Lenovo Storage S3200 and System x3550 M5 Server 750 Mailbox Resiliency Solution for Exchange Server 2013

Last update: **27 January 2016**

Describes a reference architecture for Exchange 2013 on Lenovo Storage S3200

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

Table of Contents

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 S3200. 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 my Microsoft Corporation to provide 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 running in two data centers, with each data center having two x3550 M5 servers and a dual-controller S3200 storage system.

This document provides the design considerations and best practices for implementing the described architecture to support small- to medium-sized organizations with less than 750 employees. However, the principles and techniques 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). 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 provides ports for your data and storage connections.

The 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 ten front and two rear 2.5-inch SFF HDDs or SSDs, or
- Up to four 3.5-inch HDDs

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

The Lenovo x3550 M5 server is shown in Figure 1.

Figure 1. Lenovo System x3550 M5

For more information, see the Lenovo Press product guide: http://lenovopress.com/tips1194-system-x3550-m5

2.2 Lenovo Storage S3200

The Lenovo Storage S3200 is a versatile storage system designed to provide simplicity, speed, scalability, security, and high availability for small to large businesses. The S3200 SAN array delivers enterprise-class storage management technology in a cost-effective solution with a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features.

The S3200 is a perfect fit for a wide range of workloads, from specialized workloads such as big data and analytics, video surveillance, media streaming, and private clouds, to general-purpose workloads such as file and print serving, web serving, e-mail and collaboration, and OLTP databases. The S3200 is also well-suited to secure archive storage or a consolidated backup solution.

The Lenovo Storage S3200 is shown in Figure 2.

Figure 2. Lenovo Storage S3200

Key features and benefits provided by the S3200 storage system include:

- Versatile, scalable storage with a single controller, or dual active/active controller configurations for high availability and performance.
- Flexible host connectivity to match diverse client needs with support for 6/12 Gb SAS (SAS controller) or 1/10 Gb iSCSI or 4/8/16 Gb FC connectivity (FC and iSCSI controller) or both iSCSI and FC at the same time.
- 6 Gb SAS drive-side connectivity with support for 12 x 3.5-inch large form factor (LFF) or 24 x 2.5-inch small form factor (SFF) drives in the controller enclosure. The configuration is scalable up to 96 LFF drives per system with the attachment of Lenovo Storage E1012 expansion enclosures (12 LFF drives each) or up to 192 SFF drives per system with the attachment of Lenovo Storage E1024 expansion enclosures (24 SFF drives each), to satisfy growing needs for storage capacity and performance.
- Flexibility in storing data on high-performance SAS SSDs, performance-optimized enterprise SAS HDDs, or capacity-optimized enterprise NL SAS HDDs; mixing and matching drive types and form factors within a single system to perfectly meet performance and capacity requirements for various workloads.
- Support for self-encrypting drives to allow clients to secure their sensitive data and to comply with various security regulations when required.
- Rich set of standard functions available at no additional cost, including virtualized storage pools, snapshots, thin provisioning, rapid rebuild, real-time HDD tiering, and SSD read cache.
- Optional licensed functions, including higher number of snapshots for additional scalability and real-time SSD tiering for boosting IOPS performance.
- Intuitive, web-based GUI for easy system setup and management.
- Designed to comply with NEBS and MIL-STD requirements.
- Designed for 99.999% availability.

For more information, see the Lenovo Press product guide:
http://lenovopress.com/tips1299-lenovo-storage-s3200

2.3 Emulex 16 Gb FC HBA for Lenovo System x

Streamlined installation and management, plus unrivalled scalability and industry-leading virtualization support make the Emulex 16 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 16 Gb FC dual-port HBA is shown in Figure 3.

Figure 3. Emulex 16 Gb FC dual-port 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 4 Gb, 8 Gb, or 16 Gb link attachments
- Processes over 500,000 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, see the Lenovo Press product guide: http://lenovopress.com/tips0848

2.4 Lenovo RackSwitch G8124E

The Lenovo RackSwitch G8124E is a 10 Gigabit Ethernet switch that is 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 Lenovo 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 - 8 vNICs for improved performance, availability, and security, while reducing cost and complexity.

The G8124E is cooler and implements a choice of directional cooling (front-to-back or back-to-front) 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 easier, 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 RackSwitch 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, see the Lenovo Press product guide: 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 S3200 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.

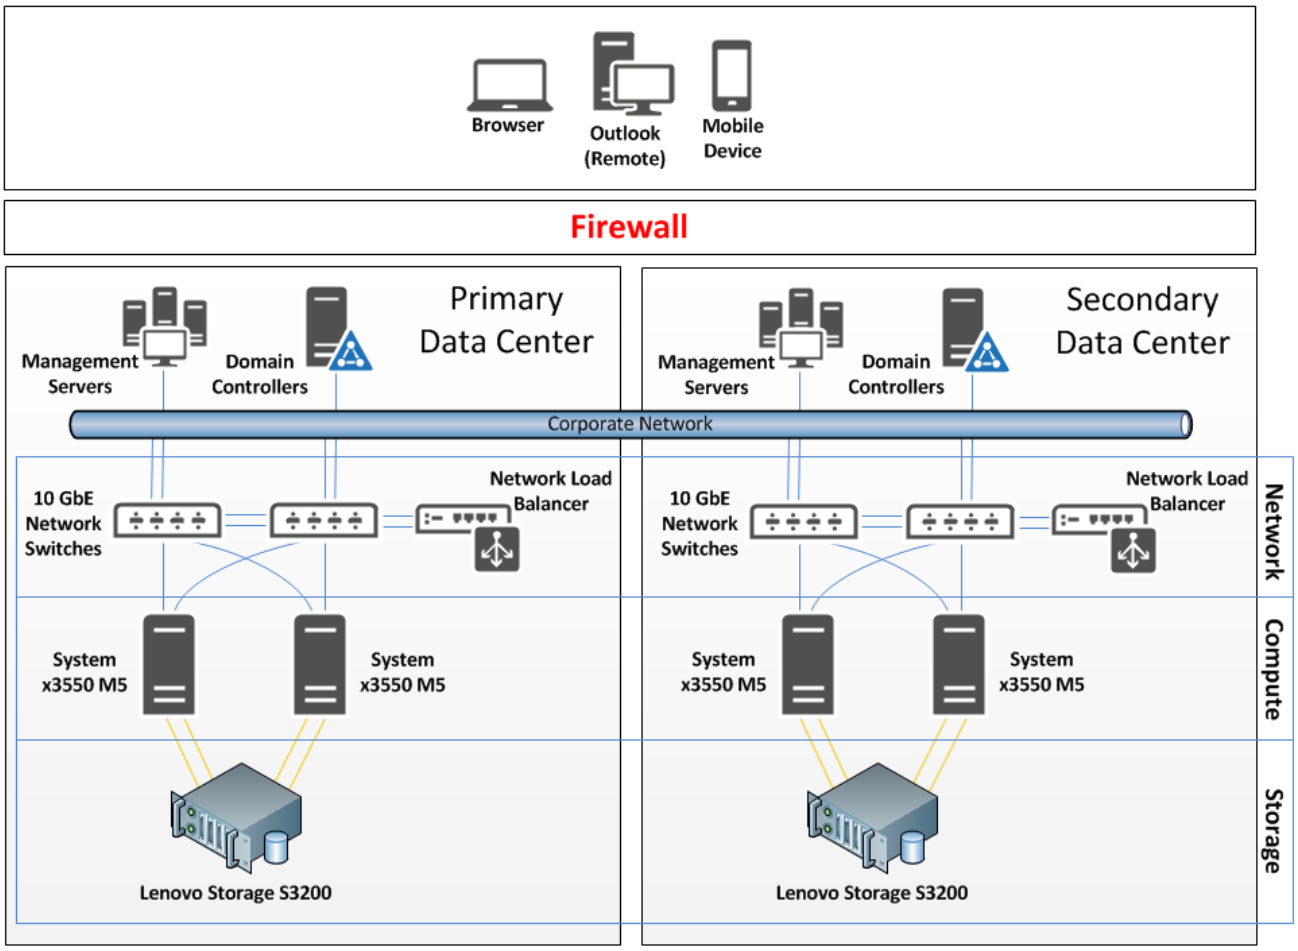

Figure 5. Architectural overview

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 disk failure or server failure).

Figure 6 shows an example of the design for a DAG. An active/passive, multi-site 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 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 in the past three days.

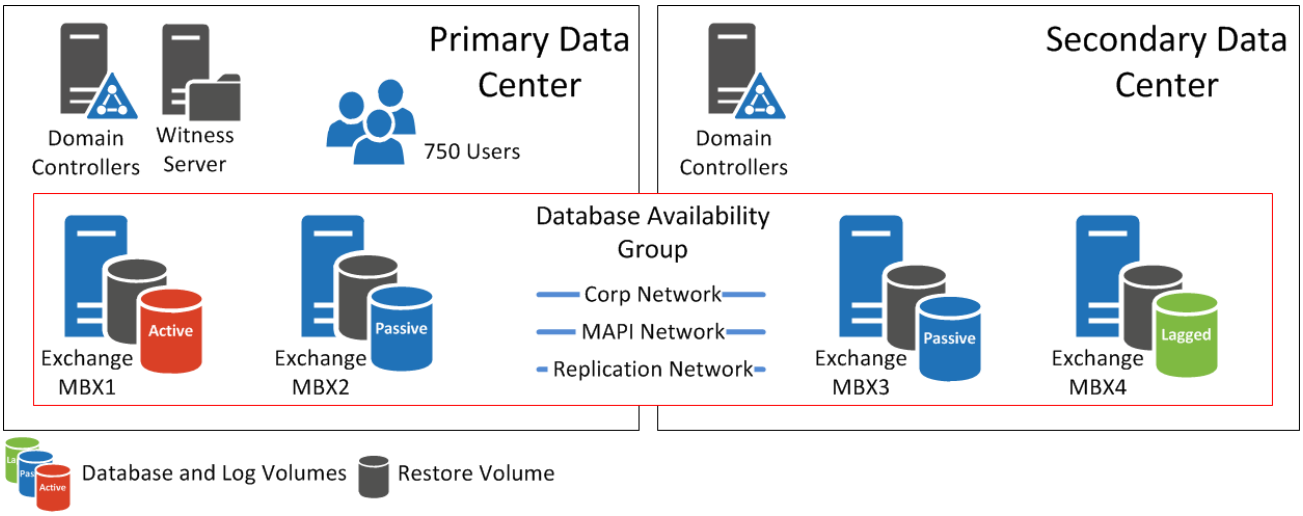

Figure 6. DAG design component diagram

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.

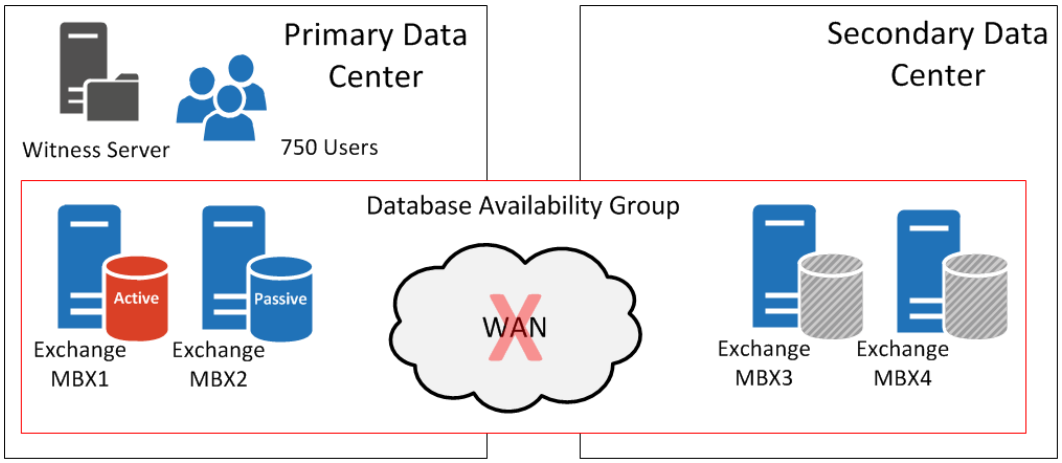

Figure 7. The Exchange environment during a WAN outage

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).

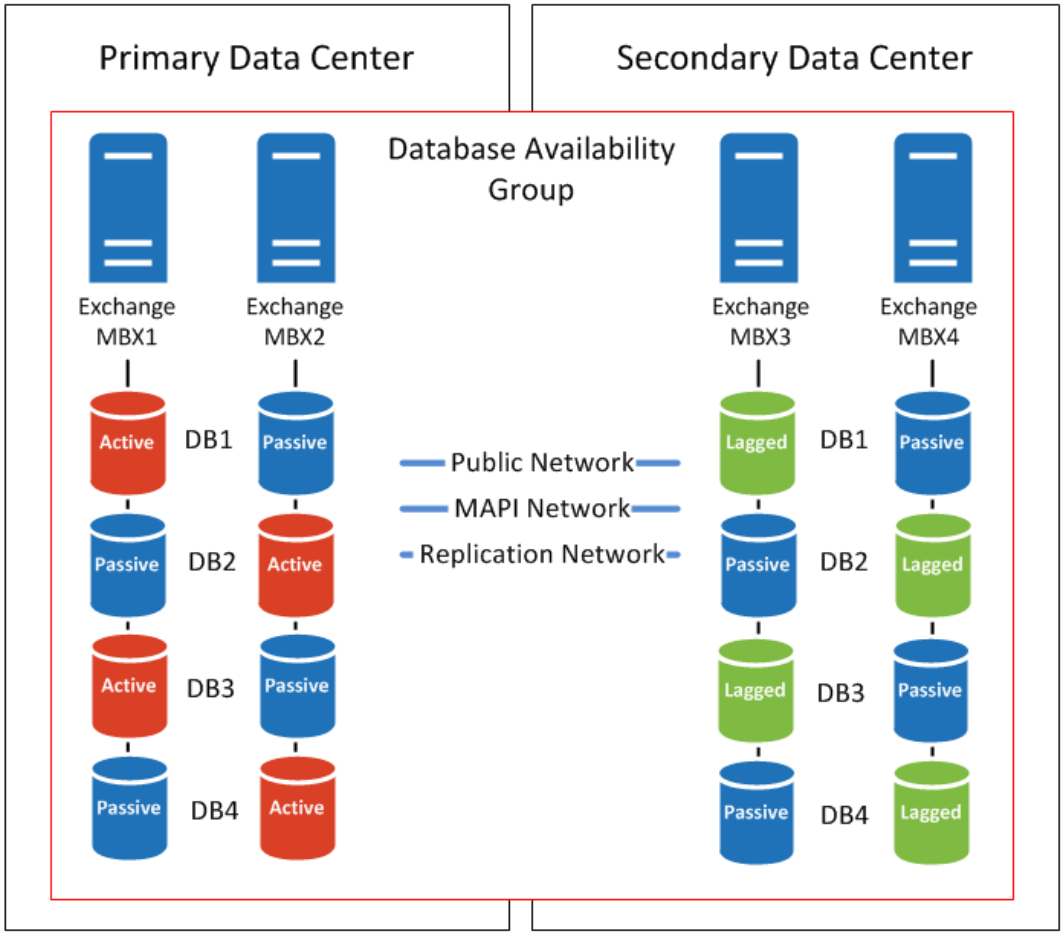

Figure 8. Exchange database distribution when all servers are operational

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.

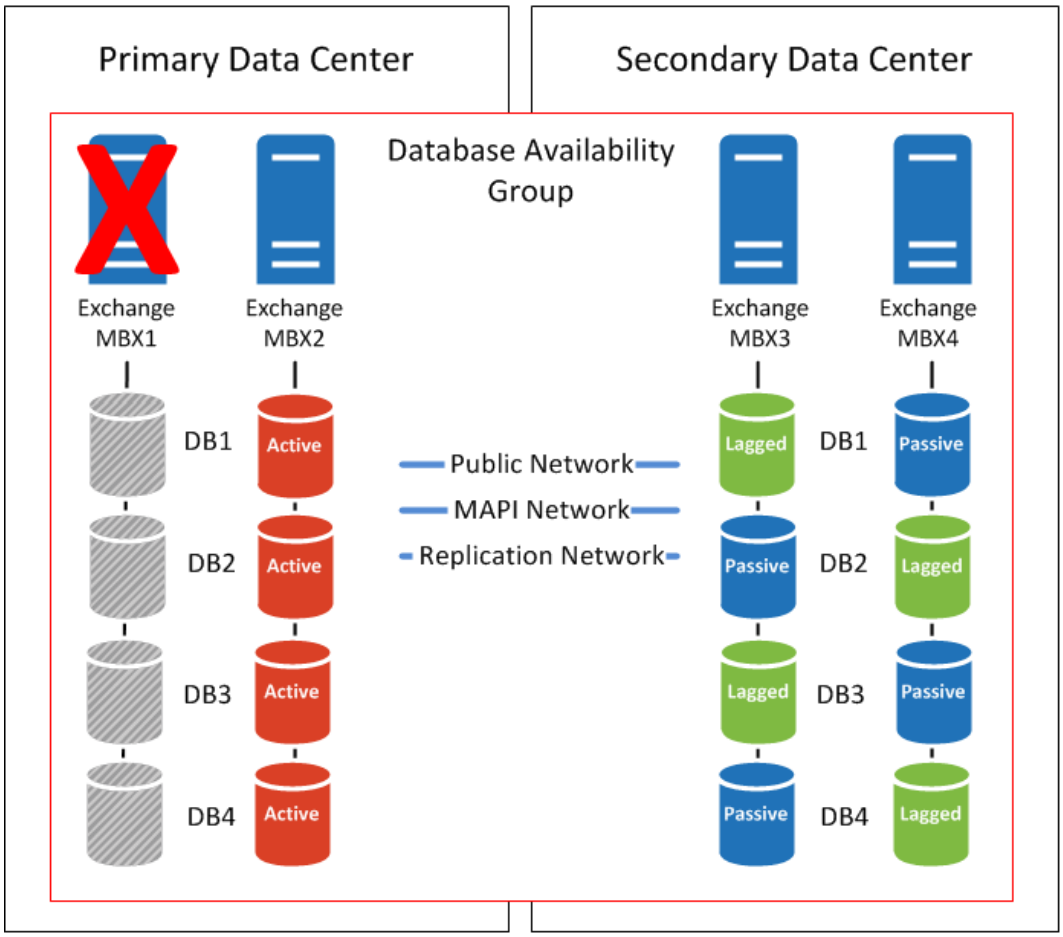

Figure 9. Exchange environment with a single server failure

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.

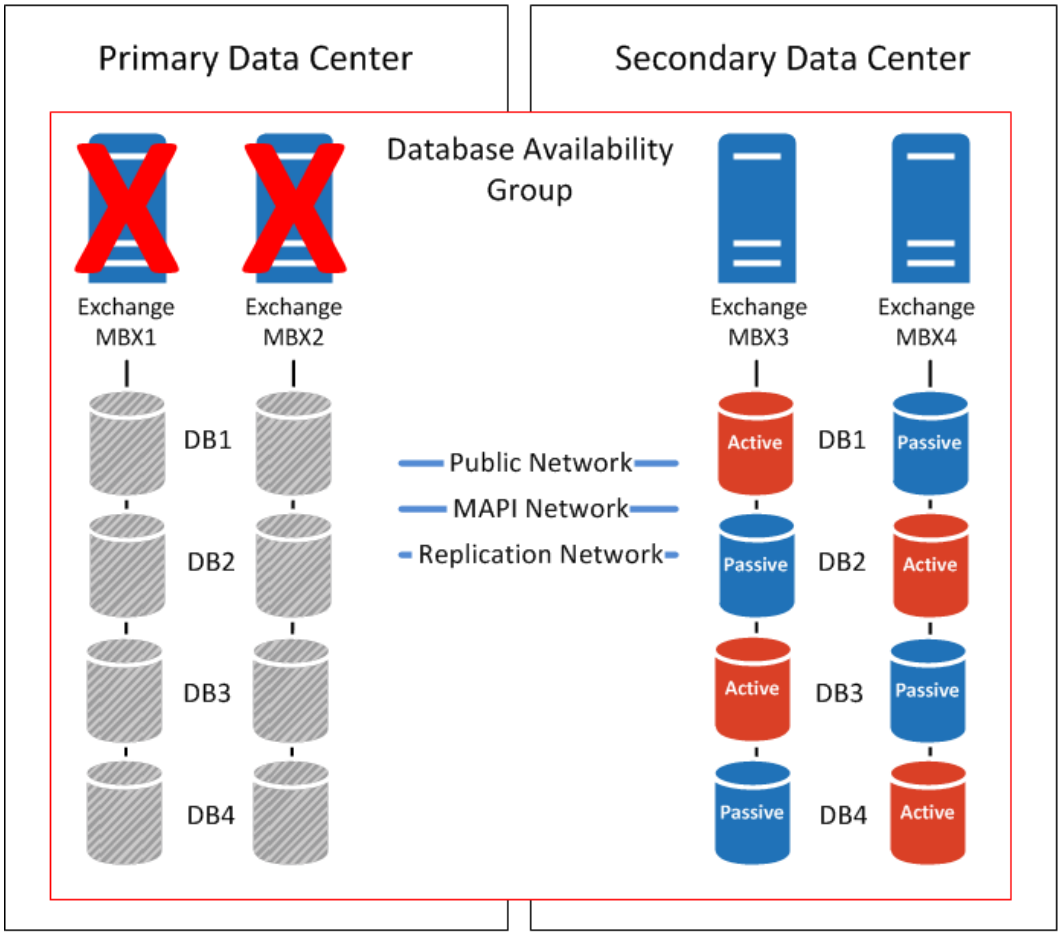

Figure 10. Exchange environment with two server failures

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.

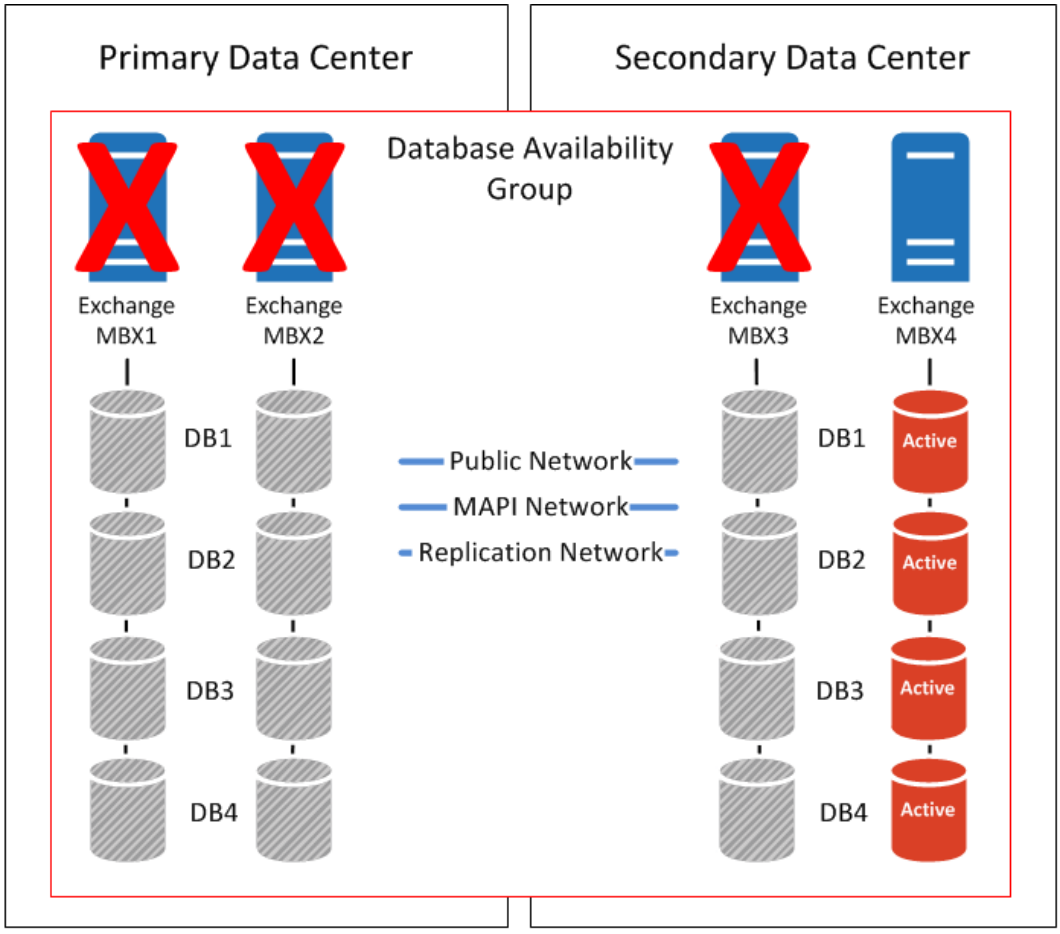

Figure 11. Exchange environment that is running with a single server

3.1.2 CAS availability

The CAS role requires consideration separate 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 cloud 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 are 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

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 (8 x 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” 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)

A minimum of two compute servers per data center are recommended.

3.3 Storage configuration

Figure 12 shows the storage network connections between the servers and the S3200 storage controllers. Each of the two x3550 M5 servers maintains one connection to each of the storage controllers in the S3200 chassis.

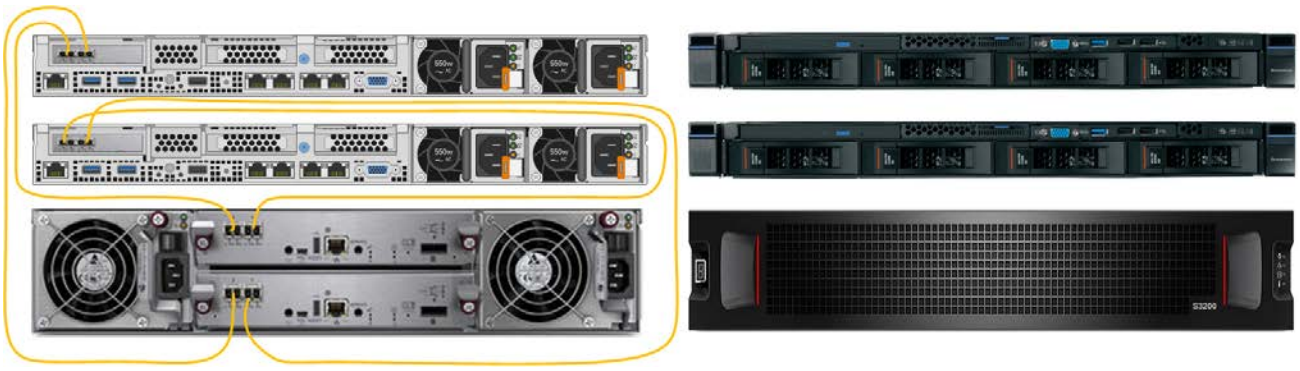

Figure 12. Storage network connections from the servers to the storage

3.3.1 Storage configuration for Exchange databases

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

Two separate storage pools are created on the S3200. Each storage pool is comprised of three two-disk RAID-1 arrays. Figure 13 shows the storage design for a *single* mailbox server.

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

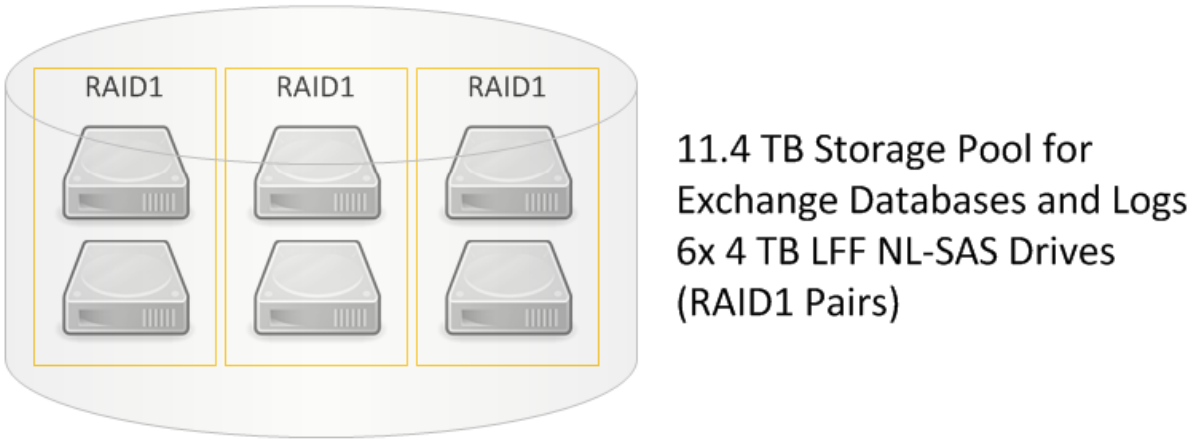

Figure 13. An example storage design for each VM in the Exchange environment

Each mailbox server requires two 2.55 TB volumes that are allocated for database and log files. 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.

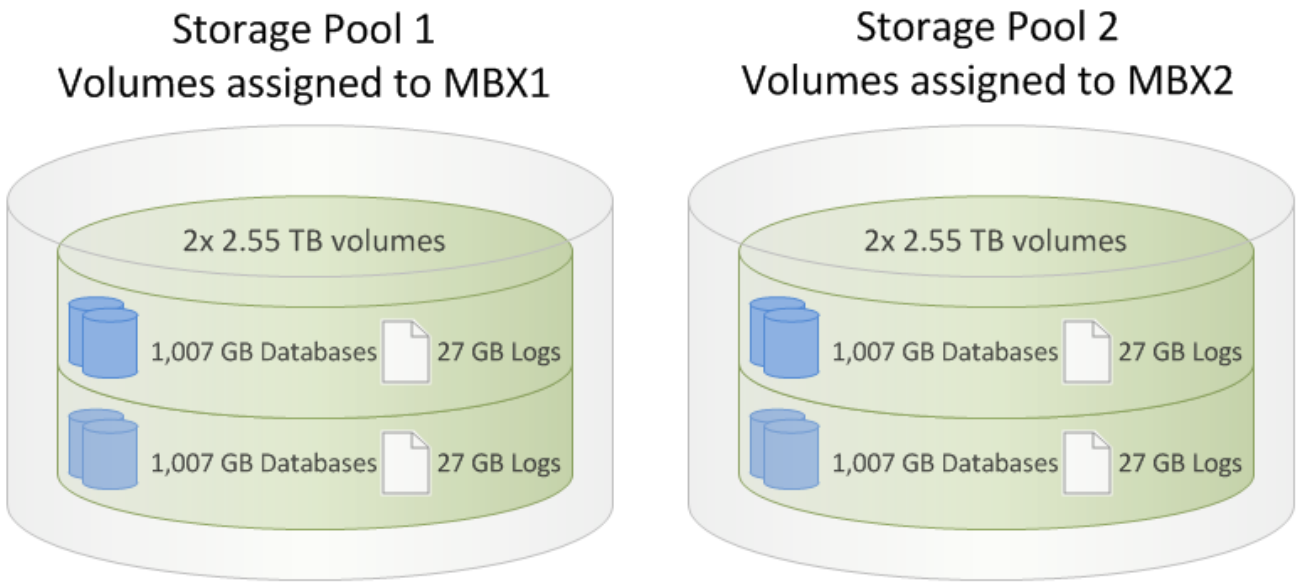

There is no single storage configuration that is appropriate for every organization. Lenovo recommends gaining a thorough understanding of the capacity needs of your organization before implementing the storage design and then monitoring the solution for bottlenecks.

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 S3200 system is designed for 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 S3200 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

The following tables 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/Os per second per mailbox (IOPS, includes 20% headroom)	0.13 (0.16 tested)
Database/Log LUN size	Two 2.55 TB LUNs per host
Total database size for performance testing	7.5 TB (960 GB x8)
% storage capacity that is used by Exchange database**	73.53% (7.5 TB total database size/10.2 TB formatted capacity)

**Storage performance characteristics change based on the percentage utilization of the individual disks. 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)	Fibre Channel
Storage model and OS/firmware revision	Lenovo Storage S3200 (Version GL210R005-02) running Management Controller Loader Code 6.16.21066 S3200 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	4096 MB/sec per host (two 16 Gb FC paths per host)
Switch type/model/firmware revision	No switch was used. Storage was directly connected.
HBA model and firmware	Emulex LPe 16002B-M6 16 Gb 2-port PCIe Fibre Channel Adapter (Firmware: 10.4.303.3)
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 drives tested in solution	12 physical HDDs (4 TB 7.2K LFF NL-SAS)
Maximum number of drives that can be hosted in the storage	96 LFF HDDs or 192 SFF HDDs (controller + 7 expansion units)

Table 3. Storage Software

HBA driver	Emulex 10.6.114.0
HBA QueueTarget Setting	0
HBA QueueDepth Setting	32
Multipathing	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)

Drive type, speed and firmware revision	4 TB LFF 7.2k RPM NL-SAS HDD, firmware version BC69
Raw capacity per drive (GB)	3,814 GB
Number of physical drives in test	12
Total raw storage capacity (GB)	45,768 GB (12 x 3,814 GB)
Disk stripe size (KB)	64 KB
Number of drives per Disk Pool	6 (three RAID-1 pairs)
Number of drives per LUN	Two LUNs are created from a single storage pool.
Raid level	RAID-1 pairs
Total formatted capacity	10.2 TB (using four 2.55 TB LUNs). Each host is assigned 2 LUNs.
Storage capacity utilization	22.82% (10.2 TB / 44.7 TB)
Database capacity utilization	16.78% (7.5 TB / 44.7 TB)

4 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, because read/write caching improves performance dramatically.
- The recommended disk stripe size for Exchange 2013 is 256 KB (or larger) or to follow the vendor's best practices. The default stripe size on the S3200 is 64 KB, which was used for our Jetstress testing.
- Troubleshooting performance on the S3200 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 information on the Lenovo Storage S3200, see the Lenovo Press product guide:
http://lenovopress.com/tips1299-lenovo-storage-s3200

5 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. Click on the underlined headings and links, below, to view the HTML report for each test.

5.1 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.

5.2 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 disk I/O latencies during the 2 hour timeframe. Each server is listed separately and the aggregate numbers across all servers are listed as well.

5.2.1 Individual server metrics

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

Table 5. Performance metrics for server 1

Database I/O	
Database Disks Transfers/sec	136.444
Database Disks Reads/sec	92.918
Database Disks Writes/sec	43.526
Average Database Disk Read Latency (ms)	7.519
Average Database Disk Write Latency (ms)	0.624
Transaction Log I/O	
Log Disks Writes/sec	11.679
Average Log Disk Write Latency (ms)	0.468

Table 6. Performance metrics for server 2

Database I/O	
Database Disks Transfers/sec	137.448
Database Disks Reads/sec	93.706
Database Disks Writes/sec	43.742
Average Database Disk Read Latency (ms)	7.513
Average Database Disk Write Latency (ms)	0.63
Transaction Log I/O	
Log Disks Writes/sec	11.787
Average Log Disk Write Latency (ms)	0.474

5.2.2 Aggregate performance across all servers metrics

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	273.892
Database Disks Reads/sec	186.624
Database Disks Writes/sec	87.268
Average Database Disk Read Latency (ms)	7.516
Average Database Disk Write Latency (ms)	0.627
Transaction Log I/O	
Log Disks Writes/sec	23.466
Average Log Disk Write Latency (ms)	0.471

5.3 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).

5.3.1 Database read-only performance

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

Table 8. Database backup read-only performance results

MB read/sec per database	174.07

MB read/sec total per server	696.26

5.3.2 Transactional log recovery/replay performance

This test measures 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

Average time to play one Log file (sec)	1.80

6 Conclusion

The Lenovo S3200 proved more than capable of handling the high IOPs generated by this ESRP configuration.

This document is developed by Lenovo and reviewed by Microsoft Exchange Product team. The test results/data presented in this document are based on the tests introduced in the ESRP test framework. A customer should not use the data directly for his/her predeployment verification. 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 determine the maximum throughput for a given solution. Rather, it is focused on producing recommendations from vendors for the Exchange application. So the data presented in this document should not be used for direct comparisons among the solutions.

7 Contact for more information

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

7.1 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.

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. Microsoft 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	10/20/2015 10:19:04 AM
Test End Time	10/21/2015 10:47:19 AM
Collection Start Time	10/20/2015 10:30:12 AM
Collection End Time	10/21/2015 10:30:06 AM
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\Stress_2015_10_20_10_19_12.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	140.729
Target Transactional I/O per Second	120
Initial Database Size (bytes)	4028083732480
Final Database Size (bytes)	4032487751680
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

Instance4496.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

nstance4496.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance4496.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance4496.4 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
Instance4496.1	7.660	0.870	23.974	11.227	33382.767	36325.949	0.000	0.700	0.000	2.768	0.000	20843.904
Instance4496.2	7.439	0.867	23.981	11.248	33434.998	36328.319	0.000	0.698	0.000	2.787	0.000	20699.887
Instance4496.3	7.086	0.406	23.916	11.204	33437.488	36392.230	0.000	0.235	0.000	2.774	0.000	20951.618
Instance4496.4	7.193	0.408	23.963	11.215	33433.515	36372.684	0.000	0.235	0.000	2.771	0.000	20866.185

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4496.1	9.157	261653.056
Instance4496.2	9.156	261666.767
Instance4496.3	9.156	261682.055
Instance4496.4	9.157	261665.761

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4496.1	0.732	94754.853
Instance4496.2	0.732	94965.213
Instance4496.3	0.736	95244.262
Instance4496.4	0.733	94792.233

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
Instance4496.1	7.660	0.870	33.131	11.227	96472.342	36325.949	0.297	0.700	0.732	2.768	94754.853	20843.904
Instance4496.2	7.439	0.867	33.137	11.248	96498.626	36328.319	0.299	0.698	0.732	2.787	94965.213	20699.887
Instance4496.3	7.086	0.406	33.072	11.204	96626.800	36392.230	0.179	0.235	0.736	2.774	95244.262	20951.618
Instance4496.4	7.193	0.408	33.119	11.215	96533.705	36372.684	0.178	0.235	0.733	2.771	94792.233	20866.185

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.086	0.000	2.060
Available MBytes	116299.022	116103.000	116376.000
Free System Page Table Entries	16398023.401	16397481.000	16398317.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	210058819.804	209940480.000	214355968.000
Pool Paged Bytes	201728088.270	201347072.000	270303232.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/20/2015 10:19:04 AM -- Preparing for testing ...

10/20/2015 10:19:08 AM -- Attaching databases ...

10/20/2015 10:19:08 AM -- Preparations for testing are complete.

10/20/2015 10:19:08 AM -- Starting transaction dispatch ..

10/20/2015 10:19:08 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

10/20/2015 10:19:08 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

10/20/2015 10:19:12 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).

10/20/2015 10:19:12 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).

10/20/2015 10:19:13 AM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/20/2015 10:19:13 AM -- Performance logging started (interval: 15000 ms).

10/20/2015 10:19:13 AM -- Attaining prerequisites:

10/20/2015 10:30:12 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 967532500.0 (lower bound: 966367600.0, upper bound: none)

10/21/2015 10:30:13 AM -- Performance logging has ended.

10/21/2015 10:47:16 AM -- JetInterop batch transaction stats: 79555, 79555, 79555 and 79554.

10/21/2015 10:47:16 AM -- Dispatching transactions ends.

10/21/2015 10:47:17 AM -- Shutting down databases ...

10/21/2015 10:47:19 AM -- Instance4496.1 (complete), Instance4496.2 (complete), Instance4496.3 (complete) and Instance4496.4 (complete)

10/21/2015 10:47:19 AM -- C:\Program Files\Exchange Jetstress\Stress_2015_10_20_10_19_12.blg has 5797 samples.

10/21/2015 10:47:19 AM -- Creating test report ...

10/21/2015 10:47:37 AM -- Instance4496.1 has 7.7 for I/O Database Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.1 has 0.7 for I/O Log Writes Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.1 has 0.7 for I/O Log Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.2 has 7.4 for I/O Database Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.2 has 0.7 for I/O Log Writes Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.2 has 0.7 for I/O Log Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.3 has 7.1 for I/O Database Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.3 has 0.2 for I/O Log Writes Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.3 has 0.2 for I/O Log Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.4 has 7.2 for I/O Database Reads Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.4 has 0.2 for I/O Log Writes Average Latency.

10/21/2015 10:47:37 AM -- Instance4496.4 has 0.2 for I/O Log Reads Average Latency.

10/21/2015 10:47:37 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/21/2015 10:47:37 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/21/2015 10:47:37 AM -- C:\Program Files\Exchange Jetstress\Stress_2015_10_20_10_19_12.xml has 5753 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	10/19/2015 5:07:00 PM
Test End Time	10/19/2015 8:17:11 PM
Collection Start Time	10/19/2015 5:18:52 PM
Collection End Time	10/19/2015 7:18:49 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_10_19_17_7_9.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	136.444
Target Transactional I/O per Second	120
Initial Database Size (bytes)	4026565394432
Final Database Size (bytes)	4027152596992
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

Instance4188.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance4188.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance4188.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance4188.4 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
Instance4188.1	7.745	0.845	23.302	11.086	33347.333	35650.662	0.000	0.696	0.000	2.947	0.000	20780.811
Instance4188.2	7.619	0.846	23.100	10.634	33399.486	35733.236	0.000	0.697	0.000	2.852	0.000	21014.660
Instance4188.3	7.303	0.402	23.103	10.714	33406.882	35591.208	0.000	0.240	0.000	2.896	0.000	21004.494
Instance4188.4	7.408	0.401	23.413	11.092	33451.969	35576.768	0.000	0.239	0.000	2.984	0.000	20578.807

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4188.1	9.156	261681.634
Instance4188.2	9.161	261528.391
Instance4188.3	9.159	261574.734
Instance4188.4	9.161	261534.637

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4188.1	0.777	100501.645
Instance4188.2	0.762	98559.584
Instance4188.3	0.769	99530.615
Instance4188.4	0.781	101453.036

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
Instance4188.1	7.745	0.845	32.458	11.086	97756.929	35650.662	0.316	0.696	0.777	2.947	100501.645	20780.811
Instance4188.2	7.619	0.846	32.261	10.634	98179.369	35733.236	0.311	0.697	0.762	2.852	98559.584	21014.660
Instance4188.3	7.303	0.402	32.262	10.714	98185.480	35591.208	0.187	0.240	0.769	2.896	99530.615	21004.494
Instance4188.4	7.408	0.401	32.574	11.092	97596.981	35576.768	0.192	0.239	0.781	2.984	101453.036	20578.807

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.090	0.000	0.472
Available MBytes	116340.135	116319.000	116438.000
Free System Page Table Entries	16397895.513	16397412.000	16398152.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	202404488.533	202371072.000	202584064.000
Pool Paged Bytes	175647138.133	175628288.000	175853568.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/19/2015 5:07:00 PM -- Preparing for testing ...

10/19/2015 5:07:05 PM -- Attaching databases ...

10/19/2015 5:07:05 PM -- Preparations for testing are complete.

10/19/2015 5:07:05 PM -- Starting transaction dispatch ..

10/19/2015 5:07:05 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

10/19/2015 5:07:05 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

10/19/2015 5:07:09 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/19/2015 5:07:09 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/19/2015 5:07:10 PM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/19/2015 5:07:10 PM -- Performance logging started (interval: 15000 ms).

10/19/2015 5:07:10 PM -- Attaining prerequisites:

10/19/2015 5:18:52 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 967249900.0 (lower bound: 966367600.0, upper bound: none)

10/19/2015 7:18:52 PM -- Performance logging has ended.

10/19/2015 8:17:09 PM -- JetInterop batch transaction stats: 10704, 10703, 10703 and 10703.

10/19/2015 8:17:09 PM -- Dispatching transactions ends.

10/19/2015 8:17:09 PM -- Shutting down databases ...

10/19/2015 8:17:11 PM -- Instance4188.1 (complete), Instance4188.2 (complete), Instance4188.3 (complete) and Instance4188.4 (complete)

10/19/2015 8:17:11 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_19_17_7_9.blg has 526 samples.

10/19/2015 8:17:11 PM -- Creating test report ...

10/19/2015 8:17:14 PM -- Instance4188.1 has 7.7 for I/O Database Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.1 has 0.7 for I/O Log Writes Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.1 has 0.7 for I/O Log Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.2 has 7.6 for I/O Database Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.2 has 0.7 for I/O Log Writes Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.2 has 0.7 for I/O Log Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.3 has 7.3 for I/O Database Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.3 has 0.2 for I/O Log Writes Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.3 has 0.2 for I/O Log Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.4 has 7.4 for I/O Database Reads Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.4 has 0.2 for I/O Log Writes Average Latency.

10/19/2015 8:17:14 PM -- Instance4188.4 has 0.2 for I/O Log Reads Average Latency.

10/19/2015 8:17:14 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/19/2015 8:17:14 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/19/2015 8:17:14 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_19_17_7_9.xml has 479 samples queried.

8.3 Streaming backup

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance800.1	961416.03	01:35:39	167.52
Instance800.2	961416.03	01:34:59	168.68
Instance800.3	961408.03	01:28:58	180.10
Instance800.4	961408.03	01:29:02	179.96
Avg			174.07
Sum			696.26

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

Instance800.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance800.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance800.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance800.4 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
Instance800.1	2.145	0.000	670.517	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance800.2	2.129	0.000	676.910	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance800.3	1.745	0.000	720.282	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance800.4	1.748	0.000	722.374	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.179	0.476	1.428
Available MBytes	117318.471	117303.000	117329.000
Free System Page Table Entries	16398098.419	16397569.000	16398329.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	210542205.990	210468864.000	210706432.000
Pool Paged Bytes	205022663.707	205008896.000	205086720.000

Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/21/2015 1:33:23 PM -- Preparing for testing ...

10/21/2015 1:33:27 PM -- Attaching databases ...

10/21/2015 1:33:27 PM -- Preparations for testing are complete.

10/21/2015 1:33:32 PM -- Performance logging started (interval: 30000 ms).

10/21/2015 1:33:32 PM -- Backing up databases ...

10/21/2015 3:09:12 PM -- Performance logging has ended.

10/21/2015 3:09:12 PM -- Instance800.1 (100% processed), Instance800.2 (100% processed), Instance800.3 (100% processed) and Instance800.4 (100% processed)

10/21/2015 3:09:12 PM -- C:\Program Files\Exchange Jetstress\DatabaseBackup_2015_10_21_13_33_27.blg has 191 samples.

10/21/2015 3:09:12 PM -- Creating test report ...

8.4 Soft recovery

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance3520.1	506	947.4279151
Instance3520.2	501	909.4977415
Instance3520.3	508	882.8165999
Instance3520.4	506	899.111295
Avg	505	909.713
Sum	2021	3638.8535515

Database Configuration

Instance3520.1 Log path: M:\LOG1 Database: M:\DB1\Jetstress001001.edb

Instance3520.2 Log path: M:\LOG2 Database: M:\DB2\Jetstress002001.edb

Instance3520.3 Log path: N:\LOG3 Database: N:\DB3\Jetstress003001.edb

Instance3520.4 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
Instance3520.1	19.335	0.746	277.043	2.105	39978.693	31580.247	19.243	0.000	2.632	0.000	200612.344	0.000
Instance3520.2	18.419	0.758	294.357	2.172	39959.748	32185.458	22.451	0.000	2.715	0.000	206342.549	0.000
Instance3520.3	17.979	0.240	305.752	2.278	40017.943	32543.048	15.249	0.000	2.847	0.000	207211.442	0.000
Instance3520.4	18.149	0.246	298.889	2.228	39996.219	32473.456	14.408	0.000	2.784	0.000	207098.998	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3520.1	0.000	0.000
Instance3520.2	0.000	0.000
Instance3520.3	0.000	0.000
Instance3520.4	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
Instance3520.1	19.335	0.746	277.043	2.105	39978.693	31580.247	19.243	0.000	2.632	0.000	200612.344	0.000
Instance3520.2	18.419	0.758	294.357	2.172	39959.748	32185.458	22.451	0.000	2.715	0.000	206342.549	0.000
Instance3520.3	17.979	0.240	305.752	2.278	40017.943	32543.048	15.249	0.000	2.847	0.000	207211.442	0.000
Instance3520.4	18.149	0.246	298.889	2.228	39996.219	32473.456	14.408	0.000	2.784	0.000	207098.998	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.616	0.000	1.981
Available MBytes	116361.989	116294.000	117299.000
Free System Page Table Entries	16398085.681	16397677.000	16398375.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	209939695.660	209924096.000	210100224.000
Pool Paged Bytes	199668932.085	199659520.000	199868416.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/19/2015 11:26:11 PM -- Preparing for testing ...

10/19/2015 11:26:15 PM -- Attaching databases ...

10/19/2015 11:26:15 PM -- Preparations for testing are complete.

10/19/2015 11:26:15 PM -- Starting transaction dispatch ..

10/19/2015 11:26:15 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)

10/19/2015 11:26:15 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)

10/19/2015 11:26:20 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/19/2015 11:26:20 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/19/2015 11:26:20 PM -- Operation mix: Sessions 2, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/19/2015 11:26:20 PM -- Performance logging started (interval: 15000 ms).

10/19/2015 11:26:20 PM -- Generating log files ...

10/20/2015 4:36:12 AM -- M:\LOG1 (101.2% generated), M:\LOG2 (100.2% generated), N:\LOG3 (101.6% generated) and N:\LOG4 (101.2% generated)

10/20/2015 4:36:12 AM -- Performance logging has ended.

10/20/2015 4:36:12 AM -- JetInterop batch transaction stats: 16803, 16803, 16803 and 16803.

10/20/2015 4:36:12 AM -- Dispatching transactions ends.

10/20/2015 4:36:12 AM -- Shutting down databases ...

10/20/2015 4:36:14 AM -- Instance3520.1 (complete), Instance3520.2 (complete), Instance3520.3 (complete) and Instance3520.4 (complete)

10/20/2015 4:36:14 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_19_23_26_20.blg has 1238 samples.

10/20/2015 4:36:14 AM -- Creating test report ...

10/20/2015 4:36:17 AM -- Instance3520.1 has 10.8 for I/O Database Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.1 has 0.7 for I/O Log Writes Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.1 has 0.7 for I/O Log Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.2 has 10.5 for I/O Database Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.2 has 0.7 for I/O Log Writes Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.2 has 0.7 for I/O Log Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.3 has 10.1 for I/O Database Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.3 has 0.2 for I/O Log Writes Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.3 has 0.2 for I/O Log Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.4 has 10.3 for I/O Database Reads Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.4 has 0.2 for I/O Log Writes Average Latency.

10/20/2015 4:36:17 AM -- Instance3520.4 has 0.2 for I/O Log Reads Average Latency.

10/20/2015 4:36:17 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.

10/20/2015 4:36:17 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.

10/20/2015 4:36:17 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_19_23_26_20.xml has 1237 samples queried.

10/20/2015 4:36:17 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_10_19_23_26_20.html was saved.

10/20/2015 8:33:22 AM -- Performance logging started (interval: 2000 ms).

10/20/2015 8:33:22 AM -- Recovering databases ...

10/20/2015 8:49:09 AM -- Performance logging has ended.

10/20/2015 8:49:09 AM -- Instance3520.1 (947.4279151), Instance3520.2 (909.4977415), Instance3520.3 (882.8165999) and Instance3520.4 (899.111295)

10/20/2015 8:49:09 AM -- C:\Program Files\Exchange Jetstress\SoftRecovery_2015_10_20_8_33_21.blg has 470 samples.

10/20/2015 8:49:09 AM -- Creating test report ...

9 Appendix B: Lenovo Bill of Materials

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

9.1 BOM for compute servers

This section contains the 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	Add'l Intel Xeon Processor E5-2603 v3 6C 1.6GHz 15MB 1600MHz 85W	1
A58Z	System x3550 M5, four 3.5-inch drive Base Chassis	1
A59V	System x3550 M5 Planar	1
A5AH	System x3550 M5 PCIe Riser 1 (one 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, four 3.5-inch 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
A2W6	Emulex 16 Gb 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 BOM for the network switches. Two networking switches per data center are required to meet the design specifications described in this document. The following table lists the BOM for a single networking switch.

Table 5: 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 BOM for the storage. One S3200 storage system per data center is required to meet the design specifications described in this document. The following table lists the BOM for a single S3200.

Table 7: Lenovo Storage S3200

Code	Description	Quantity
6411HC2	Lenovo Storage S3200 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
AT23	Lenovo Storage S3200 FC and iSCSI Controller	2
ASR7	Lenovo 3m LC-LC OM3 MMF Cable	4
AT29	Lenovo Storage S2200/S3200 16G 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, System x, TruDDR4, and VMready are trademarks of Lenovo.

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

Intel, Intel Inside (logos), and 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.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

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 addresses 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.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

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.