DELLEMC

Dell Networking MXL and M IOA FC-FlexIO Direct Connect Storage Deployment Guide

FC FlexIO Fabric Services Update - Providing F-port Connectivity to Storage

Dell Networking Solutions Engineering
February 2015

A Dell EMC Deployment Guide

Revisions

Date	Description	Authors
February 2015	Initial Release	Jim Slaughter, Kevin Locklear, Curtis Bunch

Table of contents

DELLEMC

1 Introduction

Dell Networking works to provide customers with the most efficient use of current networking equipment and the lowest cost of growth while still providing today's great new technologies focused around the explosive data growth in the industry. The emergence of SAP, Microsoft SharePoint, Virtual Desktop Infrastructure (VDI), Hadoop, larger databases, and increased usage/reliance on Microsoft Exchange Server have driven the need for increased bandwidth, lower latency and converged infrastructure.
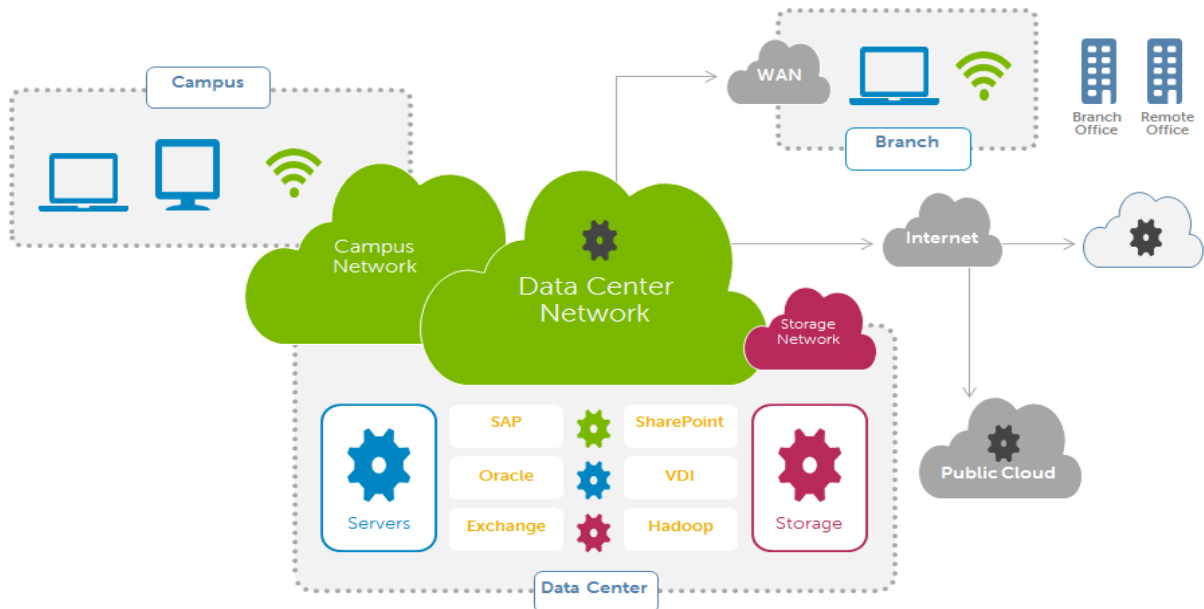
Figure 1   Networking Overview.

The focus areas of this guide are Data Center and storage networks (Figure 1). In particular, the ability the FC FlexIO module brings to the MXL and the M I/O Aggregator (M IOA) for splitting out Fibre Channel (FC) network traffic at what would seem to be the back of the blade server. The MXL and M IOA provide this magic at the back of the blade server. These switching Input/Output Modules (IOMs) allow the converged separation of the Data Center Network and the Storage Network.

DELLEMC
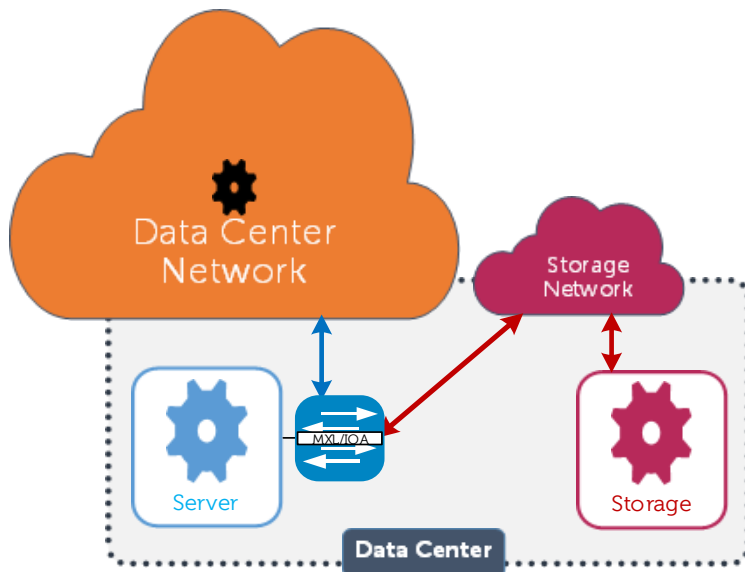
Figure 2   Converged Traffic in a Typical Environment

In a typical environment (Figure 2), converged traffic goes from the server to a Fibre Channel Forwarder (FCF) switch that de-encapsulates the two types of traffic and forwards them to their respective networks.
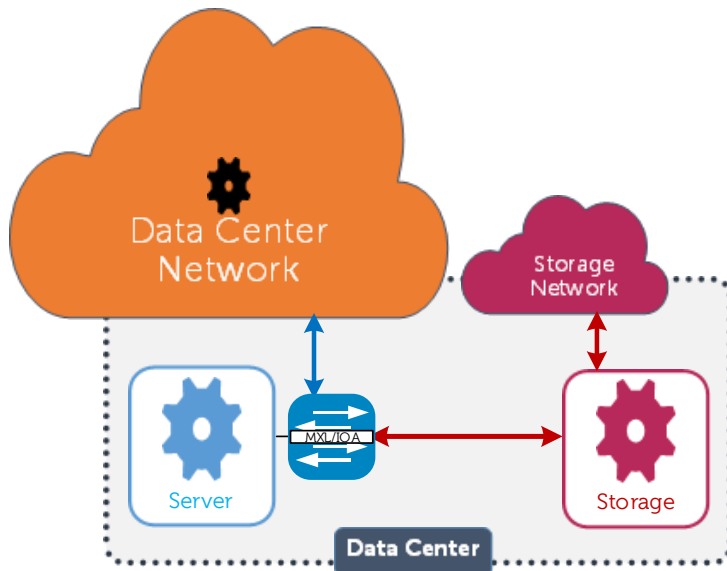
Figure 3   Converged Traffic in the Topology covered in this Guide - Direct Storage Connection

This guide discusses the new essential fabric services the Dell Networking Operating System (FTOS) 9.7 provides to the Dell Networking MXL, M IOA and their add-in FC Flex IO modules. These new fabric services enable the MXL and M IOA to have direct connectivity to FC end devices (Figure 3). In other words, this enables the FC ports on the FC Flex IO to be F-ports.

DELLEMC

1.1 Typographical Conventions

`Monospace Text`	CLI examples
<u>`Underlined Monospace Text`</u>	CLI examples that word wrap. This text should be entered as a single command.
`Italic Monospace Text`	Variables in CLI examples
`Bold monospace Text`	Commands entered at the CLI prompt

DELLEMC

2 Hardware

In this section, the hardware used to validate the topology outlined in this deployment guide is briefly discussed.

> **Note:** Refer to the Configuration Details section in the Appendix for specific firmware and driver versions.

2.1 Dell PowerEdge M1000e

The Dell PowerEdge M1000e modular blade enclosure (Figure 4) is the rock-solid foundation for Dell's blade server architecture, providing an extremely reliable and efficient platform for building any IT infrastructure. The M1000e enclosure is built from the ground up to combat data center sprawl and IT complexity, delivering one of the most energy efficient, flexible, and manageable blade server implementations on the market.

The PowerEdge M1000e chassis enclosure supports server modules, network, storage, and cluster interconnect modules (switches and pass-through modules), a high performance, and a highly available passive midplane that connects server modules to the infrastructure components, power supplies, fans, and integrated KVM and CMC. The PowerEdge M1000e uses redundant and hot-pluggable components throughout to provide maximum uptime.

Virtually unlimited in scalability, the PowerEdge M1000e chassis provides ultimate flexibility in server processor and chipset architectures. Both Intel and AMD server architectures can be supported simultaneously by the M1000e infrastructure, while cutting-edge mechanical, electrical, and software interface definitions enable multi-generational server support and expansion. For more information about the Dell PowerEdge M1000e, visit http://www.dell.com/us/business/p/poweredge-m1000e/pd.

Figure 4   Dell M1000e Module Blade Enclosure

DELLEMC

2.2 M1000e I/O Modules

M1000e blade environments can be expanded quickly and easily with Dell's complete, scale-on-demand switch designs and IO modules that support three fully redundant fabrics. IO options include Dell Networking switches with FlexIO technology that provide on-demand stacking and modular uplink scalability.

2.2.1 Dell MXL 10/40GbE Overview

The Dell Networking MXL switch (Figure 5) is a layer 2/3 blade switch with two fixed 40GbE ports on the base module and support for two optional FlexIO Expansion Modules. The MXL 10/40GbE switch runs the Dell Networking Operating System, providing switching, bridging and routing functionality for transmitting data, storage and server traffic.
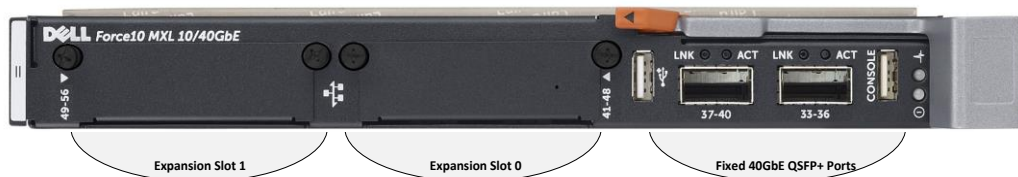
Figure 5   Dell Networking MXL

2.2.2 Dell M I/O Aggregator Overview

The M IOA (Figure 6) is a low-touch blade switch with two fixed 40 GbE ports on the base module and support for two optional FlexIO Expansion Modules. The M IOA runs the Dell Networking Operating System. In these automated modes (SMUX, VLT or Stacking) all VLANs are allowed as well as any DCBx, iSCSI or Fibre Channel over Ethernet (FCoE) settings. In addition, the external ports are all part of the same LAG, which eliminates the need for the Spanning Tree Protocol (STP) on the M IOA.
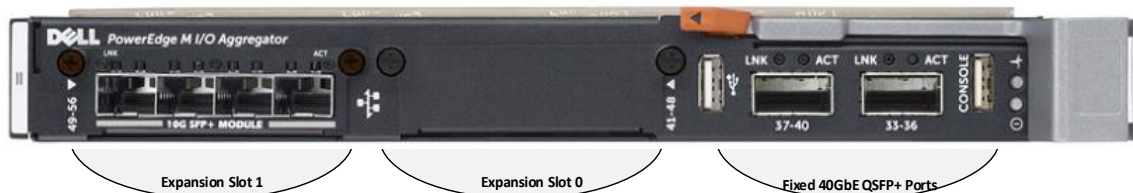
Figure 6   Dell PowerEdge M I/0 Aggregator

2.2.3 FlexIO Expansion Modules

The Dell IO Modules support a combination of four types of FlexIO Modules (Figure 7). The FlexIO modules include:

- 4-port 10Gbase-T FlexIO module (only one 10Gbase-T module can be used).
- 4-port 10Gb SFP+ FlexIO module.
- 2-port 40Gb QSFP+ FlexIO module.
- 4-port Fiber Channel 8Gb module.

DELLEMC
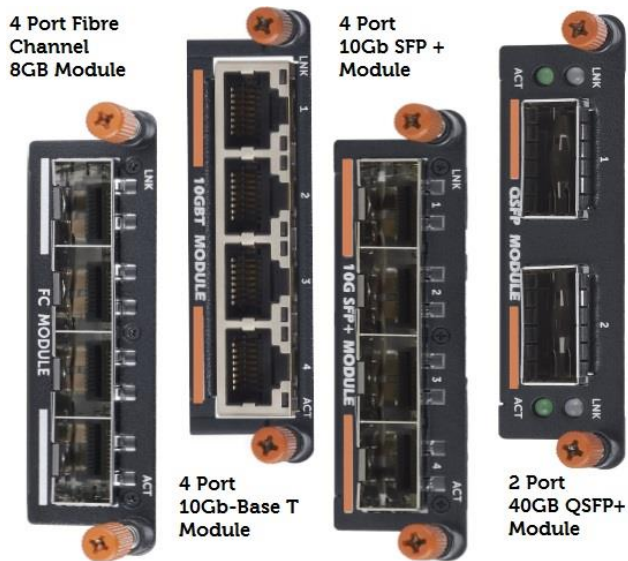
Figure 7   FlexIO Expansion Modules

In this document two 4-port Fiber Channel 8Gb modules are used in conjunction with two Dell MXLs or M IOAs.

2.3   Dell PowerEdge M620 Blade Server

The Dell PowerEdge M620 blade server (Figure 8) is a feature rich, two-socket blade server, designed for maximum performance with extreme density. Designed for taxing workloads, such as mail server, database and virtual environments, the M620 blade server is an ideal blend of density, performance, efficiency and scalability.

Figure 8   Dell PowerEdge M620 Blade Server

2.4　Dell Networking S-Series S6000 Managed Switch

The Dell Networking S6000 10/40GbE switch (Figure 9) supports Top-of-Rack (ToR), Middle-of-Row (MoR) and End-of-Row (EoR) connectivity to 10GbE servers and 40GbE fabrics for enterprise and midmarket implementations. The 1U S6000 enables streamlined connectivity and workflow operation with low latency operation and includes support for Network Virtualization with VXLAN Gateway hardware built-in. Virtual Link Trunking (VLT) eliminates spanning tree protocol (STP) blocked ports and provides fast convergence if a link or device fails.

Figure 9　　Dell Networking S6000

2.5　NetApp FAS3200 Series Data Storage System

The NetApp FAS3200 Series Data Storage System (Figure 10) is a midrange data storage platform. In this guide, the FAS3220 is used with dual control enclosures for high availability (HA) and expanded I/O for additional storage capacity. The FAS3220 can support 480 drives, with a total capacity of 1,920TB. The storage system is running Data ONTAP. For more information please visit the product page:
http://www.netapp.com/us/products/storage-systems/fas3200/

> **Note:** Any FC capable storage array can be used for this deployment. Please check with your storage array manufacturer for installation and configuration instructions.

Figure 10　NetApp FAS3200 Series Controller

DELLEMC

3 Converged Network Adapter Configuration

The solution in this guide was validated using both a Broadcom 57810 Network Daughter Card (NDC) and an Emulex OCm14102 NDC. For instructions on configuring a Broadcom or Emulex NDC, please see the relevant attachment (Table 1).

Table 1   Network Daughter Card Configuration Guides

CNA Type	Attachment Name
Broadcom 57810	CNA_Broadcom_57810-K_FCoE.pdf
Emulex OCm14102	CNA_Emulex_OCm14102_FCoE.pdf

4 Storage Configuration

The configuration of the NetApp FAS3220 is outside of the scope of this document. Almost any SAN array that can present FC storage LUNs is compatible with this deployment guide. Please refer to your specific SAN manufacture for configuration assistance.

DELLEMC

5   Switch Configurations

The CLI configuration for the MXLs and M IOAs is presented in this section. The main focus of this deployment guide is the storage configuration. The Ethernet configuration is lightly touched upon, focusing on enabling uplink interfaces to the S6000 Top of Rack (ToR) switches.

5.1   Ethernet Topology

The topology built in this guide (Figure 11) uses a single 40GbE link to connect each MXL or M IOA to Dell Networking S6000 switches, which behave as ToR switches in this environment. The S6000s are in turn connected to a collapsed core (not shown in this example). In the figure, the dashed lines from the M620 to the MXL or M IOA represent the internal, converged connections that carry Ethernet as well as FCoE traffic.
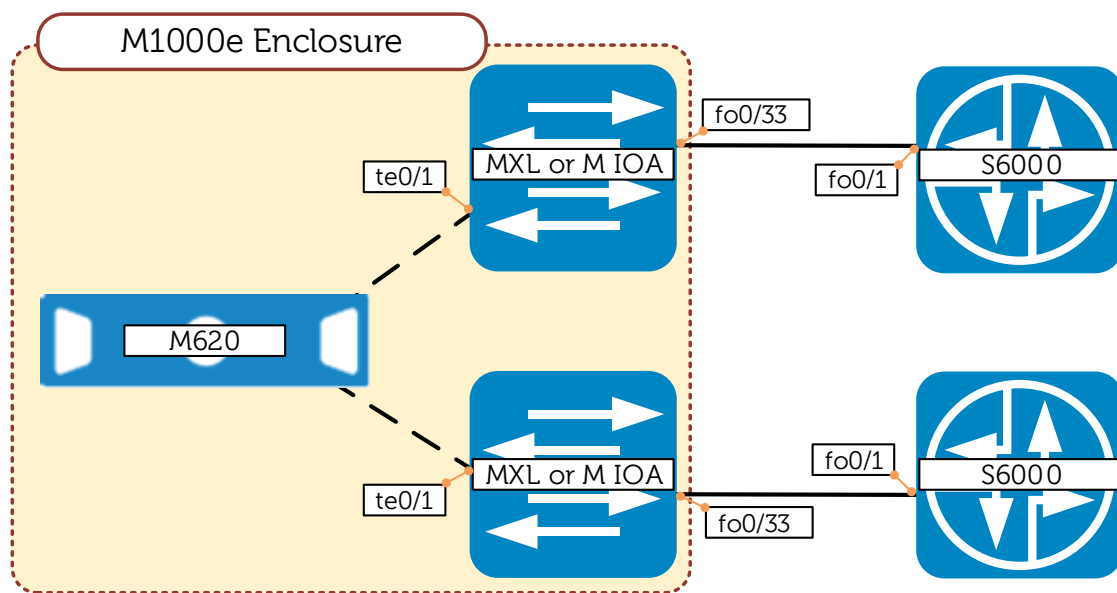
Figure 11   Ethernet Topology

5.2   Storage Topology

Figure 12 and Figure 13 show the MXL and M I/O Aggregator topologies and include the port IDs, FC Aliases and World Wide Port Names (WWPNs) used. In both figures, the dashed lines from the M620 to the MXL or M IOA represent FCoE capable links. An FCoE initialization Protocol (FIP) session is created between the NDC and the MXL or the M IOA at which point the FCoE package is decapsulated and the FC frame is sent directly to the storage array through single links from each MXL or M IOA. This topology enables redundant pathing to the storage array utilizing minimal cabling.

DELLEMC

Note: In the diagrams below a single connection is made from each MXL or M IOA to the SAN Array. This is only done to simplify the configuration for demonstration purposes. As a best practice, multiple connections from each MXL or M IOA are suggested.
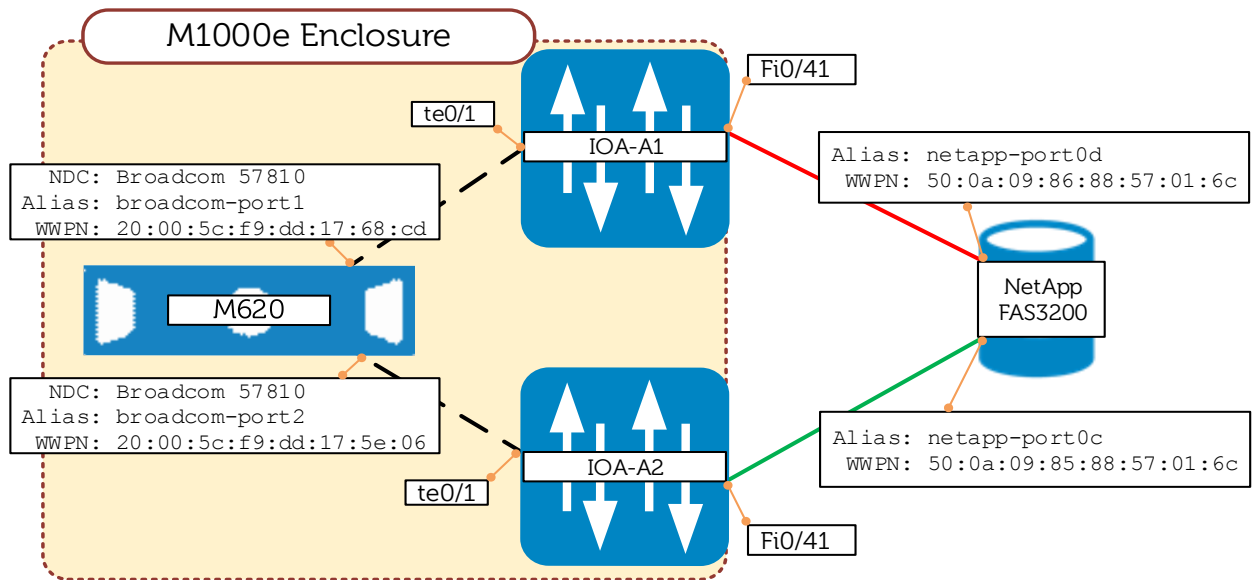
Figure 12   Storage Configuration using M IOA
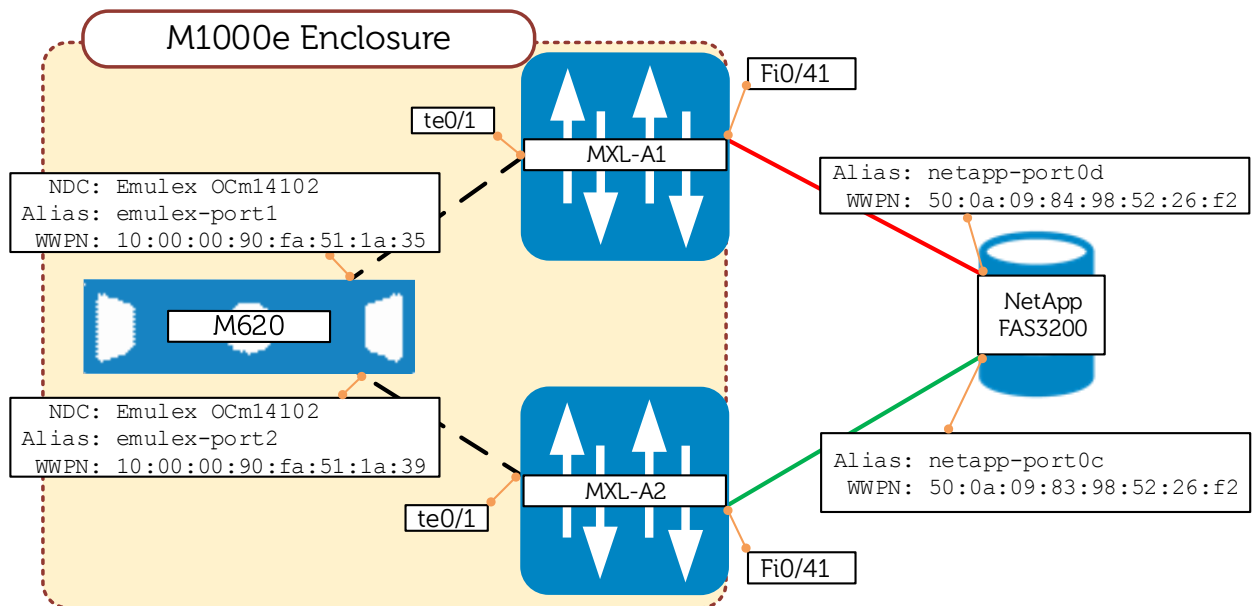
Figure 13   Storage Configuration Using MXLs

D∕ELLEMC

5.3 Switch Configurations

The following section contains the configuration steps required to directly attach storage to either a PowerEdge M IOA or a Dell Networking MXL.

> **Note:** It is recommended that the MXLs or M IOAs be returned to their factory default settings before beginning configuration. To restore factory defaults issue the command:
> `#restore factory-defaults stack-unit 0 clear-all`

Key Points on Configuration

- FC aliases for WWPNs are optional. If preferred, member WWPNs can be entered directly in the FC zone. FC aliases can be any name that makes sense to you, but must begin with a letter.
- A zoneset can contain multiple zones.
- Only one zoneset can be active at any time.
- If any changes are made to aliases, zones, or zonesets, the zoneset must be activated (or reactivated) for the changes to take effect. In our example, this would be done with the following command sequence (Figure 14) on the Dell Networking MXL or the M IOA:

```
switch1(conf)#fcoe-map SAN_FABRIC
switch1 (conf-fcoe-SAN_FABRIC)#fc-fabric
switch1(conf-fmap-SAN_FABRIC-fcfabric)#active-zoneset zoneset1
```

Figure 14   Activate or Reactivate a Zoneset

5.3.1 M I/O Aggregator

This section details configuring the M IOA. Commands to verify the configuration are in section 5.3.3.

> **Note:** The M IOAs should be in programmable MUX (PMUX) mode at this point. Please see the PowerEdge M I/O Aggregator Operational Modes section in the Appendix.

> **Note:** By default, the M IOA base-module ports operate in standalone 4x10GbE mode. To enable port 33 to operate in 40GbE mode, in configuration mode issue the CLI command `no stack-unit 0 port 33 portmode quad` and reload the switch for the changes to take effect.

DELLEMC

Configuration

IOA-A1	IOA-A2
Configure global settings	Configure global settings

```
enable
configure
hostname IOA-A1
default vlan-id 20
no feature fc
feature fc fport domain-id 1
protocol lldp
advertise management-tlv management-
address system-name
exit
```

```
enable
configure
hostname IOA-A2
default vlan-id 21
no feature fc
feature fc fport domain-id 2
protocol lldp
advertise management-tlv management-
address system-name
exit
```

Create DCB maps, FCoE VLAN,FCoE map and zoneset to be used.	Create DCB maps, FCoE VLAN,FCoE map and zoneset to be used.

```
dcb-map SAN_DCB_MAP_A1
priority-group 0 bandwidth 50 pfc off
priority-group 1 bandwidth 50 pfc on
priority-pgid 0 0 0 1 0 0 0 0

interface vlan 87
no shutdown

fc alias broadcom-port1
 member 20:01:5c:f9:dd:17:68:cd

fc alias netapp-port0d
 member 50:0a:09:85:88:57:01:6c

fc zone zone1
member broadcom-port1
member netapp-port0d

fc zoneset zoneset1
member zone1

fcoe-map SAN_FABRIC_A1
fc-map 0efc00
fabric-id 87 vlan 87

fc-fabric
active-zoneset zoneset1
end
```

```
dcb-map SAN_DCB_MAP_A2
priority-group 0 bandwidth 50 pfc off
priority-group 1 bandwidth 50 pfc on
priority-pgid 0 0 0 1 0 0 0 0

interface vlan 88
no shutdown

fc alias broadcom-port2
 member 20:01:5C:F9:DD:17:5E:06

fc alias netapp-port1d
 member 50:0a:09:86:88:57:01:6c

fc zone zone2
member broadcom-port2
member netapp-port1d

fc zoneset zoneset2
member zone2

fcoe-map SAN_FABRIC_A2
fc-map 0efc01
fabric-id 88 vlan 88

fc-fabric
active-zoneset zoneset2
end
```

Figure 15   Initial Steps in Configuring the M IOA

DELLEMC

IOA-A1	IOA-A2
Configure internal facing converged interfaces	Configure internal facing converged interfaces

IOA-A1	IOA-A2
``` configure interface tengigabitethernet 0/1 no switchport portmode hybrid switchport fcoe-map SAN_FABRIC_A1 dcb-map SAN_DCB_MAP_A1 no shutdown interface fortyGigE 0/33 mtu 12000 switchport no shutdown exit ```	``` configure interface tengigabitethernet 0/1 no switchport portmode hybrid switchport fcoe-map SAN_FABRIC_A2 dcb-map SAN_DCB_MAP_A2 no shutdown interface fortyGigE 0/33 mtu 12000 switchport no shutdown exit ```
Configure external facing FC FlexIO ports	Configure external facing FC FlexIO ports
``` interface FibreChannel 0/41 fabric SAN_FABRIC_A1 no shutdown end ```	``` interface FibreChannel 0/41 fabric SAN_FABRIC_A2 no shutdown end ```

Figure 16   Configuring MIOA Physical Interfaces

DELLEMC

5.3.2 Dell Networking MXL

This section details configuring the MXL. Commands to verify the configuration are in section 5.3.3.

Configuration

MXL-A1	MXL-A2
Configure global settings	Configure global settings

<pre>enable
stop bmp
configure
hostname MXL-A1
default vlan-id 20
feature fc fport domain-id 1
protocol lldp
advertise management-tlv management-
address system-name
exit</pre>
	<pre>enable
	stop bmp
	configure
	hostname MXL-A2
	default vlan-id 21
	feature fc fport domain-id 2
	protocol lldp
	advertise management-tlv management-
	address system-name
	exit</pre>

Create DCB maps, FCoE VLAN,FCoE map and zoneset to be used.	Create DCB maps, FCoE VLAN,FCoE map and zoneset to be used.

<pre>dcb-map SAN_DCB_MAP_A1
priority-group 0 bandwidth 50 pfc off
priority-group 1 bandwidth 50 pfc on
priority-pgid 0 0 0 1 0 0 0 0
 	
interface vlan 87
 no ip address
 no shutdown
 	
fc alias emulex-port1
 member 10:00:00:90:fa:51:1a:35
 	
fc alias netapp-port0d
 member 50:0a:09:84:98:52:26:f2
 	
fc zone zone1
member emulex-port1
member netapp-port0d
 	
fc zoneset zoneset1
 member zone1
 	
fcoe-map SAN_FABRIC_A1
 fc-map 0efc00
 fabric-id 87 vlan 87
!
 fc-fabric
  active-zoneset zoneset1
 end</pre>
	<pre>dcb-map SAN_DCB_MAP_A2
	priority-group 0 bandwidth 50 pfc off
	priority-group 1 bandwidth 50 pfc on
	priority-pgid 0 0 0 1 0 0 0 0
	
	interface vlan 88
	 no ip address
	 no shutdown
	
	fc alias emulex-port2
	 member 10:00:00:90:fa:51:1a:39
	
	fc alias netapp-port0c
	 member 50:0a:09:83:98:52:26:f2
	
	fc zone zone2
	member emulex-port2
	member netapp-port0c
	
	fc zoneset zoneset2
	 member zone2
	
	fcoe-map SAN_FABRIC_A2
	 fc-map 0efc01
	 fabric-id 88 vlan 88
	!
	 fc-fabric
	  active-zoneset zoneset2
	 end</pre>

Figure 17   Initial Steps in Configuring the MXL

MXL-A1	MXL-A2
Configure internal facing interfaces	Configure internal facing interfaces

```
configure
interface tengigabitethernet 0/1
 no ip address
 mtu 12000
 no switchport
 portmode hybrid
 switchport
 fcoe-map SAN_FABRIC_A1
 dcb-map SAN_DCB_MAP_A1
 no shutdown

interface fortyGigE 0/33
 no ip address
 mtu 12000
 switchport
 no shutdown
exit
```

```
configure
interface TenGigabitEthernet 0/1
 no ip address
 mtu 12000
 no switchport
 portmode hybrid
 switchport
 fcoe-map SAN_FABRIC_A2
 dcb-map SAN_DCB_MAP_A2
 no shutdown

interface fortyGigE 0/33
 no ip address
 mtu 12000
 switchport
 no shutdown
exit
```

Configure external facing FC FlexIO ports	Configure external facing FC FlexIO ports

```
interface FibreChannel 0/41
 fabric SAN_FABRIC_A1
 no shutdown
end
```

```
interface FibreChannel 0/41
 fabric SAN_FABRIC_A2
 no shutdown
end
```

Figure 18   Configuring MXL Physical Interfaces

5.3.3   Validation

The following section contains various show commands an administrator or network engineer can use to verify that the direct attach storage is operating as expected. The shown commands have been validated for both the MXL as well as the M IOA. The outputs of the following commands are shown below:

- show fc switch
- show fc ns switch
- show fc alias
- show fc zoneset
- show fip-snooping sessions
- show fcoe-map

Notes: In the examples below, only the output from the MXL-A1 configuration is shown.

DELLEMC

show fc switch

This command (Figure 19) displays the FC switch mode and World Wide Name (WWN). This can be useful early in the configuration to ensure that FPort mode is successfully enabled on the switch.

MXL-A1

```
MXL-A1#show fc switch
Switch Mode : FPORT
Switch WWN  : 10:00:d0:67:e5:ca:54:18
```

Figure 19   Show fc switch Command

show fc ns switch

A Fibre Channel fabric uses an addressing scheme to address the ports in the switched fabric. Each port in the switched fabric is assigned a 24-bit address by the FC switch. When the device logs into the switch fabric on a specific port with its WWN, the switch assigns the port address to that port; the switch also maintains the correlation between the port address and the WWN address of the device on that port. This function is implemented by the use of a name server database on the switch.

The `show fc ns switch` command (Figure 20) displays the objects in the name server database of currently attached devices.

MXL-A1

```
MXL-A1#show fc ns switch
Total number of devices =     2

Switch Name                   10:00:d0:67:e5:ca:54:18
Domain Id                     1
Switch Port
FC-Id                         01:01:00
Port Name                     10:00:00:90:fa:51:1a:35
Node Name                     20:00:00:90:fa:51:1a:35
Class of Service              8
Symbolic Port Name            Emulex PPN-10:00:00:90:FA:51:1A:35
Symbolic Node Name            Emulex OCm14102-U2-D FV10.2.315.37 DV10.2.261.4
WIN-V92DSLTBCQB
Port Type                     N_Port

Switch Name                   10:00:d0:67:e5:ca:54:18
Domain Id                     1
Switch Port
FC-Id                         01:29:00
Port Name                     50:0a:09:84:98:52:26:f2
Node Name                     50:0a:09:80:88:52:26:f2
Class of Service              8
Symbolic Port Name            NetApp FC Target Adapter (2432)
netappR117u26:0d
Symbolic Node Name            NetApp FAS3220 (netappR117u26/netappR117u29)
Port Type                     N_Port
```

Figure 20   Show fc ns switch Command

DELLEMC

show fc alias

The complexity of configuring zonesets can be reduced by assigning WWPN addresses to human readable names. As shown below (Figure 21) simple but useful names can be chosen to associate each address to an alias.

The `show fc alias` command can be used to identify the WWPN associated to the alias.

MXL-A1

```
MXL-A1#show fc alias

ZoneAliasName    ZoneMember
=======================================================
emulex-port1     10:00:00:90:fa:51:1a:35
netapp-port0d    50:0a:09:84:98:52:26:f2
```

Figure 21   Show fc alias Command

show fc zoneset

Zoning is a mechanism to ensure only the nodes that are part of a zone can communicate with each other. Zoning prevents unauthorized access of storage assets. In this deployment the NDC WWPN and the WWPN for the NetApp device are added to the zoneset.

Issue the command `show fc zoneset` to display the zone, the node members in the zone and the zoneset itself (Figure 22). While in this deployment there is a single zoneset (zoneset1) and a single zone (zone1), a zoneset can contain multiple zones.

MXL-A1

```
MXL-A1#show fc zoneset
ZoneSetName                     ZoneName                     ZoneMember
============================================================================
    zoneset1
                                    zone1
                                                                emulex-port1
                                                                netapp-port0d

Active Zoneset: zoneset1


ZoneName                        ZoneMember
=======================================================
zone1
                                10:00:00:90:fa:51:1a:35
                                50:0a:09:84:98:52:26:f2
```

Figure 22   Show fc zoneset Command

DELLEMC

show fip-snooping sessions

For further detail on FIP sessions the command `show fip-snooping sessions` can be issued (Figure 23). This command will list all interfaces, MAC addresses and WWPN addresses used for any server participating in a fabric login.

> **Note:** Due to width, the following output has been split into two lines.

MXL-A1

```
MXL-A1#show fip-snooping sessions
Enode MAC          Enode Intf    FCF MAC           FCF Intf       VLAN
---------          ----------    -------           --------       ----
00:90:fa:51:1a:35  Te 0/1        d0:67:e5:ca:54:1b  ~              87
```

```
FCoE MAC           FC-ID     Port WWPN              Port WWNN
--------           -----     ---------              ---------
0e:fc:00:01:01:00  01:01:00  10:00:00:90:fa:51:1a:35  20:00:00:90:fa:51:1a:35
```

Figure 23   Show fip-snooping sessions Command

DELLEMC

show fcoe-map

The show fcoe-map command is useful to identify the SAN fabric to which FCoE storage traffic is sent. As shown below (Figure 24) this command can be viewed as a summary command showing the fabric type, the fabric ID, fabric VLAN, the FC-MAP ID, the FC Domain ID and the active zoneset. Additionally the participating interfaces are also shown at the bottom of the output.

```
                              MXL-A1

MXL-A1#show fcoe-map

Fabric Name          SAN_FABRIC

Fabric Type          Fport
Fabric Id            87
Vlan Id              87
Vlan priority        3
FC-MAP               0efc00
FKA-ADV-Period       8
Fcf Priority         128
Config-State         ACTIVE
Oper-State           UP
=========================================================
Switch Config Parameters
=========================================================
DomainID             1
=========================================================
Switch Zoning Parameters
=========================================================
Default Zone Mode:       Deny
Active Zoneset:          zoneset1
=========================================================
Members
Fc 0/41
Te 0/1
```

Figure 24   Show fcoe-map command

DELLEMC

5.3.4 Optional Configuration: F-port Without Zoning

In F-port mode, the fcoe-map has the default zone mode set to deny. This setting denies all the fabric connections unless included in an active zoneset (as done above in Figure 17). To change this setting, use the `default-zone-allow` command. This change will allow all the fabric connections without zoning.

> **Note:** On PowerEdge M IOAs in standalone mode, this is the default behavior allowing all fabric connections without any additional zoning.

MXL or M IOA

```
switch(conf)#fcoe-map SAN_FABRIC
switch(conf-fcoe-SAN_FABRIC)#fc-fabric
switch(conf-fmap-SAN_FABRIC-fcfabric)#default-zone-allow all
```
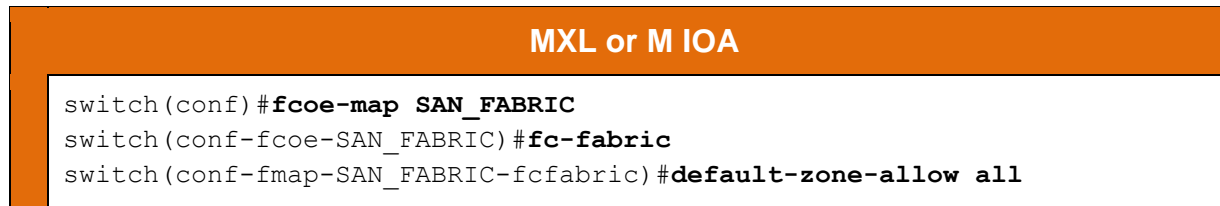
Figure 25   default-zone-allow Command

6   Server Configuration

In this section, in order to meet the needs of the environment, two different server configurations are presented. The first configuration covers installing, configuring and validating VMware ESXi. The second covers configuring and verifying the function of Multipath I/O (MPIO) inside Windows Server 2012 R2.

6.1   VMware ESXi 5.5 - Installation, Configuration and Validation

This section will guide you through downloading, installing and basic configuration of the Dell custom ESXi 5.5 Update 2 image, which can be download from support.dell.com. The ISO will need to be downloaded and either burned to a CD or a third party utility used to create a bootable USB key.

Installing Dell Custom VMware ESXi 5.5 U2

This section provides an outline of the installation process for VMware ESXI 5.5 U2. For further in-depth information on the installation of ESXi, please visit the VMware vSphere 5.5 Documentation Center at https://pubs.vmware.com/vsphere-55/index.jsp.

1. Insert the Dell custom ESXi 5.5 installation media into the server.
   This can either be a CD/DVD, a USB flash drive or mount the installation ISO through the iDRAC interface for the M620.
2. Set the BIOS to boot from the media. In most cases, this will be **Virtual CD**.
3. On the Welcome screen, press **Enter** to continue.
4. On the End User License Agreement (EULA) screen press **F11** to accept.
   At this point, the installer will scan for suitable installation targets. In this scenario, ESXi is installed on an internal SD card.
5. Select the keyboard type for the host. This will be **US Default** in most cases.
6. Enter a password for the host.
7. On the Confirm Installation window, press **Enter** to start the installation.
8. When the installation is complete, remove the installation CD, DVD, USB flash drive, or unmount the Virtual CD.
9. Press **Enter** to reboot the host.

Connecting to the ESXi Host with the vSphere Client

Once installation has been completed, access the console for the host. From here, a management NIC can be activated and an IP address assigned. Follow the steps below to complete this.

Setting up the Management NIC.

1. Press **F2** to Customize System.
2. Select Configure Management Network and press Enter.
3. Select Network Adapters and press Enter.
4. Any NIC that is currently connected with show **Connected** under Status, use the spacebar to select the appropriate management NIC.
5. Press **Enter** to exit.
6. Press **Esc**.

DELLEMC

7. Press **Y** for Yes to restart the management network.
8. Press **Esc** to exit System Customization.

Configuring a Static IP Address

1. Press **F2** to Customize System.
2. Select Configure Management Network and press Enter.
3. Select IP Configuration and press Enter.
4. Select Set static IP address and networking configuration and press the spacebar.
5. Type in an IP Address, Subnet Mask and Default Gateway and press Enter.
6. Press **Esc**.
7. Press **Y** for Yes to restart the management network.
8. Press **Esc** to exit System Customization.

Connecting to the ESXi host using the vSphere Client

1. Use a web browser to go to **http://<IP_of_ESXi_host>**.
2. Click on the link Download vSphere Client.
3. Launch the installer, agreeing to all the defaults for installation.
4. Launch the VMware vSphere Client.
5. Enter the **IP address of the ESXi host**, the **username** (root) and the **password** configured during the installation. Press **Enter**.

Verifying MPIO Through the vSphere Client

Once connected to the ESXi host through the vSphere Client you can verify that a LUN is presented to the ESXi host. Additionally you will be able to see the number of paths for each LUN, this is useful in verifying that both SAN A and SAN B are working. Figure 26 shows the final step, which presents the number of targets, number of devices and the number of paths.

1. Select your host from the left hand column
2. Click on the **Configuration** tab
3. Click **Storage Adapters** under Hardware
4. Locate your CNA in the list of storage adapters

Note: In the event the storage adapter is not listed, perform the following steps to install the software FCoE adapter.

   a. Click **Add** in the Storage Adapters window

   b. Select **Add Software FCoE Adapter**

   c. Click **OK**

5. In the Details Window, under Storage Adapters (Figure 26) you should see the two LUNs presented to your host. Make note of **Targets**, **Devices** and **Paths** as well.

Figure 26 Broadcom NetXtreme BCM57810 Software FCoE Adapter

6.2 Microsoft Server Operating System MPIO Installation and Configuration

Microsoft Multipath I/O (MPIO) must be installed and configured for the Windows operating system to properly recognize multiple paths to a storage device.

> **Note:** The following steps are specific to Windows Server 2012 R2. Microsoft MPIO is also supported on Windows Server 2012 (pre-R2), 2008 and 2008 R2. The configuration process will be similar for all versions.

1. In Windows Server 2012, select **Add Roles and features** in the Server Manager Dashboard.
2. On the Select installation type page, select **Role-based or feature-based installation** and click **Next.**
3. On the Select destination server page, choose **Select a server from the server pool**, make sure your server is highlighted, and click **Next**.
4. On the Select server roles page, click **Next** to skip.
5. On the Select features page, scroll down and check the box next to **Multipath I/O** (Figure 27).
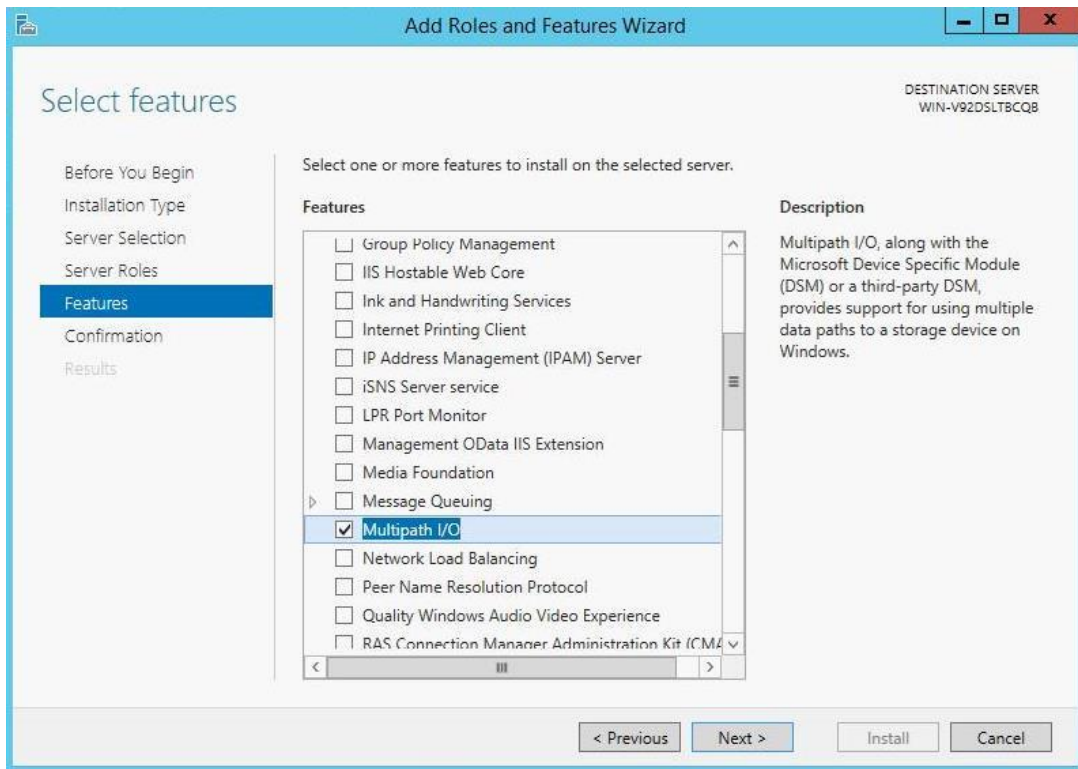
Figure 27   Windows 2012R2 – Add Roles and Features Wizard

6. Click **Next**.
7. On the Confirm installation selections page, click **Install**.
8. On the Installation progress page, click **Close** when complete.

Configuring MPIO for your Storage Array

In Windows, launch **Control Panel** or **Administrative Tools.** An MPIO icon should appear. Launch MPIO and open the **Discover Multi-Paths** tab. If your storage array is SPC-3 compliant, it should appear in the **Device Hardware Id** box (Figure 28). If your storage array does not appear, you may need to install a Device-Specific Module (DSM) from your storage vendor. Check your storage documentation for more information on the MPIO configuration requirements for your device. Highlight the **Device Hardware ID** ("NETAPP LUN" in this case) and click **Add**. Reboot your system if prompted.
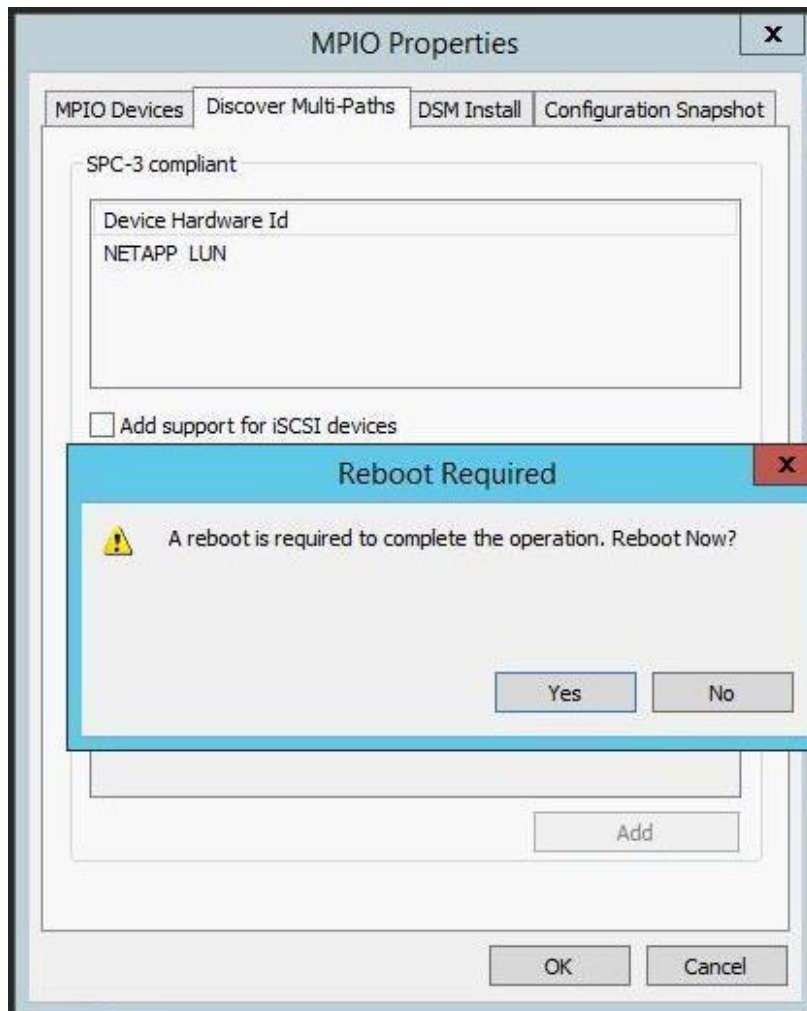
Figure 28   MPIO Reboot Prompt

After rebooting, open MPIO. The device should now appear on the MPIO Devices tab (Figure 29).
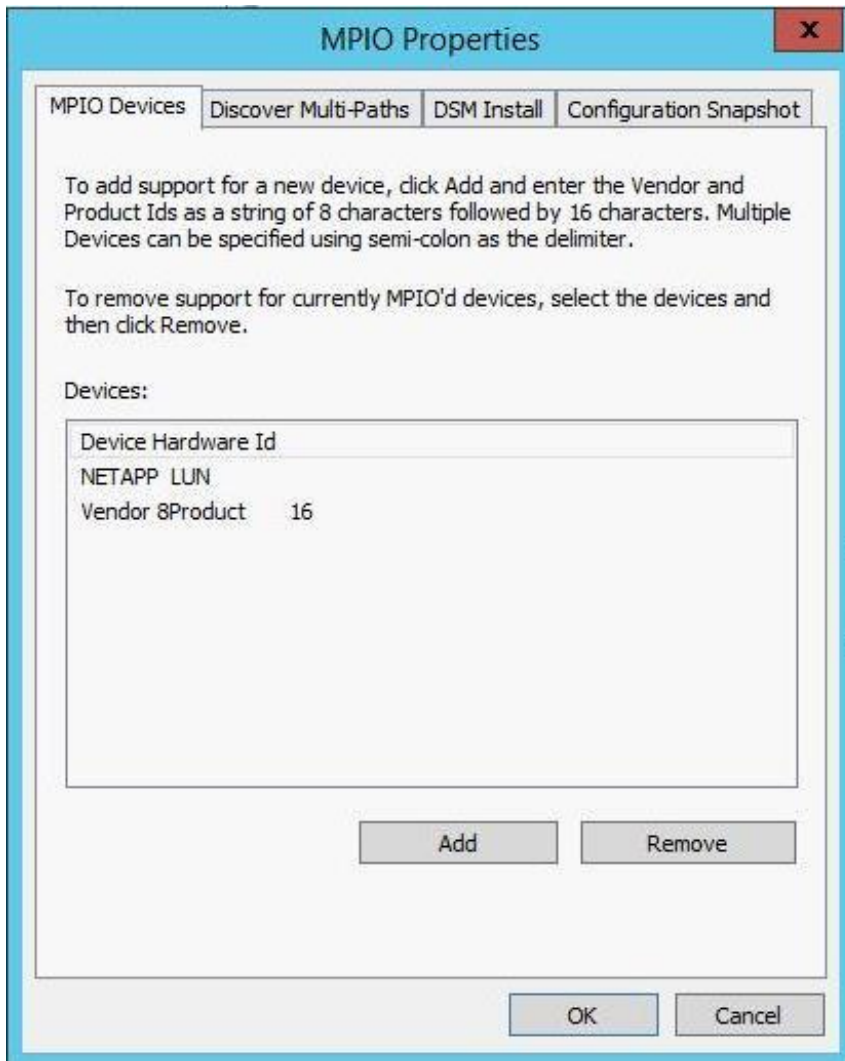
Figure 29   MPIO Confirmation

Create Disks and Format

In Windows, go **to Server Manager** > **Tools** > **Computer Management** > **Storage** > **Disk Management**. Right-click on **Disk Management** and select **Rescan Disks**. If MPIO is configured properly, only one instance of each LUN configured on your storage array should appear. If there are two instances of each LUN, there is likely an issue with Microsoft MPIO configuration. Use the **Disk Management** utility to bring the disks online, assign drive letters and format as needed.

Configure MPIO

The Windows Disk Management utility can also be used to configure MPIO. Right click on each remote disk (not the Volume) and select **Properties**.

DELLEMC
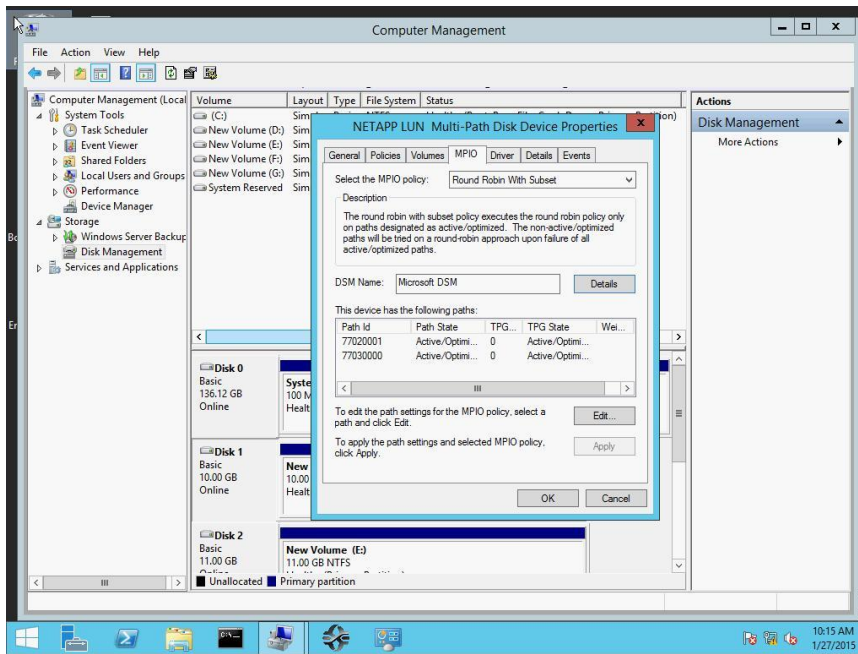
Figure 30   Windows MPIO Control Panel Applet

The **Multi-Path Disk Device Properties** window (Figure 30) is launched. Click on the **MPIO** tab. Note the option to configure the MPIO policy (**Failover Only**, **Round Robin with Subset**, etc.). The DSM Details button can be used to configure the redundant path verify period and other timer settings. Under **This device has the following paths:**, the two available paths detected by MPIO are shown.

Verify Redundant Path Functionality

Shut a port on one of the paths between the server and storage, or reload an MXL or M IOA. Verify the server can still connect to all disks. Repeat for the other path.

Note: By default, it may take Microsoft MPIO about 30 seconds to recognize a failed path. Microsoft has information about tuning MPIO parameters to reduce this time on their Technet web site, http://technet.microsoft.com.

7 Conclusion

Dell Networking MXL and M I/O Aggregator with installed FC FlexIO provide a reliable and evolutionary step in the right direction for lowering the necessary Data Center costs typically associated with storage and for providing an easy to use configuration. The Direct-connect topologies will ease switch needs and bring great benefits in today's converged infrastructure.

DELLEMC

A   Appendix

A.1   Configuration Details

This section lists all of the devices and their respective firmware/driver levels used while validating the solution in this guide.

Table 2   PowerEdge M IOA, VMware ESXi, NetApp Configuration Table

Component	Description
Operating system	VMware ESXi 5.5.0 Release Build 2068190
Server Platform	M620
BIOS Version	2.4.3
NDC	Broadcom NetExtreme II Gb Ethernet BCM57810
NDC Firmware	7.6.15
Blade Switch	PowerEdge M IOA
PowerEdge M IOA Firmware	9.7.0.0
Storage Array	NetApp FAS3220 version 8.2.1

Table 3   Dell Networking MXL, Microsoft Windows 2012R2, NetApp Configuration Table

Component	Description
Operating system	Microsoft Windows Server 2012 R2 Standard build 9600
Server Platform	M620
BIOS Version	2.4.3
NDC	Emulex OCm14102-U2-D
NDC Driver / Firmware	Driver - VJTR1 A00 (10.2.298.37), Firmware – NF2F6 A00 (10.2.315.37)
Blade Switch	PowerEdge MXL
PowerEdge MXL Firmware	9.7.0.0
Storage Array	NetApp FAS3220 version 8.2.1

DELLEMC

A.2   PowerEdge M I/O Aggregator Operational Modes

The IOA supports four operational modes: Standalone (SMUX), VLT, Stack and Programmable MUX (PMUX). See Table 4 for detailed descriptions of each mode. To enable a new operational mode the command **`stack-unit 0 iom-mode <IOA_Mode>`** is issued in configuration mode. After enabling a new operational mode, the switch must be reloaded.

> **Note:** When switching modes it is important to factory restore the switch first: **`restore factory-defaults stack-unit 0 clear-all`** and then set the switch mode accordingly.

By default, in Standalone and VLT modes all external ports are configured in a single port channel (128) and all VLANs (1-4094) are tagged on this port channel. Additionally any DCBx protocol options are allowed as well as iSCSI or FCoE settings.

Table 4   M IOA Modes and Descriptions

IOA Mode	Description
Standalone mode (SMUX)	This is the default mode for M IOA. It is a fully automated, low-touch mode, which allows VLAN memberships to be defined on the server-facing ports while all upstream ports are configured in port channel 128 (and cannot be modified).
VLT mode	This is a low-touch mode where all configurations except VLAN membership are automated. In this mode, port 9 is dedicated to VLT interconnect.
Programmable MUX mode (PMUX)	This mode provides flexibility of operation by allowing the administrator to create multiple LAGs, configure VLANs on uplinks and to configure DCB parameters on the server side.
Stack mode	This mode allows up to six M IOAs to be stacked as a single logical switch. The stack units can be in the same or different chassis. This is a low-touch mode where all configurations except VLAN membership are automated.

> **Note:** Virtual Link Trunking (VLT) allows physical links between two chassis to appear as a single virtual link to the network core or other switches (Edge, Access or ToR). VLT reduces the role of Spanning Tree protocols by allowing LAG terminations on two separate distribution or core switches, and by supporting a loop free topology. VLT provides Layer 2 multi-pathing, creating redundancy through increased bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

> **Note:** You cannot configure MXL or M IOA switches in Stacking mode if the switches contain the FC Flex IO module. Similarly, FC Flex IO Modules do not function when you insert them in to a stack of MXL or M IOA switches.

DELLEMC

A.3   PowerEdge M1000e Port Mapping

The connections are 10 Gigabit Ethernet connections for basic Ethernet traffic, iSCSI storage traffic or FCoE storage traffic. In a typical M1000e configuration of 16 half-height blade servers, ports 1-16 are used and 17 - 32 disabled. However if quad port adapters or quarter-height blade servers are used, ports 17-32 will be enabled.

Table 5 lists the port mapping for the two expansion slots on the Dell Networking MXLs and M IOAs as well as the internal 10/1 GbE interfaces on the blade servers installed in the M1000e chassis. For information on internal port mapping please see the attachment **m1000e_internal_port_mapping.pdf**

Table 5   Port-Mapping for the M1000e Blade Enclosure

Dell Networking MXL and Dell PowerEdge M I/O Aggregator – Port Mapping

	QSFP+ 2x40Gb	QSFP+ 8x10GB SFP+ (breakout)	SFP+ 4x10Gb	10G-BaseT 4x10Gb	FC8 x 4
Expansion Slot 1		56			
		55			
		54			
	53	53			
		52	52	52	52
		51	51	51	51
		50	50	50	50
	49	49	49	49	49
	QSFP+ 2 X 40Gb	QSFP+ 8 X 10GB SFP+ (breakout)	SFP+ 4 X10Gb	10G-BaseT 4 X 10Gb	FC8 x 4
Expansion Slot 0		48			
		47			
		46			
	45	45			
		44	44	44	44
		43	43	43	43
		42	42	42	42
	41	41	41	41	41
	QSFP+ 2 X 40Gb	QSFP+ 8 X 10GB SFP+ (breakout)	SFP+ 4 X10Gb	10G-BaseT 4 X 10Gb	FC8 x 4
Fixed QSFP Ports		40			
		39			
		38			
	37	37			
		36	.	.	.
		35	.	.	.
		34	.	.	.
	33	33	.	.	.
Internal 10 / 1 GB interfaces					
Internal 10/1 Gb	32	32	32	32	32
	31	31	31	31	31
	.	.	.	.	.
	.	.	.	.	.
	.	.	.	.	.
	2	2	2	2	2
	1	1	1	1	1

A.4 Fibre Channel over Ethernet and Data Center Bridging

Fibre Channel over Ethernet (FCoE) is a networking protocol that encapsulates Fibre channel frames over Ethernet networks. The FCoE protocol specification replaces the FC0 and FC1 layers of Fibre Channel stack with Ethernet. By retaining the native Fibre Channel constructs, FCoE can integrate with existing Fibre Channel fabrics and management solutions.

> **Note:** FCoE (referenced as FC-BB_E in the FC-BB-5 specifications) achieved standard status in June 2009, and is documented in the T11 publication (http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf).

FCoE operates directly above Ethernet in the network protocol stack, in contrast to iSCSI which runs on top of TCP and IP. As a consequence, FCoE cannot be routed across IP networks. Once de-encapsulation occurs the FC packets can be routed accordingly with FC switches.

Since traditional Ethernet does not provide priority-based flow control, FCoE requires modifications to the Ethernet standard to support priority-based flow control mechanisms (this reduces frame loss from congestion). The IEEE standards body added priorities via Data Center Bridging (DCB). The three primary extensions are:

- Encapsulation of native Fibre Channel frames into Ethernet frames.
- Extensions to the Ethernet protocol itself to enable Lossless Ethernet links.
- Mapping between Fibre Channel N_Port Ids (aka FCIDs) and Ethernet MAC address.

The primary purpose of the FCoE protocol is for traffic destined for FC/FCoE Storage Area Networks (SANs). FCoE enables cable reduction due to converged networking possibilities. To achieve these goals three hardware components must be in place.

- Converged Network Adapters (CNA)
- Lossless Ethernet Links (via DCB extensions)
- An FCoE capable switch, typically referred to as a Fibre Channel Forwarder (FCF)

A Fibre Channel Initialization Protocol (FIP) Snooping Bridge (FSB) is a fourth optional component which can be introduced and still allow full FCoE functionality. In traditional Fibre Channel networks, FC switches are considered trusted, while other FC devices must log directly into the switch before they can communicate with the rest of the fabric. This login process is accomplished through a protocol called FIP which operates at L2 for end point discovery and fabric association. With FCoE an Ethernet bridge typically exists between the End Node (ENode) and the FCF. This bridge prevents a FIP session from properly establishing. To allow ENodes the ability to login to the FCF, FSB is enabled on the Ethernet Bridge. By snooping on FIP packets during the discovery and login process, the intermediate bridge can implement data integrity using ACLs that permit valid FCoE traffic between the Enode and FCF.

Data Center Bridging (DCB) is a collection of mechanisms that have been added to the existing Ethernet protocol. These mechanisms allow Ethernet to become lossless which is a prerequisite for FCoE. The four additions are:

DELLEMC

- Priority-based Flow Control (PFC) (IEEE 802.1Qbb)
- Enhanced Transmission Selection (ETS) (IEEE P802.1Qaz)
- Data Center Ethernet Bridging Capability Exchange Protocol (DCBX)

Support and Feedback

Contacting Technical Support

Support Contact Information   Web: http://Support.Dell.com/

                                                            Telephone: USA: 1-800-945-3355

Feedback for this document

We encourage readers of this publication to provide feedback on the quality and usefulness of this deployment guide by sending an email to Dell_Networking_Solutions@Dell.com

DELLEMC