Reference Architecture for Workloads using the Lenovo Converged HX Series Nutanix Appliances

Last update: 05 August 2016
Version 1.1

Provides a technical overview of Lenovo Converged HX Series Nutanix appliances

Contains performance data and sizing recommendations

Shows variety of workloads that can be used in hyper-converged environment

Explains reliability and performance features of hyper-converged appliances

Mike Perks
Pawan Sharma
Frank Zimmer

Table of Contents

1 Introduction

The intended audience for this document is technical IT architects, system administrators, and managers who are interested in executing workloads on the Lenovo Converged 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. Hyper-converged means incorporating multiple components like compute and storage into a single entity through software. A hyper-converged infrastructure 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.

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.

Each of the subsequent chapters in the document describes a particular virtualized workload and provides recommendations on what appliance model to use and how to size the appliance to that workload. Some best practice recommendations are also listed. HX Series appliances are not limited to just the workloads described in this reference architecture and can execute any virtualized workload on the supported hypervisors.

This RA describes five workloads:

- Citrix XenDesktop (see chapter 3 on page 14)
- Microsoft Exchange (see chapter 4 on page 24)
- Microsoft SQL Server
- SAP Business Suite
- VMware Horizon

The final chapter provides references about the HX Series appliances and Nutanix software.

2 Technical overview of HX Series appliances

Lenovo Converged 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 and are designed to support a variety of workloads and deliver optimized performance.

2.1 Server components

The Lenovo Converged HX Series is available in three models that can be configured to your needs:

- Lenovo Converged HX3500: Optimized for compute-heavy environments such as VDI
- Lenovo Converged HX5500: Optimized for storage-heavy workloads such as file servers
- Lenovo Converged HX7500: Optimized for high-performance workloads such as databases

Figure 1 shows each of the 3 models.

Figure 1: Lenovo Converged HX Series appliances: HX3500 (top), HX5500 (middle), HX7500 (bottom)

Table 1 provides a summary of the hardware characteristics of the three HX Series appliances. The default is shown in *italics*.

Table 1: Summary of HX Series appliances

Appliance Model	Form Factor	Solid State Drives (SSD)	Hard Disk Drives (HDD)	Memory	CPU	Dual port network Interface Cards (NICs)
HX3500	2.5" drives	2 x 400 GB *2 x 800 GB*	*6 x 1TB* 6 x 2TB	128 GB 256 GB 384 GB *512 GB* 768 GB	2 x 2630 v3 2 x 2660 v3 *2 x 2680 v3* 2 x 2697 v3 2 x 2699 v3	*1* or 2
HX5500	3.5" drives	*2 x 400 GB* 2 x 800 GB	6 x 2TB *6 x 4TB* 6 x 6TB	128 GB *256 GB* 384 GB 512 GB 768 GB	2 x 2620 v3 *2 x 2630 v3* 2 x 2660 v3 2 x 2680 v3 2 x 2697 v3	*1* or 2
HX7500	2.5" drives	*4 x 400 GB* 4 x 800 GB	*20 x 1TB* 20 x 2TB	128 GB 256 GB 384 GB *512 GB* 768 GB	2 x 2643 v3 2 x 2680 v3 *2 x 2697 v3* 2 x 2699 v3	1 or *2*

For more information, see this website: http://lenovopress.com/lp0059

2.2 Software components

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

2.2.1 Hypervisor

The HX Series appliances support the following hypervisors:

- Nutanix Acropolis Hypervisor based on KVM (AHV)
- VMware ESXi 5.5 U2
- VMware ESXi 5.5 U3
- VMware ESXi 6.0 U1

The HX Series appliances come standard with the Nutanix Acropolis Hypervisor (AHV) preloaded in the factory. The VMware ESXi hypervisors are supported as a field-installable option using the Nutanix Foundation tool.

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 System x M5 and X6 rack servers, Converged HX Series appliances, and Flex System, providing automated agent-less discovery, monitoring, firmware updates, and configuration management.

Lenovo XClarity Administrator is an optional software component for the Lenovo Converged HX Series appliances 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.

Figure 2 shows the Lenovo XClarity administrator interface with some HX Series appliances.

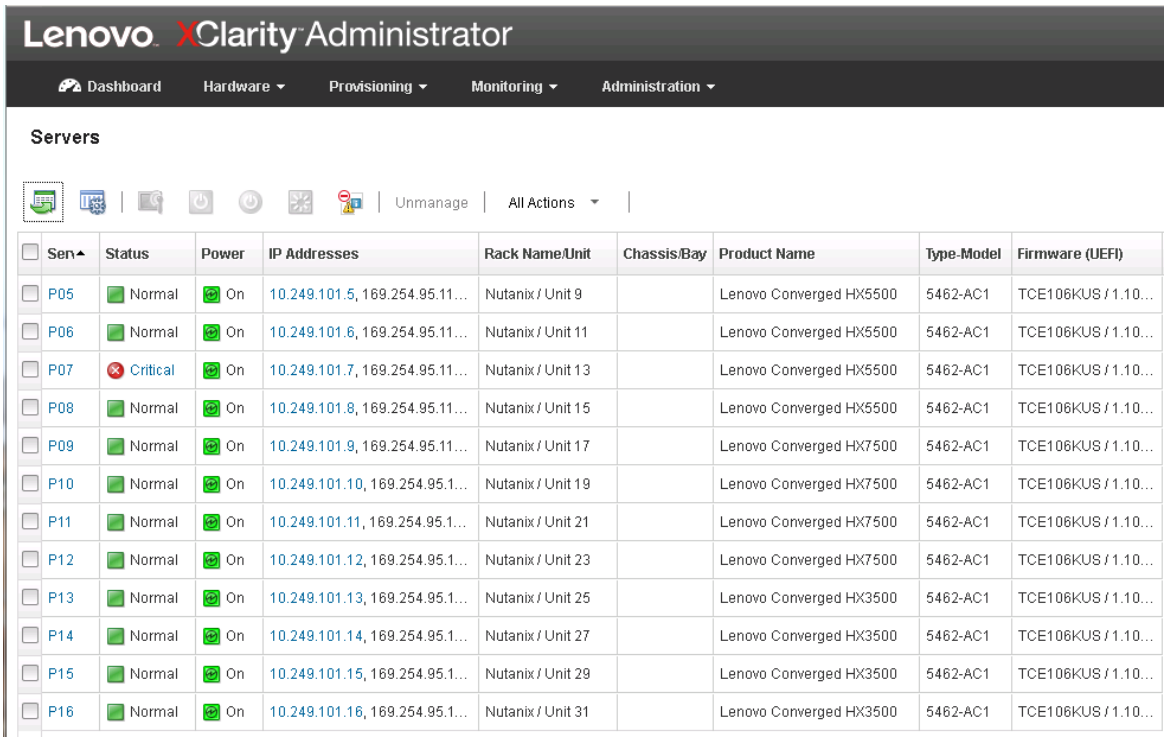

Ser	Status	Power	IP Addresses	Rack Name/Unit	Chassis/Bay	Product Name	Type-Model	Firmware (UEFI)
P05	Normal	On	10.249.101.5, 169.254.95.11...	Nutanix / Unit 9		Lenovo Converged HX5500	5462-AC1	TCE106KUS / 1.10...
P06	Normal	On	10.249.101.6, 169.254.95.11...	Nutanix / Unit 11		Lenovo Converged HX5500	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 2: XClarity Administrator interface

Lenovo XClarity Administrator is provided as a virtual appliance that can be quickly imported into a virtualized environment. Note that XClarity Administrator is not supported on the Acropolis Hypervisor (AHV).

For out-of-band management where the management subnet is different to the server subnet, then XClarity needs to be on a server that is connected into the 1GbE management network. XClarity cannot be run on one of the HX Series appliances.

For in-band management where management and servers are on the same subnet, then there is a choice depending on the data center administrator's preference. XClarity can either be installed on a separate server or a server within a Nutanix cluster running ESXi. Lenovo recommends the first option for medium to large data centers and the second option for smaller deployments to eliminate the incremental cost of another server.

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 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 3 shows the main components of the CVM.

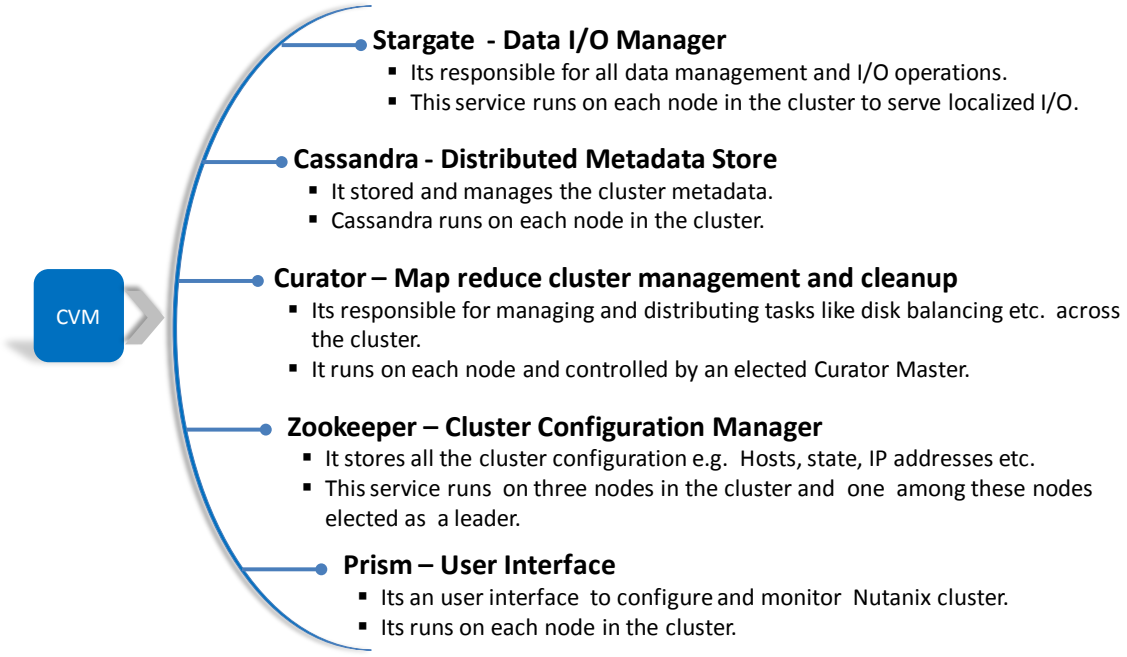

Figure 3: Controller VM components

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 4.

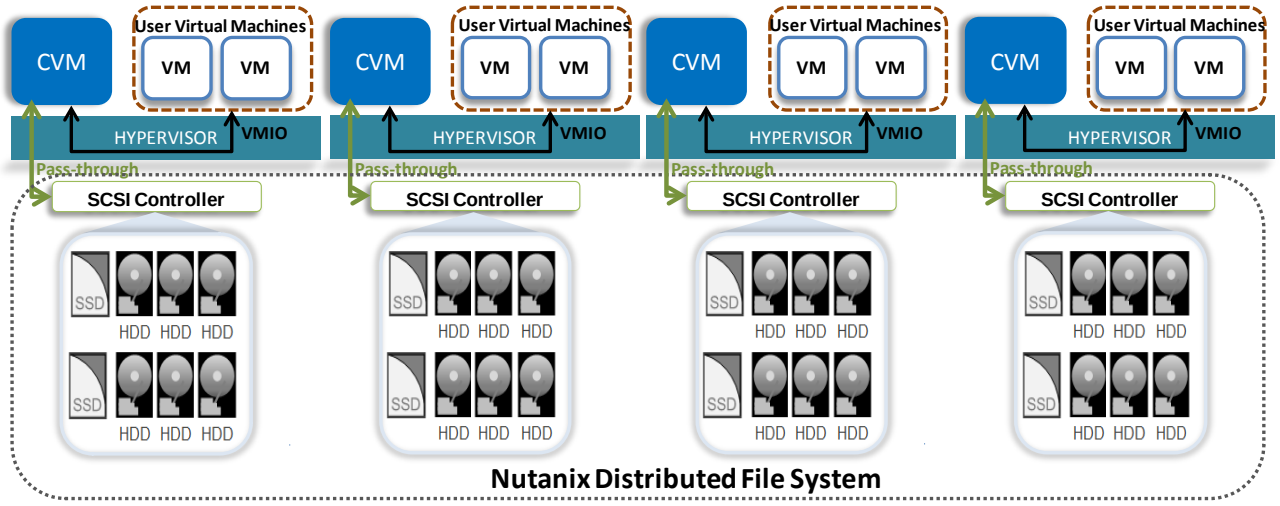

Figure 4: CVM interaction with Hypervisor and User VMs

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 2.4 for more details on the performance features of NDFS.

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.

2.3 Networking components

A cluster of HX Series appliances requires both a 10 GbE network for data and a 1GbE network for hardware management.

2.3.1 10GbE networking

Each Lenovo Converged HX Series appliance contains 1 or 2 dual-port 10GbE network adapters as well as 4 on-board 1GbE ports. The hypervisors are configured by the Nutanix software so that all of the network ports on the appliance (both 10GbE and 1 GbE) are pooled. The hypervisor VM management network should use the same 10GbE 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.

It is recommended that two top of rack (TOR) switches are used for redundancy. The second network adapter provides an additional level of redundancy in case one of the network adapters fails entirely. Redundancy across the two TORs is provided using bonded Inter-Switch-Links (ISL).

Figure 5 shows 4 HX Series appliances each with one dual-port NIC connected to two TOR switches.

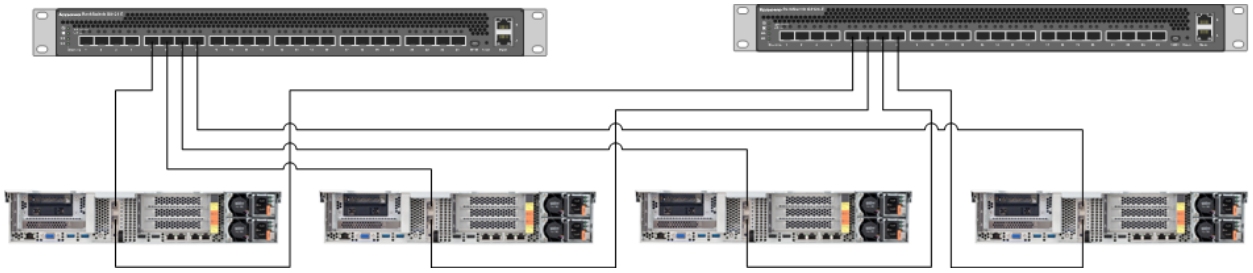

Figure 5: One NIC and two TOR switches

Figure 6 shows 4 HX Series appliances each with two dual-port NICs connected to two TOR switches.

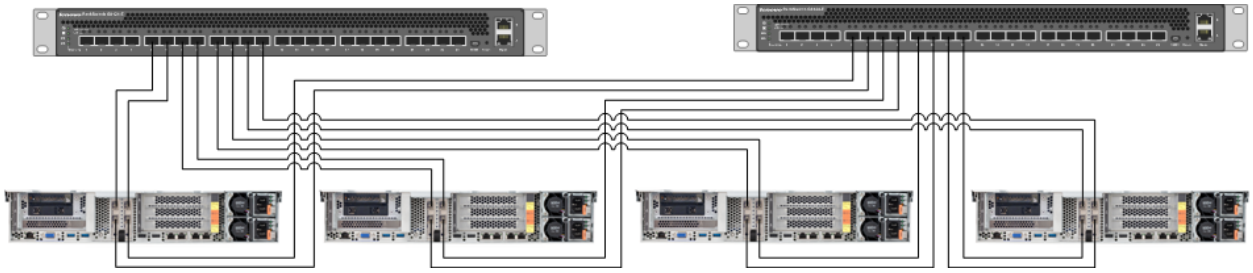

Figure 6: Two NICs and two TOR switches

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

- Lenovo RackSwitch G8124E
- Lenovo RackSwitch G8272

Lenovo RackSwitch G8124E

The Lenovo RackSwitch™ G8124E (as shown in Figure 7) is a 10 GbE switch that is specifically designed for the data center and provides a virtualized, cooler, and easier network solution. The G8124E offers 24 10 GbE ports in a 1U footprint. Designed with top performance in mind, the RackSwitch G8124E provides line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data and large data center grade buffers to keep traffic moving.

The G8124E switch is virtualized by providing rack-level virtualization of networking interfaces. The G8124E switch also supports Virtual Fabric, which allows for the distribution of a physical NIC into 2 to 8 vNICs and creates a virtual pipe between the adapter and the switch. The G8124E switch is easier to manage with server-oriented provisioning by using point-and-click management interfaces.

Figure 7: Lenovo RackSwitch G8124E

For more information, see this website: lenovopress.com/tips0787

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 8), 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 8: Lenovo RackSwitch G8272

The RackSwitch G8272 supports Lenovo Virtual Fabric, which helps clients significantly reduce cost and complexity that are related to I/O requirements of many virtualization deployments. Virtual Fabric helps reduce the number of multiple I/O adapters to a single dual-port 10 GbE adapter and the number of cables and required upstream switch ports.

By using Virtual Fabric, you can carve a dual-port 10 Gb server adapter into eight virtual network ports (vPorts) and create dedicated virtual pipes between the adapter and switch for optimal performance, higher availability, and improved security. With Virtual Fabric, you can make dynamic changes and allocate bandwidth per vPort so that you can adjust it over time without downtime.

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

2.3.2 1 GbE networking

The dedicated IMM port on all of the Lenovo Converged HX series appliances needs to be connected to a 1GbE TOR switch as shown in Figure 9.

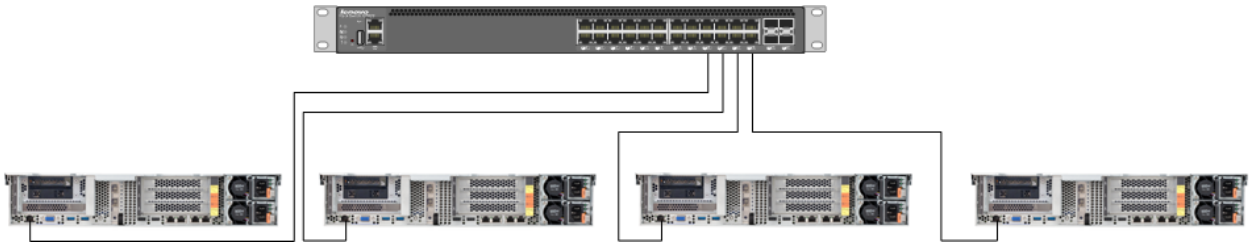

Figure 9: IMM 1GbE management network

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 10) 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 10. 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 11) 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 11: Lenovo RackSwitch G8052

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

2.3.3 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 hyerpvisor
- vSphere vMotion — Used to move VMs from one server to another.
- Fault Tolerance — Used to support the fault tolerance (FT) feature of vSphere.
- Storage network — Used for NDFS storage traffic

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 in a cluster of HX Series appliances is outside of the scope of this document.

2.4 Reliability and performance features

Reliability and excellent performance are important for any workload but particularly for hyper-converged infrastructures like the HX Series appliances.

Reliability, high availability and excellent performance are provided 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 System x servers have been rated #1 in hardware reliability for the last 2 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 HX7500 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 12 shows the high level structure of NDFS:

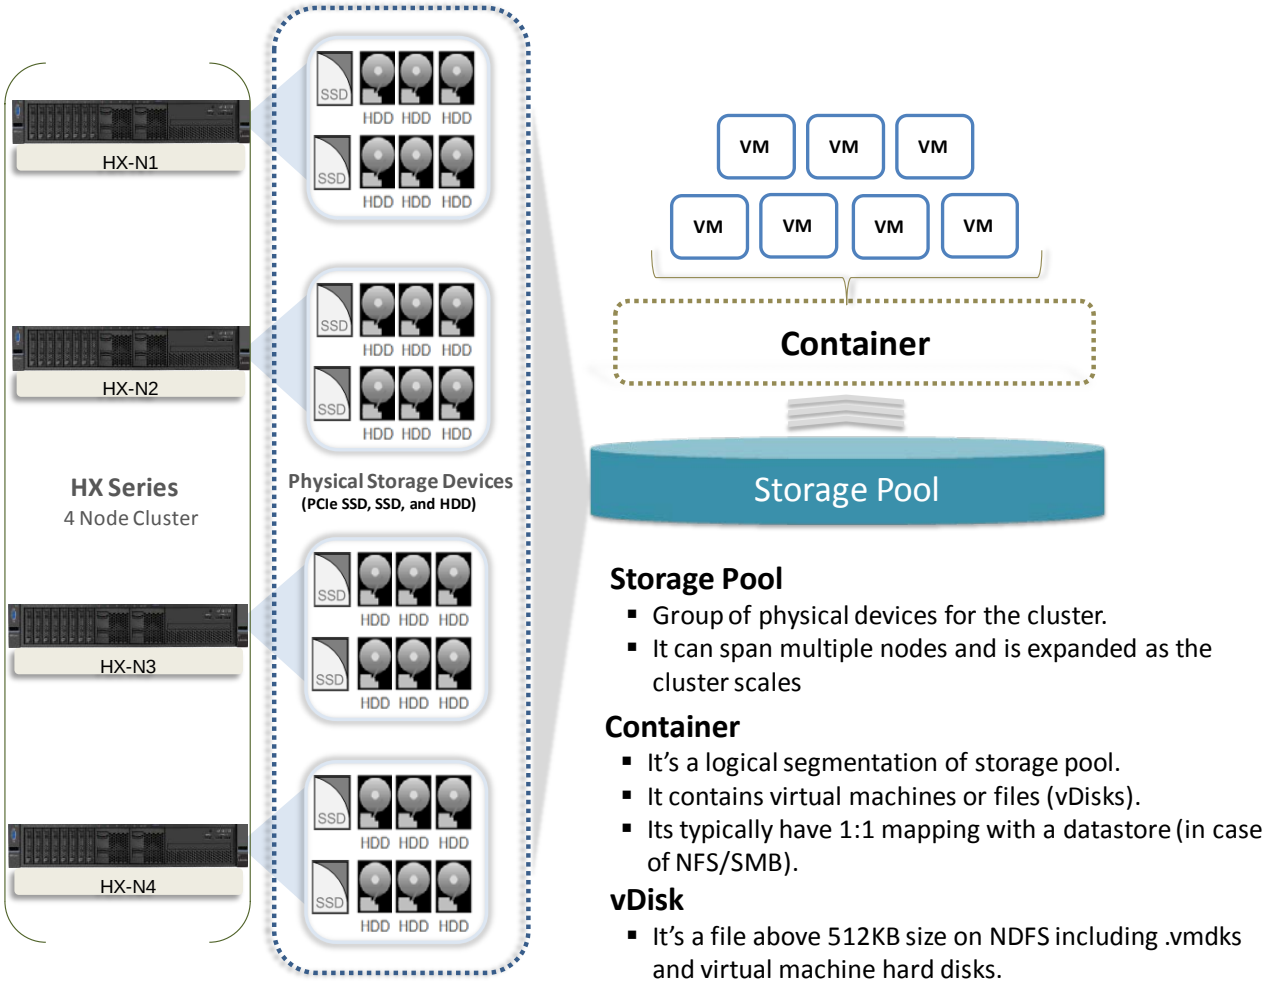

Figure 12: Nutanix Distributed File System

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/Os 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 13 and Figure 14 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.

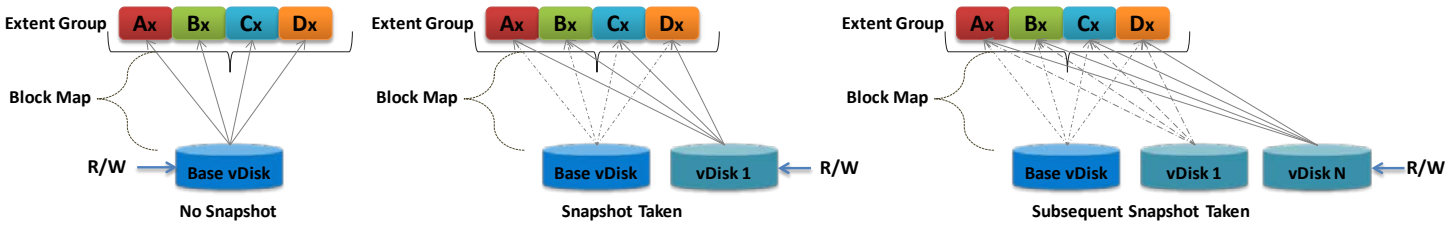

Figure 13: VM snapshots

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.

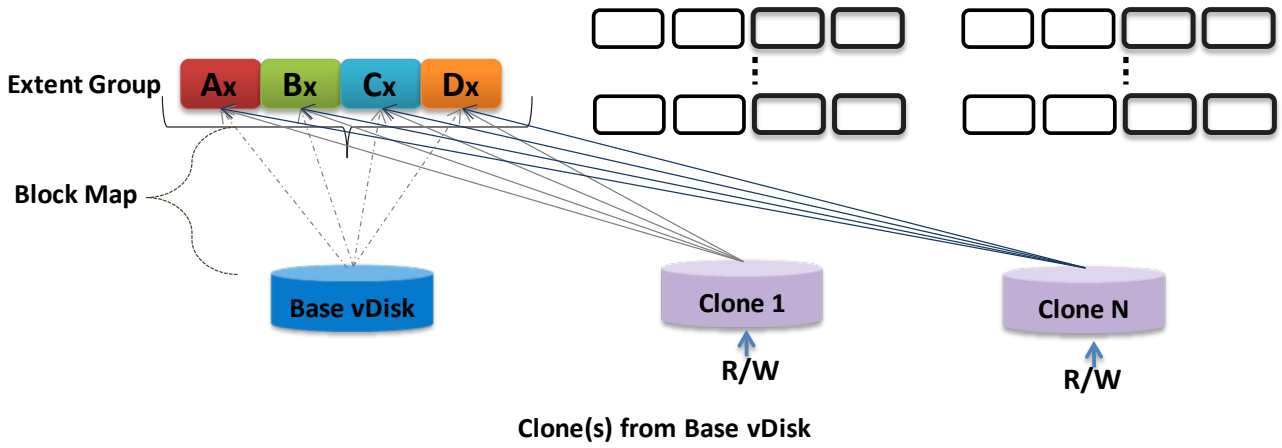

Figure 14: VM clones

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 Citrix XenDesktop

Citrix XenDesktop is a suite of virtualization software which delivers Windows virtual desktops as well as virtual applications to meet the demands of any use case. It is based on the unified FlexCast Management Architecture (FMA) platform. See this website for more details: citrix.com/products/xendesktop.

3.1 Solution overview

Figure 15 provides an architecture overview of the Citrix XenDesktop solution using Lenovo Converged HX3500 appliances. This chapter does not address the general issues of multi-site deployment and network management and limits the description to the components that are inside the customer's intranet.

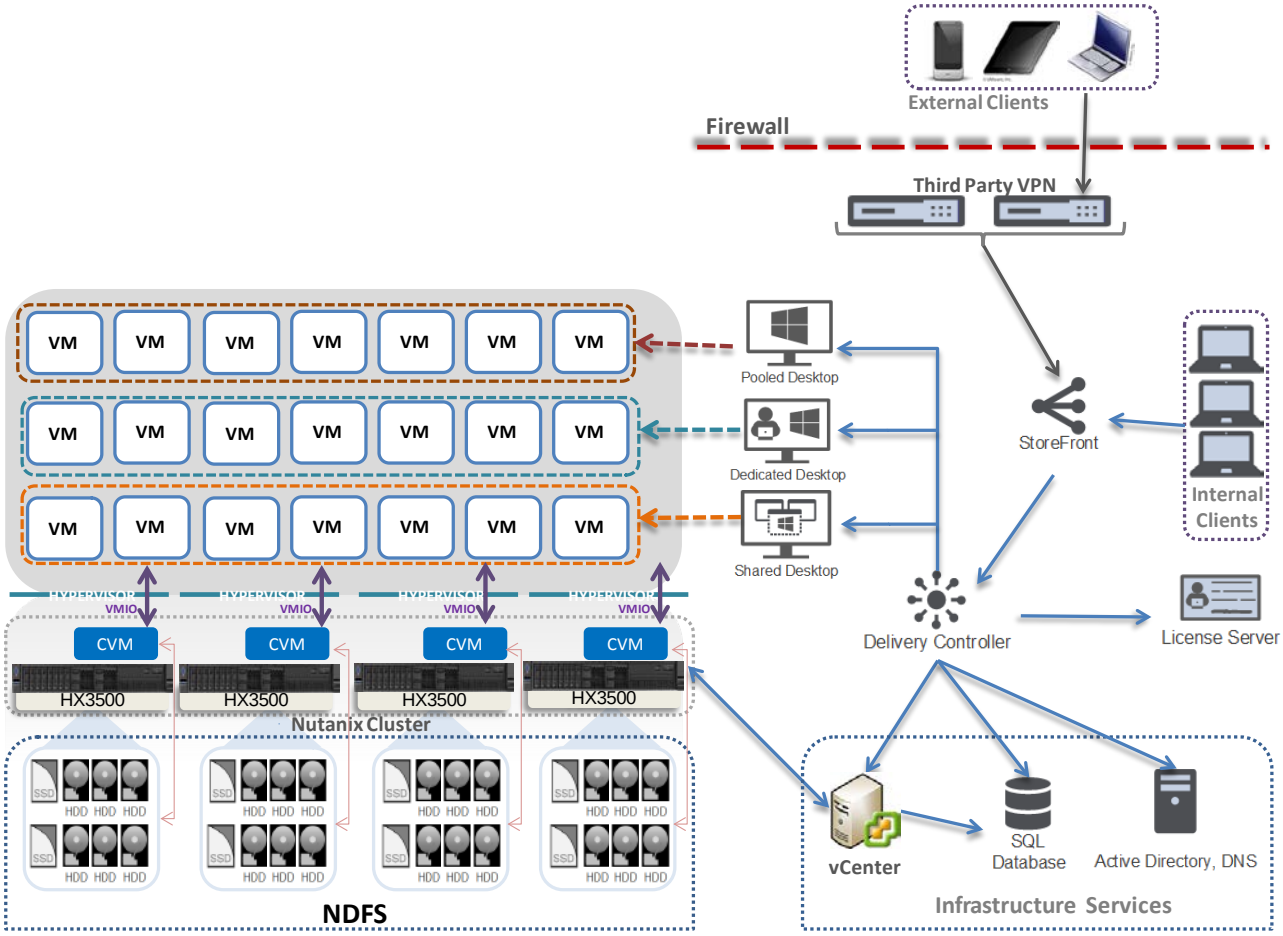

Figure 15: Lenovo Converged HX3500 solution with Citrix XenDesktop

Pooled desktops are stateless (non-persistent) virtual desktops and dedicated desktops are persistent. Shared desktops are used for hosted shared desktops or hosted shared applications.

3.2 Component model

Figure 16 is a layered component view for the Citrix XenDesktop virtualization infrastructure.

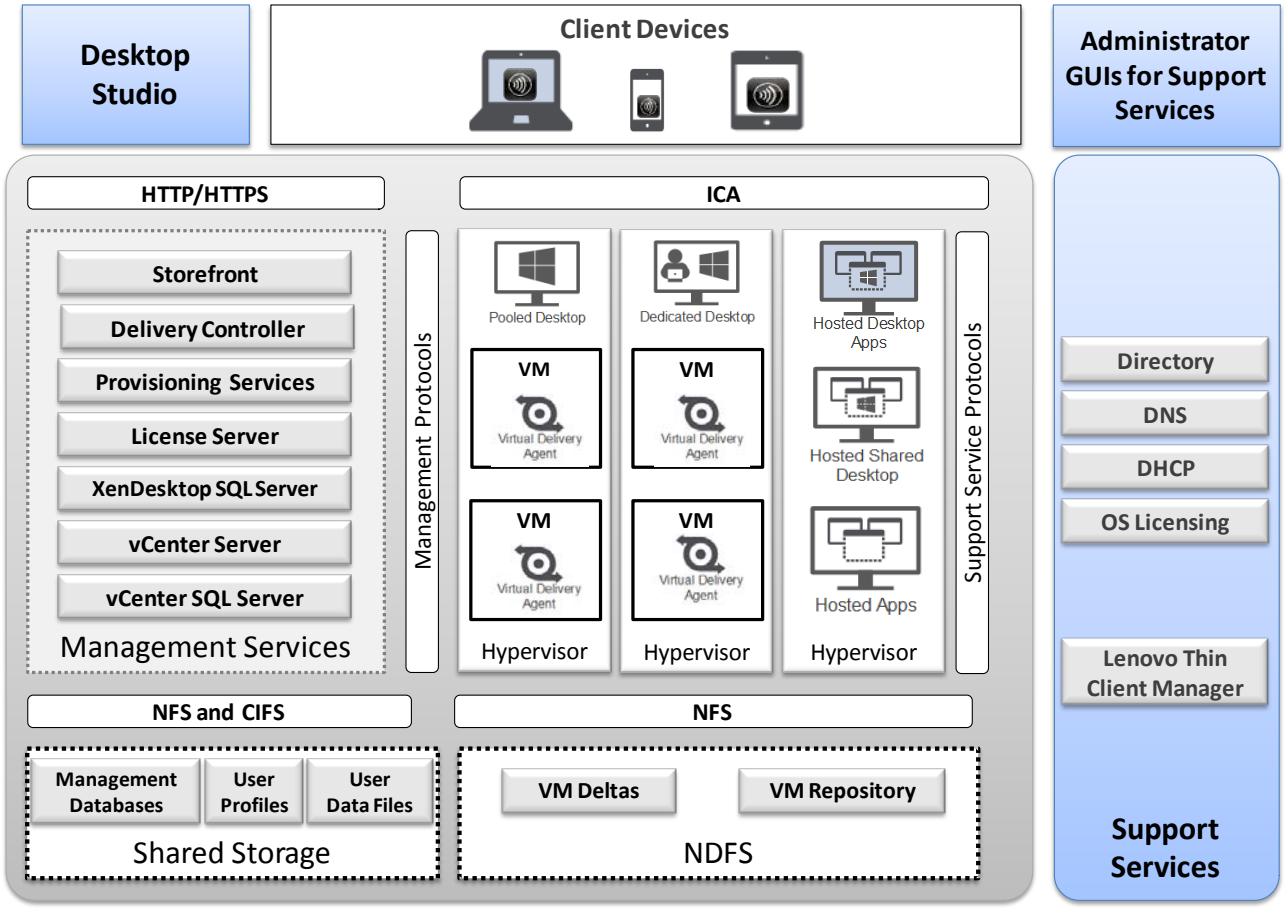

Figure 16: Component model with Citrix XenDesktop

Citrix XenDesktop features the following main components:

Desktop Studio Desktop Studio is the main administrator GUI for Citrix XenDesktop. It is used to configure and manage all of the main entities, including servers, desktop pools and provisioning, policy, and licensing.

Storefront Storefront provides the user interface to the XenDesktop environment. The Web Interface brokers user authentication, enumerates the available desktops and, upon start, delivers a *.ica* file to the Citrix Receiver on the user's local device to start a connection. The Independent Computing Architecture (ICA) file contains configuration information for the Citrix receiver to communicate with the virtual desktop. Because the Web Interface is a critical component, redundant servers must be available to provide fault tolerance.

Delivery controller	The Delivery controller is responsible for maintaining the proper level of idle desktops to allow for instantaneous connections, monitoring the state of online and connected desktops, and shutting down desktops as needed. A XenDesktop farm is a larger grouping of virtual machine servers. Each delivery controller in the XenDesktop acts as an XML server that is responsible for brokering user authentication, resource enumeration, and desktop starting. Because a failure in the XML service results in users being unable to start their desktops, it is recommended that you configure multiple controllers per farm.
PVS and MCS	Provisioning Services (PVS) is used to provision stateless desktops at a large scale. Machine Creation Services (MCS) is used to provision dedicated or stateless desktops in a quick and integrated manner. For more information, see “Citrix XenDesktop provisioning” section on page 17.
License Server	The Citrix License Server is responsible for managing the licenses for all XenDesktop components. XenDesktop has a 30-day grace period that allows the system to function normally for 30 days if the license server becomes unavailable. This grace period offsets the complexity of otherwise building redundancy into the license server.
XenDesktop SQL Server	Each Citrix XenDesktop site requires an SQL Server database that is called the *data store*, which used to centralize farm configuration information and transaction logs. The data store maintains all static and dynamic information about the XenDesktop environment. Because the XenDeskop SQL server is a critical component, redundant servers must be available to provide fault tolerance.
vCenter Server	By using a single console, vCenter Server provides centralized management of the virtual machines (VMs) for the VMware ESXi hypervisor. VMware vCenter can be used to perform live migration (called *VMware vMotion*), which allows a running VM to be moved from one physical server to another without downtime. Redundancy for vCenter Server is achieved through VMware high availability (HA). The vCenter Server also contains a licensing server for VMware ESXi.
vCenter SQL Server	vCenter Server for VMware ESXi hypervisor requires an SQL database. The vCenter SQL server might be Microsoft® Data Engine (MSDE), Oracle, or SQL Server. Because the vCenter SQL server is a critical component, redundant servers must be available to provide fault tolerance. Customer SQL databases (including respective redundancy) can be used.

Client devices	Citrix XenDesktop supports a broad set of devices and all major device operating platforms, including Apple iOS, Google Android, and Google ChromeOS. XenDesktop enables a rich, native experience on each device, including support for gestures and multi-touch features, which customizes the experience based on the type of device. Each client device has a Citrix Receiver, which acts as the agent to communicate with the virtual desktop by using the ICA/HDX protocol.
Thin-client Manager	The Lenovo Thin-client Manager (LTM) is used to manage and support Lenovo thin-client devices individually or in groups.
VDA	Each VM needs a Citrix Virtual Desktop Agent (VDA) to capture desktop data and send it to the Citrix Receiver in the client device. The VDA also emulates keyboard and gestures sent from the receiver. ICA is the Citrix remote display protocol for VDI.
Citrix Receiver	Citrix Receiver is the client software that provides access to applications, desktops and data easily and securely from any device, including smartphones, tablets, PCs and Macs

For more information, see the Lenovo Client Virtualization base reference architecture document that is available at this website: lenovopress.com/tips1275.

3.3 Citrix XenDesktop provisioning

Citrix XenDesktop features the following primary provisioning components for desktops and applications:

- Provisioning Services (PVS)
- Machine Creation Services (MCS)

3.3.1 Provisioning services

Hosted VDI desktops can be deployed with or without Citrix PVS. The advantage of PVS is that you can stream a single desktop image to create multiple virtual desktops on one or more servers in a data center. Figure 17 shows the sequence of operations that are executed by XenDesktop to deliver a hosted VDI virtual desktop.

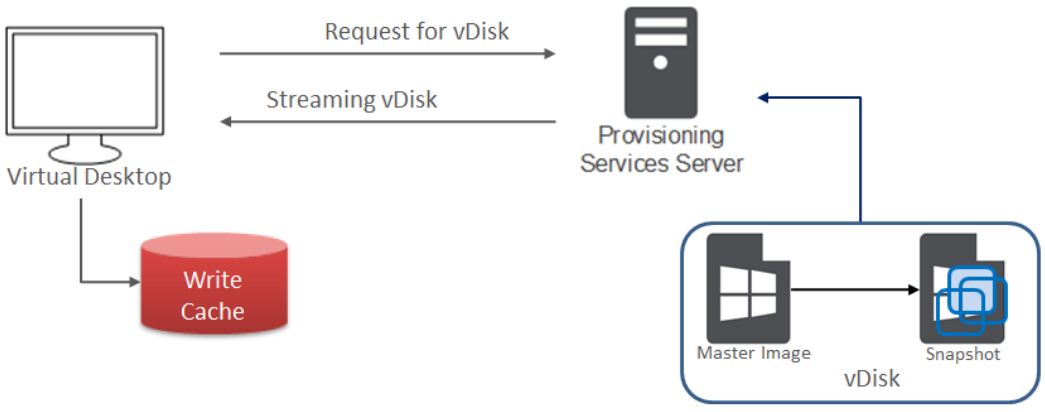

Figure 17: Using PVS for a stateless model

When the virtual disk (vDisk) master image is available from the network, the VM on a target device no longer needs its local hard disk drive (HDD) to operate; it boots directly from the network and behaves as if it were running from a local drive on the target device, which is why PVS is recommended for stateless virtual desktops. PVS often is not used for dedicated virtual desktops because the write cache is not stored on shared storage.

PVS is also used with Microsoft Roaming Profiles (MSRPs) so that the user's profile information can be separated out and reused. Profile data is available from CIFS based shared storage.

It is a best practice to use snapshots for changes to the master VM images and also keep copies as a backup.

3.3.2 Machine creation services

Unlike PVS, MCS does not require more servers. Instead, it uses integrated functionality that is built into the hypervisor and communicates through the APIs. Each desktop has one difference disk and one identity disk (as shown in Figure 18). The difference disk is used to capture any changes that are made to the master image. The identity disk is used to store information, such as device name and password.

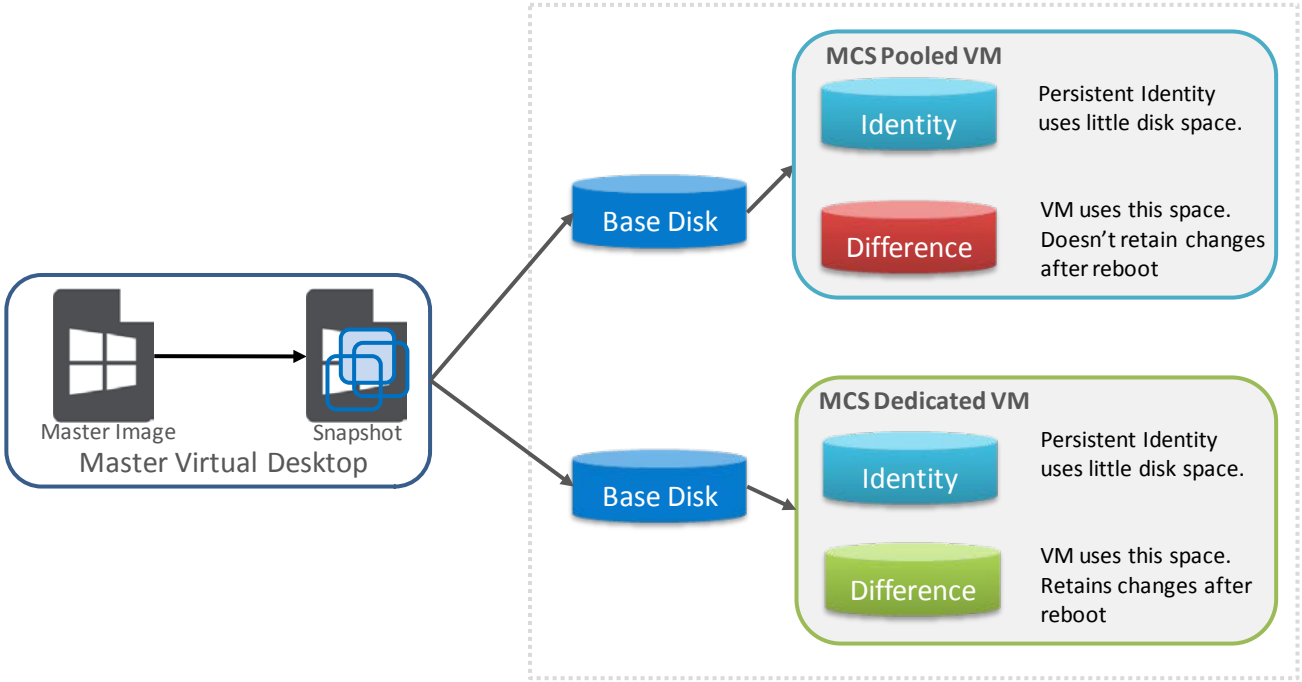

Figure 18: MCS image and difference/identity disk storage model

The following types of image assignment models for MCS are available:

- Pooled-random: Desktops are assigned randomly. When they log off, the desktop is free for another user. When rebooted, any changes that were made are destroyed.
- Pooled-static: Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless if the desktop is rebooted. During reboots, any changes that are made are destroyed.
- Dedicated: Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless if the desktop is rebooted. During reboots, any changes that are made persist across subsequent restarts.

MCS thin provisions each desktop from a master image by using built-in technology to provide each desktop with a unique identity. Only changes that are made to the desktop use more disk space.

3.4 Management VMs

A key part of the Citrix XenDesktop environment is the various management VMs used to manage the VDI infrastructure and user VMs. Table 2 lists the VM requirements and performance characteristics of each management service for Citrix XenDesktop.

Table 2: Characteristics of XenDesktop and ESXi management services

Management service VM	Virtual processors	System memory	Storage	Windows OS	HA needed	Performance characteristic
Delivery controller	4	8 GB	60 GB	2012 R2	Yes	5000 user connections
Storefront	4	4 GB	60 GB	2012 R2	Yes	30,000 connections per hour
Citrix licensing server	2	4 GB	60 GB	2012 R2	No	170 licenses per second
XenDesktop SQL server	2	8 GB	60 GB	2012 R2	Yes	5000 users
PVS servers	4	32 GB	60 GB (depends on number of images)	2012 R2	Yes	Up to 1000 desktops, memory should be a minimum of 2 GB plus 1.5 GB per image served
vCenter server	8	16 GB	60 GB	2012 R2	No	Up to 2000 desktops
vCenter SQL server	4	8 GB	200 GB	2012 R2	Yes	Double the virtual processors and memory for more than 2500 users

These management VMs can be run on separate servers from the HX series cluster or within the cluster itself. Separating out the VMs means that the management VMs can be separately managed and sized to the requirements and dedicated servers used for the user VMs. Putting all of the VMs together in one cluster means that the compute servers will execute less user VMs and additional resources are needed for the much larger and more granular management VMs. Lenovo recommends that the management and user VMs are separated for all but the smallest deployments (i.e. less than 600 users).

Table 3 lists the number of management VMs for each size of users following the recommendations for high availability and performance. The number of vCenter servers is half of the number of vCenter clusters because each vCenter server can handle two clusters of up to 1000 desktops.

Table 3: Management VMs needed

XenDesktop management service VM	**600 users**	**1500 users**	**4500 users**	**10000 users**
Delivery Controllers Includes Citrix Licensing server Includes Web server	2 (1+1) Y Y	2 (1+1) N N	2 (1+1) N N	2 (1+1) N N
Storefront	N/A	2 (1+1)	2 (1+1)	2 (1+1)
Citrix licensing servers	N/A	1	1	1
XenDesktop SQL servers	2 (1+1)	2 (1+1)	2 (1+1)	2 (1+1)
PVS servers for stateless case only	2 (1+1)	4 (2+2)	8 (6+2)	14 (10+4)
ESXi management service VM	**600 users**	**1500 users**	**4500 users**	**10000 users**
vCenter servers	1	1	3	7
vCenter SQL servers	2 (1+1)	2 (1+1)	2 (1+1)	2 (1+1)

It is assumed that common services, such as Microsoft Active Directory, Dynamic Host Configuration Protocol (DHCP), domain name server (DNS), and Microsoft licensing servers exist in the customer environment.

3.5 Performance testing

This section describes the performance benchmarking tool and the results obtained for different configurations of a cluster of 4 Lenovo Converged HX3500 appliances.

3.5.1 Login VSI benchmarking tool

Login VSI is a vendor-independent benchmarking tool that is used to objectively test and measure the performance and scalability of server-based Windows desktop environments. Leading IT analysts recognize and recommend Login VSI as an industry-standard benchmarking tool for client virtualization and can be used by user organizations, system integrators, hosting providers, and testing companies.

Login VSI provides multiple workloads to simulate real user work and suitable in performing load test, benchmarking and capacity planning for VDI environments. Table 4 lists the characteristics of the Login VSI 4.1 workloads that are used in the Lenovo testing.

Table 4. Login VSI Workload Comparison

Workload Name	**Login VSI Version**	**Apps Open**	**CPU Usage**	**Disk Reads**	**Disk Writes**	**IOPS**	**Memory**	**vCPU**
Office worker	4.1	5-8	82%	90%	101%	8.1	1.5GB	1vCPU
Knowledge worker	4.1	5-9	100%	100%	100%	8.5	1.5GB	2vCPU
Power worker	4.1	8-12	119%	133%	123%	10.8	2GB	2vCPU+

The VSImax score parameter (the number indicates user density) is used to determine the performance of a particular system configuration. The following parameters and rules are used for Login VSI tests:

- User login interval: 30 seconds per node
- Workload: Office Worker, Knowledge Worker, or Power User
- All virtual desktops were pre-booted before the tests
- The number of powered-on VMs was adjusted to stay within a 10% margin of VSImax to avoid unreasonable overhead by “idling” virtual machines
- VSImax score is derived using the “classic model” calculation

3.5.2 Performance results for virtual desktops

This section shows the performance results for Lenovo Converged HX3500 compute servers that are based on Intel Xeon E5-2600 V3 processors (Haswell).

The four HX3500 appliances used in the performance testing were configured with dual Xeon E5-2680 v3 or Xeon E5-2699 v3 processors, 512 GB of memory, two 800 GB SSDs, and six 1 TB disk drives.

Table 5 lists the Login VSI performance results of a HX Series appliance 4 node cluster using ESXi and virtual desktops.

Table 5: Virtual desktop performance with ESXi

Node	Processor	Workload	MCS Pooled	MCS Dedicated
4 X HX3500	E5-2680 v3 2.5 GHz 12C	Office worker	842 users	842 users
4 X HX3500	E5-2680 v3 2.5 GHz 12C	Knowledge worker	651 users	617 users
4 X HX3500	E5-2699 v3 2.3 GHz 16C	Knowledge worker	853 users	838 users
4 X HX3500	E5-2680 v3 2.5 GHz 12C	Power Worker	545 users	542 users
4 x HX3500	E5-2699 v3 2.3 GHz 16C	Power worker	826 users	836 users

These results indicate the comparative processor performance. The following conclusions can be drawn:

- The performance for pooled (stateless) and dedicated virtual desktops is similar.
- The Xeon E5-2699 v3 processor has significantly better performance than the Xeon E5-2680 v3 processor for knowledge workers and power workers.

3.5.3 Performance results for hosted shared desktops

This section shows the performance results for Lenovo Converged HX3500 compute servers that are based on Intel Xeon E5-2600 V3 processors (Haswell).

The four HX3500 appliances used in the performance testing were configured with dual Xeon E5-2680 v3 or Xeon E5-2699 v3 processors, 512 GB of memory, two 800 GB SSDs, six 1 TB disk drives. Each appliance had 12 Windows Server 2012 R2 VMs and each VM was configured with 4 vCPUs and 32GB of RAM.

Table 6 lists the Login VSI performance results of a HX Series appliance 4 node cluster using ESXi and hosted shared desktops.

Table 6: Hosted shared desktop performance with ESXi 6.0

Node	Processor	Workload	Hosted Shared
4 x HX3500	E5-2699v3 2.3 GHz 16C	Office worker	886 users
4 x HX3500	E5-2699v3 2.3 GHz 16C	Knowledge worker	820 users
4 x HX3500	E5-2699v3 2.3 GHz 16C	Power worker	724 users

3.6 Performance recommendations

This section provides performance recommendations for both sizing of servers and for best practices.

3.6.1 Sizing recommendations for virtual desktops

The default recommendation for virtual desktops is the Xeon E5-2680 v3 processor and 512 GB of memory because this configuration provides the best coverage for a range of user workloads and user VMs up to 3GB in size. For power users with larger VMs, Lenovo recommends the Xeon E5-2699 v3 processor and up to 768 GB of memory because this gives the best user density.

As noted in the section on performance testing, some results are constrained by I/O performance rather than processor performance. In these cases it is best to be conservative on the number of supported desktops.

Table 7 lists the recommended number of desktops for each workload type per appliance and the CPU utilization for both normal mode and failover mode. Failover mode is when an appliance is offline and users need to be moved to other appliances in the cluster. Lenovo recommends a general failover ratio of 5:1.

Table 7: Sizing recommendations for virtual desktops

Criteria	Office worker	Knowledge worker	Power worker
Processor	2 x E5-2680 v3	2 x E5-2699 v3	2 x E5-2699 v3
Normal mode – desktops per appliance and processor utilization	108 – 75%	120 – 74%	108 – 76%
Failover mode – desktops per appliance and processor utilization	144 – 87%	160 – 84%	144 – 88%
System memory for 2 GB desktops	384 GB	384 GB	384 GB
System memory for 3 GB desktops	512 GB	512 GB	512 GB
System memory for 4 GB desktops	N/A	768 GB	768 GB

3.6.2 Sizing recommendations for hosted shared desktops

The default recommendation for hosted shared desktops is the Xeon E5-2699 v3 processor with 512 GB of memory. Table 7 lists the recommended number of desktops for each workload type per appliance. Lenovo recommends a general failover ratio of 5:1.

Table 8: sizing recommendations for shared hosted desktops

Criteria	Office worker	Knowledge worker	Power worker
Processor	2 x E5-2680 v3	2 x E5-2699 v3	2 x E5-2699 v3
Normal mode – desktops per appliance	150 users	135 users	120 users
Failover mode – desktops per appliance	200 users	180 users	160 users
System memory for 12 VMs (16GB each)	384 GB	384 GB	384 GB

3.6.3 Best practices

The number of desktops that can be run on a specific server depends upon the available system memory and compute power of the processors. For a cost-effective solution, the maximum number of users should be put on each server to balance processor, memory, storage I/O, and networking.

The use of the Aero theme in Microsoft Windows® 7 or other intensive workloads has an effect on the maximum number of virtual desktops that can be supported on each compute server. Windows 10 also requires more processor resources than Windows 7, whereas little difference was observed between 32-bit and 64-bit Windows 7. Although a slower processor can be used and still not exhaust the processor power, it is a good policy to have excess capacity.

Another important consideration for compute servers is system memory. For *stateless users*, the typical range of memory that is required for each desktop is 2 GB - 4 GB. For *dedicated users*, the range of memory for each desktop is 2 GB - 6 GB. In general, power users that require larger memory sizes also require more virtual processors. This reference architecture standardizes on 2 GB per desktop as the minimum requirement of a Windows 7 desktop. The virtual desktop memory should be large enough so that swapping is not needed and vSwap can be disabled.

It is a best practice not to overcommit on memory as swapping to disk can have a severe effect on performance; a better strategy is to give each desktop more memory. Alternatively, a monitoring tool can be run to gather information about existing desktops. The desktop memory size that is required does not necessarily have to match the memory supplied in a desktop machine; it can be larger or smaller.

Lenovo recommends the use of VLANs to partition the network traffic. The following VLANs should be used:

- User (for web protocols, display protocols, and support service protocols)
- Management (for management protocols)
- Storage (for NDFS)

Lenovo recommends to always perform user virtualization, even if users have dedicated desktops. This separation of user-specific data makes it much easier to manage and perform upgrades.

Please refer to the websites below for best practices and optimizations recommended by Citrix:

- Windows 7 Optimization Guide: support.citrix.com/servlet/KbServlet/download/25161-102-665153/XD-Windows7OptimizationGuide.pdf
- Citrix Virtual Desktop Handbook 7.x: support.citrix.com/article/CTX139331

4 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.

4.1 Solution overview

Figure 19 shows the architectural overview of the Microsoft Exchange solution using Lenovo Converged HX7500 appliances. This chapter does not address integrating Exchange with unified messaging solutions and handling edge transport routing and distribution.

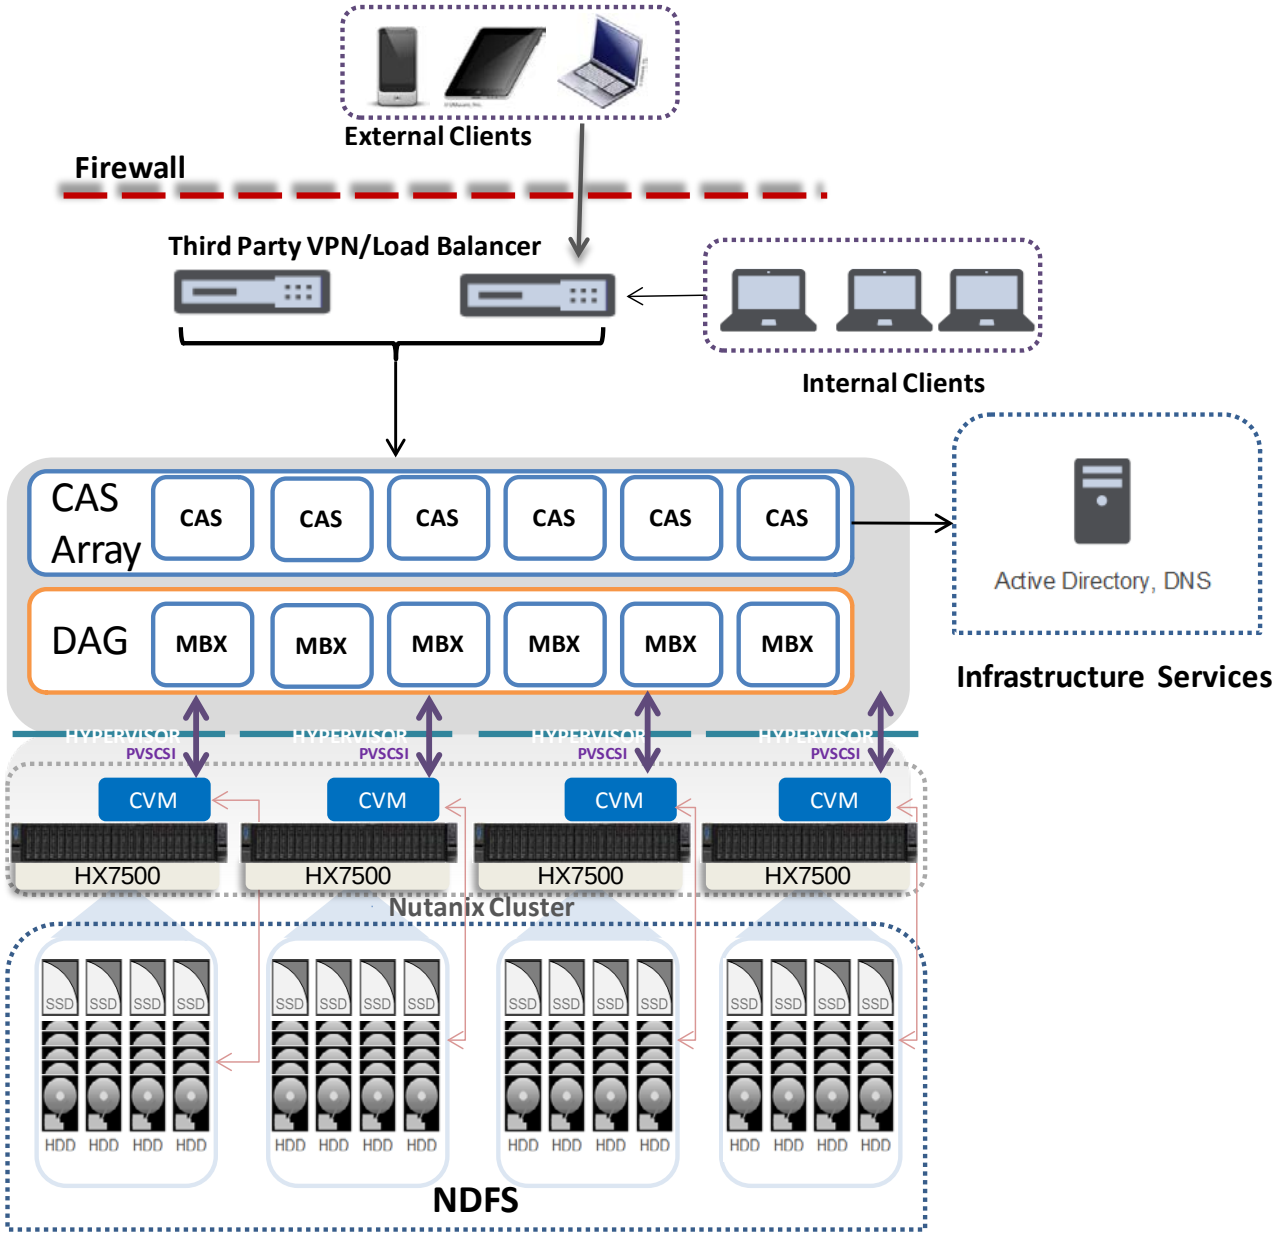

Figure 19. Lenovo Converged HX7500 solution with Microsoft Exchange

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.

4.2 Component model

This section describes the logical component view of the Exchange Server 2013 environment. Figure 20 shows a high-level component model.

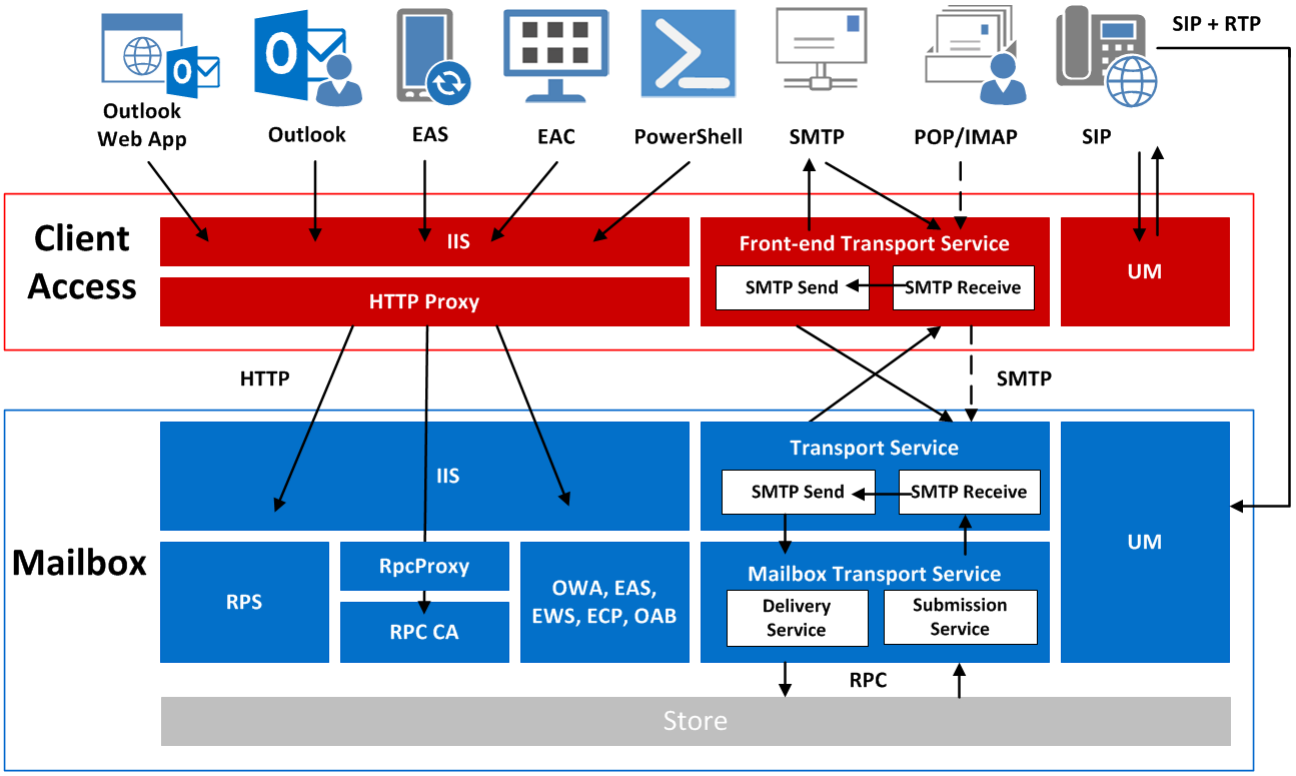

Figure 20. Exchange Server 2013 logical component view

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.

4.3 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.

4.3.1 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.

To maximize the storage performance of Exchange VMs, Lenovo recommends using the ESXi Paravirtual SCSI (PVSCSI) adapters. Each PVSCSI adapter can support up to 15 VMDKs. Lenovo recommends using 4 PVSCSI adapters with one VMDK per Exchange database or log file spread evenly across the PVSCSI adapters.

4.3.2 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 ESXi-based cluster:

- VMware vSphere high availability (HA) for failover
- VMware vSphere distributed resource scheduler (DRS) for load balancing
- Microsoft database availability groups (DAGs) for data redundancy

VMware vSphere HA pools VMs into a cluster to increase data resiliency. If a host fails, VMware HA moves the VMs to other hosts with spare capacity. Lenovo recommends enabling the “Admission Control Setting” and using the “Admission Control Policy” to set the percentage of cluster resources reserved as failover spare capacity.

VMware vSphere DRS can be used to group ESXi hosts into resource clusters to provide highly available resources and balance workloads. In order to keep the active working set for each Exchange VM local to the node, Lenovo recommends creating a host group for each node and a “should” rule that keeps each Exchange VM on a 1 to 1 ratio with the ESXi host. The hosts should be configured with a minimum of N+1 availability. For example a 4 node cluster should have 3 Exchange VMs where the fourth node is available for HA.

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.

4.3.3 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. VM snapshots capture the state of a VM while it is running. This feature enables you to take multiple snapshots of a VM and then revert the VM to any of the previous states by applying a snapshot to the VM. 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.
- Disable Hyper-threading.
- The operating system for an Exchange guest machine must use a disk that has a size equal to at least 15 GB plus the size of the virtual memory that is allocated to the guest machine. This requirement is necessary to account for the operating system and paging file disk requirements. For example, if the guest machine is allocated 16 GB of memory, the minimum disk space that is needed for the guest operating system disk is 31 GB.

4.4 Example deployment for Exchange

This section describes an example deployment for Microsoft Exchange 2013 using a cluster of 4 Lenovo Converged HX7500 appliances. Each appliance has the following configuration:

- 2 x Intel E5-2699v3 (18 cores @ 2.3 GHz) processors
- 512GB RAM
- 4 x 800GB SATA SSDs
- 20 x 1TB SATA HDDs

Each server has sufficient capacity for 4,000 mailboxes assuming 1TB per mailbox database, a data resiliency factor of 2, and a DAG configuration of 2 database copies. This means that up to 12,000 users can be supported on 3 VMs and 4 hosts.

More users or larger mailboxes databases will require 2 TB HDDs. See this website for more information on sizing Exchange databases: blogs.technet.com/b/exchange/archive/2013/05/06/ask-the-perf-guy-sizing-exchange-2013-deployments.aspx.

Each host has ESXi 6.0 U1a and is configured with HA and DRS. The cluster has 3 VMs based on Windows Server 2012 R2 with Microsoft Exchange 2013. Each VM is configured as follows:

- 24 vCPUs
- 128 GB RAM

The CVM is configured as follows:

- 10 vCPUs
- 48 GB RAM

The node configuration of processor and memory is more than sufficient for a single Exchange VM and the remaining capacity can be used for other workloads.

The Lenovo Converged HX7500 appliances use 3 HBAs for the 24 drives. Compared to other systems with less HBAs, the HX7500 has an average IOPS rate that is 23% or better and an average latency improvement of 2 milliseconds.

5 Microsoft SQL Server

Microsoft SQL Server is a database platform for large-scale online transaction processing (OLTP), data warehousing, and a business intelligence platform for data integration, analysis, and reporting solutions. It uses a common set of tools to deploy and manage databases for in-house and cloud environments.

5.1 Solution overview

Figure 21 shows high level architecture of Microsoft SQL Server on Lenovo Converged HX7500 appliances.

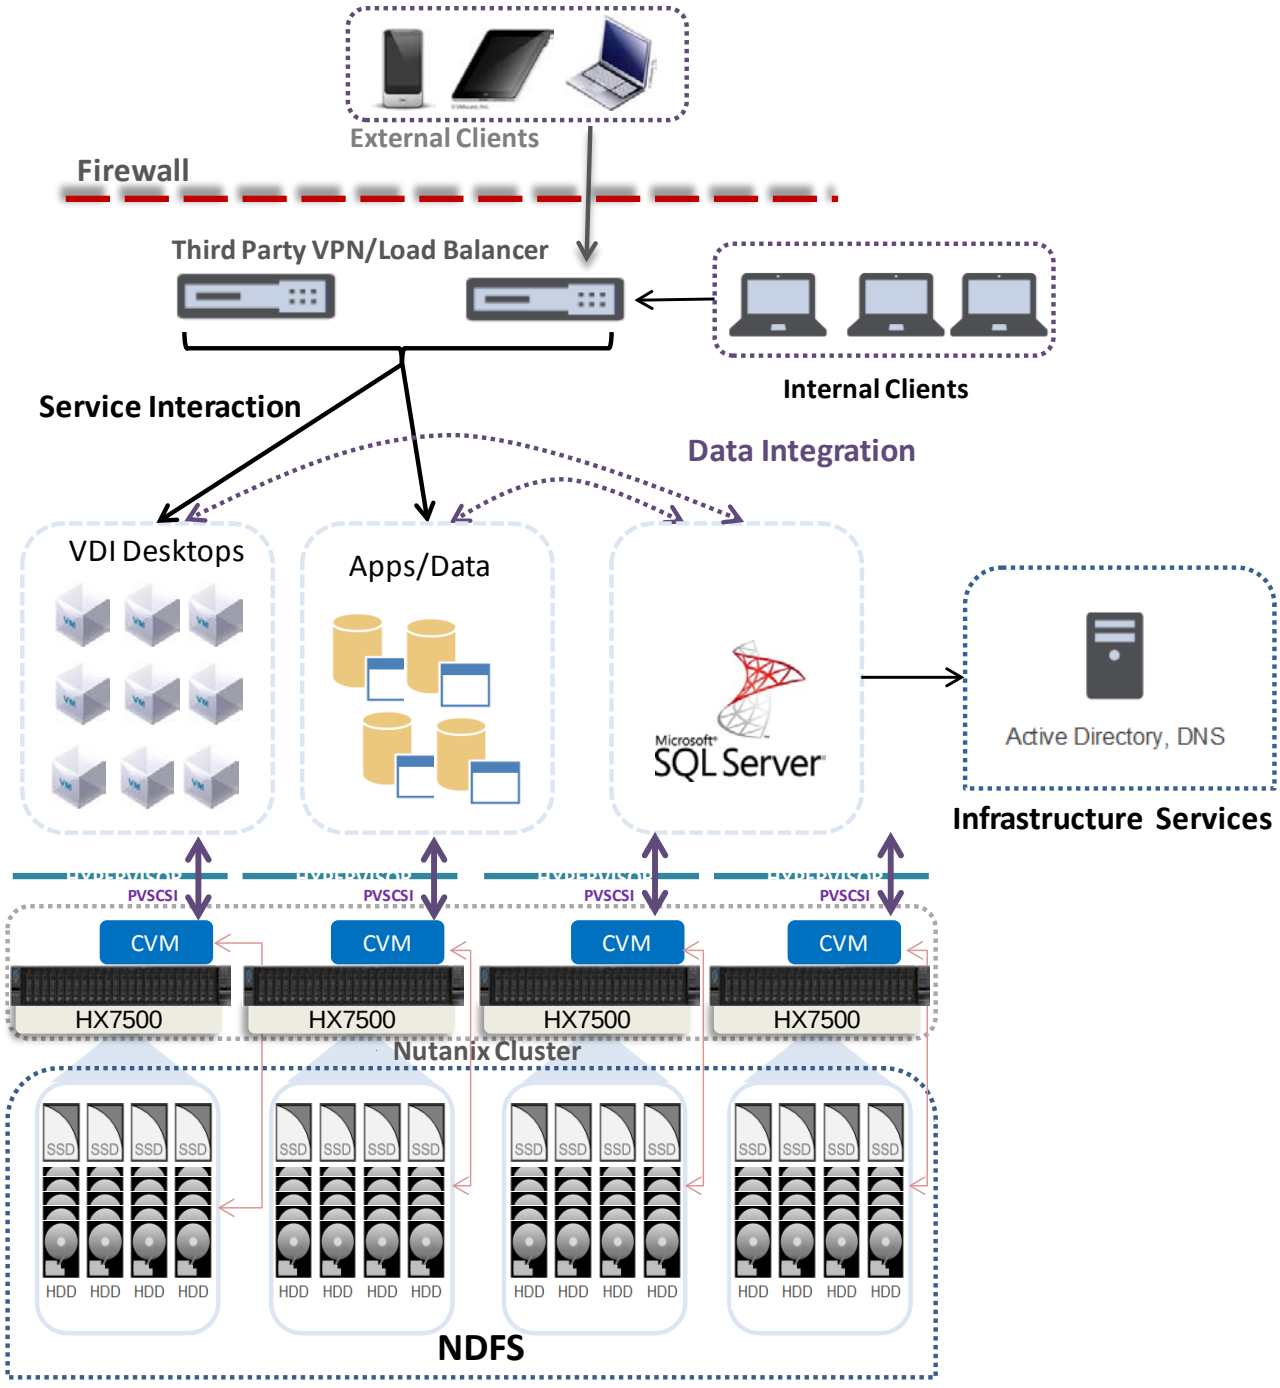

Figure 21: Lenovo Converged HX7500 solution with Microsoft SQL Server

Microsoft SQL Server 2014 can be deployed and can operate in combination with other hosted applications and provides a single scalable platform for all deployments.

5.2 Component model

Figure 22 is a layered component view for Microsoft SQL Server.

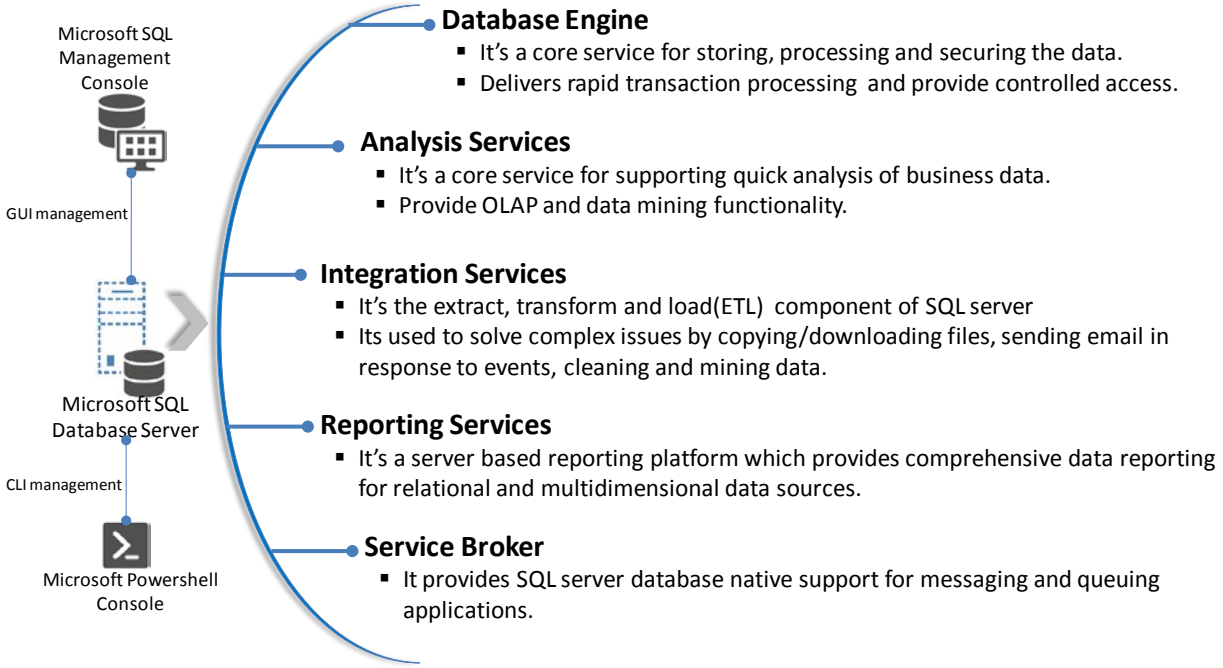

Figure 22: Component model with Microsoft SQL Server

Microsoft SQL Server features the following main components:

Database Engine This part of SQL Server actually creates and drives relational databases.

Analysis Services SQL Server Analysis Services (SSAS) is the data analysis component of SQL Server. It can create OLAP (OnLine Analytical Processing) cubes — sophisticated programming objects for organizing data inside a relational database — and do data mining (pulling relevant data out of a database in response to an ad-hoc question).

Integration Services SQL Server Integration Services (SSIS) performs the extract-transform-load (ETL) process that cleans up and formats raw data from source systems for inclusion in the database as ready-to-use information.

Reporting Services SQL Server Reporting Services (SSRS) provides reporting regardless of a database's operating system.

Service Broker SQL Server Service Broker provides native support for messaging and queuing applications which makes it easier to build distributed and reliable applications that use the Database Engine components.

5.3 SQL Server deployment best practices

This section describes recommended best practices to provide data optimization and high availability of Microsoft SQL Server. See also this website for Nutanix Best Practices: Microsoft SQL server 2014: go.nutanix.com/microsoft-sql-server-converged-infrastructure.html.

5.3.1 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 amount 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 SQL Server 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.

To maximize the storage performance of SQL Server VMs, Lenovo recommends using the ESXi Paravirtual SCSI (PVSCSI) adapters. Each PVSCSI adapter can support up to 15 VMDKs.

5.3.2 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 ESXi-based cluster:

- VMware vSphere high availability (HA) for failover
- VMware vSphere distributed resource scheduler (DRS) for load balancing
- Microsoft AlwaysOn availability groups (AAGs) for data redundancy

VMware vSphere HA pools VMs into a cluster to increase data resiliency. If a host fails, VMware HA moves the VMs to other hosts with spare capacity. Lenovo recommends enabling the “Admission Control Setting” and using the “Admission Control Policy” to set the percentage of cluster resources reserved as failover spare capacity.

VMware vSphere DRS can be used to group ESXi hosts into resource clusters to provide highly available resources and balance workloads. In order to keep the active working set for each SQL Server VM local to the node, Lenovo recommends creating a host group for each node and a “should” rule that keeps each SQL Server VM on a 1 to 1 ratio with the ESXi host. The hosts should be configured with a minimum of N+1 availability.

The Microsoft AlwaysOn availability groups (AAGs) is a recommended high availability mechanism. It uses a shared-nothing approach where transactions are replicated to other nodes so each node has a full copy of the

database and transaction logs. This provides a very fast failover mechanism. The DRS anti-affinity rules need to be used to ensure that the SQL Server VMs are placed on different physical hosts.

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

5.4 Example deployment for SQL Server

This section describes an example deployment for Microsoft SQL Server 2014 using a cluster of 4 Lenovo Converged HX7500 appliances. Each appliance had the following configuration:

- 2 x Intel E5-2699v3 (18 cores @ 2.3 GHz) processors
- 512GB RAM
- 4 x 800GB SATA SSDs
- 20 x 1TB SATA HDDs

Each host has ESXi 6.0 U1a and is configured with HA and DRS. The cluster has 2 VMs based on Windows 2012 R2 and Microsoft SQL Server 2013. Each VM is configured as follows:

- 24 vCPUs
- 128 GB RAM

The CVM is configured as follows:

- 10 vCPUs
- 48 GB RAM

The node configuration of processor and memory is more than sufficient for a single SQL Server VM. The remaining capacity can be used for workloads that use SQL Server or other independent workloads.

The Lenovo Converged HX7500 appliances use 3 HBAs for the 24 drives. Compared to other systems with less HBAs, the HX7500 has an average transaction rate of 1.2 to 2 times better with an average latency of 2 milliseconds.

6 SAP Business Suite

SAP SE provides collaborative business solutions for all types of industries. The company's flagship offering is the SAP® Business Suite®, which includes solutions for enterprise resource planning (ERP), customer relationship management (CRM), product lifecycle management (PLM), supplier relationship management (SRM), and supply chain management (SCM).

The SAP Business Suite offers organizations a choice of integrating business units that use the complete business suite or various modules separately. SAP Business Suite consists of a set of core applications, industry applications, and supplementary applications that are built upon the SAP NetWeaver technology platform. SAP NetWeaver is an open technology platform that unifies technology components in a single platform.

The SAP ERP solution addresses the core business software requirements of the most demanding organizations, in all industries and sectors. The SAP ERP solution offers complete integration for the core business applications of any size organizations. SAP ERP is an analytical tool that helps with business decisions. It also provides the capability to manage financial databases, assets, cost accounting, production operations, and corporate services of an organization.

The SAP ERP solution runs on several platforms with Windows or Linux operating systems that run on the most-used databases of the market, or any possible combination between them. SAP ERP uses the client/server model and can store, retrieve, analyze, and process corporate data in many ways.

To complement the core applications, SAP SE and partners have developed specific software that is based on the requirements of each industry.

Note that at the time of writing, SAP does not support executing SAP HANA on a Lenovo Converged HX Series Nutanix Appliance for production usage. Please contact SAP directly for updates to this statement.

6.1 SAP application architecture and virtualization

An SAP application landscape is a group of two or more SAP systems. An SAP system consists of one or more SAP instances. These SAP instances are often used for

- Development system (DEV)
- Test and Quality Assurance (QA) systems
- Training system (TRN)
- Pre-Production system (PPD)
- Production system (PRD)

Scalability is key to a successful implementation of an ERP solution. As your business grows, you expect this ability to grow with the SAP system. The SAP software is designed to scale as follows:

- Scale up (also known as vertical scaling) with a 2-tier system
- Scale out (also known as horizontal scaling) with a 3-tier system

A 2-tier system configuration has the following key advantages:

- It uses the power of 64-bit technology and the scalability of today's powerful processors in one single system. Provides the best performance, with no overhead for database connection, no network traffic, and no shadow processes.

- The 2-tier approach makes it easy to administer small installations or installations with performance demands that are not dynamic. Larger installations can also use the extra performance that can be achieved from this type of installation.

A 3-tier system configuration offers the following key advantages:

- Larger SAP application implementations use a 3-tier or multitier landscape to allow easier implementation of high availability for each component or tier. A typical 3-tier architecture consists of presentation hardware, at least one application server for load sharing and scalability, and a single database server.
- The database layer is used for storing the SAP application data that is generated by businesses and the SAP application programs. These programs are used by the database from the SAP application servers at run time. An organization has the flexibility to choose the database of its choice including Microsoft SQL Server, Oracle, IBM DB2, SAP MaxDB, SAP Sybase ASE or SAP HANA.

In addition some SAP applications require a 3-tier configuration, for example, Advanced Planner and Optimizer (APO). A 3-tier system can cover short-time load peaks easily by adding temporarily application servers.

The flexibility of the SAP software enables you to decide the best tier approach; whether that is a two-, three-, or a multitier architecture.

Traditionally the SAP application tiers are executed on physical servers. This is wasteful of resources and virtualization can be applied to consolidate servers and reduce CAPEX and OPEX. Virtualization of the SAP application tiers provide all the standard benefits including:

- Flexibility
- Load shifting
- Optimized application availability
- Rapid provisioning

Lenovo has conducted benchmarks with virtualization enabled to show the performance differences between a virtualized and non-virtualized SAP system. These benchmarks run a single guest on the physical server, by which the performance penalty for the introduction of the virtualization layer into the stack is shown. These benchmarks indicate that virtualization introduces a performance degradation of less than 10% for VMware ESXi Server. Considering the low accuracy that can be achieved with the usual sizing methodologies, this degradation is small enough to simply ignore in SAP solution sizings.

6.2 SAP Business Suite deployment best practices

This section describes recommended best practices for virtualizing SAP Business Suite. See also this website for the “SAP on VMware Best Practices”: **(**http://www.vmware.com/files/pdf/business-critical-apps/sap-on-vmware-best-practices.pdf).

SAP is certified for the Lenovo HX Series appliances. See the following website for more details: http://shop.lenovo.com/us/en/systems/solutions/alliances/sap/#tab-sap_certification.

SAP supports the ESXi hypervisor and ESXi 6.0 U1b or greater is recommended. See the following SAP notes when deploying SAP NetWeaver applications on VMware ESXi:

- Note 1492000: General Support Statement for Virtual Environments

- Note 2161991: VMware vSphere Configuration Guidelines
- Note 1501701: Single Computing Unit Performance and Sizing
- Note 1612283, Sizing SAP VMs within a NUMA node when configuring SAP application VMs

For more information about support for Windows and Linux that run on VMware and the databases that are supported, see this website: http://scn.sap.com/docs/DOC-27384.

SAP supports the following Linux guest VM operating systems to execute an SAP application tier:

- Red Hat Enterprise Linux 5
- Red Hat Enterprise Linux 6
- Red Hat Enterprise Linux 7
- SUSE Linux Enterprise Server 10
- SUSE Linux Enterprise Server 11
- SUSE Linux Enterprise Server 12

See the following SAP notes when deploying SAP NetWeaver applications on a Linux operating system:

- Note 1122387: Linux: SAP Support in Virtualized Environments
- Note 171356: SAP Software on Linux: General Information
- Note 1552925: Linux: High Availability Cluster Solutions

SAP supports the following Windows guest VM operating systems to execute an SAP application tier:

- Windows Server 2003
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

See the following SAP notes when deploying SAP NetWeaver applications on a Windows operating system:

- Note 1409608: Virtualization on Windows
- Note 1612283: Hardware Configuration Standards and Guidance
- Note 1374671: High Availability in Virtual Environment on Windows
- Note 1580509: Windows Editions Supported by SAP

6.3 Example deployment for SAP Business Suite

Choosing the correct configuration for a complex solution such as SAP is challenging. SAP systems are normally sized for 65% of the total capacity that is required.

Because I/O latencies can constrain the performance of SAP applications, Lenovo recommends the Lenovo Converged HX7500 for its consistent I/O throughput and low latency.

This section describes an example deployment for SAP Business Suite using a cluster of 4 Lenovo Converged HX7500 appliances. Each appliance had the following configuration:

- 2 x Intel E5-2699v3 (18 cores @ 2.3 GHz) processors
- 512GB RAM
- 4 x 800GB SATA SSDs
- 20 x 1TB SATA HDDs

Each host has ESXi 6.0 U1b and is configured with HA and DRS. The CVM is configured as follows:

- 4 cores
- 48 GB RAM

Each SAP virtual machine is configured individually as required by the SAP application. The node configuration of processor and memory sufficient for a 2-tier SAP application using a SQL Database in the same VM. The remaining capacity can be used other independent workloads or as an option for adoption to changing SAP workloads. This setup can be used for 3-tier implementation by separating the SAP Database from the SAP application server in a dedicated VM.

6.4 Lenovo sizing and planning questionnaire for SAP solutions

The Lenovo and SAP sizing methodology is based on SAP benchmarks, information from SAP, and actual customer experiences. Lenovo uses sizing tools and customer input to approximate the system resource requirements; however, actual customer results can vary. A sizing questionnaire is used to gather the requirements that are used to estimate the hardware resources to run SAP Business Suite on Lenovo servers.

This questionnaire was designed so that it can be answered without detailed knowledge of any SAP application. Lenovo established the Lenovo Expert Technical Sales (LETS) team to assist with queries and help size and configure a target SAP application. The LETS team was established on a geographical basis to address your questions on a local basis. Contact your local LETS team for information and questionnaires. The LETS team can be contacted at the addresses that are listed here:

- Australia/New Zealand LETSanz@lenovo.com
- ASEAN LETSasean@lenovo.com
- EMEA LETSemea@lenovo.com
- GCG LETSgcg@lenovo.com
- ISA LETSisa@lenovo.com
- Japan LETSjapan@lenovo.com
- Korea LETSkorea@lenovo.com
- LA Portuguese LETSlap@lenovo.com
- LA Spanish LETSlas@lenovo.com
- North America LETSsna@lenovo.com

A sizing estimate is an approximation of the hardware resources (processor, memory, and disk) that are required to support an SAP application or component implementation. This estimate provides an introductory understanding of the customer’s initial hardware requirements based on the information given. Your actual experiences vary from the sizing estimate for many reasons, including batch and reporting workloads and custom code. The degree of variability can range from small to significant.

Sizing the hardware requirements for each SAP implementation is an iterative process which can be refined and repeated several times. One should always check the sizing input data and estimations during the implementation project. It is important to understand that the sizing estimate is a pre-installation effort; and, therefore, only based upon standard assumptions and benchmark performance data. It cannot and should not replace capacity planning for installed systems.

7 VMware Horizon

Horizon View is a desktop virtualization product developed by VMware Inc. It provides remote desktop capabilities by using VMware virtualization technology and can deliver a consistent user experience across devices and locations while keeping corporate data secure and compliant. See this website for more details: vmware.com/products/horizon-view.

7.1 Solution overview

Figure 23 shows all of the main features of the Lenovo Hyper-converged Nutanix solution's reference architecture with VMware Horizon on VMware ESXi hypervisor. This chapter does not address the general issues of multi-site deployment and network management and limits the description to the components that are inside the customer's intranet.

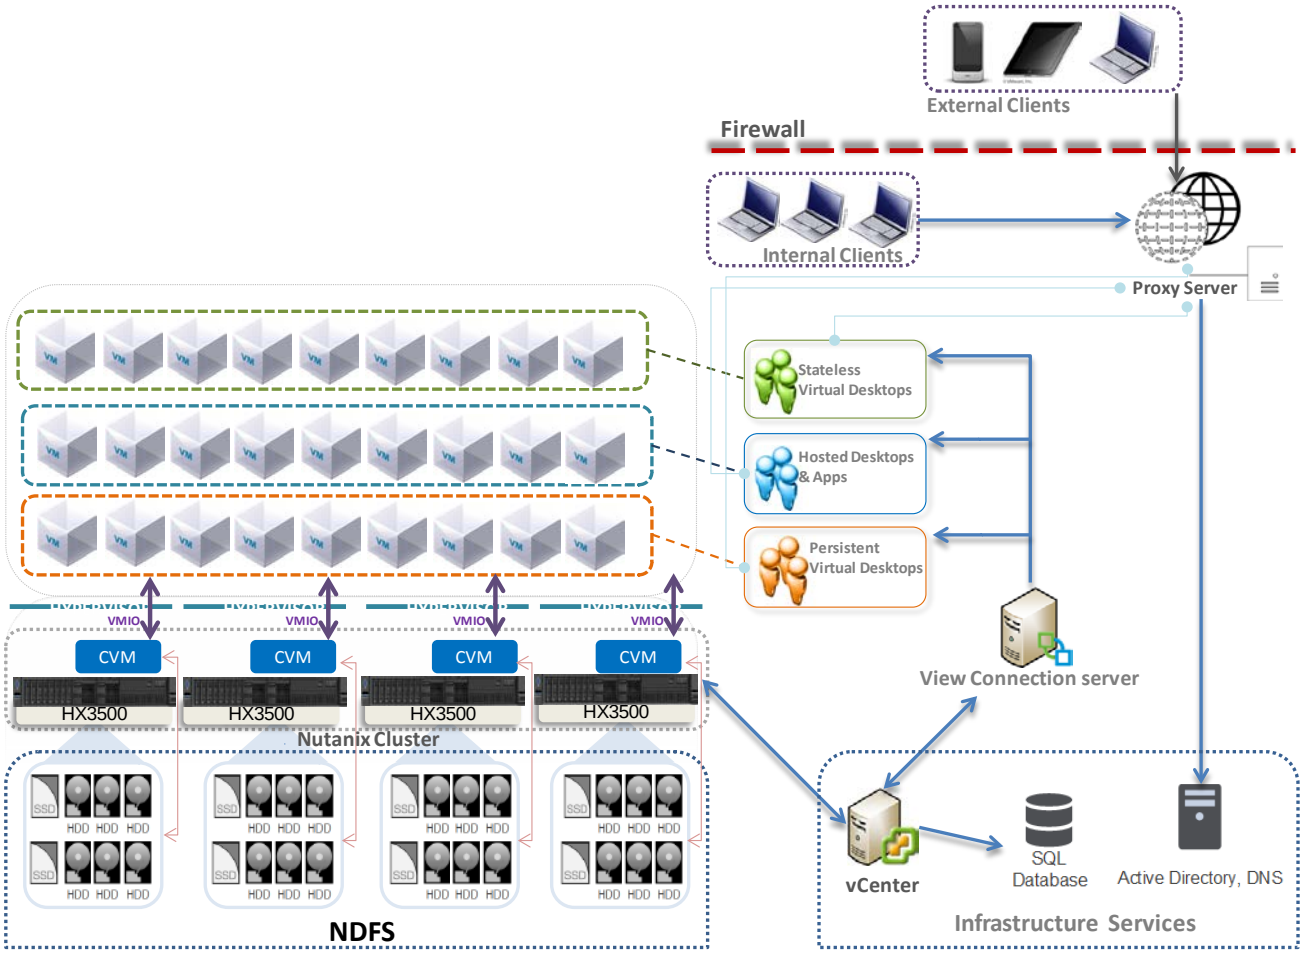

Figure 23: Lenovo Hyper-converged Nutanix solution with VMware View

7.2 Component model

Figure 24 is a layered component view for the VMware Horizon virtualization infrastructure.

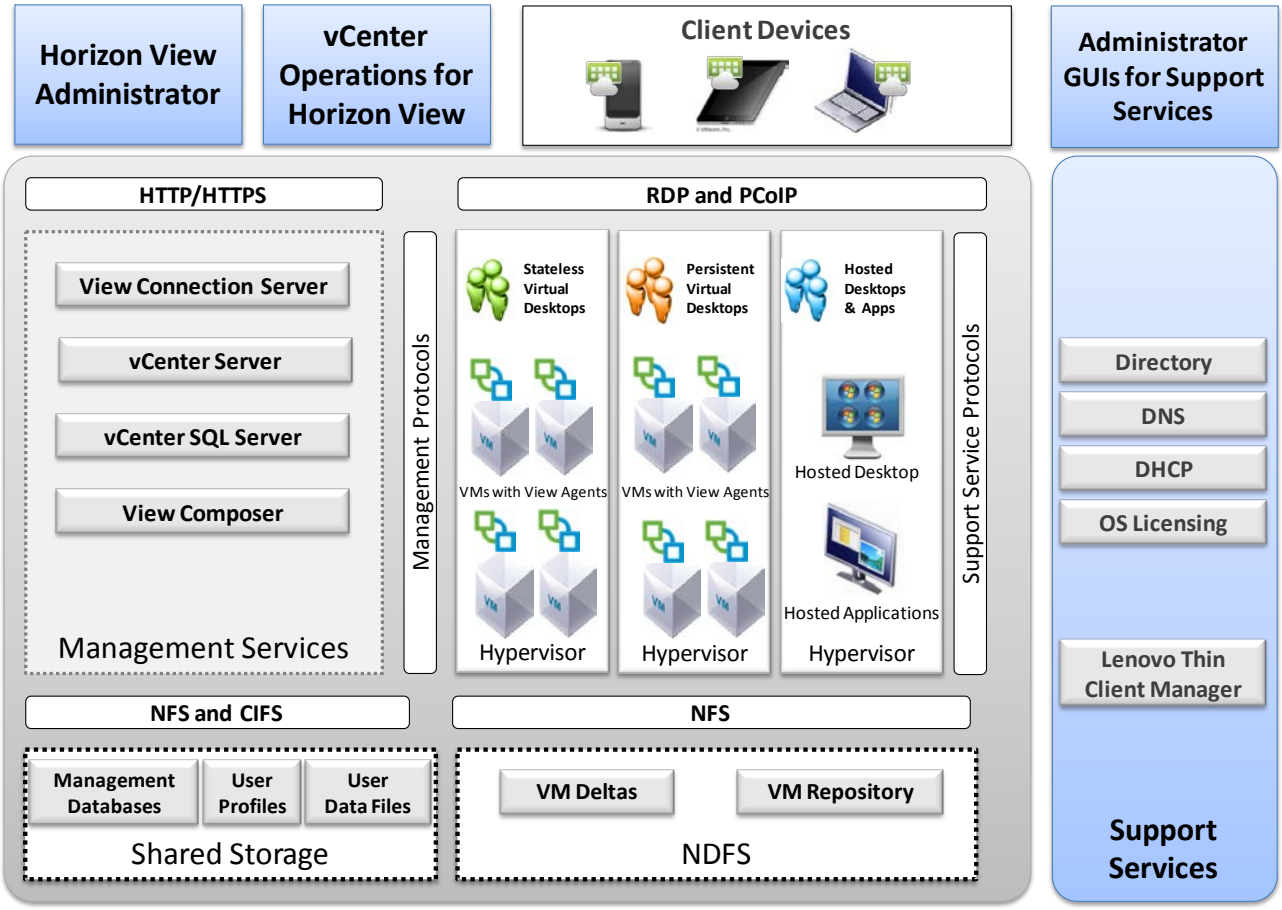

Figure 24: Component model with VMware Horizon

VMware Horizon with the VMware ESXi hypervisor features the following main components:

Horizon View Administrator — By using this web-based application, administrators can configure ViewConnection Server, deploy and manage View desktops, control user authentication, and troubleshoot user issues. It is installed during the installation of ViewConnection Server instances and is not required to be installed on local (administrator) devices.

vCenter Operations for Horizon View — This tool provides end-to-end visibility into the health, performance, and efficiency of the virtual desktop infrastructure (VDI) configuration. It enables administrators to proactively ensure the best user experience possible, avert incidents, and eliminate bottlenecks before they become larger issues.

View Connection Server	The VMware Horizon Connection Server is the point of contact for client devices that are requesting virtual desktops. It authenticates users and directs the virtual desktop request to the appropriate virtual machine (VM) or desktop, which ensures that only valid users are allowed access. After the authentication is complete, users are directed to their assigned VM or desktop. If a virtual desktop is unavailable, the View Connection Server works with the management and the provisioning layer to have the VM ready and available.
View Composer	In a VMware vCenter Server instance, View Composer is installed. View Composer is required when linked clones are created from a parent VM.
vCenter Server	By using a single console, vCenter Server provides centralized management of the virtual machines (VMs) for the VMware ESXi hypervisor. VMware vCenter can be used to perform live migration (called *VMware vMotion*), which allows a running VM to be moved from one physical server to another without downtime. Redundancy for vCenter Server is achieved through VMware high availability (HA). The vCenter Server also contains a licensing server for VMware ESXi.
vCenter SQL Server	vCenter Server for VMware ESXi hypervisor requires an SQL database. The vCenter SQL server might be Microsoft® Data Engine (MSDE), Oracle, or SQL Server. Because the vCenter SQL server is a critical component, redundant servers must be available to provide fault tolerance. Customer SQL databases (including respective redundancy) can be used.
View Event database	VMware Horizon can be configured to record events and their details into a Microsoft SQL Server or Oracle database. Business intelligence (BI) reporting engines can be used to analyse this database.
Clients	VMware Horizon supports a broad set of devices and all major device operating platforms, including Apple iOS, Google Android, and Google ChromeOS. Each client device has a VMware View Client, which acts as the agent to communicate with the virtual desktop.
Thin-client Manager	The Lenovo Thin-client Manager (LTM) is used to manage and support Lenovo thin-client devices individually or in groups.
RDP, PCoIP	The virtual desktop image is streamed to the user access device by using the display protocol. Depending on the solution, the choice of protocols available are Remote Desktop Protocol (RDP) and PC over IP (PCoIP).

For more information, see the Lenovo Client Virtualization base reference architecture document that is available at this website: lenovopress.com/tips1275.

7.3 VMware Horizon provisioning

VMware Horizon supports stateless and persistent virtual desktop models. Provisioning for VMware Horizon is a function of vCenter server and View Composer for linked clones.

vCenter Server allows for manually created pools and automatic pools. It allows for provisioning full clones and linked clones of a parent image for dedicated and stateless virtual desktops.

Because persistent virtual desktops use large amounts of storage, linked clones can be used to reduce the storage requirements. Linked clones are created from a snapshot (replica) that is taken from a golden master image. One pool can contain up to 1000 linked clones.

This document describes the use of automated pools (with linked clones) for dedicated and stateless virtual desktops. The deployment requirements for full clones using Nutanix de-duplication functionality is beyond the scope of this document.

7.4 Management VMs

A key part of the VMware Horizon environment is the various management VMs used to manage the VDI infrastructure and user VMs. Table 9 lists the VM requirements and performance characteristics of each management service.

Table 9: Characteristics of VMware Horizon management services

Management service VM	Virtual processors	System memory	Storage	Windows OS	HA needed	Performance characteristic
vCenter Server	8	12 GB	60 GB	2012 R2	No	Up to 2000 VMs.
vCenter SQL Server	4	8 GB	200 GB	2012 R2	Yes	Double the virtual processors and memory for more than 2500 users.
View Connection Server	4	16 GB	60 GB	2012 R2	Yes	Up to 2000 connections.

These management VMs can be run on separate servers from the HX series cluster or within the cluster itself. Separating out the VMs means that the management VMs can be separately managed and sized to the requirements and dedicated servers used for the user VMs. Putting all of the VMs together in one cluster means that the compute servers will execute less user VMs and need to leave enough resources for the much larger and more granular management VMs. Lenovo recommends that the management and user VMs are separated for all but the smallest deployments (i.e. less than 600 users).

Table 3 lists the number of management VMs for each size of users following the requirements for high-availability and performance. The number of vCenter servers is half of the number of vCenter clusters because each vCenter server can handle two clusters of up to 1000 desktops.

Table 10: Management VMs needed

Horizon management service VM	600 users	1500 users	4500 users	10000 users
vCenter servers	1	1	3	7
vCenter SQL servers	2 (1+1)	2 (1+1)	2 (1+1)	2 (1+1)
View Connection Server	2 (1+1)	2 (1+ 1)	4 (3+1)	7 (5+2)

It is assumed that common services, such as Microsoft Active Directory, Dynamic Host Configuration Protocol (DHCP), domain name server (DNS), and Microsoft licensing servers exist in the customer environment.

7.5 Performance testing

This section describes the performance benchmarking tool and the results obtained for different configurations of a cluster of 4 Lenovo Converged HX3500 appliances.

7.5.1 Login VSI benchmarking tool

Login VSI is a vendor-independent benchmarking tool that is used to objectively test and measure the performance and scalability of server-based Windows desktop environments. Leading IT analysts recognize and recommend Login VSI as an industry-standard benchmarking tool for client virtualization and can be used by user organizations, system integrators, hosting providers, and testing companies.

Login VSI provides multiple workloads to simulate real user work and suitable in performing load test, benchmarking and capacity planning for VDI environments. Table 4 lists the characteristics of the Login VSI 4.1 workloads that are used in the Lenovo testing.

Table 11. Login VSI Workload Comparison

Workload Name	Login VSI Version	Apps Open	CPU Usage	Disk Reads	Disk Writes	IOPS	Memory	vCPU
Office worker	4.1	5-8	82%	90%	101%	8.1	1.5GB	1vCPU
Knowledge worker	4.1	5-9	100%	100%	100%	8.5	1.5GB	2vCPU
Power worker	4.1	8-12	119%	133%	123%	10.8	2GB	2vCPU+

The VSImax score parameter (the number indicates user density) is used to determine the performance of a particular system configuration. The following parameters and rules are used for Login VSI tests:

- User login interval: 30 seconds per node
- Workload: Office Worker, Knowledge Worker, or Power User
- All virtual desktops were pre-booted before the tests
- The number of powered-on VMs was adjusted to stay within a 10% margin of VSImax to avoid unreasonable overhead by “idling” virtual machines
- VSImax score is derived using the “classic model” calculation

7.5.2 Performance results for virtual desktops

This section shows the performance results for Lenovo Converged HX3500 compute servers that are based on Intel Xeon E5-2600 V3 processors (Haswell).

The four HX3500 appliances used in the performance testing were configured with dual Xeon E5-2680 v3 or Xeon E5-2699 v3 processors, 512 GB of memory, two 800 GB SSDs, and six 1 TB disk drives.

Table 12 lists the Login VSI performance results of a HX Series appliance 4 node cluster using ESXi and virtual desktops.

Table 12: Virtual Desktop performance with ESXi 6.0

Node	Processor	Workload	Stateless Desktop	Persistent Desktop
4 X HX3500	E5-2680 v3 2.5 GHz 24C	Office worker	768 users	767 users
4 X HX3500	E5-2680 v3 2.5 GHz 24C	Knowledge worker	561 users	551 users
4 X HX3500	E5-2699 v3 2.3 GHz 32C	Knowledge worker	782 users	779 users
4 X HX3500	E5-2680 v3 2.5 GHz 24C	Power Worker	499 users	482 users
4 x HX3500	E5-2699 v3 2.3 GHz 32C	Power worker	725 users	714 users

These results indicate the comparative processor performance. The following conclusions can be drawn:

- The performance for stateless and dedicated virtual desktops is similar.
- The Xeon E5-2699 v3 processor has significantly better performance than the Xeon E5-2680 v3 processor for knowledge workers and power workers.

7.6 Performance recommendations

This section provides performance recommendations for both sizing of servers and for best practices.

7.6.1 Sizing recommendations for virtual desktops

The default recommendation for virtual desktops is the Xeon E5-2680 v3 processor and 512 GB of memory because this configuration provides the best coverage for a range of user workloads and user VMs up to 3GB in size. For power users with larger VMs, Lenovo recommends the Xeon E5-2699 v3 processor and up to 768 GB of memory because this gives the best user density.

As noted in the section on performance testing, some results are constrained by I/O performance rather than processor performance. In these cases it is best to be conservative on the number of supported desktops.

Table 7 lists the recommended number of desktops for each workload type per appliance and the CPU utilization for both normal mode and failover mode. Failover mode is when an appliance is offline and users need to be moved to other appliances in the cluster. Lenovo recommends a general failover ratio of 5:1.

Table 13: Sizing recommendations for virtual desktops

Criteria	Office worker	Knowledge worker	Power worker
Processor	2 x E5-2680 v3	2 x E5-2699 v3	2 x E5-2699 v3
Normal mode – desktops per appliance and processor utilization	102 – 81%	120 – 79%	102 – 76%
Failover mode – desktops per appliance and processor utilization	136 – 88%	160 – 89%	136 – 89%
System memory for 2 GB desktops	384 GB	384 GB	384 GB
System memory for 3 GB desktops	512 GB	512 GB	512 GB
System memory for 4 GB desktops	N/A	768 GB	768 GB

7.6.2 Best practices

The number of desktops that can be run on a specific server depends upon the available system memory and compute power of the processors. For a cost-effective solution, the maximum number of users should be put on each server to balance processor, memory, storage I/O, and networking.

The use of the Aero theme in Microsoft Windows® 7 or other intensive workloads has an effect on the maximum number of virtual desktops that can be supported on each compute server. Windows 10 also requires more processor resources than Windows 7, whereas little difference was observed between 32-bit and 64-bit Windows 7. Although a slower processor can be used and still not exhaust the processor power, it is a good policy to have excess capacity.

Another important consideration for compute servers is system memory. For *stateless users*, the typical range of memory that is required for each desktop is 2 GB - 4 GB. For *dedicated users*, the range of memory for each desktop is 2 GB - 6 GB. In general, power users that require larger memory sizes also require more virtual processors. This reference architecture standardizes on 2 GB per desktop as the minimum requirement of a Windows 7 desktop. The virtual desktop memory should be large enough so that swapping is not needed and vSwap can be disabled.

It is a best practice not to overcommit on memory as swapping to disk can have a severe effect on performance; a better strategy is to give each desktop more memory. Alternatively, a monitoring tool can be run to gather information about existing desktops. The desktop memory size that is required does not necessarily have to match the memory supplied in a desktop machine; it can be larger or smaller.

Lenovo recommends the use of VLANs to partition the network traffic. The following VLANs should be used:

- User (for web protocols, display protocols, and support service protocols)
- Management (for management protocols)
- Storage (for NDFS)

Lenovo recommends to always perform user virtualization, even if users have dedicated desktops. This separation of user-specific data makes it much easier to manage and perform upgrades.

Please refer below links for best practices and optimizations recommended by VMware:

- View Architecture Planning – VMware Horizon 6.0:

  pubs.vmware.com/horizon-view-60/topic/com.vmware.ICbase/PDF/horizon-view-60-architecture-planning.pdf

- VMware Horizon 6 with View Performance and Best Practices:

  vmware.com/files/pdf/view/vmware-horizon-view-best-practices-performance-study.pdf

- VMware View Optimization Guide for Desktops and Servers:

  vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf

Resources

- Lenovo Converged HX Series Product Guide
  lenovopress.com/lp0059
- Lenovo Converged HX Series landing page
  shop.lenovo.com/us/en/systems/converged-systems/hx_series
- Nutanix Portal (requires registration)
  portal.nutanix.com
- Nutanix Bible
  nutanixbible.com/
- Nutanix Tech Note: VMware vSphere Networking on Nutanix
  go.nutanix.com/rs/nutanix/images/Nutanix_TechNote-VMware_vSphere_Networking_with_Nutanix.pdf
- VMware vSphere
  vmware.com/products/datacenter-virtualization/vsphere
- Lenovo System x solutions for SAP environments
  http://www.lenovo.com/solutions/sap
- SAP on Linux
  http://scn.sap.com/community/linux
- SAP on Microsoft Windows
  http://scn.sap.com/community/windows

Document history

Version 1.0	31 March 2015	• Initial version
Version 1.1	5 August 2015	• Added section for SAP Business Suite

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, 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.