Dell EqualLogic PS Series

10,000-User Mailbox Resiliency Storage Solution for Microsoft Exchange Server 2010

ESRP – Storage Version 3.0

PS Series Firmware Version 5.0.2

Tested with: ESRP – Storage Version 3.0

Test Date: June 2011

Document Version: v1.0

July 2011

Table of Contents

Overview

This document provides information on the Dell storage solution for Microsoft Exchange Server, based on the *Microsoft Exchange Solution Reviewed Program (ESRP) – Storage* program[1]. For any questions or comments regarding the contents of this document, contact Dell.

Disclaimer

This document has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty, express or implied, with respect to the accuracy of the contents of this document.

The information contained in this document represents the current view of Dell Inc. on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted to be a commitment on the part of Dell, and Dell cannot guarantee the accuracy of any information presented after the date of publication.

[1] The *ESRP – Storage* program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide information on their storage solutions for Microsoft Exchange Server software. For more details on the *Microsoft ESRP – Storage* program, see the following URL: http://technet.microsoft.com/en-us/exchange/ff182054.aspx

Dell EqualLogic PS Series storage array features

With Dell™ EqualLogic™ PS Series storage arrays, your business can leverage its existing Ethernet infrastructure and deploy a comprehensive, easy-to-manage iSCSI SAN with all-inclusive enterprise-level features. The PS Series architecture was specifically designed to decrease the storage management burden on IT administrators and alleviate CIO planning and budgetary concerns. Dell chose the iSCSI protocol—with its networking and connectivity advantages—as the basis of the storage solution, and then built intelligence, automation, and redundancy into each PS Series storage array.

PS Series storage arrays use storage virtualization technology to mask the underlying complexity of the storage configuration. This virtualization occurs within and across multiple arrays that are logically grouped together, making management simple and efficient. Reliable hardware, intuitive graphical and command line user interfaces, and automated operations offer excellent productivity and service levels, while RAID configuration, storage expansion, drive sparing, and performance optimization occur automatically.

An IP network is used to connect hosts and applications to storage volumes, and also to connect arrays to each other, providing a communication mechanism that the arrays use to share configuration data and collaborate during data provisioning and load balancing. With the automated management capabilities of PS Series storage arrays, your storage administrators can provision data on-demand and make configuration changes quickly and easily—without disrupting running applications.

The PS Series storage array is a truly modular storage system. Each array contains redundant hot-swappable components for high availability and is designed for 99.999% reliability. An array does not act individually, but as part of a group of one or more arrays, accessed through a single IP address. Each array is configured with the RAID level of your choice. Performance load balancing enables volume data to be stored where the RAID configuration is optimal. When more capacity is needed, you simply add another member to the group. Capacity and performance scale automatically and linearly. Whether you have one array or many, the group provides a single management view, and the administrative effort remains the same.

Using a PS Series group for drive storage, diverse operating systems and a wide range of applications enables sharing of a reliable and high-performance storage system that can scale from hundreds of gigabytes to more than 100 terabytes. Administrators can access the group through a web browser, network connection, or serial connection. The graphical and command line user interfaces present a unified view of the storage that makes provisioning quick and easy. You can instantly create, expand, and delete volumes. Group storage space can be organized into a single pool or multiple pools for increased control and optimal flexibility. In addition, volume snapshots and replicas can be created on demand or through a schedule, providing online backup and restore capabilities with unmatched performance.

The PS Series of arrays provides for an efficient, self-regulating, tiered architecture. Without administrator intervention, data placement within a storage volume is optimized based upon latency. For applications where data becomes "hot" or most accessed, the PS Series arrays will move those pages of data to the pool member with the lowest latency, and move "cold" pages to arrays with higher latency. The result is a well-balanced, high-performing pool of storage.

Event notification mechanisms—including e-mail, syslog, and SNMP—ensure that any problems in the SAN can be quickly identified and resolved. Automatic controller failover and drive sparing mean that failures can generally be handled without user intervention. Servicing of the system (including replacing drives, controllers, fans, and power supplies) can be performed through the management interface.

For a comprehensive storage solution, Dell also provides host-based utilities that are all-inclusive in the purchase of your EqualLogic array. The Host Integration Tools enable easy point-and-click array initialization and host configuration. In addition, multipath I/O support enables you to create a reliable and high-performance I/O path between servers and PS Series group data, while Auto-Snapshot Manager (VSS provider) enables you to create snapshots that are coordinated with Windows applications.

To provide you with a truly comprehensive system, Dell includes numerous advanced features as standard functionality (no hardware add-ons or software licenses) in every PS Series storage array.

- **Modular hardware:** A PS Series group can easily grow or shrink to accommodate workload changes. Therefore, administrators can purchase only the storage they need when they need it. Future products will fully interoperate with first-generation arrays, protecting your initial investment.

- **Fully-redundant, fault-tolerant storage array.** Each array includes redundant, hot-swappable components—drives, control modules, fans, and power supplies—for a no-single-point-of-failure configuration. Components fail over automatically, without user intervention or disrupting data availability. In addition, data in each array is protected with RAID technology.
- **Support for RAID 10, RAID 5, RAID 6, RAID 6 Accelerated, and RAID 50.** You can choose to configure arrays with the appropriate RAID policy, depending on your capacity and application needs.
- **Support for a variety of drives.** Serial ATA (SATA) and Serial-Attached SCSI (SAS) provide flexibility in capacity and performance to meet your needs.
- **Automatic spare configuration and utilization.** Drive spares are automatically configured and used to replace failed drives. No user intervention is required.
- **Auto-Stat Drive Monitoring System (ADMS).** By continually monitoring drive health within a PS Series storage array or across a PS Series group, ADMS ensures optimal data availability. ADMS automatically scans drives in the background to proactively detect and correct media anomalies.
- **High-performance control modules.** Dual control modules provide support for network interface and control module failover. Nonvolatile write-back caches are mirrored across the control modules to protect data. Each control module has two, three, or four 1-Gigabit Ethernet interfaces or two 10-Gigabit Ethernet interfaces. Some control modules also have a dedicated management port.
- **Simple hardware installation.** Only a single network connection on an array is required for operation. Additional network connections can be added at any time for increased bandwidth and reliability.
- **Support for standard Ethernet networks.** Because PS Series storage arrays use standard Ethernet connections to provide access to storage, there is no need to train administrators in unfamiliar and complex technologies like Fibre Channel. Also, costs are reduced due to the high volumes and intense vendor competition among Ethernet hardware venders.

- **Easy setup and management.** A simple setup utility lets you quickly configure an array on the network and create a PS Series group. In minutes, you have a functioning iSCSI SAN. By automating complex operations like RAID configuration, drive sparing, data provisioning, and load balancing, your storage administrators can effectively manage the SAN.
  - **Graphical- and command-line user interfaces.** Password-protected management interfaces provide a single-system view of the storage. Administrators do not need multiple consoles to perform storage management tasks. Using the Group Manager graphical user interface (GUI), creating and managing volumes and configuring security, networking, and event notification are point-and-click operations. In addition, an equivalent command-line interface (CLI) can be accessed through telnet, SSH, or a serial connection, or can be used in scripts.
  - **Automatic data provisioning.** There is no need for administrators to manually create RAID sets or map data onto drives or individual controllers. Arrays in a group contribute space to a shared pool of storage, from which you create volumes. Each volume has a specific size and access controls. To increase a volume, just specify a new size. The group handles storage allocation and capacity balancing across the drives and arrays.
  - **Dynamic load balancing.** As the workload changes, data and network I/O are automatically load-balanced within and across arrays in the group—with no impact on applications and no user intervention. Thus, "hot spots" can be quickly detected and eliminated.
- **Online and seamless scalability**. Increasing array capacity is as easy as installing additional drives or adding more network connections. You can seamlessly expand overall group capacity adding another array to a PS Series group. In all cases, performance scales automatically as drive data and network I/O are load-balanced across the added resources. Processing power also increases due to the additional controllers and caches. Meanwhile, volumes remain available with no impact on hosts and applications. There is no need to open a server cabinet or reconfigure an operating system. The additional storage space and network bandwidth are immediately available for use.

  More than 1000 TB of storage can be configured in a single group. As the group expands, the management effort remains constant. A group with one array (member) is as easy to manage as a multi-member group.

Different sizes and generations of EqualLogic arrays can join into the same management group or even the same resource pool. Therefore, one does not need to throw away previous investments to fully benefit from new technology and new features.

- **Robust security for both data and management access.** Security between an iSCSI initiator (host) and iSCSI target (volume) can be based on IP address, iSCSI initiator name, or CHAP user name. This eliminates the need to understand complicated security technologies (such as Fibre Channel Switch Zoning or LUN Masking). CHAP authentication can be provided through the PS Series group itself or an external RADIUS server.

  In addition, access to the group for management purposes requires an administrative account and password. Accounts can have either read-write or read-only privileges.

- **Advanced features are standard in all arrays.** A key PS Series design principle is to include advanced functionality in all arrays. The result is a comprehensive solution with built-in intelligence and advanced features. All the features described below are standard on each array with our all-inclusive software package and require no additional software, licenses, or cost.

  - **Cloning.** A clone is an image copy of a volume. Cloning is commonly used in multiple server deployments. For example, a master image of a system can be created and then cloned for each server. Cloning can dramatically reduce overhead when deploying replicated servers, such as blade servers and web servers.
  - **Snapshots.** A snapshot quickly captures a volume's contents at a specific point in time and can be used for backups, testing, and upgrades. Both instant and scheduled snapshots are supported. Snapshots greatly simplify and improve the performance of backup and recovery operations. Consistency groups can be created for simultaneous snapshots, maintaining application synchronization across multiple data volumes.
  - **Volume Shadow Copy Service (VSS).** EqualLogic arrays are integrated with the Microsoft VSS framework, which is included with Windows Server. This feature enables turnkey snapshot backups that can offload the backup process from application servers.
  - **Virtual Drive Service (VDS).** The EqualLogic VDS provider enables you to use Microsoft Storage Manager for SANs to create and manage volumes in a PS Series group.
  - **Replication.** Using two PS Series groups, you can replicate volumes across unlimited distances to protect your data. Replication enables you to set up a simple, yet robust disaster recovery plan that guards against catastrophic events.

    A replica represents the contents of a volume at a specific point in time and is similar to a snapshot, except that it must be stored separately from the original volume. If the original volume is destroyed, you can recover data by cloning a replica. This creates a new volume containing the volume data that existed at the time the replica was created.
  - **Multipath I/O.** A redundant network path eliminates failure points between hosts and storage and improves availability. For high performance, you can load balance I/O across multiple ports (HBAs or NICs).
  - **SAN Boot.** Booting servers directly from the SAN is operationally identical to a traditional boot process, but can be accomplished easily and efficiently across hundreds of servers.
  - **Storage Pools.** With PS Series storage, you can divide SAN space into multiple storage pools in a single PS Series group to build an efficient, flexible, easy-to-manage networked storage environment. Pools can be used for segregation or tiering of data online.
  - **Tiered Storage Pools.** Automatically, without administrator intervention, each pool of storage will balance and spread data across a pool of storage or arrays, providing for linear scaling of capacity and performance. The system automatically swaps hot data with cold data between the arrays. This ensures that IO bottlenecks are avoided and both the IO performance and the capacity capabilities of different tiers or different generations of hardware can be automatically applied to maximize the ability of the solution to support application needs.

- **Wide-spread interoperability.** PS Series storage arrays are ideal for heterogeneous environments, with support for most major operating systems and cluster software.

Windows and Exchange integration

Windows integration for Dell EqualLogic PS Series SANs is provided at several levels. VDS and VSS providers are included as part of the Host Integration Tools (HIT) and provide integration with Windows® file systems and compatible backup tools. Also provided is an MPIO Device Specific Module (DSM). This provides connection awareness of the PS Series SAN to Windows hosts, simplifying configuration, enabling reliable network connections, and enhancing performance.

Auto Snapshot Manager / Microsoft Edition (ASM/ME) is an all-inclusive software package that ships with all Dell EqualLogic PS Series arrays to facilitate the deployment, ongoing management, and protection of Dell EqualLogic iSCSI SANs in your Microsoft® Windows environments. ASM/ME has specific integration with Microsoft Exchange including Exchange Server 2010.

ASM/ME leverages VSS to enable the creation of application-consistent "Smart Copies". These allow clean capture and recovery of email database information from Exchange 2010 using space-efficient EqualLogic snapshots. These point-in-time copies essentially preserve the state of the database at a specific moment. Incremental storage is then used to store changes to the database. This is more space efficient than having to keep a complete database copy online, as is required with lagged logs. ASM also allows the creation of flexible schedules for the capture of Smart Copies. Database integrity checks can be configured to run on a designated server, offloading the verification process. With these capabilities, the administrator can keep a set of point-in-time copies of the email databases, increasing the number of recovery points available. These additional Recovery Points (RPO) enhance recoverability in the event of a database corruption. Point-in-time copies of the email databases can be used for granular mailbox recovery, e-discovery, and Database reseeding. ASM/ME Smart Copy restore functions reduces the Recovery Time (RTO) associated with e-mail and database recoveries as well as database reseeds.

Solution description

The following sections outline the hardware and software environment for a Microsoft Exchange Server 2010 solution intended for small to medium sized organizations that support up to 10,000 users in a mailbox resiliency configuration utilizing the Database Availability Group (DAG) feature. The simulated solution consists of a single DAG (2 copy) solution hosted across 4 Dell R610 servers with storage provided by 4 Dell™ EqualLogic™ PS6000XV storage arrays. Connectivity between servers and storage is via iSCSI protocol and two PowerConnect 6224 switches.

The design features a configuration with no single point of failure among storage or servers and additional redundancy within each component. Each Exchange database is replicated to a different server on a separate array, where it is protected by RAID-50 redundancy. Two of the 16 SAS drives in each array are reserved as hot spares for automatic resynchronization in the case of drive failure. In addition, arrays provide redundant controllers and network interfaces with automatic failover. This is backed by dual quad port Ethernet adapters on each server providing MPIO for performance and reliability.

Further, resiliency within the virtual subdivisions of the storage is achieved by replicating each Exchange database to different pools and volumes within pools. In the event that any pool or volume is lost or destroyed (for example, by administrator error) the Exchange servers will continue serving all users.

As a whole the solution design provides high resiliency at multiple layers of hardware and software.

Hardware and software

The solution hardware environment is described in the following table.

Storage	Drives	Servers	Ethernet connections
4 PS6000XV storage arrays configured into 2 PS Series groups, each containing one data storage pool with one member (for db and log data).	64 15K-RPM 600GB Serial Attached SCSI drives	2 Dell PowerEdge™ R610 Servers, each with Intel Xeon X5620 2.4 GHz quad Core CPU and 12GB memory running MS Windows Server 2008 R2 Enterprise x64 Edition	8 gigabit Ethernet connections: 2 quad port Intel Gigabit Ethernet Adapter (model: VT Quad Port Server Adapter driver V11.4.7.0)

Dell EqualLogic PS Series storage arrays provide active load balancing of storage and connection allocation. With two arrays per pool instantiating multiple volumes, PS Series array controllers will over time select an ideal location for each page of data, optimizing performance characteristics. This feature provides balanced demand of array resources providing highest throughput and lowest latencies with no administrative intervention. Additionally, network traffic is balanced across all 8 server network interfaces through Host Integration Tools multipathing, applying intelligent automatic load-balancing to server resources as well.

PS Series topology for Exchange storage solution

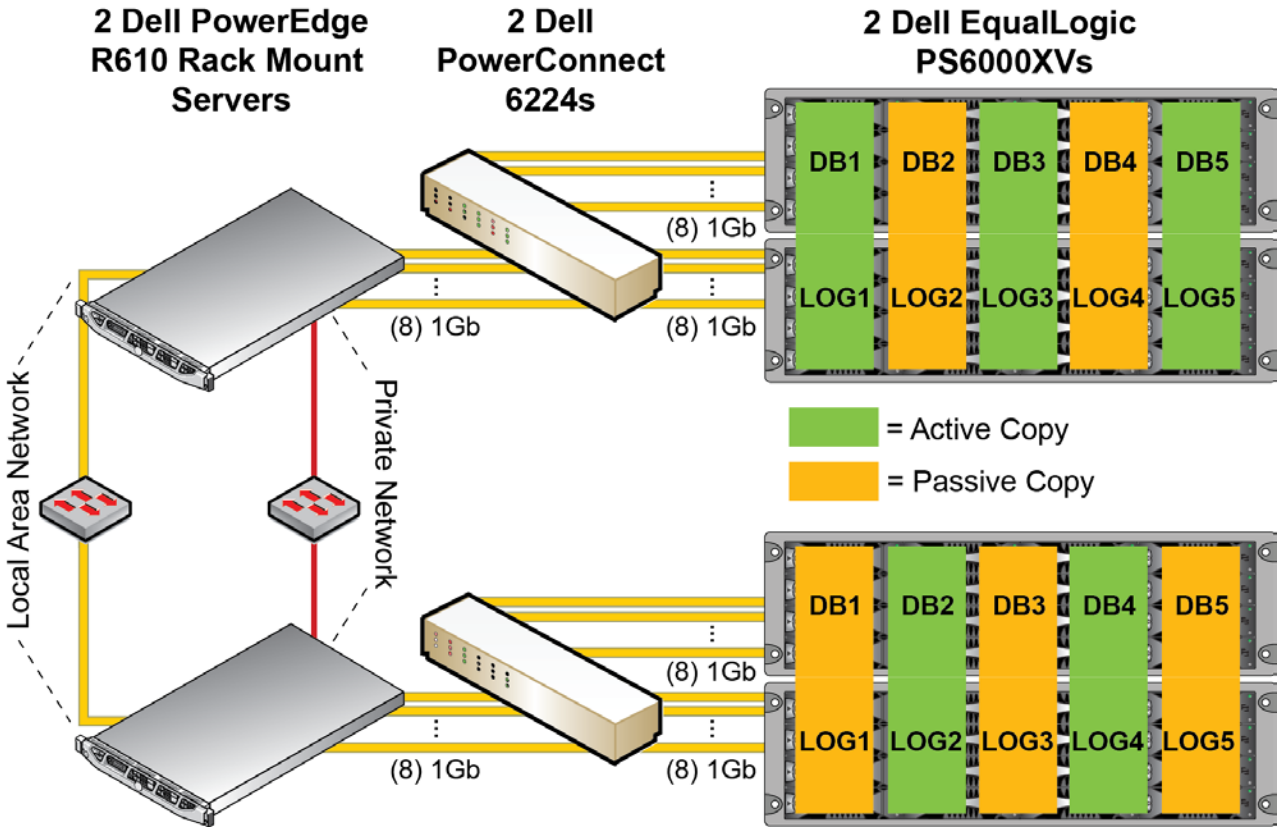

Exchange DAG architecture

This mailbox resiliency solution utilizes two PS Series array groups consisting of two Dell EqualLogic PS6000XV storage arrays per group. There is one data storage pool per group, each consisting of two PS6000XV arrays configured as RAID-50. Within each storage pool five volumes are created for a total of 10 volumes across both pools. Each volume spans each group's storage arrays in its pool, and each volume provides database and log storage for a single Microsoft Exchange server.

There are two simulated Microsoft Exchange servers in the solution. Each Exchange server uses five databases and five log folders configured in five separate volumes. Within the overall solution, one server provides three active and two passive database copies, while the second server provides 2 active and 3 passive database copies. There are a total of 5 active database copies spread across all Exchange servers and PS Series arrays. This design ensures that a fully functioning implementation shares load across all servers and storage, Yet it is capable of providing access to all mail databases and meeting service level agreements in the case that half of all server and/or storage has failed.

The following table and diagram provides details of the solutions data layout on PS Group volumes and maps the volumes to the servers in the DAG.

Table 1: Solution data layout

Server	PS Group	Pool	Volume	DB and Log Group	Active/Passive
Server 1	PS Group 1 consisting of 2 PS6000XV arrays	Data Pool 1 consisting of 2 PS6000XV arrays	Volume 1	DB/Log 1	Active (copy 1)
			Volume 2	DB/Log 2	Passive (copy 2)
			Volume 3	DB/Log 3	Active (copy 1)
			Volume 4	DB/Log 4	Passive (copy 2)
			Volume 5	DB/Log 5	Active (copy 1)
Server 2	PS Group 2 consisting of 3 PS6000XV arrays	Data Pool 2 consisting of 2 PS6000XV arrays	Volume 6	DB/Logs 1	Passive (copy 2)
			Volume 7	DB/Logs 2	Active (copy 1)
			Volume 8	DB/Logs 3	Passive (copy 2)
			Volume 9	DB/Logs 4	Active (copy 1)
			Volume 10	DB/Logs 5	Passive (copy 2)

By design the data layout of the solution provides redundancy at the pool and volume layer. Therefore, if any pool or volume is lost or destroyed (e.g. by administrator error), the solution maintains a copy of all databases and the mailbox servers would continue servicing Exchange users.

PS Series solution for Exchange Server data configuration

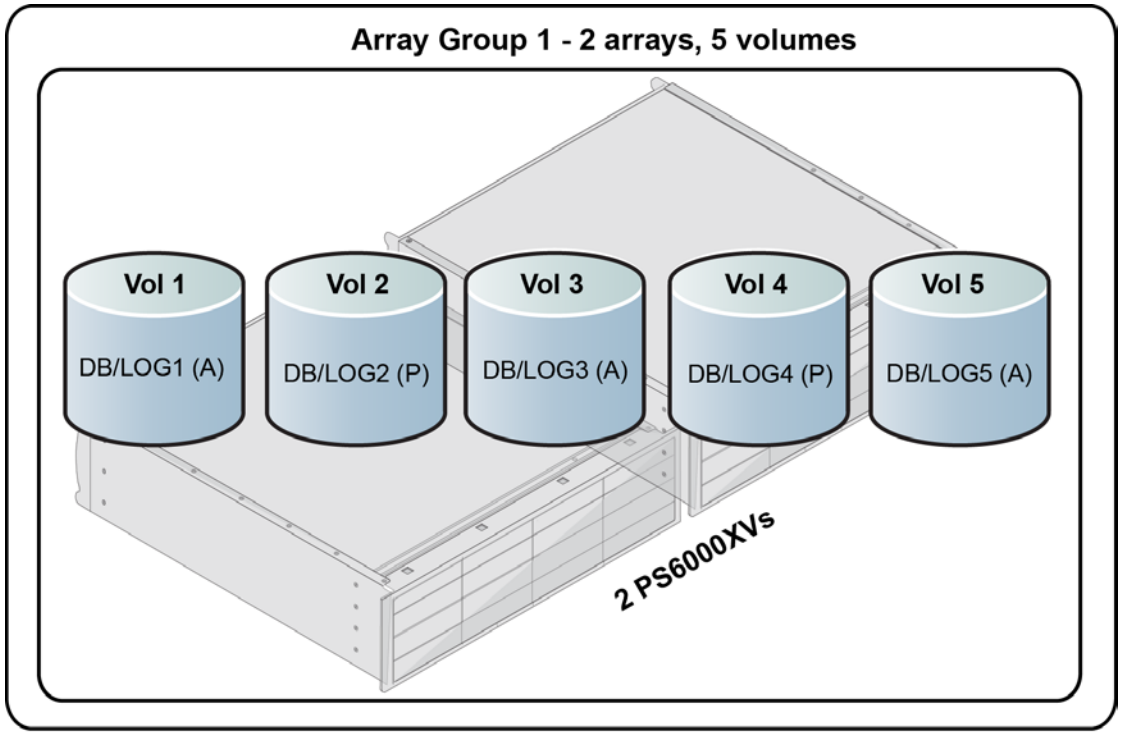

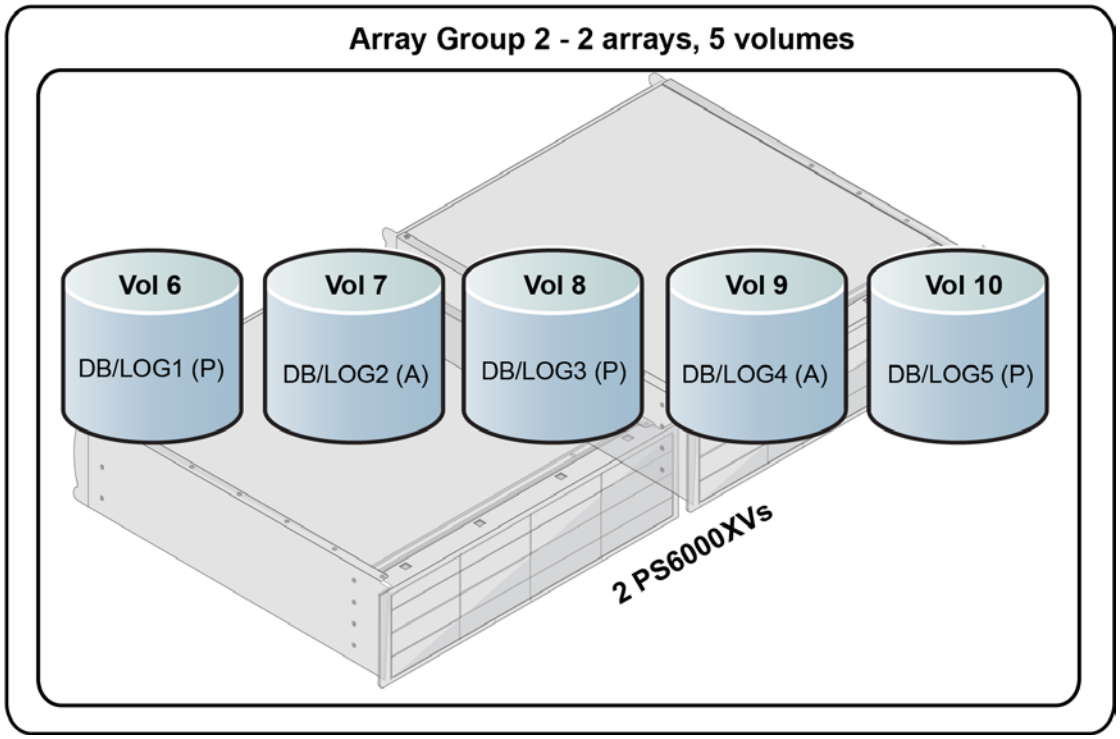

The ESRP-Storage program focuses on storage solution testing to address performance and reliability issues with storage design. However, storage is not the only factor to take into consideration when designing a scale-up Exchange solution. Other factors which affect the server scalability are:

- Server processor utilization
- Server physical and virtual memory limitations
- Resource requirements for other applications
- Directory and network service latencies
- Network infrastructure limitations
- Replication and recovery requirements
- Client usage profiles

All these factors are beyond the scope for ESRP-Storage. Therefore, the number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployment.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please see Troubleshooting Microsoft Exchange Server Performance, available at:

http://go.microsoft.com/fwlink/?LinkId=23454

Targeted customer profile

The PS Series storage solution is intended for small, medium, and large Microsoft Exchange Server 2010 organizations that want reliable, high-performance, and easy-to-manage drive storage. The tested configuration can support the following:

- 2 (1 tested) Exchange 2010 servers
- 10,000 user mailboxes
- 0.25 I/O per second per user (0.3 tested for 20% headroom)
- 1024 MB mailbox quota per user
- 5 databases per server
- 2 TB database size
- Mailbox Resiliency (2 copies) provides high availability and is the primary data protection mechanism.

Tested deployment

Since primary (active) and secondary (passive) storage in the solution is identical, the tested deployment configured primary storage only, as allowed in the ESRP Storage v3.0 program. The following tables summarize the testing environment.

Simulated Exchange configuration

Number of Exchange mailboxes simulated	10,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	2
Number of active mailboxes/server	10,000
Number of databases/host	5
Number of copies/database	2
Number of mailboxes/database	2000
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.3
Database and Log LUN size	11.92 TB (2441.61 GB x 5 LUNs)
Total database size for performance testing	10TB (2TB x 5 databases)
% storage capacity used by Exchange database[2]	83.8% (10TB/11.92TB)

Primary storage hardware

Storage Connectivity	iSCSI
Storage model and OS/firmware revision	Dell EqualLogic PS6000XV Firmware Rev: V5.0.2 http://www.dell.com/products/view.aspx?id=2509
Storage cache	8GB (2-PS6000s * 4GB = 8GB)
Number of storage controllers	2 (2 per array in active/passive configuration)
Number of storage ports	8 (4 per controller/8 per array)
Maximum bandwidth of storage connectivity to host	4 x 1 GB Ethernet per array
Switch type/model/firmware revision	PowerConnect 6224 Switch (GB Ethernet) Firmware Rev: 3.2.0.10
HBA model and firmware	Intel Gigabit VT Quad Port Server Adapter 11.4.7.0
Number of HBAs/host	2
Host server type	Dell PowerEdge R610 Intel Xeon Dual Core X5620 2.44 GHz with 12GB RAM
Total number of drives tested	32
Maximum number of spindles	32

[2] Storage performance characteristics change based on the percentage utilization of the individual disks. Tests that use a small percentage of the storage (~25%) might exhibit reduced throughput if the storage capacity utilization is significantly increased beyond what is tested in this paper.

Primary storage software

HBA driver	Intel Driver Version: 11.4.7.0
HBA QueueTarget Setting	N/A
HBA QueueDepth Setting	N/A
Multipath I/O DSM	Dell EqualLogic HIT (Host Integration Toolkit) 3.5.1
Host OS	Windows Server 2008 R2, Enterprise x64 Edition
ESE.dll file version	14.0.639.19
Replication solution name/version	N/A

Primary drive configuration (Mailbox and Log Store Drives)

Drive type, speed and firmware revision	Seagate Cheetah 15K.6 600G 15,000 rpm SAS drives Model: ST3600057SS Firmware Revision: EN00
Raw capacity per drive (GB)	558.91GB
Number of physical drives in test	32
Total raw storage capacity (GB)	17.46 TB (17885.12GB)
Drive slice size (GB)	N/A
Number of disks per LUN	Up to 32 (automatically allocated based on load)
Raid level	All storage pools configured as RAID 50
Total formatted capacity (GB)	11.92TB (12208.05GB)
Storage capacity utilization	57.27% (10TB/17.46TB)
Database capacity utilization	83.8% (10TB/11.92TB)

Best practices

Microsoft Exchange Server is a drive-intensive application. Based on the tests using the ESRP framework, Dell recommends the following best practices to help improve storage performance.

- Allow the PS Series group to automatically balance the load across arrays, caches, and network links. Automatic load balancing reduces administrator effort as Exchange workloads change over time.
- In large Exchange deployments, isolate the Exchange workload from other application workloads by creating separate storage pools for Exchange-related volumes in a PS Series group and setting up separate servers for Exchange and other applications.
- Windows NTFS allocation unit size for partitions containing Exchange 2010 databases should be set to 64k for best performance.
- Depending on the desired level of availability, you can configure Exchange using multiple storage pools in a PS Series group to provide complete resource and hardware isolation between logs and databases. You can also deploy Exchange using a single pool, which provides a high level of availability and makes provisioning simple. Performance and reliability are similar in either a single pool group or multiple pool group.
- Size and configure first for I/O performance, then for storage capacity.

- Enable Dell EqualLogic Host Integration Tools V3.5.1 on Exchange servers to ensure highly-available SAN connections with Microsoft's MPIO.
- Use Microsoft iSCSI software initiators in Exchange configurations. In these tests, the Microsoft iSCSI software initiator was used.
- Use separate volumes for Exchange databases and transaction logs to improve backup and recovery operations.
- Place SAN infrastructure on VLANs or subnets that differ from other production network traffic.
- Use non-blocking Gigabit Ethernet switches.
- Enable use of Jumbo Frames on the Intel Quadport ET devices
- Set the nic transmit/receive buffers to their maximum setting

For additional best practices on storage design in Exchange 2010, see the URL:

http://technet.microsoft.com/en-us/library/bb124518.aspx

Additional information

For more information, see the Dell website (www.dell.com). In addition, Dell EqualLogic technical documents may be useful:

http://support.dell.com/equallogic

Test result summary

This section provides a high-level summary of the test data from ESRP and the link to the detailed html reports which are generated by the ESRP testing framework. See Appendix A for detailed information about test results.

Reliability results

A number of tests in the framework check reliability, running for 24 hours. The goal is to verify the storage can handle high I/O load for a long period of time. Both log and database files are analyzed for integrity after the stress test to ensure no database or log corruption.

The following list provides an overview:

- No relevant errors were reported in the event log for the storage reliability test.
- No errors were reported by the database and log checksum process.

Storage performance results

The Primary Storage performance testing is designed to exercise the storage with maximum sustainable Exchange I/O for over two hours. The test shows how long it takes for the storage to respond to an I/O under load. The data below is the sum of all of the logical drive I/Os and the average of all the logical drives' I/O latency during the test (which was run for six hours). Each server is listed separately and the aggregate numbers across all servers are also presented.

Individual server metrics

The server metrics include the sum of I/Os across storage groups and the average latency across all storage groups on a per server basis.

Aggregate performance across all servers metrics

Aggregate performance is the sum of I/Os across all servers and the average latency across all servers.

Database I/O	
Database Drive Transfers/sec	3250.45
Database Drive Reads/sec	2081.74
Database Drive Writes/sec	1168.71
Average Database Drive Read Latency (ms)	17.88
Average Database Drive Write Latency (ms)	4.58
Transaction Log I/O	
Log Drive Writes/sec	521.45
Average Log Drive Write Latency (ms)	2.08

Database Backup/Recovery performance

This section includes two tests. The first test measures sequential read rates of the database files. The second test measures the recovery/replay performance.

Database Read-Only performance

The test measures the maximum rate to back up databases using VSS. The following table shows the average rate for a single database file.

MB read/sec per database	52.08 (Average)
MB read/sec total per server	260.42 (Sum)

Transaction Log Recovery/Replay performance

The following table shows the average rate for ~500 log files played in a single storage group. Each log file is 1 MB in size.

Average time to play one Log File (sec)	1.29 (avg. resp. to replay log / avg. # of logs replayed)

Conclusion

This document was developed by Dell Inc., and reviewed by the Microsoft Exchange Product team. The test results and data presented in this document are based on the tests introduced in the ESRP test framework. Customers should not quote the data directly for their pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

The ESRP Storage program is not designed to be a benchmarking program. Its tests are not designed for achieving the maximum throughput for a given solution. Rather, they are focused on producing recommendations from vendors for the Exchange application. Therefore, the data presented in this document should not be used for direct comparisons among the solutions.

Appendix A: Stress Testing

Microsoft Exchange ***JetStress 2010***

Stress Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	WIN-J6FRBTN5B34
Test Description	10000 mailboxes 1024MB mailbox size .3 mailbox profile 5 instances 64k CC FlowControl Jumbo Frames
Test Start Time	6/28/2011 11:54:49 PM
Test End Time	6/29/2011 11:55:59 PM
Collection Start Time	6/28/2011 11:55:50 PM
Collection End Time	6/29/2011 11:55:48 PM
Jetstress Version	14.01.0225.017
ESE Version	14.01.0218.012
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run4_24hr_16_thread\Performance_2011_6_28_23_55_0.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	3104.642
Target Transactional I/O per Second	3000
Initial Database Size (bytes)	10749040984064
Final Database Size (bytes)	10819438182400
Database Files (Count)	5

Jetstress System Parameters

Thread Count	16 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%

Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance1516.1 Log path: C:\dbs\db1\log
Database: C:\dbs\db1\db\Jetstress001001.edb

Instance1516.2 Log path: C:\dbs\db2\log
Database: C:\dbs\db2\db\Jetstress002001.edb

Instance1516.3 Log path: C:\dbs\db3\log
Database: C:\dbs\db3\db\Jetstress003001.edb

Instance1516.4 Log path: C:\dbs\db4\log
Database: C:\dbs\db4\db\Jetstress004001.edb

Instance1516.5 Log path: C:\dbs\db5\log
Database: C:\dbs\db5\db\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1516.1	19.360	4.251	396.183	223.733	32931.084	34321.751	0.000	2.028	0.000	100.998	0.000	5711.222
Instance1516.2	19.356	4.191	396.641	224.125	32935.959	34325.027	0.000	2.023	0.000	101.039	0.000	5718.673
Instance1516.3	19.363	4.148	396.039	223.711	32937.288	34325.797	0.000	2.014	0.000	101.054	0.000	5720.941
Instance1516.4	19.215	3.892	398.951	225.278	32937.528	34320.186	0.000	2.036	0.000	101.280	0.000	5716.542
Instance1516.5	19.263	3.676	396.255	223.726	32936.234	34324.290	0.000	2.028	0.000	100.995	0.000	5705.573

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1516.1	18.029	261948.184
Instance1516.2	18.176	261945.406
Instance1516.3	18.218	261945.025
Instance1516.4	19.336	261957.502
Instance1516.5	19.375	261952.615

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1516.1	2.391	232559.259
Instance1516.2	2.396	232559.510
Instance1516.3	2.397	232556.121
Instance1516.4	2.400	232562.172
Instance1516.5	2.389	232557.710

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1516.1	19.360	4.251	414.212	223.733	42899.075	34321.751	5.366	2.028	2.391	100.998	232559.259	5711.222
Instance1516.2	19.356	4.191	414.816	224.125	42970.233	34325.027	5.259	2.023	2.396	101.039	232559.510	5718.673
Instance1516.3	19.363	4.148	414.257	223.711	43008.405	34325.797	5.230	2.014	2.397	101.054	232556.121	5720.941
Instance1516.4	19.215	3.892	418.287	225.278	43524.398	34320.186	5.235	2.036	2.400	101.280	232562.172	5716.542
Instance1516.5	19.263	3.676	415.629	223.726	43611.831	34324.290	5.283	2.028	2.389	100.995	232557.710	5705.573

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	3.802	2.161	7.113
Available MBytes	13623.782	13566.000	13730.000
Free System Page Table Entries	33555609.422	33555466.000	33556122.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	345317845.274	339095552.000	347959296.000
Pool Paged Bytes	117004095.733	93835264.000	130060288.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log6/28/2011 11:54:49 PM -- Jetstress testing begins ...
6/28/2011 11:54:49 PM -- Preparing for testing ...
6/28/2011 11:54:54 PM -- Attaching databases ...
6/28/2011 11:54:54 PM -- Preparations for testing are complete.
6/28/2011 11:54:54 PM -- Starting transaction dispatch ..
6/28/2011 11:54:54 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
6/28/2011 11:54:54 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/28/2011 11:55:00 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/28/2011 11:55:00 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/28/2011 11:55:06 PM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/28/2011 11:55:06 PM -- Performance logging started (interval: 15000 ms).
6/28/2011 11:55:06 PM -- Attaining prerequisites:
6/28/2011 11:55:50 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1225318000.0 (lower bound: 1207960000.0, upper bound: none)

6/29/2011 11:55:50 PM -- Performance logging has ended.
6/29/2011 11:55:50 PM -- JetInterop batch transaction stats: 978458, 979145, 979485, 980894 and 979906.
6/29/2011 11:55:51 PM -- Dispatching transactions ends.
6/29/2011 11:55:51 PM -- Shutting down databases ...
6/29/2011 11:55:59 PM -- Instance1516.1 (complete), Instance1516.2 (complete), Instance1516.3 (complete), Instance1516.4 (complete) and Instance1516.5 (complete)
6/29/2011 11:55:59 PM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run4_24hr_16_thread\Performance_2011_6_28_23_55_0.blg has 5754 samples.
6/29/2011 11:55:59 PM -- Creating test report ...
6/29/2011 11:56:27 PM -- Instance1516.1 has 19.4 for I/O Database Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.1 has 2.0 for I/O Log Writes Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.1 has 2.0 for I/O Log Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.2 has 19.4 for I/O Database Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.2 has 2.0 for I/O Log Writes Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.2 has 2.0 for I/O Log Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.3 has 19.4 for I/O Database Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.3 has 2.0 for I/O Log Writes Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.3 has 2.0 for I/O Log Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.4 has 19.2 for I/O Database Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.4 has 2.0 for I/O Log Writes Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.4 has 2.0 for I/O Log Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.5 has 19.3 for I/O Database Reads Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.5 has 2.0 for I/O Log Writes Average Latency.
6/29/2011 11:56:27 PM -- Instance1516.5 has 2.0 for I/O Log Reads Average Latency.
6/29/2011 11:56:27 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/29/2011 11:56:27 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/29/2011 11:56:27 PM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run4_24hr_16_thread\Performance_2011_6_28_23_55_0.xml has 5751 samples queried.

Appendix B: Performance Testing

Microsoft Exchange ***Jetstress 2010***

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	WIN-J6FRBTN5B34
Test Description	10000 mailboxes 1024MB mailbox size .3 mailbox profile 5 instances 64k CC FlowControl Jumbo Frames
Test Start Time	6/28/2011 12:38:54 PM
Test End Time	6/28/2011 2:40:02 PM
Collection Start Time	6/28/2011 12:39:56 PM
Collection End Time	6/28/2011 2:39:53 PM
Jetstress Version	14.01.0225.017
ESE Version	14.01.0218.012
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run2_2hr_16_thread\Performance_2011_6_28_12_39_5.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	3250.456
Target Transactional I/O per Second	3000
Initial Database Size (bytes)	107427663 05280
Final Database Size (bytes)	107490409 84064
Database Files (Count)	5

Jetstress System Parameters

Thread Count	16 (per database)
Minimum Database Cache	160.0 MB

Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance2772.1 Log path: C:\dbs\db1\log
Database: C:\dbs\db1\db\Jetstress001001.edb

Instance2772.2 Log path: C:\dbs\db2\log
Database: C:\dbs\db2\db\Jetstress002001.edb

Instance2772.3 Log path: C:\dbs\db3\log
Database: C:\dbs\db3\db\Jetstress003001.edb

Instance2772.4 Log path: C:\dbs\db4\log
Database: C:\dbs\db4\db\Jetstress004001.edb

Instance2772.5 Log path: C:\dbs\db5\log
Database: C:\dbs\db5\db\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance2772.1	17.905	5.221	416.686	233.900	32972.086	34752.156	0.000	2.084	0.000	104.241	0.000	5900.264
Instance2772.2	17.877	4.767	413.497	232.300	32962.616	34752.572	0.000	2.099	0.000	104.013	0.000	5944.864
Instance2772.3	17.857	4.441	416.026	233.622	32989.843	34744.992	0.000	2.094	0.000	104.035	0.000	5943.064
Instance2772.4	17.885	4.435	417.731	234.450	32971.435	34751.359	0.000	2.085	0.000	104.737	0.000	5935.299
Instance2772.5	17.898	4.059	417.805	234.439	32954.097	34740.736	0.000	2.078	0.000	104.430	0.000	5891.654

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2772.1	19.732	261965.598
Instance2772.2	19.795	261970.354

Instance2772.3	19.709	261917.445
Instance2772.4	21.487	261940.355
Instance2772.5	21.415	261927.899

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2772.1	2.554	232562.158
Instance2772.2	2.569	232565.690
Instance2772.3	2.568	232563.175
Instance2772.4	2.581	232553.649
Instance2772.5	2.555	232559.277

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance2772.1	17.905	5.221	436.417	233.900	43325.455	34752.156	5.414	2.084	2.554	104.241	232562.158	5900.264
Instance2772.2	17.877	4.767	433.292	232.300	43424.861	34752.572	5.302	2.099	2.569	104.013	232565.690	5944.864
Instance2772.3	17.857	4.441	435.735	233.622	43344.386	34744.992	5.396	2.094	2.568	104.035	232563.175	5943.064
Instance2772.4	17.885	4.435	439.218	234.450	44172.721	34751.359	5.332	2.085	2.581	104.737	232553.649	5935.299
Instance2772.5	17.898	4.059	439.220	234.439	44118.124	34740.736	5.478	2.078	2.555	104.430	232559.277	5891.654

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	3.130	2.258	5.055
Available MBytes	13624.067	13612.000	13733.000
Free System Page Table Entries	33555436.521	33555026.000	33555610.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	343676245.333	338604032.000	345677824.000
Pool Paged Bytes	121374028.800	121147392.000	121466880.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log6/28/2011 12:38:54 PM -- Jetstress testing begins ...
6/28/2011 12:38:54 PM -- Preparing for testing ...
6/28/2011 12:39:00 PM -- Attaching databases ...
6/28/2011 12:39:00 PM -- Preparations for testing are complete.
6/28/2011 12:39:00 PM -- Starting transaction dispatch ..
6/28/2011 12:39:00 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)

6/28/2011 12:39:00 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/28/2011 12:39:05 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/28/2011 12:39:05 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/28/2011 12:39:12 PM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/28/2011 12:39:12 PM -- Performance logging started (interval: 15000 ms).
6/28/2011 12:39:12 PM -- Attaining prerequisites:
6/28/2011 12:39:56 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1220436000.0 (lower bound: 1207960000.0, upper bound: none)
6/28/2011 2:39:56 PM -- Performance logging has ended.
6/28/2011 2:39:56 PM -- JetInterop batch transaction stats: 87090, 87413, 87653, 87635 and 87854.
6/28/2011 2:39:57 PM -- Dispatching transactions ends.
6/28/2011 2:39:57 PM -- Shutting down databases ...
6/28/2011 2:40:02 PM -- Instance2772.1 (complete), Instance2772.2 (complete), Instance2772.3 (complete), Instance2772.4 (complete) and Instance2772.5 (complete)
6/28/2011 2:40:02 PM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run2_2hr_16_thread\Performance_2011_6_28_12_39_5.blg has 482 samples.
6/28/2011 2:40:02 PM -- Creating test report ...
6/28/2011 2:40:04 PM -- Instance2772.1 has 17.9 for I/O Database Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.1 has 2.1 for I/O Log Writes Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.1 has 2.1 for I/O Log Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.2 has 17.9 for I/O Database Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.2 has 2.1 for I/O Log Writes Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.2 has 2.1 for I/O Log Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.3 has 17.9 for I/O Database Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.3 has 2.1 for I/O Log Writes Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.3 has 2.1 for I/O Log Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.4 has 17.9 for I/O Database Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.4 has 2.1 for I/O Log Writes Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.4 has 2.1 for I/O Log Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.5 has 17.9 for I/O Database Reads Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.5 has 2.1 for I/O Log Writes Average Latency.
6/28/2011 2:40:04 PM -- Instance2772.5 has 2.1 for I/O Log Reads Average Latency.
6/28/2011 2:40:04 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/28/2011 2:40:04 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/28/2011 2:40:04 PM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run2_2hr_16_thread\Performance_2011_6_28_12_39_5.xml has 479 samples queried.

Appendix C Backup Testing

Microsoft Exchange ***Jetstress 2010***

Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance1088.1	2063908.09	11:03:56	51.81
Instance1088.2	2063932.09	11:04:27	51.77
Instance1088.3	2063932.09	11:00:18	52.09
Instance1088.4	2063972.09	11:19:20	50.64
Instance1088.5	2063900.09	10:35:43	54.11

Jetstress System Parameters

Thread Count	16 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance1088.1 Log path: C:\dbs\db1\log
Database: C:\dbs\db1\db\Jetstress001001.edb

Instance1088.2 Log path: C:\dbs\db2\log
Database: C:\dbs\db2\db\Jetstress002001.edb

Instance1088.3 Log path: C:\dbs\db3\log
Database: C:\dbs\db3\db\Jetstress003001.edb

Instance1088.4 Log path: C:\dbs\db4\log
Database: C:\dbs\db4\db\Jetstress004001.edb

Instance1088.5 Log path: C:\dbs\db5\log
Database: C:\dbs\db5\db\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1088.1	7.976	0.000	207.232	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance1088.2	7.956	0.000	207.041	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance1088.3	7.908	0.000	208.394	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance1088.4	8.180	0.000	202.528	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance1088.5	7.606	0.000	216.448	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	2.411	1.061	3.230
Available MBytes	14725.206	14709.000	14750.000
Free System Page Table Entries	33555857.954	33555450.000	33556106.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	215318555.166	214515712.000	217309184.000
Pool Paged Bytes	162100362.093	157868032.000	164335616.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log6/30/2011 6:27:04 PM -- Jetstress testing begins ...
6/30/2011 6:27:04 PM -- Preparing for testing ...
6/30/2011 6:27:09 PM -- Attaching databases ...
6/30/2011 6:27:09 PM -- Preparations for testing are complete.
6/30/2011 6:27:18 PM -- Performance logging started (interval: 30000 ms).
6/30/2011 6:27:18 PM -- Backing up databases ...
7/1/2011 5:46:39 AM -- Performance logging has ended.
7/1/2011 5:46:39 AM -- Instance1088.1 (100% processed), Instance1088.2 (100% processed), Instance1088.3 (100% processed), Instance1088.4 (100% processed) and Instance1088.5 (100% processed)
7/1/2011 5:46:39 AM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run6_db_backup\DatabaseBackup_2011_6_30_18_27_9.blg has 1357 samples.
7/1/2011 5:46:39 AM -- Creating test report ...

Appendix D Soft Recovery Testing

Microsoft Exchange ***Jetstress 2010***

SoftRecovery Test Result Report

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance1088.1	504	661.4099617
Instance1088.2	504	647.6195374
Instance1088.3	501	650.5367426
Instance1088.4	516	651.5975444
Instance1088.5	511	663.5315654

Database Configuration

Instance1088.1 Log path: C:\dbs\db1\log
Database: C:\dbs\db1\db\Jetstress001001.edb

Instance1088.2 Log path: C:\dbs\db2\log
Database: C:\dbs\db2\db\Jetstress002001.edb

Instance1088.3 Log path: C:\dbs\db3\log
Database: C:\dbs\db3\db\Jetstress003001.edb

Instance1088.4 Log path: C:\dbs\db4\log
Database: C:\dbs\db4\db\Jetstress004001.edb

Instance1088.5 Log path: C:\dbs\db5\log
Database: C:\dbs\db5\db\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1088.1	9.523	1.086	688.285	4.582	40731.722	32768.000	10.545	0.000	6.873	0.000	232492.851	0.000
Instance1088.2	9.490	0.934	696.522	4.679	40572.554	32768.000	10.156	0.000	7.018	0.000	232534.355	0.000
Instance1088.3	9.510	1.048	689.405	4.624	40603.397	32768.000	10.411	0.000	6.932	0.000	232555.490	0.000
Instance1088.4	9.288	1.059	696.392	4.752	40695.795	32768.000	10.620	0.000	7.128	0.000	232526.023	0.000
Instance1088.5	9.246	1.034	675.316	4.619	40836.480	32768.000	10.881	0.000	6.928	0.000	232504.507	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==>	Database Maintenance IO	Database Maintenance IO Reads

Instances	Reads/sec	Average Bytes
Instance1088.1	28.375	262031.412
Instance1088.2	28.401	261925.190
Instance1088.3	28.326	262014.657
Instance1088.4	29.068	261979.526
Instance1088.5	29.116	262025.186

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1088.1	9.523	1.086	716.660	4.582	49493.678	32768.000	10.545	0.000	6.873	0.000	232492.851	0.000
Instance1088.2	9.490	0.934	724.923	4.679	49244.722	32768.000	10.156	0.000	7.018	0.000	232534.355	0.000
Instance1088.3	9.510	1.048	717.731	4.624	49341.649	32768.000	10.411	0.000	6.932	0.000	232555.490	0.000
Instance1088.4	9.288	1.059	725.460	4.752	49562.375	32768.000	10.620	0.000	7.128	0.000	232526.023	0.000
Instance1088.5	9.246	1.034	704.431	4.619	49978.677	32768.000	10.881	0.000	6.928	0.000	232504.507	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	3.230	1.816	5.935
Available MBytes	13346.409	13301.000	14515.000
Free System Page Table Entries	33555568.646	33555474.000	33556066.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	218599573.854	218296320.000	220811264.000
Pool Paged Bytes	165996893.659	165986304.000	166031360.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log7/1/2011 6:45:35 AM -- Jetstress testing begins ...
7/1/2011 6:45:35 AM -- Preparing for testing ...
7/1/2011 6:45:40 AM -- Attaching databases ...
7/1/2011 6:45:40 AM -- Preparations for testing are complete.
7/1/2011 6:45:40 AM -- Starting transaction dispatch ..
7/1/2011 6:45:40 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
7/1/2011 6:45:40 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
7/1/2011 6:45:46 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
7/1/2011 6:45:46 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
7/1/2011 6:45:50 AM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
7/1/2011 6:45:50 AM -- Performance logging started (interval: 15000 ms).
7/1/2011 6:45:50 AM -- Generating log files ...
7/1/2011 7:16:21 AM -- C:\dbs\db1\log (100.8% generated), C:\dbs\db2\log (100.8% generated), C:\dbs\db3\log (100.2% generated), C:\dbs\db4\log (103.0% generated) and C:\dbs\db5\log (102.2% generated)

7/1/2011 7:16:21 AM -- Performance logging has ended.
7/1/2011 7:16:21 AM -- JetInterop batch transaction stats: 21636, 21571, 21424, 21830 and 21737.
7/1/2011 7:16:22 AM -- Dispatching transactions ends.
7/1/2011 7:16:22 AM -- Shutting down databases ...
7/1/2011 7:16:25 AM -- Instance1088.1 (complete), Instance1088.2 (complete), Instance1088.3 (complete), Instance1088.4 (complete) and Instance1088.5 (complete)
7/1/2011 7:16:25 AM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run7_soft_recovery\Performance_2011_7_1_6_45_46.blg has 121 samples.
7/1/2011 7:16:25 AM -- Creating test report ...
7/1/2011 7:16:26 AM -- Instance1088.1 has 17.0 for I/O Database Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.1 has 2.0 for I/O Log Writes Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.1 has 2.0 for I/O Log Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.2 has 17.0 for I/O Database Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.2 has 1.9 for I/O Log Writes Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.2 has 1.9 for I/O Log Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.3 has 16.9 for I/O Database Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.3 has 1.9 for I/O Log Writes Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.3 has 1.9 for I/O Log Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.4 has 16.8 for I/O Database Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.4 has 1.9 for I/O Log Writes Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.4 has 1.9 for I/O Log Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.5 has 16.8 for I/O Database Reads Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.5 has 2.0 for I/O Log Writes Average Latency.
7/1/2011 7:16:26 AM -- Instance1088.5 has 2.0 for I/O Log Reads Average Latency.
7/1/2011 7:16:26 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
7/1/2011 7:16:26 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
7/1/2011 7:16:26 AM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run7_soft_recovery\Performance_2011_7_1_6_45_46.xml has 120 samples queried.
7/1/2011 7:16:26 AM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run7_soft_recovery\Performance_2011_7_1_6_45_46.html was saved.
7/1/2011 7:16:30 AM -- Performance logging started (interval: 4000 ms).
7/1/2011 7:16:30 AM -- Recovering databases ...
7/1/2011 7:27:34 AM -- Performance logging has ended.
7/1/2011 7:27:34 AM -- Instance1088.1 (661.4099617), Instance1088.2 (647.6195374), Instance1088.3 (650.5367426), Instance1088.4 (651.5975444) and Instance1088.5 (663.5315654)
7/1/2011 7:27:34 AM -- C:\JetStress\1024MB_results\10k_tests\5_volumes\build6\run7_soft_recovery\SoftRecovery_2011_7_1_7_16_26.blg has 164 samples.
7/1/2011 7:27:34 AM -- Creating test report ...