Lenovo Storage S2200 and System x3550 M5 Server 750 Mailbox Resiliency Solution for Exchange Server 2013

Last update: **04 January 2016**

Describes a reference architecture for Exchange 2013 on Lenovo Storage S2200

Contains performance data for sizing recommendations

Includes details about database distribution and high-availability features

Contains a detailed bill of materials for servers, networking, and storage

Roland Mueller

LENOVO PRESS

1 Introduction

This document describes a low-cost solution for Microsoft Exchange Server 2013 that is based on Lenovo System x3550 M5 servers and Lenovo Storage S2200. The intended audience of this document is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing Microsoft Exchange Server 2013. The solution is based on the Microsoft Exchange Solution Reviewed Program (ESRP) that was developed by Microsoft Corporation. This program provides a common storage testing framework for vendors to provide information on its storage solutions for Microsoft Exchange Server software.

For more details on the Microsoft ESRP Storage program, visit the following website:
technet.microsoft.com/en-us/exchange/ff182054.aspx

Microsoft Exchange Server 2013 is the market leader in the enterprise messaging and collaboration market. Exchange Server 2013 builds upon the Exchange Server 2010 architecture and is redesigned for simplicity of scale, improved hardware utilization, and increased failure isolation. Exchange Server 2013 brings a rich set of technologies, features, and services to the Exchange Server product line. Its goal is to support people and organizations as their work habits evolve from a communication focus to a collaboration focus. At the same time, Exchange Server 2013 helps lower the total cost of ownership.

This solution targets small- to medium-sized organizations that are implementing Microsoft Exchange Server 2013. The solution described in this document provides a site resilient, highly available infrastructure that runs in two data centers, with each data center having two x3550 M5 servers and a dual-controller S2200 storage system.

This document provides the design considerations and best practices for implementing the described architecture to support small- to medium-sized organizations with fewer than 750 employees. However, the principles and techniques that are described throughout this document can be expanded upon to support much larger user populations with the addition of storage and compute resources.

1.1 Disclaimer

This document was 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 content of this document.

Please see Lenovo’s disclaimer at the end of this document.

2 Features of the solution

This section describes the hardware featured in this solution.

2.1 Lenovo System x3550 M5

The Lenovo System x3550 M5 server delivers the performance and reliability required for business-critical applications (such as Exchange Server 2013). Lenovo System x3550 M5 servers can be equipped with up to two 18-core E5-2600 v3 series processors and up to 1.5 TB of TruDDR4 memory. Up to three PCIe 3.0 expansion slots, four integrated 1 gigabit Ethernet (GbE) network ports, and an optional embedded dual-port 10/40 GbE network adapter provide ports for your data and storage connections.

The Lenovo System x3550 M5 includes an on-board RAID controller and the choice of hot-swap SAS or SATA hard disk drives (HDDs) and small form factor (SFF) hot-swap solid-state drives (SSDs). The x3550 M5 supports a maximum of 24 TB of internal storage.

The x3550 M5 supports the following components:

- Up to 10 front- and two rear-mounted SFF HDDs or SSDs
- Up to four 3.5-inch HDDs

The x3550 M5 also supports remote management via the Lenovo Integrated Management Module v2 (IMM2), which enables continuous management capabilities. All of these key features and others help solidify the dependability Lenovo customers are accustomed to with System x servers.

The Lenovo x3550 M5 server is shown in Figure 1.

Figure 1. Lenovo System x3550 M5

For more information, go to: http://lenovopress.com/tips1194-system-x3550-m5

2.2 Lenovo Storage S2200

The Lenovo Storage S2200 storage area network (SAN) array offers simplicity, speed, and scalability. With its easy-to-use management interface, the S2200 makes complex administrative storage tasks simple.

There's no need for a dedicated resource to manage this array. Setup is simple using the intuitive user interface. The Lenovo S2200 offers a great balance between simplicity and performance.

It supports up to 96 drives and can have mixed large form factor (LFF) and SFF enclosures in the same array, enabling this device to provide you with the flexibility you need in your organization. And the S2200 provides in-chassis controller upgrades without data migration. Just swap out the S2200 controller with the S3200 controller and you can experience greater performance, increased fabric support, and greater capacity. The S2200 is designed not only to fit into your infrastructure, but to grow with it.

The Lenovo Storage S2200 is shown in Figure 2.

Figure 2. Lenovo Storage S2200

The Lenovo Storage S2200 includes the following benefits:

- Entry SAN with enterprise features
  - o Simple management graphical user interface (GUI); no storage administration needed
  - o Scale-out options with Thin Provisioning allows you to grow your SAN in pace with your business
  - o High-performance storage to support demanding applications at entry price points
- Fits into your existing environment with flexible configurations and options
  - o Supports up to 96 drives
  - o Fits into various environments with Fibre Channel or iSCSI fabrics
  - o Mix LFF and SFF enclosures in the same array
- Optimized to grow with your business
  - o In-chassis controller upgrades without data migrations or downtime
  - o Designed to be always on with 99.999% availability
  - o High availability configurations: dual controllers, redundant fans, hot-swap drives and fans

For more information, go to: http://lenovopress.com/tips1298-lenovo-storage-s2200

2.3 Emulex 8 Gb FC HBA for Lenovo System x

Streamlined installation and management, plus unrivalled scalability and industry-leading virtualization support make the Emulex 8 Gb Fibre Channel single-port and dual-port host bus adapters ideal solutions for enterprise and mixed operating system environments. With powerful management tools, they deliver maximum performance in the widest range of applications and environments.

The Emulex 8 Gb FC HBA (dual-port and single-port models) is shown in Figure 3.

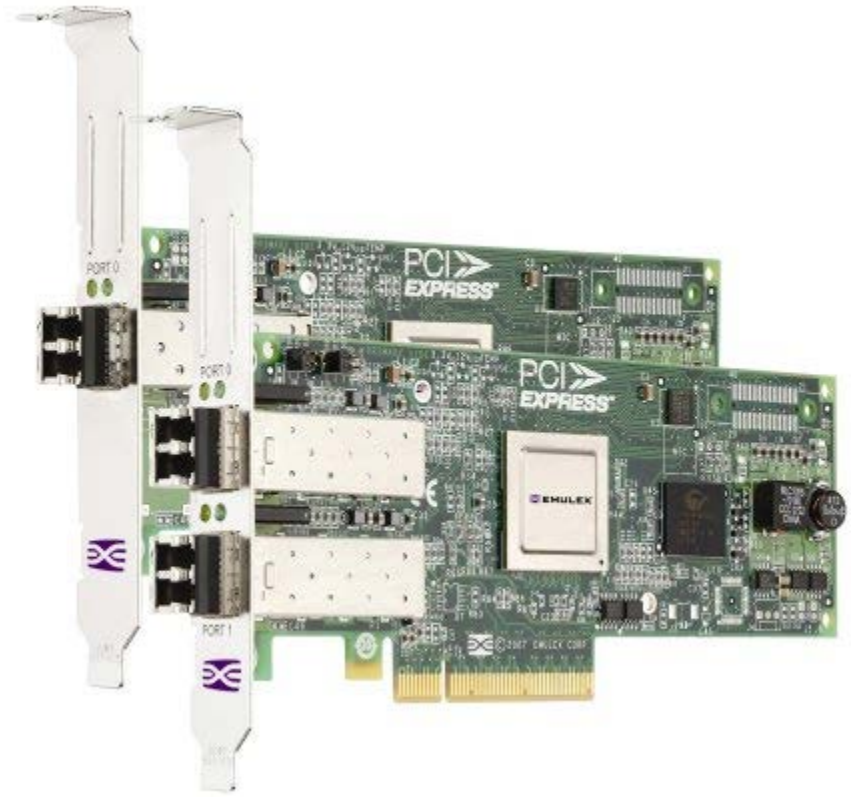

Figure 3. Emulex 8 Gb FC HBA for Lenovo System x servers

The adapters have the following features:

- Standard PCI Express card with low-profile form factor
- One or two independent Fibre Channel ports with optical modules included
- PCI Express 2.0 host interface
- Support for standard and low-profile PCI Express slots
- Auto-negotiation between 2 Gb, 4 Gb, and 8 Gb link attachments
- Processes over 200K I/O operations per second (IOPS), per port
- Support for up to 1600 MB per second maximum at full-duplex, per port

For more information, go to: http://lenovopress.com/tips0727-emulex-8gb-fc-hba

2.4 Lenovo RackSwitch G8124E

The Lenovo RackSwitch G8124E is a 10 Gigabit Ethernet switch specifically designed for the data center and provides a virtualized, cooler, and easier network solution. The G8124E offers 24 10 GbE ports in a high-density, 1U footprint.

Designed with ultra-low latency and top performance in mind, the RackSwitch G8124E provides line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data-center-grade buffers keep traffic moving. The G8124E also supports Converged Enhanced Ethernet (CEE) and Data Center Bridging for support of FCoE and can be used for NAS or iSCSI.

The G8124E is virtualized and supports VMready technology, which is an innovative, standards-based solution to manage VMs in small to large-scale data center and cloud environments. VMready works with all leading

VM providers. The G8124E also supports Virtual Fabric, which allows for the carving up of a physical NIC into 2-to-8 vNICs for improved performance, availability, and security, while reducing cost and complexity.

The G8124E runss cool and implements a choice of directional cooling to maximize data center layout and provisioning. Its superior airflow design complements the hot-aisle and cold-aisle data center cooling model.

The G8124E is easy to use, with server-oriented provisioning via point-and-click management interfaces. Its industry-standard CLI and easy interoperability simplifies configuration for those familiar with Cisco environments.

Figure 4 shows the Lenovo RackSwitch G8124E.

Figure 4. Lenovo RackSwitch G8124E

The G8124E includes the following benefits:

- A total of 24 SFP+ ports that operate at 10 Gb or 1 Gb Ethernet speeds
- Optimal for high-performance computing and applications that require high bandwidth and low latency
- All ports are nonblocking 10 Gb Ethernet with deterministic latency of 570 nanoseconds
- VMready helps reduce configuration complexity and improves security levels in virtualized environments
- Variable-speed fans automatically adjust as needed, which helps to reduce energy consumption
- Easy, standards-based integration into Cisco and other networks helps reduce downtime and learning curve

For more information, go to: http://lenovopress.com/tips0787

3 Solution description

Figure 5 shows the architectural overview of a multisite Exchange Server 2013 deployment with two data centers. Each data center has two x3550 M5 servers, two 10 GbE network switches, one Layer 4 network load balancer, and one S2200 storage system. Multiple paths connect the servers to the networking and storage infrastructure to maintain access to critical resources if there is a planned or unplanned outage.

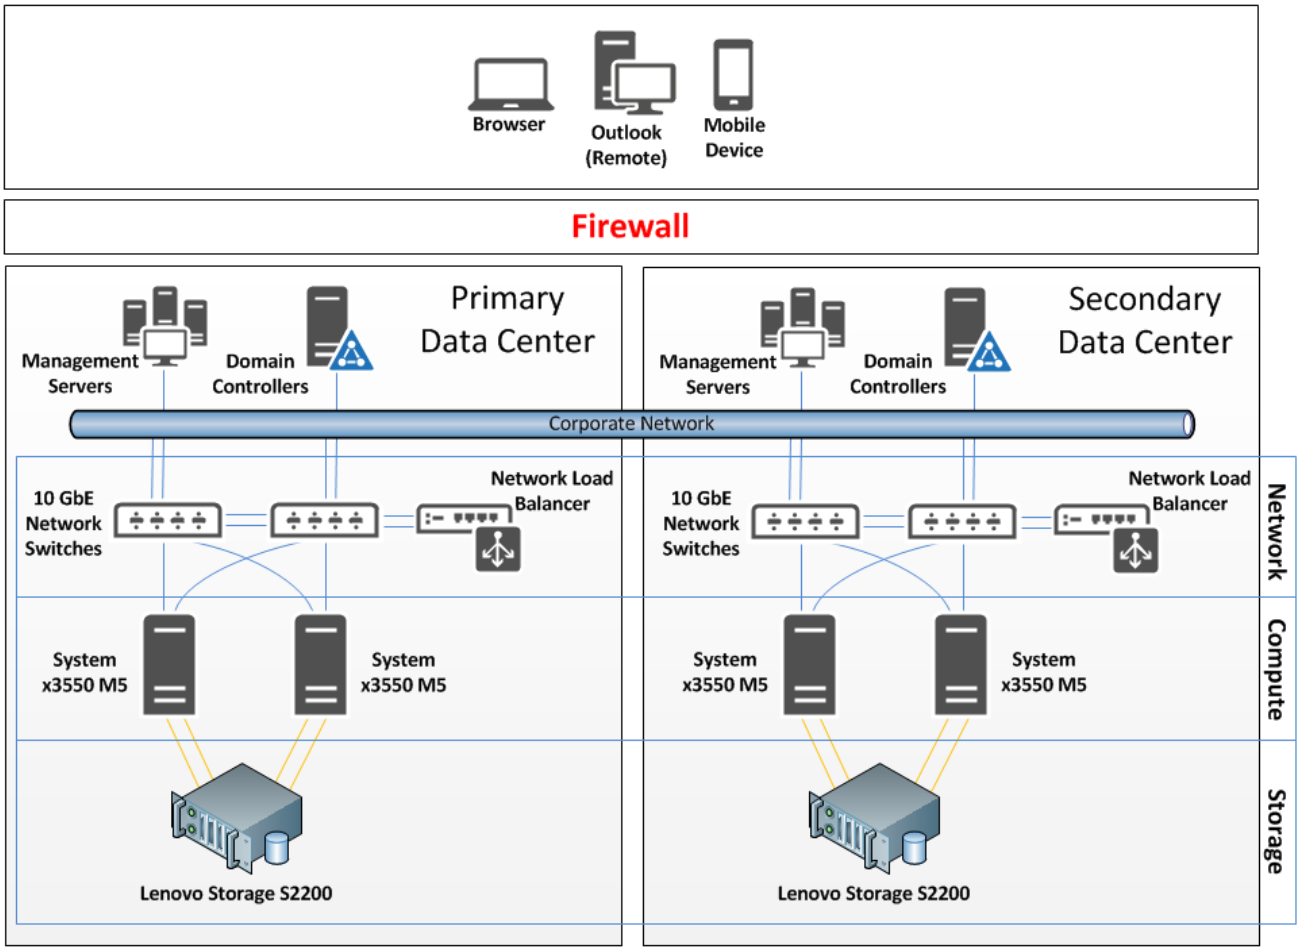

Figure 5. Architectural overview

3.1 High availability

A database availability group (DAG) is the base component of the high availability and site resilience framework that is built into Microsoft Exchange Server 2013. A DAG is a group of up to 16 mailbox servers that hosts a set of mailbox databases and provides automatic database-level recovery from failures that affect individual servers or databases.

A DAG is a boundary for mailbox database replication, database and server switchovers, failovers, and an internal component called *Active Manager*. Active Manager, which runs on every server in a DAG, manages switchovers and failovers.

Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it works with the other servers in the DAG to provide automatic recovery from failures that affect mailbox databases (such as a drive failure or server failure).

Figure 6 shows an example of design for a DAG. An active/passive, multisite implementation with a user population of 750 requires a DAG that spans both data centers. Two mailbox servers with the Mailbox and the Client Access Server (CAS) roles installed are required to host the active copy, the two passive copies, and a lagged copy of each mailbox database.

A lagged mailbox database copy is a mailbox database copy configured with a replay lag time value greater than 0. A lagged database copy is one that is not updated by replaying transactions as they become available. Instead, the transaction logs are kept for a certain period and then replayed. Because the lagged database copy retains logs for a longer period of time, more space is required to house them. This configuration retains logs for 72 hours, which allows administrators to recover from logical corruption that occurred up to three days in the past.

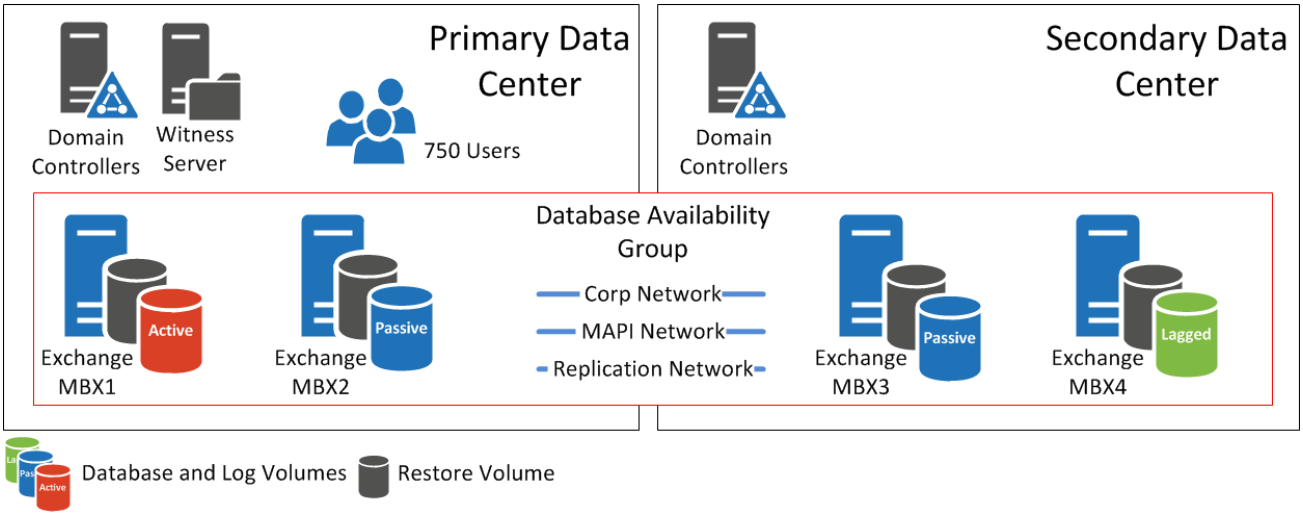

Figure 6. DAG design component diagram

The active database copies are hosted by the two mailbox servers in the primary data center because the primary data center is closest to most of the user population. Hosting the active databases close to the users prevents users from losing access to their email if there is a wide area network (WAN) outage.

The DAG is assigned a witness server (a file server) that is used as another vote to maintain quorum if there is a WAN outage. The DAG's witness server should be placed in the same data center that hosts the active database copies during normal runtime, as shown in Figure 6. For example, the DAG's witness server is located in the primary data center because the active databases are hosted there and most of the user population is located closest to that data center.

If there is a WAN failure, the primary data center has three quorum votes (the two mailbox servers and the witness server), but the secondary data center has only two quorum votes (the two mailbox servers). Therefore, the databases that are hosted in the primary data center remain active and the databases in the secondary data center are taken offline. This prevents the databases in the secondary data center from becoming out of sync with the database copies in the primary data center. Because the DAG's user population is located near the primary data center the users do not lose access to their mailboxes because they do not have to traverse the WAN.

Figure 7 shows the environment during a WAN outage.

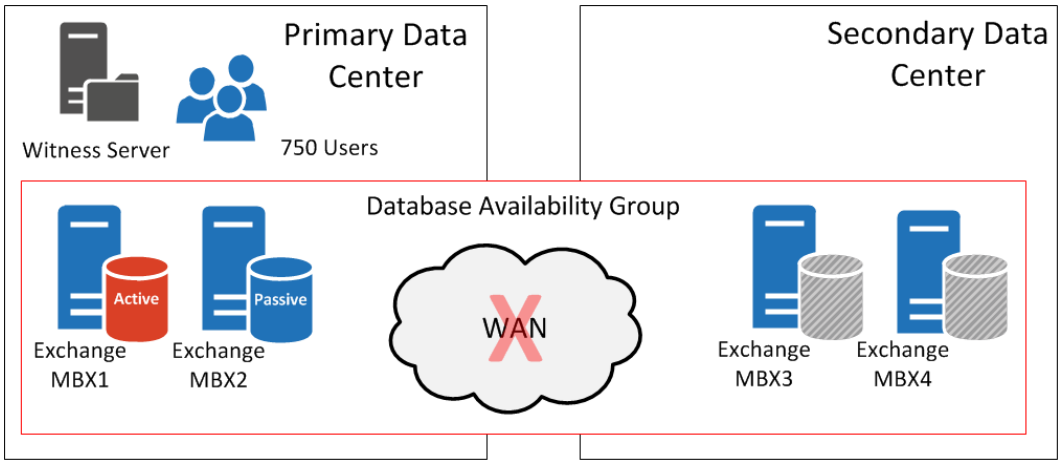

Figure 7. The Exchange environment during a WAN outage

3.1.1 Exchange database distribution

The DAG consists of four mailbox servers (two in each data center). Four mailbox databases are required to meet the needs of the 750 users. The two mailbox servers in the primary data center each host two active mailbox databases. Also, each server hosts the passive copies of the active databases that are hosted by the other mailbox server. (For example, if MBX1 hosts the active copy of a database, MBX2 hosts the passive copy of that database.) At the secondary data center, two mailbox servers each host two passive mailbox database copies and two lagged mailbox database copies of the databases that are active in the primary data center.

Figure 8 shows a detailed view of the database layout for the DAG in the Exchange environment at normal runtime (all mailbox servers are operational).

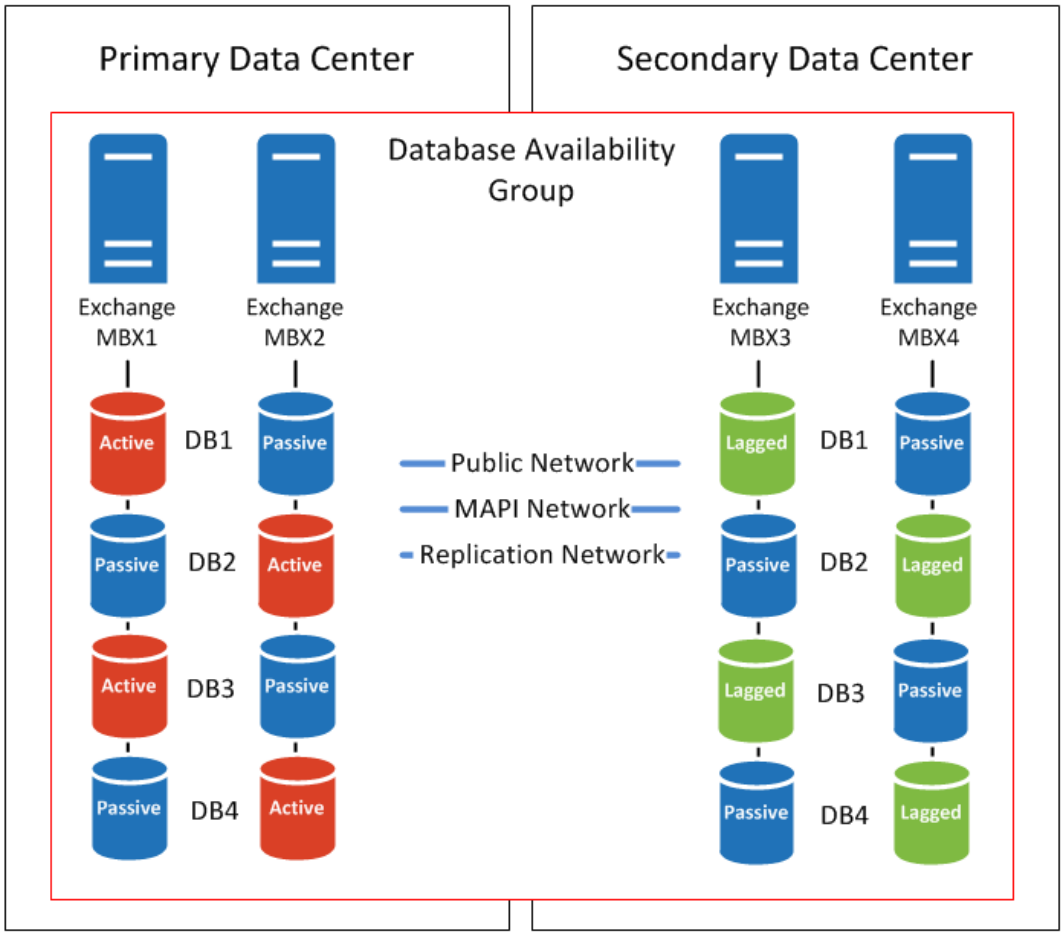

Figure 8. Exchange database distribution when all servers are operational

If there is a server failure or a mailbox server is taken offline for maintenance, the active copies of the databases that are hosted on the affected server go offline and the passive copies that are hosted by the other mailbox server in the same data center become active, as shown in Figure 9.

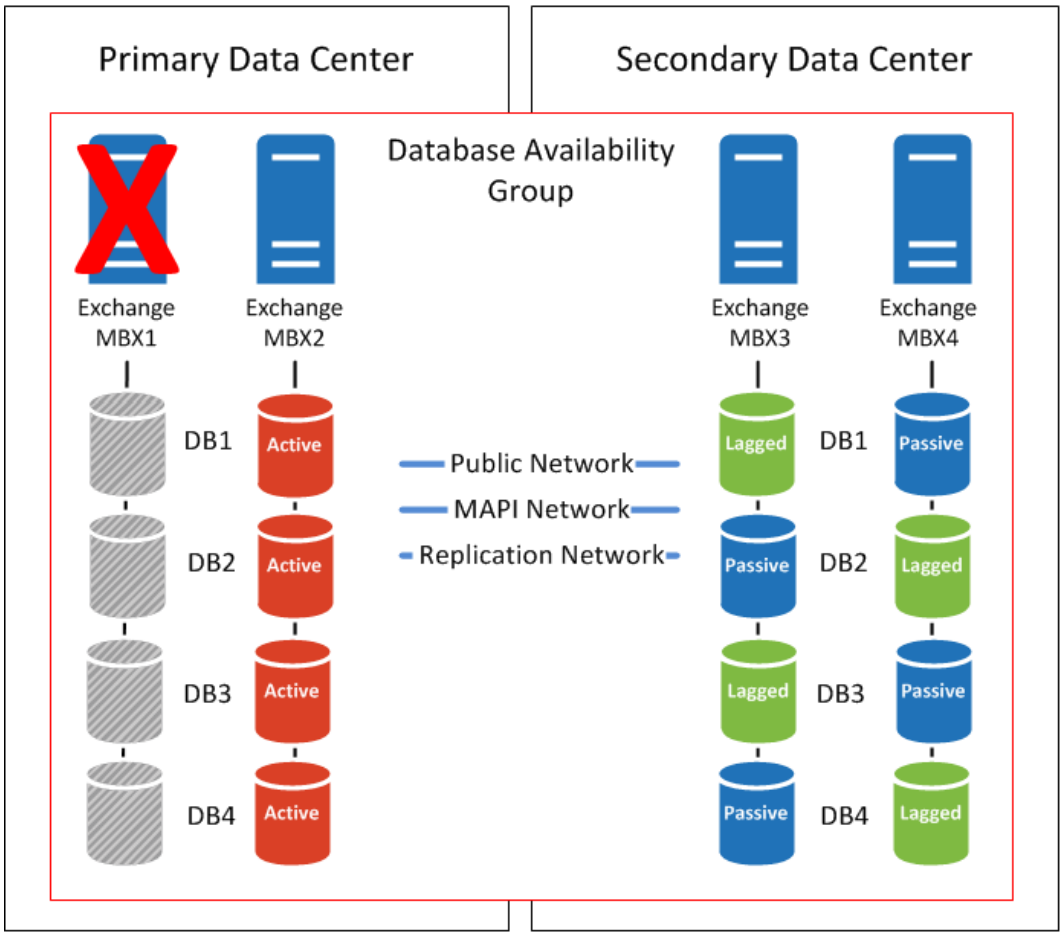

Figure 9. Exchange environment with a single server failure

If there is a second server failure in the primary data center, the passive copies that are hosted by the mailbox servers in the secondary data center become active. At this point, Lenovo recommends playing the logs forward on the lagged database copies to convert them to a highly-available passive copy of the database rather than a lagged copy. Doing so prevents a disruption of mail service if the environment sustains a third server failure, as shown in Figure 10.

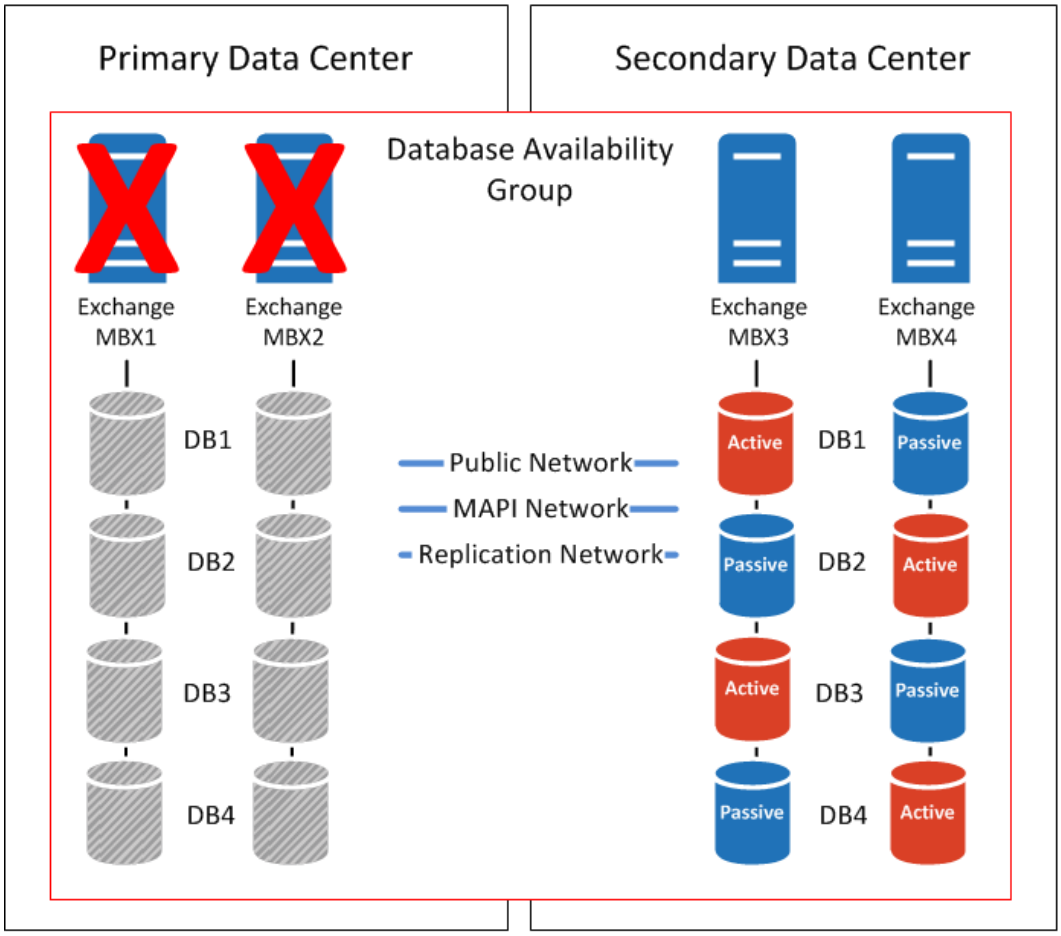

Figure 10. Exchange environment with two server failures

Finally, if there is a third server failure, the passive database copies on the remaining mailbox server become active to support the entire user population of 750, as shown in Figure 11.

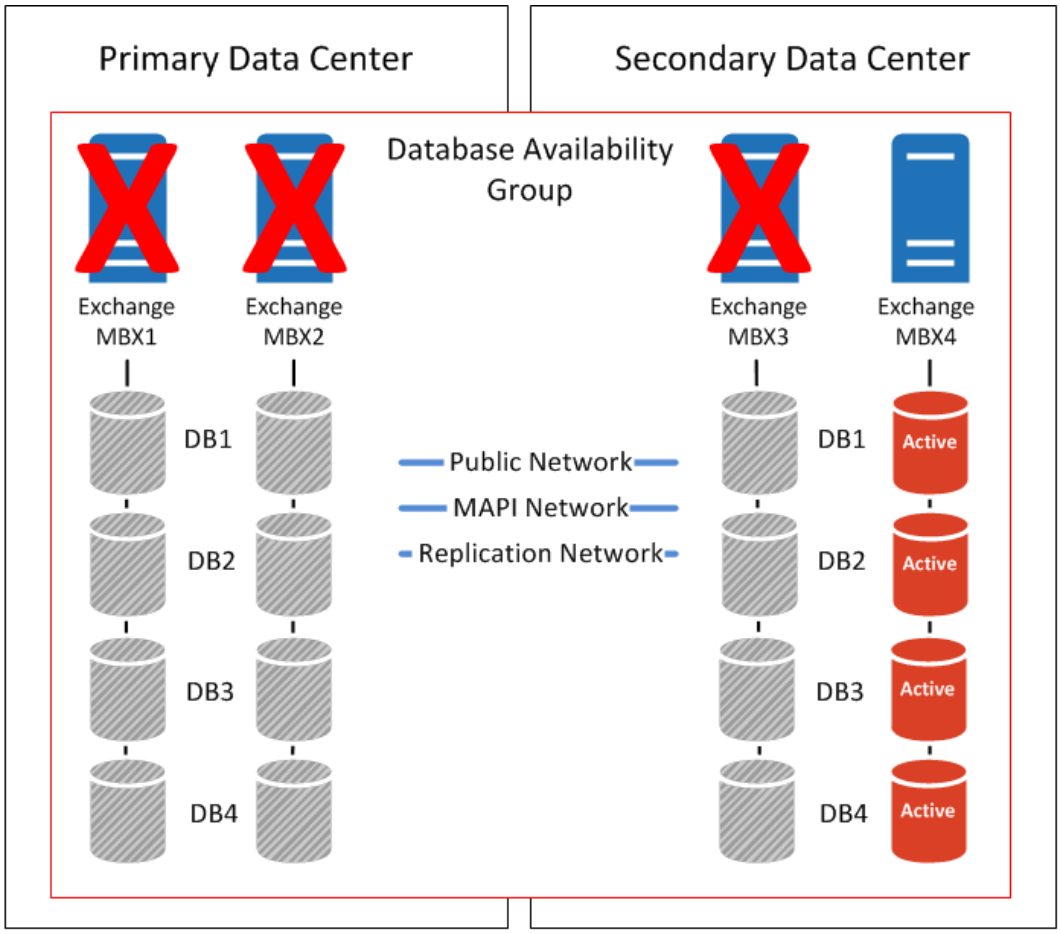

Figure 11. Exchange environment running with a single server

3.1.2 CAS availability

The CAS role requires separate consideration from the Exchange mailbox databases for high-availability. In Exchange Server 2010, CAS server high availability was achieved by using an array of CAS servers that were load balanced by using a network load balancer. In Exchange Server 2013, the CAS array was replaced with the concept of a single namespace for Outlook connectivity.

In a default installation, each CAS server registers its fully qualified domain name (FQDN) as its internal host name in Outlook Anywhere. When an Outlook client makes a connection to a CAS server, it connects to the server's registered internal host name. If the server fails, the connection times out and Outlook Anywhere automatically discovers an available CAS server and creates a connection. However, this process is slow and can leave an Outlook client disconnected for some time. To reduce the time required to create a connection, each CAS server can be configured to use a single namespace as its internal host name in Outlook Anywhere. This configuration requires registering the single namespace as the internal host name for each CAS server and creating a dynamic name system (DNS) record on the DNS server that points to the single namespace. This technique ensures Outlook clients take less time to re-establish connectivity to one of the other IP addresses to which the shared namespace resolves.

Note: DNS round-robin can be used for loud distribution, but a network load balancer is a better option because it provides faster switching capabilities. When a network load balancer is used, the single namespace

resolves to the virtual IP (VIP) that is defined by the network load balancer rather than the IP address of a CAS server. When the network load balancer detects a server failure, it redirects incoming connections to CAS servers that remain online.

3.1.3 Backup strategy

Exchange native data protection features–such as multiple database copies and lagged database copies–and RAID are used in replace of traditional backups. This methodology is used to reduce the overall cost and complexity of the solution.

3.2 Compute server sizing

The recommended minimum of two Lenovo x3550 M5 servers are installed in each of the data centers. The recommended configuration for the compute servers includes the following components:

- Minimum of 48 GB RAM (eight 8GB TruDDR4 RDIMMs recommended)
- Two Intel Xeon E5-2603 v3 (Haswell) 1.6 GHz 6-core processors
- Four 4 TB 7.2K NL-SAS 3.5-inch HDDs (one RAID-1 pair for the OS and one RAID-1 pair for the restore volume)
- One Emulex VFA5 ML2 dual port 10 GbE SFP+ network adapter
- One Emulex 8 Gb FC dual-port host bust adapter (HBA)

3.3 Storage configuration

Figure 12 shows the storage network connections between the servers and the S2200 storage controllers. Each of the two x3550 M5 servers maintains one connection to each of the storage controllers in the S2200 chassis.

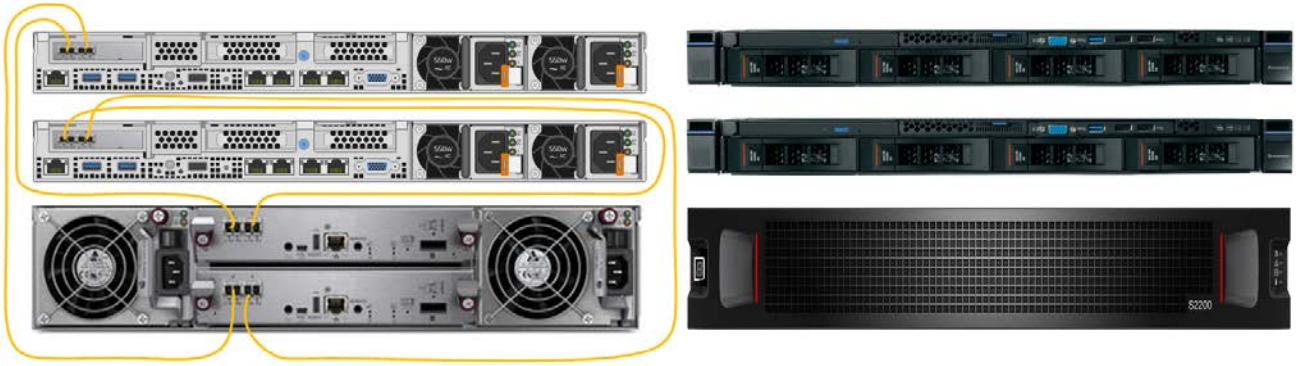

Figure 12. Storage network connections from the servers to the storage

3.3.1 Storage configuration for Exchange databases

Each mailbox server requires two volumes for the mailbox databases and logs. Also, a two-drive RAID-1 volume can be partitioned using local drives on the server for temporary storage for mailbox databases reseeding (Restore volume).

Two separate storage pools are created on the S2200. Each storage pool consists of three, two-drive RAID-1 arrays. Figure 13 shows the storage design for a *single* mailbox server.

Lenovo performed cost-based analysis and determined that 3.5-inch 4 TB NL-SAS drives are the least expensive option for the number of drives and storage subsystems required. Therefore, this design is based on 4 TB 7.2k NL-SAS drives.

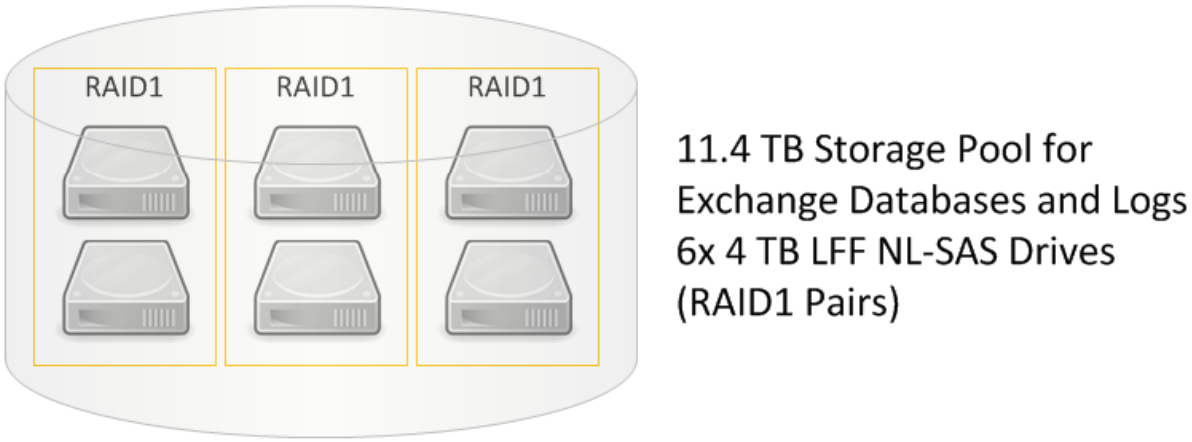

Figure 13. An example of storage design for each VM in the Exchange environment

Each mailbox server requires two 2.55 TB volumes that are allocated for database and log files, as shown in *Figure 14*. Also, a single 2.55 TB volume should be allocated and assigned using the local storage available on each mailbox server for temporary storage space for when a failed database is reseeded.

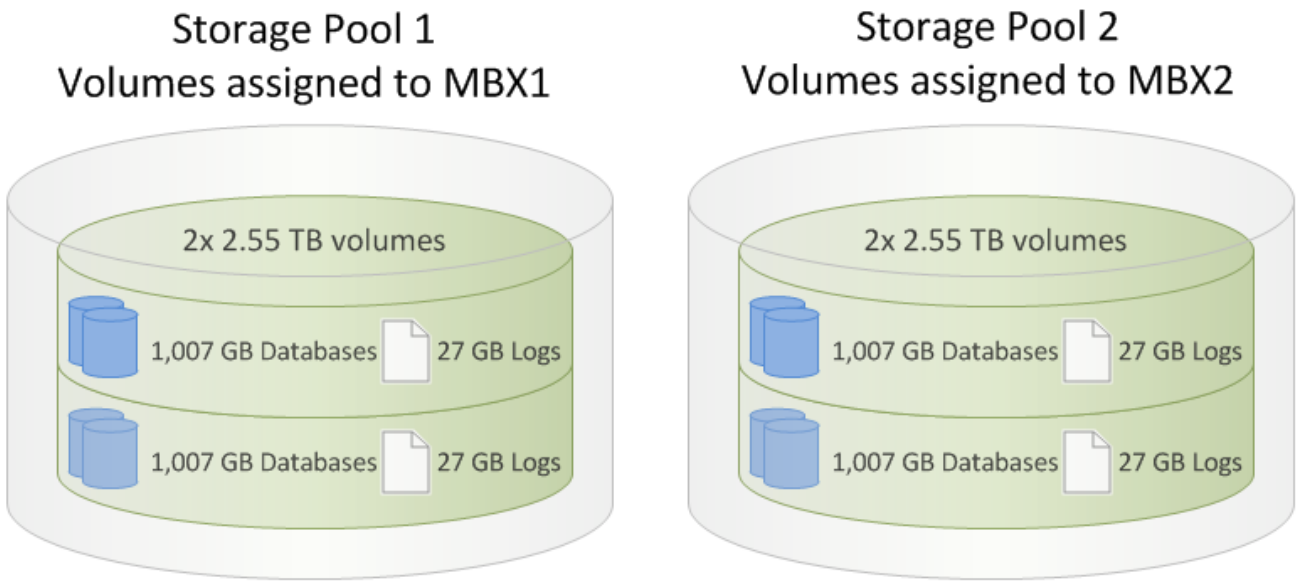

Figure 14. Allocating two 2.5 TB volumes per server for database and log files

There is no single storage configuration that is appropriate for every organization. Lenovo recommends gaining a thorough understanding of the capacity needs of your organization before implementing the storage design and then monitoring the solution for bottlenecks.

Microsoft provides comprehensive guidance on sizing and capacity planning for Exchange. For more information about Exchange sizing and capacity planning, see the TechNet article at: technet.microsoft.com/en-us/library/dn879075(v=exchg.150).aspx

3.4 Targeted customer profile

The Lenovo Storage S2200 system targets entry-level customers, with the ability to scale out as their business grows. The targeted profile for this solution includes:

- Four Exchange Server 2013 servers (two tested, simulating four with database copies)
- Two identical S2200 storage systems with 4 TB HDDs
- 750 mailboxes
- 0.13 IOPS user profile (0.16 tested for 20% growth)
- 5 GB mailbox size
- Background Database Maintenance was enabled during all Jetstress tests
- DAG Mailbox Resiliency with four database copies provides high availability and is the primary data protection mechanism.

3.5 Tested deployment

Tables 1, 2, 3, and 4 summarize the testing environment.

Table 1. Simulated Exchange Configuration

Number of Exchange mailboxes simulated	750
Number of Database Availability Groups (DAGs)	1 Active/Passive DAG
Number of servers/DAG	4 (2 per site)
Number of active mailboxes/server	375 per server at site A during normal runtime
Number of databases/host	4
Number of copies/database	4
Number of mailboxes/database	188
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	0.13 (0.16 tested)
Database/Log LUN size	Two 2.55 TB LUNs per host
Total database size for performance testing	7.36 TB (942 GB x8)
% storage capacity that is used by Exchange database**	72.16% (7.36 TB total database size/10.2 TB formatted capacity)

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

Table 2. Storage Hardware

Storage Connectivity (Fibre Channel, SAS, SATA, iSCSI)	Fiber Channel
Storage model and OS/firmware revision	Lenovo Storage S2200 (Version GL210R005-02) running Management Controller Loader Code 6.16.21066 S2200 WHCL Link
Storage cache	12 GB (6 GB per controller)
Number of storage controllers	One dual-controller unit
Number of storage expansions	0
Maximum bandwidth of storage connectivity to host	2048 MB/sec per host (two 8 Gb FC paths per host)
Switch type/model/firmware revision	No switch was used. Storage was directly connected.
HBA model and firmware	Emulex LPe 12002-M8 8 Gb 2-port PCIe Fibre Channel Adapter (Firmware: 2.02X11)
Number of HBA's/host	One Dual Port HBA
Host server type	Lenovo System x3550 M5, two 12-core E5-2690 v3 2.6 GHz Intel Xeon processors, 128 GB RAM
Total number of disks tested in solution	12 Physical Drives (4 TB 7.2K LFF NL-SAS)
Maximum number of spindles can be hosted in the storage	48 LFF Drives or 96 SFF Drives (Controller + 3 Expansion Units)

Table 3. Storage Software

HBA driver	Emulex 10.6.114.0
HBA QueueTarget Setting	0
HBA QueueDepth Setting	32
Multi-Pathing	Microsoft MPIO
Host OS	Windows Server 2012 R2
ESE.dll file version	15.00.0995.000

Table 4. Storage Disk Configuration (Mailbox Store and Transactional Log Disks)

Disk type, speed and firmware revision	4 TB LFF 7.2k RPM NL-SAS, FW version (BC69)
Raw capacity per disk (GB)	3814 GB
Number of physical disks in test	12
Total raw storage capacity (GB)	45,768 GB (12 x 3,814 GB)
Disk stripe size (KB)	64 KB
Number of disks per Disk Pool	6 (three RAID-1 pairs)
Number of disks per LUN	Two LUNs are created from a single storage pool
Raid level	RAID-1 pairs
Total formatted capacity	10.2 TB (four 2.55 TB LUNs); each host is assigned two LUNs
Storage capacity utilization	22.82% (10.2 TB/44.7 TB)
Database capacity utilization	16.47% (7.36 TB/44.7 TB)

4 Contact for more information

To contact Lenovo, please visit the customer portal at:
shop.lenovo.com/us/en/landingpage/contact/

4.1 About the author

Roland G. Mueller works for Lenovo in Kirkland, Washington. He has a second office at the Microsoft main campus in Redmond, Washington to facilitate close collaboration with Microsoft. He specializes in Exchange Server infrastructure sizing, design and performance testing. Before Lenovo, Roland worked for IBM, specializing in various technologies, including virtualization, bare-metal server deployment, and Microsoft Exchange Server.

5 Best practices

Exchange Server 2013 is a disk-intensive application. Based on the testing run using the ESRP framework, Lenovo recommends the following to improve the storage performance.

- Format the database and log volumes at 64k allocation unit size, as recommended by Microsoft.
- Isolate Exchange database and log files from other disk intensive application workloads to avoid performance conflicts. During our testing, the storage subsystems were dedicated to Exchange. Sharing the storage with other applications can negatively impact Exchange I/O performance.
- Log truncation - Enable circular logging for deployments that use Exchange native data protection features.
- Enable read/write caching on all LUNs as read/write caching improves performance dramatically.
- The recommended disk stripe size for Exchange 2013 is 256 KB (or larger), or follow the vendor's best practices. The default stripe size on the S2200 is 64 KB, which was used for our Jetstress testing.
- Troubleshooting performance on the S2200 storage system is performed using standard tools (such as Performance Monitor, or the monitoring features of the HBA software). These tools help determine where the heaviest disk I/O is occurring, and the data transfer rates. Assuming all storage components are functioning and configured correctly, most Exchange Server 2013 storage latency issues are caused by an insufficient number of drives for the I/O load.
- The number of IOPS hitting a specific drive can be measured using the Performance Monitor Logical Disk object, Disk Transfers/sec counter. Add the data from all database volumes to find the total IOPS being generated.
- The average disk latency should be less than 20ms, with the maximum value no higher than 100ms. Disk latency can be measured using the Performance monitor Logical Disk object, Avg. Disk sec/Read.

For information about Exchange 2013 storage design, please visit:
technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx

For more best practices, tuning, and troubleshooting recommendations for the Lenovo Storage S2200 please view the Product Guide: lenovopress.com/tips1298-lenovo-storage-s2200

6 Test result summary

This section provides a high-level summary of test data from the ESRP and a reference to the detailed HTML reports generated by the ESRP testing framework. Please click on the underlined headings and links below to view the html report for each test.

6.1 Reliability

Several of the tests in the framework check reliability and run for a duration of 24 hours. The test objective is to verify that the storage can handle high I/O workloads for extensive periods. Log and database files are analyzed for integrity after the stress test to ensure there is no database or log corruption.

The following list provides an overview: (clicking on the underlined word below shows the HTML report for each test)

- There were no errors reported in the saved eventlog file.
- There were no errors reported during the database and log checksum process.

6.2 Storage performance results

The Primary Storage performance testing is designed to exercise the storage with a maximum sustainable Exchange I/O pattern for 2 hours. The testing purpose is to reveal how long it takes for the storage to respond to I/O operations under a load. The following data are the sum of all logical disk I/Os and the average of all logical disks I/O latencies during the 2 hour timeframe. Each server is listed separately and the aggregate numbers across all servers are listed as well.

6.2.1 Individual server metrics

The sum of I/O operations and the average of latency across all databases per server.

Table 5. Performance metrics for server 1

Database I/O	
Database Disks Transfers/sec	137.747
Database Disks Reads/sec	93.416
Database Disks Writes/sec	44.331
Average Database Disk Read Latency (ms)	7.6165
Average Database Disk Write Latency (ms)	0.85775
Transaction Log I/O	
Log Disks Writes/sec	10.462
Average Log Disk Write Latency (ms)	0.6215

Table 6. Performance metrics for server 2

Database I/O	
Database Disks Transfers/sec	141.918
Database Disks Reads/sec	96.559
Database Disks Writes/sec	45.359
Average Database Disk Read Latency (ms)	7.378
Average Database Disk Write Latency (ms)	0.791
Transaction Log I/O	
Log Disks Writes/sec	10.734
Average Log Disk Write Latency (ms)	0.575

6.2.2 Aggregate performance across all servers metrics

Table 7 shows the sum of I/O operations across servers in the solution and the average latency across all servers in the solution.

Table 7. Aggregate performance metrics for both servers

Database I/O	
Database Disks Transfers/sec	279.665
Database Disks Reads/sec	189.975
Database Disks Writes/sec	89.69
Average Database Disk Read Latency (ms)	7.49725
Average Database Disk Write Latency (ms)	0.824375
Transaction Log I/O	
Log Disks Writes/sec	21.196
Average Log Disk Write Latency (ms)	0.59825

6.3 Database backup/recovery performance

There are two test reports in this section. The first one measures the sequential read rate of the database files and the second measures the recovery/replay performance (playing transaction logs into the database).

6.3.1 Database read-only performance

This test measures the maximum rate the databases could be backed up via VSS. Table 8 shows the average rate for a single database file.

Table 8. Database backup read-only performance results

MB read/sec per database	104.03
MB read/sec total per server	416.11

6.3.2 Transactional log recovery/replay performance

This test measures the maximum rate at which the log files can be played against the databases. Table 9 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

Average time to play one Log file (sec)	2.03

7 Conclusion

The Lenovo S2200 proved more than capable of handling the high IOPs generated by this ESRP configuration.

This document is developed by Lenovo and reviewed by the Microsoft Exchange Product team. The test results/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 predeployment verifications. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

The ESRP program is not designed to be a benchmarking program; tests are not designed to get the maximum throughput for a giving solution. Rather, the program is 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.

8 Appendix A: Test results

This section includes test reports for Stress, Performance, Streaming Backup, and Soft Recovery. Each server's test results were reviewed by Microsoft, and experienced comparable performance. We include results from only one of the servers here, to make the section easier to follow.

8.1 Stress testing

Overall Test Result	**Pass**
Machine Name	LENOVOX3650M5-1
Test Description	750 Users, 5GB Mailbox, 150 Msg, 4 DBs (2 per volume), .16 IOPS
Test Start Time	10/1/2015 4:27:33 PM
Test End Time	10/4/2015 8:07:18 PM
Collection Start Time	10/1/2015 4:39:37 PM
Collection End Time	10/2/2015 4:39:24 PM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0847.030
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)

Database Sizing and Throughput

Achieved Transactional I/O per Second	131.909
Target Transactional I/O per Second	120
Initial Database Size (bytes)	4045154549760
Final Database Size (bytes)	4057276088320
Database Files (Count)	4

Jetstress System Parameters

Thread Count	2
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	4

Database Configuration

Instance4504.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance4504.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance4504.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance4504.4 Log path: N:\LOG4 Database: N:\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
Instance4504.1	8.084	0.623	22.392	10.600	33485.477	34762.087	0.000	0.371	0.000	2.475	0.000	20856.597
Instance4504.2	7.734	0.618	22.396	10.594	33505.659	34742.099	0.000	0.369	0.000	2.481	0.000	20833.444
Instance4504.3	7.945	1.150	22.410	10.586	33509.050	34753.381	0.000	0.900	0.000	2.475	0.000	20777.459
Instance4504.4	7.999	1.141	22.364	10.568	33484.142	34763.109	0.000	0.912	0.000	2.479	0.000	20844.707

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4504.1	9.159	261435.799
Instance4504.2	9.158	261447.587
Instance4504.3	9.154	261438.622
Instance4504.4	9.154	261444.519

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4504.1	0.653	84703.638
Instance4504.2	0.652	84672.456
Instance4504.3	0.650	84346.509
Instance4504.4	0.654	84816.707

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
Instance4504.1	8.084	0.623	31.550	10.600	99656.514	34762.087	0.278	0.371	0.653	2.475	84703.638	20856.597
Instance4504.2	7.734	0.618	31.555	10.594	99663.429	34742.099	0.278	0.369	0.652	2.481	84672.456	20833.444
Instance4504.3	7.945	1.150	31.564	10.586	99612.651	34753.381	0.394	0.900	0.650	2.475	84346.509	20777.459
Instance4504.4	7.999	1.141	31.518	10.568	99695.576	34763.109	0.398	0.912	0.654	2.479	84816.707	20844.707

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.059	0.000	2.571
Available MBytes	116276.251	116054.000	116355.000
Free System Page Table Entries	16394881.631	16394157.000	16395405.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	192278079.699	191864832.000	193589248.000
Pool Paged Bytes	154448480.663	153972736.000	160505856.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/1/2015 4:27:33 PM -- Preparing for testing ...

10/1/2015 4:27:37 PM -- Attaching databases ...

10/1/2015 4:27:37 PM -- Preparations for testing are complete.

10/1/2015 4:27:37 PM -- Starting transaction dispatch ..

10/1/2015 4:27:38 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

10/1/2015 4:27:38 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

10/1/2015 4:27:42 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).

10/1/2015 4:27:42 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).

10/1/2015 4:27:43 PM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/1/2015 4:27:43 PM -- Performance logging started (interval: 15000 ms).

10/1/2015 4:27:43 PM -- Attaining prerequisites:

10/1/2015 4:39:37 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 966475800.0 (lower bound: 966367600.0,

upper bound: none)

10/2/2015 4:39:37 PM -- Performance logging has ended.

10/4/2015 8:07:14 PM -- JetInterop batch transaction stats: 216216, 216216, 216216 and 216216.

10/4/2015 8:07:14 PM -- Dispatching transactions ends.

10/4/2015 8:07:14 PM -- Shutting down databases ...

10/4/2015 8:07:18 PM -- Instance4504.1 (complete), Instance4504.2 (complete), Instance4504.3 (complete) and Instance4504.4 (complete)

10/4/2015 8:07:18 PM -- C:\Program Files\Exchange Jetstress\Stress_2015_10_1_16_27_42.blg has 5794 samples.

10/4/2015 8:07:18 PM -- Creating test report ...

10/4/2015 8:07:46 PM -- Instance4504.1 has 8.1 for I/O Database Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.1 has 0.4 for I/O Log Writes Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.1 has 0.4 for I/O Log Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.2 has 7.7 for I/O Database Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.2 has 0.4 for I/O Log Writes Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.2 has 0.4 for I/O Log Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.3 has 7.9 for I/O Database Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.3 has 0.9 for I/O Log Writes Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.3 has 0.9 for I/O Log Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.4 has 8.0 for I/O Database Reads Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.4 has 0.9 for I/O Log Writes Average Latency.

10/4/2015 8:07:46 PM -- Instance4504.4 has 0.9 for I/O Log Reads Average Latency.

10/4/2015 8:07:46 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/4/2015 8:07:46 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/4/2015 8:07:46 PM -- C:\Program Files\Exchange Jetstress\Stress_2015_10_1_16_27_42.xml has 5746 samples queried.

8.2 Performance testing

Overall Test Result	**Pass**
Machine Name	LENOVOX3650M5-1
Test Description	750 Users, 5GB Mailbox, 150 Msg, 4 DBs (2 per volume), .16 IOPS
Test Start Time	9/30/2015 1:42:09 PM
Test End Time	10/1/2015 1:13:32 PM
Collection Start Time	9/30/2015 1:53:46 PM
Collection End Time	9/30/2015 3:53:37 PM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0847.030
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2015_9_30_13_42_18.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 137.748

Target Transactional I/O per Second	120
Initial Database Size (bytes)	4041195126784
Final Database Size (bytes)	4045154549760
Database Files (Count)	4

Jetstress System Parameters

Thread Count	2
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	4

Database Configuration

Instance5108.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance5108.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance5108.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance5108.4 Log path: N:\LOG4 Database: N:\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
Instance5108.1	7.708	0.590	23.512	11.252	33433.879	34915.661	0.000	0.355	0.000	2.635	0.000	20651.929
Instance5108.2	7.448	0.588	23.407	11.235	33397.376	34971.897	0.000	0.362	0.000	2.648	0.000	21301.658
Instance5108.3	7.625	1.126	23.300	11.008	33425.529	34955.267	0.000	0.882	0.000	2.612	0.000	20915.220
Instance5108.4	7.685	1.127	23.197	10.836	33468.406	34961.431	0.000	0.887	0.000	2.567	0.000	20883.866

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance5108.1	9.158	261610.744
Instance5108.2	9.153	261791.827
Instance5108.3	9.155	261711.505
Instance5108.4	9.152	261797.282

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance5108.1	0.695	90008.219
Instance5108.2	0.706	91467.812
Instance5108.3	0.688	90085.490
Instance5108.4	0.680	88062.096

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
Instance5108.1	7.708	0.590	32.670	11.252	97395.296	34915.661	0.295	0.355	0.695	2.635	90008.219	20651.929
Instance5108.2	7.448	0.588	32.560	11.235	97598.488	34971.897	0.301	0.362	0.706	2.648	91467.812	21301.658
Instance5108.3	7.625	1.126	32.455	11.008	97821.539	34955.267	0.419	0.882	0.688	2.612	90085.490	20915.220
Instance5108.4	7.685	1.127	32.349	10.836	98066.418	34961.431	0.410	0.887	0.680	2.567	88062.096	20883.866

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.059	0.000	0.160
Available MBytes	116414.635	116405.000	116509.000
Free System Page Table Entries	16394936.198	16394563.000	16395280.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	191739147.223	191680512.000	191893504.000
Pool Paged Bytes	153551190.046	153538560.000	153575424.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 9/30/2015 1:42:09 PM -- Preparing for testing ...

9/30/2015 1:42:14 PM -- Attaching databases ...

9/30/2015 1:42:14 PM -- Preparations for testing are complete.

9/30/2015 1:42:14 PM -- Starting transaction dispatch ..

9/30/2015 1:42:14 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

9/30/2015 1:42:14 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

9/30/2015 1:42:18 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

9/30/2015 1:42:18 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

9/30/2015 1:42:19 PM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

9/30/2015 1:42:19 PM -- Performance logging started (interval: 15000 ms).

9/30/2015 1:42:19 PM -- Attaining prerequisites:

9/30/2015 1:53:46 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 968265700.0 (lower bound: 966367600.0,

upper bound: none)

9/30/2015 3:53:47 PM -- Performance logging has ended.

10/1/2015 1:13:29 PM -- JetInterop batch transaction stats: 70584, 70584, 70583 and 70583.

10/1/2015 1:13:29 PM -- Dispatching transactions ends.

10/1/2015 1:13:29 PM -- Shutting down databases ...

10/1/2015 1:13:32 PM -- Instance5108.1 (complete), Instance5108.2 (complete), Instance5108.3 (complete) and Instance5108.4 (complete)

10/1/2015 1:13:32 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_9_30_13_42_18.blg has 524 samples.

10/1/2015 1:13:32 PM -- Creating test report ...

10/1/2015 1:13:35 PM -- Instance5108.1 has 7.7 for I/O Database Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.1 has 0.4 for I/O Log Writes Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.1 has 0.4 for I/O Log Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.2 has 7.4 for I/O Database Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.2 has 0.4 for I/O Log Writes Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.2 has 0.4 for I/O Log Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.3 has 7.6 for I/O Database Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.3 has 0.9 for I/O Log Writes Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.3 has 0.9 for I/O Log Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.4 has 7.7 for I/O Database Reads Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.4 has 0.9 for I/O Log Writes Average Latency.

10/1/2015 1:13:35 PM -- Instance5108.4 has 0.9 for I/O Log Reads Average Latency.

10/1/2015 1:13:35 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/1/2015 1:13:35 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/1/2015 1:13:35 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_9_30_13_42_18.xml has 478 samples queried.

8.3 Streaming backup

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance6012.1	967320.03	02:21:00	114.33
Instance6012.2	967312.03	02:20:53	114.43
Instance6012.3	967328.03	02:52:08	93.65
Instance6012.4	967328.03	02:52:04	93.70
Avg			104.03
Sum			416.11

Jetstress System Parameters

Thread Count	2
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%

Database Configuration

Instance6012.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance6012.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance6012.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance6012.4 Log path: N:\LOG4 Database: N:\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
Instance6012.1	3.776	0.000	457.809	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance6012.2	3.782	0.000	458.127	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance6012.3	4.705	0.000	374.696	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance6012.4	4.699	0.000	374.861	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	0.472	0.177	0.703
Available MBytes	117386.359	117380.000	117388.000
Free System Page Table Entries	16395208.461	16394820.000	16395516.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	192863467.848	192749568.000	193245184.000
Pool Paged Bytes	157701958.904	157691904.000	157765632.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/4/2015 11:16:38 PM -- Preparing for testing ...

10/4/2015 11:16:43 PM -- Attaching databases ...

10/4/2015 11:16:43 PM -- Preparations for testing are complete.

10/4/2015 11:16:48 PM -- Performance logging started (interval: 30000 ms).

10/4/2015 11:16:48 PM -- Backing up databases ...

10/5/2015 2:08:57 AM -- Performance logging has ended.

10/5/2015 2:08:57 AM -- Instance6012.1 (100% processed), Instance6012.2 (100% processed), Instance6012.3 (100% processed) and Instance6012.4 (100% processed)

10/5/2015 2:08:57 AM -- C:\Program Files\Exchange Jetstress\DatabaseBackup_2015_10_4_23_16_43.blg has 343 samples.

10/5/2015 2:08:57 AM -- Creating test report ...

8.4 Soft recovery

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance4788.1	501	1065.8090385
Instance4788.2	506	1006.1469579
Instance4788.3	501	1008.9280517
Instance4788.4	505	1010.2087464
Avg	503	1022.773
Sum	2013	4091.0927945

Database Configuration

Instance4788.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance4788.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance4788.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance4788.4 Log path: N:\LOG4 Database: N:\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
Instance4788.1	20.329	0.378	276.392	1.845	39423.930	29328.611	16.783	0.000	2.307	0.000	184954.789	0.000
Instance4788.2	18.777	0.405	302.747	2.005	39566.242	31642.634	19.921	0.000	2.506	0.000	201413.530	0.000
Instance4788.3	18.936	1.017	300.077	1.969	39386.251	31909.161	19.691	0.000	2.461	0.000	201694.076	0.000
Instance4788.4	18.982	1.013	301.208	1.991	39525.073	31910.889	19.846	0.000	2.491	0.000	201919.169	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4788.1	0.000	0.000
Instance4788.2	0.000	0.000
Instance4788.3	0.000	0.000
Instance4788.4	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
Instance4788.1	20.329	0.378	276.392	1.845	39423.930	29328.611	16.783	0.000	2.307	0.000	184954.789	0.000
Instance4788.2	18.777	0.405	302.747	2.005	39566.242	31642.634	19.921	0.000	2.506	0.000	201413.530	0.000
Instance4788.3	18.936	1.017	300.077	1.969	39386.251	31909.161	19.691	0.000	2.461	0.000	201694.076	0.000
Instance4788.4	18.982	1.013	301.208	1.991	39525.073	31910.889	19.846	0.000	2.491	0.000	201919.169	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.391	0.000	1.755
Available MBytes	116399.802	116323.000	117359.000
Free System Page Table Entries	16395083.533	16394774.000	16395298.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	192721270.491	192679936.000	192856064.000
Pool Paged Bytes	158618473.813	158597120.000	158793728.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/5/2015 7:56:49 AM -- Preparing for testing ...

10/5/2015 7:56:53 AM -- Attaching databases ...

10/5/2015 7:56:53 AM -- Preparations for testing are complete.

10/5/2015 7:56:53 AM -- Starting transaction dispatch ..

10/5/2015 7:56:53 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

10/5/2015 7:56:53 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

10/5/2015 7:56:57 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/5/2015 7:56:57 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/5/2015 7:56:58 AM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/5/2015 7:56:58 AM -- Performance logging started (interval: 15000 ms).

10/5/2015 7:56:58 AM -- Generating log files ...

10/5/2015 1:21:29 PM -- M:\LOG1 (100.2% generated), M:\LOG2 (101.2% generated), N:\LOG3 (100.2% generated) and N:\LOG4 (101.0% generated)

10/5/2015 1:21:29 PM -- Performance logging has ended.

10/5/2015 1:21:29 PM -- JetInterop batch transaction stats: 16741, 16741, 16741 and 16741.

10/5/2015 1:21:29 PM -- Dispatching transactions ends.

10/5/2015 1:21:29 PM -- Shutting down databases ...

10/5/2015 1:21:33 PM -- Instance4788.1 (complete), Instance4788.2 (complete), Instance4788.3 (complete) and Instance4788.4 (complete)

10/5/2015 1:21:33 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_5_7_56_57.blg has 1295 samples.

10/5/2015 1:21:33 PM -- Creating test report ...

10/5/2015 1:21:38 PM -- Instance4788.1 has 10.1 for I/O Database Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.1 has 0.3 for I/O Log Writes Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.1 has 0.3 for I/O Log Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.2 has 9.7 for I/O Database Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.2 has 0.3 for I/O Log Writes Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.2 has 0.3 for I/O Log Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.3 has 9.9 for I/O Database Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.3 has 0.9 for I/O Log Writes Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.3 has 0.9 for I/O Log Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.4 has 9.9 for I/O Database Reads Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.4 has 0.9 for I/O Log Writes Average Latency.

10/5/2015 1:21:38 PM -- Instance4788.4 has 0.9 for I/O Log Reads Average Latency.

10/5/2015 1:21:38 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/5/2015 1:21:38 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/5/2015 1:21:38 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_5_7_56_57.xml has 1294 samples queried.

10/5/2015 1:21:38 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_5_7_56_57.html was saved.

10/5/2015 1:47:01 PM -- Performance logging started (interval: 2000 ms).

10/5/2015 1:47:01 PM -- Recovering databases ...

10/5/2015 2:04:48 PM -- Performance logging has ended.

10/5/2015 2:04:48 PM -- Instance4788.1 (1065.8090385), Instance4788.2 (1006.1469579), Instance4788.3 (1008.9280517) and Instance4788.4 (1010.2087464)

10/5/2015 2:04:48 PM -- C:\Program Files\Exchange Jetstress\SoftRecovery_2015_10_5_13_47_0.blg has 525 samples.

10/5/2015 2:04:48 PM -- Creating test report ...

9 Appendix B: Lenovo Bill of Materials

This appendix features the Bill of Materials (BOMs) for the ESRP. The BOM lists in this appendix are not meant to be exhaustive and must always be rechecked with the configuration tools. Any information about pricing, support, and maintenance options is outside the scope of this document.

9.1 BOM for compute servers

This section contains BOM for the servers. Two servers per data center are required to meet the design specifications described in this document. The following table lists the BOM for a single server.

Table 4: Lenovo System x3550 M5

Code	Description	Quantity
5463AC1	System x3550 M5	1
A5BF	Intel Xeon Processor E5-2603 v3 6C 1.6GHz 15MB Cache 1600MHz 85W	1
A5BV	Second Intel Xeon Processor E5-2603 v3 6C 1.6GHz 15MB 1600MHz 85W	1
A58Z	System x3550 M5 4x 3.5" Base Chassis	1
A59V	System x3550 M5 Planar	1
A5AH	System x3550 M5 PCIe Riser 1 (1x ML2 x16 CPU0)	1
A5AX	System x 550W High Efficiency Platinum AC Power Supply	2
6204	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
A1ML	Lenovo Integrated Management Module Advanced Upgrade	1
A5AD	System x3550 M5 PCIe Riser 2, 1-2 CPU (FHHL x16 CPU1 + LP x16 CPU0)	1
A5AB	System x Advanced LCD Light Path Kit	1
A597	LCD OP Cable	1
A595	ODD Filler	1
A5A4	System x3550 M5 4x 3.5" HS HDD Kit	1
A3YZ	ServeRAID M5210 SAS/SATA Controller for System x	1
A3Z2	ServeRAID M5200 Series 2GB Flash/RAID 5 Upgrade	1
A3W9	4TB 7.2k 6Gbps NL SATA 3.5” G2HS HDD	4
A5B8	8GB TruDDR4 Memory (2Rx8, 1.2V) PC4-17000 CL15 2133MHz LP RDIMM	8
A40Q	Emulex VFA5 ML2 Dual Port 10GbE SFP+ Adapter	1
3581	Emulex 8Gb FC Dual-port HBA	1
AFAS	ML2 Bracket for Emulex VFA5 ML2 Dual Port 10GbE SFP+ Adapter	1
4048	2U bracket for Emulex 8Gb FC Dual-port HBA for System x	1
A1PJ	3m Passive DAC SFP+ Cable	2
88Y6854	5m LC-LC fiber cable (networking)	2

9.2 BOM for networking

This section features the BOMs for the network switches. Two networking switches per data center are required to meet the design specifications described in this document. Table 10 lists the BOM for a single networking switch.

Table 10: Lenovo RackSwitch G8124E

Code	Description	Quantity
7159BR6	Lenovo System Networking RackSwitch G8124E (Rear to Front)	1
90Y9430	3m IBM Passive DAC SFP+ Cable	2
00D6185	Adjustable 19" 4 Post Rail Kit	1

9.3 BOM for storage

This section features the BOMs for the storage. One S2200 storage system per data center is required to meet the design specifications described in this document. Table 11 lists the BOM for a single S2200.

Table 11: Lenovo Storage S2200

Code	Description	Quantity
6411HC1	Lenovo Storage S2200 LFF Chassis	1
AT1W	Lenovo Storage 3.5" 4TB 7.2k NL-SAS HDD	12
AT1F	Lenovo USB A Male-to-Mini-B 1.5m cable	1
AT25	Lenovo Storage S2200 FC and iSCSI Controller	2
ASR7	Lenovo 3m LC-LC OM3 MMF Cable	4
AT28	Lenovo Storage S2200/S3200 8G Fibre Channel SFP+ Module 1 pack	4
AT2K	Lenovo Storage SAN Rack Mount Kit-Rails 25"-36"	1
6263	4.3m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2

Trademarks and special notices

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, RackSwitch, and System x are trademarks of Lenovo.

Microsoft, Windows, Exchange Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos),Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information may address anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.