Dell EqualLogic PS 6210E Series 5,000-User Mailbox Resiliency Storage Solution for Microsoft Exchange Server 2013

Microsoft ESRP 4.0

Daniel J. Curran: Principal Engineer
Dell Storage Engineering
Tested Date: May 2014

Revisions

Date	Description
July 2014	Initial release

Table of contents

Executive summary

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.

[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

1 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 hundreds of 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), Serial-Attached SCSI (SAS) and Solid State Drives (SSD) 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.

The Dell EqualLogic PS6210 arrays offer IT generalists the ability to manage more data with fewer resources. The family has the capability to tightly integrate with common application environments and the flexibility to be used with various operating systems. The PS6210 Series is designed to:

- Deploy in virtualized VMware or Hyper-V environments, Exchange and SQL server applications and distributed DAS implementations
- Support block and file data along with the FS Series through a single intuitive interface
- Support server and desktop virtualization
- Promote server/storage/network consolidation to ease management of the virtual ecosystem
- Support multi-way replication for robust disaster recovery

EqualLogic PS6210 Series Arrays provide enhanced storage performance for the small-to-medium enterprise with the simplicity of the EqualLogic product line. Dell EqualLogic PS6210 arrays can drive up to approximately 2GB/sec of throughput per array for sequential, large-block workloads. In addition, flash-enabled EqualLogic PS6210 arrays provide up to 3 times the random performance of prior-generation arrays. The full line of PS6210 Series arrays brings 10GbE iSCSI SAN speed and efficiency to real-world applications, featuring:

- Dual controllers, each with 16GB non-volatile cache of memory
- Two 10GBASE-T RJ45 auto-sensing (10Gb/1Gb/100Mb) ports
- Two 10GbE SFP+ ports for fiber or copper cabling
- Up to 24 hot-pluggable drives, including SAS, NL-SAS and SSD
- Model options for 2.5" drives and 3.5" drives
- Up to 96TB capacity per array

2 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 2013.

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 2013 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.

3 Simulated environment

The solution presented in this document is designed to simulate a medium-sized number of mailboxes hosted on highly redundant hardware. Application level redundancy is augmented with redundant storage to create a highly available and fault tolerant solution.

The Mailbox Resiliency features of Exchange 2013 have greatly enhanced the availability of Exchange Server, while also improving I/O performance. The solution presented here is a Mailbox Resiliency solution utilizing 1 Database Availability Group (DAG) and 2 copies of every database. The tested environment simulates all users in this DAG running on a single PS6210E storage array, or half of the storage solution. The number of users simulated was 5,000 on one server or half of the server solution. The mailbox size was 2GB per user. The server has 6 databases, all active simulating a failover (half of all server and or storage had failed). See Figure 2 – PS Series topology for Exchange storage solution as tested.

The replication mechanism is the native Exchange 2013 DAG database replication engine. This is a very efficient and reliable replication mechanism and is the recommended method for providing highly-available and redundant Exchange solutions.

4 Solution description

The following sections outline the hardware and software environment for a Microsoft Exchange Server 2013 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 R720 servers with storage provided by 2 Dell EqualLogic PS6210E storage arrays. Connectivity between servers and storage is via iSCSI protocol and two Dell Force10 S4810 switches. See Figure 1 – PS Series topology for Exchange storage solution best practice.

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-10 redundancy. Two of the 24 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 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.

For information about compatibility please use the following link:

http://windowsservercatalog.com/item.aspx?idItem=27e784c3-e90f-6503-7374-6e2fcbce1f9f&bCatID=1282

5 Hardware and software

The solution hardware environment is described in the following table.

Table 1 Solution hardware environment

Storage	Drives	Servers	Ethernet connections
2 PS6210E storage arrays configured into 2 PS Series groups, each containing one data storage pool with one member (for dB and log data).	48 7200-RPM 4TB Serial Attached SCSI drives (22 active & 2 spare drives per Array)	2 Dell PowerEdge R720 Server, each with 2 Intel Xeon E5-2640 v2 @ 2.0 GHz 8 Core CPU and 192GB memory running MS Windows Server 2012 Datacenter Editionx64	Broadcom BCM57810 NetXtreme II 10 GigE

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 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 (6) to support 5,000 users, dividing the workload equally between the six databases.

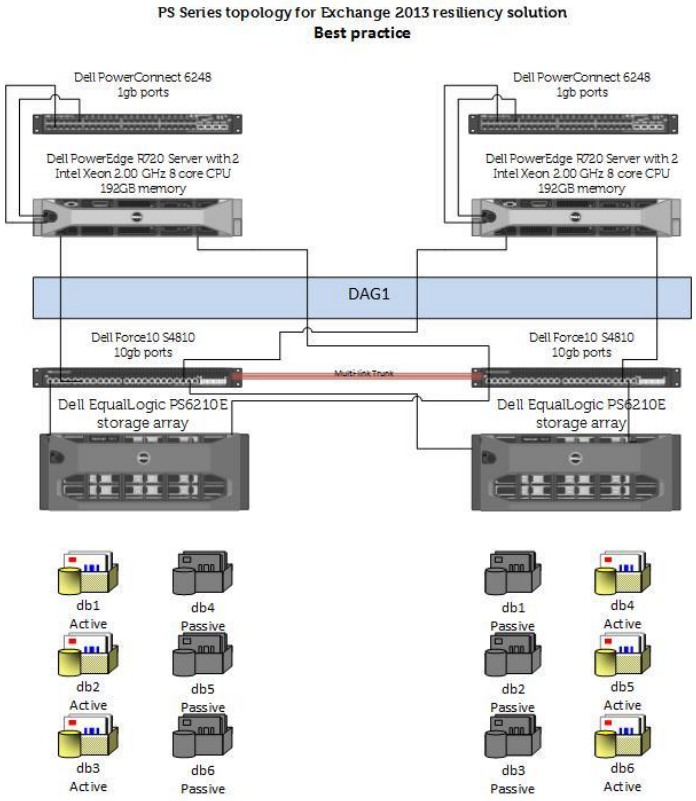

Figure 1 PS Series topology for Exchange storage solution best practice

6 Exchange DAG architecture

This mailbox solution utilizes two PS Series array groups consisting of 1 Dell EqualLogic PS6210E storage array per group. There is one data storage pool per group, each consisting of 1 PS6210E array configured as RAID 10. Within each storage pool six volumes are created for a total of 12 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 6 databases and 6 log folders configured in six separate volumes. Within the overall solution, 1 server provides 3 active and 3 passive database copies. While the second server provides 3 active and 3 passive database copies. There are a total of 6 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 fail.

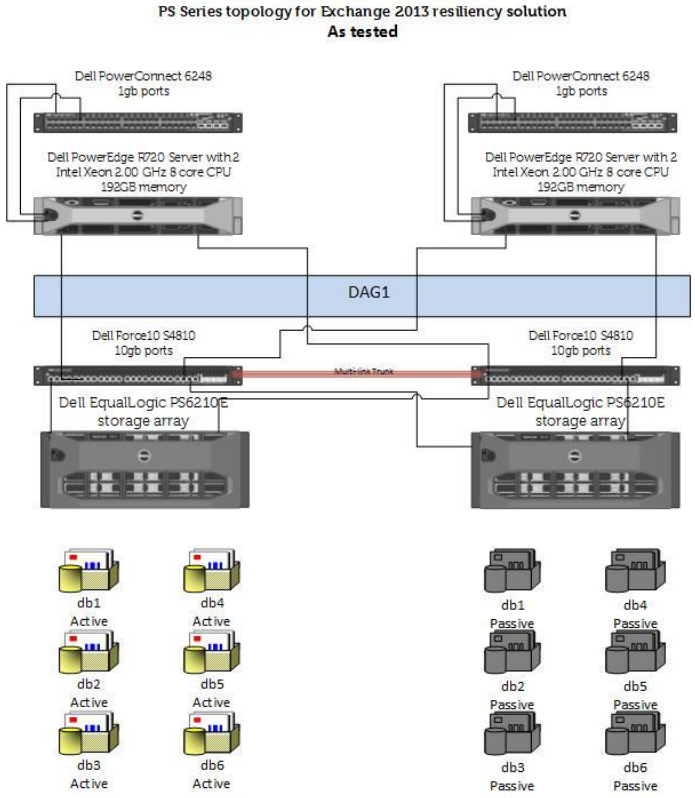

Figure 2 PS Series topology for Exchange storage solution as tested

The following table and diagram (See Table 2: Solution data layout) provides details of the solutions data layout on PS Group volumes and maps the volumes to the servers in the DAG.

Table 2 Solution data layout

Server	PS Group	Pool	Volume	DB and Log Group	Active/Passive
Server 1	PS Group 1 consisting of 1 PS6210E array	Data Pool 1 consisting of 1 PS6210E	Volume 1	DB1/Log	Active (copy 1)
			Volume 2	DB2/Log	Active (copy 1)
			Volume 3	DB3/Log	Active (copy 1)
			Volume 4	DB4/Logs	Passive (copy 2)
			Volume 5	DB5/Logs	Passive (copy 2)
			Volume 6	DB6/Logs	Passive (copy 2)
Server 2	PS Group 2 consisting of 1 PS6210E array	Data Pool 2 consisting of 1 PS6210E	Volume 7	DB1/Log	Passive (copy 2)
			Volume 8	DB2/Log	Passive (copy 2)
			Volume 9	DB3/Log	Passive (copy 2)
			Volume 10	DB4/Logs	Active (copy 1)
			Volume 11	DB5/Logs	Active (copy 1))
			Volume 12	DB6/Logs	Active (copy 1)

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.

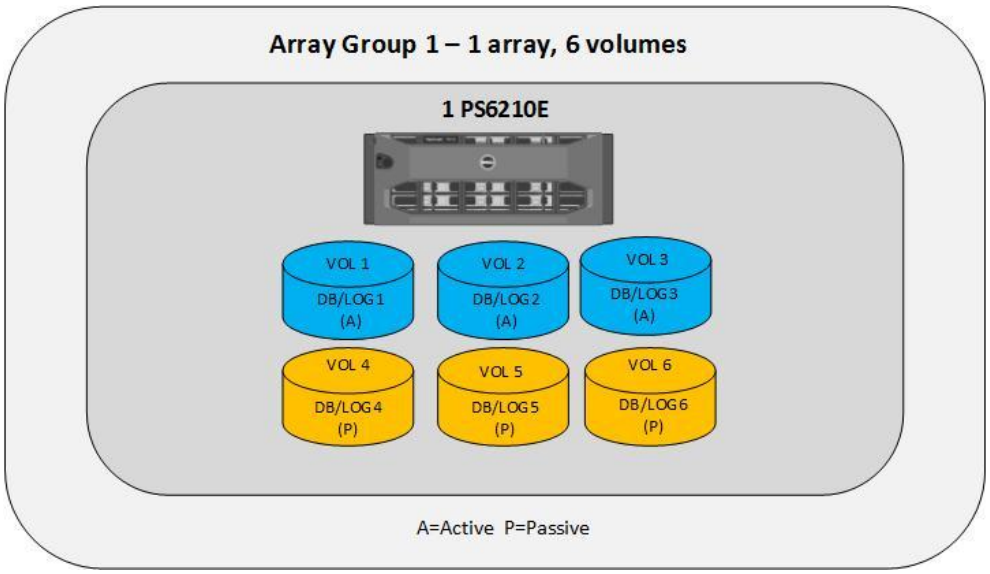

Figure 3 Array Group 1

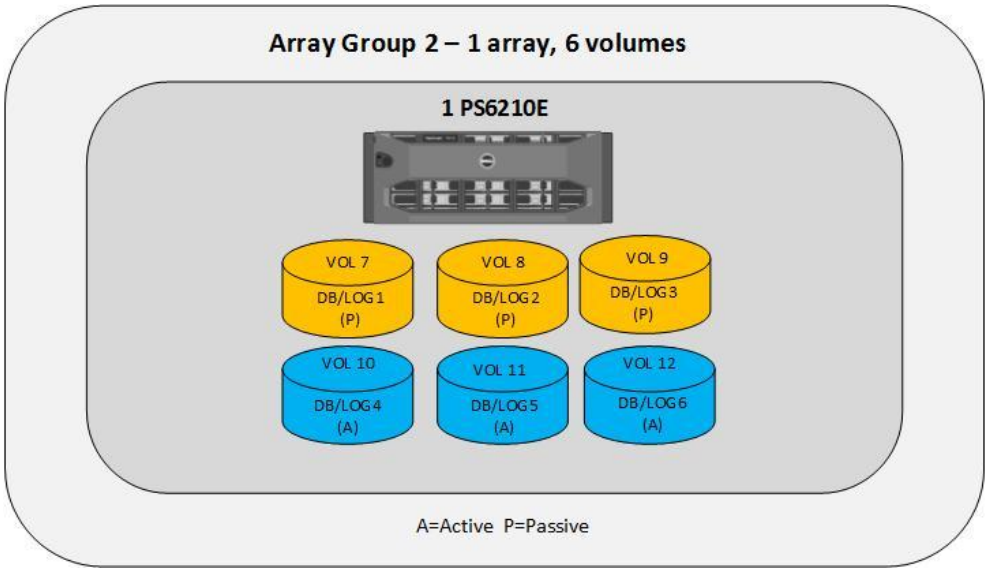

Figure 4 Array Group 2

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

7 Targeted customer profile

The PS Series storage solution is intended for small, medium, and large Microsoft Exchange Server 2013 organizations that want reliable, high-performance, and easy-to-manage drive storage. The tested configuration can support the following:

- 2 (1 tested) Exchange 2013 servers
- 5,000 user mailboxes
- 0.12 I/O per second per user (includes 20% headroom)
- 2048 MB mailbox quota per user
- 6 databases per server
- 1.62 TB database size
- Mailbox Resiliency (2 copies) provides high availability and is the primary data protection mechanism.

7.1 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 v4.0 program. The following tables summarize the testing environment.

8 Simulated Exchange configuration

Table 3

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	6
Number of copies/database	2
Number of mailboxes/database	833
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.12
Database and Log LUN size	11.47 TB (1958 GB x 6 LUNs)
% storage capacity used by Exchange database[2]	34.40% (11.47TB/39.46TB)

[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.

9 Primary storage

Table 4 Hardware

Storage Connectivity	iSCSI
Storage model and OS/firmware revision	Dell EqualLogic PS6210E Firmware Rev: V7.0.3 http://www.dell.com/us/business/p/equallogic-ps6210-series/pd
Storage cache	Dual controllers with 16GB of memory each
Number of storage controllers	2 (2 per array in active/passive configuration)
Number of storage ports	Two 10GBASE-T with RJ45 and two 10GbE SFP+ for fibre or twin-ax copper cabling and one (1) 100BASE-TX dedicated management port per controller
Maximum bandwidth of storage connectivity to host	2x10GbE per Array
Switch type/model/firmware revision	Dell Force10 S4810 Switch (10GbE) Firmware Rev: 9.3.0.0
HBA model and firmware	Broadcom BCM57810 NetXtreme II 10 GigE
Number of HBAs/host	2 used for storage in this test, 6 in all
Host server type	Dell PowerEdge R720 Server, each with 2 Intel Xeon E5-2640 v2 @ 2.0 GHz 8 Core CPU
Total number of drives tested	24 (22 with 2 Hot Spares)
Maximum number of spindles	24

Table 5 Primary storage software

HBA driver	Driver Version: 7.0.5.43
HBA QueueTarget Setting	N/A
HBA QueueDepth Setting	N/A
Multipath I/O DSM	Dell EqualLogic HIT (Host Integration Toolkit) 4.7.0
Host OS	Windows Server 2012, Datacenter x64 Edition
ESE.dll file version	15.0.712.8
Replication solution name/version	Microsoft Exchange Server 2013 DAG replication

10 Primary drive configuration (Mailbox and Log Store Drives)

Table 6

Drive type, speed and firmware revision	**Seagate 4TB 72,00 rpm SAS 6Gb/s drives Model: ST4000NM0023 Firmware Revision: GE09**
Raw capacity per drive (GB)	3717.12 GB
Number of physical drives in test	22 (22 active & 2 hot spares)
Total raw storage capacity (GB)	79.86TB (81776.64 GB)
Drive slice size (GB)	N/A
Number of disks per LUN	Up to 22 (automatically allocated based on load)
Raid level	All storage pools configured as RAID 10
Total formatted capacity (GB)	39.46TB (40407 GB)
Storage capacity utilization	49%
Database capacity utilization	34.40%

11 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 2013 databases must be set to 64k for best performance.
- Size and configure first for I/O performance, then for storage capacity.
- Enable Dell EqualLogic Host Integration Tools V4.7 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.
- Place SAN infrastructure on VLANs or subnets that differ from other production network traffic.
- Enable use of Jumbo Frames.

For server sizing please refer to the Microsoft Exchange Server Role Calculator at http://gallery.technet.microsoft.com/Exchange-2013-Server-Role-f8a61780

General sizing and requirements can be located at http://technet.microsoft.com/en-us/library/aa996719.aspx

Additional best practices on storage design in Exchange 2013 are found at http://technet.microsoft.com/library/dd638104%28EXCHG.150%29#StoreReq

12 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

http://en.community.dell.com/techcenter/extras/m/white_papers/20438088.aspx

13 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.

13.1 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:

1. No relevant errors were reported in the event log for the storage reliability test.
2. No errors were reported by the database and log checksum process.
3. No errors were reported during either the backup or restore process.

13.2 Storage performance results

The Primary Storage performance testing is designed to exercise the storage with maximum sustainable Exchange I/O for 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/O's and the average of all the logical drives' I/O latency in the 2 hours test duration. Each server is listed separately and the aggregate numbers across all servers are also presented.

13.3 Individual server metrics

The server metrics include the sum of I/O's across all mailbox databases and the average latency across all databases on a per server basis.

Table 7

Database I/O	
Database Drive Transfers/sec	1056.71
Database Drive Reads/sec	741.49
Database Drive Writes/sec	315.22
Average Database Drive Read Latency (ms)	16.97
Average Database Drive Write Latency (ms)	1.55
Transaction Log I/O	
Log Drive Writes/sec	75.52
Average Log Drive Write Latency (ms)	.394

14 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 (playing transaction logs into the database).

14.1 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.

Table 8

MB read/sec per database	96.19(Average)
MB read/sec total per server	577.17(Sum)

14.2 Transaction Log Recovery/Replay performance

The purpose of this test is to measure the maximum rate at which the log files can be played against the databases. The following table shows the average rate for 500 log files played in a single database. Each log file is 1 MB in size.

Table 9

Average time to play one Log File (sec)	2.02 (avg. resp. to replay log / avg. # of logs replayed)

15 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.

A Appendix A: Stress Testing

Microsoft Exchange Jetstress 2013

Stress Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	EQLESRP5
Test Description	Stress 26 thread
Test Start Time	5/17/2014 7:24:34 PM
Test End Time	5/18/2014 7:26:43 PM
Collection Start Time	5/17/2014 7:26:41 PM
Collection End Time	5/18/2014 7:26:38 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0712.008
Operating System	Windows Server 2012 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Stress_2014_5_17_19_24_48.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	1009.779
Target Transactional I/O per Second	600
Initial Database Size (bytes)	10756580311040
Final Database Size (bytes)	10784807976960
Database Files (Count)	6

Jetstress System Parameters

Thread Count	26
Minimum Database Cache	192.0 MB
Maximum Database Cache	1536.0 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

Instance1604.1	Log path: F:\log Database: F:\db\Jetstress001001.edb
Instance1604.2	Log path: G:\log Database: G:\db\Jetstress002001.edb
Instance1604.3	Log path: H:\log Database: H:\db\Jetstress003001.edb
Instance1604.4	Log path: I:\log Database: I:\db\Jetstress004001.edb
Instance1604.5	Log path: J:\log Database: J:\db\Jetstress005001.edb
Instance1604.6	Log path: K:\log Database: K:\db\Jetstress006001.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
Instance1604.1	17.751	0.628	117.826	50.470	32831.939	35335.647	0.000	0.330	0.000	11.826	0.000	20351.126
Instance1604.2	17.765	1.050	117.816	50.445	32832.507	35334.590	0.000	0.407	0.000	11.813	0.000	20366.039
Instance1604.3	17.792	1.483	117.799	50.347	32833.009	35337.535	0.000	0.419	0.000	11.774	0.000	20352.744
Instance1604.4	17.811	1.899	117.908	50.432	32834.109	35324.532	0.000	0.431	0.000	11.795	0.000	20340.113
Instance1604.5	17.764	2.300	117.877	50.493	32833.876	35324.138	0.000	0.406	0.000	11.802	0.000	20369.186
Instance1604.6	17.766	2.707	117.899	50.466	32834.208	35334.729	0.000	0.422	0.000	11.791	0.000	20359.234

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1604.1	8.563	261903.522
Instance1604.2	8.564	261914.387
Instance1604.3	8.570	261915.893
Instance1604.4	8.558	261906.347
Instance1604.5	8.567	261906.946
Instance1604.6	8.572	261905.189

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1604.1	1.030	232527.539
Instance1604.2	1.030	232563.133
Instance1604.3	1.026	232529.835
Instance1604.4	1.027	232563.220
Instance1604.5	1.029	232564.965
Instance1604.6	1.028	232555.505

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
Instance1604.1	17.751	0.628	126.389	50.470	48352.127	35335.647	5.808	0.330	1.030	11.826	232527.539	20351.126
Instance1604.2	17.765	1.050	126.380	50.445	48356.100	35334.590	6.406	0.407	1.030	11.813	232563.133	20366.039
Instance1604.3	17.792	1.483	126.369	50.347	48368.459	35337.535	7.507	0.419	1.026	11.774	232529.835	20352.744
Instance1604.4	17.811	1.899	126.466	50.432	48335.142	35324.532	6.913	0.431	1.027	11.795	232563.220	20340.113
Instance1604.5	17.764	2.300	126.445	50.493	48354.953	35324.138	7.024	0.406	1.029	11.802	232564.965	20369.186
Instance1604.6	17.766	2.707	126.471	50.466	48360.054	35334.729	6.827	0.422	1.028	11.791	232555.505	20359.234

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.851	0.291	2.999
Available MBytes	185765.398	185518.000	185866.000
Free System Page Table Entries	33555858.221	33555787.000	33555861.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	349577417.092	348659712.000	350392320.000
Pool Paged Bytes	158988675.209	158515200.000	159596544.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```
5/17/2014 7:24:34 PM -- Preparing for testing ...
5/17/2014 7:24:41 PM -- Attaching databases ...
5/17/2014 7:24:41 PM -- Preparations for testing are complete.
5/17/2014 7:24:41 PM -- Starting transaction dispatch ..
5/17/2014 7:24:41 PM -- Database cache settings: (minimum: 192.0 MB, maximum: 1.5 GB)
5/17/2014 7:24:41 PM -- Database flush thresholds: (start: 15.3 MB, stop: 30.7 MB)
5/17/2014 7:24:48 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
5/17/2014 7:24:48 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
5/17/2014 7:24:49 PM -- Operation mix: Sessions 26, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
5/17/2014 7:24:49 PM -- Performance logging started (interval: 15000 ms).
5/17/2014 7:24:49 PM -- Attaining prerequisites:
5/17/2014 7:26:41 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1458823000.0 (lower bound: 1449551000.0, upper bound:
none)
5/18/2014 7:26:41 PM -- Performance logging has ended.
5/18/2014 7:26:41 PM -- JetInterop batch transaction stats: 331795, 331795, 331795, 331795, 331794 and 331794.
5/18/2014 7:26:41 PM -- Dispatching transactions ends.
5/18/2014 7:26:41 PM -- Shutting down databases ...
5/18/2014 7:26:43 PM -- Instance1604.1 (complete), Instance1604.2 (complete), Instance1604.3 (complete), Instance1604.4 (complete),
Instance1604.5 (complete) and Instance1604.6 (complete)
5/18/2014 7:26:43 PM -- C:\Program Files\Exchange Jetstress\Stress_2014_5_17_19_24_48.blg has 5751 samples.
5/18/2014 7:26:43 PM -- Creating test report ...
5/18/2014 7:27:06 PM -- Instance1604.1 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.1 has 0.3 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.1 has 0.3 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.2 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.2 has 0.4 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.2 has 0.4 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.3 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.3 has 0.4 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.3 has 0.4 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.4 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.4 has 0.4 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.4 has 0.4 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.5 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.5 has 0.4 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.5 has 0.4 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.6 has 17.8 for I/O Database Reads Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.6 has 0.4 for I/O Log Writes Average Latency.
5/18/2014 7:27:06 PM -- Instance1604.6 has 0.4 for I/O Log Reads Average Latency.
5/18/2014 7:27:06 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
5/18/2014 7:27:06 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
5/18/2014 7:27:06 PM -- C:\Program Files\Exchange Jetstress\Stress_2014_5_17_19_24_48.xml has 5743 samples queried.
```

B Appendix B: Performance Testing

Microsoft Exchange Jetstress 2013

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	EQLESRP5
Test Description	
Test Start Time	5/6/2014 3:30:51 PM
Test End Time	5/6/2014 5:32:57 PM
Collection Start Time	5/6/2014 3:32:55 PM
Collection End Time	5/6/2014 5:32:41 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0712.008
Operating System	Windows Server 2012 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2014_5_6_15_31_5.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	1056.711
Target Transactional I/O per Second	600
Initial Database Size (bytes)	10748594356224
Final Database Size (bytes)	10751119327232
Database Files (Count)	6

Jetstress System Parameters

Thread Count	26
Minimum Database Cache	192.0 MB
Maximum Database Cache	1536.0 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

Instance5004.1 Log path: F:\log
Database: F:\db\Jetstress001001.edb

Instance5004.2 Log path: G:\log
Database: G:\db\Jetstress002001.edb

Instance5004.3 Log path: H:\log
Database: H:\db\Jetstress003001.edb

Instance5004.4 Log path: I:\log
Database: I:\db\Jetstress004001.edb

Instance5004.5 Log path: J:\log
Database: J:\db\Jetstress005001.edb

Instance5004.6 Log path: K:\log
Database: K:\db\Jetstress006001.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
Instance5004.1	16.962	0.779	123.928	52.734	32823.898	35939.070	0.000	0.344	0.000	12.592	0.000	20225.485
Instance5004.2	16.974	1.128	123.808	52.528	32824.874	35947.329	0.000	0.364	0.000	12.535	0.000	20266.775
Instance5004.3	16.948	1.427	123.854	52.833	32822.402	35948.278	0.000	0.363	0.000	12.630	0.000	20216.649
Instance5004.4	16.957	1.646	123.230	52.506	32828.801	36023.114	0.000	0.429	0.000	12.675	0.000	20406.211
Instance5004.5	17.012	1.977	123.180	52.310	32828.956	35967.730	0.000	0.427	0.000	12.606	0.000	20347.929
Instance5004.6	16.963	2.318	123.490	52.311	32823.487	35996.869	0.000	0.440	0.000	12.484	0.000	20327.961

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance5004.1	8.605	261903.102
Instance5004.2	8.601	261895.222
Instance5004.3	8.612	261893.915
Instance5004.4	8.617	261854.804
Instance5004.5	8.620	261906.745
Instance5004.6	8.611	261928.672

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance5004.1	1.090	232575.117
Instance5004.2	1.088	232576.198
Instance5004.3	1.093	232546.519
Instance5004.4	1.108	232583.578
Instance5004.5	1.099	232558.104
Instance5004.6	1.086	232571.295

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
Instance5004.1	16.962	0.779	132.532	52.734	47696.637	35939.070	5.597	0.344	1.090	12.592	232575.117	20225.485
Instance5004.2	16.974	1.128	132.409	52.528	47704.886	35947.329	5.833	0.364	1.088	12.535	232576.198	20266.775
Instance5004.3	16.948	1.427	132.466	52.833	47715.595	35948.278	6.514	0.363	1.093	12.630	232546.519	20216.649
Instance5004.4	16.957	1.646	131.846	52.506	47796.295	36023.114	6.471	0.429	1.108	12.675	232583.578	20406.211
Instance5004.5	17.012	1.977	131.800	52.310	47811.394	35967.730	6.584	0.427	1.099	12.606	232558.104	20347.929
Instance5004.6	16.963	2.318	132.101	52.311	47758.535	35996.869	6.463	0.440	1.086	12.484	232571.295	20327.961

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.891	0.156	1.472
Available MBytes	185810.983	185784.000	185843.000
Free System Page Table Entries	33555858.490	33555843.000	33555861.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	348302828.720	347512832.000	348549120.000
Pool Paged Bytes	132903520.402	132775936.000	133054464.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

5/6/2014 3:30:51 PM -- Preparing for testing ...
5/6/2014 3:30:58 PM -- Attaching databases ...
5/6/2014 3:30:58 PM -- Preparations for testing are complete.
5/6/2014 3:30:58 PM -- Starting transaction dispatch ..
5/6/2014 3:30:58 PM -- Database cache settings: (minimum: 192.0 MB, maximum: 1.5 GB)
5/6/2014 3:30:58 PM -- Database flush thresholds: (start: 15.3 MB, stop: 30.7 MB)
5/6/2014 3:31:05 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
5/6/2014 3:31:05 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
5/6/2014 3:31:06 PM -- Operation mix: Sessions 26, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
5/6/2014 3:31:06 PM -- Performance logging started (interval: 15000 ms).
5/6/2014 3:31:06 PM -- Attaining prerequisites:
5/6/2014 3:32:55 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1471398000.0 (lower bound: 1449551000.0, upper bound: none)
5/6/2014 5:32:56 PM -- Performance logging has ended.
5/6/2014 5:32:56 PM -- JetInterop batch transaction stats: 29733, 29733, 29732, 29732, 29732 and 29732.
5/6/2014 5:32:56 PM -- Dispatching transactions ends.
5/6/2014 5:32:56 PM -- Shutting down databases ...
5/6/2014 5:32:57 PM -- Instance5004.1 (complete), Instance5004.2 (complete), Instance5004.3 (complete), Instance5004.4 (complete), Instance5004.5 (complete) and Instance5004.6 (complete)
5/6/2014 5:32:57 PM -- C:\Program Files\Exchange Jetstress\Performance_2014_5_6_15_31_5.blg has 485 samples.
5/6/2014 5:32:58 PM -- Creating test report ...
5/6/2014 5:33:00 PM -- Instance5004.1 has 17.0 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.1 has 0.3 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.1 has 0.3 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.2 has 17.0 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.2 has 0.4 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.2 has 0.4 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.3 has 16.9 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.3 has 0.4 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.3 has 0.4 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.4 has 17.0 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.4 has 0.4 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.4 has 0.4 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.5 has 17.0 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.5 has 0.4 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.5 has 0.4 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.6 has 17.0 for I/O Database Reads Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.6 has 0.4 for I/O Log Writes Average Latency.
5/6/2014 5:33:00 PM -- Instance5004.6 has 0.4 for I/O Log Reads Average Latency.
5/6/2014 5:33:00 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
5/6/2014 5:33:00 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
5/6/2014 5:33:00 PM -- C:\Program Files\Exchange Jetstress\Performance_2014_5_6_15_31_5.xml has 477 samples queried.

C Appendix C Backup Testing

Microsoft Exchange Jetstress 2013

Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance4604.1	1714203.09	04:56:43	96.28
Instance4604.2	1714179.09	04:56:02	96.50
Instance4604.3	1714187.09	04:57:22	96.07
Instance4604.4	1714203.09	04:57:07	96.15
Instance4604.5	1714195.09	04:57:13	96.12
Instance4604.6	1714179.09	04:57:30	96.03
Avg			96.19
Sum			577.17

Jetstress System Parameters

Thread Count	26
Minimum Database Cache	192.0 MB
Maximum Database Cache	1536.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance4604.1 Log path: F:\log
Database: F:\db\Jetstress001001.edb

Instance4604.2 Log path: G:\log
Database: G:\db\Jetstress002001.edb

Instance4604.3 Log path: H:\log
Database: H:\db\Jetstress003001.edb

Instance4604.4 Log path: I:\log
Database: I:\db\Jetstress004001.edb

Instance4604.5 Log path: J:\log
Database: J:\db\Jetstress005001.edb

Instance4604.6 Log path: K:\log
Database: K:\db\Jetstress006001.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
Instance4604.1	4.098	0.000	384.964	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance4604.2	4.054	0.000	385.936	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance4604.3	4.082	0.000	384.269	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance4604.4	4.066	0.000	384.392	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance4604.5	4.075	0.000	384.237	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance4604.6	4.098	0.000	383.805	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	3.752	3.150	4.148
Available MBytes	187512.099	187504.000	187519.000
Free System Page Table Entries	33555834.316	33555747.000	33555861.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	349335620.956	349278208.000	349540352.000
Pool Paged Bytes	163204437.333	163033088.000	163569664.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

5/19/2014 11:54:29 AM -- Preparing for testing ...
5/19/2014 11:54:36 AM -- Attaching databases ...
5/19/2014 11:54:36 AM -- Preparations for testing are complete.
5/19/2014 11:54:43 AM -- Performance logging started (interval: 30000 ms).
5/19/2014 11:54:43 AM -- Backing up databases ...
5/19/2014 4:52:14 PM -- Performance logging has ended.
5/19/2014 4:52:14 PM -- Instance4604.1 (100% processed), Instance4604.2 (100% processed), Instance4604.3 (100% processed), Instance4604.4 (100% processed), Instance4604.5 (100% processed) and Instance4604.6 (100% processed)
5/19/2014 4:52:14 PM -- C:\Program Files\Exchange Jetstress\DatabaseBackup_2014_5_19_11_54_36.blg has 594 samples.
5/19/2014 4:52:14 PM -- Creating test report ...

D Appendix D Soft Recovery Testing

Microsoft Exchange Jetstress 2013

SoftRecovery Test Result Report

Soft-Recovery Statistics – All

Database Instance	Log files replayed	Elapsed seconds
Instance4660.1	501	1014.1696251
Instance4660.2	501	1011.5958485
Instance4660.3	501	1012.1149727
Instance4660.4	501	1008.2606752
Instance4660.5	502	1008.2606752
Instance4660.6	502	1014.9383029
Avg	501	1011.557
Sum	3008	6069.3400996

Database Configuration

Instance4660.1 Log path: F:\log

Database: F:\db\Jetstress001001.edb

Instance4660.2 Log path: G:\log

Database: G:\db\Jetstress002001.edb

Instance4660.3 Log path: H:\log

Database: H:\db\Jetstress003001.edb

Instance4660.4 Log path: I:\log

Database: I:\db\Jetstress004001.edb

Instance4660.5 Log path: J:\log

Database: J:\db\Jetstress005001.edb

Instance4660.6 Log path: K:\log

Database: K:\db\Jetstress006001.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
Instance4660.1	18.476	0.269	316.004	1.976	39878.666	32768.000	14.694	0.000	2.470	0.000	209665.691	0.000
Instance4660.2	18.362	0.256	325.423	1.980	39866.192	32768.000	14.622	0.000	2.475	0.000	209578.031	0.000
Instance4660.3	18.448	0.258	320.963	1.980	39856.789	32768.000	13.938	0.000	2.484	0.000	209564.513	0.000
Instance4660.4	18.366	0.252	323.432	1.988	39914.343	32636.402	15.061	0.000	2.485	0.000	208737.471	0.000
Instance4660.5	18.470	0.262	321.175	1.992	39866.502	32768.000	14.036	0.000	2.495	0.000	209729.053	0.000
Instance4660.6	18.295	0.259	321.727	1.980	39789.586	32636.928	14.090	0.000	2.475	0.000	208787.603	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4660.1	0.000	0.000
Instance4660.2	0.000	0.000
Instance4660.3	0.000	0.000
Instance4660.4	0.000	0.000
Instance4660.5	0.000	0.000
Instance4660.6	0.000	0.000

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
Instance4660.1	18.476	0.269	316.004	1.976	39878.666	32768.000	14.694	0.000	2.470	0.000	209665.691	0.000
Instance4660.2	18.362	0.256	325.423	1.980	39866.192	32768.000	14.622	0.000	2.475	0.000	209578.031	0.000
Instance4660.3	18.448	0.258	320.963	1.980	39856.789	32768.000	13.938	0.000	2.484	0.000	209564.513	0.000
Instance4660.4	18.366	0.252	323.432	1.988	39914.343	32636.402	15.061	0.000	2.485	0.000	208737.471	0.000
Instance4660.5	18.470	0.262	321.175	1.992	39866.502	32768.000	14.036	0.000	2.495	0.000	209729.053	0.000
Instance4660.6	18.295	0.259	321.727	1.980	39789.586	32636.928	14.090	0.000	2.475	0.000	208787.603	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.629	0.000	3.164
Available MBytes	185834.251	185760.000	187348.000
Free System Page Table Entries	33555857.279	33555835.000	33555861.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	351048936.542	350515200.000	351424512.000
Pool Paged Bytes	163920288.127	163889152.000	163954688.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

5/20/2014 1:07:58 PM -- Preparing for testing ...
5/20/2014 1:08:05 PM -- Attaching databases ...
5/20/2014 1:08:05 PM -- Preparations for testing are complete.
5/20/2014 1:08:05 PM -- Starting transaction dispatch ..
5/20/2014 1:08:05 PM -- Database cache settings: (minimum: 192.0 MB, maximum: 1.5 GB)
5/20/2014 1:08:05 PM -- Database flush thresholds: (start: 15.3 MB, stop: 30.7 MB)
5/20/2014 1:08:11 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
5/20/2014 1:08:11 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
5/20/2014 1:08:12 PM -- Operation mix: Sessions 26, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
5/20/2014 1:08:12 PM -- Performance logging started (interval: 15000 ms).
5/20/2014 1:08:12 PM -- Generating log files ...
5/20/2014 2:18:28 PM -- F:\log (100.2% generated), G:\log (100.2% generated), H:\log (100.2% generated), I:\log (100.2% generated), J:\log (100.4% generated) and K:\log (100.4% generated)
5/20/2014 2:18:28 PM -- Performance logging has ended.
5/20/2014 2:18:28 PM -- JetInterop batch transaction stats: 16833, 16833, 16833, 16832, 16832 and 16832.
5/20/2014 2:18:28 PM -- Dispatching transactions ends.
5/20/2014 2:18:29 PM -- Shutting down databases ...
5/20/2014 2:18:30 PM -- Instance4660.1 (complete), Instance4660.2 (complete), Instance4660.3 (complete), Instance4660.4 (complete), Instance4660.5 (complete) and Instance4660.6 (complete)
5/20/2014 2:18:30 PM -- C:\Program Files\Exchange Jetstress\Performance_2014_5_20_13_8_11.blg has 280 samples.
5/20/2014 2:18:30 PM -- Creating test report ...
5/20/2014 2:18:31 PM -- Instance4660.1 has 17.5 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.1 has 0.4 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.1 has 0.4 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.2 has 17.4 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.2 has 0.3 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.2 has 0.3 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.3 has 17.4 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.3 has 0.4 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.3 has 0.4 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.4 has 17.4 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.4 has 0.4 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.4 has 0.4 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.5 has 17.4 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.5 has 0.4 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.5 has 0.4 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.6 has 17.4 for I/O Database Reads Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.6 has 0.4 for I/O Log Writes Average Latency.
5/20/2014 2:18:31 PM -- Instance4660.6 has 0.4 for I/O Log Reads Average Latency.
5/20/2014 2:18:31 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
5/20/2014 2:18:31 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
5/20/2014 2:18:31 PM -- C:\Program Files\Exchange Jetstress\Performance_2014_5_20_13_8_11.xml has 279 samples queried.
5/20/2014 2:18:32 PM -- C:\Program Files\Exchange Jetstress\Performance_2014_5_20_13_8_11.html was saved.
5/20/2014 2:18:33 PM -- Performance logging started (interval: 4000 ms).
5/20/2014 2:18:33 PM -- Recovering databases ...
5/20/2014 2:35:28 PM -- Performance logging has ended.
5/20/2014 2:35:28 PM -- Instance4660.1 (1014.1696251), Instance4660.2 (1011.5958485), Instance4660.3 (1012.1149727), Instance4660.4 (1008.2606752), Instance4660.5 (1008.2606752) and Instance4660.6 (1014.9383029)
5/20/2014 2:35:28 PM -- C:\Program Files\Exchange Jetstress\SoftRecovery_2014_5_20_14_18_32.blg has 251 samples.
5/20/2014 2:35:28 PM -- Creating test report ...