Dell Storage PS 4210E Series 2,500-User Mailbox Resiliency Storage Solution for Microsoft Exchange Server 2013

Microsoft ESRP 4.0

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

Revisions

Date	Description
October 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 Storage PS Series storage array features

With Dell Storage 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

For a comprehensive storage solution, Dell also provides host-based utilities that are all-inclusive in the purchase of your Dell Storage 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. 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 Dell Storage 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)**. Dell Storage 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 Dell Storage 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 Storage PS4210 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 PS4210 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
- Dell Storage Host Software to extend the functionality of the array-based software and enable cooperation with host operating systems, hypervisors and applications
- Host Integration Tools for Microsoft, VMware and Linux products
- Promote server/storage/network consolidation to ease management of the virtual ecosystem
- Support multi-way replication for robust disaster recovery
- SAN HQ management software with an interface for monitoring groups of PS Series arrays to ensure storage resources are utilized for maximum benefit and integrated into Dell Services with SupportAssist
- Dell Storage Update Manager for simplified firmware updates of PS Series platforms

Dell Storage PS4210 Series Arrays provide enhanced storage performance for the small-to-medium enterprise with the simplicity of the Dell Storage product line. Dell Storage PS4210 arrays can drive up to approximately 2GB/sec of throughput per array for sequential, large-block workloads. In addition, flash-enabled Dell Storage PS4210 arrays provide up to 3 times the random performance of prior-generation arrays. The full line of PS4210 Series arrays brings 10GbE iSCSI SAN speed and efficiency to real-world applications, featuring:

- Dual controllers, each with 8GB non-volatile cache
- Two 10GBASE-T RJ45 auto-sensing (10Gb/1Gb/100Mb) ports
- Two 10GbE SFP+ ports for fibre or copper cabling
- Auto negotiation from 1G to 10G to make your environment easier to manage
- Up to 24 hot-pluggable drives, including SAS, NL-SAS and SSD
- Model options for 2.5" drives and 3.5" drives
- Up to 48TB capacity per array

2 Windows and Exchange integration

Windows integration for Dell Storage 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 Storage PS Series arrays to facilitate the deployment, ongoing management, and protection of Dell Storage 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 Dell Storage 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 Solution description

This solution is intended for small-to-medium enterprises that are planning to deploy and virtualize Microsoft Exchange Server 2013 on Microsoft Hyper-V and Dell Storage. It uses the Exchange Server 2013 Database Availability Group (DAG) feature to provide mailbox resiliency and high availability for Exchange users. The solution design represents a virtualized Exchange Server 2013 environment supporting 2,500 users in a mailbox resiliency configuration across two PS4210E storage arrays. See Figure 1 – Conceptual topology.

In this solution, 2,500 users are distributed across two mailbox server virtual machines (vm's) each on a separate Dell PowerEdge R620 Hypervisor host in a single DAG. The DAG has two RAID-protected copies of every Exchange database—a primary (active) copy and a secondary (passive) copy—that are evenly split between hypervisor hosts and storage pools on two PS4210E arrays. Each Exchange database replicates to an alternate mailbox server that resides on a different Hyper-V host through the use of Exchange native DAG host-based log shipping mechanism.

Each mailbox server virtual machine is configured to support up to 2,500 users (1,250 active and 1,250 passive) with a 2 GB mailbox capacity and 0.084 I/O operations per second (IOPS) per user (that includes an additional 20 percent I/O headroom). This user profile corresponds to about 100 messages per user per day.

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 12 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. Connectivity between servers and storage is via iSCSI protocol and two Dell Force10 S4810 switches. See Figure 2 – PS Series topology for Exchange storage solution best practice.

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://www.windowsservercatalog.com/item.aspx?idItem=3063c717-dafe-a93f-567d-2bab1f68c99e&bCatID=1338

4 Simulated environment

The solution presented in this document is designed to simulate a small to 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 PS4210E storage array, or half of the storage solution. The number of users simulated was 2,500 on one server or half of the server solution. The mailbox size was 2GB per user. The server has 4 databases, all active simulating a failover (half of all server and or storage had failed). See Figure 3 – 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.

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 PS4210E storage arrays configured into 2 PS Series groups, each containing one data storage pool with one member (for dB and log data).	24 7200-RPM 2TB Serial Attached SCSI drives (10 active & 2 spare drives per Array)	2 Dell PowerEdge R620 Server, each with 2 Intel Xeon E5-2640 v2 @ 2.5 GHz 8 Core CPU and 164GB memory running MS Windows Server 2012 R2 Datacenter Editionx64	Intel 10 GigE 2P X520

Dell Storage PS Series storage arrays provide active load balancing of storage and connection allocation. 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.

6 Conceptual system design

The elements of the infrastructure supporting the simulated environment, their main relationships and connectivity links are represented in the conceptual diagram in Figure 1.

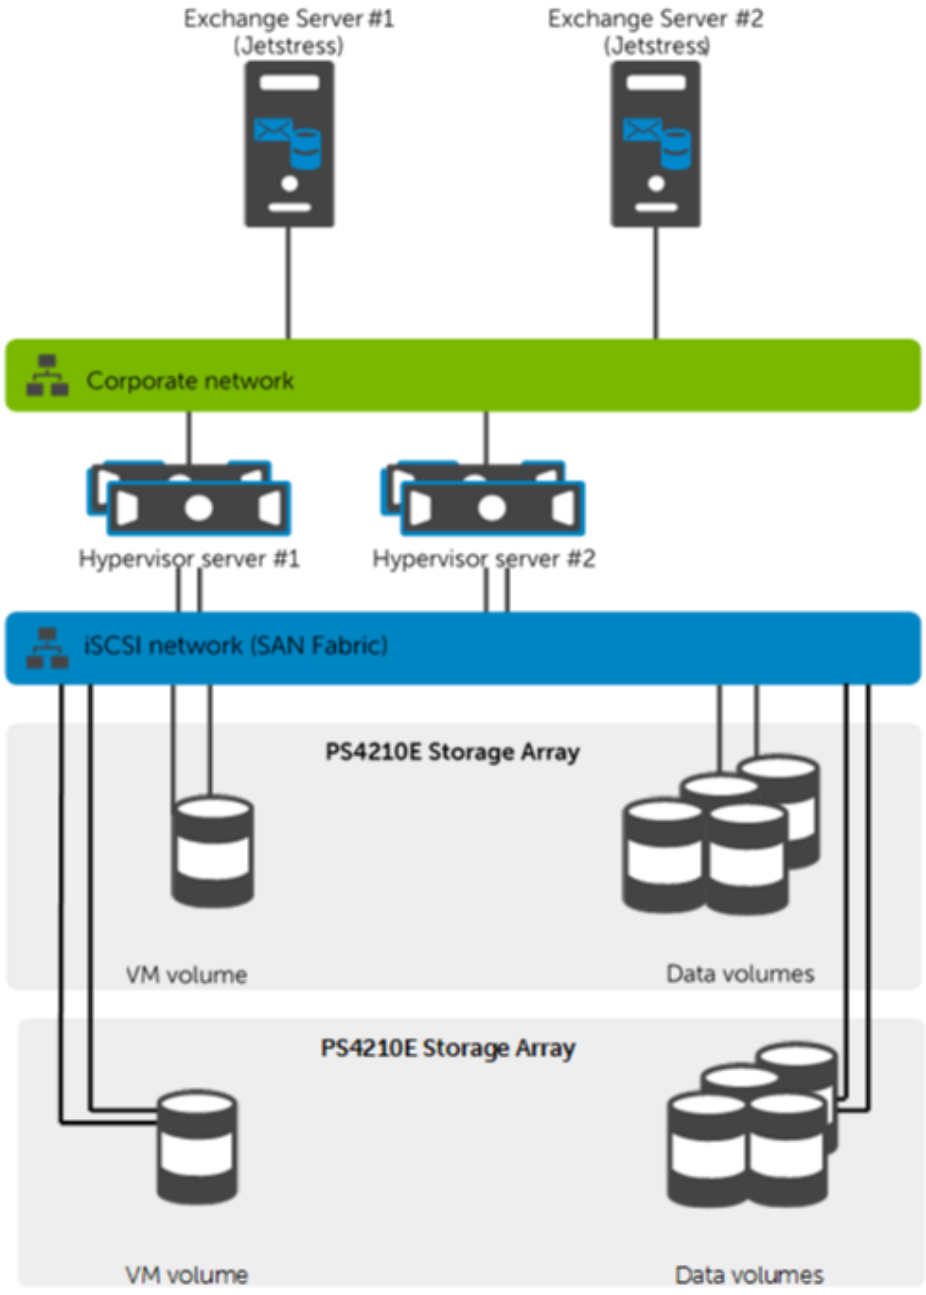

Figure 1 Conceptual topology

7 Exchange DAG physical architecture

This physical architecture utilizes two PS Series array groups consisting of 1 Dell Storage PS4210E storage array per group. There is one data storage pool per group, each consisting of 1 PS4210E array configured as RAID 10. Within each storage pool four volumes are created for a total of 8 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 server virtual machines (VM) in the solution. Each Exchange server VM uses 4 databases and 4 log folders configured in 4 separate volumes. Within the overall solution, 1 server provides 2 active and 2 passive database copies. While the second server provides 2 active and 2 passive database copies. There are a total of 4 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.

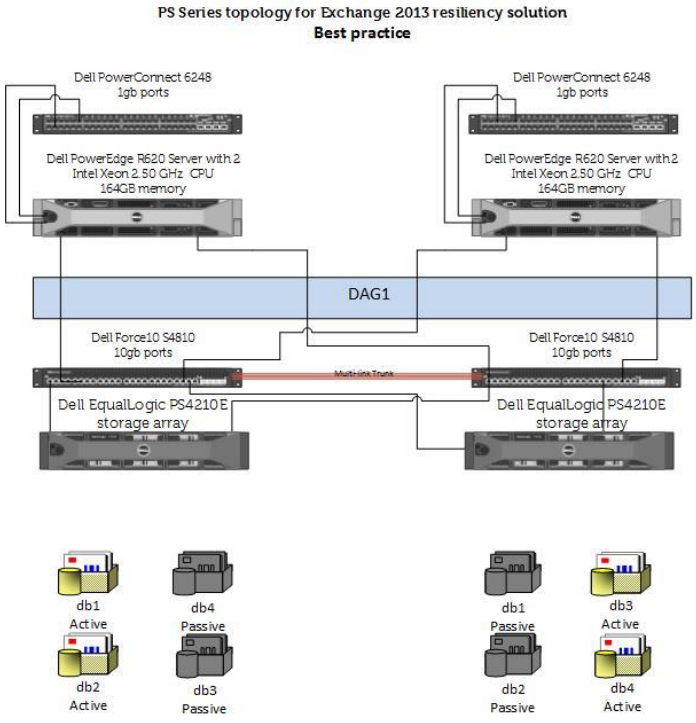

Figure 2 PS Series topology for Exchange storage solution best practice

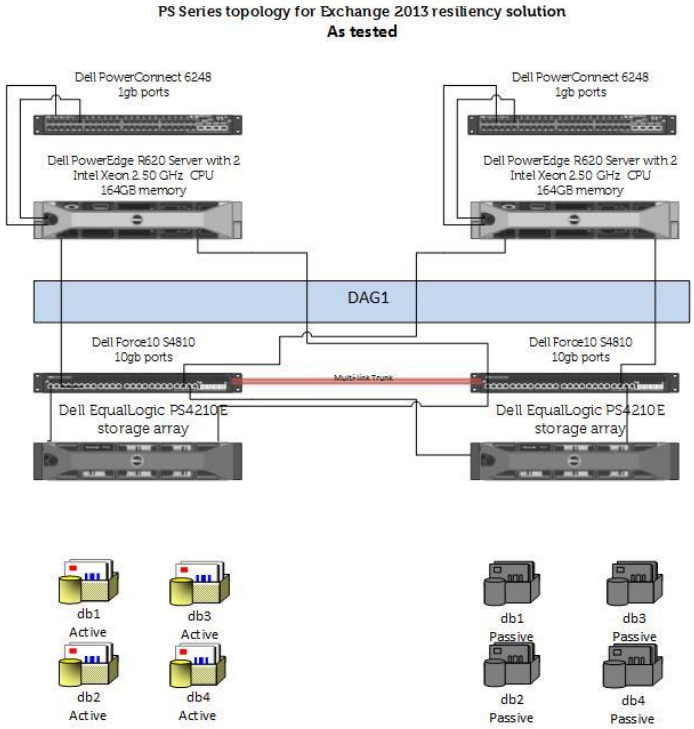

Figure 3 PS Series topology for Exchange storage solution as tested

The following table and figures (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 (see figure 4 and 5).

Table 2 Solution data layout

Server	PS Group	Pool	Volume	DB and Log Group	Active/Passive
Server 1	PS Group 1 consisting of 1 PS4210E array	Data Pool 1 consisting of 1 PS4210E	Volume 1	DB1/Log	Active (copy 1)
			Volume 2	DB2/Log	Active (copy 1)
			Volume 3	DB3/Log	Passive (copy 1)
			Volume 4	DB4/Log	Passive (copy 1)
Server 2	PS Group 2 consisting of 1 PS4210E array	Data Pool 2 consisting of 1 PS4210E	Volume 5	DB1/Log	Passive (copy 2)
			Volume 6	DB2/Log	Passive (copy 2)
			Volume 7	DB3/Log	Active (copy 2)
			Volume 8	DB4/Log	Active (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.

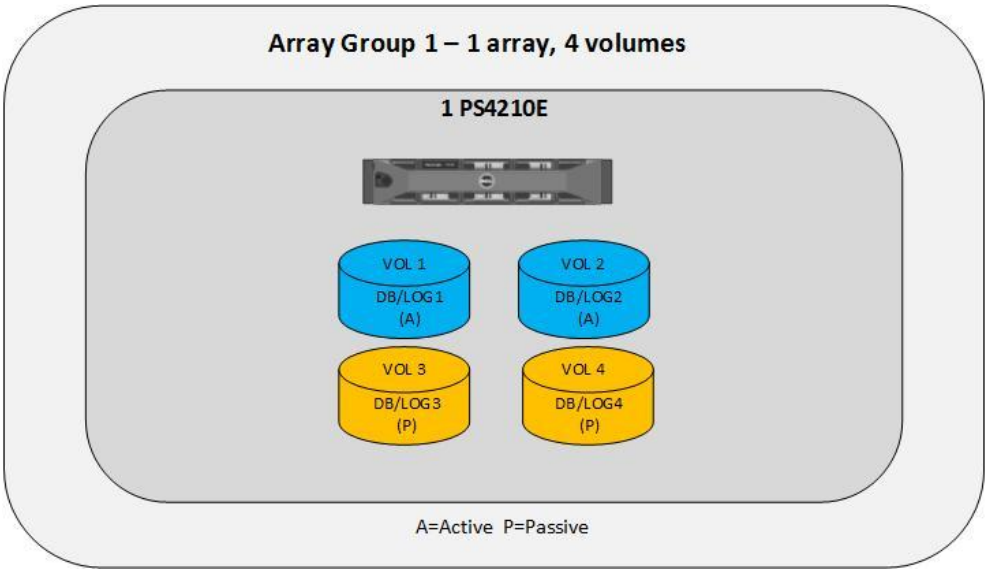

Figure 4 Array Group 1 Exchange volumes

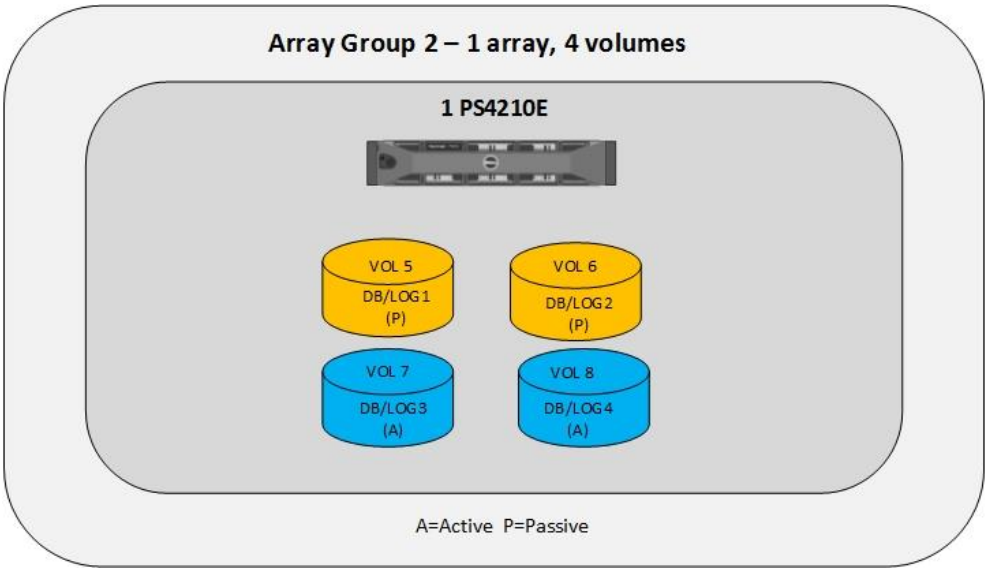

Figure 5 Array Group 2 Exchange volumes

The number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployments.

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

8 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
- 2,500 user mailboxes
- .084 I/O per second per user (includes 20% headroom)
- 2048 MB mailbox quota per user
- 4 databases per server
- 1.62 TB database size
- Mailbox Resiliency (2 copies) provides high availability and is the primary data protection mechanism.

8.1 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.

9 Simulated Exchange configuration

Table 3

Number of Exchange mailboxes simulated	**2,500**
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	2
Number of active mailboxes/server	2500
Number of databases/host	4
Number of copies/database	2
Number of mailboxes/database	625
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.084
Database and Log LUN size	7.8 TB (1950 GB x 4 LUNs)
% storage capacity used by Exchange database[2]	88%

[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.

10 Storage Hardware

Table 4 Hardware

Storage Connectivity	iSCSI
Storage model and OS/firmware revision	Dell Storage PS4210E Firmware Rev:V7.1.0 http://www.windowsservercatalog.com/item.aspx?idItem=08d2110c-d053-5608-fc12-2e2ba0322bfd&bCatID=1338
Storage cache	Dual controllers with 8GB 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	Intel 2P x520 10 GigE
Number of HBAs/host	2
Host server type	Dell PowerEdge R620 Server, each with Intel Xeon E5-2640 v2 @ 2.5 GHz
Total number of drives tested	10 (10 plus 2 Hot Spares)
Maximum number of spindles	10

11 Storage Software

Table 5 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 Storage HIT (Host Integration Toolkit)
Host OS	Windows Server 2012 R2, Datacenter x64 Edition
ESE.dll file version	15.0.847.30
Replication solution name/version	Microsoft Exchange Server 2013 DAG replication

12 Storage disk configuration (Mailbox and Log Store Drives)

Table 6 Storage disk

Drive type, speed and firmware revision	**Seagate 2TB 7200 rpm SAS 6Gb/s drives Model: ST3200044SS Firmware Revision: KN66**
Raw capacity per drive (GB)	1820 GB
Number of physical drives in test	10 (10 active & 2 hot spares)
Total raw storage capacity (GB)	18.2TB (18200 GB)
Drive slice size (GB)	N/A
Number of disks per LUN	Up to 10 (automatically allocated based on load)
Raid level	All storage pools configured as RAID 10
Total formatted capacity (GB)	8.86TB (8860 GB)
Storage capacity utilization	49%
Database capacity utilization	34.40%

13 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 Storage 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

14 Additional information

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

http://support.dell.com/Dell Storage

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

15 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.

15.1 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.

15.2 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.

15.3 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	255.24
Database Drive Reads/sec	176.67
Database Drive Writes/sec	78.57
Average Database Drive Read Latency (ms)	15.58
Average Database Drive Write Latency (ms)	.622
Transaction Log I/O	
Log Drive Writes/sec	19.25
Average Log Drive Write Latency (ms)	.611

16 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).

16.1 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	59.39(Average)
MB read/sec total per server	237.57(Sum)

16.2 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)	2.45 (avg. resp. to replay log / avg. # of logs replayed)

17 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

Test Summary

Overall Test Result	**Pass**
Machine Name	R6U29R620DC
Test Description	Stress
Test Start Time	10/17/2014 2:17:49 PM
Test End Time	10/18/2014 2:23:36 PM
Collection Start Time	10/17/2014 2:23:34 PM
Collection End Time	10/18/2014 2:23:30 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0847.030
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\ESRP\Stress_2014_10_17_14_17_58.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	240.945
Target Transactional I/O per Second	210
Initial Database Size (bytes)	5380469948416
Final Database Size (bytes)	5387415715840
Database Files (Count)	4

Jetstress System Parameters

Thread Count	6
Minimum Database Cache	128.0
MB **Maximum Database Cache**	1024.0
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

Instance984.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance984.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance984.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance984.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance984.1	16.001	0.518	41.453	18.812	33930.995	34825.675	0.000	0.859	0.000	4.419	0.000	20584.577
Instance984.2	15.123	0.534	41.458	18.737	33993.342	34810.967	0.000	0.864	0.000	4.400	0.000	20563.336
Instance984.3	15.070	0.536	41.541	18.787	33968.123	34810.845	0.000	0.837	0.000	4.406	0.000	20514.798
Instance984.4	15.095	0.544	41.448	18.709	33931.902	34836.993	0.000	0.869	0.000	4.389	0.000	20565.477

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance984.1	8.667	261842.176
Instance984.2	8.717	261785.315
Instance984.3	8.728	261818.568
Instance984.4	8.720	261812.215

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance984.1	0.386	150578.939
Instance984.2	0.384	149678.228
Instance984.3	0.383	149164.600
Instance984.4	0.384	149646.091

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
Instance984.1	16.001	0.518	50.120	18.812	73341.933	34825.675	3.918	0.859	0.386	4.419	150578.939	20584.577
Instance984.2	15.123	0.534	50.175	18.737	73568.221	34810.967	6.472	0.864	0.384	4.400	149678.228	20563.336
Instance984.3	15.070	0.536	50.269	18.787	73529.920	34810.845	6.599	0.837	0.383	4.406	149164.600	20514.798
Instance984.4	15.095	0.544	50.168	18.709	73542.530	34836.993	6.804	0.869	0.384	4.389	149646.091	20565.477

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.857	0.164	6.478
Available MBytes	163627.015	163549.000	163699.000
Free System Page Table Entries	16494747.365	16494345.000	16494935.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	617158461.664	616894464.000	618299392.000
Pool Paged Bytes	109799809.112	109199360.000	114884608.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

10/17/2014 2:17:49 PM -- Preparing for testing ...
10/17/2014 2:17:53 PM -- Attaching databases ...
10/17/2014 2:17:53 PM -- Preparations for testing are complete.
10/17/2014 2:17:53 PM -- Starting transaction dispatch ..
10/17/2014 2:17:53 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
10/17/2014 2:17:53 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
10/17/2014 2:17:58 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
10/17/2014 2:17:58 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
10/17/2014 2:17:58 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/17/2014 2:17:58 PM -- Performance logging started (interval: 15000 ms).
10/17/2014 2:17:58 PM -- Attaining prerequisites:
10/17/2014 2:23:34 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 968609800.0 (lower bound: 966367600.0, upper bound: none)
10/18/2014 2:23:34 PM -- Performance logging has ended.
10/18/2014 2:23:34 PM -- JetInterop batch transaction stats: 124392, 124391, 124391 and 124391.
10/18/2014 2:23:34 PM -- Dispatching transactions ends.
10/18/2014 2:23:35 PM -- Shutting down databases ...
10/18/2014 2:23:36 PM -- Instance984.1 (complete), Instance984.2 (complete), Instance984.3 (complete) and Instance984.4 (complete)
10/18/2014 2:23:36 PM -- C:\ESRP\Stress_2014_10_17_14_17_58.blg has 5776 samples.
10/18/2014 2:23:37 PM -- Creating test report ...
10/18/2014 2:24:22 PM -- Instance984.1 has 16.0 for I/O Database Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.1 has 0.9 for I/O Log Writes Average Latency.
10/18/2014 2:24:22 PM -- Instance984.1 has 0.9 for I/O Log Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.2 has 15.1 for I/O Database Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.2 has 0.9 for I/O Log Writes Average Latency.
10/18/2014 2:24:22 PM -- Instance984.2 has 0.9 for I/O Log Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.3 has 15.1 for I/O Database Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.3 has 0.8 for I/O Log Writes Average Latency.
10/18/2014 2:24:22 PM -- Instance984.3 has 0.8 for I/O Log Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.4 has 15.1 for I/O Database Reads Average Latency.
10/18/2014 2:24:22 PM -- Instance984.4 has 0.9 for I/O Log Writes Average Latency.
10/18/2014 2:24:22 PM -- Instance984.4 has 0.9 for I/O Log Reads Average Latency.
10/18/2014 2:24:22 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/18/2014 2:24:22 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/18/2014 2:24:22 PM -- C:\ESRP\Stress_2014_10_17_14_17_58.xml has 5753 samples queried.

B Appendix B: Performance Testing

Microsoft Exchange Jetstress 2013

Test Summary

Overall Test Result	**Pass**
Machine Name	R6U29R620DC
Test Description	Performance 2h Thread count equals 6
Test Start Time	10/8/2014 8:21:35 AM
Test End Time	10/8/2014 10:27:22 AM
Collection Start Time	10/8/2014 8:27:20 AM
Collection End Time	10/8/2014 10:27:08 AM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0847.030
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\ESRP\Performance_2014_10_8_8_21_44.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	255.235
Target Transactional I/O per Second	210
Initial Database Size (bytes)	5372182003712
Final Database Size (bytes)	5372844703744
Database Files (Count)	4

Jetstress System Parameters

Thread Count	6
Minimum Database Cache	128.0
MB **Maximum Database Cache**	1024.0
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

Instance3488.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance3488.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance3488.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance3488.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance3488.1	16.340	0.648	44.113	19.678	33031.000	36065.317	0.000	0.590	0.000	4.829	0.000	20493.021
Instance3488.2	15.290	0.624	44.250	19.717	33045.023	36056.293	0.000	0.587	0.000	4.823	0.000	20428.696
Instance3488.3	15.370	0.585	44.205	19.661	33010.994	36130.281	0.000	0.608	0.000	4.820	0.000	20533.994
Instance3488.4	15.312	0.632	44.102	19.509	33026.240	36105.762	0.000	0.659	0.000	4.773	0.000	20394.257

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3488.1	8.588	261780.294
Instance3488.2	8.709	261811.376
Instance3488.3	8.689	261906.623
Instance3488.4	8.710	261764.906

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3488.1	0.421	165028.219
Instance3488.2	0.419	163500.918
Instance3488.3	0.420	163486.337
Instance3488.4	0.414	160540.587

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
Instance3488.1	16.340	0.648	52.701	19.678	70308.684	36065.317	6.248	0.590	0.421	4.829	165028.219	20493.021
Instance3488.2	15.290	0.624	52.958	19.717	70663.776	36056.293	6.491	0.587	0.419	4.823	163500.918	20428.696
Instance3488.3	15.370	0.585	52.894	19.661	70612.211	36130.281	6.601	0.608	0.420	4.820	163486.337	20533.994
Instance3488.4	15.312	0.632	52.812	19.509	70749.586	36105.762	7.158	0.659	0.414	4.773	160540.587	20394.257

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.912	0.416	2.862
Available MBytes	164125.259	164107.000	164185.000
Free System Page Table Entries	16496248.775	16496047.000	16496387.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	68343581.395	68227072.000	68481024.000
Pool Paged Bytes	144968633.453	144769024.000	145207296.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

10/8/2014 8:21:35 AM -- Preparing for testing ...
10/8/2014 8:21:40 AM -- Attaching databases ...
10/8/2014 8:21:40 AM -- Preparations for testing are complete.
10/8/2014 8:21:40 AM -- Starting transaction dispatch ..
10/8/2014 8:21:40 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
10/8/2014 8:21:40 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
10/8/2014 8:21:44 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
10/8/2014 8:21:44 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
10/8/2014 8:21:45 AM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/8/2014 8:21:45 AM -- Performance logging started (interval: 15000 ms).
10/8/2014 8:21:45 AM -- Attaining prerequisites:
10/8/2014 8:27:20 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 969179100.0 (lower bound: 966367600.0, upper bound: none)
10/8/2014 10:27:21 AM -- Performance logging has ended.
10/8/2014 10:27:21 AM -- JetInterop batch transaction stats: 11868, 11868, 11867 and 11867.
10/8/2014 10:27:21 AM -- Dispatching transactions ends.
10/8/2014 10:27:21 AM -- Shutting down databases ...
10/8/2014 10:27:22 AM -- Instance3488.1 (complete), Instance3488.2 (complete), Instance3488.3 (complete) and Instance3488.4 (complete)
10/8/2014 10:27:22 AM -- C:\ESRP\Performance_2014_10_8_8_21_44.blg has 501 samples
10/8/2014 10:27:22 AM -- Creating test report ...
10/8/2014 10:27:48 AM -- Instance3488.1 has 16.3 for I/O Database Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.1 has 0.6 for I/O Log Writes Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.1 has 0.6 for I/O Log Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.2 has 15.3 for I/O Database Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.2 has 0.6 for I/O Log Writes Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.2 has 0.6 for I/O Log Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.3 has 15.4 for I/O Database Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.3 has 0.6 for I/O Log Writes Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.3 has 0.6 for I/O Log Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.4 has 15.3 for I/O Database Reads Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.4 has 0.7 for I/O Log Writes Average Latency.
10/8/2014 10:27:48 AM -- Instance3488.4 has 0.7 for I/O Log Reads Average Latency.
10/8/2014 10:27:48 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/8/2014 10:27:48 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/8/2014 10:27:48 AM -- C:\ESRP\Performance_2014_10_8_8_21_44.xml has 478 samples queried.

C Appendix C Backup Testing

Microsoft Exchange Jetstress 2013

Database backup Test Result Report

Database Backup Statistics - All

Microsoft Exchange Jetstress 2013 Database backup Test Result Report

Database Backup Statistics - All

Database	Database Size	Elapsed Backup	MBytes
Instance504.	1282584.03	05:28:44	65.03
Instance504.	1282576.03	06:11:20	57.57
Instance504.	1282568.03	06:12:16	57.42
Instance504.	1282560.03	06:11:22	57.56
Avg			59.39
Sum			237.57

Jetstress System Parameters

Thread Count 6

Minimum Database Cache 128.0 MB

Maximum Database Cache 1024.0 MB

Insert Operations 40%

Delete Operations 20%

Replace 5%

Read Operations 35%

Lazy Commits 70%

Database Configuration

Instance504.1 Log path: E:\log1

Database: E:\db1\Jetstress001001.edb

Instance504.2 Log path: F:\log2

Database: F:\db2\Jetstress002001.edb

Instance504.3 Log path: G:\log3

Database: G:\db3\Jetstress003001.edb

Instance504.4 Log path: H:\log4

Database: H:\db4\Jetstress004001.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
Instance504.1	6.668	0.000	260.093	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance504.2	7.618	0.000	230.240	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance504.3	7.651	0.000	229.614	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance504.4	7.621	0.000	230.182	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	6.156	3.617	7.578
Available MBytes	164962.821	164941.000	164995.000
Free System Page Table Entries	16494719.895	16494404.000	16494896.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	309379556.473	309174272.000	309575680.000
Pool Paged Bytes	148992203.699	148819968.000	149323776.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```
10/14/2014 7:58:15 AM -- Preparing for testing ...
10/14/2014 7:58:20 AM -- Attaching databases ...
10/14/2014 7:58:20 AM -- Preparations for testing are complete.
10/14/2014 7:58:24 AM -- Performance logging started (interval: 30000 ms).
10/14/2014 7:58:24 AM -- Backing up databases ...
10/14/2014 2:10:41 PM -- Performance logging has ended.
10/14/2014 2:10:41 PM -- Instance504.1 (100% processed), Instance504.2 (100% processed), Instance504.3 (100% processed) and Instance504.4
(100% processed)
10/14/2014 2:10:41 PM -- C:\ESRP\DatabaseBackup_2014_10_14_7_58_20.blg has 744 samples.
10/14/2014 2:10:41 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
Instance1984.1	503	1283.3589814
Instance1984.2	511	1217.8080929
Instance1984.3	502	1214.5363269
Instance1984.4	501	1214.033728
Avg	504	1232.434
Sum	2017	4929.7371292

Database Configuration

Instance1984.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance1984.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance1984.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance1984.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance1984.1	37.246	0.332	167.366	1.565	38580.421	26276.226	20.826	0.000	1.957	0.000	166648.164	0.000
Instance1984.2	33.687	0.370	173.661	1.674	38666.356	29024.636	25.582	0.000	2.100	0.000	181911.299	0.000
Instance1984.3	33.572	0.365	174.366	1.650	38714.787	28903.336	24.828	0.000	2.063	0.000	181176.780	0.000
Instance1984.4	33.778	0.366	174.105	1.647	38566.656	28631.176	23.770	0.000	2.066	0.000	180600.836	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1984.1	0.000	0.000
Instance1984.2	0.000	0.000
Instance1984.3	0.000	0.000
Instance1984.4	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
Instance1984.1	37.246	0.332	167.366	1.565	38580.421	26276.226	20.826	0.000	1.957	0.000	166648.164	0.000
Instance1984.2	33.687	0.370	173.661	1.674	38666.356	29024.636	25.582	0.000	2.100	0.000	181911.299	0.000
Instance1984.3	33.572	0.365	174.366	1.650	38714.787	28903.336	24.828	0.000	2.063	0.000	181176.780	0.000
Instance1984.4	33.778	0.366	174.105	1.647	38566.656	28631.176	23.770	0.000	2.066	0.000	180600.836	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	2.280	0.000	11.611
Available MBytes	163941.058	163876.000	164949.000
Free System Page Table Entries	16494951.559	16494803.000	16495093.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	307626266.926	307515392.000	307974144.000
Pool Paged Bytes	148119359.096	148115456.000	148127744.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

10/15/2014 8:45:00 AM -- Preparing for testing ...
10/15/2014 8:45:05 AM -- Attaching databases ...
10/15/2014 8:45:05 AM -- Preparations for testing are complete.
10/15/2014 8:45:05 AM -- Starting transaction dispatch ..
10/15/2014 8:45:05 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
10/15/2014 8:45:05 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
10/15/2014 8:45:09 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
10/15/2014 8:45:09 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
10/15/2014 8:45:10 AM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/15/2014 8:45:10 AM -- Performance logging started (interval: 15000 ms).
10/15/2014 8:45:10 AM -- Generating log files ...
10/15/2014 11:40:44 AM -- E:\log1 (100.6% generated), F:\log2 (102.2% generated), G:\log3 (100.4% generated) and H:\log4 (100.2% generated)
10/15/2014 11:40:44 AM -- Performance logging has ended.
10/15/2014 11:40:44 AM -- JetInterop batch transaction stats: 16833, 16833, 16833 and 16833.
10/15/2014 11:40:44 AM -- Dispatching transactions ends.
10/15/2014 11:40:44 AM -- Shutting down databases ...
10/15/2014 11:40:45 AM -- Instance1984.1 (complete), Instance1984.2 (complete), Instance1984.3 (complete) and Instance1984.4 (complete)
10/15/2014 11:40:45 AM -- C:\ESRP\Performance_2014_10_15_8_45_9.blg has 701 samples.
10/15/2014 11:40:45 AM -- Creating test report ...
10/15/2014 11:41:13 AM -- Instance1984.1 has 16.6 for I/O Database Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.1 has 0.7 for I/O Log Writes Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.1 has 0.7 for I/O Log Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.2 has 15.7 for I/O Database Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.2 has 0.6 for I/O Log Writes Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.2 has 0.6 for I/O Log Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.3 has 15.7 for I/O Database Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.3 has 0.6 for I/O Log Writes Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.3 has 0.6 for I/O Log Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.4 has 15.7 for I/O Database Reads Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.4 has 0.7 for I/O Log Writes Average Latency.
10/15/2014 11:41:13 AM -- Instance1984.4 has 0.7 for I/O Log Reads Average Latency.
10/15/2014 11:41:13 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/15/2014 11:41:13 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/15/2014 11:41:14 AM -- C:\ESRP\Performance_2014_10_15_8_45_9.xml has 700 samples queried.
10/15/2014 11:41:14 AM -- C:\ESRP\Performance_2014_10_15_8_45_9.html was saved.
10/15/2014 11:41:15 AM -- Performance logging started (interval: 2000 ms).
10/15/2014 11:41:15 AM -- Recovering databases ...
10/15/2014 12:02:39 PM -- Performance logging has ended.
10/15/2014 12:02:39 PM -- Instance1984.1 (1283.3589814), Instance1984.2 (1217.8080929), Instance1984.3 (1214.5363269) and Instance1984.4 (1214.033728)
10/15/2014 12:02:39 PM -- C:\ESRP\SoftRecovery_2014_10_15_11_41_14.blg has 637 samples.
10/15/2014 12:02:39 PM -- Creating test report ...