Reference Architecture utilizing the PowerEdge M1000e Blade Enclosure, SQL Server and Exchange Server

Dell Engineering
October 2014

A Dell EMC Technical White Paper

Revisions

Date	Revision	Description	Author
October 2014	1.0	Initial release	Ed Blazek

Table of contents

1 Introduction

The small or remote data center of today requires an infrastructure with enterprise class features that support high-end applications, while lowering the total cost of ownership with efficient use of space, cooling and power consumption. Ease of deployment and simplified management that suit the skill-sets and resources available at these data centers are considered essential.

This document presents a consolidated infrastructure solution, based on an M1000e chassis containing Dell blade servers and blade arrays, which is designed to meet the needs of large enterprise branch offices or small and medium enterprise data centers. This solution provides excellent performance, ease of deployment and enhanced data center efficiency with the shared power, cooling, networking and management infrastructure of the PowerEdge M1000e Blade Enclosure.

This solution is designed to provide an infrastructure based on the enterprise class computing Dell PowerEdge M1000e Blade Enclosure, Dell Networking MXL switches, EqualLogic PS-M4110 Blade Arrays and Dell PowerEdge M620 blade servers running Windows 2012 R2, Microsoft Exchange Server 2013 and Microsoft SQL Server 2014. The Microsoft Exchange Server 2013 environment consists of 6000 mailboxes using the Data Availability Group (DAG) feature. The Microsoft SQL Server 2014 environment was tested using the SQLIO Disk Subsystem Benchmark Tool. This complete Data Center solution is contained in an M1000e chassis that utilizes a single management interface. This architecture has been designed and validated by Dell Network Engineering.

The extensive design and engineering work put into this solution allows customers to quickly and confidently deploy this architecture into production environments, helping to eliminate costly and time consuming trial-and-error work often encountered during complex deployments. The information provided in this white paper will aid customers in sizing their solution, planning for appropriate use-cases, and preparing for deployment.

Audience

IT administrators and IT managers who are planning to purchase a small or remote data center can use this document to understand the design elements, components, and the overall architecture of the solution.

2 Overview

This consolidated infrastructure solution combines EqualLogic PS-M4110 Blade Arrays, PowerEdge M-Series servers and Dell Networking MXL switches in a single PowerEdge M1000e blade chassis (Figure 1). This solution helps solve some of the challenges facing IT and data center managers today, which include: increasing costs for power and cooling, rising costs for data center operations, growing costs for the data center footprint and increased complexity of platforms and tools. This solution addresses these challenges in the following ways:

Challenge: Increasing Costs for Power and Cooling
By consolidating server, storage and networking in the M1000e chassis, we are able to make the most of shared power and cooling resources for an efficient data center. Specifically, we are able to take advantage of the design efficiencies of the M1000e including:

- The ability to support more users per watt than comparable solutions.
- Up to six hot-swap ultra-efficient power supplies.
- Nine hot-swap redundant fan modules with dynamic power-efficient fans.
- Optimized airflow design to efficiently cool the enclosure and enable exceptional performance in a low power envelope.
- Next-generation technology to ensure that the absolute minimum amount of fresh air is needed to effectively cool the enclosure.

Challenge: Rising Costs for Data Center Operations and Increasing Cost of the Data Center Footprint
Reduce the cost of data center operations by simplifying deployment, management and maintenance.

Simplify deployment

- By consolidating servers, storage, and network switches in the M1000e, the physical installation of the hardware as well as the configuration (using OpenManage Deployment Toolkit (DTK) and the M1000e Chassis Management Controller (CMC)) requires a fraction of the time compared to setting up the individual components.
- When compared to comparable standalone components this solution drastically reduces the number and complexity of both network and power cabling.
- Other features that simplify deployment include I/O virtualization, chassis LCD display, embedded Hypervisor and an integrated keyboard, video and mouse (KVM) switch.

Simplify Management
The M1000e helps reduce the cost and complexity of managing computing resources with features such as:

- Centralized Chassis Management Controller (CMC) modules provide a redundant central point for infrastructure monitoring, alerting, inventory, and control via a single secure interface.
- Dynamic power management enables you to set high and low power thresholds to help ensure that blades operate efficiently within your power envelope.
- Secure web (SSL) and Telnet/SSH (CLI) interfaces; support for multiple user roles and permissions, including integration into Microsoft Active Directory Service

Simplify Maintenance

- Dell's FlexAddress technology allows the M1000e enclosure to assign the World Wide Name (WWN) or Media Access Control (MAC) address of Fibre Channel, Ethernet and iSCSI controllers to a blade slot instead of directly to the blade. By removing the network and storage identity from the server hardware, customers are now able to upgrade and replace components or even an entire blade server without being forced to change the identity on the network or rezoning switches.
- The M1000e enclosure also includes fully redundant power, cooling and I/O; and hot plug and swappable PSUs, HDDs and fans.
- Other features that simplify maintenance include the iDRAC8 Express or Enterprise with Lifecycle Controller, OpenManage and a dual SD card option for failsafe Hypervisors.

Reducing the Data Center Footprint

By combining servers, storage and network switches in a single M1000e chassis, this solution can help reduce the data center footprint. Specifically, the M620 blade servers in this solution allow increased computing density with more computational power per U when compared to rack servers. Since extra rack space is not required for KVM switches or additional management infrastructure, additional data center real estate is saved.

Challenge: Increased Platform & Tool Complexity

The PowerEdge M1000e modular server enclosure solution supports server modules, network, storage, and cluster interconnect modules (switches and pass-through modules) in a modular, easy-to-use package.

Designed from the ground up to support current and future generations of server, storage, networking, and management technologies, the PowerEdge M1000e includes the headroom necessary to scale for the future.

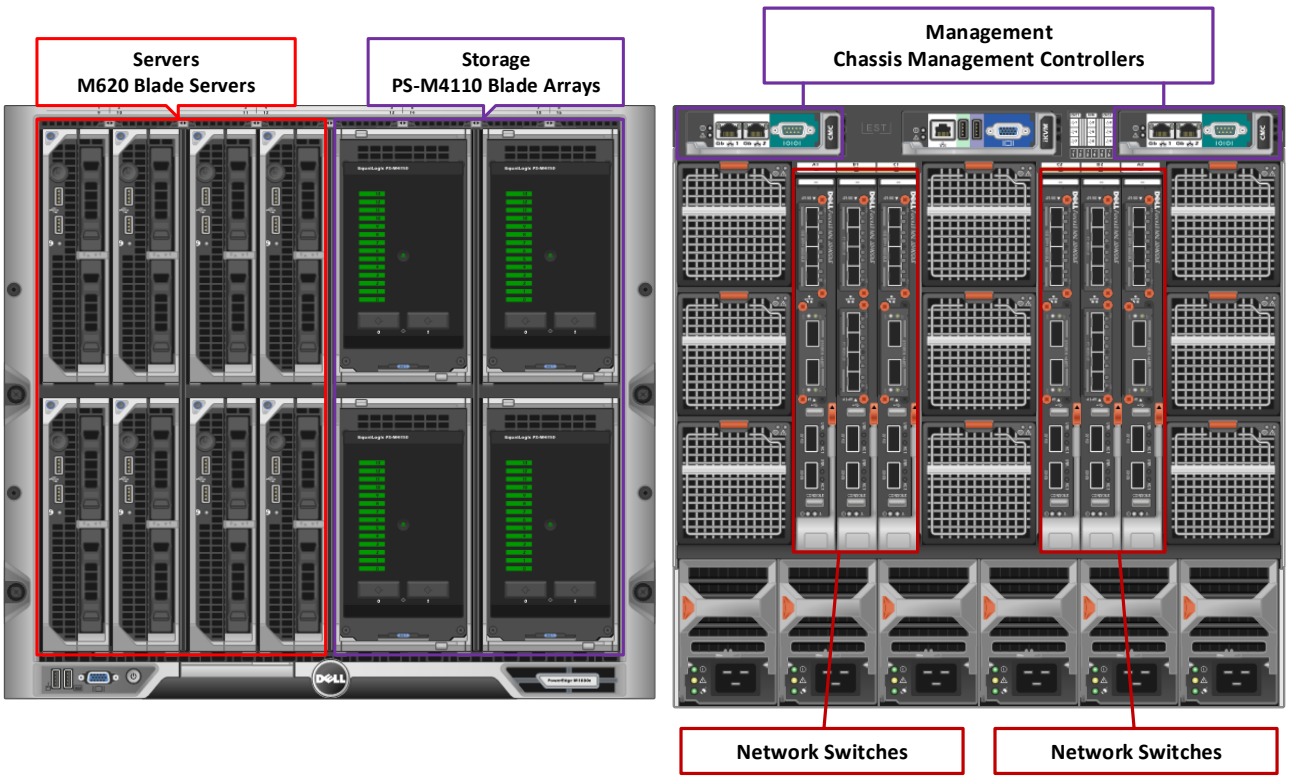

Figure 1 Dell PowerEdge M1000e Blade Enclosure - Front and Rear Views

2.1 Chassis - PowerEdge M1000e Blade Enclosure

The Dell PowerEdge M1000e modular blade enclosure is the rock-solid foundation for Dell EMC'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 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.

2.2 Server – PowerEdge M620 Blade Server

The Dell PowerEdge M620 blade server (Figure 2) is a feature rich, 2-socket blade server, designed for maximum performance with extreme density.

Figure 2 M620 Blade Server

Designed for taxing workloads, such as email, database and virtual environments, the M620 blade server is an ideal blend of density, performance, efficiency and scalability. The M620 delivers unprecedented memory density and superb performance without compromising enterprise-class features.

- Intel Xeon processor E5-2600 and E5-2600 v2 product families. Supporting up to twelve cores per processor.
- Memory
  - Up to 768GB (24 DIMM slots): 2GB/4GB/8GB/16GB/32GB DDR3 up to 1866MT/s.
  - Up to 1.5TB (24 DIMM slots): 64GB DDR3 LRDIMM up to 1600MT/s (with Intel Xeon processor E5-2600 v2 product family only).
- Support for a failsafe hypervisor. Protect against hardware failure and maximize virtualization uptime by running hypervisor on the optional SD card and installing a backup copy on the second mirrored SD card.
- The M620 blade server takes advantage of the shared power, cooling and networking infrastructure of the M1000e blade enclosure coupled with the Dell Chassis Management Controller to manage individual or groups of M620 blade servers.

2.3 Networking Switches – Dell Networking MXL Blade Switch

Improve infrastructure performance

The Dell Networking MXL blade switch helps ensure high-speed support for the growing number of required virtual machines (VMs) per physical server with standards-based 10/40GbE connectivity.

- Replace multiple GbE NICs, mezzanine cards and switches with fewer, higher-bandwidth 10GbE NICs and I/O modules (IOMs) per blade chassis.
- Reduce physical port connectivity requirements between top-of-rack (ToR) and blade servers with local switching, and help improve east-west traffic performance within the M1000e chassis.

Improve system flexibility

Improve data center flexibility and utilization with 10GBase-T, SFP+, and QSFP+ FlexIO modules for the Dell Networking MXL switch.

- Dell Networking MXL includes 32 internal server-facing ports, 2 fixed 40GbE QSFP+ ports and 2 optional expansion slots.
- Optional FlexIO modules enable rapid network upgrades when business demands require higher data center performance (Dell Networking MXL supports up to six external 40GbE ports or 24 external 10GbE SFP+ ports).

Maximize functionality

Virtualize, automate and orchestrate networking functions and services with the reliable and feature-rich Dell Networking Operating System (FTOS).

- Leverages a distributed multiprocessor architecture that offers best-in-class performance, stability and reliability.
- Provides software portability and modularity for high-performance application features.
- Hardware Abstraction Layer (HAL) makes FTOS applications portable across products.
- Includes industry standard Command Line Interface (CLI).
- Interoperable with third-party vendor products.

2.4 Storage - PS-M4110 Blade Array

The PS-M4110 blade array is a doublewide, half-height array that can be installed in any slot in the PowerEdge M1000e chassis (up to four arrays per chassis). The PS-M4110 blade array features:

- 14x 2.5” hot-swappable drives in a drawer-in-drawer blade design, with options for 7.2K NL SAS, 10K or 15K SAS, or a hybrid model combining 10K SAS and Solid State Drives (SSDs).
- A 6Gbps SAS backplane.
- Dual, hot-pluggable 10GbE controllers for link, switch-failure and controller redundancy.
- Up to 2GB of memory per controller for improved performance.
- Cache-to-flash memory de-staging designed for improved data-access reliability.
- Available storage configuration through Fabrics A or B (Fabric B was utilized in this document).

Dedicated Management port
The PS-M4110 provides a dedicated management port for ongoing administration with EqualLogic Group Manager through either an out-of-band management network or an in-band SAN network.

Expand your storage with ease
Save time scaling storage capacity and performance simultaneously and non-disruptively with outstanding EqualLogic peer-scaling architecture.

- Grow system I/O and processor bandwidth automatically as you add capacity.
- Install up to four PS-M4110 arrays per blade chassis.
- Use two PS-M4110 arrays to make up an EqualLogic storage group, and scale up to 16 arrays per group outside the chassis using additional PS6210 family arrays.
- Install up to two EqualLogic groups per chassis, with a maximum of two members each.

Management
Help simplify management and mitigate risk with centralized monitoring and reporting as well as powerful data protection with the all-inclusive enhanced software features of the EqualLogic PS-M4110 Blade Array.

Note: Drive capacities are subject to change with future releases.

Storage Models

- **PS-M4110E**
  - High capacity at an affordable price.
  - 14 x 2.5", 7.2K RPM NL-SAS or NL-SAS SED HDDs 1TB per drive.
  - Up to 14TB per array, up to 28TB per group inside the blade chassis.
- **PS-M4110X**
  - Balance of capacity and performance.
    14 x 2.5", 10K RPM SAS or SAS SED (900GB drives only) HDDs
    600GB or 900GB or 1.2TB per drive.
  - Up to 16.8TB per array, up to 33.6TB per group inside the blade chassis.
- **PS-M4110XV**
  - High-performance option for business-critical applications.
  - 14 x 2.5", 15K RPM SAS or SAS SED HDDs 300GB per drive.
  - Up to 4.2TB per array, up to 8.4TB per group inside the blade chassis.
- **PS-M4110XS**
  - Hybrid configuration with SSD performance and HDD capacity.
  - 9 x 2.5", 10K RPM SAS HDDs and 5 x 2.5" SSDs
    600GB and 1.2TB HDD and 400GB and 800GB SSD.
  - Up to 14.8TB per array, up to 29.6TB per group inside the blade chassis.

Note: This white paper utilized two PS-M4110XV Storage Models for SQL and two PS-M4110E storage models for Exchange

3 Software

This section includes the software used for the end solution. Other software (Jetstress and SQLIO) were used to validate the solution and are discussed in the appropriate sections.

3.1 Microsoft Windows Server 2012 R2

Windows Server 2012 R2, the latest release of the server OS from Microsoft, includes many new features and enhancements. Windows Server 2012 R2 focuses on four core areas:

Beyond virtualization: Windows Server 2012 R2 provides a robust and dynamic virtualization platform through Hyper-V technologies, and includes features for cloud services.

The power of many servers, the simplicity of one: Windows Server 2012 R2 offers features that allow for high availability and ease of management for multiple-server infrastructures.

Every app, any cloud: Windows Server 2012 R2 delivers a scalable and flexible web and application platform by providing a consistent and open set of tools and frameworks that apply to applications on-premises, in the cloud, or in a hybrid environment.

Modern work style, enabled: Microsoft Windows Server 2012 R2 empowers users and IT staff with remote access to data, applications, and simpler management tools while strengthening security and compliance.

3.2 Microsoft Exchange Server 2013

Exchange Server 2013 equips you with a robust messaging platform that can give your users access anywhere at any time. Exchange Server 2013 builds upon the Exchange Server 2010 architecture, and has been redesigned for simplicity of scale, hardware utilization, and failure isolation. Some of the new features in Exchange Server 2013 include:

- Smart Search, which is a new search feature that learns from users' communication and collaboration behavior to enhance and prioritize search results in Exchange. Users can also merge contacts from multiple sources to provide a single view of a person.
- Improved search and indexing. Exchange 2013 helps you to find and search data not only in Exchange, but also across your organization. In addition to Exchange 2013, you can also search across Lync 2013, SharePoint 2013, and Windows file servers. This not only increases user's productivity, but also helps with compliance and eDiscovery.
- Data Loss Prevention (DLP) helps keep your organization safe from users mistakenly sending sensitive information to unauthorized people. DLP also helps you identify, monitor, and protect sensitive data through deep content analysis.
- Greater integration with Microsoft SharePoint 2013 and Microsoft Lync 2013 through site mailboxes and in-place eDiscovery. Together, these products offer a suite of features that make scenarios such as enterprise eDiscovery and collaboration using site mailboxes possible.

For more information about Microsoft Exchange Server 2013, visit, http://www.office.microsoft.com/en-us/exchange/.

3.3 Microsoft SQL Server 2014

SQL Server 2014 enables customers to build mission-critical applications and big data solutions using high-performance, in-memory technology across Online Transaction Processing (OLTP), data warehousing, Business Intelligence (BI) and analytic workloads without having to buy expensive add-ons or high-end appliances. SQL Server 2014 uses a common set of tools to deploy and manage databases both on-premises and in the cloud, making it easier for customers to take advantage of the cloud with existing skills.

Benefits

SQL Server 2014 makes it easier and more cost effective to build high-performance mission-critical applications, enterprise ready big data assets, and BI solutions that help employees make better decisions, faster. These solutions have the flexibility of being deployed on premises, in the cloud or in a hybrid environment, and can be managed through a common and familiar tool set.

Mission-critical performance

SQL Server 2014 accelerates reliable, mission critical applications with a new in-memory OLTP engine that can deliver on average 10x, and up to 30x transactional performance gains. For data warehousing, the new updatable in-memory column store can query 100x faster than legacy solutions.

Faster insights on any data

Get to insights faster with a complete BI platform that speeds up how you access, analyze, clean and shape both internal and external data. With SQL Server 2014 and Power BI for Office 365, it is easy to connect every user in an organization to the right data they need to make better decisions, faster.

3.4 Multi-Path I/O

There are generally two types of multi-path access for communicating from a host to an external device. For general networking communications, the preferred method of providing redundant connections is teaming multiple NICs into a single, virtual network connection entity. For storage, the preferred method is the use of Multi-Path I/O (MPIO). Though some storage solutions support either method for iSCSI connectivity, EqualLogic requires the use of MPIO to enable multiple NIC/HBA connections to an EqualLogic SAN.

The MPIO Toolkits can be downloaded from: https://eqlsupport.dell.com/support/download.aspx.

3.4.1 EqualLogic MPIO Requirements

The following host port requirements are required to use MPIO with EqualLogic SANs:

- At least two Ethernet ports are required on each host.
- The host operating system must have a supported MPIO driver or service available.
- The ports used for MPIO cannot be teamed with other ports.
- The ports must be the same speed.
- The ports must have assigned IP addresses on the same subnet.

3.4.2 EqualLogic MPIO General Recommendations

Follow this general set of guidelines for configuring MPIO on a host:

- Configure volume access controls to use standard iSCSI Qualified Name (IQN).
- For a more secure configuration, use the IQN name plus the CHAP security ID.
- Install the Dell EMC-provided MPIO extension features (if available for the host operating system).
- For Microsoft Windows, install the Device Specific Module (DSM) found in the Host Integration Toolkit for Windows.

4 Exchange Server 2013

This section describes a tested and validated solution for a 6,000-mailbox Exchange 2013 environment with Data Availability Group (DAG). DAG is a high availability mechanism in Microsoft Exchange. This model of mailbox resiliency supports multiple copies (up to 16) of the Exchange database in a DAG. There is only one active copy of a given Exchange 2013 database at any given time. Secondary copies are periodically synced with the primary copy. E-Mail clients access the primary (active) copy, and database changes to the primary copy are copied to the secondary (passive) copies in the form of transaction logs. The copied log records are played on the secondary copy to keep the secondary database copies consistent with the primary copy. All hosts within a DAG are configured to be identical in terms of CPU and memory as well as storage resources for Exchange 2013 databases and logs. The primary and secondary copies do not share any storage resources and reside on their own dedicated storage resources as discussed later.

In this configuration, each of the two Dell PowerEdge M620 servers (Figure 3) hosts four active databases with 750 users per database, or 3000 users per server (6000 active users) with 1 GB mailbox each. The tested mailbox user profile was 150 messages a day or 0.101 IOPS (Input Output per Second) per user. Figure 3 shows the servers and storage used for the Exchange environment.

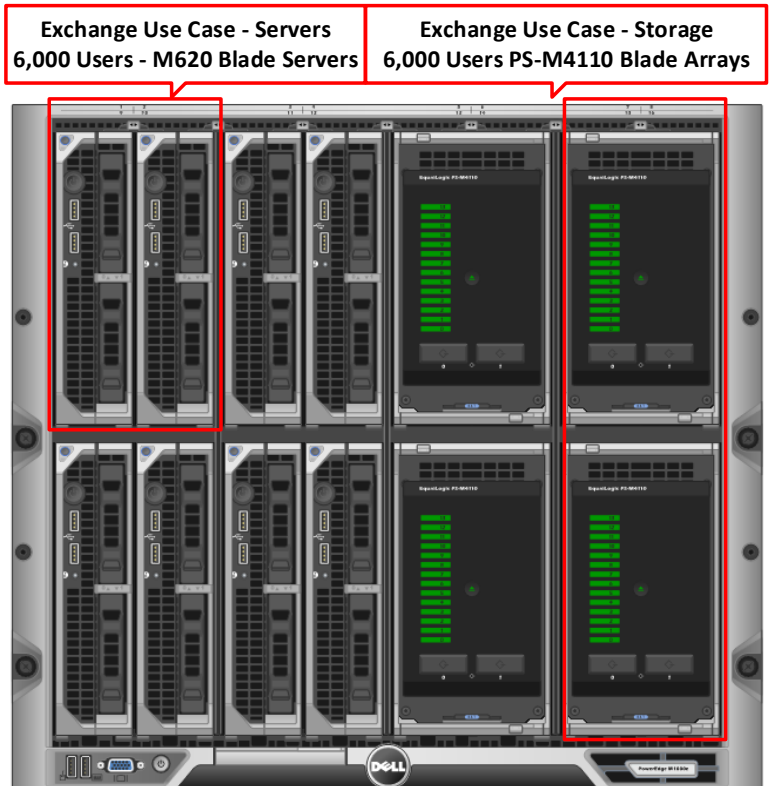

Figure 3 Exchange Servers and Storage

4.1 Targeted Customer Profile

- 6,000 concurrent users, 3,000 users from each of the two servers.
- Four Databases/four volumes per server (750 users per database).
- Workload of 150 messages per day per mailbox and a 1 GB mailbox per user.
- Two PS-M4110E (1TB 7.2k) in a group with two pools segregating the two sets of volumes (RAID-6).
- Windows Server 2012 R2 with EqualLogic ASM/ME 4.7.

4.2 Tested Deployment Utilizing Jetstress

The tested deployment simulated a failure scenario where one of the Exchange Server 2013 DAG members is offline and the passive copies on the surviving DAG member are activated to provide mailbox service continuity. Therefore, the IOPS simulated mimicked that of 6000 mailboxes on the same Exchange 2013 Server. The target IOPS for this profile was 605. The achieved IOPS were 645, which exceeded the target while maintaining read and write latencies well within the recommended thresholds. Table 1 provides a summary of the testing environment.

Table 1 Simulated Exchange Configuration

Feature	Specification
Number of Exchange mailboxes simulated	6,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	2
Number of active mailboxes/server	3,000
Number of databases/server	4
Number of copies/database	2
Number of mailboxes/database	750
Simulated profile: IOPS/mailbox	.101 (150 messages / day)
Database/Log size	1024 GB
Total database size for performance testing	4.8 TB Total and 1.2 TB per DB

4.3 Test Information

The test runs in this document were ran for 8 hours. The goal was to verify that the storage could handle high I/O load for 8 hours. Both log and database files were analyzed for integrity after the stress test to ensure no database/log corruption occurred. There were no errors reported in the event log file or during the database and log checksum process.

The Microsoft Exchange Jetstress 2013 Performance Test Result Reports for Server 1 and Server 2 are attached to this document.

4.4 Best Practices

The following best practices should be kept in mind when configuring Exchange:

- Exchange server is a disk I/O intensive application. The disk subsystem should be capable of supporting both capacity and I/O throughput demands of the application.
- Sharing Exchange 2013 storage resources with other applications may negatively affect the performance of Exchange 2013 deployment and therefore is not recommended.

  Note: During testing, the database and log folders shared the same physical disk. Other testing indicated that separating the database folders from log folders on two different set of disks does not provide a noticeable performance advantage.

- In an Exchange Server 2013 resiliency solution, separating the database and log folders is no longer a required best practice.
- Windows NTFS allocation unit size for Exchange 2013 database partitions should be set to 64K for best performance.
- Exchange Server 2013 storage latencies are most often related to the number of disks available for given a workload. Windows Performance Monitor can be used to monitor Exchange Server 2013 database counters. Average database read latencies (Avg. Disk sec/Read) should not exceed 20ms.

For Exchange 2013 best practices on storage design follow the link: http://technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx.

4.5 Exchange Conclusion

The storage performance testing is designed to exercise the storage with the maximum sustainable Exchange type of I/O for two hours. The tests show how long it takes the storage to respond to an I/O under load. The target for the stress test was 605 IOPS; the achieved results were 645 IOPS, which exceeded the target.

The test results/data presented in this document are based on a lab configuration. Customers should not quote the data directly for pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

The tests ran were not to benchmark a design or to get the maximum throughput for a given solution. Rather, they focus on producing recommendations from customers utilizing Exchange. So, the data presented in this document should not be used for direct comparisons among solutions.

5 Microsoft SQL Server

This section describes validating the consolidated solution for SQL Server. Since the I/O system performance is so important to the performance of SQL Server, SQLIO was used to simulate the I/O activity of a typical SQL deployment. SQLIO is available for download from the Microsoft Download Center. Figure 4 shows the servers and storage used for the SQL Server environment and Table 2 provides a summary of the configuration of the testing environment.

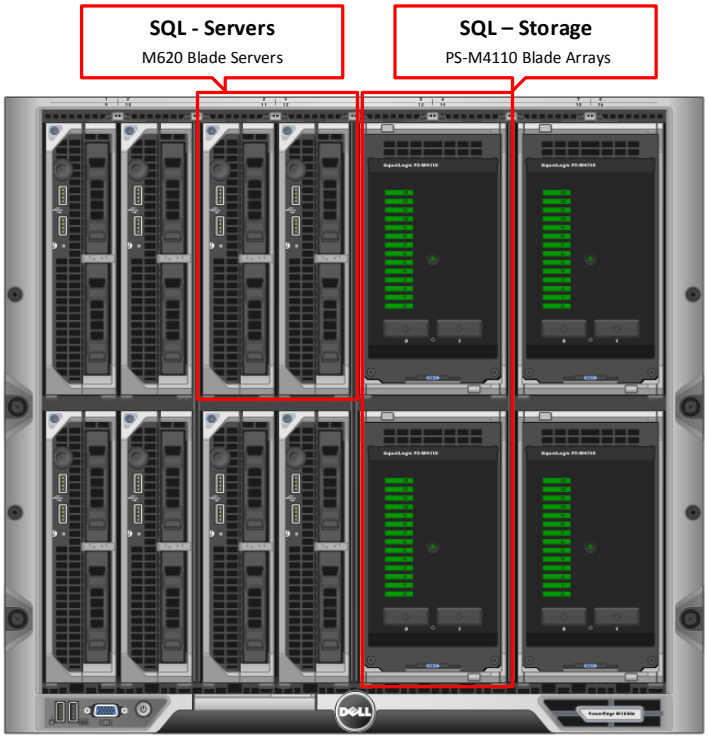

Figure 4 SQL Servers and Storage

Table 2 Simulated SQL Configuration

Description	Specification
Pool Name	Group 5
Pool Members	PS-M4110XV-Slot5 PS-M4110XV-Slot13
Raid Policy	Raid 6
Disks per enclosure / Model Number	14 / 15K-SAS-ST9146853SS
Number of Volumes on Server	5 Volumes (4 volumes at 150 GB and 1 Volume at 100 GB
Server Drive / Size of databases	D:\152576 MB E:\152576 MB F:\152576 MB G:\152576 MB H:\101376 MB

5.1 Test Results Overview

Table 3 presents an overview of the test results. The I/O was tested with both random and sequential read/writes with 8K and 256K block sizes. The 8K files represent small blocks (transaction logs) while the 256K blocks represent the largest I/O block a non-Enterprise Edition of SQL will issue. Terms used in the test results include:

- **Sequential Read/Write Test** - Contiguous data is written and read to and from adjacent locations on the storage device (i.e. blocks are read or written one after another).
- **Random Read/Write Test** - Data is written and read to and from random locations on the storage device.
- **IOPs (Input Output per Second**) - Number of operations performed per second.
- **MB/s** - Throughput (Amount of data moved around per second).

Table 3 Test Results Overview

PS-M4110XV x2			
Random 8K			
	Reads	9973.72 IOPS	17.70 ms Latency
	Writes	3073.25 IOPS	13.64ms Latency
Sequential 256K			
	Reads	552.12 MB/s (Avg.)	
	Writes	308.78 MB/s (Avg.)	

5.2 Test Results Detail

The following sections present detailed information for each test, including graphs detailing the reads and writes.

5.2.1 8K Random Reads and Writes

The two graphs in Figure 5 show the results of the 8k random read and write test.

8K Random Reads	8K Random Writes
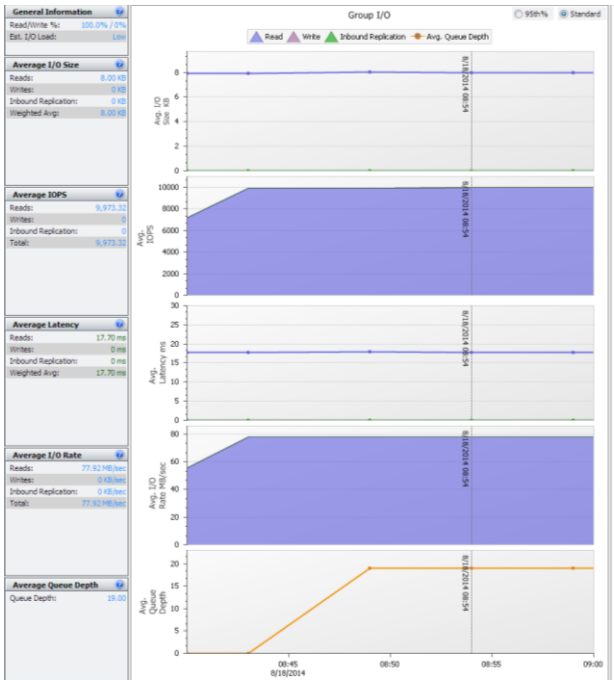	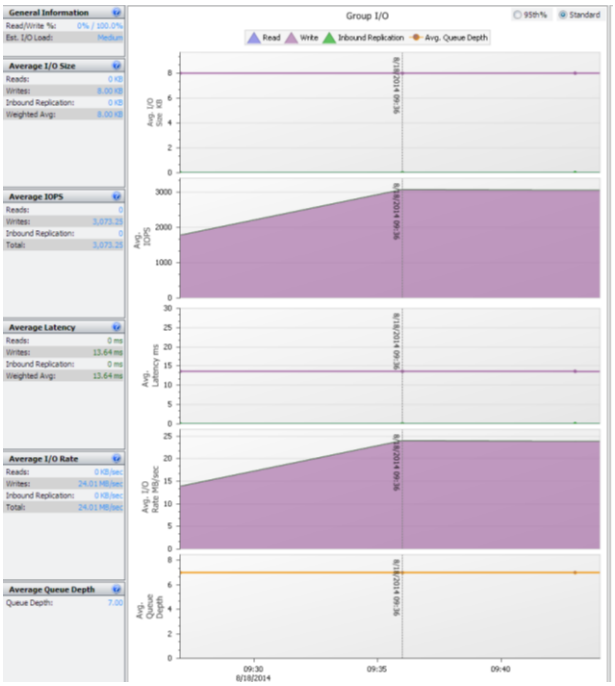

Figure 5 8K Random Read and Write Test Results

5.2.2 256K Random Reads and Writes

The two graphs in Figure 6 show the results of the 256K random read and write test.

256K Random Reads	256K Random Writes
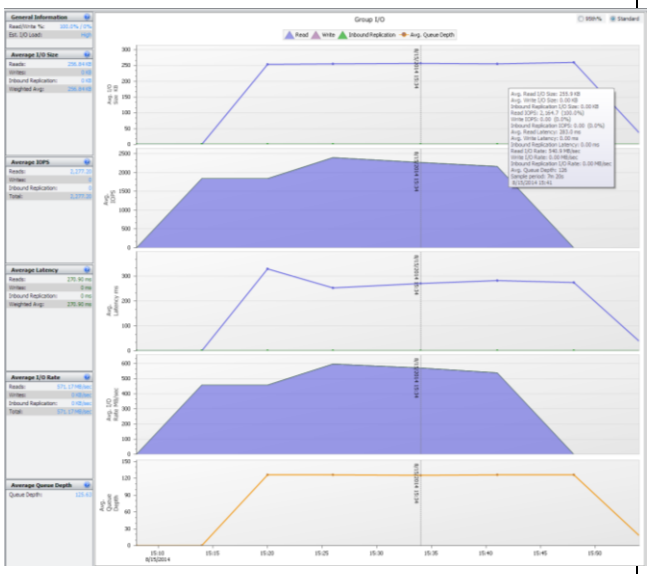	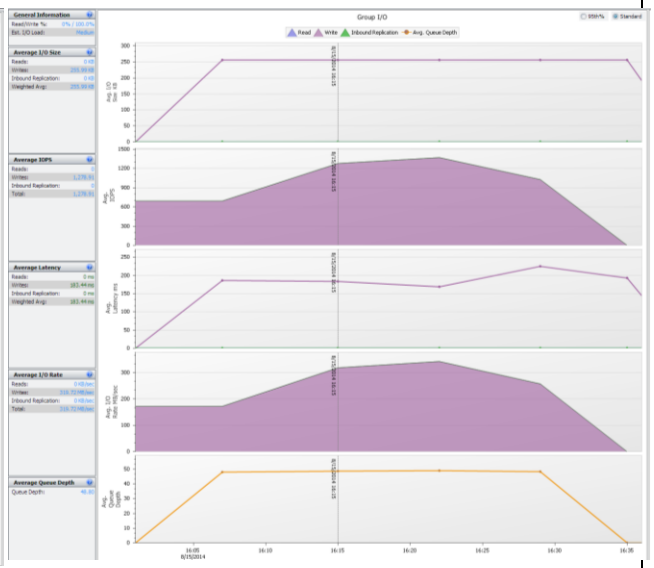

Figure 6 256K Random Read and Write Test Results

5.3 Microsoft SQL Server 2014 Conclusion

Dell EMC, in partnership with Microsoft, enables customers to enhance the Return of Investment on data warehouse systems. The end-to-end database best practices and recommendations enable the customer to achieve a balanced data warehouse environment with greater performance benefits than traditional data warehouse systems.

The Dell EMC-Microsoft Architecture provides the following benefits to customers:

- Delivers a tested and validated configuration with proven methodology and performance behavior.
- Achieves a balanced and optimized system at all the levels of the stack by following the best practices of hardware and software components.
- Avoids over-provisioning of hardware resources.
- Offers high availability at all the levels of setup (host, switches, and storage).
- Helps customers avoid the pitfalls of an improperly designed and configured system.
- Reduces future support costs by limiting solution re-architect efforts because of scalability challenges.

6 Networking Configuration

This section describes the networking configuration for end users (Fabric A) and for storage (Fabric B). Figure 7 shows the location of each of the switches in the M1000e chassis.

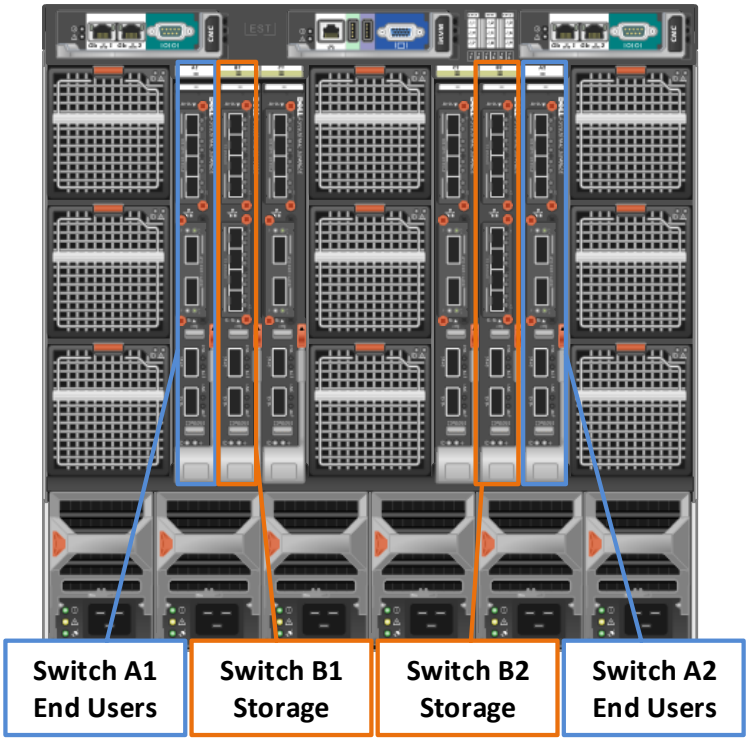

Figure 7 Dell PowerEdge M1000e Blade Enclosure - Rear View

6.1 Physical Cable Connections

Table 4 and Table 5 present the physical connections for the end users network (Fabric A) and the storage network (Fabric B) respectively.

Table 4 Physical Cable Connections for the End User Network

MXL A1 and A2	
From Switch / Port	To Switch / Port
MXL_A1 / fortyGigE 0/33	S4810_STACK2-0 / fortyGigE 0/48
MXL_A1 / fortyGigE 0/37	S4810_STACK1-0 / fortyGigE 0/48
MXL_A1 / fortyGigE 0/41	MXL_A2 / fortyGigE 0/41
MXL_A1 / fortyGigE 0/45	MXL_A2 / fortyGigE 0/45
MXL_A2 / fortyGigE 0/33	S4810_STACK2-1 / fortyGigE 0/48
MXL_A2 / fortyGigE 0/37	S4810_STACK1-1 / fortyGigE 0/48
S4810 (Stacking)	
From Switch / Port	To Switch / Port
S4810_STACK1-0 / fortyGigE 0/52	S4810_STACK1-1 / fortyGigE 0/60
S4810_STACK1-0 / fortyGigE 0/60	S4810_STACK1-1 / fortyGigE 0/52
S4810_STACK2-0 / fortyGigE 0/52	S4810_STACK2-1 / fortyGigE 0/60
S4810_STACK2-0 / fortyGigE 0/60	S4810_STACK2-1 / fortyGigE 0/52

Table 5 Physical Cable Connections for the Storage Network (Fabric B)

MXL B1 and B2	
MXL_B1 / fortyGigE 0/33	MXL_B2 / fortyGigE 0/33
MXL_B1 / fortyGigE 0/37	MXL_B2 / fortyGigE 0/37

6.2 Network Fabric A1 and A2 Configuration (End User)

The two networking switches in Fabrics A1 and A2 (Figure 8) are used for end user traffic. Two cables are required to LAG these switches together in order to create a Virtual Link Trunk (VLT) between Fabric A1 and A2. Fabric A1 and A2 are then connected to the S4810 stacks using a port-channel (consisting of two 40GbE ports) from each Fabric as shown in (Figure 9).

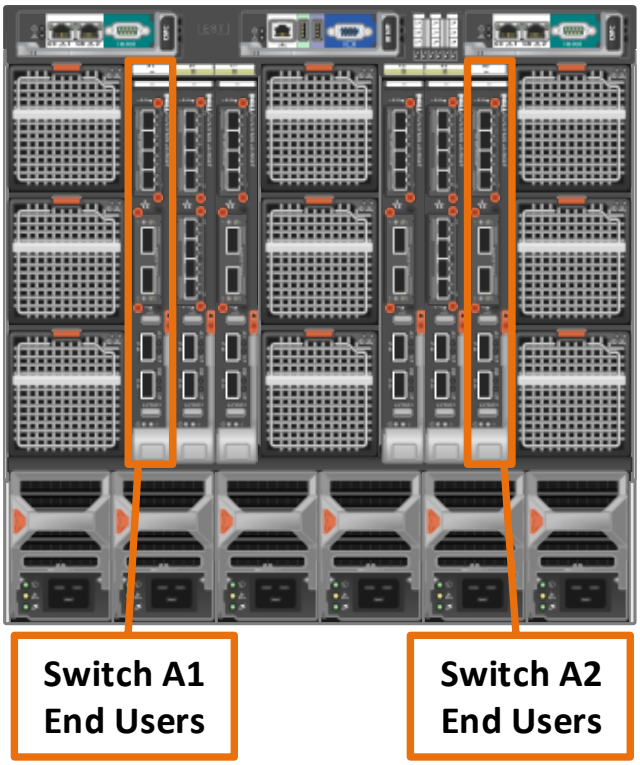

Figure 8 Dell PowerEdge M1000e Blade Enclosure

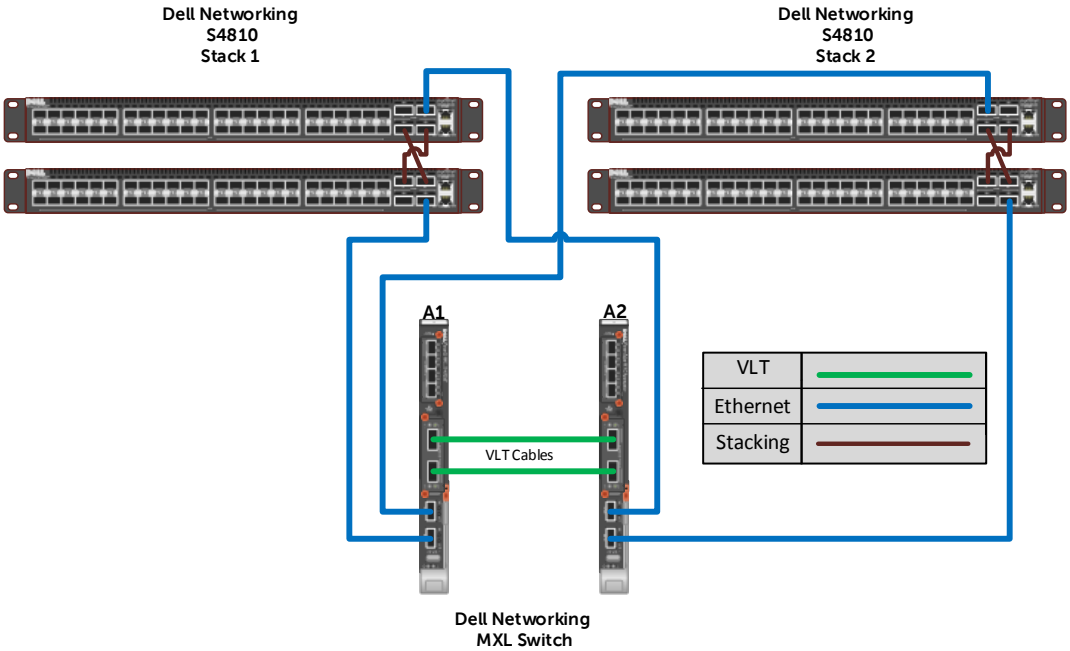

Figure 9 Dell Networking MXL Switch's A Fabric Cables

6.2.1 Configure the First Switch (A1)

Perform the following steps to configure the first switch (A1).

1. Login to the Dell Networking MXL 10/40GbE Blade Switch.
   Perform the following steps to log in to the blade switch.
   a. Power on the chassis.
   b. Login to the chassis CMC.
   c. From the CMC CLI, use the following command to connect to the Dell Networking MXL 10/40GbE Blade Switch.

```
$connect switch-n
```

Note: Where *n* is the number of the switch (1 through 6). For example, the switch in the B1 slot is number 3, and the switch in the B2 slot is number 4,

2. Delete the startup configuration.
   Use the following commands to delete the startup configuration.

```
MXL_1>enable
MXL_1#delete startup-config
Proceed to delete startup-config [confirm yes/no] yes
MXL_1#reload
System configuration has been modified. Save? [yes/no] no
Proceed with reload [confirm yes/no] yes
```

Note: The switch will reboot.

3. Configure the Hostname and Out Of Band (OOB) management port.
   Use the following commands to configure the host name and OOB management port.

```
FTOS>enable
FTOS#configure
FTOS(conf)#hostname MXL_A1
MXL_A1(conf)#interface ManagementEthernet 0/0
MXL_A1(conf-if-ma-0/0)#no shutdown
MXL_A1(conf-if-ma-0/0)#ip address ipaddress mask
Proceed with Static IP [confirm yes/no]: yes
MXL_A1(conf-if-ma-0/0)#management route W.X.Y.Z /24 A.B.C.D
MXL_A1(conf-if-ma-0/0)#exit
```

Note: *W.X.Y.Z* is the network your management system is connecting from, and *A.B.C.D* is the gateway for the switch. If your management system is on the same subnet as the switch, this step may be omitted. The example above assumes a class C subnet mask.

4. Configure login credentials.
   Use the following commands to configure login credentials.

```
MXL_A1(conf)#username admin privilege 15 password 0 yourpassword
MXL_A1(conf)#enable password level 15 0 yourpassword
MXL_A1(conf)#exit
```

5. Enable the switch ports.
   There are two ways to enable the switch ports. **Option 1** enables the ports individually and **Option 2** enables multiple ports using the *range* parameter.

   Option 1
   Use the following commands to enable ports individually by entering the port number.

```
MXL_A1#configure
MXL_A1(conf)#interface tengigabitethernet 0/1
MXL_A1(conf-if-te-0/1)#switchport
MXL_A1(conf-if-te-0/1)#no shutdown
MXL_A1(conf-if-te-0/1)#exit
MXL_A1(conf)#exit
```

   Option 2
   Use the following commands to enable multiple ports at once using the *range* parameter. *The following example assumes the base MXL switch with one 2-port 40GB QSFP+ IO module installed. Actual configuration may vary depending on the modules installed.*

```
MXL_A1#configure
MXL_A1(conf)#interface range tengigabitethernet 0/1 – 32 ,
tengigabitethernet 0/49 – 56 , fortyGigE 0/33 - 45
% Warning: Non-existing ports (not configured) are ignored by interface-
range
MXL_A1(conf-if-range-te-0/1-32,te-0/49-56,fo-0/3...)#switchport
MXL_A1(conf-if-range-te-0/1-32,te-0/49-56,fo-0/3...)#no shutdown
MXL_A1(conf-if-range-te-0/1-32,te-0/49-56,fo-0/3...)#end
```

6. Configure spanning tree.
   Use the following commands to configure spanning tree.

```
MXL_A1#configure
MXL_A1(conf)#protocol spanning-tree rstp
MXL_A1(conf-rstp)#no disable
MXL_A1(conf-rstp)#bridge-priority 4096
MXL_A1(conf-rstp)#exit
MXL_A1(conf)#interface range tengigabitethernet 0/1 – 32
MXL_A1(conf-if-range-te-0/1-32)#spanning-tree rstp edge-port
MXL_A1(conf-if-range-te-0/1-32)#end
```

7. Configure port channels for the Link Aggregation Groups (LAGs).
   Use the following commands to configure the switch interconnects as a LAG.

```
MXL_A1#configure
MXL_A1(conf)#interface port-channel 33
MXL_A1(conf-if-po-33)#exit
MXL_A1(conf)#interface port-channel 37
MXL_A1(conf-if-po-37)#exit
MXL_A1(conf)#interface port-channel 100
MXL_A1(conf-if-po-100)#description 40gig VLT interlink
MXL_A1(conf-if-po-100)#no ip address
MXL_A1(conf-if-po-100)#channel-member fortyGigE 0/41,45
MXL_A1(conf-if-po-100)#no shutdown
MXL_A1(conf-if-po-100)#exit
MXL_A1(conf)#interface range port-channel 33 , port-channel 37
MXL_A1(conf-if-range-po-33,po-37)#no ip address
MXL_A1(conf-if-range-po-33,po-37)#switchport
MXL_A1(conf-if-range-po-33,po-37)#no shutdown
MXL_A1(conf-if-range-po-33,po-37)#exit
```

8. Configure QSFP ports for LAG.
   Use the following commands to configure and assign the 40 GB QSFP ports to the port channel.

```
MXL_A1(conf)#interface fortyGigE 0/33
MXL_A1(conf-if-fo-0/33)#no ip address
MXL_A1(conf-if-fo-0/33)#no switchport
MXL_A1(conf-if-fo-0/33)#port-channel protocol lacp
MXL_A1(conf-if-fo-0/33-lacp)#port-channel 33 mode active
MXL_A1(conf-if-fo-0/33-lacp)#exit
MXL_A1(conf-if-fo-0/33)#protocol lldp
MXL_A1(conf-if-fo-0/33-lldp)#no disable
MXL_A1(conf-if-fo-0/33-lldp)#advertise management-tlv management-address system-
name
MXL_A1(conf-if-fo-0/33-lldp)#exit
MXL_A1(conf-if-fo-0/33)#no shutdown
MXL_A1(conf-if-fo-0/33)#exit
MXL_A1(conf)#interface fortyGigE 0/37
MXL_A1(conf-if-fo-0/37)#no ip address
MXL_A1(conf-if-fo-0/37)#no switchport
MXL_A1(conf-if-fo-0/37)#port-channel protocol lacp
MXL_A1(conf-if-fo-0/37-lacp)#port-channel 37 mode active
MXL_A1(conf-if-fo-0/37-lacp)#exit
MXL_A1(conf-if-fo-0/37)#protocol lldp
MXL_A1(conf-if-fo-0/37-lldp)#no disable
MXL_A1(conf-if-fo-0/37-lldp)#advertise management-tlv management-address system-
name
MXL_A1(conf-if-fo-0/37-lldp)#exit
MXL_A1(conf-if-fo-0/37)#no shutdown
MXL_A1(conf-if-fo-0/37)#exit
MXL_A1(conf)#interface range fortyGigE 0/41 , fortyGigE 0/45
MXL_A1(conf-if-range-fo-0/41,fo-0/45)#no ip address
MXL_A1(conf-if-range-fo-0/41,fo-0/45)#no switchport
MXL_A1(conf-if-range-fo-0/41,fo-0/45)#protocol lldp
MXL_A1(conf-if-range-fo-0/41,fo-0/45-lldp)#no disable
MXL_A1(conf-if-range-fo-0/41,fo-0/45-lldp)#advertise management-tlv management-
address system-name
MXL_A1(conf-if-range-fo-0/41,fo-0/45-lldp)#exit
MXL_A1(conf-if-range-fo-0/41,fo-0/45)#no shutdown
MXL_A1(conf-if-range-fo-0/41,fo-0/45)#exit
```

9. Configure VLT.
   Use the following commands to configure VLT.

```
MXL_A1(conf)#vlt domain 100
MXL_A1(conf-vlt-domain)#peer-link port-channel 100
MXL_A1(conf-vlt-domain)#back-up destination 172.25.188.42
MXL_A1(conf-vlt-domain)#primary priority 1
MXL_A1(conf-vlt-domain)#exit
MXL_A1(conf)#interface port-channel 33
MXL_A1(conf-if-po-33)#vlt-peer-lag port-channel 33
MXL_A1(conf-if-po-33)#exit
MXL_A1(conf)#interface port-channel 37
MXL_A1(conf-if-po-37)#vlt-peer-lag port-channel 37
MXL_A1(conf-if-po-37)#end
```

10. Save the configuration.
    Use the following command to save the configuration.

```
MXL_A1#copy running-config startup-config
```

6.2.2 Configure the Second Switch (A2)

Repeat the commands from the section above to configure the second switch, with the following changes:

- Configure spanning tree RSTP with bridge-priority 8192.
- Configure VLT domain 100 with back-up destination of 172.25.188.41 (management IP of MXL_A1).
- Use switch VLT *default* priority (Do not configure VLT priority).

Note: The preceding procedure places all switch ports in the default VLAN. If you prefer to place ports in a non-default VLAN, refer to the documentation for the switch.

6.3 Network Fabric B1 and B2 Configuration (Storage)

The two networking switches in Fabrics B1 and B2 are used for storage. Two cables are required to LAG these two switches together (Figure 10). By using FC Flex I/O modules, this configuration allows additional external storage to be added in the future.

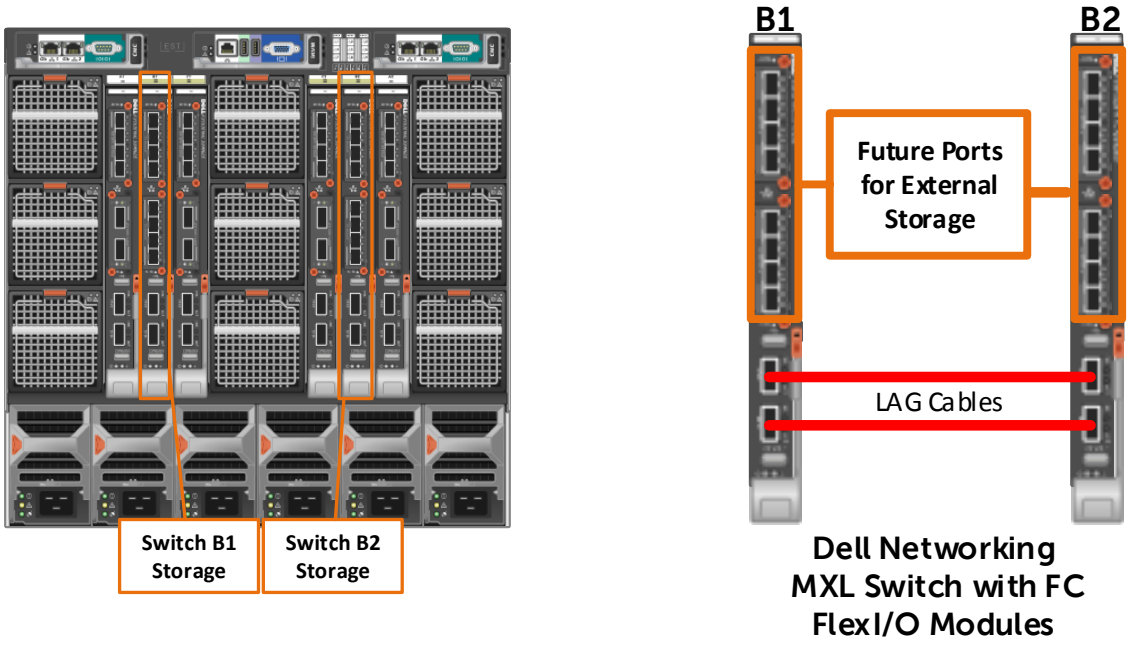

Figure 10 Dell Networking MXL Switch's B Fabric Cables

6.3.1 Configure the First Switch (B1)

Perform the following steps to configure the first switch (B1).

1. Login to the Dell Networking MXL 10/40GbE Blade Switch.
   Perform the following steps to log in to the blade switch.

a. Power on the chassis.

b. Login to the chassis CMC.
   From the CMC CLI, connect to the Dell Networking MXL 10/40GbE Blade Switch using the following command:

```
$connect switch-n
```

Note: Where n is the number of the switch (1 through 6). For example, the switch in the B1 slot is number 3, and the switch in the B3 slot is number 4.

11. Delete the startup configuration.
    Use the following commands to delete the startup configuration.

```
MXL_B1>enable
MXL_1#delete startup-config
Proceed to delete startup-config [confirm yes/no] yes
MXL_B1#reload
System configuration has been modified. Save? [yes/no] no
Proceed with reload [confirm yes/no] yes
```

Note: The switch will reboot.

12. Disable iSCSI optimization.

Note: As a best practice, Dell EMC recommends disabling iSCSI session monitoring for larger SAN deployments. ISCSI session monitoring is enabled by default as part of iSCSI optimization. These steps show how to disable iSCSI optimization and manually configure the switches for use with EqualLogic iSCSI storage.

Use the following commands to disable iSCSI optimization.

```
FTOS# configure
FTOS(conf)# no iscsi enable
FTOS(conf)# exit
```

Note: The switch will reboot after the preceding commands.

13. Configure the Hostname and the OOB Management Port.
    Use the following commands to configure the host name and the OOB management port.

```
FTOS>enable
FTOS#configure
FTOS(conf)#hostname MXL_B1
MXL_B1(conf)#interface ManagementEthernet 0/0
MXL_B1(conf-if-ma-0/0)#no shutdown
MXL_B1(conf-if-ma-0/0)#ip address ipaddress mask
Proceed with Static IP [confirm yes/no]: yes
MXL_B1(conf-if-ma-0/0)#management route W.X.Y.Z /24 A.B.C.D
MXL_B1(conf-if-ma-0/0)#exit
```

Note: *W.X.Y.Z* is the network your management system is connecting from and *A.B.C.D* is the gateway for the switch. If your management system is on the same subnet as the switch, this step may be omitted. The example above assumes a class C subnet mask.

14. Configure login credentials.
    Use the following commands to configure login credentials.

```
MXL_B1(conf)#username admin privilege 15 password 0 yourpassword
MXL_B1(conf)#enable password level 15 0 yourpassword
MXL_B1(conf)#exit
```

15. Enable switch ports.
    There are two ways to enable the switch ports. **Option 1** enables the ports individually and **Option 2** enables multiple ports using the *range* parameter.

    Option 1
    Use the following commands to enable ports individually by entering the port number.

```
MXL_B1#configure
MXL_B1(conf)#interface tengigabitethernet 0/1
MXL_B1(conf-if-te-0/1)#switchport
MXL_B1(conf-if-te-0/1)#no shutdown
MXL_B1(conf-if-te-0/1)#exit
MXL_B1(conf)#exit
```

    Option 2
    Use the following commands to enable multiple ports at once using the *range* parameter. The following example assumes the base MXL switch with two optional 4-port SFP+ Flex IO modules installed. Actual configuration may vary depending on the modules installed.

```
MXL_B1#configure
MXL_B1(conf)#interface range tengigabitethernet 0/1 – 32 ,
tengigabitethernet 0/41 - 56
MXL_B1(conf-if-range-te-0/1-32,te-0/41-56)#switchport
MXL_B1(conf-if-range-te-0/1-32,te-0/41-56)#no shutdown
MXL_B1(conf-if-range-te-0/1-32,te-0/41-56)#exit
```

16. Enable jumbo frames.
    Use the following commands to enable jumbo frames.

```
MXL_B1# configure
MXL_B1(conf)# interface range tengigabitethernet 0/1 – 32 ,
tengigabitethernet 0/41 - 56
MXL_B1(conf-if-range-te-0/1-32,te-0/41-56)#mtu 12000
```

17. Configure flow control.
    Use the following command to configure flow control.

```
MXL_B1(conf-if-range-te-0/1-32,te-0/41-56)#flowcontrol rx on tx off
```

18. Configure spanning tree on edge ports.
    Use the following commands to configure Spanning tree on edge ports.

```
MXL_B1#configure
MXL_B1(conf)#protocol spanning-tree rstp
MXL_B1(conf-rstp)#no disable
MXL_B1(conf-rstp)#exit
MXL_B1(conf)#interface range tengigabitethernet 0/1 – 32
MXL_B1(conf-if-range-te-0/1-32)#spanning-tree rstp edge-port
MXL_B1(conf-if-range-te-0/0-32)#end
```

19. Configure port channel for LAG.
    Use the following commands to configure the switch interconnect as a LAG.

```
MXL_B1#configure
MXL_B1(conf)#interface Port-channel 1
MXL_B1(conf-if-po-1)#mtu 12000
MXL_B1(conf-if-po-1)#switchport
MXL_B1(conf-if-po-1)#no shutdown
MXL_B1(conf-if-po-1)#exit
```

20. Configure QSFP ports for LAG.
    Use the following commands to assign the 40 GB QSFP ports to the port channel.

```
MXL_B1(conf)#interface range fortyGigE 0/33 , fortyGigE 0/37
MXL_B1(conf-if-range-fo-0/33,fo-0/37)#mtu 12000
MXL_B1(conf-if-range-fo-0/33,fo-0/37)#no shutdown
MXL_B1(conf-if-range-fo-0/33,fo-0/37)#flowcontrol rx on tx off
MXL_B1(conf-if-range-fo-0/33,fo-0/37)#port-channel-protocol lacp
MXL_B1(conf-if-range-fo-0/33,fo-0/37-lacp)#port-channel 1 mode active
MXL_B1(conf-if-range-fo-0/33,fo-0/37-lacp)#exit
MXL_B1(conf-if-range-fo-0/33,fo-0/37)#exit
MXL_B1(conf)# exit
```

21. Save the configuration.
    Use the following commands to save the configuration.

```
MXL_B1# copy running-config startup-config
```

6.3.2 Configure the Second Switch (B2)

Repeat the commands from the section above to configure the second switch.

Note: The preceding procedure places all switch ports in the default VLAN. If you prefer to place ports in a non-default VLAN, refer to the documentation for your switch.

6.4 S4810 Configuration (End User)

Two switch stacks comprised of two S4810 switches per stack are used to pass end user's traffic to Fabric A1 and A2. A port-channel consisting of two 40GbE ports from each stack is connected to Fabric A1 and A2.

6.4.1 Configure the First S4810 Switch Stack

Perform the following commands to configure the first S4810 switch stack.

1. Delete the startup configuration.
   Use the following commands to delete the startup configuration.

```
FTOS>enable
FTOS#delete startup-config
Proceed to delete startup-config [confirm yes/no] yes
FTOS#reload
System configuration has been modified. Save? [yes/no] no
Proceed with reload [confirm yes/no] yes
```

Note: The switch will reboot.

22. Confirm the switch OS version.

Note: Each stack unit must be running the same Dell Networking OS version prior to joining the stack.

Use the following commands to confirm the OS version.

```
FTOS>enable
FTOS#show version
Dell Real Time Operating System Software
Dell Operating System Version:  2.0
Dell Application Software Version:  9.5(0.0P2)
Copyright (c) 1999-2014 by Dell Inc. All Rights Reserved.
Build Time: Thu Jul  3 15:45:57 2014
Build Path: /sites/eqx/work/swbuild01_1/build04/E9-5-0/SW/SRC
Dell Networking OS uptime is 6 day(s), 2 hour(s), 43 minute(s)

System image file is "system://B"

System Type: S4810
Control Processor: Freescale QorIQ P2020 with 2 Gbytes (2147483648 bytes)
of memory, cores(s) 1.

128M bytes of boot flash memory.

  1 52-port GE/TE/FG (SE)
 48 Ten GigabitEthernet/IEEE 802.3 interface(s)
  4 Forty GigabitEthernet/IEEE 802.3 interface(s)
```

23. Configure the hostname and OOB management port.
    Use the following commands to configure the hostname and the OOB management port.

```
FTOS>enable
FTOS#configure
FTOS(conf)#hostname S4810_STACK1-0
S4810_STACK1-0(conf)#interface ManagementEthernet 0/0
S4810_STACK1-0(conf-if-ma-0/0)#no shutdown
S4810_STACK1-0(conf-if-ma-0/0)#ip address ipaddress mask
Proceed with Static IP [confirm yes/no]: yes
S4810_STACK1-0(conf-if-ma-0/0)#management route W.X.Y.Z /24 A.B.C.D
S4810_STACK1-0(conf-if-ma-0/0)#exit
```

Note: *W.X.Y.Z* is the network your management system is connecting from and *A.B.C.D* is the gateway for the switch. If your management system is on the same subnet as the switch, this step may be omitted. The example above assumes a class C subnet mask.

24. Configure the login credentials.
    Use the following commands to configure login credentials.

```
S4810_STACK1-0(conf)#username admin privilege 15 password 0 yourpassword
S4810_STACK1-0(conf)#enable password level 15 0 yourpassword
S4810_STACK1-0(conf)#exit
```

25. Enable the switch ports.
    There are two ways to enable the switch ports. **Option 1** enables the ports individually and **Option 2** enables multiple ports using the *range* parameter.

    Option 1
    Use the following commands to enable the ports individually by entering the port number.

```
S4810_STACK1-0#configure
S4810_STACK1-0(conf)#interface tengigabitethernet 0/0
S4810_STACK1-0(conf-if-te-0/0)#switchport
S4810_STACK1-0(conf-if-te-0/0)#no shutdown
S4810_STACK1-0(conf-if-te-0/0)#exit
S4810_STACK1-0(conf)#exit
```

Option 2

Use the following commands to enable multiple ports at once using the *range* parameter.

```
S4810_STACK1-0#configure
S4810_STACK1-0(conf)#interface range tengigabitethernet 0/0 – 47 ,
fortyGigE 0/48 - 60
% Warning: Non-existing ports (not configured) are ignored by interface-
range
S4810_STACK1-0(conf-if-range-te-0/0-47,fo-0/48-60)#switchport
S4810_STACK1-0(conf-if-range-te-0/0-47,fo-0/48-60)#no shutdown
S4810_STACK1-0(conf-if-range-te-0/0-47,fo-0/48-60)#exit
```

26. Configure stacking and reload.

    Use the following commands to configure stacking and reload.

```
S4810_STACK1-0>enable
S4810_STACK1-0#stack-unit 0 priority
S4810_STACK1-0#configure
S4810_STACK1-0(conf)#stack-unit 0 stack-group 13
02:39:12: %STKUNIT4-M:CP %IFMGR-6-STACK_PORTS_ADDED: Ports Fo 0/52 have been
configured as stacking ports. Please save and reload for config to take effect.
S4810_STACK1-0(conf)#stack-unit 0 stack-group 15
02:39:15: %STKUNIT4-M:CP %IFMGR-6-STACK_PORTS_ADDED: Ports Fo 0/60 have been
configured as stacking ports. Please save and reload for config to take effect.
S4810_STACK1-0(conf)#end
S4810_STACK1-0#write memory
S4810_STACK1-0#reload
```

Note: On S4810, Port Fo 0/52 corresponds to stack-group 13, and Port Fo 0/60 corresponds to stack-group 15. For S4810 physical ports to stack-group mapping, consult the S4810 switch configuration guide on the Dell EMC support website.

27. Rename the second S4810 intended to be the standby unit as **Stack-unit 1** using the *`stack-unit 0 renumber 1`* command.
28. Repeat the steps 1 to 6 on the second S4810 (Standby Unit); **Stack-unit 0** should be replaced with **Stack-unit 1**.
29. Connect the second switch to the management unit (**Stack-unit 0**) using 40GbE data cables, and then reload the second switch to join the stack.

30. Use the following commands to verify the stack configuration.

```
S4810_STACK1-0>enable
S4810_STACK1-0#show system brief
Stack MAC : 00:01:e8:8b:34:d9
Reload-Type                :   normal-reload [Next boot : normal-reload]

--  Stack Info  --
Unit  UnitType     Status         ReqTyp       CurTyp          Version     Ports
-------------------------------------------------------------------------------
  0   Management   online         S4810        S4810           9.5(0.0P2)  64
  1   Standby      online         S4810        S4810           9.5(0.0P2)  64
  2   Member       not present
  3   Member       not present
  4   Member       not present
  5   Member       not present
  6   Member       not present
  7   Member       not present
  8   Member       not present
  9   Member       not present
 10   Member       not present
 11   Member       not present

--  Power Supplies  --
Unit   Bay   Status       Type    FanStatus
-------------------------------------------------------------------------
  0     0    up           AC      up
  0     1    up           AC      up
  1     0    up           AC      up
  1     1    up           AC      up

--  Fan  Status  --
Unit Bay   TrayStatus  Fan0    Speed   Fan1    Speed
-------------------------------------------------------------------------------
-
 0    0     up          up      9360    up      9600
 0    1     up          up      9360    up      9360
 1    0     up          up      9600    up      9360
 1    1     up          up      9360    up      9360

Speed in RPM
```

The following commands should be entered on the management switch.

31. Configure port channel for LAG.
    Use the following commands to configure the uplink LAG from the stack to the MXL.

```
S4810_STACK1-0#configure
S4810_STACK1-0(conf)#interface Port-channel 48
S4810_STACK1-0(conf-if-po-48)#no ip address
S4810_STACK1-0(conf-if-po-48)#switchport
S4810_STACK1-0(conf-if-po-48)#no shutdown
S4810_STACK1-0(conf-if-po-48)#exit
```

32. Configure QSFP+ ports for LAG.
    Use the following commands to assign the 40 GB QSFP+ ports to the port channel.

```
S4810_STACK1-0(conf)#interface range fortyGigE 0/48 , fortyGigE 1/48
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#no ip address
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#no switchport
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#port-channel-protocol lacp
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48-lacp)#port-channel 48 mode
active
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48-lacp)#exit
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#protocol lldp
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48-lldp)#no disable
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48-lldp)#advertise management-tlv
management-address system-name
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48-lldp)#exit
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#no shutdown
S4810_STACK1-0(conf-if-range-fo-0/48,fo-1/48)#exit
```

33. Configure spanning tree.
    Use the following commands to configure spanning tree on the stack.

```
S4810_STACK1-0(conf)#protocol spanning-tree rstp
S4810_STACK1-0(conf-rstp)#no disable
S4810_STACK1-0(conf-rstp)#end
```

34. Save the configuration.
    Use the following commands to save the configuration.

```
S4810_STACK1-0#copy running-config startup-config
```

6.4.2 Configure the Second S4810 Switch Stack

Repeat the commands from the section above to configure the second S4810 switch stack.

Note: The preceding procedures places all switch ports in the default VLAN. If you prefer to place ports in a non-default VLAN, refer to the documentation for the switch.

7 PS-M4110 Storage Configuration

7.1 PS-M4110 Controller

The PS-M4110 controller (Table 6) is based on a modified version of the PS-4100 Controller. Host and SAS cards are combined to form a single unit fixed I/O module, connecting to the M1000e chassis infrastructure.

Table 6 PS-M4110 Storage Controller

Controller Type	Network Interfaces	Storage Type	Notes
Type 13	2 x 10Gb Ports, (One Per Controller), Connected through the Backplane Each of the two ports has an active link and a standby link to the fabric switches in the backplane of the M1000e chassis.	SAS NL-SAS	• Dual, hot-pluggable 10GbE controllers • 2GB of memory per controller, Cache to Flash Design • 1 x dedicated management port – accessible via CMC • 14x 2.5” 6Gb/s SAS disk drives

Failover behavior in the PS-M4110

Each PS-M4110 array is configured with dual controllers, active and passive connections, and redundant fabrics to provide redundancy and prevent volume connections between the hosts and SAN from being dropped in the event of a failure.

Three failure scenarios protected by the array are:

- In the event of a link failure, the active link of the active port goes down, and the standby link for the same port takes over.
- In the event of a switch failure, both active and passive ports automatically link to the redundant fabric switch.
- If the active controller fails, the passive controller takes over.

7.2 PS-M4110 Setup

Figure 11 shows the four storage arrays used in this solution, the two arrays used for the SQL environment are installed in slots 5 and 13 and the two arrays used for Exchange are in slots 7 and 15.

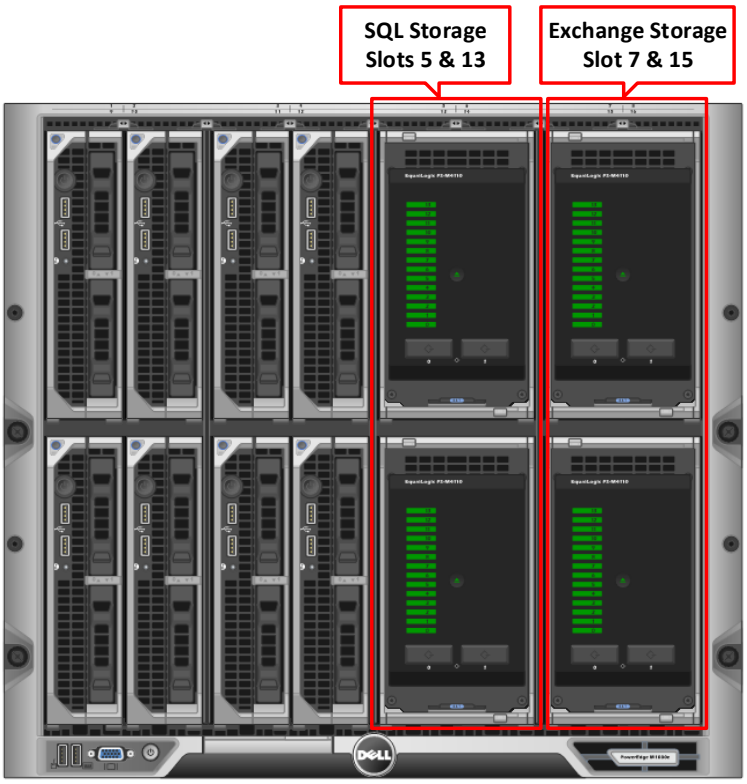

Figure 11 PS-M4110 Storage Arrays

PS-M4110 Setup Overview

- Configure one EqualLogic Group with two pools.
- Add each of the two PS-M4110E blade arrays to the respective pool.
- Install and configure the EqualLogic Host Integration Toolkit (Microsoft Edition for the hosts).
- Configure RAID-6 on both arrays (in the Web interface, right-click the member, and choose Modify Raid Configuration).
- Create four volumes of 1.2 TB on each pool. The volumes should be thick volumes, no snapshot reserve, default of 512 blocks per sector and access restricted to the one server accessing that pool.
- Connect the volumes of Array 1 to Server 1, and volumes from Array 2 to Server 2 respectively.
- When the disks are writable and online, convert them to GPT (GUID Partition Table), and create one primary partition for each volume.
- Format the disks with NTFS and 64k unit size.
- Format the disk sequentially within each pool, i.e. not parallel.
- After configuration, use SAN Head Quarters (SANHQ) to monitor your SAN for any issues.

Perform the following steps to set up the PS-4110.

1. Login to the PS-M4110
   Perform the following steps to log in to the Storage Blade.

a. Power on the chassis.
b. Login to the chassis CMC.
c. From the CMC CLI, use the following command to connect to the Dell Storage Array.

```
$connect server-n
```

Note: Where *n* is the slot number for the M4110 storage device. In this environment, the SQL storage is in slots 5, 13, and the Exchange storage is in slots 7, 15.

The following example shows the command to connect to the Dell Storage Array in slot 5.

```
Example
$connect server-5.
Escape character is '^\'.
```

35. Login to and configure the Storage Array
    Use the following commands to login to and configure the storage array.

```
login: grpadmin
Password: grpadmin
           Welcome to Group Manager
        Copyright 2001-2014 Dell Inc.
Local authentication. User privilege is group-admin.
It appears that the storage array has not been configured.
Would you like to configure the array now? (y/n) [n]y
Group Manager Setup Utility

Do you want to proceed (yes | no ) [no]: yes
Initializing.  This may take several minutes to complete.
 Enter the network configuration for the array.
   Member name []: PS-M4110E-Slot5
   Network interface [eth0]:
   IP address for network interface []: 10.10.10.11
   Netmask [255.255.255.0]: 255.255.255.0
   Default gateway [10.10.10.1]:

Initializing interface eth0.  This may take a minute.....
Enter the IP address and name of the group that the array will join.
      Group name []: 10.10.10.10
Valid responses include alphanumeric characters along with middle and end
      Group name []: Trinity
      Group IP address [10.10.10.10]:
```

```
 Searching to see if the group exists.  This may take a few minutes.
   Password for managing group membership:
   Do you want to add this array to group Trinity (yes | no ) [yes]: yes
 Synchronizing configuration data with the group.  This may take
Saving the configuration...
Waiting for configuration to become active.......Done

 The array has successfully become a member of group Trinity.
 You must now configure the RAID set using the Group manager GUI or CLI.
 Configuration of the RAID set requires that the raid-level be selected.
```

A Component Tables

The following tables (Table 7 through Table 11) present the components used in this solution.

Table 7 M1000e Chassis

Component	Description//Firmware Version
PowerEdge M1000e Chassis Additional Information	
Chassis Management Controller (CMC) Firmware	4.50, A00
CMC Hardware Version	A03
iKVM Switch	01.00.01.01, A03
Midplane Version	1.1
Flash Media	FlexAddress, Extended Storage

Table 8 Server

Components	Description//Firmware Version
Dell PowerEdge M620 Additional Information	
Memory	96 GB DDR3
Processors	2 x Intel(R) Xeon(R) CPU E5-2670 @ 2.60GHz
Operating System	Microsoft Windows Server 2012 R2 Datacenter Edition
Sever BIOS	2.2.10
Integrated Dell Remote Access Controller (iDRAC)	1.57.57
Lifecycle Controller (LC)	1.4.2.12
32-Bit Diagnostics	4241A0
OS Driver Pack	7.4.2.05
Broadcom NetXtreme II 57810S-k Dual-port 10GbE NDC	7.8.53
RAID Controller (PERC H310 Mini)	20.13.0-0007
CPLD	1.0.7

Table 9 Switch Modules

Component	Description//Firmware Version
MXL Switch Module Additional Information	
Fabric A Switch Modules - MXL Switch Module	9.5
Fabric B Switch Modules - MXL Switch Module	9.5
Fabric C Switch Modules - Empty	N/A

Table 10 Storage Array

Component	Description//Firmware Version
Dell EqualLogic PS-M4110 storage array Additional Information (PS-M4110XV for SQL and PS-M4110E for Exchange)	
Storage Array Firmware	Version 7.0.7
RAID Configuration	Raid 6
Disk Model (SQL storage)	ST9146853SS (15K SAS)
Disk module (Exchange storage)	ST91000640SS (7.2K SAS)
Drive Firmware (SQL storage)	YEOA
Drive Firmware (Exchange storage)	AEOD

Table 11 Software and Software Utilities

Component	Description
Microsoft Windows Server 2012 R2 Additional Information	
Microsoft Exchange Jetstress 2013	15.00.0775.000
Microsoft SLQIO	1.0.0
Microsoft Exchange Server 2013 Additional Information	2013
Microsoft SQL Server 2014 Additional Information	2014
EqualLogic SAN Headquarters	3.0.1.7331
EqualLogic Host integration Tools	4.7.1
iSCSI Initiator	Windows Server 2012 R2 Software Initiator
OpenManage System Management Software	7.4.0.1
M620 Chipset Drivers	9.3.2.1017,G00
Network Adapter Drivers – Broadcom NetXtreme II 57810S-k Dual-port 10GbE	7.0.5.43
Broadcom Driver and Management Applications	16.4.5.1

B Attachments

The following Microsoft Exchange Jetstress 2013 Performance Test Result Reports are attached to this document:

- Server-1-Jetstress_Performance_Test.html
- Server-2-Jetstress_Performance_Test.html

This following switch configuration files are attached to this document:

- MXL_A1_Config.pdf
- MXL_A2_Config.pdf
- MXL_B1_Config.pdf
- MXL_B2_Config.pdf
- S4810_STACK1-0.pdf
- S4810_STACK2-0.pdf

Note: To view the attachments in Adobe Reader or Adobe Acrobat click on the Attachments button (Paperclip) in the left side navigation panel of the main Reader window or from the Menu bar, select **View** > **Show/Hide** > **Navigation Panes** > **Attachments**.

C Reference Information

- Microsoft Exchange Server Jetstress 2013 Tool
  http://www.microsoft.com/en-us/download/details.aspx?id=36849

- SQLIO Disk Subsystem Benchmark Tool
  http://www.microsoft.com/en-us/download/details.aspx?id=20163

- Dell TechCenter Networking Publications
  http://en.community.dell.com/techcenter/networking/p/guides

- Dell TechCenter Storage Publications
  http://en.community.dell.com/techcenter/storage/w/wiki/2660.equallogic-technical-content

- Performance Test Dell PS-M4110 vs HP Blade
  http://www.dell.com/learn/ng/en/ngdhs1/shared-content~data-sheets~en/documents~principled_technologies_datacenter_in_a_box_solution_performance_test_dell_ps_m4110_vs_hp_blade.pdf