30,000 Mailbox Resiliency Solution for Microsoft Exchange 2013 using Lenovo ThinkAgile HX7520 Appliances and AHV

Last update: 04 December 2017

Provides a technical overview of Lenovo ThinkAgile HX Solution for Microsoft Exchange 2013

Contains performance results for 30,000 Mailboxes using Microsoft ESRP Storage Program

Uses HX7520 appliances with Nutanix Acropolis Hypervisor (AHV)

Shows class-leading hyper-converged performance of HX7520 appliance

Kenny Bain
Pawan Sharma
Mike Perks

LENOVO PRESS

Table of Contents

1 Overview

This document provides information on Lenovo ThinkAgile HX7520 solution for Microsoft Exchange Server, based the Microsoft Exchange *Solution Reviewed Program (ESRP) – Storage* program[1].

The intended audience for this document is technical IT architects, system administrators, and managers who are interested in executing virtualized Microsoft Exchange workloads on the Lenovo ThinkAgile HX Series Nutanix Appliances. The abbreviated term of HX Series is used in the remainder of this document.

HX Series appliances provide a hyper-converged infrastructure which seamlessly pools compute and storage to deliver high performance for the virtual workloads and provides flexibility to combine the local storage using a distributed file system to eliminate shared storage such as SAN or NAS. These factors make the solution cost effective without compromising the performance.

This document describes the solution for executing 30,000 mailboxes with Microsoft Exchange Server 2013 on Lenovo ThinkAgile HX7520 appliances using the Nutanix Acropolis Hypervisor (AHV). The mailbox resiliency data group with two database copies executes on 2 four node clusters.

The cluster of 4 Lenovo ThinkAgile HX7520 appliances proved more than capable of handling the high IOPs generated by 30,000 mailboxes. Part of the reason for this is because each Lenovo ThinkAgile HX7520 appliance uses 3 HBAs for the 24 drives. As a consequence the cluster has an average IOPS rate that is 23% better and an average latency improvement of 2 milliseconds over other systems with less HBAs.

Chapter 2 provides a technical overview of the HX Series appliances and explains why the combination of Lenovo servers and Nutanix software provides best of breed system performance and reliability.

Chapter 2.2 provides a solution overview of Microsoft Exchange on Lenovo ThinkAgile HX7520 appliances with the deployment best practices, recommended appliance configuration, and the performance test results from the Microsoft ESRP storage test. The detailed log files are in appendix A.

1.1 Disclaimer

This document has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty, express or implied, with respect to, the accuracy of the contents of this document.

The information contained in this document represents the current view of Lenovo on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted to be a commitment on the part of Lenovo, and Lenovo cannot guarantee the accuracy of any information presented after the date of publication.

---

[1] The *ESRP – Storage* program was developed by 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, use the following website: http://technet.microsoft.com/en-us/exchange/ff182054.aspx.

2 Technical overview of HX Series appliances

Lenovo ThinkAgile HX Series appliances are designed to help you simplify IT infrastructure, reduce costs, and accelerate time to value. These hyper-converged appliances from Lenovo combine industry-leading hyper-convergence software from Nutanix with Lenovo enterprise platforms.

2.1 Server Components

The Lenovo ThinkAgile HX series appliances are available in a variety of Xeon Scalable processor-based models to support different workloads. In particular the Lenovo ThinkAgile HX7520 is recommended for Microsoft Exchange and is shown below

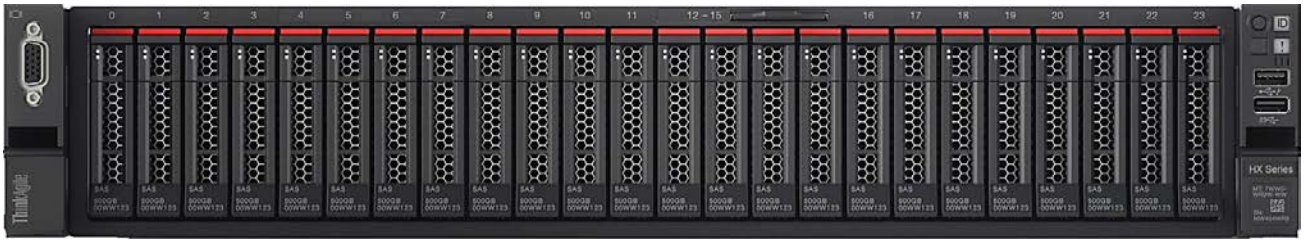

The table below provides a summary of the configuration options for the HX7520.

Appliance Model	Form Factor	Solid State Drives (SSD)	Hard Disk Drives (HDD)	Memory	CPU	2 or 4 Port LOM	GPUs
HX7520	2.5" drives 2U	4 x 1.92 TB 4 x 3.84 TB	20 x 1TB 20 x 2TB	384 GB 768 GB 1536 GB	2 x G 6126 2 x G 6152 2 x P 8153 2 x P 8176	1 x 10Gb RJ45 1 x 10Gb SFP+	N/A

For more information see the product guide for Lenovo ThinkAgile HX7520: lenovopress.com/lp0730

2.2 Software components

This section gives an overview of the software components used in the solution.

2.2.1 Hypervisor

The ThinkAgile HX Series appliances support the following hypervisors:

- Nutanix Acropolis Hypervisor based on KVM (AHV)
- VMware ESXi 6.0 U3
- VMware ESXi 6.5 U1

The HX Series appliances come standard with the Nutanix Acropolis Hypervisor (AHV) or VMware ESXi hypervisors preloaded in the factory.

2.2.2 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized systems management solution that helps administrators deliver infrastructure faster. This solution integrates easily with Lenovo servers, ThinkAgile HX Series appliances, and Flex System, providing automated agent-less discovery, monitoring, firmware updates, and configuration management.

Lenovo XClarity Pro goes one step further and provides entitlement to additional functions such as XClarity Integrators for Microsoft System Center and VMware vCenter, XClarity Administrator Configuration Patterns and Service and Support.

Lenovo XClarity Administrator is an optional software component for the Lenovo ThinkAgile HX Series appliances but Lenovo XClarity Pro licenses are always provided with ThinkAgile XSN models.

Lenovo XClarity Administrator which can be used to manage firmware upgrades outside of the Nutanix Prism web console. Note that XClarity should not be used to install hypervisors and Nutanix Foundation should be used instead.

Lenovo XClarity Administrator is provided as a virtual appliance that can be quickly imported into a virtualized environment. XClarity can either be installed on a separate server or a server within a Nutanix cluster providing that the hardware management network with the server IMMs is routable from the server hosting the XClarity VM.

Figure 1 shows the Lenovo XClarity administrator interface with some HX Series appliances.

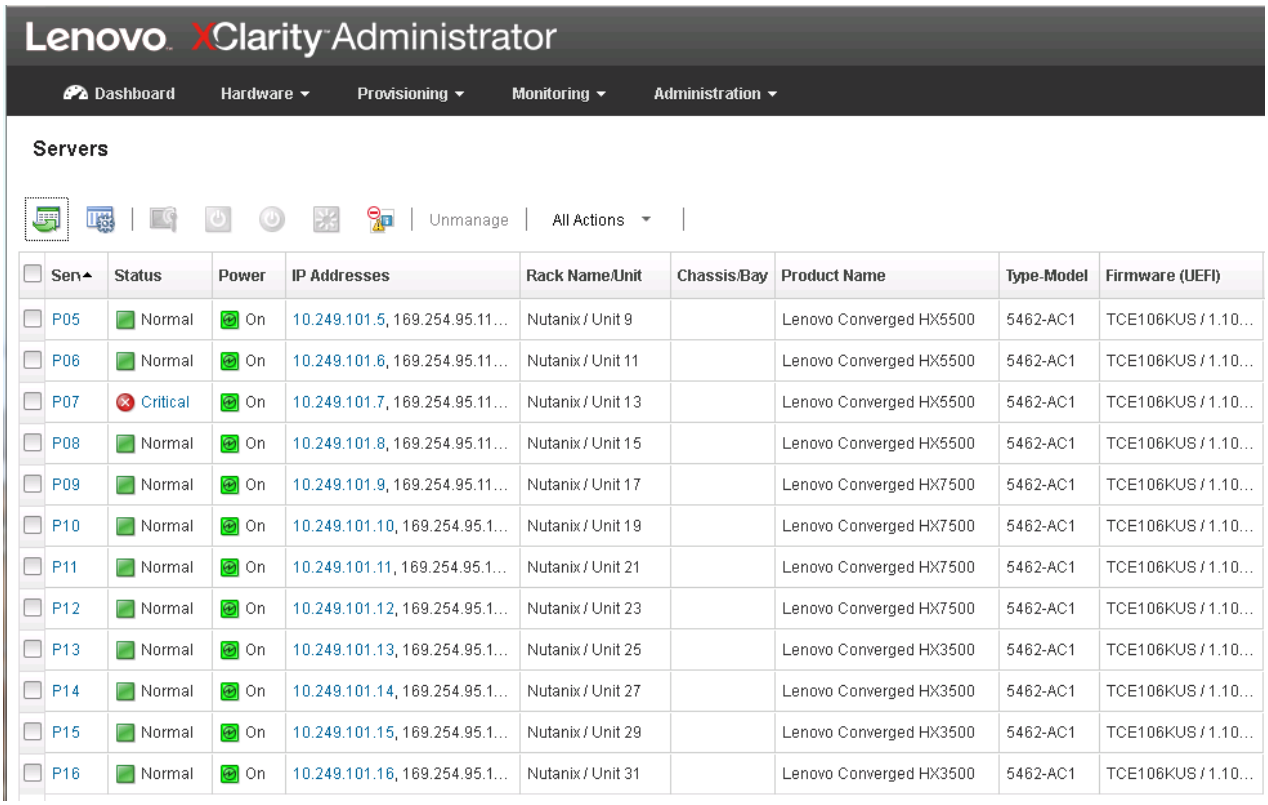

Ser	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product Name	Type-Model	Firmware (UEFI)
P05	Normal	On	10.249.101.5, 169.254.95.11...	Nutanix / Unit 9		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P06	Normal	On	10.249.101.6, 169.254.95.11...	Nutanix / Unit 11		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P07	Critical	On	10.249.101.7, 169.254.95.11...	Nutanix / Unit 13		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P08	Normal	On	10.249.101.8, 169.254.95.11...	Nutanix / Unit 15		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P09	Normal	On	10.249.101.9, 169.254.95.11...	Nutanix / Unit 17		Lenovo Converged HX7500	5462-AC1	TCE106KUS / 1.10...
P10	Normal	On	10.249.101.10, 169.254.95.1...	Nutanix / Unit 19		Lenovo Converged HX7500	5462-AC1	TCE106KUS / 1.10...
P11	Normal	On	10.249.101.11, 169.254.95.1...	Nutanix / Unit 21		Lenovo Converged HX7500	5462-AC1	TCE106KUS / 1.10...
P12	Normal	On	10.249.101.12, 169.254.95.1...	Nutanix / Unit 23		Lenovo Converged HX7500	5462-AC1	TCE106KUS / 1.10...
P13	Normal	On	10.249.101.13, 169.254.95.1...	Nutanix / Unit 25		Lenovo Converged HX3500	5462-AC1	TCE106KUS / 1.10...
P14	Normal	On	10.249.101.14, 169.254.95.1...	Nutanix / Unit 27		Lenovo Converged HX3500	5462-AC1	TCE106KUS / 1.10...
P15	Normal	On	10.249.101.15, 169.254.95.1...	Nutanix / Unit 29		Lenovo Converged HX3500	5462-AC1	TCE106KUS / 1.10...
P16	Normal	On	10.249.101.16, 169.254.95.1...	Nutanix / Unit 31		Lenovo Converged HX3500	5462-AC1	TCE106KUS / 1.10...

Figure 1: XClarity Administrator interface

2.2.3 Nutanix Prism

Nutanix Prism gives administrators a simple and elegant way to manage virtual environments. Powered by advanced data analytics and heuristics, Prism simplifies and streamlines common workflows within a data center.

Nutanix Prism is a part of the Nutanix software preloaded on the appliances and offers the following features:

- Single point of control
    - Accelerates enterprise-wide deployment
    - Manages capacity centrally
    - Adds nodes in minutes
    - Supports non-disruptive software upgrades with zero downtime
    - Integrates with REST APIs and PowerShell
- Monitoring and alerting
    - Tracks infrastructure utilization (storage, processor, memory)
    - Centrally monitors multiple clusters across multiple sites
    - Monitors per virtual machine (VM) performance and resource usage
    - Checks system health
    - Generates alerts and notifications
- Integrated data protection
    - Offers customizable RPO/RTO and retention policies
    - Supports configurable per-VM replication (1:1, 1:many and many:1)
    - Provides efficient VM recovery
    - Deploys affordable data recovery (DR) and backup to the cloud
- Diagnostics and troubleshooting
    - Provides time-based historical views of VM activity
    - Performs proactive alert analysis
    - Correlates alerts and events to quickly diagnose issues
    - Generates actionable alerts and reduces resolution times
    - Analyzes trending patterns for accurate capacity planning

2.2.4 ThinkAgile Network Orchestrator for Nutanix

The Lenovo® ThinkAgile™ Network Orchestrator is a unique feature of the Lenovo RackSwitch CNOS (Cloud Network OS) network switch firmware that automatically provisions the switches as needed, on-the-fly, to support changes in the virtual network, such as the creation, moving, and shutdown of virtual machines, as well as manipulation of guest virtual machines on VLANs. These tasks are performed dynamically in response to PRISM commands, with the switches configured to detect the changes and act upon them. This software capability simplifies the server administrator's tasks by reducing the need to provision the switches, reducing

maintenance windows, reducing human error, and dramatically saving time and administrative costs. The switches learn of changes in the environment from PRISM and dynamically modify their configurations as needed.

The ThinkAgile Network Orchestrator capability is supported with the CNOS version 10.3.2.0 or higher and the Acropolis hypervisor AOS version 5.0.2 and higher. See the following paper for more details:

lenovopress.com/lp0604-thinkagile-network-orchestrator-for-nutanix

2.2.5 Nutanix Foundation

Nutanix Foundation is a separate utility that you use to orchestrate the installation of hypervisors and Nutanix software on one or more nodes. The maximum number of nodes that can be deployed at one time is 20.

Foundation is available both as a stand-alone VM and also integrated into the CVM. Because CVM is pre-installed in the factory, the CVM integration of Foundation simplifies the deployment and cluster creation of new servers delivered from the factory.

2.2.6 Nutanix Controller VM

The Nutanix Controller VM (CVM) is the key to hyper-converged capability and each node in a cluster has its own instance. Figure 2 shows the main components of the CVM.

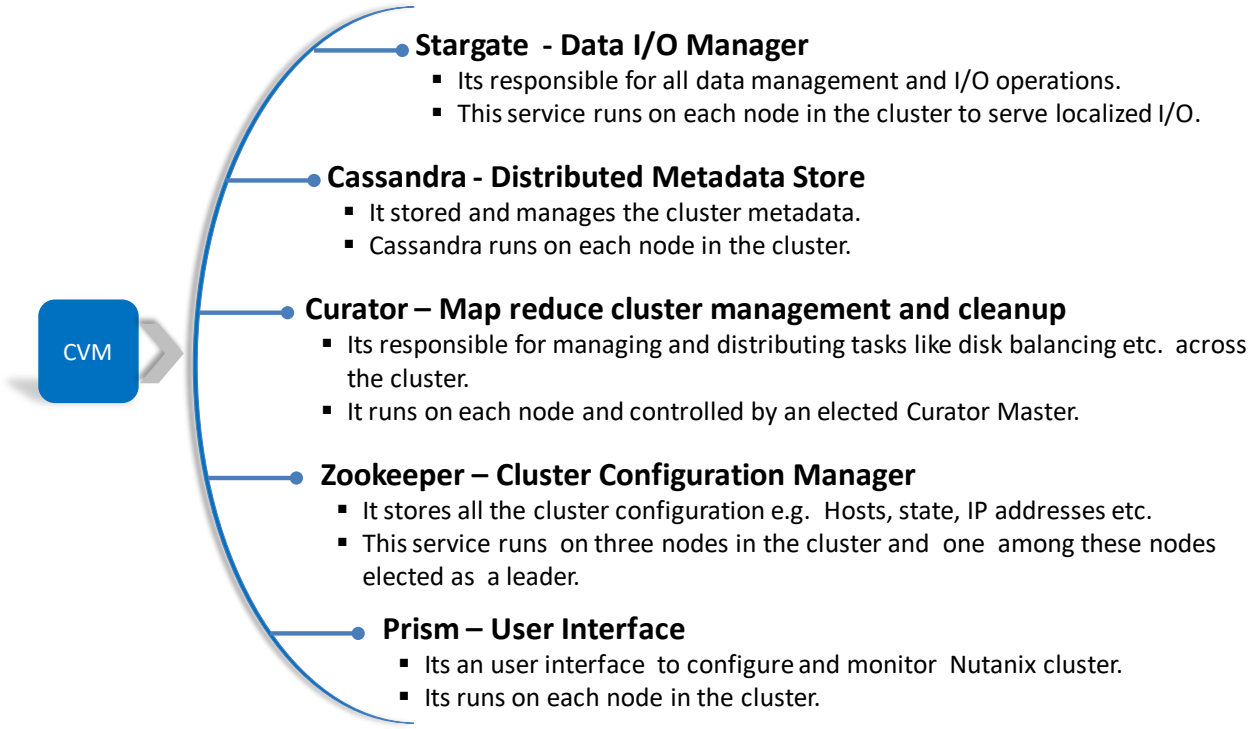

Figure 2: Controller VM components

The CVM works as interface between the storage and hypervisor to manage all I/O operations for the hypervisor and user VMs running on the nodes as shown in Figure 3.

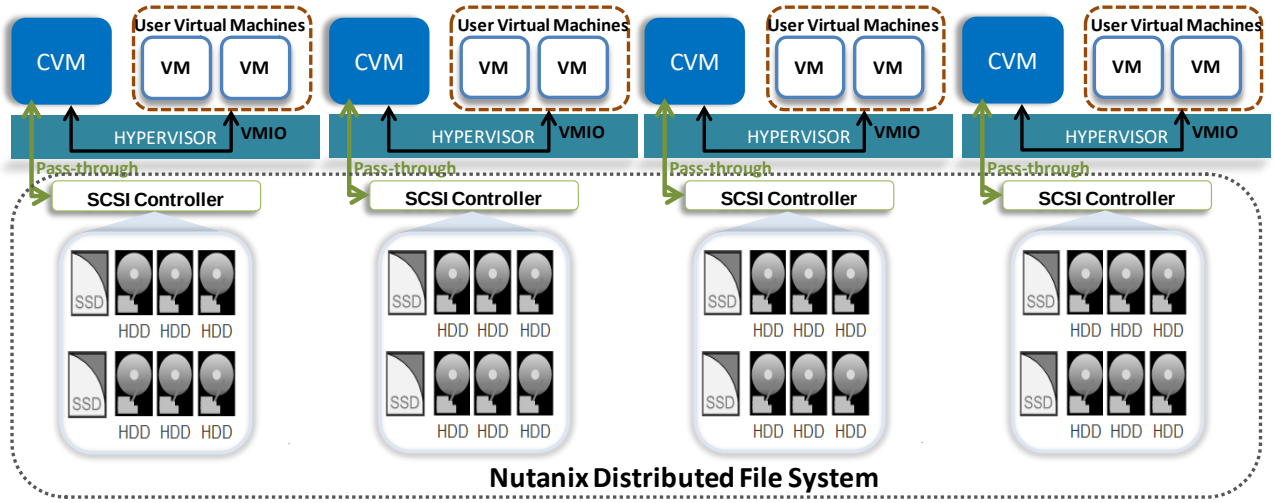

Figure 3: CVM interaction with Hypervisor and User VMs

CVM virtualizes all the local storage attached to each node in a cluster and presents it as centralized storage array using Nutanix Distributed File System (NDFS). All I/O operations are handled locally to provide the highest performance. See section 2.5 for more details on the performance features of NDFS.

2.3 Data network components

The data network is the fabric that carries all inter-node storage I/O traffic for the shared Lenovo HX distributed file system, in addition to the user data traffic via the virtual Network Interface Cards (NICs) exposed through the hypervisor to the virtual machines.

Each HX Series appliance contains between zero and two dual-port 10GbE network adapters as well as 4 on-board 1GbE ports. The hypervisors are configured by the Nutanix software so that the fastest network ports on the appliance are pooled for the data network. The hypervisor VM management network should use the same network. Because all of the network ports are pooled, each appliance only needs two network IP addresses; one for the hypervisor and one for the Nutanix CVM. These IP addresses should be all on the same subnet.

All storage I/O for virtual machines (VMs) running on a HX Series appliance node is handled by the hypervisor on a dedicated private network. The I/O request is handled by the hypervisor, which then forwards the request to the private IP on the local controller VM (CVM). The CVM then performs the remote data replication with other nodes in the cluster using its external IP address. In most cases, read request traffic is served locally and does not enter the data network. This means that the only traffic in the public data network is remote replication traffic and VM network I/O (i.e. user data). In some cases, the CVM will forward requests to other CVMs in the cluster, such as if a CVM is down or data is remote. Also, cluster-wide tasks, such as disk balancing, temporarily generate I/O traffic on the data network.

For more information on the network architecture see nutanixbible.com.

2.3.1 Data network switches

The following Lenovo 10GbE TOR switches are recommended for use in a HX Series cluster:

- Lenovo ThinkSystem NE1032 RackSwitch
- Lenovo RackSwitch G8272

Lenovo ThinkSystem NE1032 RackSwitch

The Lenovo ThinkSystem NE1032 RackSwitch (as shown in Figure 4) is a 1U rack-mount 10 Gb Ethernet switch that delivers lossless, low-latency performance with feature-rich design that supports virtualization, Converged Enhanced Ethernet (CEE), high availability, and enterprise class Layer 2 and Layer 3 functionality. The switch delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data.

The NE1032 RackSwitch has 32x SFP+ ports that support 1 GbE and 10 GbE optical transceivers, active optical cables (AOCs), and direct attach copper (DAC) cables. The switch helps consolidate server and storage networks into a single fabric, and it is an ideal choice for virtualization, cloud, and enterprise workload solutions.

Figure 4: Lenovo ThinkSystem NE1032 RackSwitch

For more information, see this website: lenovopress.com/lp0605

Lenovo RackSwitch G8272

The Lenovo RackSwitch G8272 uses 10Gb SFP+ and 40Gb QSFP+ Ethernet technology and is specifically designed for the data center. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic.

The RackSwitch G8272 (shown in Figure 5), is ideal for latency sensitive applications, such as high-performance computing clusters and financial applications. In addition to the 10 Gb Ethernet (GbE) and 40 GbE connections, the G8272 can use 1 GbE connections.

Figure 5: Lenovo RackSwitch G8272

For more information, see this website: lenovopress.com/tips1267

2.3.2 VLANs

It is a networking best practice to use VLANs to logically separate different kinds of network traffic. The following standard VLANs are recommended:

- Management Used for all management traffic for the hypervisor
- Storage network Used for NDFS storage traffic

The following ESXi specific VLANs are recommended:

- vSphere vMotion Used to move VMs from one server to another.
- Fault Tolerance Used to support the fault tolerance (FT) feature of vSphere.

In addition, each workload application might require one or more VLANs for its logical networks. For larger networks with many workloads, it is easy to run out of unique VLANs. In this case, VXLANs could be used.

The procedure for configuring VLANs for HX Series appliances is outside of the scope of this document.

2.3.3 Redundancy

It is recommended that two top of rack (TOR) switches are used for redundancy in the data network. It is recommended to use two dual-port 10Gbps network adapters in the HX Series appliances for workloads that require high throughput on the network or scale-out cluster deployments. This will effectively provide two redundant, bonded links per host for 20Gbps of bandwidth per logical link. Note that by default, the bonding configuration in the HX appliances is active/passive, but this can be changed to active/active with the proper configuration on the hypervisor host and switch side.

In order to support the logical pairing of the network adapter ports and to provide automatic failover of the switches, the Lenovo ThinkSystem NE1032 RackSwitch and G8272 support virtual link aggregation groups (VLAGs). When VLAG is enabled over the inter-switch link (ISL) trunk, it enables logical grouping of these switches. When one of the switches is lost, or the uplink from the host to the switch is lost, the connectivity is automatically maintained over the other switch.

Figure 6 shows the two scenarios of single port and dual-port connectivity using the Lenovo RackSwitch G8272. Note the connections into the customer switch and also the extra link between the data switches and the management switch that is required for initial setup only.

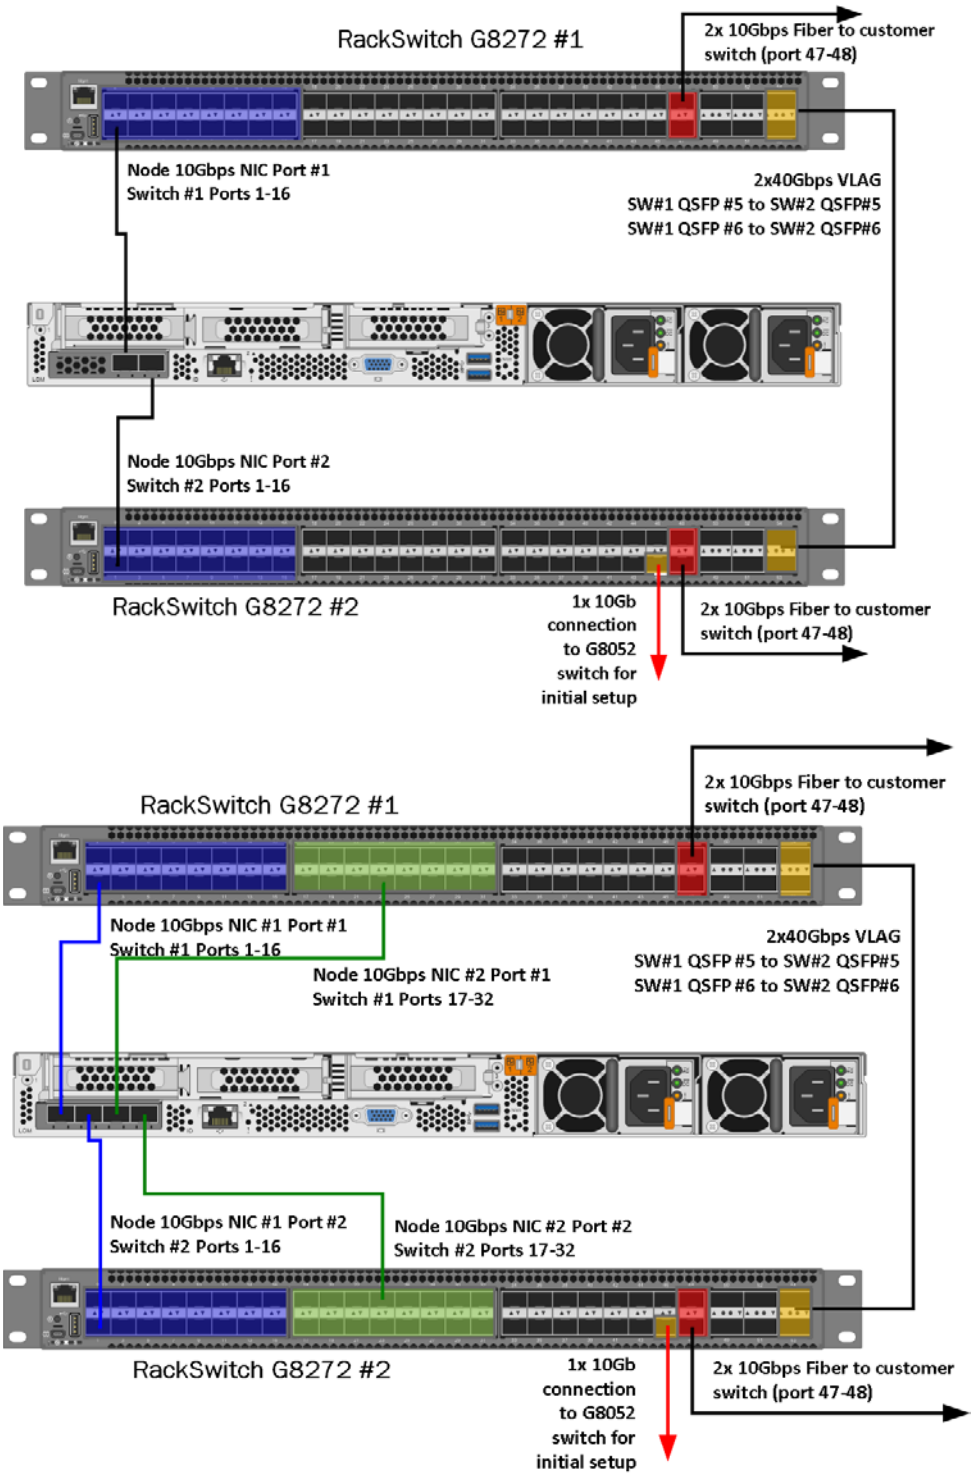

Figure 6: Data network with single and dual adapters

In addition, the Lenovo Cloud Network Operating System (CNOS) should be used on the G8272 switches. A detailed description of the CNOS operating system and its application to the HX Series deployments is provided in the following paper:

Networking Guide for Lenovo ThinkAgile HX Series Appliances (CNOS Switch Firmware)

lenovopress.com/lp0595-networking-guide-for-lenovo-converged-hx-series-nutanix-cnos.

2.4 Hardware management network components

The hardware management network is used for out-of-band access to the HX appliances via the optional Lenovo XClarity Administrator. It may also be needed to re-image an appliance. All systems management for HX appliances is handled in-band via Intelligent Platform Management Interface (IPMI) commands.

The dedicated Integrated Management Module (IMM) port on all of the Lenovo ThinkAgile HX series appliances needs to be connected to a 1GbE TOR switch as shown in Figure 7.

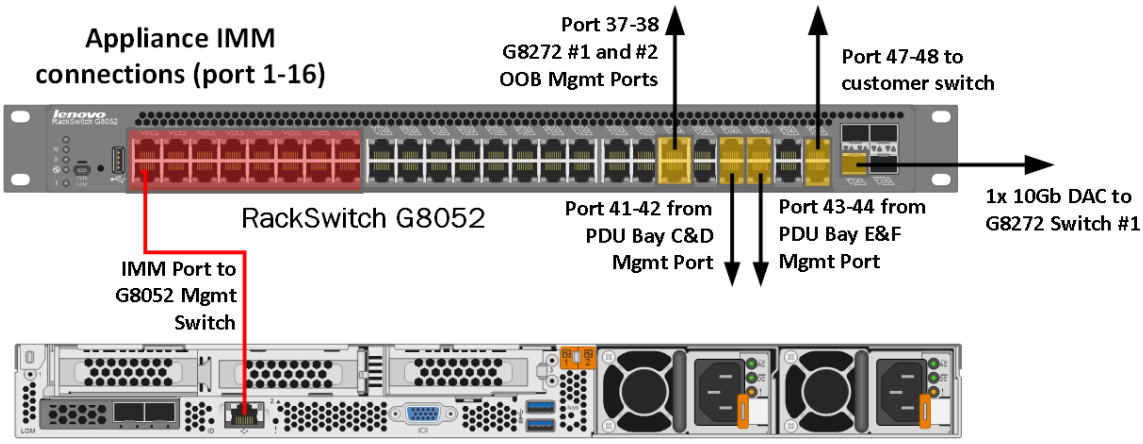

Figure 7: IMM 1GbE management network

2.4.1 Hardware management switches

The following Lenovo 1GbE TOR switches are recommended for use in a HX Series cluster:

- Lenovo RackSwitch G7028
- Lenovo RackSwitch G8052

Lenovo RackSwitch G7028

The Lenovo RackSwitch G7028 (as shown in Figure 8) is a 1 Gb top-of-rack switch that delivers line-rate Layer 2 performance at an attractive price. G7028 has 24 10/100/1000BASE-T RJ45 ports and four 10 Gb Ethernet SFP+ ports. It typically uses only 45 W of power, which helps improve energy efficiency.

Figure 8. Lenovo RackSwitch G7028

For more information, see this website: lenovopress.com/tips1268

Lenovo RackSwitch G8052

The Lenovo System Networking RackSwitch G8052 (as shown in Figure 9) is an Ethernet switch that is designed for the data center and provides a virtualized, cooler, and simpler network solution. The Lenovo RackSwitch G8052 offers up to 48 1 GbE ports and up to four 10 GbE ports in a 1U footprint. The G8052 switch is always available for business-sensitive traffic by using redundant power supplies, fans, and numerous high-availability features.

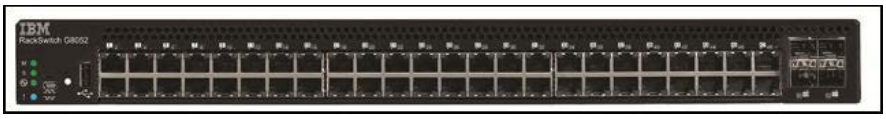

Figure 9: Lenovo RackSwitch G8052

For more information, see this website: lenovopress.com/tips0813

2.5 Reliability and performance features

Reliability and excellent performance are important for any workload but particularly for hyper-converged infrastructures like the HX Series appliances. These requirements are met through the following design features of Nutanix software combined with Lenovo Servers.

Hardware reliability

Lenovo uses the highest quality hardware components combined with firmware that is thoroughly tested. As a consequence Lenovo servers have been rated #1 in hardware reliability for the last 3 years. This is important as it lowers the frequency of a server failure which in turn lowers OPEX.

A HX appliance has redundant hardware components by including two power supplies, multiple chassis fans, two Intel CPUs, multiple memory DIMMs, multiple SSDs and HDDs, and optionally up to two dual-port network interface cards.

Hardware performance

The HX Series appliances have been carefully designed for performance. In addition to all of the usual attributes like processors and memory, the 24 drive HX7520 uses three HBA controllers instead of the one. As a consequence the latency is halved for some workloads that heavily utilize the cold tier. This allows a higher throughput and improved transaction rates.

Distributed file system

The Nutanix Distributed file system (NDFS) is an intelligent file system which virtualizes the local attached storage (SSD/HDD) on all the nodes in a cluster and presents it as single storage entity to cluster. Figure 10 shows the high level structure of NDFS:

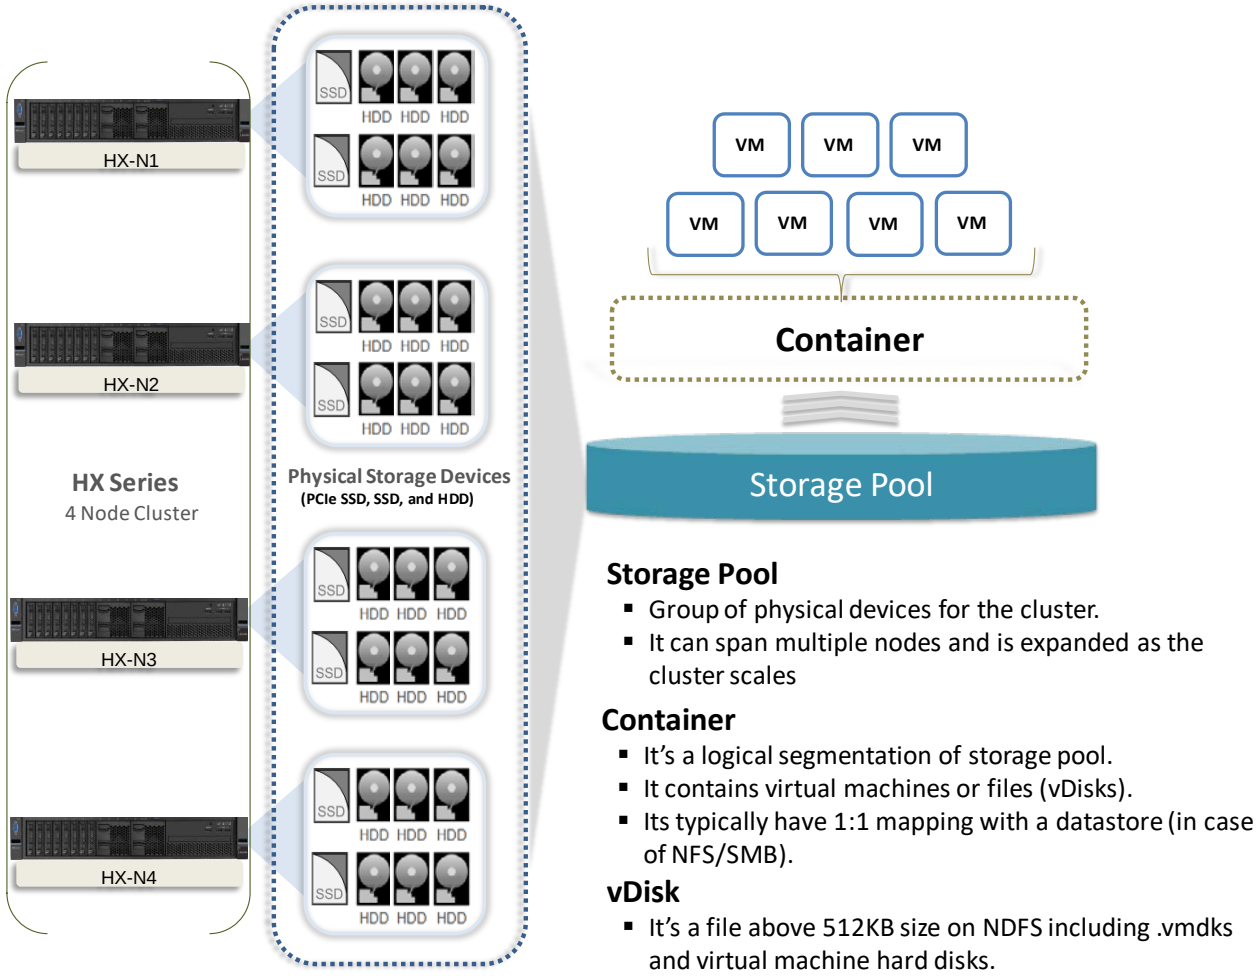

Figure 10: Nutanix Distributed File System

Data protection via replication

The Nutanix platform replication factor (RF) and checksum is used to ensure data redundancy and accessibility in the event of a node or disk failure or corruption. It uses an OpLog which acts as a staging area for incoming writes on low latency SSDs which are then replicated to the OpLogs for one or two other Controller VMs before acknowledging a successful write. This approach ensures that data available in at least two to three different locations and is fault tolerant. While the data is being written a checksum is calculated and stored as part of its metadata.

In the case of a drive or node failure, that data is replicated out to more nodes to maintain the replication factor. A checksum is computed every time the data is read to ensure the data validity. If the checksum and data mismatch, then the data replica is read to replace the invalid copy.

Performance with data tiering

Nutanix uses a disk tiering concept in which disk resources (SSD and HDD) are pooled together to form a cluster wide storage tier. This tier can be accessed by any node within the cluster for data placement and can leverage the full tier capacity.

The following data tiering functions are provided:

- The SSD on a local node always has the highest tier priority for write I/O.
- If the local node's SSD is full then the other SSDs in the cluster are used for I/O.
- The NDFS Information Lifecycle Management (ILM) component migrates cold data from the local SSD to HDD to free up SSD space. It also moves heavily accessed data to the local SSD to provide high performance.

Performance by data locality

Data locality is a crucial factor for cluster and VM performance. In order to minimize latency the CVM will work to ensure that all I/O happens locally. This ensures optimal performance and provides very low latencies and high data transfer speeds that cannot be achieved easily with shared storage arrays, even if all-flash.

The following occurs in case of a VM migration or high availability event that moves a VM from Node-A to Node-B:

- The VM's data is provided by the CVM running on Node-B.
- All write I/O requests occur locally i.e. to the local storage of Node-B.
- When a request comes for reading old data, the I/O request is forwarded by Node-B to Node-A. NDFS detects that the I/O request originated from different node and migrates the data locally in the background i.e. from Node-A to Node-B so that all subsequent read I/O operations are served locally. This approach (migration only on a read) helps to avoid network flooding.

Performance of snapshots and clones

NDFS provides support for offloaded snapshots and clones using a redirect-on-write algorithm. When a snapshot or clone is created, the base vDisk is marked as read only and another vDisk is created with read/write permissions as shown in Figure 11 and Figure 12 below.

At this point both vDisks have the same block map - a metadata mapping of the vDisk to its corresponding extents. This approach reduces the overhead of creating snapshots and allows snapshots to be taken very quickly with little performance impact.

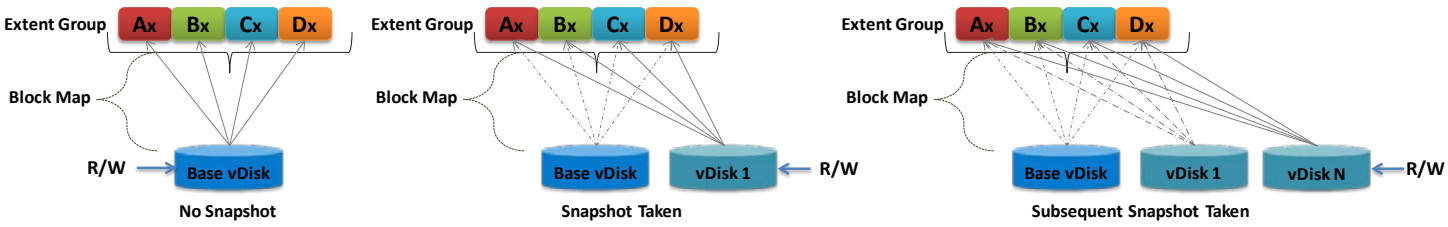

Figure 11: VM snapshots

When a VM is cloned the current block map is locked and then clones are created. These updates are metadata only so again no actual I/O takes place. The logic applies for clones of clones as well where a previously cloned VM acts as a base vDisk. All the clones inherit the prior block map and any new writes take place on the individual block maps.

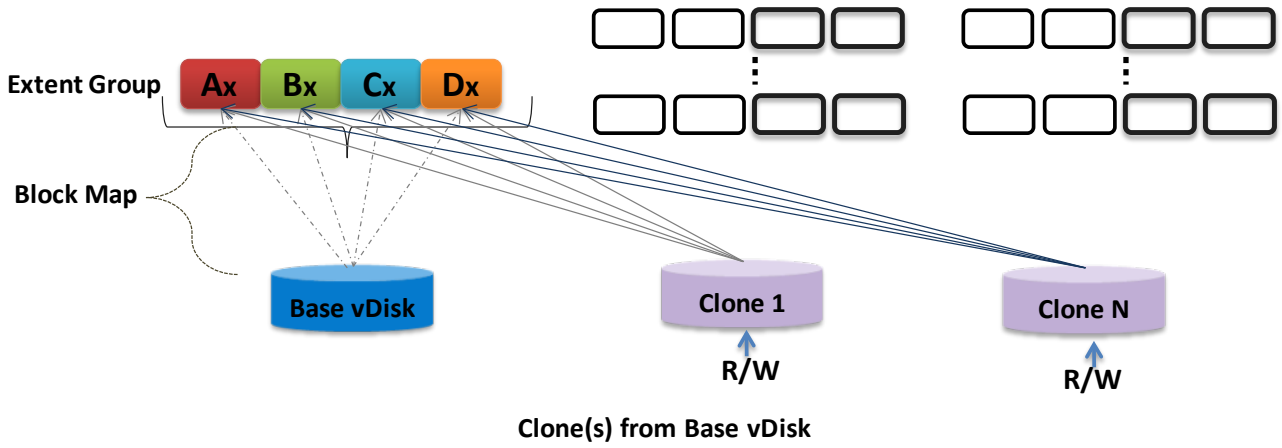

Figure 12: VM clones

Storage reduction via De-duplication and Compression

The Nutanix elastic de-duplication engine increases the effective capacity of a disk, as well as the RAM and cache of the system by removing duplicate data. It's an intelligent technology which performs following actions to increase storage efficiency:

- Sequential streams of data fingerprinted at 4K granularity
- Single instance of the shared VM data is loaded into the cache upon read
- Each node in a cluster performs its own fingerprinting and deduplication

The Nutanix capacity optimization engine is responsible for performing data transformations and compression to achieve data optimization. NDFS provides following compression methods:

- In-line compression sequential streams of data or large I/O sizes are compressed in memory before written to the disk
- Post-process compression whereby data is written in an uncompressed state and the curator framework is used to compress the data in a cluster wide manner

The Nutanix capacity optimization engine uses the Google snappy compression library to deliver good compression ratios with minimal compute overhead and very fast compression or decompression rates.

Elimination of "split-brain" errors

In a distributed system it is possible for one participant to become disconnected which will cause differences in the stored data. NDFS uses the proven "Paxos" algorithm to eliminate these "split-brain" issues by reaching a consensus (quorum) among the participants in a distributed system before the writes are made.

Drive reliability via active monitoring

The CVM actively monitors the performance of every drive in a node. The deterioration of a drive's performance may indicate that the drive is about to fail. The CVM proactively moves data off the drive before it fails and marks the drive offline and in need to replacement. The idea is to avoid the expensive data transfers to maintain data redundancy and possible loss of data.

3 Microsoft Exchange

Microsoft Exchange Server 2013 is the market leader in enterprise messaging and collaboration. Exchange Server 2013 builds upon the Exchange Server 2010 architecture and was redesigned for simplicity of scale, improved hardware utilization, and increased failure isolation. The goal of Exchange Server 2013 is to support people and organizations as their work habits evolve from a communication focus to a collaboration focus.

3.1 Solution overview

Figure 13 shows the architectural overview of the Microsoft Exchange solution using Lenovo ThinkAgile HX7520 appliances. This chapter does not address integrating Exchange with unified messaging solutions and handling edge transport routing and distribution.

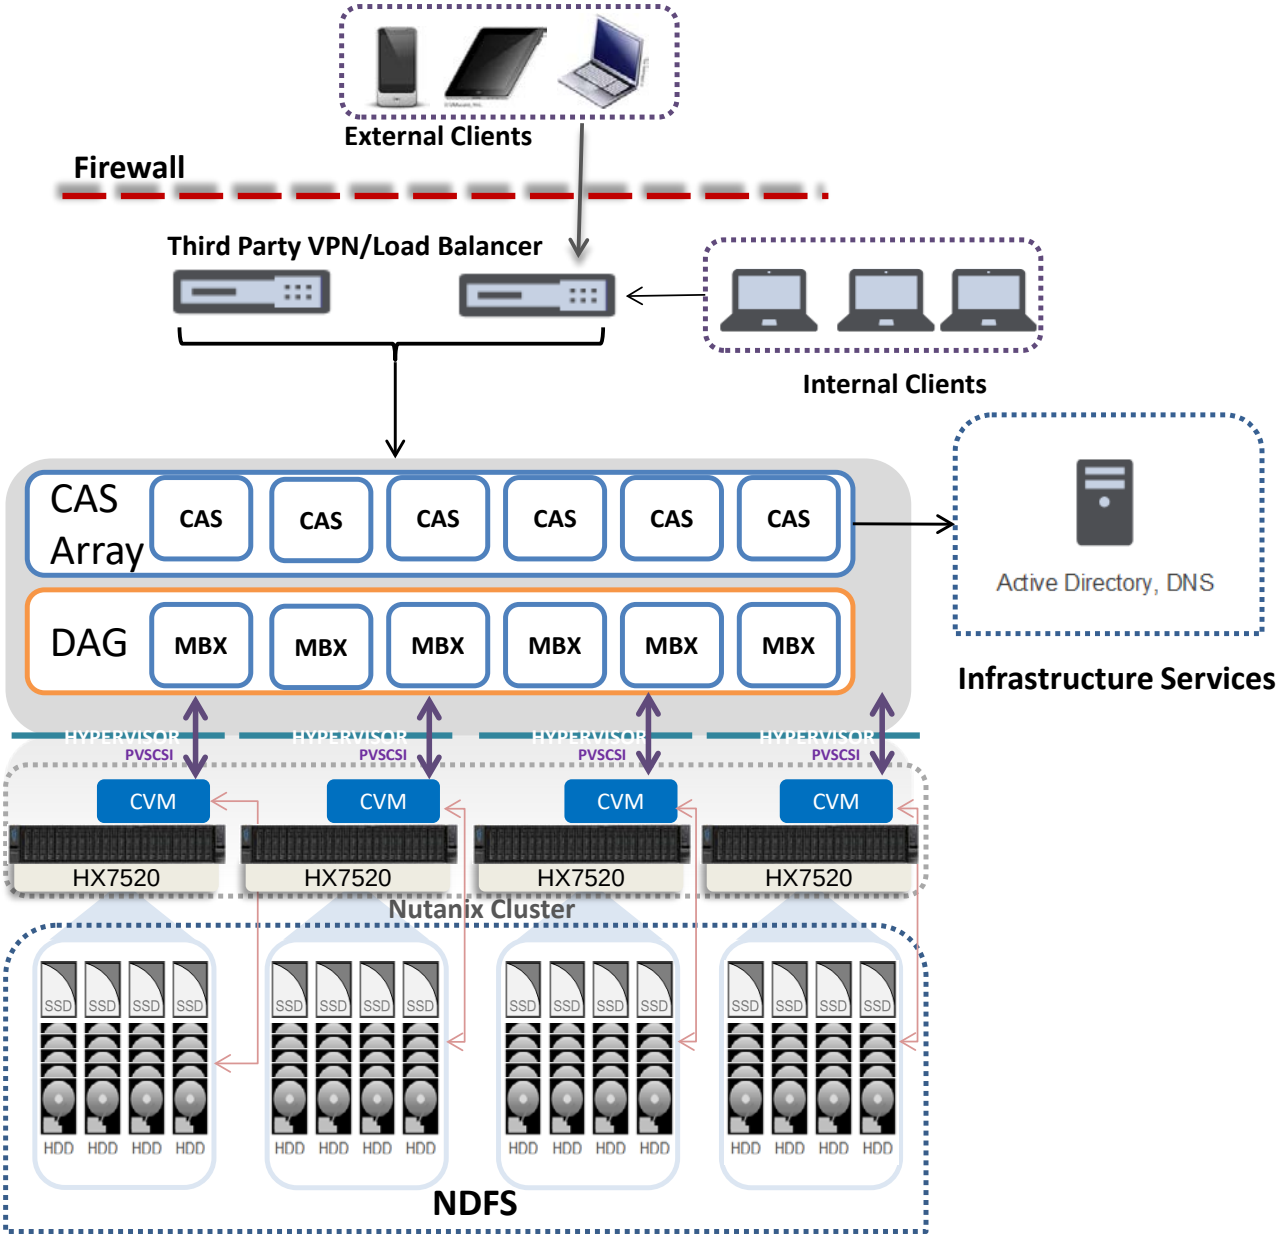

Figure 13. Lenovo ThinkAgile HX Series solution with Microsoft Exchange

The Client Access Server (CAS) role provides client protocols, SMTP, and unified messaging support. The Mailbox Server (MBX) role provides all of the data processing services. Lenovo recommends that these roles are combined into a multi-role server.

For load balancing into the CAS layer either a network load balancer can be used with a CAS array object or a layer 4 or layer 7 load balancer can be used without the need for configuring a CAS array.

3.2 Component model

This section describes the logical component view of the Exchange Server 2013 environment. Figure 14 shows a high-level component model.

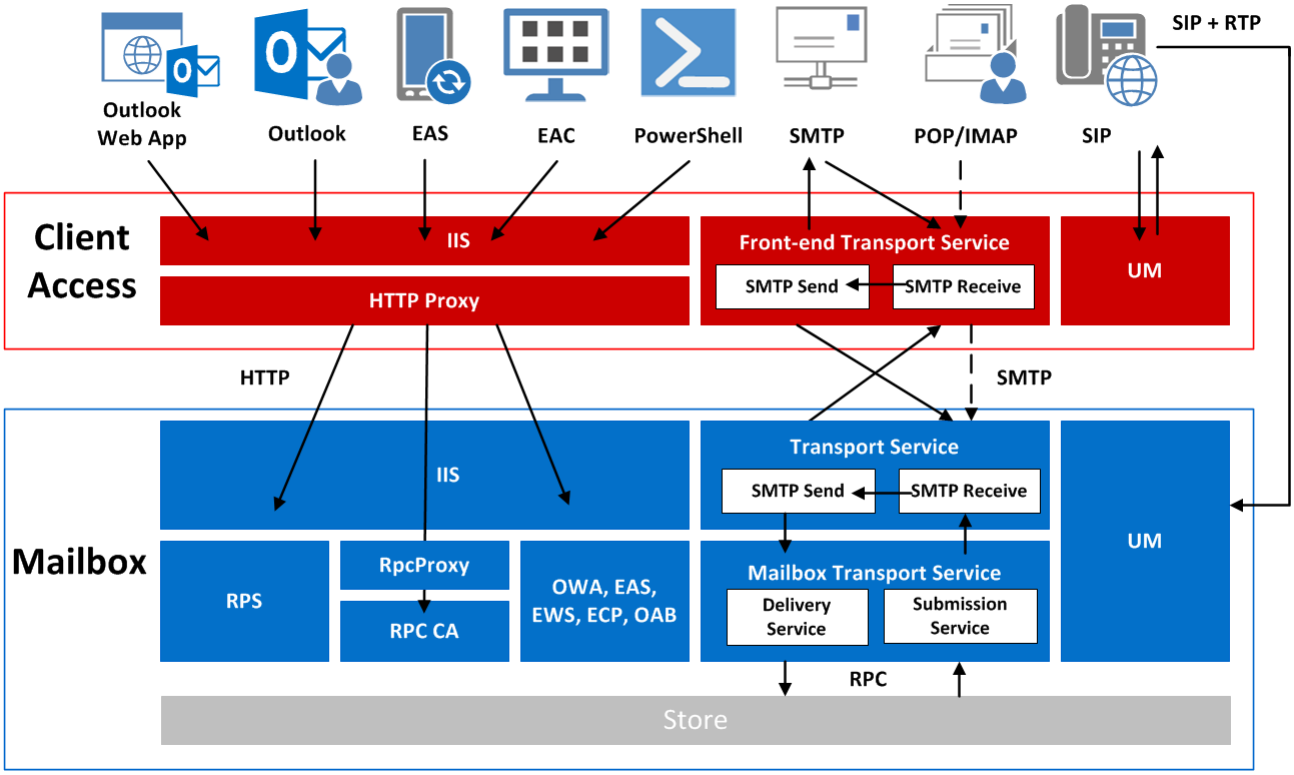

Figure 14. Exchange Server 2013 logical component view

The following basic concepts and terminology are used throughout this section:

Exchange Admin Center (EAC) – The EAC is the web-based management console in Microsoft Exchange Server 2013 that is optimized for on-premises, online, and hybrid Exchange deployments. The EAC replaces the Exchange Management Console (EMC) and the Exchange Control Panel (ECP), which were the two interfaces used to manage Exchange Server 2010.

Exchange Control Panel (ECP) – The ECP is a web application that runs on a Client Access Server and provides services for the Exchange organization.

Exchange Web Services (EWS) – EWS provides the functionality to enable client applications to communicate with the Exchange server.

Internet Information Services (IIS) – IIS is an extensible web server that was created by Microsoft for use with Windows NT family.

Internet Message Access Protocol (IMAP) – IMAP is a communications protocol for email retrieval and storage developed as an alternative to POP.

Microsoft Exchange ActiveSync (EAS) – EAS is a communications protocol that is designed for the synchronization of email, contacts, calendar, tasks, and notes from a messaging server to a smartphone or other mobile device.

Microsoft Outlook® Web App (OWA) – OWA (formerly Outlook Web Access) is a browser-based email client with which users can access their Microsoft Exchange Server mailbox from almost any web browser.

Offline Address Book (OAB) – The OAB is a copy of an address list collection that was downloaded so a Microsoft Outlook user can access the address book while disconnected from the server. Microsoft Exchange generates the new OAB files and then compresses the files and places them on a local share.

Outlook Anywhere – Outlook Anywhere is a service that provides RPC/MAPI connectivity for Outlook clients over HTTP or HTTPS by using the Windows RPC over HTTP component. In previous versions of Exchange Server, this function was used for remote or external access only. However, in Exchange Server 2013, all Outlook connectivity is via HTTP/HTTPS (even for internal clients).

Post Office Protocol (POP) – The POP is an application-layer Internet standard protocol that is used by local email clients to retrieve email from a remote server over a TCP/IP connection

Real-time Transport Protocol (RTP) – RTP is a network protocol for delivering audio and video over IP networks.

Remote PowerShell (RPS) – RPS allows you to use Windows PowerShell on your local computer to create a remote Shell session to an Exchange server if you do not have the Exchange management tools installed.

RPC Client Access (RPC) – In Microsoft Exchange Server 2007, the Client Access server role was introduced to handle incoming client connections to Exchange mailboxes. Although most types of client connections were made to the Client Access server, Microsoft Office Outlook still connected directly to the Mailbox server when it was running internally with the MAPI protocol.

A new service was introduced with Exchange Server 2010 to allow these MAPI connections to be handled by the Client Access server. The RPC Client Access service provides data access through a single, common path of the Client Access server, with the exception of public folder requests (which are still made directly to the Mailbox server). This change applies business logic to clients more consistently and provides a better client experience when failover occurs.

Remote Procedure Call over HTTP – The RPC over HTTP component wraps RPCs in an HTTP layer that allows traffic to traverse network firewalls without requiring RPC ports to be opened. In Exchange 2013, this feature is enabled by default because Exchange 2013 does not allow direct RPC connectivity.

Session Initiation Protocol (SIP) – SIP is a protocol that is used for starting, modifying, and ending an interactive user session that involves multimedia elements, such as video, voice, and instant messaging.

Simple Mail Transfer Protocol (SMTP) – SMTP is an Internet standard for email transmission.

Unified Messaging (UM) – UM allows an Exchange Server mailbox account that was enabled for UM to receive email, voice, and fax messages in the Inbox.

3.3 Exchange deployment best practices

This section describes recommended best practices for Microsoft Exchange mailboxes. See also this website for Nutanix Best Practices Guide: Virtualizing Microsoft Exchange: go.nutanix.com/virtualizing-microsoft-exchange-converged-infrastructure.html.

3.3.1 Data optimization

By default all Nutanix storage containers are thin provisioned which reduces unused capacity and automatically provisions additional storage capacity when needed. It is also very easy to add additional storage capacity for mailboxes by simply adding nodes to the cluster. It is also possible to set a storage reservation to guarantee a minimum amount of storage capacity.

Data compression can be used to further increase data capacity especially for data that is less frequently accessed. Lenovo recommends enabling compression with a delay of 1440 minutes (1 day) which minimizes the performance impact on I/O writes.

Data de-duplication is not recommended and should be disabled for active Exchange mailboxes because of the frequency of changes. Note that de-duplication may be beneficial for backup volumes which are not changed very often.

A resiliency factor of 2 is the default. This provides a minimum level of data redundancy but a resiliency factor of 3 might be important in some environments. Using erasure coding saves significant storage capacity but it only recommended for archive data.

3.3.2 Cluster high availability

The minimum number of nodes in each cluster is 3 and should be at least 4 to provide failover. The following high availability features are recommended for an AHV-based cluster:

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).

Lenovo recommends a DAG configuration of 2 database copies and optionally one lagged copy. With a data resiliency factor of 2, the effective number of copies of each mailbox is 4 and this allows two disk failures without losing data.

DR across datacenters can also be done using DAGs assuming there is sufficient band-width between the sites. The scenarios for active-active and active-passive DR sites using DAGs are outside the scope of this document.

3.3.3 Other best practices

Consider the following points regarding virtualizing Exchange:

- All Exchange 2013 server roles should be supported in a single VM.
- Some hypervisors include features for taking snapshots of VMs. However, VM snapshots are not application aware, and the use of snapshots can have unintended and unexpected consequences for a server application that maintains state data, such as Exchange. Therefore, making VM snapshots of an Exchange guest VM is not supported.

3.4 DAG architecture overview

The following section illustrates the Lenovo ThinkAgile HX7520 based Exchange 2013 mailbox resiliency solution Lenovo implemented for ESRP testing.

Compute per node: 2 x Intel(R) Xeon(R) Platinum 8170 CPU @ 2.10GHz Processors

RAM per node: 768GB

Raw storage per node: SSD: 4 x 1.92TB SATA SSDs – 7.68TB
HDD: 20 x 2TB SATA HDDs – 40TB

Raw storage per cluster: SSD: 16 x 1.92 TB SATA SSDs – 31TB
HDD: 80 x 2TB SATA HDDs – 160 TB

Figure 15 describes the high level DAG architecture of the 30,000 mailbox virtualized Exchange 2013 mailbox resiliency solution. This solution comprises of two Lenovo ThinkAgile HX7520 clusters with 4 nodes per cluster. Each node uses the Nutanix Acropolis Hypervisor (AHV).

The DAG has six Exchange 2013 mailbox servers and two database copies. The two database copies were placed on two physically isolated clusters.

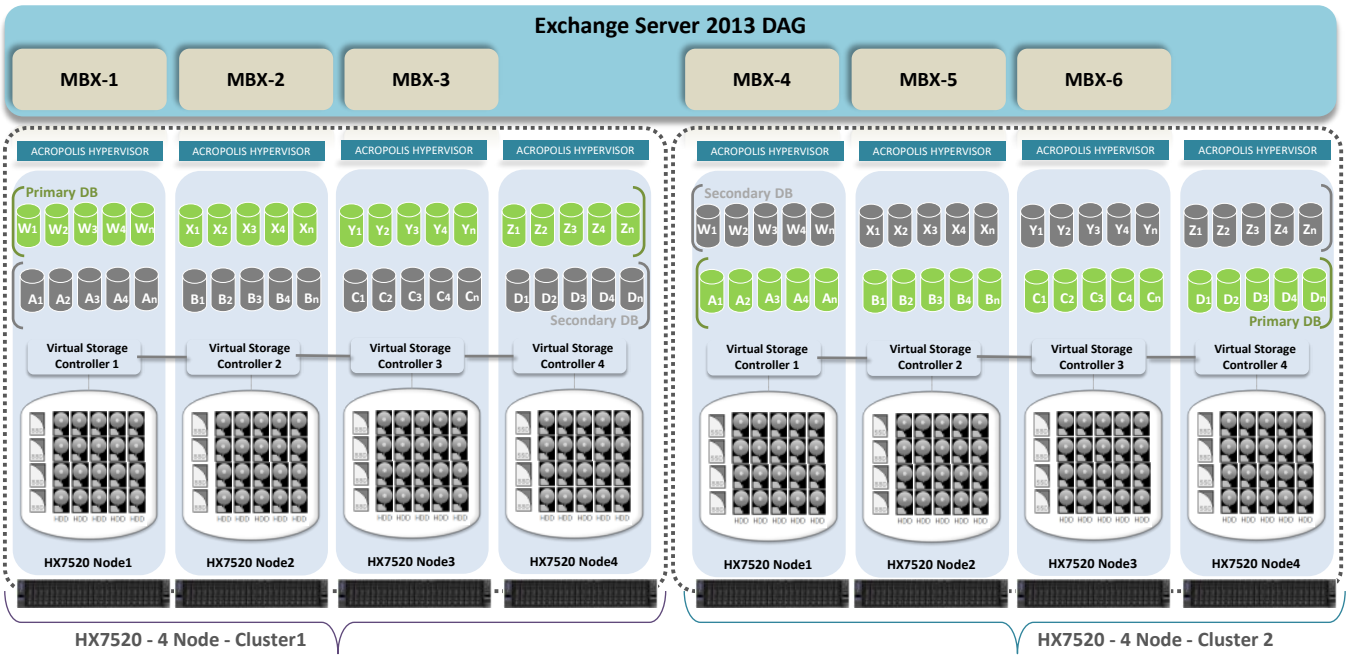

Figure 15. DAG Architecture using Exchange 2013 and Lenovo ThinkAgile HX7520 (30,000 Mailboxes)

On both Cluster1 and Cluster2, the Mailbox server (MBX) VMs were created on three nodes. All the Database/Logs volumes were connected to Mailbox servers using the iSCSI protocol. The primary database copy and secondary database copy are stored on two physically separated and isolated clusters. The two clusters can be located at the same datacentre or two different datacenters.

Each Mailbox Server VM is configured as follows:

- 40 vCPUs
- 512 GB RAM
- Windows Server 2016
- Microsoft Exchange 2013 with 10,000 mailboxes

The ESRP-Storage program focuses on storage solution testing to address performance and reliability issues with storage design. However, storage is not the only factor to take into consideration when designing a scale up Exchange solution. Other factors which affect the server scalability are: server processor utilization, server physical and virtual memory limitations, resource requirements for other applications, directory and network service latencies, network infrastructure limitations, replication and recovery requirements, and client usage profiles. All these factors are beyond the scope for ESRP-Storage. Therefore, the number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployment.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please refer to Microsoft's Troubleshooting Microsoft Exchange Server Performance, available at http://technet.microsoft.com/en-us/library/dd335215.aspx.

3.5 Targeted customer profile

The target customer profile for a medium enterprise Microsoft Exchange 2013 environment is as follows:

- 30,000 mailboxes of 1GB
- 6x Exchange 2013 servers (3x Tested)
- 0.06 IOPS per mailbox
- 24/7 background database maintenance
- Mailbox resiliency factor of 2
- 10 databases per host

3.6 Tested deployment environment

The section describes the tested deployment environment.

3.6.1 Simulated exchange configuration

The following table summarizes the simulated Exchange configuration.

Number of Exchange mailboxes simulated	30,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	6 (3 tested)

Number of active mailboxes/server	5,000 (5,000 active and 5,000 passive mailboxes per server and tested 10,000 active mailboxes per server)
Number of databases/host	10
Number of copies/database	2
Number of mailboxes/database	1000
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	0.06 IOPs / Mailbox
Database/LUN size	1.5 TB
Total database size for performance testing	30 TB
% storage capacity used by Exchange database[2]	38.3%

3.6.2 Storage hardware

The following table summarizes the storage hardware.

Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI)	iSCSI
Storage model and OS/firmware revision	HX7520 running Acropolis 5.1.3
Storage cache	47.69 GB per node
Number of storage controllers	4x virtual controller virtual machines
Number of storage ports	4 x 10 Gbe Port
Maximum bandwidth of storage connectivity to host	40 Gbps per node
Switch type/model/firmware revision	Lenovo RackSwitch G8272 (10GbE) Firmware version: 7.7.5
HBA model and firmware	Lenovo ThinkSystem 430-8i HBA
Number of HBA's/host	3
Host server type	3x Lenovo ThinkAgile HX7520 (2 x Intel(R) Xeon(R) Platinum 8170 CPU @ 2.10GHz) 768 GB RAM

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

Total number of disks tested in solution	96 (4 node cluster)
Maximum number of spindles can be hosted in the storage	96 (cluster can be scaled to 40+ nodes)

3.6.3 Storage software

The following table summarizes the storage software.

HBA driver	Nutanix Virt I/O SCSI Pass-thru Driver 62.62.101.5800
HBA QueueTarget Setting	N/A
HBA QueueDepth Setting	N/A
Hypervisor	Nutanix Acropolis Hypervisor (AHV)
Exchange VM guest OS	Windows Server 2016
ESE.dll file version	15.00.0847.030
Replication solution name/version	N/A

3.6.4 Storage disk configuration (mailbox store disks)

The following table summarizes the storage disk configuration.

Disk type, speed and firmware revision	Per Node: 4x Intel 1.92TB 6Gbps SATA G3HS 2.5" SSD 20x 2 TB 7.2K 6Gbps NL SATA 2.5" G3HS 512e HDD
Raw capacity per disk (GB)	2048 GB (2 TB)
Number of physical disks in test	96 (16x 1.92 TB + 80x 2 TB)
Total raw storage capacity (GB)	195,297GB
Disk slice size (GB)	N/A
Number of slices or disks per LUN	N/A
Raid level	Nutanix Replication Factor 2 (RAID 1)
Total formatted capacity	78.33 TB
Storage capacity utilization	41%
Database capacity utilization	15.7%

3.7 Performance test results

This section provides a high-level summary of the results of executing the Microsoft ESRP storage test version 4.0 on the configuration of 4 Lenovo ThinkAgile HX7520 appliances as described in the previous section. ESRP storage test results include reliability, storage performance, and database backup/restore.

Note that the ESRP program is not designed to be a benchmarking program and 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 and customers should not quote the data directly for their pre-deployment verifications. It is recommended that a proof of concept is carried out to validate the storage design for a specific customer environment.

The results in this section were developed by Lenovo and reviewed by the Microsoft Exchange Product team.

3.7.1 Reliability

Several of the tests in the ESP test framework are used to check reliability and run for 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.

Executing this test on the Lenovo ThinkAgile HX7520 appliances showed:

- No errors reported in the saved event log file.
- No errors reported during the database and log checksum process.

For detailed log results see “24-Hour stress/reliability test results” on page 26.

3.7.2 Storage performance results

The primary storage performance test in the ESP test framework is designed to exercise the storage with a maximum sustainable Exchange I/O pattern for 2 hours. The purpose is to reveal how long it takes for the storage to respond to I/O operations under a load.

For detailed log results see “2-Hour performance test results” on page 28.

Individual server metrics

Table 1 shows the sum of I/O’s and the average latency across all storage groups on a per server basis.

Table 1: Individual Server Performance

	Node 1	**Node 2**	**Node 3**
Database I/O			
Database Disks Transfers/sec	1538.06	1493.00	1470.82
Database Disks Reads/sec	1001.24	1016.35	1001.24
Database Disks Write/sec	459.57	476.65	469.57
Average Database Disk Read Latency (ms)	7.26	6.80	6.97
Average Database Disk Write Latency (ms)	4.61	4.83	4.66
Transaction Log I/O			
Log Disks Writes/sec	11.37	11.62	11.39
Average Log Disk Write Latency(ms)	0.99	1.01	0.99

Aggregate performance metrics across all servers

Table 2 shows the sum of I/O's and the average latency across the 3 primary servers in the solution.

Table 2: Aggregate Server Performance

Database I/O	
Database Disks Transfers/sec	4501.88
Database Disks Reads/sec	3018.13
Database Disks Writes/sec	1405.79
Average Database Disk Read Latency (ms)	7.01
Average Database Disk Write Latency (ms)	4.7
Transaction Log I/O	
Log Disks Writes/sec	34.38
Average Log Disk Write Latency (ms)	1.00

3.7.3 Database backup/recovery performance

Several of the tests in the ESP test framework are used to measure the sequential read rate of the database files and the recovery/replay performance (playing transaction logs into the database).

The database read-only performance test measures the maximum rate at which databases could be backed up using Microsoft Volume Shadow Copy Service (VSS). Table 3 shows the average read performance for a backing up a single database file and all ten database files on a single node.

Table 3: Database backup read-only performance results

MB read/sec per database	70.32
MB read/sec total per node (10 databases)	703.16

For detailed log results see "Database backup test results" on page 30.

3.7.4 Transaction log recovery/replay performance

The test is to measure the maximum rate at which the log files can be played against the databases. Table 4 shows the average rate for 500 log files played in a single storage group. Each log file is 1 MB in size.

Table 4: Transaction Log Recovery/Replay Performance

Average time to play one Log file (sec)	0.41

For detailed log results see "Soft recovery test results" on page 32.

4 Conclusion

This document is developed by storage solution providers, and reviewed by Microsoft Exchange Product team. The test results/data presented in this document is based on the tests introduced in the ESRP test framework. Customer should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

ESRP program is not designed to be a benchmarking program; tests are not designed to getting the maximum throughput for a giving solution. Rather, it is focused on producing recommendations from vendors for Exchange application. So the data presented in this document should not be used for direct comparisons among the solutions.

The cluster of 4 Lenovo ThinkAgile HX7520 appliances proved more than capable of handling the high IOPs generated by 30,000 mailboxes. Part of the reason for this is because each Lenovo ThinkAgile HX7520 appliance uses 3 HBAs for the 24 drives.

For best performance, it is recommended to use the Nutanix Windows drivers instead of the default Windows drivers for iSCSI. See download.nutanix.com/mobility/1.1.1/Nutanix-VirtIO-1.1.1.msi.

Appendix A – Test results

This section provides test results from 1 of the 3 primary mailbox servers under test: VM1. All server's test results are comparable to each another.

24-Hour stress/reliability test results

Checksum Statistics - All

Database	Seen pages	Bad pages	Correctable pages	Wrong page-number pages	File length / seconds taken
D:\Jetstress001001.edb	32220928	0	0	0	1006904 MB/17202 sec
E:\Jetstress002001.edb	32221184	0	0	0	1006912 MB/13510 sec
F:\Jetstress003001.edb	32221696	0	0	0	1006928 MB/12235 sec
G:\Jetstress004001.edb	32220672	0	0	0	1006896 MB/12280 sec
H:\Jetstress005001.edb	32220160	0	0	0	1006880 MB/12275 sec
I:\Jetstress006001.edb	32220416	0	0	0	1006888 MB/12224 sec
J:\Jetstress007001.edb	32221184	0	0	0	1006912 MB/12274 sec
K:\Jetstress008001.edb	32221440	0	0	0	1006920 MB/12250 sec
L:\Jetstress009001.edb	32220672	0	0	0	1006896 MB/12219 sec
M:\Jetstress010001.edb	32221184	0	0	0	1006912 MB/12252 sec
(Sum)	322209536	0	0	0	10069048 MB/17202 sec

Disk Subsystem Performance (of checksum)

LogicalDisk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Read
D:	0.037	0.000	937.460	0.000	65535.979
E:	0.033	0.000	1193.071	0.000	65536.000
F:	0.028	0.000	1315.782	0.000	65536.000
G:	0.028	0.000	1312.761	0.000	65536.000
H:	0.028	0.000	1308.407	0.000	65536.000
I:	0.028	0.000	1318.555	0.000	65536.000
J:	0.028	0.000	1308.860	0.000	65536.000
K:	0.028	0.000	1316.282	0.000	65536.000
L:	0.028	0.000	1319.770	0.000	65536.000
M:	0.028	0.000	1315.733	0.000	65536.000

Memory System Performance (of checksum)

Counter	Average	Minimum	Maximum
% Processor Time	2.253	0.407	4.475
Available MBytes	516703.834	516690.000	516723.000
Free System Page Table Entries	12295294.415	12294870.000	12296023.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	142419363.965	141787136.000	142938112.000
Pool Paged Bytes	228281081.298	228257792.000	228347904.000

Test Log

```
11/28/2017 2:23:38 PM -- Preparing for testing ...
11/28/2017 2:23:48 PM -- Attaching databases ...
11/28/2017 2:23:48 PM -- Preparations for testing are complete.
11/28/2017 2:23:48 PM -- Starting transaction dispatch ..
11/28/2017 2:23:48 PM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
11/28/2017 2:23:48 PM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
11/28/2017 2:23:59 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
11/28/2017 2:23:59 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
11/28/2017 2:24:00 PM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%
11/28/2017 2:24:00 PM -- Performance logging started (interval: 15000 ms).
11/28/2017 2:24:00 PM -- Attaining prerequisites:
11/28/2017 2:26:39 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2416669000.0 (lower bound: 2415919000.0, upper bound: none)
11/29/2017 2:26:40 PM -- Performance logging has ended.
11/29/2017 2:26:40 PM -- JetInterop batch transaction stats: 283260, 283260, 283260, 283259, 283259, 283259, 283259, 283259, 283259 and 283259.
11/29/2017 2:26:40 PM -- Dispatching transactions ends.
11/29/2017 2:26:40 PM -- Shutting down databases ...
11/29/2017 2:26:43 PM -- Instance744.1 (complete), Instance744.2 (complete), Instance744.3 (complete), Instance744.4 (complete), Instance744.5 (complete), Instance744.6
(complete), Instance744.7 (complete), Instance744.8 (complete), Instance744.9 (complete) and Instance744.10 (complete)
11/29/2017 2:26:43 PM -- C:\Program Files\Exchange Jetstress\24HR Test R1\Stress_2017_11_28_14_23_59.blg has 5758 samples.
11/29/2017 2:26:43 PM -- Creating test report ...
11/29/2017 2:27:35 PM -- Instance744.1 has 10.9 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.1 has 2.2 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.1 has 2.2 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.2 has 7.4 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.2 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.2 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.3 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.3 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.3 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.4 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.4 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.4 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.5 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.5 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.5 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.6 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.6 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.6 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.7 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.7 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.7 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.8 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.8 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.8 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.9 has 7.5 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.9 has 2.0 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.9 has 2.0 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.10 has 7.4 for I/O Database Reads Average Latency.
11/29/2017 2:27:35 PM -- Instance744.10 has 1.9 for I/O Log Writes Average Latency.
11/29/2017 2:27:35 PM -- Instance744.10 has 1.9 for I/O Log Reads Average Latency.
11/29/2017 2:27:35 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
11/29/2017 2:27:35 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
11/29/2017 2:27:35 PM -- C:\Program Files\Exchange Jetstress\24HR Test R1\Stress_2017_11_28_14_23_59.xml has 5747 samples queried.
```

2-Hour performance test results

Test Summary

Overall Test Result	Pass
Machine Name	EX1
Test Description	
Test Start Time	11/27/2017 10:27:41 AM
Test End Time	11/27/2017 12:30:01 PM
Collection Start Time	11/27/2017 10:29:58 AM
Collection End Time	11/27/2017 12:29:49 PM
Jetstress Version	15.01.1019.000
ESE Version	15.00.0847.030
Operating System	Windows Server 2016 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\2HR Test R1\Performance_2017_11_27_10_28_2.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	1538.058
Target Transactional I/O per Second	600
Initial Database Size (bytes)	10508769951744
Final Database Size (bytes)	10512712597504
Database Files (Count)	10

Jetstress System Parameters

Thread Count	16
Minimum Database Cache	320.0 MB
Maximum Database Cache	2560.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance3192.1	Log path: D:\Logs Database: D:\Jetstress001001.edb
Instance3192.2	Log path: E:\Logs Database: E:\Jetstress002001.edb
Instance3192.3	Log path: F:\Logs Database: F:\Jetstress003001.edb
Instance3192.4	Log path: G:\Logs Database: G:\Jetstress004001.edb
Instance3192.5	Log path: H:\Logs Database: H:\Jetstress005001.edb
Instance3192.6	Log path: I:\Logs Database: I:\Jetstress006001.edb
Instance3192.7	Log path: J:\Logs Database: J:\Jetstress007001.edb
Instance3192.8	Log path: K:\Logs Database: K:\Jetstress008001.edb
Instance3192.9	Log path: L:\Logs Database: L:\Jetstress009001.edb
Instance3192.10	Log path: M:\Logs Database: M:\Jetstress010001.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
Instance3192.1	10.017	3.009	104.435	49.355	33383.061	34940.825	0.000	1.712	0.000	11.781	0.000	20319.482
Instance3192.2	6.651	3.107	104.389	49.104	33394.476	34926.588	0.000	1.536	0.000	11.732	0.000	20505.499
Instance3192.3	6.609	3.158	104.496	49.472	33427.742	34948.016	0.000	1.520	0.000	11.848	0.000	20428.248
Instance3192.4	6.635	3.310	104.617	49.228	33398.621	34942.324	0.000	1.441	0.000	11.745	0.000	20327.622
Instance3192.5	6.664	3.438	104.015	48.645	33425.576	34987.862	0.000	1.491	0.000	11.765	0.000	20362.181
Instance3192.6	6.597	3.593	104.603	49.499	33390.924	34959.246	0.000	1.468	0.000	11.877	0.000	20266.893
Instance3192.7	6.597	3.955	104.130	49.205	33404.345	34944.162	0.000	1.555	0.000	11.809	0.000	20473.626
Instance3192.8	6.647	3.903	104.808	49.972	33436.308	34920.587	0.000	1.454	0.000	11.924	0.000	20249.548
Instance3192.9	6.608	4.047	104.307	49.047	33400.884	34929.324	0.000	1.460	0.000	11.763	0.000	20365.405
Instance3192.10	6.590	4.348	104.836	49.895	33405.630	34917.371	0.000	1.480	0.000	11.983	0.000	20195.532

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3192.1	9.068	261708.514
Instance3192.2	9.074	261697.790
Instance3192.3	9.079	261668.254
Instance3192.4	9.081	261690.634
Instance3192.5	9.081	261640.998
Instance3192.6	9.073	261793.272
Instance3192.7	9.086	261790.801
Instance3192.8	9.071	261767.807
Instance3192.9	9.076	261764.549
Instance3192.10	9.079	261654.597

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3192.1	1.025	231574.182
Instance3192.2	1.027	232558.405
Instance3192.3	1.031	232075.247
Instance3192.4	1.021	232542.097
Instance3192.5	1.025	232559.148
Instance3192.6	1.031	232561.778
Instance3192.7	1.033	232561.778
Instance3192.8	1.031	232074.452
Instance3192.9	1.025	231591.989
Instance3192.10	1.035	232075.247

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
Instance3192.1	10.017	3.009	113.503	49.355	51623.986	34940.825	1.678	1.712	1.025	11.781	231574.182	20319.482
Instance3192.2	6.651	3.107	113.463	49.104	51653.173	34926.588	1.704	1.536	1.027	11.732	232558.405	20505.499
Instance3192.3	6.609	3.158	113.575	49.472	51672.978	34948.016	1.690	1.520	1.031	11.848	232075.247	20428.248
Instance3192.4	6.635	3.310	113.698	49.228	51632.927	34942.324	1.709	1.441	1.021	11.745	232542.097	20327.622
Instance3192.5	6.664	3.438	113.096	48.645	51749.528	34987.862	1.667	1.491	1.025	11.765	232559.148	20362.181
Instance3192.6	6.597	3.593	113.676	49.499	51621.094	34959.246	1.680	1.468	1.031	11.877	232561.778	20266.893
Instance3192.7	6.597	3.955	113.216	49.205	51732.382	34944.162	1.692	1.555	1.033	11.809	232561.778	20473.626
Instance3192.8	6.647	3.903	113.879	49.972	51623.871	34920.587	1.711	1.454	1.031	11.924	232074.452	20249.548
Instance3192.9	6.608	4.047	113.383	49.047	51681.015	34929.324	1.695	1.460	1.025	11.763	231591.989	20365.405
Instance3192.10	6.590	4.348	113.915	49.895	51596.770	34917.371	1.714	1.480	1.035	11.983	232075.247	20195.532

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.740	0.325	1.422
Available MBytes	514113.804	514084.000	514226.000
Free System Page Table Entries	12295359.422	12294747.000	12295693.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	139806792.685	139632640.000	139984896.000
Pool Paged Bytes	220286488.585	220073984.000	220426240.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```
11/27/2017 10:27:41 AM-- Preparing for testing ...
11/27/2017 10:27:52 AM-- Attaching databases ...
11/27/2017 10:27:52 AM-- Preparations for testing are complete.
11/27/2017 10:27:52 AM-- Starting transaction dispatch ..
11/27/2017 10:27:52 AM-- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
11/27/2017 10:27:52 AM-- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
11/27/2017 10:28:02 AM-- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
11/27/2017 10:28:02 AM-- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
11/27/2017 10:28:03 AM-- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%
11/27/2017 10:28:03 AM-- Performance logging started (interval: 15000 ms).
11/27/2017 10:28:03 AM-- Attaining prerequisites:
11/27/2017 10:29:58 AM-- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2430362000.0 (lower bound: 2415919000.0, upper bound: none)
11/27/2017 12:29:59 PM-- Performance logging has ended.
11/27/2017 12:29:59 PM-- JetInterop batch transaction stats: 28141, 28141, 28141, 28141, 28141, 28141, 28141, 28141, 28141 and 28140.
11/27/2017 12:29:59 PM-- Dispatching transactions ends.
11/27/2017 12:29:59 PM-- Shutting down databases ...
11/27/2017 12:30:01 PM-- Instance3192.1 (complete), Instance3192.2 (complete), Instance3192.3 (complete), Instance3192.4 (complete), Instance3192.5 (complete), Instance3192.6
(complete), Instance3192.7 (complete), Instance3192.8 (complete), Instance3192.9 (complete) and Instance3192.10 (complete)
11/27/2017 12:30:01 PM-- C:\Program Files\Exchange Jetstress\2HR Test R1\Performance_2017_11_27_10_28_2.blg has 486 samples.
11/27/2017 12:30:01 PM-- Creating test report ...
11/27/2017 12:30:12 PM-- Instance3192.1 has 10.0 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.1 has 1.7 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.1 has 1.7 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.2 has 6.7 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.2 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.2 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.3 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.3 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.3 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.4 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.4 has 1.4 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.4 has 1.4 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.5 has 6.7 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.5 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.5 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.6 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.6 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.6 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.7 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.7 has 1.6 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.7 has 1.6 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.8 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.8 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.8 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.9 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.9 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.9 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.10 has 6.6 for I/O Database Reads Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.10 has 1.5 for I/O Log Writes Average Latency.
11/27/2017 12:30:12 PM-- Instance3192.10 has 1.5 for I/O Log Reads Average Latency.
11/27/2017 12:30:12 PM-- Test has 0 Maximum Database Page Fault Stalls/sec.
11/27/2017 12:30:12 PM-- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
11/27/2017 12:30:12 PM-- C:\Program Files\Exchange Jetstress\2HR Test R1\Performance_2017_11_27_10_28_2.xml has 478 samples queried.
```

Database backup test results

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance5572.1	1002568.03	05:17:59	52.55
Instance5572.2	1002576.03	04:00:35	69.45
Instance5572.3	1002568.03	03:49:01	72.96
Instance5572.4	1002560.03	03:49:43	72.74
Instance5572.5	1002544.03	03:52:44	71.79
Instance5572.6	1002552.03	03:49:04	72.94
Instance5572.7	1002568.03	03:52:52	71.75
Instance5572.8	1002576.03	03:48:56	72.99
Instance5572.9	1002560.03	03:48:27	73.14
Instance5572.10	1002552.03	03:49:21	72.85
Avg			70.32
Sum			703.16

Jetstress System Parameters

Thread Count	16
Minimum Database Cache	320.0 MB
Maximum Database Cache	2560.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance5572.1 Log path: D:\Logs
Database: D:\Jetstress001001.edb

Instance5572.2 Log path: E:\Logs
Database: E:\Jetstress002001.edb

Instance5572.3 Log path: F:\Logs
Database: F:\Jetstress003001.edb

Instance5572.4 Log path: G:\Logs
Database: G:\Jetstress004001.edb

Instance5572.5 Log path: H:\Logs
Database: H:\Jetstress005001.edb

Instance5572.6 Log path: I:\Logs
Database: I:\Jetstress006001.edb

Instance5572.7 Log path: J:\Logs
Database: J:\Jetstress007001.edb

Instance5572.8 Log path: K:\Logs
Database: K:\Jetstress008001.edb

Instance5572.9 Log path: L:\Logs
Database: L:\Jetstress009001.edb

Instance5572.10 Log path: M:\Logs
Database: M:\Jetstress010001.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
Instance5572.1	5.637	0.000	209.912	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.2	4.489	0.000	276.949	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.3	5.487	0.000	291.350	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.4	5.496	0.000	290.938	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.5	4.808	0.000	287.065	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.6	5.493	0.000	291.229	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.7	4.811	0.000	286.616	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.8	5.486	0.000	291.632	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.9	5.484	0.000	292.328	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance5572.10	5.472	0.000	291.112	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	1.814	0.804	4.860
Available MBytes	516791.431	516780.000	516796.000
Free System Page Table Entries	12295853.213	12295268.000	12300633.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	140234287.572	140136448.000	140361728.000
Pool Paged Bytes	220894022.551	220852224.000	220987392.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

11/27/2017 6:28:45 PM -- Preparing for testing ...
11/27/2017 6:28:56 PM -- Attaching databases ...
11/27/2017 6:28:56 PM -- Preparations for testing are complete.
11/27/2017 6:29:07 PM -- Performance logging started (interval: 30000 ms).
11/27/2017 6:29:07 PM -- Backing up databases ...
11/27/2017 11:47:06 PM -- Performance logging has ended.
11/27/2017 11:47:06 PM -- Instance5572.1 (100% processed), Instance5572.2 (100% processed), Instance5572.3 (100% processed), Instance5572.4 (100% processed), Instance5572.5 (100% processed), Instance5572.6 (100% processed), Instance5572.7 (100% processed), Instance5572.8 (100% processed), Instance5572.9 (100% processed) and Instance5572.10 (100% processed)
11/27/2017 11:47:06 PM -- C:\Program Files\Exchange Jetstress\DB Backup R1\DatabaseBackup_2017_11_27_18_28_56.blg has 635 samples.
11/27/2017 11:47:06 PM -- Creating test report ...

Soft recovery test results

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance1640.1	506	229.4692523
Instance1640.2	501	194.6868829
Instance1640.3	504	205.4692788
Instance1640.4	504	199.2806169
Instance1640.5	504	203.859903
Instance1640.6	510	234.0473634
Instance1640.7	506	209.2505271
Instance1640.8	506	206.547416
Instance1640.9	509	203.3130329
Instance1640.10	513	203.5942824
Avg	506	208.952
Sum	5063	2089.5185557

Database Configuration

Instance1640.1 Log path: D:\Logs
Database: D:\Jetstress001001.edb

Instance1640.2 Log path: E:\Logs
Database: E:\Jetstress002001.edb

Instance1640.3 Log path: F:\Logs
Database: F:\Jetstress003001.edb

Instance1640.4 Log path: G:\Logs
Database: G:\Jetstress004001.edb

Instance1640.5 Log path: H:\Logs
Database: H:\Jetstress005001.edb

Instance1640.6 Log path: I:\Logs
Database: I:\Jetstress006001.edb

Instance1640.7 Log path: J:\Logs
Database: J:\Jetstress007001.edb

Instance1640.8 Log path: K:\Logs
Database: K:\Jetstress008001.edb

Instance1640.9 Log path: L:\Logs
Database: L:\Jetstress009001.edb

Instance1640.10 Log path: M:\Logs
Database: M:\Jetstress010001.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
Instance1640.1	7.119	1.294	1462.292	8.957	39711.141	32182.857	3.098	0.000	11.197	0.000	205970.286	0.000
Instance1640.2	5.115	1.355	1725.869	10.374	39732.372	32768.000	1.829	0.000	12.967	0.000	209700.475	0.000
Instance1640.3	5.039	1.184	1673.160	9.961	39726.295	32112.640	1.895	0.000	12.463	0.000	205520.896	0.000
Instance1640.4	4.829	1.363	1666.700	10.179	39824.763	32768.000	1.909	0.000	12.723	0.000	209725.834	0.000
Instance1640.5	5.031	1.502	1646.303	9.951	39834.592	32768.000	1.771	0.000	12.439	0.000	209720.574	0.000
Instance1640.6	4.506	1.420	1384.515	8.767	39695.805	32193.123	1.800	0.000	10.959	0.000	206035.986	0.000
Instance1640.7	5.036	1.814	1659.306	9.834	39724.902	32125.490	1.855	0.000	12.293	0.000	205603.137	0.000
Instance1640.8	4.859	1.419	1654.900	9.912	39832.521	32768.000	1.797	0.000	12.390	0.000	209725.348	0.000
Instance1640.9	4.956	1.505	1649.979	10.073	39614.564	32768.000	1.912	0.000	12.591	0.000	209714.652	0.000
Instance1640.10	4.867	2.411	1649.649	10.134	39837.803	32768.000	1.829	0.000	12.668	0.000	209715.200	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1640.1	0.000	0.000
Instance1640.2	0.000	0.000
Instance1640.3	0.000	0.000
Instance1640.4	0.000	0.000
Instance1640.5	0.000	0.000
Instance1640.6	0.000	0.000
Instance1640.7	0.000	0.000
Instance1640.8	0.000	0.000
Instance1640.9	0.000	0.000
Instance1640.10	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
Instance1640.1	7.119	1.294	1462.292	8.957	39711.141	32182.857	3.098	0.000	11.197	0.000	205970.286	0.000
Instance1640.2	5.115	1.355	1725.869	10.374	39732.372	32768.000	1.829	0.000	12.967	0.000	209700.475	0.000
Instance1640.3	5.039	1.184	1673.160	9.961	39726.295	32112.640	1.895	0.000	12.463	0.000	205520.896	0.000
Instance1640.4	4.829	1.363	1666.700	10.179	39824.763	32768.000	1.909	0.000	12.723	0.000	209725.834	0.000
Instance1640.5	5.031	1.502	1646.303	9.951	39834.592	32768.000	1.771	0.000	12.439	0.000	209720.574	0.000
Instance1640.6	4.506	1.420	1384.515	8.767	39695.805	32193.123	1.800	0.000	10.959	0.000	206035.986	0.000
Instance1640.7	5.036	1.814	1659.306	9.834	39724.902	32125.490	1.855	0.000	12.293	0.000	205603.137	0.000
Instance1640.8	4.859	1.419	1654.900	9.912	39832.521	32768.000	1.797	0.000	12.390	0.000	209725.348	0.000
Instance1640.9	4.956	1.505	1649.979	10.073	39614.564	32768.000	1.912	0.000	12.591	0.000	209714.652	0.000
Instance1640.10	4.867	2.411	1649.649	10.134	39837.803	32768.000	1.829	0.000	12.668	0.000	209715.200	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	4.050	0.649	7.558
Available MBytes	514499.190	514042.000	516656.000
Free System Page Table Entries	12295633.276	12295332.000	12295767.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	141676332.138	141512704.000	141881344.000
Pool Paged Bytes	221901647.448	221827072.000	221966336.000
Database Page Fault Stalls/sec	0.048	0.000	2.739

Test Log

```
8/12/2016 3:35:45 AM -- Preparing for testing ...
8/12/2016 3:35:55 AM -- Attaching databases ...
8/12/2016 3:35:55 AM -- Preparations for testing are complete.
8/12/2016 3:35:55 AM -- Starting transaction dispatch ..
8/12/2016 3:35:55 AM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
8/12/2016 3:35:55 AM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
8/12/2016 3:36:06 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
8/12/2016 3:36:06 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
8/12/2016 3:36:07 AM -- Operation mix: Sessions 22, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/12/2016 3:36:07 AM -- Performance logging started (interval: 15000 ms).
8/12/2016 3:36:07 AM -- Generating log files ...
8/12/2016 5:10:02 AM -- F:\Logs (100.4% generated), G:\Logs (101.8% generated), H:\Logs (100.2% generated), I:\Logs (101.6% generated), J:\Logs (101.2% generated),
K:\Logs (100.6% generated), L:\Logs (100.8% generated), M:\Logs (100.2% generated), N:\Logs (101.2% generated) and O:\Logs (101.6% generated)
8/12/2016 5:10:02 AM -- Performance logging has ended.
8/12/2016 5:10:02 AM -- JetInterop batch transaction stats: 17120, 17120, 17120, 17120, 17120, 17120, 17119, 17119, 17119 and 17119.
8/12/2016 5:10:02 AM -- Dispatching transactions ends.
8/12/2016 5:10:02 AM -- Shutting down databases ...
8/12/2016 5:10:05 AM -- Instance3488.1 (complete), Instance3488.2 (complete), Instance3488.3 (complete), Instance3488.4 (complete), Instance3488.5 (complete),
Instance3488.6 (complete), Instance3488.7 (complete), Instance3488.8 (complete), Instance3488.9 (complete) and Instance3488.10 (complete)
8/12/2016 5:10:05 AM -- C:\Program Files\Exchange Jetstress\Soft_Recovery\Performance_2016_8_12_3_36_6.blg has 375 samples.
8/12/2016 5:10:05 AM -- Creating test report ...
8/12/2016 5:10:08 AM -- Instance3488.1 has 13.9 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.1 has 2.9 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.1 has 2.9 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.2 has 12.7 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.2 has 2.8 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.2 has 2.8 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.3 has 12.6 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.3 has 3.0 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.3 has 3.0 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.4 has 12.6 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.4 has 2.7 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.4 has 2.7 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.5 has 12.5 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.5 has 2.7 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.5 has 2.7 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.6 has 12.6 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.6 has 2.7 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.6 has 2.7 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.7 has 12.6 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.7 has 2.7 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.7 has 2.7 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.8 has 12.5 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.8 has 2.7 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.8 has 2.7 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.9 has 12.5 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.9 has 2.6 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.9 has 2.6 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.10 has 12.6 for I/O Database Reads Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.10 has 2.6 for I/O Log Writes Average Latency.
8/12/2016 5:10:08 AM -- Instance3488.10 has 2.6 for I/O Log Reads Average Latency.
8/12/2016 5:10:08 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/12/2016 5:10:08 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/12/2016 5:10:08 AM -- C:\Program Files\Exchange Jetstress\Soft_Recovery\Performance_2016_8_12_3_36_6.xml has 374 samples queried.
8/12/2016 5:10:08 AM -- C:\Program Files\Exchange Jetstress\Soft_Recovery\Performance_2016_8_12_3_36_6.html was saved.
8/12/2016 5:10:08 AM -- Performance logging started (interval: 4000 ms).
8/12/2016 5:10:08 AM -- Recovering databases ...
8/12/2016 5:35:08 AM -- Performance logging has ended.
```

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, ThinkSystem, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

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

Nutanix is a trademark of Nutanix, Inc. 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.