60,000 Mailbox Resiliency Solution for Microsoft Exchange 2016 using All Flash Lenovo ThinkAgile HX7520 Appliances and AHV

Last update: 29 January 2018

Provides a technical overview of Lenovo ThinkAgile HX all flash Solution for Microsoft Exchange 2016

Contains performance results for 60,000 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

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 60,000 mailboxes with Microsoft Exchange Server 2016 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 60,000 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

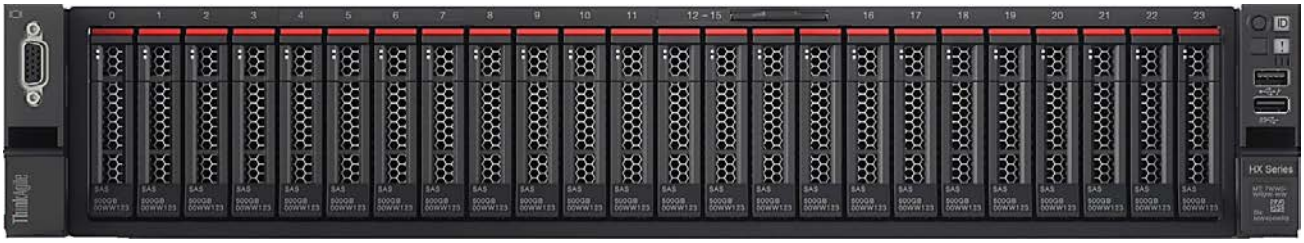

The table below provides a summary of the configuration options for the HX7520.

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	24 x 1.92 TB 0 x 3.84 TB	0 x 1TB 0 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.

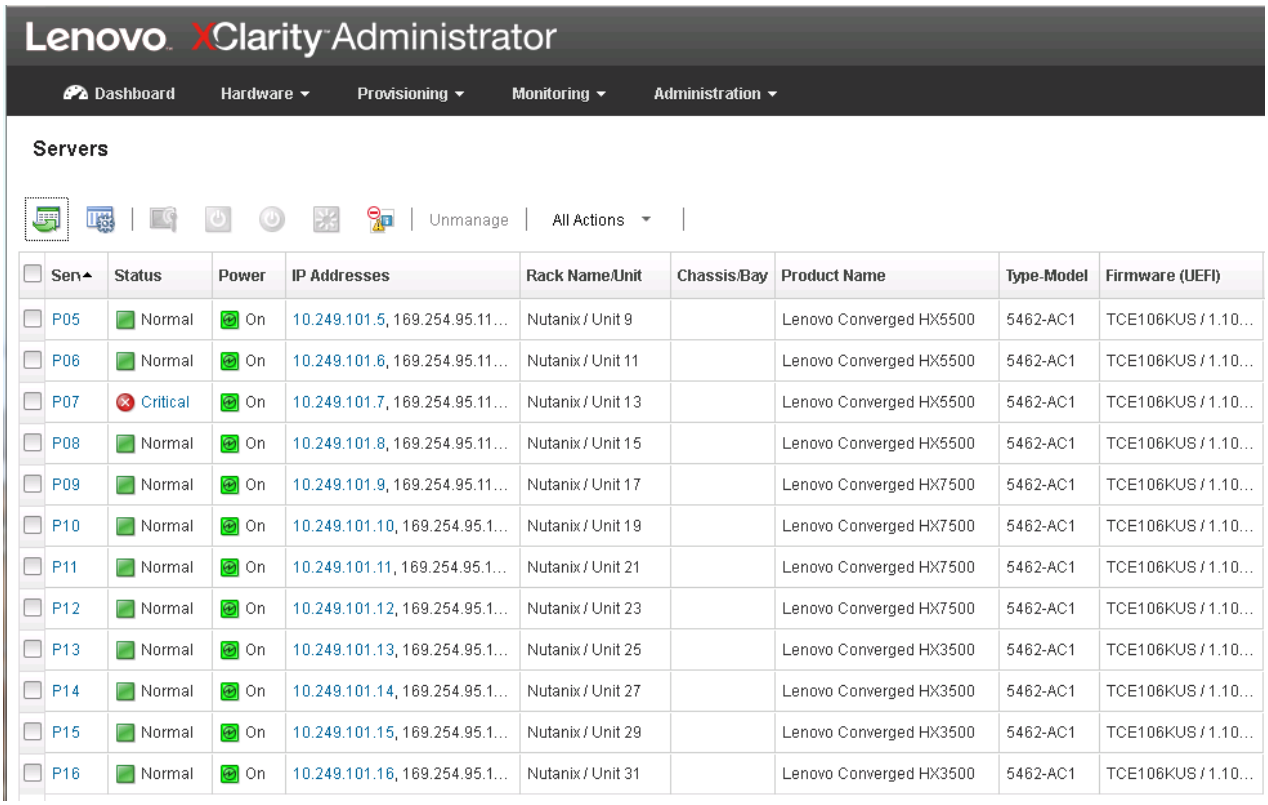

Ser	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product Name	Type-Model	Firmware (UEFI)
P05	Normal	On	10.249.101.5, 169.254.95.11...	Nutanix / Unit 9		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P06	Normal	On	10.249.101.6, 169.254.95.11...	Nutanix / Unit 11		Lenovo Converged HX5500	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.

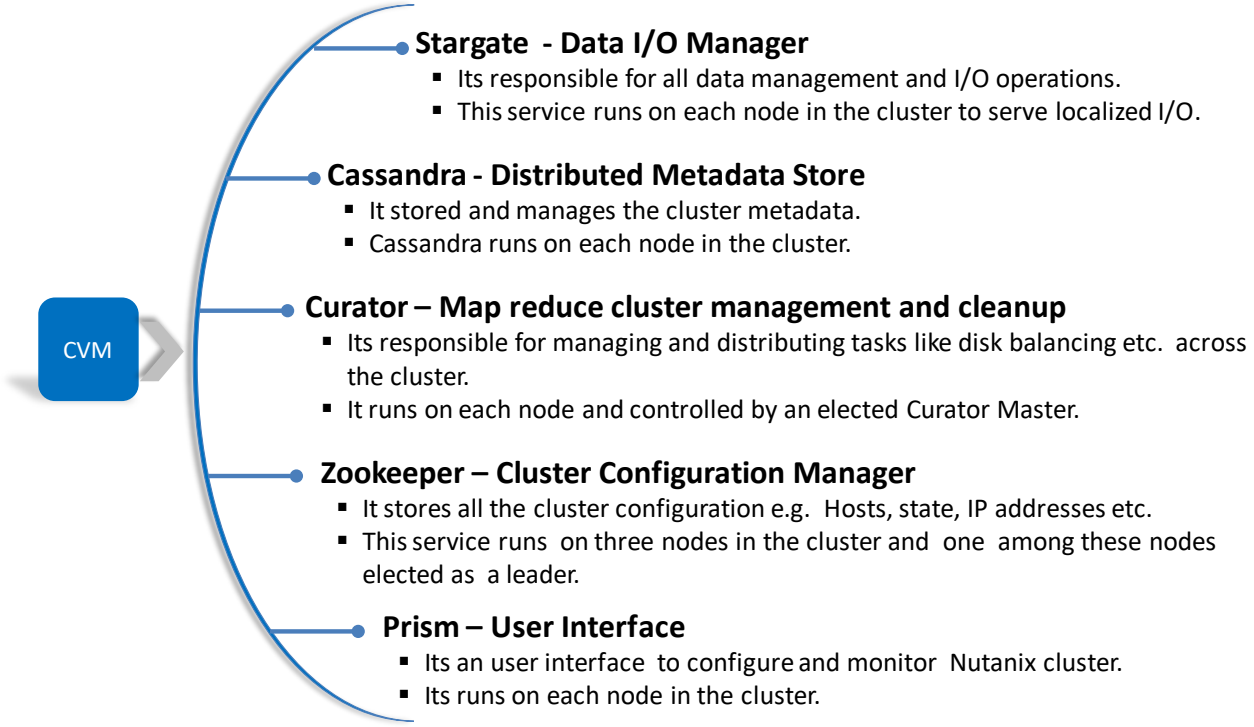

Figure 2: Controller VM components

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.

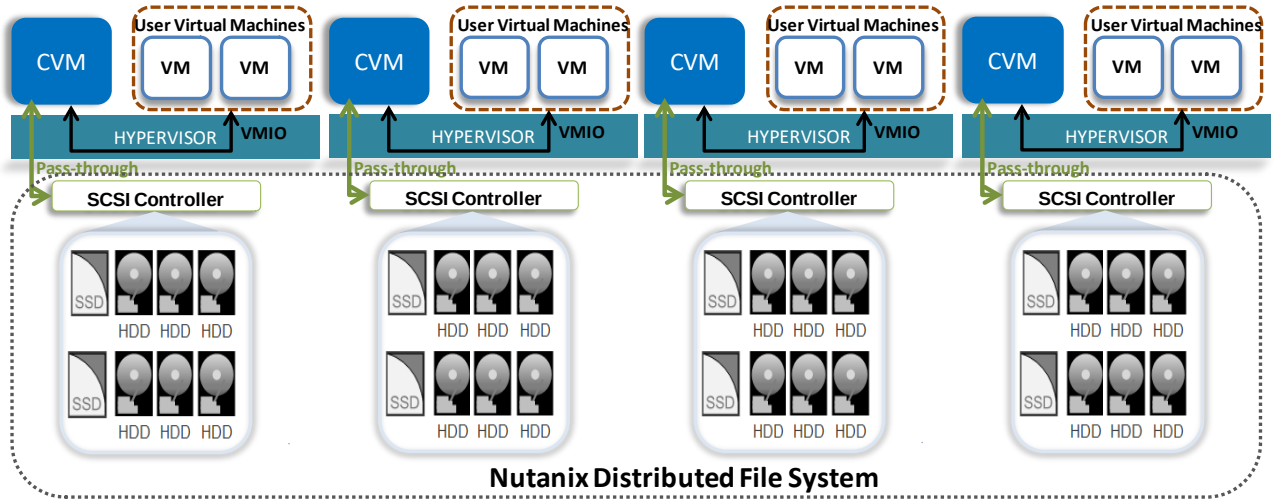

Figure 3: CVM interaction with Hypervisor and User VMs

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.

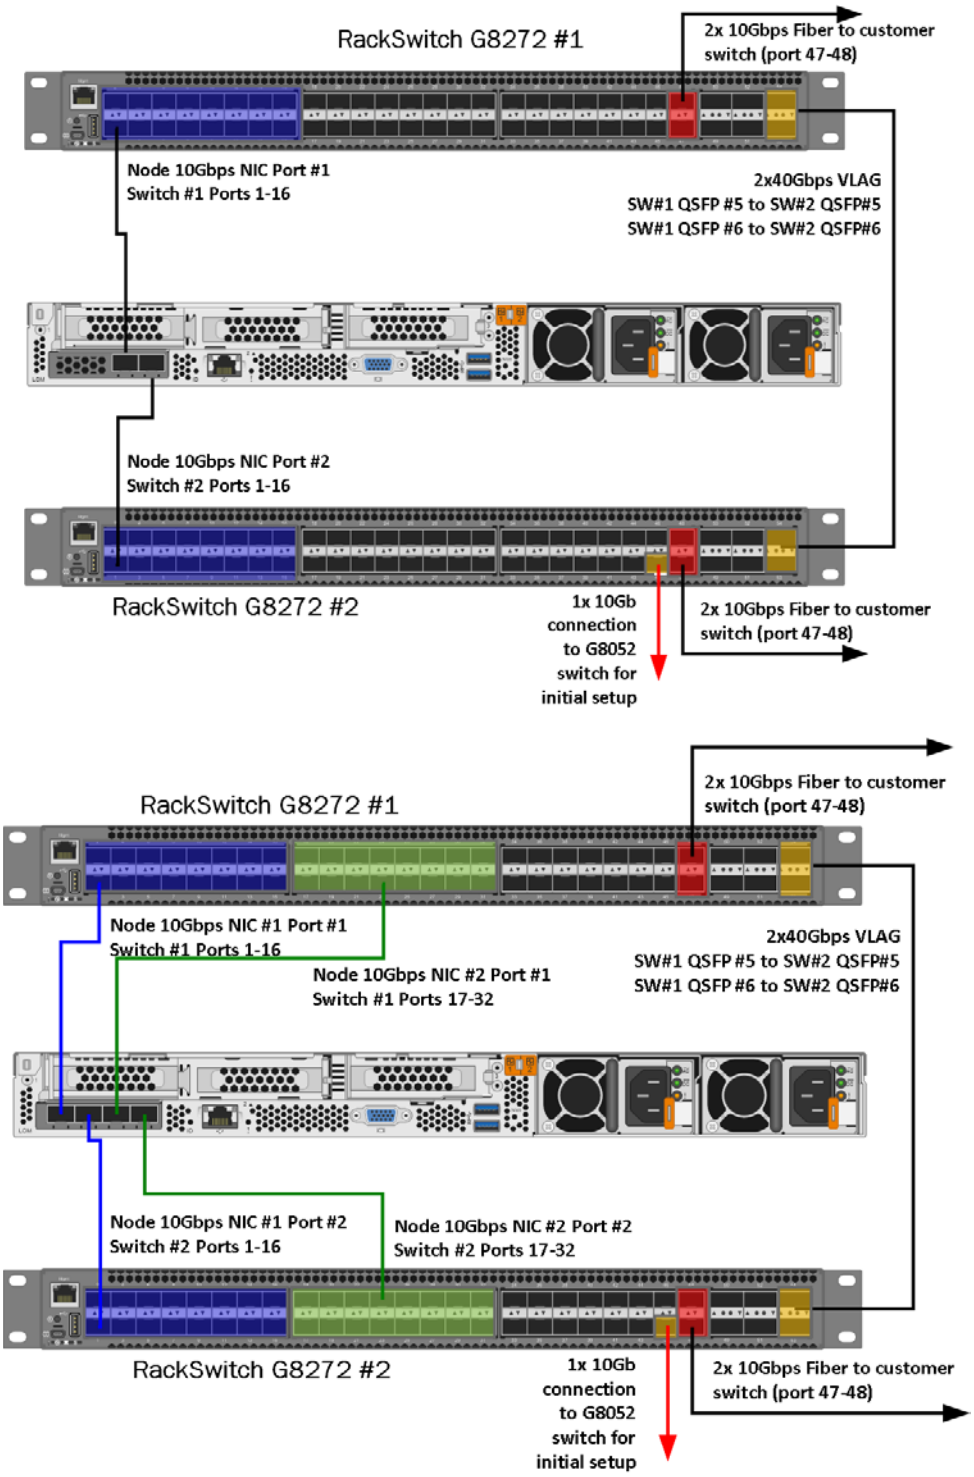

Figure 6: Data network with single and dual adapters

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.

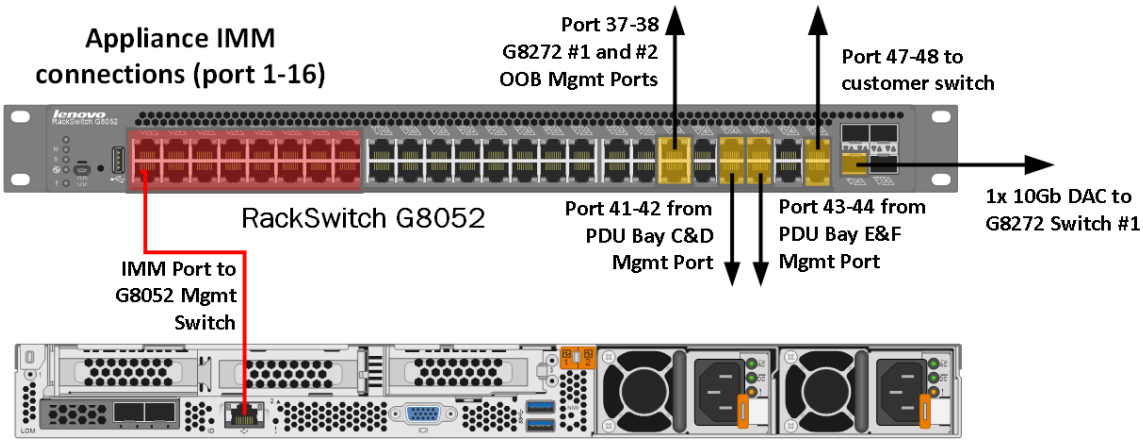

Figure 7: IMM 1GbE management network

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.

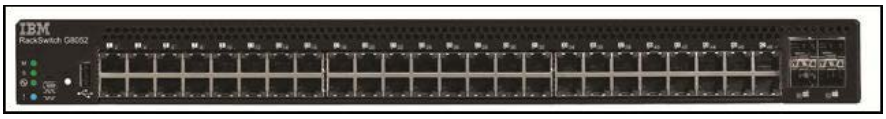

Figure 9: Lenovo RackSwitch G8052

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:

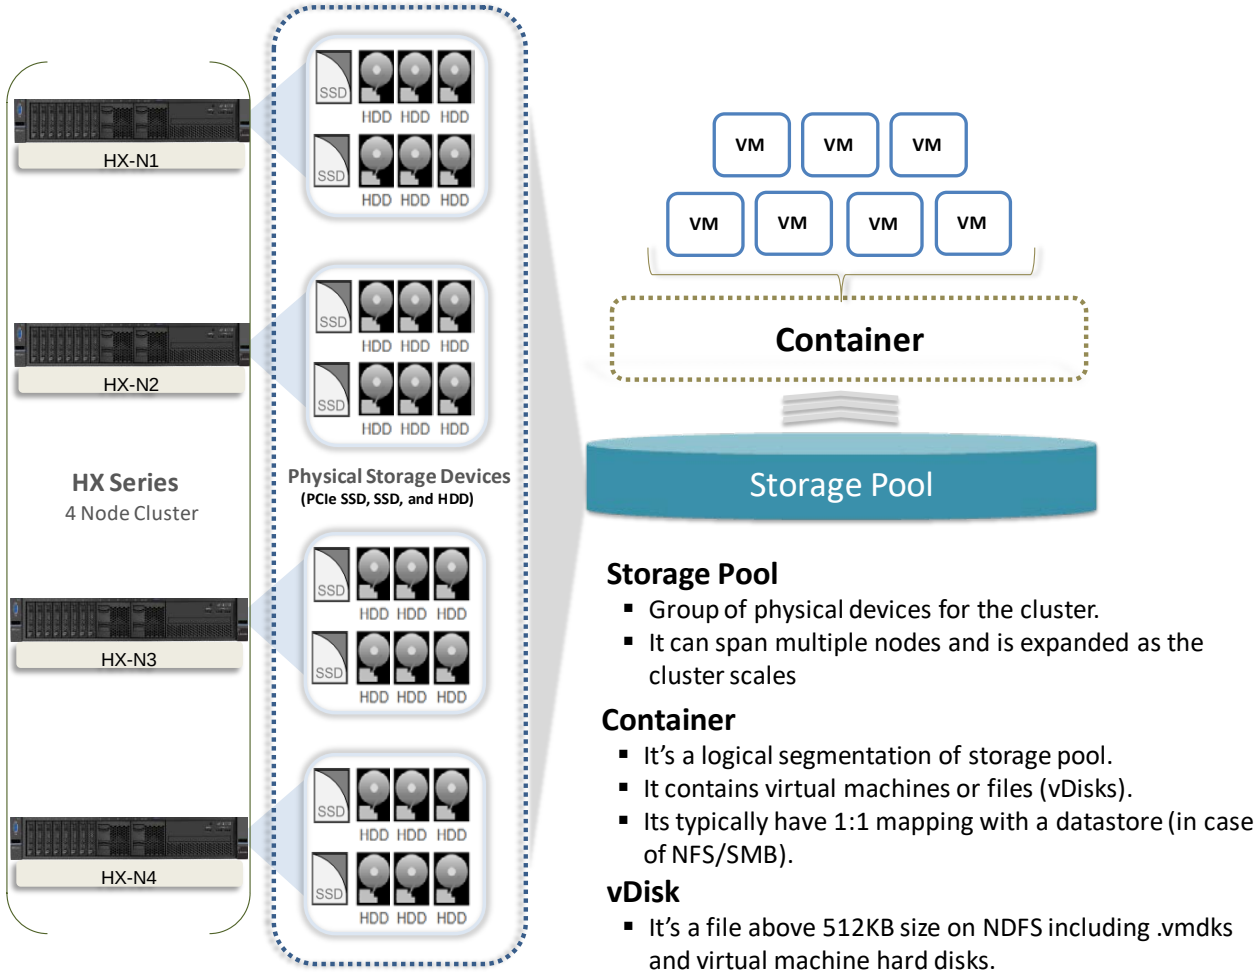

Figure 10: Nutanix Distributed File System

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.

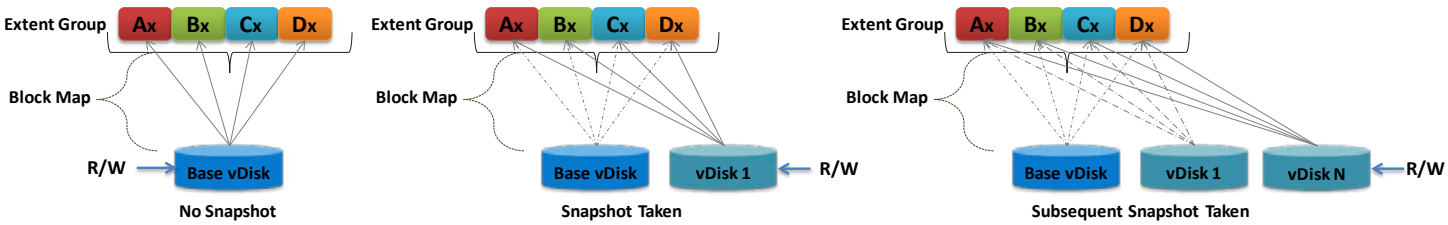

Figure 11: VM snapshots

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.

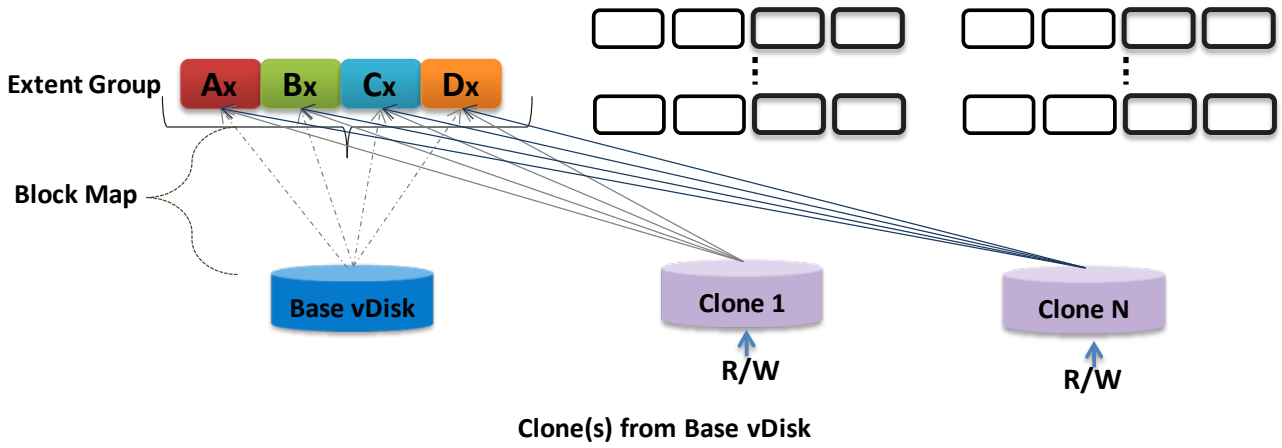

Figure 12: VM clones

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 2016 is the market leader in enterprise messaging and collaboration. With increasing processor performance, the primary design goal for Exchange 2016 is simplicity of scale, hardware utilization, and failure isolation. With Exchange 2016, the number of server roles is reduced to two: Edge Transport server and the Mailbox server which includes client access and mailbox services.

3.1 Solution overview

Figure 13 shows the architectural overview of the Microsoft Exchange 2016 solution using Lenovo ThinkAgile HX7520 appliances. This chapter does not address integrating Exchange with unified messaging solutions and handling edge transport routing and distribution.

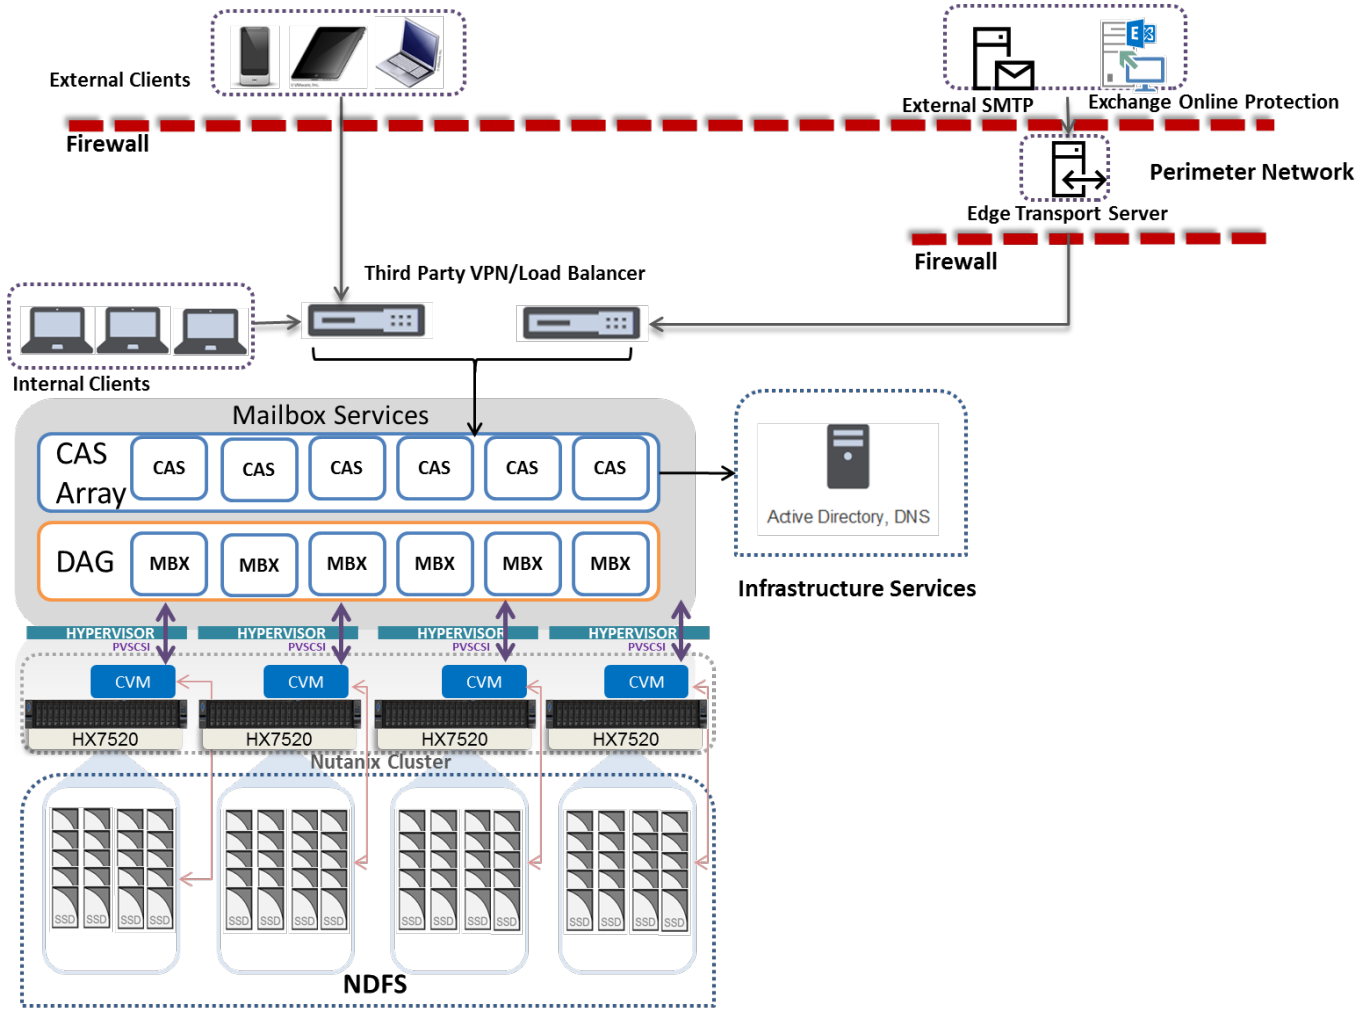

Figure 13. Lenovo ThinkAgile HX Series solution with Microsoft Exchange

The Client Access Server (CAS) role which is no longer a separate entity and is automatically installed with the Mailbox Server role still provides client protocols and unified messaging support. The Mailbox Server (MBX) role provides all of the data processing services. All external mail through SMTP is now handled via a separately installed Edge Transport Server, which usually resides within the perimeter network.

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 2016 environment. Figure 14 shows a high-level component model.

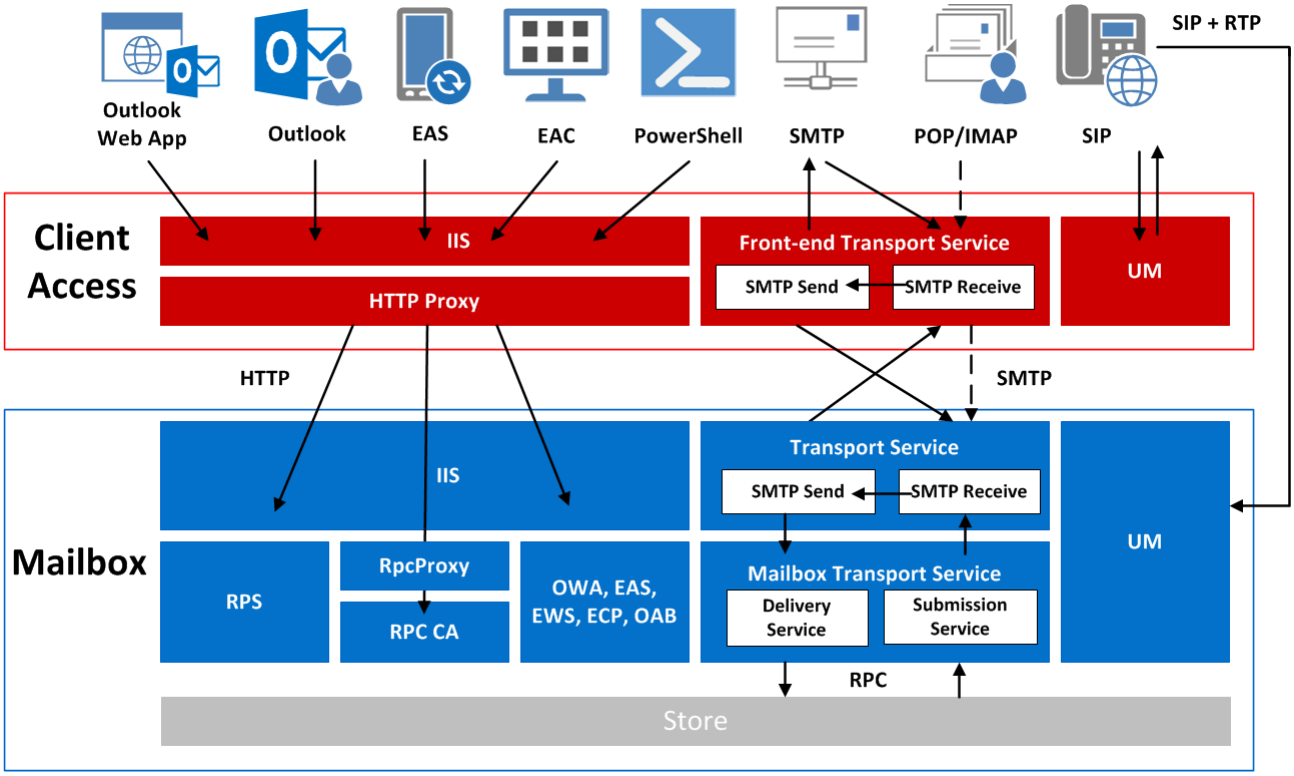

Figure 14. Exchange Server 2016 logical component view

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 2016 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 2016, 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 – In Exchange 2016 this feature has been replaced by **MAPI over HTTP** offering improvements over the traditional Outlook anywhere (RPC over HTTP). In Exchange 2016 the MAPI over HTTP feature is enabled by default because Exchange 2016 which 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 is 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 2016. 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 2016 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 2016 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: 24 x 1.92TB SATA SSDs – 46.08TB

Raw storage per cluster: SSD: 96 x 1.92 TB SATA SSDs – 184.32TB

Figure 15 describes the high level DAG architecture of the 60,000 mailbox virtualized Exchange 2016 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).

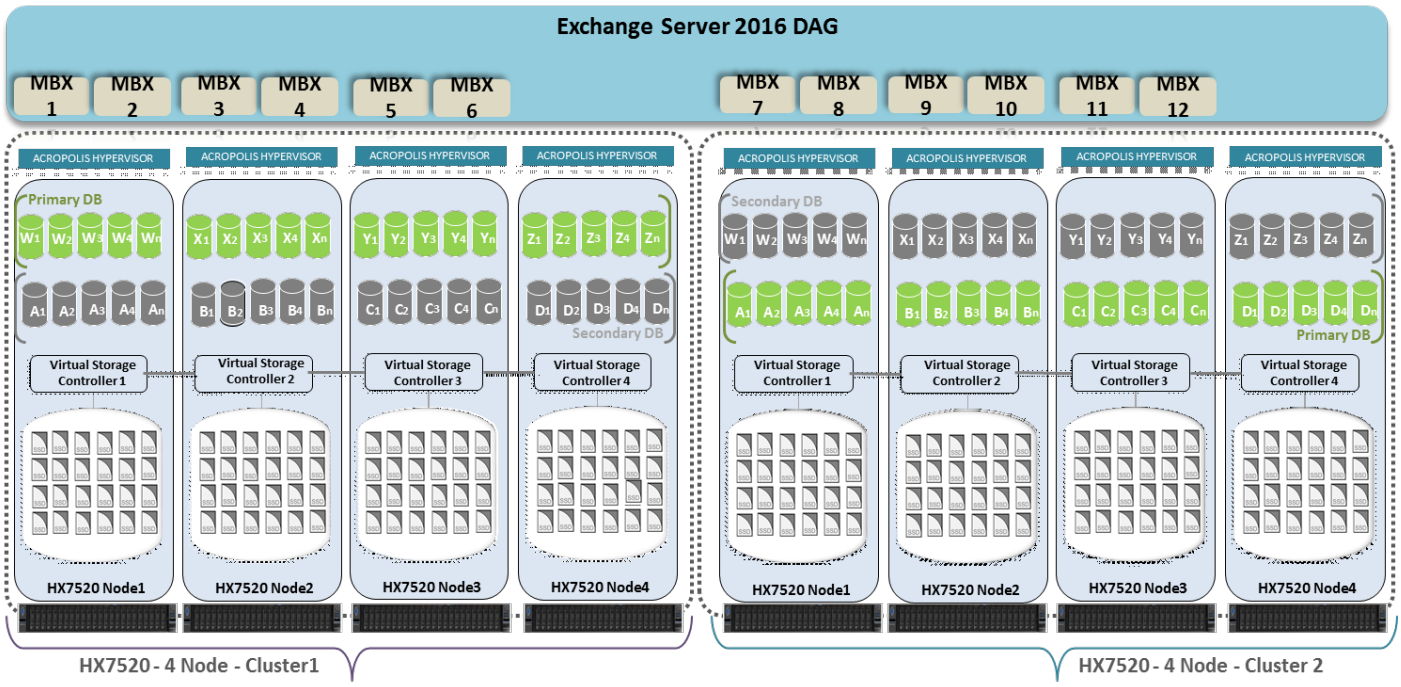

Figure 15. DAG Architecture using Exchange 2016 and Lenovo ThinkAgile HX7520 (60,000 Mailboxes)

The DAG has twelve Exchange 2016 mailbox servers and two database copies. The two database copies were placed on two physically isolated clusters. 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
- 96 GB RAM
- Windows Server 2016
- Microsoft Exchange 2016 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 2016 environment is as follows:

- 60,000 mailboxes of 0.75GB
- 12x Exchange 2016 servers (6x 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	60,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	12 (6 tested)
Number of active mailboxes/server	10,000 (10,000 active and 10,000 passive mailboxes per server and tested 20,000 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	800 GB
Total database size for performance testing	32 TB
% storage capacity used by Exchange database[2]	44.8%

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	N/A
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
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)

[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.

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: 24 x Intel 1.92TB 6Gbps SATA G3HS 2.5" SSD
Raw capacity per disk (GB)	1966 GB (1.92 TB)
Number of physical disks in test	96 x 1.92 TB
Total raw storage capacity (GB)	188,736 GB
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	71.33 TB
Storage capacity utilization	39%
Database capacity utilization	17.6%

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	3559.83	3441.53	3684.05
Database Disks Reads/sec	2530.87	2445.83	2514.00
Database Disks Write/sec	1210.10	1171.54	1260.44
Average Database Disk Read Latency (ms)	6.44	7.56	6.40
Average Database Disk Write Latency (ms)	8.14	10.24	8.26
Transaction Log I/O			
Log Disks Writes/sec	15.65	15.12	15.97
Average Log Disk Write Latency(ms)	2.58	2.31	2.63

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	10685.41
Database Disks Reads/sec	7490.70
Database Disks Writes/sec	3642.08
Average Database Disk Read Latency (ms)	6.8
Average Database Disk Write Latency (ms)	8.88
Transaction Log I/O	
Log Disks Writes/sec	46.74
Average Log Disk Write Latency (ms)	2.51

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 20 database files on a single node.

Table 3: Database backup read-only performance results

MB read/sec per database	133.32
MB read/sec total per node (10 databases)	1330.05

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)	3.25

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 60,000 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 6 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
e:\Jetstress001001.edb	24158976	0	0	0	754968 MB/8917 sec
f:\Jetstress002001.edb	24159744	0	0	0	754992 MB/8910 sec
g:\Jetstress003001.edb	24160256	0	0	0	755008 MB/8873 sec
h:\Jetstress004001.edb	24160000	0	0	0	755000 MB/8871 sec
i:\Jetstress005001.edb	24158976	0	0	0	754968 MB/8868 sec
j:\Jetstress006001.edb	24159744	0	0	0	754992 MB/8865 sec
k:\Jetstress007001.edb	24159488	0	0	0	754984 MB/8904 sec
l:\Jetstress008001.edb	24159232	0	0	0	754976 MB/8859 sec
m:\Jetstress009001.edb	24159488	0	0	0	754984 MB/8850 sec
n:\Jetstress010001.edb	24159744	0	0	0	754992 MB/8901 sec
(Sum)	241595648	0	0	0	7549864 MB/8917 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
E:	0.028	0.000	1348.411	0.000	65536.000
F:	0.029	0.000	1351.555	0.000	65536.000
G:	0.028	0.000	1360.243	0.000	65536.000
H:	0.028	0.000	1360.611	0.000	65536.000
I:	0.028	0.000	1361.514	0.000	65536.000
J:	0.028	0.000	1362.484	0.000	65536.000
K:	0.029	0.000	1354.019	0.000	65536.000
L:	0.028	0.000	1362.398	0.000	65536.000
M:	0.028	0.000	1364.083	0.000	65536.000
N:	0.029	0.000	1355.564	0.000	65536.000

Memory System Performance (of checksum)

Counter	Average	Minimum	Maximum
% Processor Time	0.471	0.000	2.168
Available MBytes	95735.864	95713.000	95748.000
Free System Page Table Entries	12293580.373	12293168.000	12294133.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	110251907.037	110116864.000	110481408.000
Pool Paged Bytes	200873116.203	200826880.000	200962048.000

Test Log

1/16/2018 7:35:48 PM -- Preparing for testing ...
1/16/2018 7:35:58 PM -- Attaching databases ...
1/16/2018 7:35:58 PM -- Preparations for testing are complete.
1/16/2018 7:35:58 PM -- Starting transaction dispatch ..
1/16/2018 7:35:58 PM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
1/16/2018 7:35:58 PM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
1/16/2018 7:36:09 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
1/16/2018 7:36:09 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
1/16/2018 7:36:10 PM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
1/16/2018 7:36:10 PM -- Performance logging started (interval: 15000 ms).
1/16/2018 7:36:10 PM -- Attaining prerequisites:
1/16/2018 7:37:27 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2428822000.0 (lower bound: 2415919000.0, upper bound: none)
1/17/2018 7:37:27 PM -- Performance logging has ended.
1/17/2018 7:37:27 PM -- JetInterop batch transaction stats: 336888, 336888, 336888, 336888, 336888, 336888, 336888, 336888, 336888 and 336887.
1/17/2018 7:37:27 PM -- Dispatching transactions ends.
1/17/2018 7:37:27 PM -- Shutting down databases ...
1/17/2018 7:37:30 PM -- Instance3452.1 (complete), Instance3452.2 (complete), Instance3452.3 (complete), Instance3452.4 (complete), Instance3452.5 (complete), Instance3452.6 (complete), Instance3452.7 (complete), Instance3452.8 (complete), Instance3452.9 (complete) and Instance3452.10 (complete)
1/17/2018 7:37:30 PM -- C:\Program Files\Exchange Jetstress\24HR_AF-R1-96GB-20000-750MB\Stress_2018_1_16_19_36_9.blg has 5752 samples.
1/17/2018 7:37:30 PM -- Creating test report ...
1/17/2018 7:38:46 PM -- Instance3452.1 has 7.6 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.1 has 3.3 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.1 has 3.3 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.2 has 8.0 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.2 has 3.2 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.2 has 3.2 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.3 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.3 has 3.3 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.3 has 3.3 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.4 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.4 has 3.3 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.4 has 3.3 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.5 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.5 has 3.3 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.5 has 3.3 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.6 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.6 has 3.2 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.6 has 3.2 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.7 has 8.0 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.7 has 3.2 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.7 has 3.2 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.8 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.8 has 3.4 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.8 has 3.4 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.9 has 7.4 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.9 has 3.2 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.9 has 3.2 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.10 has 8.0 for I/O Database Reads Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.10 has 3.2 for I/O Log Writes Average Latency.
1/17/2018 7:38:46 PM -- Instance3452.10 has 3.2 for I/O Log Reads Average Latency.
1/17/2018 7:38:46 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
1/17/2018 7:38:46 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
1/17/2018 7:38:46 PM -- C:\Program Files\Exchange Jetstress\24HR_AF-R1-96GB-20000-750MB\Stress_2018_1_16_19_36_9.xml has 5746 samples queried.
1/17/2018 7:38:46 PM -- C:\Program Files\Exchange Jetstress\24HR_AF-R1-96GB-20000-750MB\Stress_2018_1_16_19_36_9.html was saved.
1/17/2018 7:38:46 PM -- Performance logging started (interval: 30000 ms).
1/17/2018 7:38:46 PM -- Verifying database checksums ...
1/17/2018 10:07:24 PM -- E: (100% processed), F: (100% processed), G: (100% processed), H: (100% processed), I: (100% processed), J: (100% processed), K: (100% processed), L: (100% processed), M: (100% processed) and N: (100% processed)
1/17/2018 10:07:24 PM -- Performance logging has ended.
1/17/2018 10:07:24 PM -- C:\Program Files\Exchange Jetstress\24HR_AF-R1-96GB-20000-750MB\DBChecksum_2018_1_17_19_38_46.blg has 295 samples.

2-Hour performance test results

Test Summary

Overall Test Result	**Pass**
Machine Name	EX-1
Test Description	
Test Start Time	1/18/2018 6:50:57 AM
Test End Time	1/18/2018 8:52:36 AM
Collection Start Time	1/18/2018 6:52:31 AM
Collection End Time	1/18/2018 8:52:20 AM
Jetstress Version	15.01.1019.000
ESE Version	15.01.0225.037
Operating System	Windows Server 2016 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\2Hr-AF-Storage-R1-96GB-20000-750MB\Performance_2018_1_18_6_51_18.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	2139.156
Capacity Percentage	100%
Throughput Percentage	100%
Initial Database Size (bytes)	7918745288704
Final Database Size (bytes)	7923677790208
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

Instance3884.1	Log path: e:\logs Database: e:\Jetstress001001.edb
Instance3884.2	Log path: f:\logs Database: f:\Jetstress002001.edb
Instance3884.3	Log path: g:\logs Database: g:\Jetstress003001.edb
Instance3884.4	Log path: h:\logs Database: h:\Jetstress004001.edb
Instance3884.5	Log path: i:\logs Database: i:\Jetstress005001.edb
Instance3884.6	Log path: j:\logs Database: j:\Jetstress006001.edb
Instance3884.7	Log path: k:\logs Database: k:\Jetstress007001.edb
Instance3884.8	Log path: l:\logs Database: l:\Jetstress008001.edb
Instance3884.9	Log path: m:\logs Database: m:\Jetstress009001.edb
Instance3884.10	Log path: n:\logs Database: n:\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
Instance3884.1	6.820	10.624	138.879	75.754	36586.522	34263.687	0.000	3.340	0.000	15.723	0.000	18783.706
Instance3884.2	7.195	10.305	139.099	76.003	36659.785	34267.343	0.000	3.072	0.000	15.807	0.000	18763.972
Instance3884.3	6.773	10.528	138.574	75.198	36661.489	34272.678	0.000	2.808	0.000	15.661	0.000	18919.762
Instance3884.4	6.654	11.001	138.473	75.264	36630.221	34282.891	0.000	3.060	0.000	15.783	0.000	18719.753
Instance3884.5	6.698	12.395	138.368	74.508	36612.827	34262.980	0.000	3.548	0.000	15.626	0.000	18803.068
Instance3884.6	6.711	12.135	138.140	75.046	36652.188	34256.896	0.000	3.179	0.000	15.662	0.000	18910.086
Instance3884.7	7.132	12.448	139.062	75.769	36683.548	34244.099	0.000	3.183	0.000	15.693	0.000	18683.141
Instance3884.8	6.660	12.107	138.625	75.298	36683.765	34278.609	0.000	2.887	0.000	15.693	0.000	18843.991
Instance3884.9	6.747	13.904	138.427	75.380	36657.097	34275.349	0.000	3.314	0.000	15.647	0.000	18877.262
Instance3884.10	7.097	13.406	138.292	74.996	36558.670	34263.365	0.000	2.951	0.000	15.691	0.000	18651.657

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3884.1	8.748	261579.508
Instance3884.2	8.742	261465.442
Instance3884.3	8.733	261648.453
Instance3884.4	8.754	261600.367
Instance3884.5	8.746	261560.081
Instance3884.6	8.759	261547.685
Instance3884.7	8.739	261651.881
Instance3884.8	8.756	261597.510
Instance3884.9	8.741	261538.650
Instance3884.10	8.744	261588.053

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3884.1	1.212	229645.865
Instance3884.2	1.221	228965.057
Instance3884.3	1.215	231569.035
Instance3884.4	1.216	230615.654
Instance3884.5	1.209	228171.179
Instance3884.6	1.215	229642.592
Instance3884.7	1.202	230129.123
Instance3884.8	1.212	230615.654
Instance3884.9	1.209	229630.772
Instance3884.10	1.209	230127.850

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
Instance3884.1	6.820	10.624	147.627	75.754	49919.241	34263.687	2.073	3.340	1.212	15.723	229645.865	18783.706
Instance3884.2	7.195	10.305	147.841	76.003	49952.185	34267.343	2.285	3.072	1.221	15.807	228965.057	18763.972
Instance3884.3	6.773	10.528	147.307	75.198	49999.834	34272.678	2.209	2.808	1.215	15.661	231569.035	18919.762
Instance3884.4	6.654	11.001	147.227	75.264	50006.932	34282.891	2.162	3.060	1.216	15.783	230615.654	18719.753
Instance3884.5	6.698	12.395	147.114	74.508	49985.514	34262.980	2.147	3.548	1.209	15.626	228171.179	18803.068
Instance3884.6	6.711	12.135	146.898	75.046	50061.242	34256.896	2.146	3.179	1.215	15.662	229642.592	18910.086
Instance3884.7	7.132	12.448	147.801	75.769	49985.375	34244.099	2.147	3.183	1.202	15.693	230129.123	18683.141
Instance3884.8	6.660	12.107	147.381	75.298	50046.332	34278.609	2.451	2.887	1.212	15.693	230615.654	18843.991
Instance3884.9	6.747	13.904	147.168	75.380	50013.552	34275.349	2.296	3.314	1.209	15.647	229630.772	18877.262
Instance3884.10	7.097	13.406	147.036	74.996	49941.146	34263.365	2.101	2.951	1.209	15.691	230127.850	18651.657

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.975	0.280	1.779
Available MBytes	93072.063	93041.000	93182.000
Free System Page Table Entries	12293913.365	12293379.000	12294212.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	110643666.038	110505984.000	110759936.000
Pool Paged Bytes	201465471.198	201379840.000	201637888.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```
1/18/2018 6:50:57 AM -- Preparing for testing ...
1/18/2018 6:51:07 AM -- Attaching databases ...
1/18/2018 6:51:07 AM -- Preparations for testing are complete.
1/18/2018 6:51:07 AM -- Starting transaction dispatch ..
1/18/2018 6:51:07 AM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
1/18/2018 6:51:07 AM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
1/18/2018 6:51:18 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
1/18/2018 6:51:18 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
1/18/2018 6:51:19 AM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
1/18/2018 6:51:19 AM -- Performance logging started (interval: 15000 ms).
1/18/2018 6:51:19 AM -- Attaining prerequisites:
1/18/2018 6:52:31 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 2418287000.0 (lower bound: 2415919000.0, upper bound: none)
1/18/2018 8:52:32 AM -- Performance logging has ended.
1/18/2018 8:52:32 AM -- JetInterop batch transaction stats: 35234, 35234, 35233, 35233, 35233, 35233, 35233, 35233, 35233 and 35233.
1/18/2018 8:52:32 AM -- Dispatching transactions ends.
1/18/2018 8:52:32 AM -- Shutting down databases ...
1/18/2018 8:52:36 AM -- Instance3884.1 (complete), Instance3884.2 (complete), Instance3884.3 (complete), Instance3884.4 (complete), Instance3884.5 (complete), Instance3884.6 (complete), Instance3884.7 (complete), Instance3884.8 (complete), Instance3884.9 (complete) and Instance3884.10 (complete)
1/18/2018 8:52:36 AM -- C:\Program Files\Exchange Jetstress\2Hr-AF-Storage-R1-96GB-20000-750MB\Performance_2018_1_18_6_51_18.blg has 483 samples.
1/18/2018 8:52:36 AM -- Creating test report ...
1/18/2018 8:52:42 AM -- Instance3884.1 has 6.8 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.1 has 3.3 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.1 has 3.3 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.2 has 7.2 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.2 has 3.1 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.2 has 3.1 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.3 has 6.8 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.3 has 2.8 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.3 has 2.8 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.4 has 6.7 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.4 has 3.1 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.4 has 3.1 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.5 has 6.7 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.5 has 3.5 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.5 has 3.5 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.6 has 6.7 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.6 has 3.2 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.6 has 3.2 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.7 has 7.1 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.7 has 3.2 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.7 has 3.2 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.8 has 6.7 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.8 has 2.9 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.8 has 2.9 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.9 has 6.7 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.9 has 3.3 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.9 has 3.3 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.10 has 7.1 for I/O Database Reads Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.10 has 3.0 for I/O Log Writes Average Latency.
1/18/2018 8:52:42 AM -- Instance3884.10 has 3.0 for I/O Log Reads Average Latency.
1/18/2018 8:52:42 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
1/18/2018 8:52:42 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
1/18/2018 8:52:42 AM -- C:\Program Files\Exchange Jetstress\2Hr-AF-Storage-R1-96GB-20000-750MB\Performance_2018_1_18_6_51_18.xml has 478 samples queried.
```

Database backup test results

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance2504.1	755632.03	01:35:53	131.34
Instance2504.2	755664.03	01:34:39	133.05
Instance2504.3	755672.03	01:34:34	133.17
Instance2504.4	755664.03	01:34:34	133.16
Instance2504.5	755632.03	01:34:28	133.29
Instance2504.6	755664.03	01:34:28	133.30
Instance2504.7	755648.03	01:34:27	133.33
Instance2504.8	755648.03	01:34:18	133.53
Instance2504.9	755648.03	01:34:14	133.63
Instance2504.10	755656.03	01:34:12	133.69
Avg			133.15
Sum			1331.51

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

Instance2504.1 Log path: e:\logs
Database: e:\Jetstress001001.edb

Instance2504.2 Log path: f:\logs
Database: f:\Jetstress002001.edb

Instance2504.3 Log path: g:\logs
Database: g:\Jetstress003001.edb

Instance2504.4 Log path: h:\logs
Database: h:\Jetstress004001.edb

Instance2504.5 Log path: i:\logs
Database: i:\Jetstress005001.edb

Instance2504.6 Log path: j:\logs
Database: j:\Jetstress006001.edb

Instance2504.7 Log path: k:\logs
Database: k:\Jetstress007001.edb

Instance2504.8 Log path: l:\logs
Database: l:\Jetstress008001.edb

Instance2504.9 Log path: m:\logs
Database: m:\Jetstress009001.edb

Instance2504.10 Log path: n:\logs
Database: n:\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
Instance2504.1	3.074	0.000	524.032	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.2	3.026	0.000	531.742	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.3	3.012	0.000	531.987	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.4	3.014	0.000	531.904	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.5	3.003	0.000	532.621	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.6	2.983	0.000	532.669	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.7	2.986	0.000	532.781	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.8	2.969	0.000	533.709	0.000	262144.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.9	2.950	0.000	534.148	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2504.10	2.943	0.000	534.411	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	2.408	0.394	2.902
Available MBytes	95686.555	95655.000	95740.000
Free System Page Table Entries	12294256.691	12293820.000	12294555.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	110951311.414	110923776.000	111026176.000
Pool Paged Bytes	202280059.309	202264576.000	202342400.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

1/18/2018 12:50:02 PM -- Preparing for testing ...
1/18/2018 12:50:12 PM -- Attaching databases ...
1/18/2018 12:50:12 PM -- Preparations for testing are complete.
1/18/2018 12:50:23 PM -- Performance logging started (interval: 30000 ms).
1/18/2018 12:50:23 PM -- Backing up databases ...
1/18/2018 2:26:17 PM -- Performance logging has ended.
1/18/2018 2:26:17 PM -- Instance2504.1 (100% processed), Instance2504.2 (100% processed), Instance2504.3 (100% processed), Instance2504.4 (100% processed), Instance2504.5 (100% processed), Instance2504.6 (100% processed), Instance2504.7 (100% processed), Instance2504.8 (100% processed), Instance2504.9 (100% processed) and Instance2504.10 (100% processed)
1/18/2018 2:26:17 PM -- C:\Program Files\Exchange Jetstress\DBBackup_AF-R1-96GB-20000-750MB\DatabaseBackup_2018_1_18_12_50_12.blg has 191 samples.
1/18/2018 2:26:17 PM -- Creating test report ...

Soft recovery test results

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance3064.1	504	1853.9829033
Instance3064.2	508	1869.0297954
Instance3064.3	501	1843.2016125
Instance3064.4	511	1863.9672992
Instance3064.5	504	1863.6860467
Instance3064.6	505	1853.4516371
Instance3064.7	502	1854.7953762
Instance3064.8	502	1855.0610084
Instance3064.9	503	1850.4828894
Instance3064.10	507	1870.3735519
Avg	504	1857.803
Sum	5047	18578.0321201

Database Configuration

Instance3064.1 Log path: e:\logs
Database: e:\Jetstress001001.edb

Instance3064.2 Log path: f:\logs
Database: f:\Jetstress002001.edb

Instance3064.3 Log path: g:\logs
Database: g:\Jetstress003001.edb

Instance3064.4 Log path: h:\logs
Database: h:\Jetstress004001.edb

Instance3064.5 Log path: i:\logs
Database: i:\Jetstress005001.edb

Instance3064.6 Log path: j:\logs
Database: j:\Jetstress006001.edb

Instance3064.7 Log path: k:\logs
Database: k:\Jetstress007001.edb

Instance3064.8 Log path: l:\logs
Database: l:\Jetstress008001.edb

Instance3064.9 Log path: m:\logs
Database: m:\Jetstress009001.edb

Instance3064.10 Log path: n:\logs
Database: n:\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
Instance3064.1	2.252	1.489	121.273	2.657	37184.959	77987.042	1.566	0.000	1.357	0.000	206516.929	0.000
Instance3064.2	2.375	1.619	122.596	2.658	37258.597	78058.726	1.402	0.000	1.357	0.000	206991.626	0.000
Instance3064.3	2.239	1.522	122.481	2.638	37268.598	78658.390	1.461	0.000	1.355	0.000	206955.789	0.000
Instance3064.4	2.297	1.828	121.917	2.736	37236.500	77795.399	1.473	0.000	1.368	0.000	207931.974	0.000
Instance3064.5	2.195	1.892	121.046	2.598	37379.461	79111.880	1.446	0.000	1.349	0.000	206530.804	0.000
Instance3064.6	2.183	1.521	121.790	2.667	37272.553	78916.007	1.466	0.000	1.365	0.000	207919.615	0.000
Instance3064.7	2.259	1.400	122.730	2.637	37390.270	79064.188	1.470	0.000	1.352	0.000	207887.617	0.000
Instance3064.8	2.240	1.569	122.372	2.652	37456.310	78257.944	1.434	0.000	1.352	0.000	206973.825	0.000
Instance3064.9	2.246	1.411	121.661	2.648	37248.259	78543.668	1.430	0.000	1.357	0.000	206967.839	0.000
Instance3064.10	2.215	1.569	122.025	2.644	37297.844	79003.941	1.495	0.000	1.354	0.000	207903.406	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3064.1	0.000	0.000
Instance3064.2	0.000	0.000
Instance3064.3	0.000	0.000
Instance3064.4	0.000	0.000
Instance3064.5	0.000	0.000
Instance3064.6	0.000	0.000
Instance3064.7	0.000	0.000
Instance3064.8	0.000	0.000
Instance3064.9	0.000	0.000
Instance3064.10	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
Instance3064.1	2.252	1.489	121.273	2.657	37184.959	77987.042	1.566	0.000	1.357	0.000	206516.929	0.000
Instance3064.2	2.375	1.619	122.596	2.658	37258.597	78058.726	1.402	0.000	1.357	0.000	206991.626	0.000
Instance3064.3	2.239	1.522	122.481	2.638	37268.598	78658.390	1.461	0.000	1.355	0.000	206955.789	0.000
Instance3064.4	2.297	1.828	121.917	2.736	37236.500	77795.399	1.473	0.000	1.368	0.000	207931.974	0.000
Instance3064.5	2.195	1.892	121.046	2.598	37379.461	79111.880	1.446	0.000	1.349	0.000	206530.804	0.000
Instance3064.6	2.183	1.521	121.790	2.667	37272.553	78916.007	1.466	0.000	1.365	0.000	207919.615	0.000
Instance3064.7	2.259	1.400	122.730	2.637	37390.270	79064.188	1.470	0.000	1.352	0.000	207887.617	0.000
Instance3064.8	2.240	1.569	122.372	2.652	37456.310	78257.944	1.434	0.000	1.352	0.000	206973.825	0.000
Instance3064.9	2.246	1.411	121.661	2.648	37248.259	78543.668	1.430	0.000	1.357	0.000	206967.839	0.000
Instance3064.10	2.215	1.569	122.025	2.644	37297.844	79003.941	1.495	0.000	1.354	0.000	207903.406	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.640	0.049	3.200
Available MBytes	93195.662	93064.000	95591.000
Free System Page Table Entries	12294090.011	12293659.000	12294385.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	112026959.448	111878144.000	112205824.000
Pool Paged Bytes	202934510.345	202883072.000	204034048.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```
1/18/2018 4:48:39 PM -- Preparing for testing ...
1/18/2018 4:48:49 PM -- Attaching databases ...
1/18/2018 4:48:49 PM -- Preparations for testing are complete.
1/18/2018 4:48:49 PM -- Starting transaction dispatch ..
1/18/2018 4:48:49 PM -- Database cache settings: (minimum: 320.0 MB, maximum: 2.5 GB)
1/18/2018 4:48:49 PM -- Database flush thresholds: (start: 25.6 MB, stop: 51.2 MB)
1/18/2018 4:49:00 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
1/18/2018 4:49:00 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
1/18/2018 4:49:01 PM -- Operation mix: Sessions 16, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
1/18/2018 4:49:01 PM -- Performance logging started (interval: 15000 ms).
1/18/2018 4:49:01 PM -- Generating log files ...
1/18/2018 5:35:12 PM -- e:\logs (100.8% generated), f:\logs (101.6% generated), g:\logs (100.2% generated), h:\logs (102.2% generated), i:\logs (100.8% generated), j:\logs (101.0% generated), k:\logs (100.4% generated), l:\logs (100.4% generated), m:\logs (100.6% generated) and n:\logs
(101.4% generated)
1/18/2018 5:35:13 PM -- Performance logging has ended.
1/18/2018 5:35:13 PM -- JetInterop batch transaction stats: 18036, 18036, 18036, 18036, 18036, 18036, 18036, 18036, 18036 and 18035.
1/18/2018 5:35:13 PM -- Dispatching transactions ends.
1/18/2018 5:35:13 PM -- Shutting down databases ...
1/18/2018 5:35:17 PM -- Instance3064.1 (complete), Instance3064.2 (complete), Instance3064.3 (complete), Instance3064.4 (complete), Instance3064.5 (complete), Instance3064.6 (complete), Instance3064.7 (complete), Instance3064.8 (complete), Instance3064.9 (complete) and Instance3064.10
(complete)
1/18/2018 5:35:17 PM -- C:\Program Files\Exchange Jetstress\SW\Recovery_AF-R1-96GB-20000-750MB\Performance_2018_1_18_16_49_0.blg has 184 samples.
1/18/2018 5:35:17 PM -- Creating test report ...
1/18/2018 5:35:19 PM -- Instance3064.1 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.1 has 2.4 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.1 has 2.4 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.2 has 5.9 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.2 has 2.0 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.2 has 2.0 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.3 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.3 has 3.1 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.3 has 3.1 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.4 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.4 has 2.2 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.4 has 2.2 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.5 has 5.7 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.5 has 2.9 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.5 has 2.9 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.6 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.6 has 2.3 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.6 has 2.3 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.7 has 5.9 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.7 has 2.5 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.7 has 2.5 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.8 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.8 has 2.2 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.8 has 2.2 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.9 has 5.6 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.9 has 3.4 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.9 has 3.4 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.10 has 5.9 for I/O Database Reads Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.10 has 2.2 for I/O Log Writes Average Latency.
1/18/2018 5:35:19 PM -- Instance3064.10 has 2.2 for I/O Log Reads Average Latency.
1/18/2018 5:35:19 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
1/18/2018 5:35:19 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
1/18/2018 5:35:19 PM -- C:\Program Files\Exchange Jetstress\SW\Recovery_AF-R1-96GB-20000-750MB\Performance_2018_1_18_16_49_0.xml has 183 samples queried.
1/18/2018 5:35:19 PM -- C:\Program Files\Exchange Jetstress\SW\Recovery_AF-R1-96GB-20000-750MB\Performance_2018_1_18_16_49_0.html was saved.
1/18/2018 5:35:20 PM -- Performance logging started (interval: 4000 ms).
1/18/2018 5:35:20 PM -- Recovering databases ...
1/18/2018 6:06:31 PM -- 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.