for the Management port by entering the **interface ManagementEthernet** *slot/port* command in CONFIGURATION mode.

2 Assign an IP address to the interface by entering the **ip address** *ip-address/mask* in INTERFACE mode.

3 Enable the interface by entering the **no shutdown** command in INTERFACE mode.

Configuring the Default VLAN IP Address

Assign IP addresses to the default VLAN in order to access the system remotely.

1 Enter INTERFACE mode for the VLAN by entering the **interface vlan 1** command in CONFIGURATION mode.

2 Assign an IP address to the interface by entering the **ip address** *ip-address/mask* in INTERFACE mode.

3 Enable the interface by entering the **no shutdown** command in INTERFACE mode.

Configuring the Management Route

Define a path from the MXL switch to the network from which you are accessing the MXL remotely. Management routes are separate from IP routes and are used to manage the MXL switch through the management port.

1 Configure a management route to the network from which you are accessing the system by entering the **management route** *ip-address/mask gateway* command in CONFIGURATION mode.

For more information about how to perform the initial configuration using the CLI, refer to the *Dell Networking OS Configuration Guide for the Dell Networking MXL 10/40GbE Switch IO Module*.

Assembling a Switch Stack

After you complete the initial switch configuration, the MXL 10/40GbE Switch is powered up and operational. Stacking is supported on the 40GbE ports on the base module or a 2-Port 40GbE QSFP+ module to connect up to six MXL 10/40GbE Switches in a single stack.

Figure 1-10 shows an example using six MXL 10/40GbE Switches in a chassis. The MXL 10/40GbE Switches are connected to operate as a single stack in a ring topology using only the 40GbE ports on the base modules. You can use the 40GbE ports on the base module and plug-in modules to create a stack in either a ring or daisy-chain topology.

> ✐ **NOTE:** Power up all MXL 10/40GbE Switches in the stack should be powered up with the initial configuration before you attach the cables.

Figure 1-10. Six Stacked MXL 10/40GbE Switches
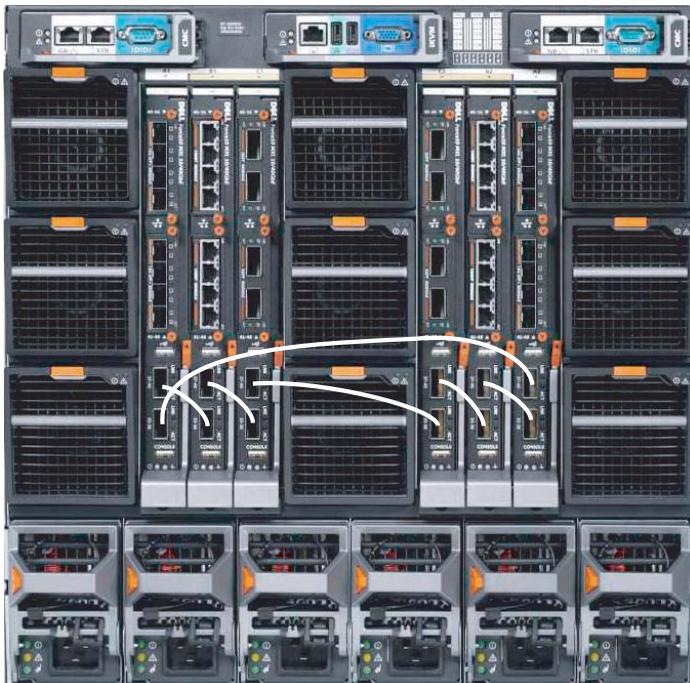
Use only QSFP transceivers and QSFP cables (separately purchased) to connect stacking ports as follows:

1 Insert a QSFP cable in the bottom stacking port on the rightmost switch.

2 Connect the upper stacking port on the next switch to the left.

3 Continue connecting each switch to the next in this way until you reach the leftmost switch in the stack.

4 On the leftmost switch, connect the bottom stacking port to the upper stacking port on the rightmost switch to create a loop.

NOTE: The resulting ring topology allows the entire stack to function as a single switch with resilient fail-over capabilities.
If you do not connect the leftmost switch to the rightmost switch (Step 4), the stack operates in a daisy-chain topology with less resiliency. Any failure in a non-edge stack unit causes a split stack.

Configuring and Bringing Up a Stack

NOTE: Although stacking is supported on 40GbE ports on the base module or a 2-Port 40GbE QSFP+ module, this section shows how to configure stacking only on the base-module ports (Figure 1-10).
To convert the 40GbE ports on the 2-Port QSFP+ module from their default 4x10GbE mode to 40GbE mode to configure a stack, refer to FlexIOTM Plug-in Modules. After the converting the ports to 40GbE, you do not have to reload the switch because a switch reload is required as part of the stack configuration procedure.

After you attach the QSFP+ cables in a stack of MXL 10/40GbE Switches, to configure and bring up the stack, follow these steps:

1 Connect the terminal to the console port on an MXL 10/40GbE Switch. Enter the following commands to access the CLI and configure the two 40GbE ports on the base module for stacking mode:

```
Login: username
Password: *****
Dell> enable
Dell# configure
Dell(conf)# stack-unit 0 stack-group 0
Dell(conf)# stack-unit 0 stack-group 1
```

Where **stack-unit 0** defines the default stack-unit number in the initial configuration of a switch; **stack-group 0** defines the stack group for the

lower 40GbE base-module port and **stack-group 1** defines the stack group for the upper 40GbE base-module port. To display the ports in each stack group, enter the **show system stack-unit** *unit-number* **stack-group** command.

2 Save the stacking configuration on the 40GbE ports:

```
Dell# write memory
```

3 Repeat Steps 1 and 2 on each MXL 10/40GbE Switch in the stack by entering the **stack-unit 0 stack-group 0** and **stack-unit 0 stack-group 1** commands and saving the configuration.

4 Reboot each switch by entering the **reload** command in EXEC Privilege mode:

```
Dell# reload
```

If the stacked switches all reboot at approximately the same time, the switch with the highest MAC address is automatically elected as the Master Switch. The switch with the next highest MAC address is elected as Standby Master.

To configure the stack so that the roles are assigned according to pre-determined priorities, enter the **stack-unit priority** command as described in the following Note.

To reload the stack, connect the terminal to the Master Switch and enter the **reload** command in EXEC Privilege mode. If you connect the terminal to a member (non-Master) switch, you will not be able to access the CLI.

To determine which switch is the Stack Master, enter the **show system** command at the terminal.

To remove a port from stacking mode, use the **no** form of the **stack-unit stack-group** command; for example, **no stack-unit 0 stack-group 0**. After entering the command, save the configuration and reload the switch for the change to take effect. When the reload completes, the port comes up in 40GbE mode if it is on the base module and in 4x10GbE (quad) mode if the port is on a plug-in module.

> **NOTE:** You can manually configure the switch that will be Master by entering the **stack-unit** *unit-number* **priority** *number* command in CONFIGURATION mode on each stacked switch, where:
> **stack-unit** *unit-number* identifies the switch in the stack.
> **priority** *number* specifies the management priority. Range: 1-14. Default: 0.

The switch with the highest priority number is elected Master. The switch with the next highest priority number is elected Standby Master and takes over stack management if the Master Switch fails. For example:

```
Dell> enable
Dell# configure
Dell(conf)# stack-unit 0 priority 14
```

The **no** form of the **stack-unit** *unit-number* **priority** *number* command reverts the management priority of a stack unit to the default value of 0.

Managing a Stack

Master and Member Switches

You can manage a stack of MXL 10/40GbE Switches as a single entity when connected together. Manage the stack from the CLI through the serial console connection or a remote Telnet session over the OOB management IP address.

When you create a stack is created, one switch automatically becomes the Master Switch and another switch is elected Standby Master. The Master Switch maintains stack operation with minimal impact in the event of:

- Switch failure
- Inter-switch stacking link failure
- Switch insertion
- Switch removal

If the Master Switch goes off line, the Standby Master replaces it as the new master and the switch with the next highest MAC address or priority becomes Standby Master.

Stack Startup

Topology Discovery

When a stack is formed, a topology discovery process builds up a database that contains information about all of the switches in the stack, including the Dell Networking OS version, hardware version, management priority, and switch MAC address. To view this information, use the **show system** command.

Auto Stack Number Assignment

During the stack formation process, a unique stack-unit number is assigned if the same number is assigned to more than one switch. After assignment is complete, each switch saves its stack-unit number. To view stack-unit numbers, enter the **show system** command.

Dell Networking OS Version Checking

Following the stack-unit number assignment, the Master Switch performs a consistency check to make sure that all switches in the stack are running the same Dell Networking OS version.

If the Master Switch determines that all switches are not running the same Dell Networking OS version, the ports on switches with the incorrect version are disabled.

To download the required Dell Networking OS image from the Master Switch and reload a member switch so that it joins the stack, enter the following command in EXEC Privilege mode:
```
Dell# upgrade system stack-unit unit-number partition
```

Where **stack-unit** *unit-number* identifies the switch whose Dell Networking OS version needs to be upgraded; *partition* identifies the partition on the Master Switch from which the Dell Networking OS image boots up. For example:
```
Dell# upgrade system stack-unit 3 a:
```

To display the boot partition used on the Master Switch, enter the **show version** command.

To ensure that a stack unit boots from partition a:, enter the commands:
```
Dell# configure
Dell(conf)# boot system stack-unit unit-number primary system a:
Dell(conf)# end
Dell# write memory
Dell# power-cycle stack-unit unit-number
```

> **NOTE:** When you stack an MXL 10/40GbE Switch, booting is supported only from flash memory; it is not supported over the network via an IP address.

> **NOTE:** To upgrade all switches in a stack with the same Dell Networking OS version in a specified partition, enter the following command in EXEC Privilege mode:
> ```
> Dell# upgrade system {ftp: | tftp:} partition
> ```

This command uses an interactive CLI that requests the server IP address and image filename, and prompts you to upgrade all member stack units. After upgrading all switches in the stack, save the configuration (using the **write memory** command) and reload the stack to activate the new Dell Networking OS version (using the **reload** command).

System Initialization

The Master Switch initializes the stack using the last saved system configuration file.

If you change the stack configuration, be sure to save the configuration file. The Master Switch automatically distributes the configuration file to the member switches. If the Master Switch later becomes unavailable, the Standby Master becomes the new Master Switch and configures the stack with the latest configuration synchronized from the Master Switch.

In case of a split stack in which the Master Switch in one of the new stacks was previously a member switch, the last saved configuration on the original stack is used to configure the new stack.

Splitting 40GbE QSFP+ Ports into 10GbE SFP+ Ports

The MXL 10/40GbE Switch supports splitting a 40GbE port on the base module or a 2-Port 40GbE QSFP+ module into four 10GbE SFP+ ports using a 4x10G breakout cable.

> **NOTE:** By default, the 40GbE ports on a 2-Port 40GbE QSFP+ module come up in 4x10GbE (quad) mode as eight 10GbE ports. On the base module, you must convert the 40GbE ports to 4x10GbE mode as described below.

Before you attach a 4x10G breakout cable to a 40GbE QSFP+ port, ensure that the following requirements are met:

• A 4x10GbE QSFP+ port cannot be used for stacking.

• In order for a 40GbE port to operate in 4x10GbE mode, you must enter the **stack-unit port portmode quad** command.

    Dell(conf)# stack-unit *unit-number* port *port-number*
    portmode quad

stack-unit *unit-number*: Enter the number of the stack unit to be reset. Range: 0-5. To display the stack-unit number, enter the **show system brief** command.

port *port-number*: Enter the port number of the 40GbE QSFP+ port to be split. Valid values on base module: 33 or 37; slot 0: 41 or 45; slot 1: 49 or 53.

portmode quad: Identifies the uplink port as a split 10GbE SFP+ port.

Then save the configuration and reload the switch.

```
Dell# write memory
Dell# reload
```

To change a port from 4x10GbE to 40GbE mode of operation, enter the **no stack-unit port portmode quad** command, save the configuration, and reload the switch.

To display the current 40GbE or 4x10GbE mode of port operation, enter the **show running-config** command.

Switch Configuration

After you complete the initial switch configuration, the MXL 10/40GbE Switch is operational. You must manually configure the switch to enable software features, such as:

- Data center bridging capability exchange protocol (DCBX)
- Fiber Channel over Ethernet (FCoE) connectivity
- FCoE initiation protocol (FIP) snooping
- iSCSI optimization
- Internet group management protocol (IGMP) snooping
- Remote authentication dial-in service (RADIUS) support
- Terminal access controller access control system plus (TACACS+) client

For information about how to configure switch software, refer to the *Configuration Guide for the Dell Networking MXL 10/40 GbE Switch IO Module* on the Dell Support website at http://support.dell.com/manuals.

DCB Support

DCB enhancements for data center networks are supported to eliminate packet loss and provision links with required bandwidth. On the MXL 10/40GbE Switch, you must manually configure DCBX port roles to enable the switch to auto-configure its DCB settings to match the DCB configuration in the ToR switches to which it connects.

FCoE Connectivity

Many data centers use Fiber Channel (FC) in storage area networks (SANs). FCoE encapsulates Fiber Channel frames over Ethernet networks.

On a MXL 10/40GbE Switch, the internal ports of the MXL 10/40GbE Switch support FCoE connectivity and connect to the converged network adapter (CNA) in blade servers. FCoE allows Fiber Channel to use 10-Gigabit Ethernet networks while preserving the Fiber Channel protocol.

You must manually configure the FCoE settings on the MXL 10/40GbE Switch to match the FCoE settings used in the ToR switches to which it connects.

iSCSI Operation

Support for iSCSI traffic is turned on by default when the MXL 10/40GbE Switch powers up.

The MXL 10/40GbE Switch monitors known transmission control protocol (TCP) ports for iSCSI sessions.When a session is detected, an entry is created and monitored as long as the session is active.

The MXL 10/40GbE Switch also detects EqualLogic iSCSI storage devices on all interfaces and auto-configures to optimize performance. Performance optimization operations, such as Jumbo frame size support, spanning tree protocol (STP) port-state fast, and disabling of storm control on interfaces connected to an EqualLogic iSCSI storage device, are applied automatically.

CLI configuration is necessary only when the configuration includes iSCSI storage devices that cannot be automatically detected and when non-default quality of service (QoS) handling is required.

Next Steps

If you installed the MXL 10/40GbE Switch in a stack, you can configure additional settings for switch stacking.

For information about how to create different stacking scenarios, refer to the *Ethernet Stacking* white paper on the Dell Support and Dell Networking White Papers websites at:
http://support.dell.com/manuals
http://en.community.dell.com/techcenter/networking/w/wiki/networking-whitepapers.aspx

For information about how to configure stacking, refer to the *Configuration Guide for the Dell Networking MXL 10/40 GbE Switch IO Module* on the Dell Support website at http://support.dell.com/manuals.

- You can customize the MXL 10/40GbE Switch for use in your data center network by configuring software features as necessary. To perform additional switch configuration, do one of the following:

  - For remote out-of-band management, enter the OOB management interface IP address into a Telnet or SSH client and log in to the switch using the user ID and password to access the CLI.

  - For local management, use the attached console connection to log in to the CLI.

  - For remote in-band management from a network management station, enter the IP address of the default virtual local area network (VLAN) and log in to the switch to access the CLI.

For information about how to configure software settings, refer to the *Configuration Guide for the Dell Networking MXL 10/40 GbE Switch IO Module* on the Dell Support website at http://support.dell.com/manuals.

Technical Specifications

The MXL 10/40GbE Switch is an I/O module and installed with Server (model: PowerEdge M1000e) for communication.

NOTE: The product must be operated at an ambient temperature of 50°C.

![CAUTION icon] **CAUTION: A Lithium Battery Caution: There is a danger of explosion if the battery is incorrectly replaced.**

Replace only with same or equivalent type. Dispose batteries according to manufacturer's instructions.

Chassis Physical Design

Parameter	Specifications
Height	1.32 inches (33.45 mm)
Width	10.81 inches (274.75 mm)
Depth	12.17 inches (309.24 mm)

Environmental Parameters

Parameter	Specifications
Operating temperature	32° to 104°F (0° to 40°C)
Operating humidity	10 to 85% (RH), non-condensing
Storage temperature	–40° to 158°F (–40° to 70°C)
Storage humidity	5 to 95% (RH), non-condensing
Maximum thermal output	419.7 BTU/hr

Printed in the U.S.A.

0WQMDP6A02