



Dell EqualLogic Technical Article

Dell EqualLogic PS-M4110 Blade Array Technical Article

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell. Dell, the DELL logo, and the DELL badge, PowerConnect™, EqualLogic™, Force10™, PowerEdge™ and PowerVault™ are trademarks of Dell Inc. Broadcom® is a registered trademark of Broadcom Corporation. Brocade is a registered trademark and VCS is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Intel® is a registered trademark of Intel Corporation in the U.S. and other countries. Microsoft®, Windows®, Windows Server®, and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Qlogic is a registered trademark of QLogic Corporation.

Table of contents

1	Executive summary	5
2	Audience	5
3	Introduction	5
4	Reference architecture overview	6
4.1	PowerEdge M-Series blade servers.....	6
4.1.1	Additional configuration information	7
4.1.2	Network adapter information.....	7
4.2	IO module information.....	8
4.3	Array information.....	8
4.4	Reference architecture network configuration	8
4.4.1	SAN infrastructure.....	8
5	Switch configuration	10
5.1	Force10 MXL.....	10
5.1.1	Global Switch settings:.....	10
5.1.2	Array and Host Port configuration:.....	10
5.1.3	LACP configuration (Non-Stack configurations):.....	10
5.1.4	Stack configuration (Non-LAG configurations):	11
5.2	PowerConnect M8024-K	12
5.2.1	iSCSI Auto-Configuration.....	12
5.2.2	Global Switch settings:.....	12
5.2.3	Host Port configuration:.....	13
5.2.4	LACP configuration (Non-Stacked configurations):.....	13
5.2.5	Stack Configuration (Non-LAG Configurations):	14
5.3	PowerConnect M8428-k.....	15
5.3.1	Array and Host Port configuration:.....	15
5.3.2	LACP Member Port configuration:.....	16
6	General recommendations and guidelines.....	18
A.	Port Map	19
B.	Additional resources	22

Acknowledgements

This whitepaper was produced by the PG Technical Marketing for EqualLogic Infrastructure and Networking.

The team that created this whitepaper:

Guy Westbrook and Tony Ansley

1 Executive summary

This paper provides guidance for incorporating a data center-in-a-box solution with the new Dell EqualLogic™ PS-M4110 Blade Array, PowerConnect™ or Force10™ Blade IO Modules, and PowerEdge M-Series blade servers to form a reliable, stable, and well performing consolidated solution. This paper focuses on the networking recommendations for creating a self-contained, dedicated SAN network within a single M1000 blade chassis. Additional information regarding scalable M1000 blade chassis solutions can be found in the [EqualLogic Configuration Guide](#). See the Appendix B of this document for links to additional resources.

2 Audience

This white paper is primarily intended for those who are involved in defining, implementing, or supporting an EqualLogic storage network using the new EqualLogic PS-M4110 blade array as part of a strategy that deploys storage with the purpose of providing data for use by the blade servers located in the same Dell M1000e Blade Enclosure.

This document assumes the reader is familiar with EqualLogic storage operation and general networking fundamentals as well as being familiar with the M1000e chassis architecture and its configuration.

3 Introduction

Dell's EqualLogic product suite continues to drive innovation by delivering all the functionality and enterprise-class features of its traditional rack-based arrays in a new blade chassis form factor product. The newest member of the Dell EqualLogic PS Series family is the scalable and easy-to-manage Dell EqualLogic PS-M4110 blade array designed for use with the Dell PowerEdge™ M1000e blade chassis product.

The PS-M4110 is targeted at small to medium blade chassis customers demanding shared storage with advanced features at lower cost and with lower space, heating, and electrical requirements than other storage products on the market. The tight integration of PS-M4110 along with the Dell M-Series chassis and blade servers presents some new and unique use cases for shared storage.

Of particular interest is a use case for a completely converged server, network, and storage infrastructure called data center-in-a-box. In this deployment model, the M1000e, along with integrated servers, networking, and shared storage, becomes a modular building block for customers wanting to grow their infrastructure – virtual or physical – as they need additional capacity while minimizing the data center impact in terms of power, cooling, and rack space.

4 Reference architecture overview

The Storage Interoperability Lab conducted preliminary testing on a self-contained M1000e based data center-in-a-box architecture that included testing of multiple networking components as well as using blade servers from the latest two generations of M-Series blade servers. The following sections describe the various sub-systems used:

4.1 PowerEdge M-Series blade servers

The blade servers used servers from the last two generations of PowerEdge M-Series servers. Each server is running Windows Server 2008R2 with Service Pack 1. Detailed M1000e and Server configuration information is provided below.

Figure 1 Tested server configurations

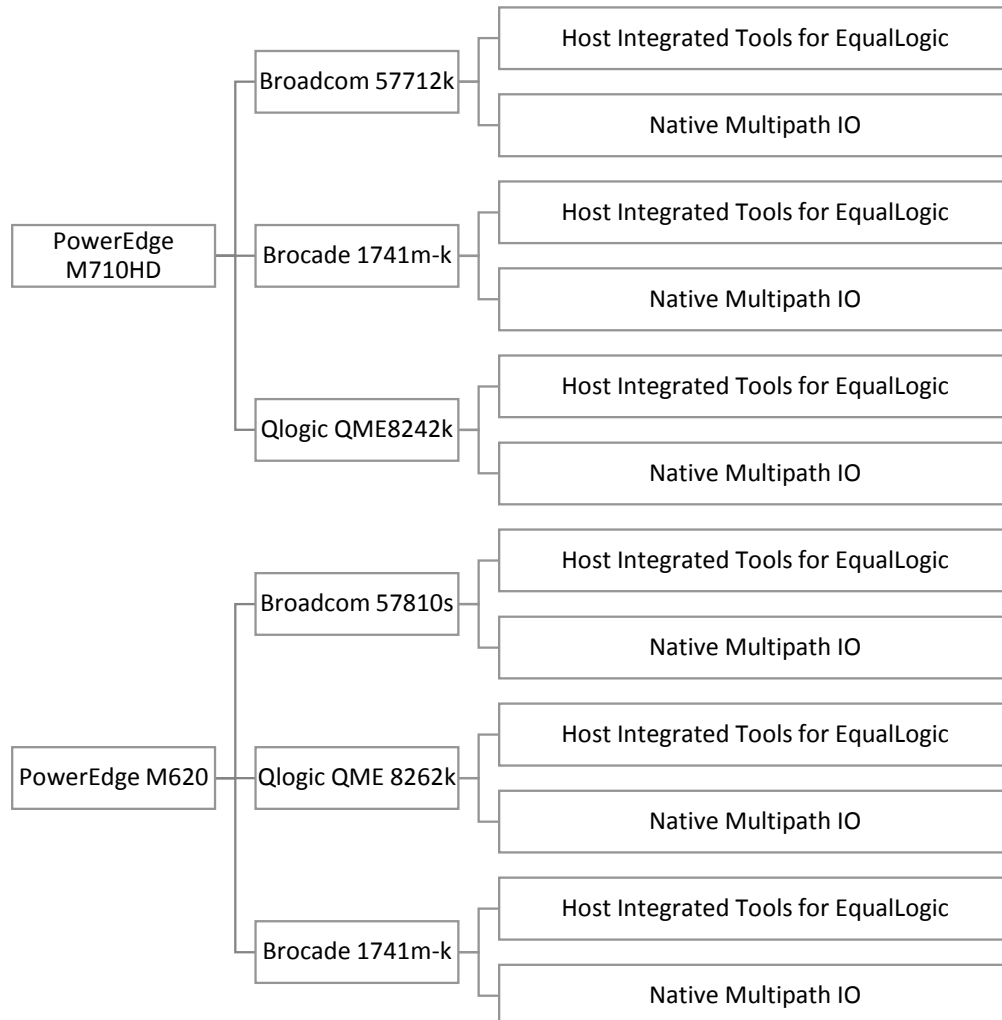



Figure 2 Blade Chassis configuration (Front)



4.1.1 Additional configuration information

Chassis Management Console (CMC) version: 4.1
 EqualLogic Host Integration Tools version: 4.0
 Microsoft Native MPIO – V6.1 Build 7601

4.1.2 Network adapter information

Table 1 NIC/CNA and Driver version tested

Network adapter	Driver version
QLogic QME8242k	4.5.14.1020
Brocade® 1741m-k	3.0.1.0
Broadcom® 57712k	7.2.8.0
Broadcom 57810S-k	7.2.8.0
QLogic® QME8262k	4.6.12.217

4.2 IO module information

LAN IO modules were installed in Fabric Slots A1 and A2 for all tests.

SAN IO modules were installed in Fabric Slots B1 and B2 for all tests (except for tests involving the Broadcom 57712k which is a Network Daughter Card (NDC) where Fabric A was used). The IO modules tested are listed in Table 2 below.

Table 2 IOMs and firmware version tested

IOM	Interconnect modes	Firmware version
Force10 MXL	Stack & LAG	8.3.16.1
PowerConnect M8024-k	Stack & LAG	4.2.2.3
PowerConnect M8428-k	LAG	6.3.1_dcb1

4.3 Array information

Two PS-M4110 arrays running Array Software version 6.0.1 were installed in the M1000e slots 5/6 and 13/14 for all tests.

4.4 Reference architecture network configuration


4.4.1 SAN infrastructure

A pair of Force10 or PowerConnect 10Gb Ethernet IO modules reside in Fabric “B” (Slots “B1” and “B2”). These IO modules are interconnected using either native stacking or industry standard link aggregation based on IEEE Link Aggregation Control Protocol. Either option may be used based on the needs of the customer or pre-defined networking policy. By interconnecting these IO modules, a single, layer 2 Ethernet network fabric is created to meet the networking requirements of the EqualLogic PS Series arrays. The PS-M4110 supports the use of networks in either Fabric “A” or Fabric “B” of the M1000e chassis.

A second set of IO modules reside in Fabric “A” and are used to host LAN or Client traffic. They need to be uplinked to an access/aggregation layer or top-of-rack switch to allow the blade servers to provide client-server communications with the rest of the data center.

The reference architecture is illustrated in Figure 3 below.

Figure 3 Sample MXL datacenter in a box solution



5 Switch configuration

Any of the currently available 10Gb Ethernet IO modules can be used as SAN switches within this reference architecture. The following sections provide the detailed list of CLI commands used to configure each switch model for use as a SAN switch to support the PS-M4110 Storage Blades as defined in the reference architecture described in Section 4.

5.1 Force10 MXL

5.1.1 Global Switch settings:

```
FTOS>en

FTOS#

FTOS#conf

FTOS (conf) #no dcb enable

FTOS (conf) #no iscsi enable
```

Note: The above step is required even though iSCSI is enabled by default

```
FTOS (conf) #iscsi enable

FTOS (conf) #exit

FTOS# wr mem
```

5.1.2 Array and Host Port configuration:

```
FTOS (conf) #interface range tengigabitethernet 1/1 - 16

FTOS (conf-if-range-te-1/1-16) #switchport

FTOS (conf-if-range-te-1/1-16) #no shutdown

FTOS (conf-if-range-te-1/1-16) #exit

FTOS (conf) #exit

FTOS#wr mem
```

5.1.3 LACP configuration (Non-Stack configurations):

```
FTOS(conf)#interface range fortyGigE 1/33 - 37

    % Warning: Non-existing ports (not configured) are ignored
    by interface-range

FTOS (conf-if-range-fo-1/33-37) #mtu 9216
```

```
FTOS (conf-if-range-fo-1/33-37) port-channel-protocol LACP
FTOS (conf-if-range-fo-1/33-37-lacp) port-channel 1 mode active
FTOS (conf-if-range-fo-1/33-37) #no shutdown
FTOS (conf-if-range-fo-1/33-37) #exit
FTOS (conf) #exit
FTOS#wr mem
FTOS#conf
FTOS (conf) #interface Port-channel 1
FTOS (conf-if-po-1) # mtu 9216
FTOS (conf-if-po-1) # switchport
FTOS (conf-if-po-1) # no shutdown
FTOS (conf-if-po-1) # exit
FTOS (conf) # exit
FTOS# wr mem
```

Note: Repeat the same steps on the second switch

5.1.4 Stack configuration (Non-LAG configurations):

```
FTOS>en
FTOS#
FTOS#conf
FTOS (conf) #stack-unit 1 stack-group 0
FTOS (conf) #stack-unit 1 stack-group 1
```

Note: Repeat the same steps on the second switch. Stack-unit number may differ as per the user settings

5.2 PowerConnect M8024-K

5.2.1 iSCSI Auto-Configuration

The PowerConnect M8024-k IO module supports the “iSCSI Auto-configuration” feature. When enabled, the switches will automatically enable link level flow control on all ports of the switch and configure the switch for jumbo frames (mtu 9216).

Monitoring for EqualLogic storage arrays via LLDP is also enabled by this command. Upon detection of an EqualLogic array, the specific interface involved will have spanning-tree portfast enabled and unicast storm control disabled. These changes appear in the running config.

Disabling iSCSI Optimization does not disable flow control, portfast, or storm control configuration applied as a result of enabling iSCSI Optimization.

Note: iSCSI Auto-configuration only configures the global switch settings as described here and the individual port settings for ports that have active EqualLogic iSCSI port connections. Host port connections must still be configured manually.

```
console>enable
console#config
console(config)#iscsi enable
```

5.2.2 Global Switch settings:

Note: Flow control is enabled by default on the PowerConnect M8024k.

If the state of the flow control setting is unknown, use the following commands to enable:

```
console>enable
console#config
console(config) #no mode simple
<Followed by reboot>
console>enable
console#config
console(config)#flowcontrol
console(config)#exit
console# copy running-config startup-config
console#exit
```

Spanning Tree is a global configuration setting for the PowerConnect M8024k.

Spanning Tree is enabled by default on the PowerConnect M8024k.

Rapid Spanning Tree (rstp) is the default mode for spanning tree for the PowerConnect M8024k.

If the state of the spanning-tree settings is unknown, use the following commands to enable spanning tree:

```
console>enable
console#config
console(config)#spanning-tree
console(config)#exit
console# copy running-config startup-config
console#exit
```

5.2.3 Host Port configuration:

```
console>enable
console#config
console(config)#interface range te 1/0/1-1/0/16
console(config-if)#spanning-tree portfast
console(config-if)#mtu 9216
console(config-if)#no lldp transmit
console(config-if)#no shutdown
console(config-if)#exit
console(config)#exit
console# copy running-config startup-config
console#exit
```

5.2.4 LACP configuration (Non-Stacked configurations):

```
console>enable
console#config
console(config)#interface range tengigabitethernet 1/0/17-1/0/20
```

```
console(config-if)#channel-group 1 mode auto
console(config-if)#mtu 9216
console(config-if)#exit
console(config)#interface port-channel 1
console(config-if-ch1)#mtu 9216
console(config)#exit
console# copy running-config startup-config
console#exit
```

Note: Repeat the same steps on the second switch

5.2.5 Stack Configuration (Non-LAG Configurations):

```
console#configure
console(config)#stack
console(config-stack)#stack-port tengigabitethernet 1/0/17 stack
console(config-stack)#stack-port tengigabitethernet 1/0/17 stack
console(config-stack)#stack-port tengigabitethernet 1/0/17 stack
console(config-stack)#stack-port tengigabitethernet 1/0/17 stack
console(config-stack)#exit
console(config)#exit
console# copy running-config startup-config
console#reload
Are you sure you want to reload the stack? (y/n) y
```

Note: On switch two, you must renumber the switch ID (default switch ID is 1). Renumber the switch to switch ID 2 before creating and adding to the stack.

To change the switch ID number perform the following command on IO module #2:

```
console#configure
console(config)#switch 1 renumber 2
Are you sure you want to renumber? (y/n)y
```

Create the stack on switch 2:

Put ports into stacking mode. Ports 17-20 were used as stacking ports.

```
console#configure
console(config)#stack
console(config-stack)#stack-port tengigabitethernet 2/0/17 stack
console(config-stack)#stack-port tengigabitethernet 2/0/18 stack
console(config-stack)#stack-port tengigabitethernet 2/0/19 stack
console(config-stack)#stack-port tengigabitethernet 2/0/20 stack
console(config-stack)#exit
console(config)#exit
```

5.3 PowerConnect M8428-k

5.3.1 Array and Host Port configuration:

```
M8428-k>
M8428-k>en
M8428-k#conf t
M8428-k (conf) #interface intengigabitethernet 0/1
M8428-k (conf-if-te-0/1) # mtu 9208
M8428-k (conf-if-te-0/1) # switchport
M8428-k (conf-if-te-0/1) # switchport mode access
M8428-k (conf-if-te-0/1) # lldp disable
M8428-k (conf-if-te-0/1) # qos flowcontrol tx on rx on
M8428-k (conf-if-te-0/1) # no shutdown
M8428-k (conf-if-te-0/1) # end
M8428-k# wr mem
```

Note: There is no range command available for M8428-k. So repeat the same steps for intengigabitethernet 0/2 – 0/16 one by one.

5.3.2 LACP Member Port configuration:

```
M8428-k>
M8428-k>en
M8428-k#conf t
M8428-k (conf) #interface Extengigabitethernet 0/17
M8428-k (conf-if-te-0/17)# channel-group 2 mode active type
standard
M8428-k (conf-if-te-0/17)# no shutdown
M8428-k (conf-if-te-0/17)# lldp disable
M8428-k (conf-if-te-0/17)# lacp timeout long
M8428-k (conf-if-te-0/17)# qos flowcontrol tx on rx on
```

Note : Repeat the same steps above for all LACP Ports.

LACP Port-Channel configuration:

Repeat the same steps above for all LACP Ports.

```
M8428-k>en
M8428-k#conf t
M8428-k (conf) #interface Port-channel 2
M8428-k (conf-if-po-2) # mtu 9208
M8428-k (conf-if-po-2) # switchport
M8428-k (conf-if-po-2) # switchport mode trunk
M8428-k (conf-if-po-2) # switchport trunk allowed vlan all
M8428-k (conf-if-po-2) # no shutdown
M8428-k (conf-if-po-2) # end
M8428-k# wr mem
```

Repeat the same steps on the second switch

6 General recommendations and guidelines

- Recommend for use with general purpose small-medium sized consolidated loads.
- Not recommended for purpose built heavy I/O applications.
- Do not directly attach 10GbE arrays to 1GbE IOMs.
- IOMs are just like any other switch and must follow the same configuration recommendations as other Top of Rack (ToR) switches.
- PS-M4110 blade array requires Fabric A or Fabric B (default).
- IOMs must be interconnected.

A. Port Map

This appendix contains the port mapping information for the IOMs used during testing.

During testing of each M-Series IO Module, the blade servers and PS-M4110 blade arrays remained in the same server slots, resulting in an internal port mapping that was similar for all three IOMs. The Force10 MXL has 32 internal network connections and the per server slot connection map is illustrated in Figure 4 below. The PowerConnect M8024-k IO module and the PowerConnect M8428-k IO modules both have 16 internal network connections and the common per server slot connection map is illustrated in Figure 6.

Where each IO module differs is in the inter-switch connection between the two IO modules in the fabric. Since this set of connections uses the external ports available, each IO module’s configuration as tested is provided in Figure 5, 0, and Figure 8.

In all configurations tested, the IO modules illustrated below were installed into the Fabric Slots “B1” and “B2”.

Figure 4 Force10 MXL “Front-end” port map for IOM testing

M620-1P1	P1	P17	
M620-2P1	P2	P18	
	P3	P19	
	P4	P20	
PSM-COP0	P5	P21	
PSM-C1P0	P6	P22	
	P7	P23	
M620-8P1	P8	P24	
	P9	P25	
	P10	P26	
	P11	P27	
	P12	P28	
PSM-COP0	P13	P29	
PSM-C1P0	P14	P30	
	P15	P31	
	P16	P32	

M620-1P2	P1	P17	
M620-2P2	P2	P18	
	P3	P19	
	P4	P20	
PSM-COP1	P5	P21	
PSM-C1P1	P6	P22	
	P7	P23	
M620-8P2	P8	P24	
	P9	P25	
	P10	P26	
	P11	P27	
	P12	P28	
PSM-COP1	P13	P29	
PSM-C1P1	P14	P30	
	P15	P31	
	P16	P32	

			Server Connections
			Array Connections

Figure 5 Force10 MXL External port map




Figure 6 PowerConnect M8024-k and M8428-k common internal port map

M620-1P1	P1	P1	M620-1P2
M620-2P1	P2	P2	M620-2P2
	P3	P3	
	P4	P4	
PSM-C0P0	P5	P5	PSM-C0P1
PSM-C1P0	P6	P6	PSM-C1P1
	P7	P7	
M620-8P1	P8	P8	M620-8P2
	P9	P9	
	P10	P10	
	P11	P11	
	P12	P12	
PSM-C0P0	P13	P13	PSM-C0P1
PSM-C1P0	P14	P14	PSM-C1P1
	P15	P15	
	P16	P16	

Server Connections
 Array Connections

Figure 7 PowerConnect M8024-k External port map





Figure 8 PowerConnect M8428-k External port map



B. Additional resources

- Dell Tech Center EqualLogic Whitepapers:
<http://en.community.dell.com/techcenter/storage/w/wiki/2632.storage-infrastructure-and-solutions-team-publications.aspx>
- Dell EqualLogic Compatibility Matrix:
<http://en.community.dell.com/techcenter/storage/w/wiki/2661.equallogic-compatibility-matrix.aspx>
- Dell EqualLogic Configuration Guide:
<http://en.community.dell.com/techcenter/storage/w/wiki/2639.equallogic-configuration-guide.aspx>
- Dell Rapid EqualLogic Configuration Portal:
<http://en.community.dell.com/techcenter/storage/w/wiki/2632.storage-infrastructure-and-solutions-team-publications.aspx>
- Dell EqualLogic Support Site (Login ID and Password Required):
<https://support.equallogic.com/secure/login.aspx>

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.