

# Dell EqualLogic PS Series

## 5,000-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

Copyright © 2011 Dell Inc. All Rights Reserved.

EqualLogic is a registered trademark of Dell Inc.

Dell is a trademark of Dell Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Information in this document is subject to change without notice.

Reproduction in any manner whatsoever without the written permission of Dell is strictly forbidden.

July 2011

# Table of Contents

Overview.....	4
Disclaimer.....	4
Dell EqualLogic PS Series storage array features.....	5
Windows and Exchange integration.....	8
Solution description.....	8
Hardware and software.....	8
Exchange DAG architecture.....	9
Targeted customer profile.....	12
Tested deployment.....	12
Simulated Exchange configuration.....	12
Primary storage hardware.....	13
Primary storage software.....	13
Primary drive configuration (Mailbox and Log Store Drives).....	14
Best practices.....	14
Additional information.....	15
Test result summary.....	15
Reliability results.....	15
Storage performance results.....	15
Database Backup/Recovery performance.....	16
Conclusion.....	16
Appendix A: Stress Testing.....	17
Stress Test Result Report.....	17
Appendix B: Performance Testing.....	20
Performance Test Result Report.....	20
Appendix C Backup Testing.....	24
Database backup Test Result Report.....	24
Appendix D Soft Recovery Testing.....	26
SoftRecovery Test Result Report.....	26

## 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<sup>1</sup>. 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.

---

<sup>1</sup> 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 vendors.
- **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 5,000 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 2 Dell R610 servers with storage provided by 2 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

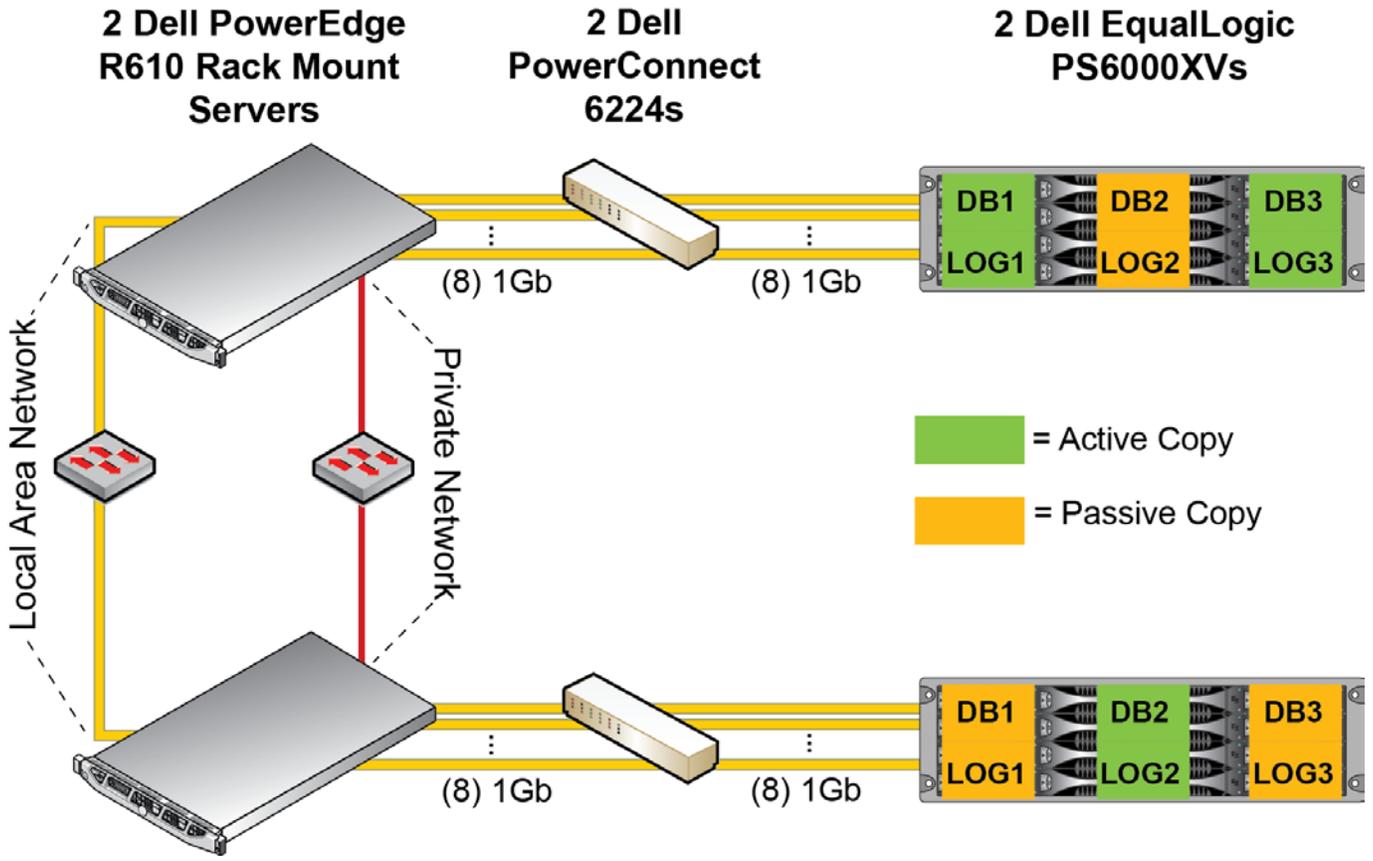
The solution hardware environment is described in the following table.

Storage	Drives	Servers	Ethernet connections
2 PS6000XV storage arrays configured into 2 PS Series groups, each containing one data storage pool with one member (for db and log data).	32 15K-RPM 600GB Serial Attached SCSI drives	2 Dell PowerEdge™ R610 Server, each with Intel Xeon 5620 2.40 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 Adapters (model: VT Quad Port, 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.

This solution used the minimum number of databases to support 5,000 users (3), dividing the workload equally between the three databases.

**PS Series topology for Exchange storage solution**



**Exchange DAG architecture**

This mailbox resiliency solution utilizes two PS Series array groups consisting of 1 Dell EqualLogic PS6000XV storage array per group. There is one data storage pool per group, each consisting of 1 PS6000XV array configured as RAID-50. Within each storage pool three volumes are created for a total of 6 volumes across both pools. Each volume spans each group’s storage array in its pool, and each volume provides database and log storage for a single Microsoft Exchange server.

There are 2 simulated Microsoft Exchange servers in the solution. Each Exchange server uses 3 databases and 3 log folders configured in three separate volumes. Within the overall solution, 1 server provides 2 active and 1 passive database copies, while the second server provides 1 active and 2 passive database copies. There are a total of 3 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.

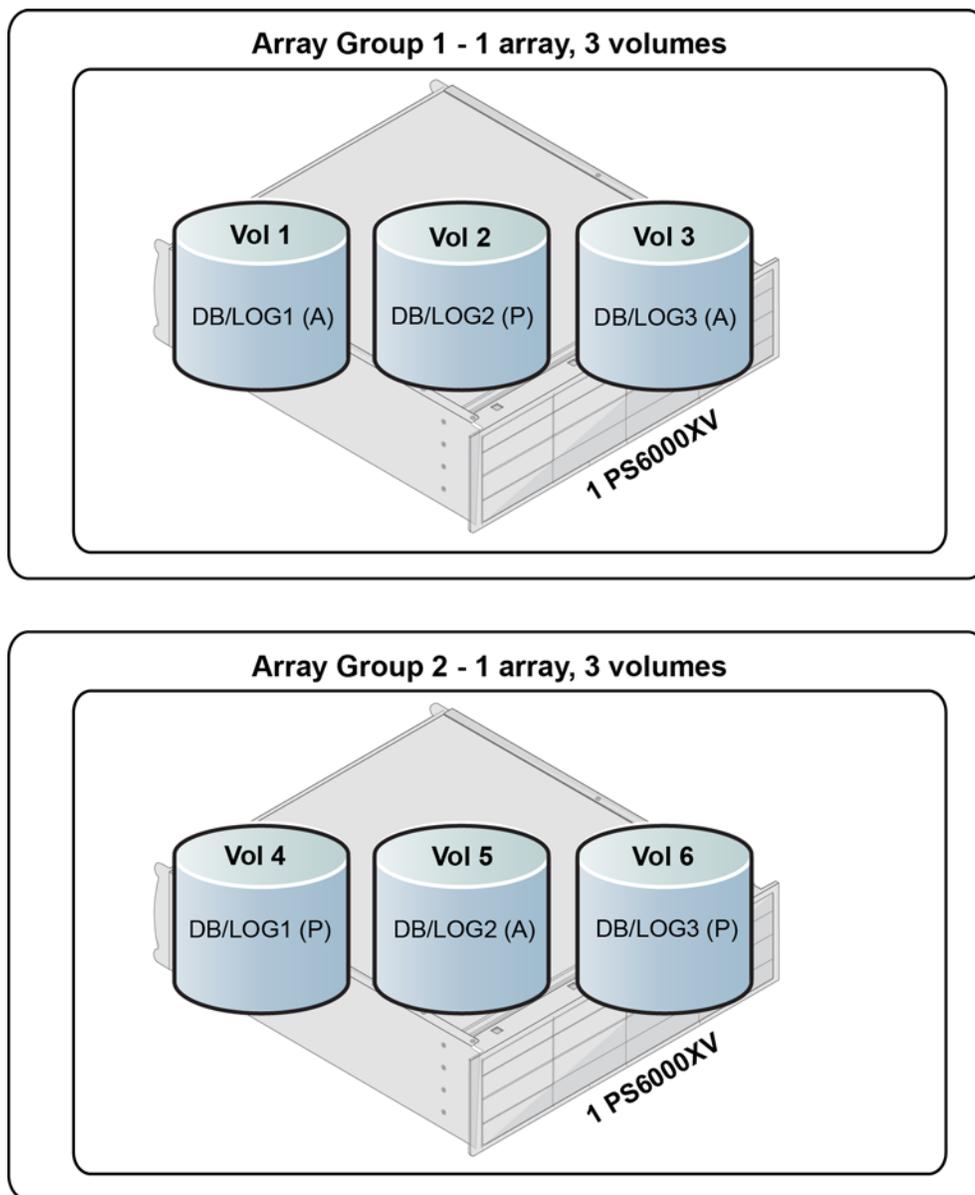
The following table and diagram provides details of the solutions data layout on PS Group volumes and maps the volumes to the servers in the DAG.

**Table 1: Solution data layout**

Server	PS Group	Pool	Volume	DB and Log Group	Active/Passive
Server 1	PS Group 1 consisting of 1 PS6000XV array	Data Pool 1 consisting of 1 PS6000XV	Volume 1	DB/Log 1	Active (copy 1)
			Volume 2	DB/Log 2	Passive (copy 2)
			Volume 3	DB/Log 3	Active (copy 1)
Server 2	PS Group 2 consisting of 1 PS6000XV array	Data Pool 2 consisting of 1 PS6000XV	Volume 4	DB/Logs 1	Passive (copy 2)
			Volume 5	DB/Logs 2	Active (copy 1)
			Volume 6	DB/Logs 3	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
- 5,000 user mailboxes
- 0.25 I/O per second per user (0.3 tested for 20% headroom)
- 1024 MB mailbox quota per user
- 3 databases per server
- 1.62 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	5,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	2
Number of active mailboxes/server	5000
Number of databases/host	3
Number of copies/database	2
Number of mailboxes/database	1667
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.3
Database and Log LUN size	5.96TB (2037.65GB x 3 LUNs )
Total database size for performance testing	4.86 TB (1.62TB x 3 databases)
% storage capacity used by Exchange database <sup>2</sup>	81.54% (4.86TB/5.96TB)

---

<sup>2</sup> 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 hardware

Storage Connectivity	iSCSI
Storage model and OS/firmware revision	Dell EqualLogic PS6000XV Firmware Rev: V5.0.2 <a href="http://dell.com/products/viewwww..aspx?id=2509">http://dell.com/products/viewwww..aspx?id=2509</a>
Storage cache	4 GB (1-PS6000 * 4 GB = 4 GB)
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
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 2.40 GHz with 12GB RAM
Total number of drives tested	16
Maximum number of spindles	16

## 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	16
Total raw storage capacity (GB)	8.73 TB (8942.56GB)
Drive slice size (GB)	N/A
Number of disks per LUN	Up to 16 (automatically allocated based on load)
Raid level	All storage pools configured as RAID 50
Total formatted capacity (GB)	5.96 TB (6112.95GB)
Storage capacity utilization	55.67% (4.86TB/8.73TB)
Database capacity utilization	81.54% (4.86TB/5.96TB)

## 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 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](http://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.

<b>Database I/O</b>	
Database Drive Transfers/sec	1754.65
Database Drive Reads/sec	1124.50
Database Drive Writes/sec	630.16
Average Database Drive Read Latency (ms)	16.52
Average Database Drive Write Latency (ms)	4.62
<b>Transaction Log I/O</b>	
Log Drive Writes/sec	91.69
Average Log Drive Write Latency (ms)	2.29

## 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	59.88(Average)
MB read/sec total per server	179.66(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)	1.34 (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

<b>Overall Test Result</b>	<b>Pass</b>
<b>Machine Name</b>	WIN-J6FRBTN5B34
<b>Test Description</b>	5000 1024MB .3iop 3 instances 64k CC Flowcontrol Jumbo Frames
<b>Test Start Time</b>	6/8/2011 1:15:32 PM
<b>Test End Time</b>	6/9/2011 1:16:33 PM
<b>Collection Start Time</b>	6/8/2011 1:16:28 PM
<b>Collection End Time</b>	6/9/2011 1:16:26 PM
<b>Jetstress Version</b>	14.01.0225.017
<b>ESE Version</b>	14.01.0218.012
<b>Operating System</b>	Windows Server 2008 R2 Enterprise (6.1.7600.0)
<b>Performance Log</b>	<a href="C:\JetStress\1024MB_results\3_luns\build3\run5_12thrs_24hrs_jumbo_flowC\Performance_2011_6_8_13_15_40.blg">C:\JetStress\1024MB_results\3_luns\build3\run5_12thrs_24hrs_jumbo_flowC\Performance_2011_6_8_13_15_40.blg</a>

Database Sizing and Throughput

<b>Achieved Transactional I/O per Second</b>	1658.045
<b>Target Transactional I/O per Second</b>	1500
<b>Initial Database Size (bytes)</b>	5379145793536
<b>Final Database Size (bytes)</b>	5416240218112
<b>Database Files (Count)</b>	3

Jetstress System Parameters

<b>Thread Count</b>	12 (per database)
<b>Minimum Database Cache</b>	96.0 MB
<b>Maximum Database Cache</b>	768.0 MB
<b>Insert Operations</b>	40%
<b>Delete Operations</b>	20%
<b>Replace Operations</b>	5%

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

Database Configuration

**Instance4988.1** Log path: C:\dbs\db1\log  
 Database: C:\dbs\db1\db\Jetstress001001.edb

**Instance4988.2** Log path: C:\dbs\db2\log  
 Database: C:\dbs\db2\db\Jetstress002001.edb

**Instance4988.3** Log path: C:\dbs\db3\log  
 Database: C:\dbs\db3\db\Jetstress003001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O DB Reads Avg Latency (msec)	I/O DB Writes Avg Latency (msec)	I/O DB Reads/sec	I/O DB Writes/sec	I/O DB Reads Avg Bytes	I/O DB Writes Average Bytes	I/O Log Reads Avg Latency (msec)	I/O Log Writes Avg Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance4988.1	17.992	4.855	353.608	198.009	32926.638	34257.184	0.000	2.722	0.000	84.499	0.000	5963.293
Instance4988.2	16.872	4.351	354.181	198.922	32967.306	34252.214	0.000	2.394	0.000	87.262	0.000	5801.469
Instance4988.3	17.112	4.752	354.481	198.845	32952.195	34247.711	0.000	2.760	0.000	84.492	0.000	5962.180

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4988.1	21.146	261926.872
Instance4988.2	21.128	261895.828
Instance4988.3	21.127	261915.190

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4988.1	2.098	232555.463
Instance4988.2	2.105	232558.869
Instance4988.3	2.098	232556.990

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Avg (msec)	I/O Database Writes Avg (msec)	I/O DB Reads/sec	I/O DB Writes/sec	I/O DB Reads Avg Bytes	I/O DB Writes Avg Bytes	I/O Log Reads Avg (msec)	I/O Log Writes Avg (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance4988.1	17.992	4.855	374.753	198.009	45848.080	34257.184	5.491	2.722	2.098	84.499	232555.463	5963.293

Instance4988.2	16.872	4.351	375.309	198.92	45855.089	34252.214	4.728	2.394	2.105	87.262	232558.869	5801.469
Instance4988.3	17.112	4.752	375.608	198.84	45830.878	34247.711	4.880	2.760	2.098	84.492	232556.990	5962.180

### Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.946	1.129	4.231
Available MBytes	12659.342	12635.000	12727.000
Free System Page Table Entries	33555487.676	33555018.000	33556618.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	205227638.209	204718080.000	207314944.000
Pool Paged Bytes	146262295.855	145121280.000	147337216.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

### Test Log

6/8/2011 1:15:32 PM -- Jetstress testing begins ...  
6/8/2011 1:15:32 PM -- Preparing for testing ...  
6/8/2011 1:15:36 PM -- Attaching databases ...  
6/8/2011 1:15:36 PM -- Preparations for testing are complete.  
6/8/2011 1:15:36 PM -- Starting transaction dispatch ..  
6/8/2011 1:15:36 PM -- Database cache settings: (minimum: 96.0 MB, maximum: 768.0 MB)  
6/8/2011 1:15:36 PM -- Database flush thresholds: (start: 7.7 MB, stop: 15.3 MB)  
6/8/2011 1:15:40 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).  
6/8/2011 1:15:40 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).  
6/8/2011 1:15:44 PM -- Operation mix: Sessions 12, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.  
6/8/2011 1:15:44 PM -- Performance logging started (interval: 15000 ms).  
6/8/2011 1:15:44 PM -- Attaining prerequisites:  
6/8/2011 1:16:28 PM -- \\MSEXCHANGE Database(JetstressWin)\Database Cache Size, Last: 738947100.0 (lower bound: 724775700.0, upper bound: none)  
6/9/2011 1:16:29 PM -- Performance logging has ended.  
6/9/2011 1:16:29 PM -- JetInterop batch transaction stats: 861959, 863291 and 861843.  
6/9/2011 1:16:29 PM -- Dispatching transactions ends.  
6/9/2011 1:16:29 PM -- Shutting down databases ...  
6/9/2011 1:16:33 PM -- Instance4988.1 (complete), Instance4988.2 (complete) and Instance4988.3 (complete)  
6/9/2011 1:16:33 PM --  
[C:\JetStress\1024MB\\_results\3\\_luns\build3\run5\\_12thrs\\_24hrs\\_jumbo\\_flowC\Performance\\_2011\\_6\\_8\\_13\\_15\\_40.blg](#) has 5754 samples.  
6/9/2011 1:16:33 PM -- Creating test report ...  
6/9/2011 1:16:57 PM -- Instance4988.1 has 18.0 for I/O Database Reads Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.1 has 2.7 for I/O Log Writes Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.1 has 2.7 for I/O Log Reads Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.2 has 16.9 for I/O Database Reads Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.2 has 2.4 for I/O Log Writes Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.2 has 2.4 for I/O Log Reads Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.3 has 17.1 for I/O Database Reads Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.3 has 2.8 for I/O Log Writes Average Latency.  
6/9/2011 1:16:57 PM -- Instance4988.3 has 2.8 for I/O Log Reads Average Latency.  
6/9/2011 1:16:57 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.  
6/9/2011 1:16:57 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.  
6/9/2011 1:16:57 PM --  
[C:\JetStress\1024MB\\_results\3\\_luns\build3\run5\\_12thrs\\_24hrs\\_jumbo\\_flowC\Performance\\_2011\\_6\\_8\\_13\\_15\\_40.xml](#) has 5751 samples queried.

# Appendix B: Performance Testing

Microsoft Exchange *Jetstress 2010*

## Performance Test Result Report

Test Summary

<b>Overall Test Result</b>	<b>Pass</b>
<b>Machine Name</b>	WIN-J6FRBTN5B34
<b>Test Description</b>	5000 1024MB .3iop 3 instances 64k CC Flowcontrol Jumbo Frames
<b>Test Start Time</b>	6/7/2011 7:15:44 PM
<b>Test End Time</b>	6/7/2011 9:16:45 PM
<b>Collection Start Time</b>	6/7/2011 7:16:40 PM
<b>Collection End Time</b>	6/7/2011 9:16:40 PM
<b>Jetstress Version</b>	14.01.0225.017
<b>ESE Version</b>	14.01.0218.012
<b>Operating System</b>	Windows Server 2008 R2 Enterprise (6.1.7600.0)
<b>Performance Log</b>	<a href="C:\JetStress\1024MB_results\3_luns\build3\run4_12thrds_2hrs_jumbo_flowC\Performance_2011_6_7_19_15_52.blg">C:\JetStress\1024MB_results\3_luns\build3\run4_12thrds_2hrs_jumbo_flowC\Performance_2011_6_7_19_15_52.blg</a>

Database Sizing and Throughput

**Achieved Transactional I/O per Second** 1754.655

**Target Transactional I/O per Second** 1500

**Initial Database Size (bytes)** 5375823904768

**Final Database Size (bytes)** 5379145793536

**Database Files (Count)** 3

Jetstress System Parameters

**Thread Count** 12 (per database)

**Minimum Database Cache** 96.0 MB

**Maximum Database Cache** 768.0 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

**Instance4988.1** Log path: C:\dbs\db1\log  
 Database: C:\dbs\db1\db\Jetstress001001.edb

**Instance4988.2** Log path: C:\dbs\db2\log  
 Database: C:\dbs\db2\db\Jetstress002001.edb

**Instance4988.3** Log path: C:\dbs\db3\log  
 Database: C:\dbs\db3\db\Jetstress003001.edb

Transactional I/O Performance

MSEExchange Database ==> Instances	I/O DB Reads Avg Latency (msec)	I/O DB Writes Avg Latency (msec)	I/O DB Reads/sec	I/O DB 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 Avg Bytes	I/O Log Writes Avg Bytes
Instance4988.1	17.008	4.549	376.500	210.695	32941.105	34452.609	0.000	2.261	0.000	91.757	0.000	5918.073
Instance4988.2	16.271	4.714	373.051	208.980	32973.554	34450.172	0.000	2.312	0.000	91.304	0.000	5917.770
Instance4988.3	16.295	4.600	374.949	210.481	32976.280	34452.388	0.000	2.297	0.000	92.034	0.000	5910.065

Background Database Maintenance I/O Performance

MSEExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4988.1	22.147	261963.451
Instance4988.2	22.128	261901.392
Instance4988.3	22.123	261932.461

Log Replication I/O Performance

MSEExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4988.1	2.258	232553.968
Instance4988.2	2.246	232557.680
Instance4988.3	2.260	232564.345

Total I/O Performance

<b>MSExchange Database =&gt; Instances</b>	<b>I/O DB Reads Avg (msec)</b>	<b>I/O DB Writes Avg (msec)</b>	<b>I/O DB Reads/sec</b>	<b>I/O DB Writes/sec</b>	<b>I/O DB Reads Avg Bytes</b>	<b>I/O DB Writes Avg Bytes</b>	<b>I/O Log Reads Avg Latency (msec)</b>	<b>I/O Log Writes Avg Latency (msec)</b>	<b>I/O Log Reads/sec</b>	<b>I/O Log Writes/sec</b>	<b>I/O Log Reads Average Bytes</b>	<b>I/O Log Writes Average Bytes</b>
Instance4988.1	17.008	4.549	398.646	210.695	45664.276	34452.609	5.116	2.261	2.258	91.757	232553.968	5918.073
Instance4988.2	16.271	4.714	395.179	208.980	45792.509	34450.172	4.823	2.312	2.246	91.304	232557.680	5917.770
Instance4988.3	16.295	4.600	397.073	210.481	45732.886	34452.388	5.077	2.297	2.260	92.034	232564.345	5910.065

### Host System Performance

<b>Counter</b>	<b>Average</b>	<b>Minimum</b>	<b>Maximum</b>
<b>% Processor Time</b>	2.184	1.531	4.062
<b>Available MBytes</b>	12669.640	12660.000	12722.000
<b>Free System Page Table Entries</b>	33555407.350	33555042.000	33556106.000
<b>Transition Pages RePurposed/sec</b>	0.000	0.000	0.000
<b>Pool Nonpaged Bytes</b>	205478869.333	205041664.000	205643776.000
<b>Pool Paged Bytes</b>	144621977.600	144515072.000	144728064.000
<b>Database Page Fault Stalls/sec</b>	0.000	0.000	0.000

### Test Log

6/7/2011 7:15:44 PM -- Jetstress testing begins ...  
6/7/2011 7:15:44 PM -- Preparing for testing ...  
6/7/2011 7:15:48 PM -- Attaching databases ...  
6/7/2011 7:15:48 PM -- Preparations for testing are complete.  
6/7/2011 7:15:48 PM -- Starting transaction dispatch ..  
6/7/2011 7:15:48 PM -- Database cache settings: (minimum: 96.0 MB, maximum: 768.0 MB)  
6/7/2011 7:15:48 PM -- Database flush thresholds: (start: 7.7 MB, stop: 15.3 MB)  
6/7/2011 7:15:52 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).  
6/7/2011 7:15:52 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).  
6/7/2011 7:15:59 PM -- Operation mix: Sessions 12, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.  
6/7/2011 7:15:59 PM -- Performance logging started (interval: 15000 ms).  
6/7/2011 7:15:59 PM -- Attaining prerequisites:  
6/7/2011 7:16:40 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 726077400.0 (lower bound: 724775700.0, upper bound: none)  
6/7/2011 9:16:41 PM -- Performance logging has ended.  
6/7/2011 9:16:41 PM -- JetInterop batch transaction stats: 77818, 77483 and 77497.  
6/7/2011 9:16:41 PM -- Dispatching transactions ends.  
6/7/2011 9:16:41 PM -- Shutting down databases ...  
6/7/2011 9:16:45 PM -- Instance4988.1 (complete), Instance4988.2 (complete) and Instance4988.3 (complete)  
6/7/2011 9:16:45 PM --  
C:\JetStress\1024MB\_results\3\_luns\build3\run4\_12thrds\_2hrs\_jumbo\_flowC\Performance\_2011\_6\_7\_19\_15\_52.blg has 482 samples.  
6/7/2011 9:16:45 PM -- Creating test report ...  
6/7/2011 9:16:46 PM -- Instance4988.1 has 17.0 for I/O Database Reads Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.1 has 2.3 for I/O Log Writes Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.1 has 2.3 for I/O Log Reads Average Latency.

6/7/2011 9:16:46 PM -- Instance4988.2 has 16.3 for I/O Database Reads Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.2 has 2.3 for I/O Log Writes Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.2 has 2.3 for I/O Log Reads Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.3 has 16.3 for I/O Database Reads Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.3 has 2.3 for I/O Log Writes Average Latency.  
6/7/2011 9:16:46 PM -- Instance4988.3 has 2.3 for I/O Log Reads Average Latency.  
6/7/2011 9:16:46 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.  
6/7/2011 9:16:46 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.  
6/7/2011 9:16:46 PM --  
C:\JetStress\1024MB\_results\3\_luns\build3\run4\_12thrs\_2hrs\_jumbo\_flowC\Performance\_2011\_6\_7\_19\_15\_52.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
<b>Instance4988.1</b>	1721755.09	08:15:55	57.86
<b>Instance4988.2</b>	1721779.09	07:48:47	61.21
<b>Instance4988.3</b>	1721771.09	07:53:34	60.59

Jetstress System Parameters

**Thread Count** 12 (per database)

**Minimum Database Cache** 96.0 MB

**Maximum Database Cache** 768.0 MB

**Insert Operations** 40%

**Delete Operations** 20%

**Replace Operations** 5%

**Read Operations** 35%

**Lazy Commits** 70%

Database Configuration

**Instance4988.1** Log path: C:\dbs\db1\log  
Database: C:\dbs\db1\db\Jetstress001001.edb

**Instance4988.2** Log path: C:\dbs\db2\log  
Database: C:\dbs\db2\db\Jetstress002001.edb

**Instance4988.3** Log path: C:\dbs\db3\log  
Database: C:\dbs\db3\db\Jetstress003001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O DB Writes Average Latency (msec)	I/O DB Reads/sec	I/O DB Writes/sec	I/O Database Reads Average Bytes	I/O DB Writes Avg 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 Avg Bytes	I/O Log Writes Avg Bytes
<b>Instance4988.1</b>	7.650	0.000	231.523	0.000	262144.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000
<b>Instance4988.2</b>	6.664	0.000	244.886	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
<b>Instance4988.3</b>	7.326	0.000	242.416	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	1.633	0.032	2.360
Available MBytes	13480.989	13463.000	13514.000
Free System Page Table Entries	33555776.317	33555482.000	33557098.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	204761828.590	204754944.000	204853248.000
Pool Paged Bytes	150000224.194	149291008.000	150425600.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log6/9/2011 11:17:30 PM -- Jetstress testing begins ...

6/9/2011 11:17:30 PM -- Preparing for testing ...

6/9/2011 11:17:34 PM -- Attaching databases ...

6/9/2011 11:17:34 PM -- Preparations for testing are complete.

6/9/2011 11:17:40 PM -- Performance logging started (interval: 30000 ms).

6/9/2011 11:17:40 PM -- Backing up databases ...

6/10/2011 7:33:36 AM -- Performance logging has ended.

6/10/2011 7:33:36 AM -- Instance4988.1 (100% processed), Instance4988.2 (100% processed) and Instance4988.3 (100% processed)

6/10/2011 7:33:36 AM -- [C:\JetStress\1024MB\\_results\3\\_luns\build3\run6\\_backup\DatabaseBackup\\_2011\\_6\\_9\\_23\\_17\\_34.blg](#) has 990 samples.

6/10/2011 7:33:36 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
<b>Instance4988.1</b>	508	708.0852436
<b>Instance4988.2</b>	505	651.5975444
<b>Instance4988.3</b>	501	673.8743836

Database Configuration

<b>Instance4988.1</b>	Log path: C:\dbs\db1\log Database: C:\dbs\db1\db\Jetstress001001.edb
<b>Instance4988.2</b>	Log path: C:\dbs\db2\log Database: C:\dbs\db2\db\Jetstress002001.edb
<b>Instance4988.3</b>	Log path: C:\dbs\db3\log Database: C:\dbs\db3\db\Jetstress003001.edb

Transactional I/O Performance

MSEExchange Database ==> Instances	I/O DB Reads Average Latency (msec)	I/O DB Writes Average Latency (msec)	I/O DB Reads/sec	I/O DB Writes/sec	I/O DB Reads Average Bytes	I/O DB Writes Average Bytes	I/O Log Reads Avg Latency (msec)	I/O Log Writes Avg Latency (msec)	I/O Log Read s/sec	I/O Log Write s/sec	I/O Log Reads Average Bytes	I/O Log Writes Avg Bytes
Instance4988.1	10.895	1.793	621.691	4.304	41053.329	32579.678	10.324	0.037	6.462	0.017	231019.860	1.471
Instance4988.2	9.826	1.567	677.082	4.649	41089.850	32665.600	10.797	0.002	6.973	0.005	231624.407	1.600
Instance4988.3	10.331	1.763	654.232	4.459	41137.441	32471.009	11.194	0.004	6.688	0.004	230225.122	1.547

Background Database Maintenance I/O Performance

MSEExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
<b>Instance4988.1</b>	29.508	262026.952
<b>Instance4988.2</b>	29.558	261899.358
<b>Instance4988.3</b>	29.416	261915.885

Total I/O Performance

MSEExchange Database	I/O DB Reads Avg	I/O DB Writes Avg	I/O DB Reads	I/O DB Writes/se	I/O Database Reads	I/O Database	I/O Log Reads	I/O Log Writes	I/O Log	I/O Log	I/O Log Reads	I/O Log Writes
----------------------	------------------	-------------------	--------------	------------------	--------------------	--------------	---------------	----------------	---------	---------	---------------	----------------

==> Instances	Latency (msec)	Latency (msec)	/sec	c	Average Bytes	Writes Average Bytes	Average Latency (msec)	Avg Latency (msec)	Read s/sec	Write s/sec	Average Bytes	Avg Bytes
Instance4988.1	10.895	1.793	651.199	4.304	51066.263	32579.678	10.324	0.037	6.462	0.017	231019.860	1.471
Instance4988.2	9.826	1.567	706.640	4.649	50326.048	32665.600	10.797	0.002	6.973	0.005	231624.407	1.600
Instance4988.3	10.331	1.763	683.647	4.459	50637.070	32471.009	11.194	0.004	6.688	0.004	230225.122	1.547

## Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	2.016	0.000	7.291
Available MBytes	12672.401	12631.000	13446.000
Free System Page Table Entries	33555411.040	33555026.000	33556106.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	205581526.189	205193216.000	206086144.000
Pool Paged Bytes	150511334.327	150474752.000	150605824.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log6/10/2011 11:06:15 AM -- Jetstress testing begins ...

6/10/2011 11:06:15 AM -- Preparing for testing ...

6/10/2011 11:06:18 AM -- Attaching databases ...

6/10/2011 11:06:18 AM -- Preparations for testing are complete.

6/10/2011 11:06:18 AM -- Starting transaction dispatch ..

6/10/2011 11:06:18 AM -- Database cache settings: (minimum: 96.0 MB, maximum: 768.0 MB)

6/10/2011 11:06:18 AM -- Database flush thresholds: (start: 7.7 MB, stop: 15.3 MB)

6/10/2011 11:06:23 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

6/10/2011 11:06:23 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

6/10/2011 11:06:26 AM -- Operation mix: Sessions 12, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

6/10/2011 11:06:26 AM -- Performance logging started (interval: 15000 ms).

6/10/2011 11:06:26 AM -- Generating log files ...

6/10/2011 11:39:03 AM -- C:\dbs\db1\log (101.6% generated), C:\dbs\db2\log (101.0% generated) and C:\dbs\db3\log (100.2% generated)

6/10/2011 11:39:03 AM -- Performance logging has ended.

6/10/2011 11:39:03 AM -- JetInterop batch transaction stats: 21450, 21169 and 21417.

6/10/2011 11:39:03 AM -- Dispatching transactions ends.

6/10/2011 11:39:03 AM -- Shutting down databases ...

6/10/2011 11:39:06 AM -- Instance4988.1 (complete), Instance4988.2 (complete) and Instance4988.3 (complete)

6/10/2011 11:39:06 AM -- [C:\JetStress\1024MB\\_results\3\\_luns\build3\run7\\_softrecov\Performance\\_2011\\_6\\_10\\_11\\_6\\_23.blg](C:\JetStress\1024MB_results\3_luns\build3\run7_softrecov\Performance_2011_6_10_11_6_23.blg) has 130 samples.

6/10/2011 11:39:06 AM -- Creating test report ...

6/10/2011 11:39:07 AM -- Instance4988.1 has 16.4 for I/O Database Reads Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.1 has 2.9 for I/O Log Writes Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.1 has 2.9 for I/O Log Reads Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.2 has 15.3 for I/O Database Reads Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.2 has 2.5 for I/O Log Writes Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.2 has 2.5 for I/O Log Reads Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.3 has 15.7 for I/O Database Reads Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.3 has 3.0 for I/O Log Writes Average Latency.

6/10/2011 11:39:07 AM -- Instance4988.3 has 3.0 for I/O Log Reads Average Latency.

6/10/2011 11:39:07 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.

6/10/2011 11:39:07 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

6/10/2011 11:39:07 AM -- [C:\JetStress\1024MB\\_results\3\\_luns\build3\run7\\_softrecov\Performance\\_2011\\_6\\_10\\_11\\_6\\_23.xml](C:\JetStress\1024MB_results\3_luns\build3\run7_softrecov\Performance_2011_6_10_11_6_23.xml) has 129 samples queried.

6/10/2011 11:39:07 AM -- [C:\JetStress\1024MB\\_results\3\\_luns\build3\run7\\_softrecov\Performance\\_2011\\_6\\_10\\_11\\_6\\_23.html](C:\JetStress\1024MB_results\3_luns\build3\run7_softrecov\Performance_2011_6_10_11_6_23.html) was saved.  
6/10/2011 11:39:11 AM -- Performance logging started (interval: 2000 ms).  
6/10/2011 11:39:11 AM -- Recovering databases ...  
6/10/2011 11:50:59 AM -- Performance logging has ended.  
6/10/2011 11:50:59 AM -- Instance4988.1 (708.0852436), Instance4988.2 (651.5975444) and Instance4988.3 (673.8743836)  
6/10/2011 11:51:00 AM -- [C:\JetStress\1024MB\\_results\3\\_luns\build3\run7\\_softrecov\SoftRecovery\\_2011\\_6\\_10\\_11\\_39\\_7.blg](C:\JetStress\1024MB_results\3_luns\build3\run7_softrecov\SoftRecovery_2011_6_10_11_39_7.blg) has 349 samples.  
6/10/2011 11:51:00 AM -- Creating test report ...