

Microsoft® Exchange Server 2007 on Dell™ EqualLogic™ Storage – Best Practices

End to End Solutions
Engineering

By Mahmoud Ahmadian, Farrukh Noman,
Casey Birch
Dell Product Group
October 2008



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.

© 2008 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell, Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, *PowerEdge*, and *EqualLogic* are trademarks of Dell Inc. *Microsoft* is either a trademark or registered trademark of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Table of Contents

INTRODUCTION	4
WHY ISCSI SAN?	5
BENEFITS OF DEPLOYING EXCHANGE ON DELL EQUALLOGIC PS SERIES ISCSI SAN	5
EASY TO DEPLOY	5
RELIABILITY	6
DATA PROTECTION	6
HIGH PERFORMANCE	7
EASY EXPANSION	7
THIN PROVISIONING	7
BEST PRACTICES FOR DELL EQUALLOGIC PS SERIES ISCSI SAN	8
ISCSI INFRASTRUCTURE BEST PRACTICES	8
<i>Network Recommendations</i>	<i>8</i>
<i>Host/Initiator Recommendations</i>	<i>9</i>
<i>Storage Connections Recommendations</i>	<i>10</i>
STORAGE DESIGN BEST PRACTICES	10
<i>Configuring Exchange Volumes</i>	<i>11</i>
<i>Sizing for Capacity and Performance</i>	<i>12</i>
CONCLUSION	15

INTRODUCTION

During the last ten years, messaging services have transformed the way business and IT is done, from providing basic functionality such as e-mail to providing rich collaboration features. Microsoft® Exchange as one of the most popular and widely used e-mail applications today. The rich features in these applications has significantly changed the messaging requirements and needs of today's information worker from basic desktop e-mail access to collaboration from anywhere with any device. This has resulted in ever increasing demand on storage systems both in terms of capacity and bandwidth. The IT administrators are also facing the increasingly complex task of coping with new requirements and maintaining messaging systems.

Many organizations are rethinking their storage strategies to meet the demanding criteria and handle the future requirements. Scalable storage solutions like Storage area network (SAN) are becoming increasingly complex and expensive, especially with the fiber-channel environment. A virtualized storage solution is required that should be easy to manage and flexible to meet the security, availability and scalability needs of an organization. Dell EqualLogic iSCSI solution incorporates features that enable businesses to effectively meet those challenges. This whitepaper highlights some of these features and benefits of Dell EqualLogic PS storage and recommends best practices for a high performing, scalable, flexible and reliable Microsoft Exchange 2007 environment.

Why iSCSI SAN?

Direct Attached Storage (DAS) has been widely used for Exchange deployments. Although DAS may be adequate for some businesses in terms of performance but it may not be the best fit in terms of consolidation and scalable needs of large or growing organizations. Each DAS storage system is isolated and as a result managed and backed up separately. Even the most routine maintenance activities such as adding new storage require planned down time. Unless Microsoft Exchange Continuous Cluster Replication (CCR) is deployed, some DAS storage solutions are susceptible to unplanned down time with potentially harmful business results.

A Storage Area Network or SAN overcomes the potential problems and limitation posed by DAS environments. Fiber-channel (FC) SANs have been used in the industry for quite a long time but add more complexity to the overall solution. They provide scalability and performance but they are typically higher priced due to the additional FC infrastructure requirements and they are complex to configure and manage.

Internet SCSI (iSCSI) technology has matured over the past few years and delivers SAN like SCSI block traffic over an existing LAN based TCP/IP infrastructure. The iSCSI SAN requires no additional networking infrastructure while delivering FC SAN like capabilities at a more attractive price point. With its ease of management and leveraging the existing IP networking expertise, iSCSI SAN are getting popular within IT organizations. Even in performance, iSCSI SAN deployed in a dedicated Gigabit environment has proved to be on par with FC SAN and thus makes it a good candidate for critical applications like Exchange.

Benefits of Deploying Exchange on Dell EqualLogic PS Series iSCSI SAN

Dell EqualLogic PS Series iSCSI Storage solution takes the management capability in iSCSI SAN to the next level by driving intelligence into the Storage system. The unique management set provides seamless integration of multiple EqualLogic arrays, transparent load balancing, provisioning and automation of tasks internally with minimal intervention. By deploying Exchange Server with a PS Series SAN, businesses can combine the industry-leading messaging and collaboration application with reliable, scalable, and high-performance disk storage to meet the ever-expanding needs of e-mail users. All features and capabilities come standard with the PS Series iSCSI SAN at no additional costs and greatly simplify IT for the following reasons.

Easy to Deploy

PS Series Storage automates operations such as RAID configuration, assigning hot spares as well as data provisioning and load balancing. Traditionally, these functions have required storage and SAN experience. The Dell EqualLogic iSCSI SAN enables customers new to SAN to configure PS Series storage and create volumes for Exchange store in a few short steps.

The EqualLogic PS Series supports tiered storage model which allows greater flexibility of allocating disk resources. Multiple tiers can be configured to satisfy different organizational needs. In an EqualLogic iSCSI SAN environment, an entity called "group" provides a virtualized access through a single IP address of all PS Series arrays (members) within it. One or more members comprise a "storage pool" from which volumes are created. The underlying RAID type for each member is chosen before it is added to the storage pool. The storage volumes are virtualized and can be automatically tuned for optimal performance by moving them to best RAID type based on I/O characteristics or can be manually configured to use a defined RAID set.

Host access to volumes is controlled by one of three methods: authentication using Challenge Handshake Authentication Protocol CHAP; limit access by IP address; limit access to iSCSI initiator name. Device Specific Module (DSM) shipped with PS Series arrays integrates with Microsoft MPIO to make configuring host access simple.

Reliability

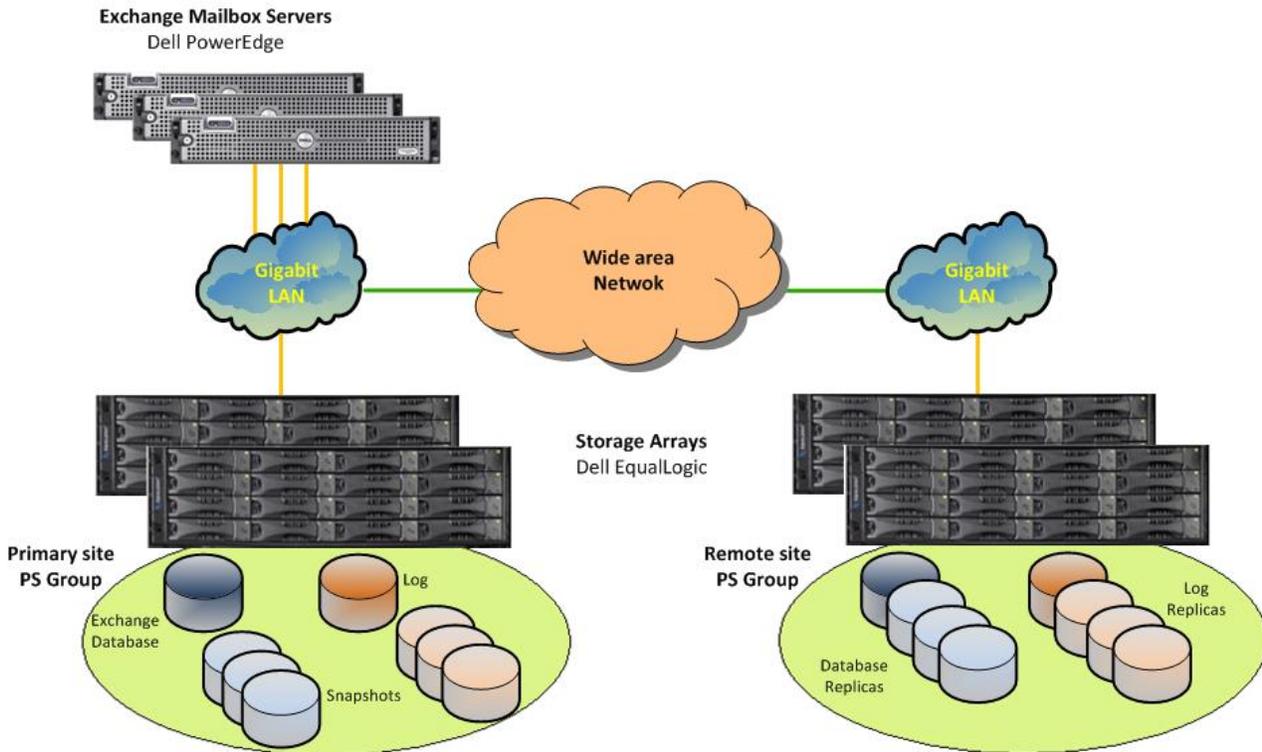
The PS Series arrays come with redundant dual controllers, power supplies and fans. In the event of component failures, applications are not affected and continue to stay online without disruption even when these hardware components are being serviced. In addition, a background disk monitoring service on the PS Series array known as Auto-Stat Disk Monitoring System (ADMS) scans disks to ensure optimal data availability. On the host side, the reliable access to the storage volumes is provided through multiple NICs or iSCSI HBAs and MPIO software. Dell EqualLogic DSM with multi-pathing feature enabled on Microsoft iSCSI Initiator provides failover and load balancing between multiple connections.

Data Protection

To guard against data loss, the first level of data protection is provided with RAID and hot spares. In addition, volume snapshot and replication functionality are built into PS Series arrays to provide an added level of data protection. Dell EqualLogic PS Series host software includes Volume Shadow Copy VSS provider and Auto-Snapshot Manager, enabling tight integration with Windows 2003 for frequent and fast backups using point-in-time-copies or snapshots. Snapshots are fast, space efficient and allow more frequent backup operations. There is no disruption in access to the production volume during snapshot creation. They capture volume changes in the snapshot reserve space. The unchanged pages are shared between the base volume and the snapshot.

Dell EqualLogic iSCSI SAN volume replication is another method of data protection where the primary and remote site groups are configured as replication partners for purposes of site protection and disaster recovery. The first replica is complete copy of the base volume. The second replica captures only the incremental changes similar to a snapshot process. Replica volume is formed from the first replica and the changed data from the second replica. This process is continued as additional replicas are generated. In the event of a site failure, the secondary site and the replica can be brought online. Once the primary site is restored, the only the changes at the secondary site need to be synched back to the primary before the primary site can be brought online.

Figure 1. Data Replication with Snapshots



High Performance

The deployment of PS Series in Exchange environment provides a comprehensive implementation with high performance that matches the expectation of email services. Dell EqualLogic storage technology provides an enhanced process to boost the disk I/Os through efficient buffering and optimized seek mechanism on the RAID controllers. The storage volumes are virtualized and accessed through single group IP while the internal group connections are dynamically load-balanced between members to provide maximum performance. All active disks within a member are always utilized spreading the data across all disks within a RAID set.

Easy Expansion

For additional storage capacity and performance, PS Series array members may be added to an online group without disrupting applications. A group can include up to 12 members and each storage pool can have up to 8 members. There is no degradation in random I/O performance when multiple arrays are added to a group. The performance scales linearly because all the arrays are equally accessible with the same bandwidth and the load balancing of data and network I/O occurs automatically.

Thin Provisioning

Storage space provisioning often presents a challenge to many IT operations. To avoid running out of application disk space and risking of downtime, excessive amounts of disk space are allocated for application data. Over-allocation can degrade performance and compromise scalability and leads to poor resource utilization. To take the guess-work out of the equation, Dell EqualLogic PS array thin provisioning method assigns a virtual amount of storage to a volume. It allows maintaining a single free space pool to service the immediate

growth requirements of applications. Compared to a standard provisioning method, thin provisioning reduces the power and cost due to over-allocated disk resources.

Best Practices for Dell EqualLogic PS Series iSCSI SAN

A PS series group is a single entity which consists of one or more Dell EqualLogic storage arrays. The group IP address is used to manage a group and provide iSCSI access to the volumes. To get the best performance and availability, this section outlines the best practices for the EqualLogic array infrastructure and storage design parameters in Exchange environment.

iSCSI Infrastructure Best Practices

PS Series iSCSI SAN utilizes standard IP gear and interconnect without costly additions to the networking infrastructure. However, to meet performance, reliability and scalability goals, iSCSI network components configuration should observe the following recommendations for optimal results.

Network Recommendations

To achieve good performance for an Exchange solution, it is essential that iSCSI SAN traffic be separated from the client traffic through dedicated network or VLANs. The data network configuration should also be optimal to utilize the full bandwidth and take advantage of EqualLogic array's failover capabilities. Some of the network parameters that require careful consideration for the Dell EqualLogic iSCSI setup are briefly discussed below.

- **Switches and Inter-links**

The iSCSI nodes/initiators and Dell EqualLogic array connections are recommended to use a dedicated gigabit network environment. As the connections in iSCSI network are full duplex, so each port on the switch should be able to handle 2Gbps traffic. Switches with higher bandwidth and larger buffer space should be preferred.

When providing high-availability or expanding EqualLogic iSCSI network environment, proper consideration is required while interlinking multiple switches together. Multiple inter-switch links can give rise to multiple loops and requires spanning tree protocol to figure out the usable links. For data-intensive applications like Exchange, the loop detection process can impact the performance and increase database latency delays. The inter-switch links should be minimized using large switches instead and if deployed should be correctly provisioned with multiple trunks to avoid bandwidth congestion between them.

The types of switches also play an important part in the network configuration. Layer-2 switches as mentioned earlier should be checked for loops and inter-switch bandwidth links. Layer-3 switches which are usually deployed between different subnets for routing needs should be monitored to avoid bottleneck between linking subnets.

- **Flow Control**

The flow control mechanism in an IP network allows the source and target to control the flow of packets and avoid buffer overflow on the receiving side. This eventually avoids retransmission of lost packets. In an EqualLogic iSCSI SAN, the flow control mechanism should be enabled on all the devices from host nodes to arrays. The switches should have enough buffer space and bandwidth available for the individual connecting ports to handle the bursts of iSCSI traffic. The buffer space is usually shared between all the online ports and therefore limits the capacity reserved for each port. So it is strongly recommended to enable flow control on the switch ports to

avoid any packet loss and retransmissions. Having the flow control process between connecting devices greatly reduces the delays and improves overall performance.

- **Unicast Flow Control**

The iSCSI traffic is unicast traffic and typically it can storm the entire link. That is one of the reasons why iSCSI data traffic needs to be isolated from the client network traffic. If the physical isolation is not possible then VLANs should be configured on the switches. It is recommended to disable the unicast storm control on the switches and enable the broadcast and multicast storm control if applicable.

- **Jumbo Frames**

The jumbo feature extends the size of standard Ethernet frame from 1500 to 9000 bytes, allowing six times the size of standard payload in each frame. Jumbo feature provides improvement because of less fragmentation and reduction in overall network frames. The success of this process relies on the fact that the jumbo feature should be enabled on the entire path from source to target including all the intermediate devices. The EqualLogic arrays are automatically enabled to use jumbo frames and it is recommended to enable this feature on the host nodes and the connecting switches.

Host/Initiator Recommendations

An iSCSI host is configured with initiator to access iSCSI target on the storage side. The affordable infrastructure of iSCSI with existing IP network provides flexibility to use regular gigabit Network Interface card (NIC) with software initiator or iSCSI Host bus adapter (HBA) with hardware initiator. With a software initiator and NIC, the TCP/IP and iSCSI protocol stacks are processed in the host software and therefore consumes more CPU cycles for the network traffic. On the other-hand a hardware initiator offloads the TCP/IP and iSCSI process onto the HBA and thus has less impact on the host CPU. Hardware initiators are strongly recommended for larger or heavy Exchange deployments where mailbox server CPU consumption is already high due to heavy database transactions.

The type of PCI bus technology and slots used for NICs or iSCSI HBAs are carefully considered to avoid any performance impacts. The 32-bit, 33MHz PCI bus technology limits the bandwidth to less than 132MB/s and more than one PCI slots on a bus share the bandwidth and reduces it even more. It is recommended to use 64-bit and higher speed PCI slots to achieve good performance.

The iSCSI host can be configured with multiple physical connections to provide path failover and load-balancing features. Dell EqualLogic Multipath I/O (MPIO) DSM works in conjunction with Microsoft MPIO driver to enable those features. When the MPIO DSM is enabled, the multiple sessions for each target volume are configured by default to use the Least Queue Depth load-balancing policy. The default policy can be changed to any of the following supported policies to better tune a specific topology.

- a) **Fail Over only:** This policy does not require MPIO DSM and is by default available to activate the standby path in case of active path failure.
- b) **Round Robin:** Round-robin policy makes effective use of all the available paths by transmitting data in rotating fashion.
- c) **Least Queue Depth:** This policy serves the data through the connection which is least busy and has the least number of outstanding I/Os to serve.

The initiator nodes are provided access to the storage volumes through volume Access Control list (ACL). Each storage volume which is also an iSCSI target, maintains its own access list. To prevent unauthorized nodes from discovering the volume targets, it is

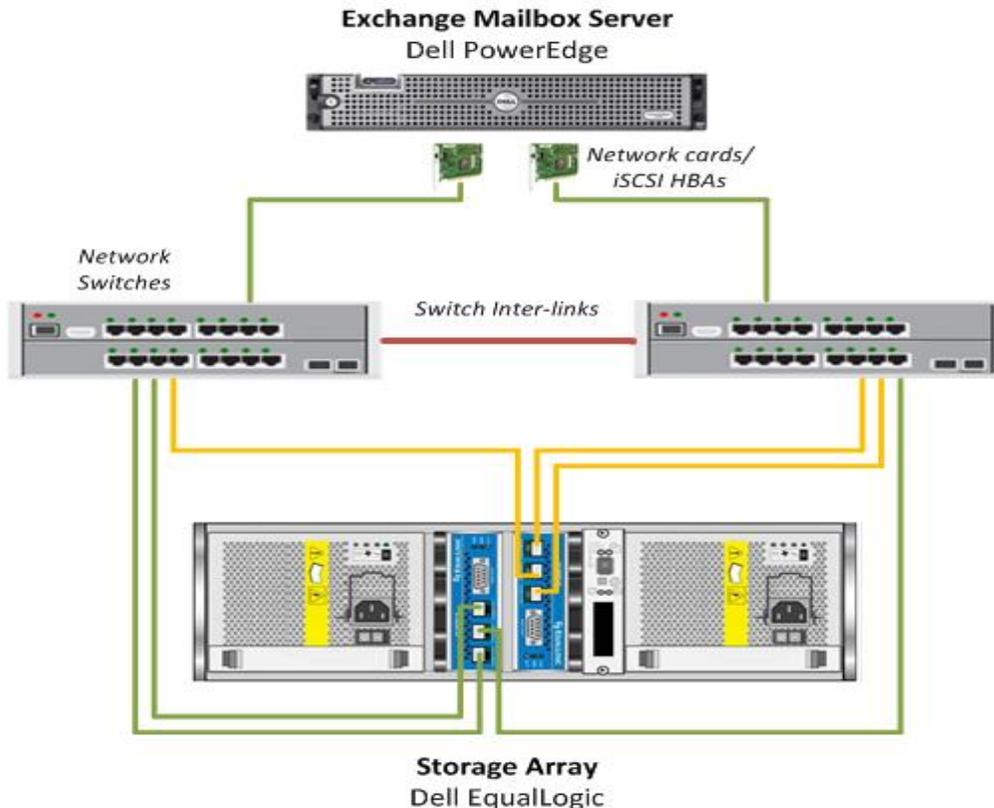
recommended to provide explicit access to the intended iSCSI hosts only and not to the complete subnets.

The volumes should be configured to align disk partition with the RAID stripe size of the array. For critical applications like Exchange, the partition alignment can improve the overall performance. Dell EqualLogic arrays have a default RAID stripe size of 64K bytes, equivalent to 128 sectors. The volumes can be formatted to skip first 64K bytes and have the usable disk space beginning 129th sector of the partition for proper alignment.

Storage Connections Recommendations

The Dell EqualLogic storage system has two high performance controller modules with three gigabit Ethernet interfaces/ports in each module. Two controllers have three pairs of interfaces which can be configured to provide path redundancy and data load-balancing. The controller modules' cache is mirrored for complete redundancy and automatic failover in case the primary module goes down. Therefore, it is strongly recommended to connect both modules to the data network to avoid a single point of failure. To maintain high-availability of network paths, the Ethernet interfaces of a module should be connected to different switches if possible.

Figure2. Dell EqualLogic Storage array with redundant switch connections



Storage Design Best Practices

To realize optimal performance, and full benefits of deploying Dell EqualLogic, arrays design best practices for deploying Exchange log and database volumes must be adhered to. Proper sizing in terms of capacity and performance is critical to any Exchange storage design. Once specific capacity and IO requirements are clearly defined, the following sections can be consulted for a discussion on how these requirements can be met by the PS 5000 series array. For information pertaining to determining your specific sizing requirements refer to this

Microsoft TechNet article on Exchange 2007 best practices: [http://technet.microsoft.com/en-us/library/bb738147\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738147(EXCHG.80).aspx)

The following section also includes a discussion and recommendations regarding deploying Exchange database and log volumes on a virtualized storage platform as many organizations may have typically only deployed Exchange on traditional DAS and SAN solutions in the past.

Configuring Exchange Volumes

In an EqualLogic iSCSI SAN environment, a group provides virtualized access through a single IP address to all of the PS Series arrays/members within it. A storage pool, comprised of one or more arrays within a group, forms a RAID subsystem from which volumes are created. The disk resources of the member array(s) within a storage pool are considered 'virtualized' in that data from volumes created on the storage pool is dynamically load balanced across all disk resources in that pool to maximize performance and efficiency. There is no way to simply dedicate certain physical spindles within a storage pool to specific volumes as with traditional storage arrays. As discussed above, this approach avoids hot spots and performance bottlenecks, but what ramifications, if any, does this have on Exchange design?

In Microsoft's documented best practices for Exchange 2007 (mentioned in TechNet article referenced above) it is recommended that transaction log volumes and database log volumes be implemented on separate physical disks. The reason for this is two-fold:

- 1) In *some* storage arrays mixing the sequential IO of logs and random IO of the databases can result in significant performance degradation.
- 2) In the case of multiple disk failures, assuming both logs and databases share the same disks, email from the last good backup to the point of failure may not be recoverable.

As any given EqualLogic array can only be a member of a single storage pool, many deployments of one or just a few arrays may require the Exchange database and log volumes all be placed within a single pool, which provides the potential for logs and databases to share physical spindles. Even in larger deployments of several arrays where it is possible to separate the Exchange log and databases onto separate pools, a storage or Exchange administrator may choose not to do so in favor of the simplicity gained by having all Exchange volumes on a single pool. In these cases where both log and database volumes are on a single storage pool, and therefore the potential exists for both volumes to share physical disk resources, there are certain technological and design considerations to keep in mind that alleviate concerns of going against Microsoft recommendations.

- **Performance Considerations**

As the Microsoft best practices state, only some, not all, storage arrays suffer when combining sequential (transaction log) and random (database) IO on the same physical resources. Dell EqualLogic technology allows the mixing of random and sequential workloads without impacting performance. The virtualization features of the array actually give EqualLogic arrays some of their performance advantage over competing products. As mentioned above, advanced buffering and optimized seek mechanisms give Dell EqualLogic arrays a performance boost above their competitors.

- **Recoverability Considerations**

Utilizing RAID protects Exchange data from single disk failures. However, losing more than one disk simultaneously within a mirror set can cause the loss of data. How failures impact different RAID types is discussed further below. If the resulting data loss from a multi-disk failure affects only one of either the Exchange databases or its transaction logs, it will be possible to restore data up to the point of failure with

recent valid backups. However, if both database and transaction log data are lost in a multi-disk failure it is likely that data will only be able to be recovered from the time of the last good backup, barring the use of some forms of data replication as discussed below. It is this scenario that Microsoft intends customers to avoid with their stated best practice.

EqualLogic PS Series arrays are built with redundant components, built-in monitoring and alerting, and recommended configurations include multiple hot spares that auto-rebuild in case of a failure; these features provide robust layers of protection against data loss. Additionally, the conditions for a multiple disk failure to cause data loss requires that a single mirrored pair within a RAID 10 set, or two disks within the same RAID 5 subset of a RAID 50 set fail simultaneously. Given the current manufacturer MTBF (Mean Time Between Failure) and AFR (Annual Failure Rate) rates for hard drives the risk of data loss in a RAID configuration with two hot spares is extremely low. Collectively, these measures make a multiple disk failure that causes irretrievable data loss highly unlikely within an EqualLogic array in a RAID 10 or 50 configuration with two hot spares.

Continuous replication of the Exchange data to a different storage pool alleviates any remaining concerns of data loss. Among high availability models for Exchange, Clustered Continuous Replication (CCR) is the most recommended because it protects not only the server but also the data by replicating it to a passive cluster node. Replicating the Exchange data from the active CCR node on one storage pool, to the passive CCR node on a separate storage pool protects Exchange from any physical failures on the active node. See the following Dell white paper on Exchange 2007 High Availability for further information on CCR:
http://www.dell.com/downloads/global/power/Dell_ExchSP1_HA.pdf

Sizing for Capacity and Performance

Exchange storage design is governed by space and I/O requirements while meeting expected performance targets. The number of disks required for Exchange volumes depends on 5 factors: disk speed, disk size, the RAID type, effective storage space required and IOs per second (IOPs) demand. Exchange storage requirement is capacity-bound if the dominating factor is space which is usually the case with larger mailboxes. In such scenarios, larger capacity and lower speed disk drives may be utilized to meet larger space requirements. The solution is IOPS-bound if the prevailing criterion is I/Os which is usually the case with heavy profile mailboxes. Under these circumstances, high speed disk drives may be utilized to meet performance requirements.

- **RAID Type**

The RAID engine technology in Dell EqualLogic array is designed to handle random I/Os of Exchange database and distribute it effectively across all storage arrays within a PS pool. With multiple arrays and different RAID types in a pool, the automatic load balancing algorithm moves volume between arrays members depending on the best RAID fit and performance. Selecting a particular RAID level for Exchange deployment depends on the workload, I/O performance requirement and disk failure rates in the storage array. In general, RAID-10 is more reliable and provides higher I/O per second but it requires more spindles. On the other hand, RAID-50 comparably provides less performance but is more cost-effective in terms of providing higher storage capacity. The performance is also impacted with disk failure rate and rebuild-time in a particular RAID type. When a disk fails in RAID-10, the read performance only decreases on the RAID-1 set with the failed disk and so the overall impact is minimal. The RAID-10 rebuild also involves only two disks; the active disk in RAID-1 set and spare disk. But in case of RAID-50, one disk failure impacts the reads and writes on all the disks. The write I/Os are more affected than reads due to the parity calculations. RAID-50 is comparatively more robust than

RAID 5 due to higher availability with more sets and capacity to handle more disk failures. Factors like IOPS requirements, degraded mode performance and RAID rebuild period largely influence the selection criteria of particular RAID type in an Exchange deployment.

- **Disk Drives Size and Performance**

Dell EqualLogic PS series products are tiered based on disk technology. The PS5000XV utilizes the 15K RPM SAS drives up to 450GB per disk targeting the most demanding IOPS-bound Exchange environments. These arrays provide higher IOPS and are usually deployed to satisfy heavy mailbox profile requirements. The PS5000X is based on the 10K RPM SAS drives targeting moderate to high IOPS-bound Exchange workloads with up to 400GB per disk drive. For capacity-bound Exchange deployments the PS5000E supporting up to 1 TB, 7200 RPM SATA drives is the most likely choice. The SATA arrays provide lower IOPS per spindle and therefore preferred for light to medium profiles with large mailboxes.

To better understand the performance and capacity of Dell EqualLogic arrays, as it relates specifically to Exchange 2007, Dell has used JetStress to characterize Exchange 2007 performance in various RAID and drive configurations. JetStress is a Microsoft developed tool for simulating Exchange disk IO load and is used in the ESRP program (<http://technet.microsoft.com/en-us/exchange/bb412164.aspx>). The results in the table below (Table 1) are to be used as a guide on how to size Dell EqualLogic arrays for your Exchange 2007 deployment once you have determined your exact IO and capacity requirements.

Table 1 – Sizing Guide for EqualLogic Arrays

EqualLogic Array	RAID Type	Random IOPS (approx. database IOPS)	Raw Capacity (after RAID)
PS5000 E (using 14x1TB, SATA II drives + 2 hot spares)	RAID 10	600	7 TB
	RAID 50	300	12 TB
PS5000 X (using 14x400GB, 10K RPM drives +2 hot spares)	RAID 10	1700	2.8 TB
	RAID 50	1200	4.8 TB
PS5000 XV (using 14x300GB 15K RPM drives + 2 hot spares)	RAID 10	2300	2.1 TB
	RAID 50	1700	3.6 TB

Please note the following key points related to the IOPS numbers in the above table:

1. The IOPS numbers are the result of thorough testing with JetStress and are a somewhat conservative guideline as these numbers were often exceeded in lab testing.

2. The IOPS numbers only take into account the random IO generated by database reads and writes. Database IO makes up the large majority of Exchange disk IO and is typically the bottleneck in terms of IO. Disks are typically capable of handling larger amounts of sequential IO (as generated by transaction logs).
3. These numbers assume a Read/Write ratio of 50/50 (i.e. 50% reads and 50% writes), as observed when Exchange 2007 mailbox servers have the recommended amount of memory for caching. A significantly higher percentage of reads or very large heavy mailboxes could result in lower random IOPS.
4. The IOPS numbers reflect the perspective of the server or application, as any write penalties accrued because of RAID type are already accounted for.

Conclusion

Dell EqualLogic PS Storage systems provide an intelligent and virtualized storage solution that enables businesses to suitably meet their ever growing messaging needs. It overcomes much of the challenges addressed with a traditional FC SAN through its easy management and enhanced backup methodologies. This whitepaper highlighted some of these features and benefits of Dell EqualLogic PS storage and recommended best practices for a high performing, scalable, flexible and reliable Microsoft Exchange 2007 environment.

Dell PowerEdge™ Servers and Dell EqualLogic Storage provide a standards-based hardware platform for deploying Exchange Server 2007 messaging solutions. More information can be obtained at www.dell.com/exchange. Dell Services include assessment, design and implementation tailored for those messaging deployments.