Reference Architecture: Lenovo Client Virtualization with VMware Horizon on ThinkSystem and ThinkAgile VX Servers

Last update: 05 March 2024
Version 1.11

Reference Architecture for VMware Horizon (with View)

Contains performance data and sizing recommendations for servers, storage, and networking

Describes variety of storage models including SAN storage and hyper-converged vSAN

Contains detailed bill of materials for servers, storage and networking

Chandrakandh Mouleeswaran

Cristian Ghetau

Vinay Kulkarni

LENOVO PRESS

Table of Contents

1 Introduction

The intended audience for this document is technical IT architects, system administrators, and managers who are interested in server-based desktop virtualization and server-based computing (terminal services or application virtualization) that uses VMware Horizon™ (with View). In this document, the term client virtualization is used as to refer to all of these variations. Compare this term to server virtualization, which refers to the virtualization of server-based business logic and databases.

This document gives an architecture overview and logical component model of VMware Horizon. The document also provides the operational model of VMware Horizon by combining Lenovo® hardware platforms such as ThinkSystem Servers, ThinkAgile VX Servers and RackSwitch networking with storage such as ThinkSystem DE Series and DM Series storage or VMware vSAN. The operational model presents performance benchmark measurements and discussion, sizing guidance, and some example deployment models. The last section contains detailed bill of material configurations for each piece of hardware.

See also the “Reference Architecture for Workloads using the Lenovo ThinkAgile HX Series appliances” for VMware Horizon specific information. This document is available at this website: lenovopress.com/lp0665.

2 Architectural overview

Figure 1 shows all of the main features of the Lenovo Client Virtualization reference architecture with VMware Horizon on VMware ESXi hypervisor. It also shows remote access, authorization, and traffic monitoring. This reference architecture does not address the general issues of multi-site deployment and network management.

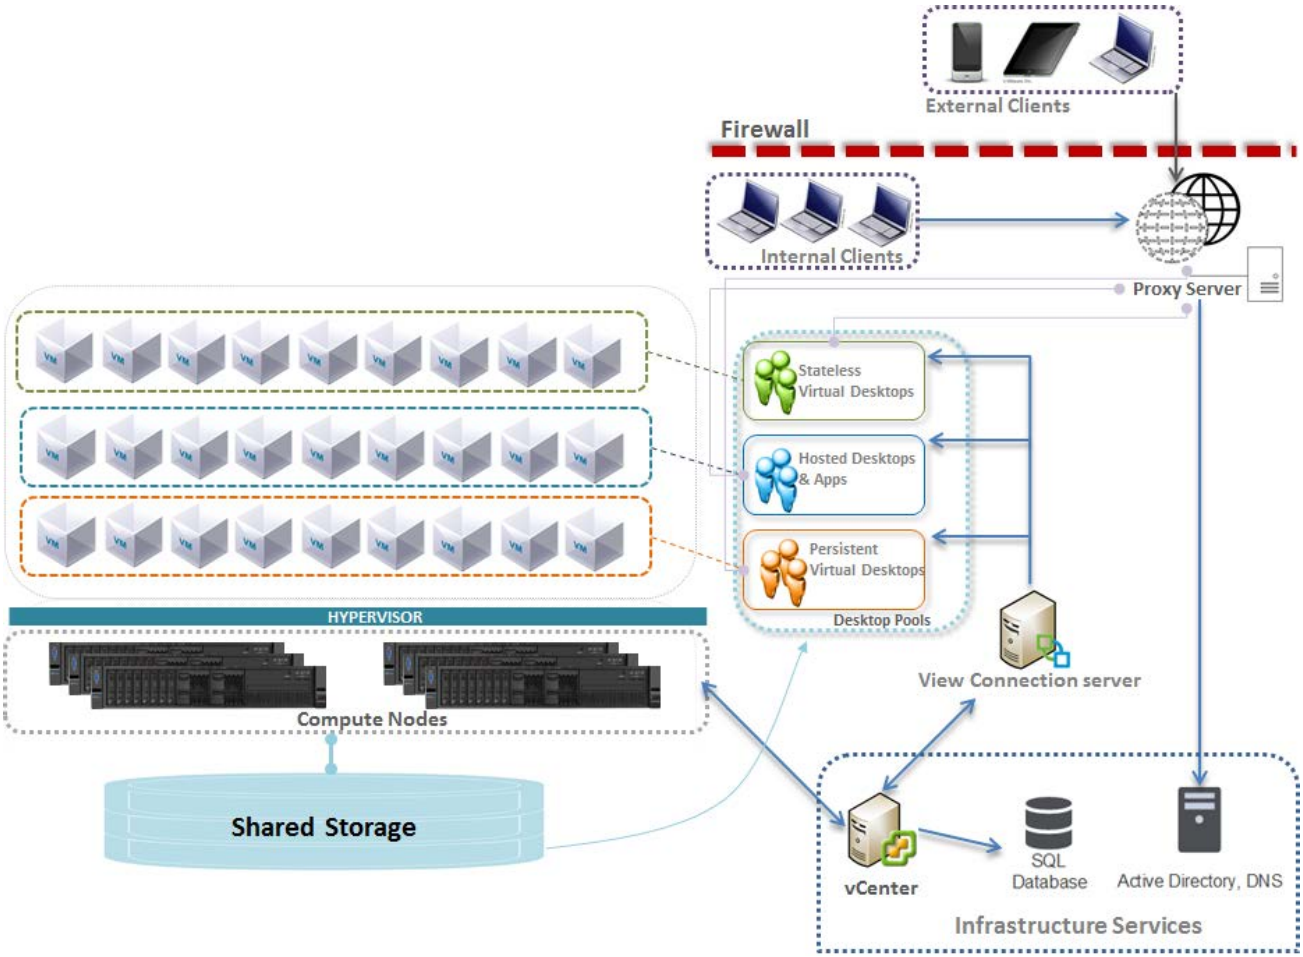

Figure 1: LCV reference architecture with VMware Horizon

3 Component model

Figure 2 is a layered view of the LCV solution that is mapped to the VMware Horizon virtualization infrastructure.

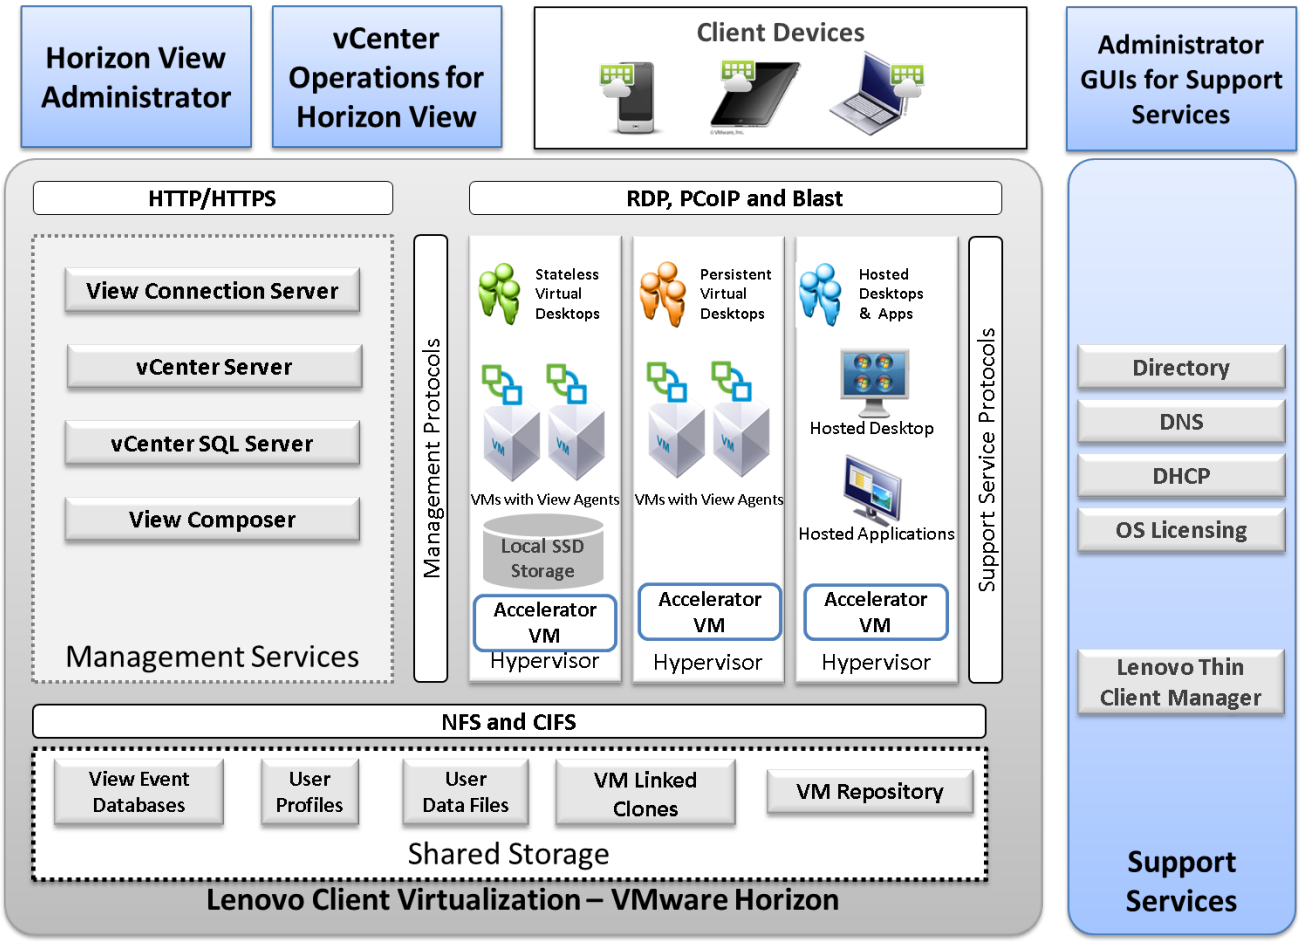

Figure 2: Component model with VMware Horizon

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 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 analyze 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.
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).
Hypervisor ESXi	ESXi is a bare-metal hypervisor for the compute servers. ESXi also contains support for vSAN storage. For more information, see “VMware vSAN” on page 6.
Accelerator VM	The optional accelerator VM.
Shared storage	Shared storage is used to store user profiles and user data files. Depending on the provisioning model that is used, different data is stored for VM images. For more information, see “Storage model”. Shared storage can also be distributed storage as used in hyper-converged systems.

For more information, see the Lenovo Client Virtualization base reference architecture document that is available at this website: lenovopress.com/lp0756.

3.1 VMware Horizon provisioning

VMware Horizon supports stateless and dedicated 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, linked clones of a parent image and Instant clone of running parent virtual machine for dedicated and stateless virtual desktops. Refer here for more details on Instant Clone technology https://blogs.VMware.com/euc/2016/02/horizon-7-view-instant-clone-technology-linked-clone-just-in-time-desktop.html

Because dedicated 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. The golden master image and replica should be on shared storage area network (SAN) storage. One pool can contain up to 2000 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 are beyond the scope of this document.

3.2 Storage model

This section describes the different types of shared or distributed data stored for stateless and dedicated desktops. Stateless and dedicated virtual desktops should have the following common shared storage items:

- The paging file (or vSwap) is transient data that can be redirected to Network File System (NFS) storage. In general, it is recommended to disable swapping, which reduces storage use (shared or local). The desktop memory size should be chosen to match the user workload rather than depending on a smaller image and swapping, which reduces overall desktop performance.
- User profiles (from Microsoft Roaming Profiles) are stored by using Common Internet File System (CIFS).
- User data files are stored by using CIFS.

Dedicated virtual desktops or stateless virtual desktops that need mobility require the following items to be on NFS or block I/O shared storage:

- NFS or block I/O is used to store all virtual desktops' associated data, such as the master image, replicas, and linked clones.
- NFS is used to store View Composer persistent disks when View Persona Management is used for user profile and user data files. This feature is not recommended.

* For more information, see the following VMware knowledge base article about creating linked clones on NFS storage: http://kb.vmware.com/kb/2046165

- NFS is used to store all virtual images for linked clones. The replicas and linked clones can be stored on local solid-state drive (SSD) storage. These items are discarded when the VM is shut down.

3.3 VMware vSAN

VMware vSAN is a Software Defined Storage (SDS) solution embedded in the ESXi hypervisor. vSAN pools flash caching devices and magnetic disks across three or more 10 GbE connected servers into a single shared datastore that is resilient and simple to manage.

vSAN can be scaled to 64 servers, with each server supporting up to 5 disk groups, with each disk group consisting of a single flash caching device (SSD) and up to 7 HDDs. Performance and capacity can easily be increased simply by adding more components: disks, flash or servers.

The flash cache is used to accelerate both reads and writes. Frequently read data is kept in read cache; writes are coalesced in cache and de-staged to disk efficiently, greatly improving application performance.

vSAN manages data in the form of flexible data containers that are called objects and the following types of objects for VMs are available:

- VM Home
- VM swap (`.vswp`)
- VMDK (`.vmdk`)
- Snapshots (`.vmsn`)

Internally, VM objects are split into multiple components that are based on performance and availability requirements that are defined in the VM storage profile. These components are distributed across multiple hosts in a cluster to tolerate simultaneous failures and meet performance requirements. vSAN uses a distributed RAID architecture to distribute data across the cluster. Components are distributed with the use of the following two storage policies:

- Number of stripes per object. It uses RAID 0 method.
- Number of failures to tolerate. It uses either RAID-1 or RAID-5/6 method. RAID-5/6 is currently supported for the all flash configuration only.

For more information about VMware Horizon virtual desktop types, objects, and components, see “VMware vSAN Design and Sizing Guide for Horizon Virtual Desktop Infrastructures”, which is available at this website: VMware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf

3.3.1 VMware vSAN Architecture

VMware vSAN supports two types of storage architectures leveraging latest hardware technologies and to provide scalable solution with different drive options to meet performance and capacity requirements.

vSAN Original Storage Architecture(OSA)

vSAN Original Storage Architecture comprises dedicated cache and capacity tiers and disk groups. All Flash configuration uses flash for both cache and capacity tier and hybrid configuration uses flash or NVMe drives for cache and HDDs for capacity tier. This architecture supports maximum 5 disk groups and maximum of 7

capacity drives per disk group. vSAN OSA does support different SSDs and HDDs type and sizes to create flexible configurations.

vSAN Express Storage Architecture(ESA)

vSAN Express Storage Architecture is a single tier storage solution and it is supported from vSAN 8 onwards. vSAN ESA requires high performance NVMe drives and leverages 25GbE ethernet links to provide superior performance. It uses improved erasure coding which further reduces performance overhead and enables customers to achieve RAID 5/6 at the performance of RAID 1.

3.3.2 VMware vSAN storage policies

VMware vSAN uses the Storage Policy-based Management (SPBM) function in vSphere to enable policy driven virtual machine provisioning, and uses vSphere APIs for Storage Awareness (VASA) to expose vSAN's storage capabilities to vCenter.

This approach means that storage resources are dynamically provisioned based on requested policy, and not pre-allocated as with many traditional storage solutions. Storage services are precisely aligned to VM boundaries; change the policy, and vSAN will implement the changes for the selected VMs.

VMware vSAN has the following storage policies:

Storage Policy	Description	Default	Maximum
Failure Tolerance Method	Defines a method used to tolerate failures. RAID-1 uses mirroring and RAID 5/6 uses parity blocks (erasure encoding) to provide space efficiency. RAID-5/6 is supported only for All Flash configurations. RAID 5 requires minimum 4 hosts and RAID 6 requires minimum 6 hosts. When RAID 5/6 is chosen, RAID 5 is used when FTT=1 and RAID 6 is used when FTT=2.	RAID-1	N/A
Number of failures to tolerate	Defines the number of host, disk, or network failures a VM object can tolerate. For n failures tolerated, n+1 copies of the VM object are created and 2n+1 hosts with storage are required. For example with a FTT=1, RAID-1 uses 2x the storage and RAID-5/6 uses 1.33x the storage. When FTT=2, RAID-1 uses 3x the storage and RAID-5/6 uses 1.5x the storage.	1	3
Number of disk stripes per object	The number of HDDs across which each replica of a VM object is striped. A value higher than 1 might result in better performance, but can result in higher use of resources.	1	12
Object space reservation	Percentage of the logical size of the object that should be reserved (or thick provisioned) during VM creation. The rest of the storage object is thin provisioned. If your disk is thick provisioned, 100% is	0%	100%

	reserved automatically. When deduplication and compression is enabled, this should be set to either 0% (do not apply) or 100%.		
Flash read cache reservation	SSD capacity reserved as read cache for the VM object. Specified as a percentage of the logical size of the object. Should be used only to address read performance issues. Reserved flash capacity cannot be used by other objects. Unreserved flash is shared fairly among all objects.	0%	100%
Force provisioning	If the option is set to Yes, the object is provisioned, even if the storage policy cannot be satisfied by the data store. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is no longer possible. The default of No is acceptable for most production environments.	No	N/A
IOPS limit for object	Defines IOPS limit for a disk and assumes a default block size of 32 KB. Read, write and cache operations are all considered equivalent. When the IOPS exceeds the limit, then IO is throttled.	0	User Defined
Disable object checksum	Detects corruption caused by hardware/software components including memory, drives, etc. during the read or write operations. Object checksums carry a small disk IO, memory and compute overhead and can be disabled on a per object basis.	No	Yes

The following storage policies and default values are recommended for VMware Horizon linked clones and full clones. Table 1 lists the default storage policies for linked clones.

Table 1: Default storage policy values for linked clones

	Dedicated				**Stateless**		
Storage Policy	**VM_HOME**	**Replica**	**OS Disk**	**Persistent Disk**	**VM_HOME**	**Replica**	**OS Disk**
Number of disk stripes per object	1	1	1	1	1	1	1
Flash-memory read cache reservation	0%	10 %	0%	0%	0%	10 %	0%
Number of failures to tolerate (FTT)	1	1	1	1	1	1	1
Failure Tolerance Method	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Force provisioning	No	No	No	No	No	No	No
Object-space reservation	0%	0%	0%	100%	0%	0%	0%
Disable object Checksum	N/A	N/A	N/A	N/A	N/A	N/A	N/A
IOPS limit for object	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Table 2 lists the default storage policies for full clones.

Table 2: Default storage policy values for full clones

Storage Policy	Dedicated	
	VM_HOME	Full Clone Disk
Number of disk stripes per object	1	1
Flash-memory read cache reservation	0%	0%
Number of failures to tolerate (FTT)	1	1
Failure Tolerance Method	N/A	N/A
Force provisioning	No	No
Object-space reservation	0%	100%
Disable object Checksum	N/A	N/A
IOPS limit for object	N/A	N/A

4 Operational model

This section describes the options for mapping the logical components of a client virtualization solution onto hardware and software. The “Operational model scenarios” section gives an overview of the available mappings and has pointers into the other sections for the related hardware. Each subsection contains performance data, has recommendations on how to size for that particular hardware, and a pointer to the BOM configurations that are described in section 5 on page 30. The last part of this section contains some deployment models for example customer scenarios.

4.1 Operational model scenarios

Figure 3 shows the following operational models (solutions) in Lenovo Client Virtualization: enterprise, small-medium business (SMB), and hyper-converged.

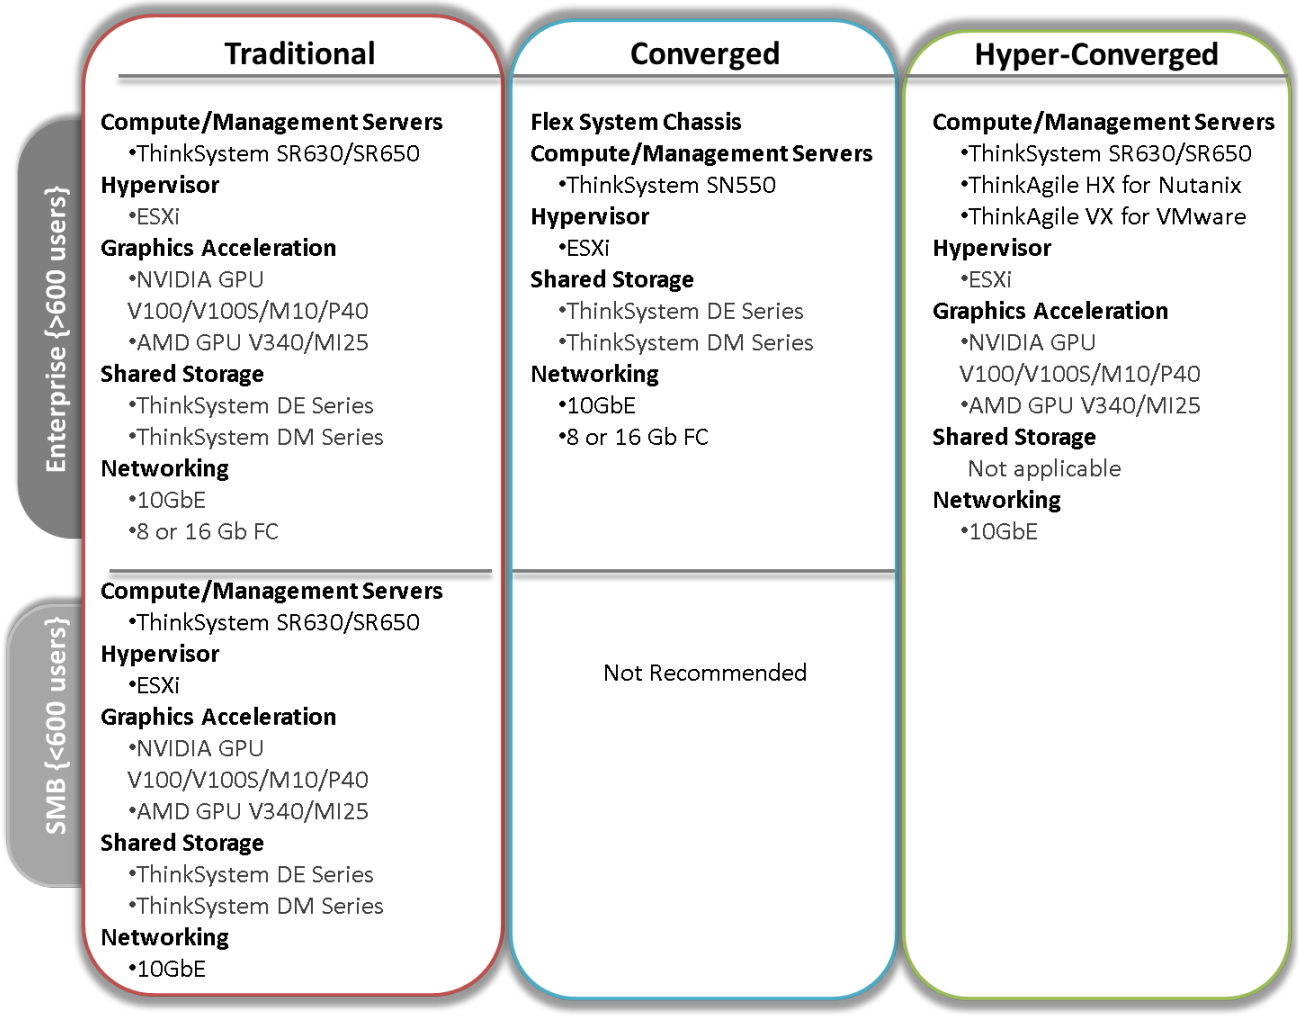

Figure 3: Operational model scenarios

The vertical axis is split into two halves: greater than 600 users is termed Enterprise and less than 600 is termed SMB. The 600 user split is not exact and provides rough guidance between Enterprise and SMB. The last column in Figure 3 (labelled “hyper-converged”) spans both halves because a hyper-converged solution can be deployed in a linear fashion from a small number of users (100) up to a large number of users (>4000).

The horizontal axis is split into three columns. The left-most column represents traditional rack-based systems with top-of-rack (TOR) switches and shared storage. The middle column represents converged systems where the compute, networking, and sometimes storage are converged into a chassis, such as the Flex System. The right-most column represents hyper-converged systems and the software that is used in these systems. For the purposes of this reference architecture, the traditional and converged columns are merged for enterprise solutions; the only significant differences are the networking, form factor, and capabilities of the compute servers.

Converged systems are not generally recommended for the SMB space because the converged hardware chassis can be more overhead when only a few compute nodes are needed. Other compute nodes in the converged chassis can be used for other workloads to make this hardware architecture more cost-effective.

4.1.1 Enterprise operational model

For the enterprise operational model, see the following sections for more information about each component, its performance, and sizing guidance:

- 4.2 Operating System and Hypervisor support
- 4.3 Compute servers for virtual desktops
- 4.5 Graphics acceleration
- 4.6 Management servers
- 4.7 Systems management
- 4.8 Shared storage
- 4.10 Networking
- 4.11 Racks

To show the enterprise operational model for different sized customer environments, four different sizing models are provided for supporting 300, 600, 1500, and 3000 users.

The SMB model is the same as the Enterprise model for traditional systems.

4.1.2 Hyper-converged operational model

For the hyper-converged operational model, see the following sections for more information about each component, its performance, and sizing guidance.

To show the hyper-converged operational model for different sized customer environments, four different sizing models are provided for supporting 300, 600, 1500, and 3000 users. The management server VMs for a hyper-converged cluster can either be in a separate hyper-converged cluster or on traditional shared storage.

4.2 Operating System and Hypervisor support

This reference architecture was tested with VMware ESXi 7.0 U3 and 8.0 U1 hypervisor with different processor generations from Intel and AMD.

The hypervisor is convenient because it can start from a M.2 boot drive and does not require any additional local storage. Alternatively, ESXi can be booted from SAN shared storage.

To use the latest version ESXi, download the Lenovo custom image from the following website:

https://VMware.lenovo.com/content/custom_iso/

4.3 Compute servers for virtual desktops

This section describes stateless and dedicated virtual desktop models. In some customer environments, stateless and dedicated desktop models might be required, which requires a hybrid implementation.

Compute servers are servers that run a hypervisor and host virtual desktops. There are several considerations for the performance of the compute server, including the processor family and clock speed, the number of processors, the speed and size of main memory, and local storage options.

For stateless users, the typical range of memory that is required for each desktop is 2 GB – 6 GB. For dedicated users, the range of memory for each desktop is 2 GB – 8 GB. Designers and engineers that require graphics acceleration might need 8 GB – 16 GB of RAM per desktop. In general, power users that require larger memory sizes also require more virtual processors. The virtual desktop memory should be large enough so that swapping is not needed and vSwap can be disabled. This reference architecture standardizes on 2 GB per desktop as the absolute minimum requirement of a Windows 10 desktop.

Windows 10 version 1909 was used for all of the performance testing. In general Windows 10 requires 10% to 20% more compute power than Windows 7. In addition, each Windows 10 versions 1803/1809/1909 have many new features and enhancements which will have an impact on the performance for login time and application loading time based on the particular applications running in the desktops. The following optimizations were applied to Windows 10 version 1909 base image:

- Used VMware latest OS Optimization Tool version 1151 and applied templates specific to Windows 10 version.
- Applied tuning parameters as per Microsoft recommendations: https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds_vdi-recommendations-1909 .
- Applied tuning scripts from: https://github.com/TheVDIGuys/Windows_10_VDI_Optimize
- Set Adobe acrobat as a default app for PDF files using steps in following webpage: https://www.adobe.com/devnet-docs/acrobatetk/tools/AdminGuide/pdfviewer.html

For more information, see “BOM for enterprise and SMB compute servers” section on page 30.

4.3.1 Login Enterprise Benchmark for Virtual Desktops

Login Enterprise is tool for benchmarking virtual desktops and applications. VSIMax is the metric used to measure the density of virtual desktops and EUX score metric is used to measure application response time. Login Enterprise comes with knowledge worker workload which uses standard office applications such as MicroSoft office applications and Edge/Chrome browser. It also provides response time for user login, application open and closing time. Also note that the previous LoginVSI benchmark is withdrawn from the market and Login Enterprise will be a replacement tool for virtual desktop benchmark. The Login Enterprise workload is 60-70% heavier than previous LoginVSI and it matches real world requirements.

4.3.2 ThinkSystem V2 Servers 3rd Gen Intel Xeon Processors

Table 3 shows the Login Enterprise 4.9.10 benchmark results for knowledge worker workload on Windows 10. The test was done with 280 desktops on ThinkSystem SR650 V2 server with local NVMe drives, 1.5 TB memory, 25GbE NIC and vSphere 8.0U1. As shown in Figure 4, the end user experience (EUX score) decreases when the total number of desktops increases beyond VSIMax due to increase in the response time (as shown in Figure 5) of applications due to resource contention.

Table 3: Login Enterprise Knowledge Worker Results for Windows 10

Server	Processor	Knowledge Worker	Total Desktops	VSIMax	EUX Score
ThinkSystem SR650V2	2x Gold 6338N @ 2.20GHz, 32C	3 vcpu+2GB+40GB disk	280	267	6.2
ThinkSystem SR650V2	2x Gold 6338N @ 2.20GHz, 32C	3 vcpu+2GB+40GB disk	250	>250	6.4

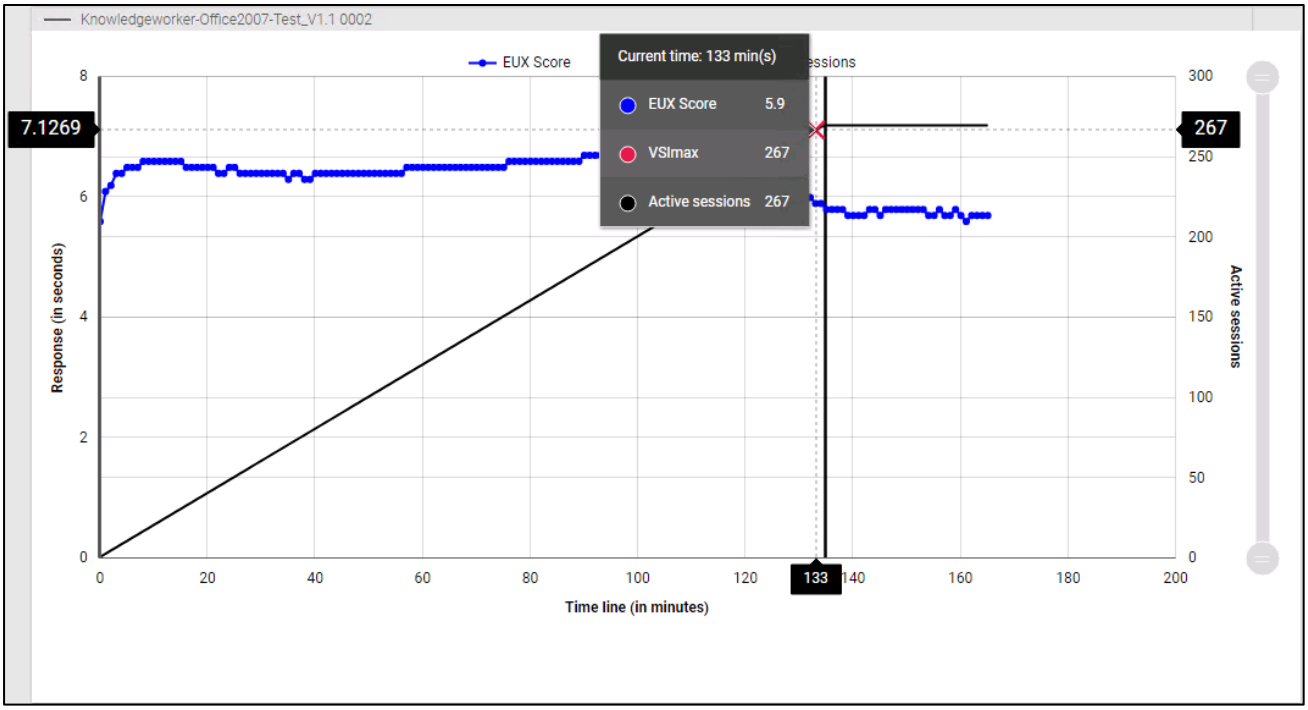

Figure 4: Login Enterprise VSIMax and EUX Score

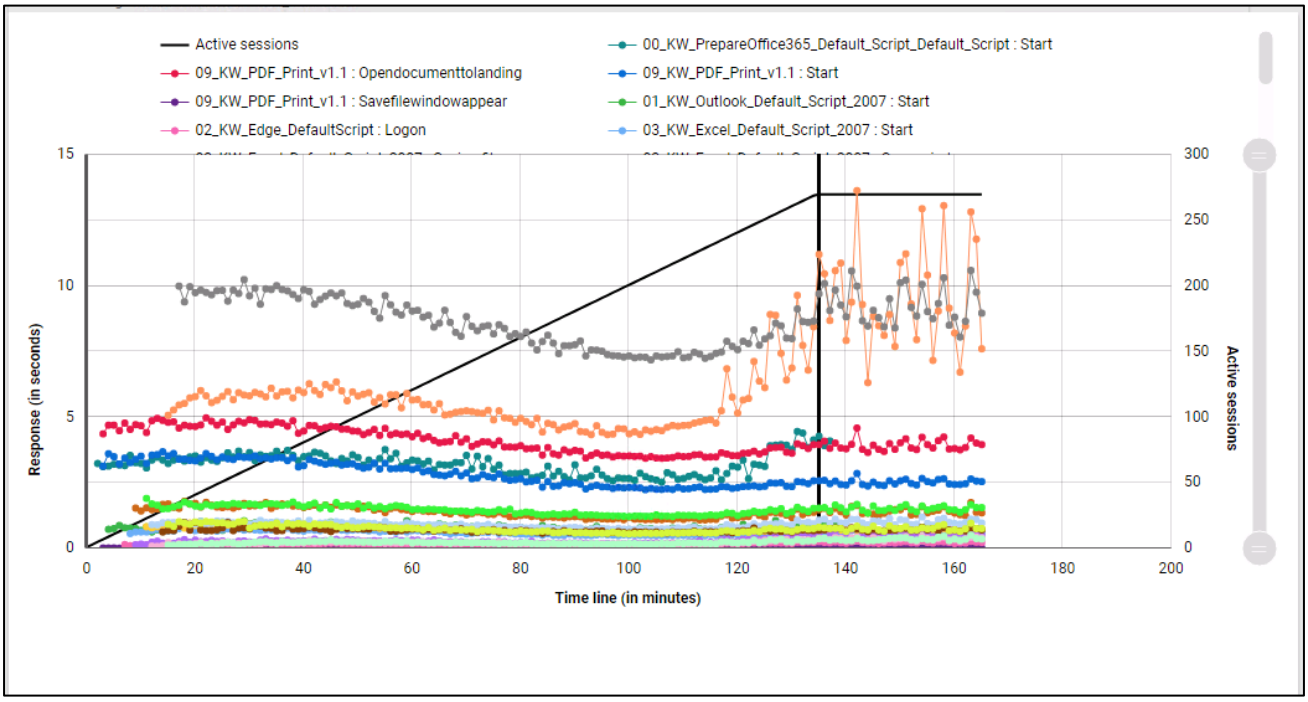

Figure 5: Login Enterprise Application Response Time with 3rd Gen Intel Xeon Scalable Processors

Table 4 lists the recommended number of virtual desktops per server for different workload types and VM memory sizes. The number of users is reduced in some cases to fit within the available memory and still maintain a reasonably balanced system of compute and memory.

Table 4: Recommended number of virtual desktops per server

Workload	Knowledge worker	Power worker
vSphere Cluster Nodes	8	8
Processor	2 x 6338 N	2 x 6338 N
VCPU	2	4
VM memory size	4 GB	8 GB
System memory	768 GB	1024 GB
Memory overhead	32 GB	32 GB
Desktops per server (normal mode)	165	125
Desktops per server (failover mode)	190	143

Table 5 lists the approximate number of compute servers that are needed for different numbers of users and Knowledge worker workloads.

Table 5: Compute servers needed for Knowledge workers and different numbers of users

Knowledge workers	300 users	600 users	1200 users	3000 users

Compute servers @165 users (normal)	3	4	7	18
Compute servers @190 users (failover)	2	3	6	16

4.3.3 ThinkSystem V3 Servers 4th Gen Intel Xeon Scalable Processors

Table 6 shows the Login Enterprise 5.1.2 benchmark results for knowledge worker workload on Windows 10. The test was done with 280 desktops on ThinkSystem SR850 V3 4 socket server with local NVMe drives, 1.5 TB memory, 10GbE NIC and vSphere 8.0U1.

Table 6: Login Enterprise Knowledge Worker Results on 4th Gen Intel Xeon Scalable Processors

Server	Processor	Knowledge Worker	Total Desktops	VSIMax	EUX Score
ThinkSystem SR850 V3	4 x 8444H 16C	3 vcpu+2GB+40GB disk	350	292	6.2

4.3.4 ThinkSystem V3 Servers 5th Gen Intel Xeon Scalable Processors

Table 7 shows the Login Enterprise 5.1.2 benchmark results for knowledge worker workload on Windows 10. The test was done with ThinkSystem SR650 V3 2-socket server with local NVMe drives, 2TB memory, 25GbE NIC and vSphere 8.0U1.

Table 7: Login Enterprise Knowledge Worker Results on 5th Gen Intel Scalable Processors

Server	Processor	Knowledge Worker	Total Desktops	VSIMax	EUX Score
ThinkSystem SR650 V3	2 x Intel Xeon Gold 6548Y+ 32C 2.5 GHz	3 vcpu+2GB+40GB disk	420	401	6.2
ThinkSystem SR650 V3	2x Intel Xeon Platinum 8580 60C 2.0 GHz	3 vcpu+2GB+40GB disk	620	598	5.9

The 5th Gen Intel Xeon Scalable Processors provide better density and improved performance in application response time than 3rd Gen processors. All the office applications response time is less than two seconds, and it is consistent during peak load and boot storm scenarios.

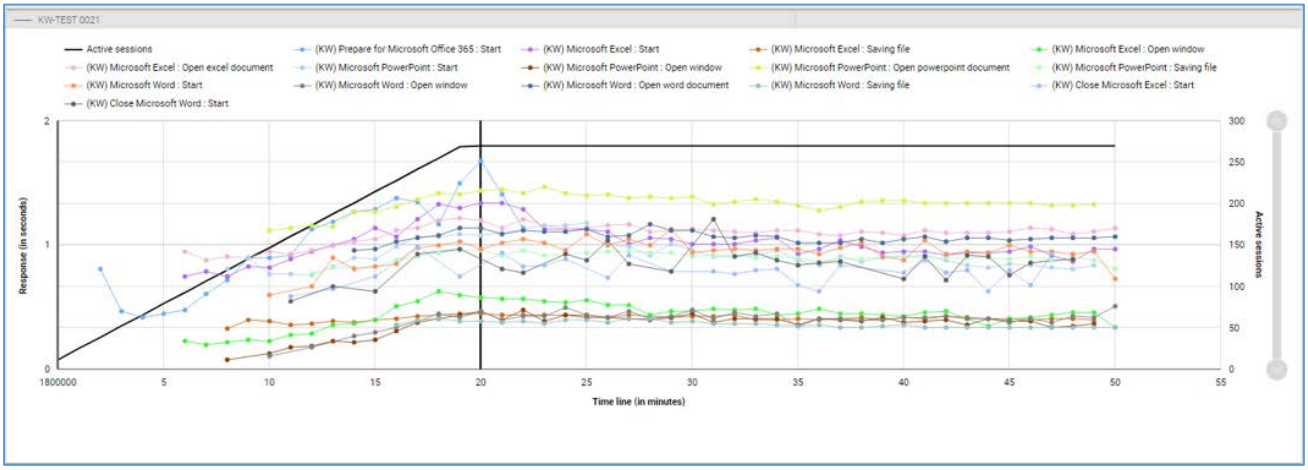

Figure 6: Login Enterprise Application Response Time with 5th Gen Intel Xeon Scalable Processors

4.3.5 ThinkSystem SR665 with 3rd Generation AMD EPYC Processors (Milan)

This section presents the single node and dual node performance for Lenovo ThinkAgile hyper-converged servers based on 3rd Generation AMD EPYC processors in conjunction with VMware ESXi 7.0 U3 and VMware Horizon 7.13. The tests were done using single and dual socket nodes Lenovo ThinkSystem SR665 with 32 cores AMD EPYC 7543 CPU and 16 core AMD EPYC 7313 CPU with 1536GB RAM. The VSAN storage consists of 2 x800GB SAS SSD cache drivers and 6 x 7.68 TB SAS SSD capacity drives.

Table 16: Login Enterprise Performance AMD EPYC 7003

Server	Processor	Workload	Users (@100% CPU)	Users (@75% CPU)
ThinkSystem SR665	2x AMD EPYC 7543 32C @ 2.8GHz	Knowledge worker	258	170
ThinkSystem SR665	2x AMD EPYC 7313 16C @ 2.8GHz	Knowledge worker	172	110

4.4 Compute servers for VMware vSAN

This section presents the compute servers for the hyper-converged system using VMware vSAN 8 with ESXi 8.0 U1. Additional processing and memory are required to support vSAN and OSA clusters with two disk groups need approximately 32 GB and those with 4 disk groups need 50GB. See the following for more details on the calculation of vSAN system memory:
kb.VMware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=2113954

As the price per GB for flash memory continues to reduce, there is a trend to use all SSDs to provide the overall best performance for vSAN.

vSAN 8 supports Express Storage Architecture (ESA) which require NVMe drives only capacity tiler and there is no cache tier required. It reduces storage footprint and provide better performance.

For more information, see “BOM for hyper-converged compute servers” on page 32.

4.4.1 ThinkAgile VX V2 with 3rd Generation Intel Xeon Scalable Processors (Ice Lake)

This section shows the performance results and sizing guidance for Lenovo ThinkAgile VX compute servers based on the 3rd Generation Intel Scalable processor (Whitley Ice Lake). Table 8 lists the VSI max performance results for Login Enterprise knowledge worker with ESXi 8.0 and Windows 10 version 1909.The nodes with Gold 6338N processors were tested with VSAN All Flash (OSA - Original Storage Architecture) configuration. The 3rd Gen Intel Xeon processors are with more number of cores and it provides higher density than previous generations.

Table 8: Login Enterprise performance Intel SP3

Server	Processor	Workload	Users (@100% CPU)	Users (@75% CPU
ThinkAgile VX7530	2x Gold 6338N @ 2.20GHz, 32C	Knowledge worker	275 users	180 users

4.4.2 ThinkAgile VX V3 with 4th Generation Intel Xeon Processors and vSAN 8 ESA

The benchmark was performed with 4 node ThinkAgile VX650 V3 cluster using vSAN 8 Express Storage Architecture configuration. The testing was done using Horizon 8 instant clones. The server is configured with 2x Intel(R) Xeon(R) Platinum 8468 48C 2.1Ghz processors, 16x 64GB DDR5 4800 MHz, 4x 3.20TB NVMe SSD and Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapter. The testing was performance by enabling/disabling RDMA. The test was done with 1460 desktops and achieved VSIMax 1425 with RDMA enabled. The test with RDMA disabled resulted ~5% lesser number and it shows there is slight improvement when using RDMA for vSAN ESA. Figure 7 shows VDI performance on Intel 4th Gen Xeon SP processors provide better EUX score from the beginning of the test and higher VSIMax than previous generation results described in section 4.3.2

Table 9: Login Enterprise 5.1.2 Knowledge Worker Results with vSAN ESA

Server	Processor	RDMA	Workload	Total Desktops	VSIMax	EUX Score	Users @75% CPU
ThinkAgile VX650 V3	2x Intel Xeon Platinum	Enabled	Login Enterprise Knowledge worker (3 vcpu+2GB+40GB disk)	1460	1425	6.3	270

	8468 48C 2.10Ghz						
ThinkAgile VX650 V3	2x Intel Xeon Platinum 8468 48C 2.10Ghz	Disabled	Login Enterprise Knowledge worker (3 vcpu+2GB+40GB disk)	1460	1353	6.2	260

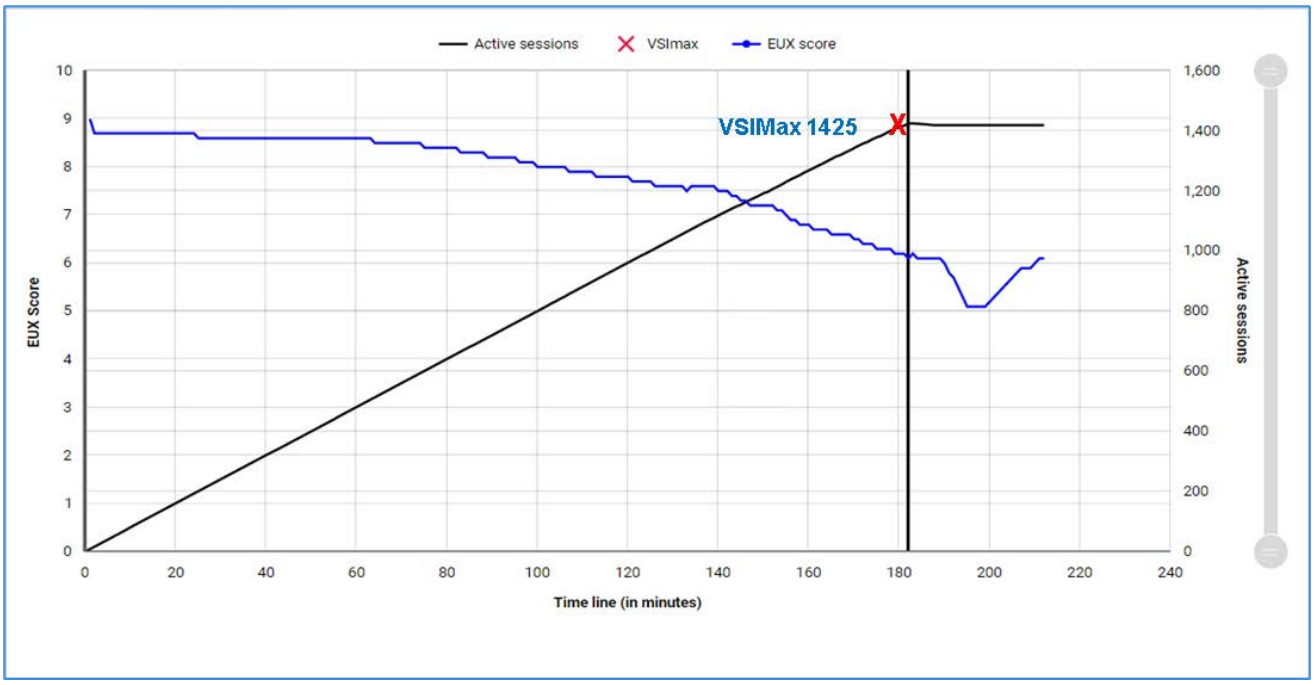

Figure 7: Login Enterprise VSIMax on 4th Gen Intel Xeon SP with vSAN ESA

4.4.3 Windows 11 Performance with Lenovo ThinkAgile VX

This section shows the performance results of Windows 11 on Lenovo ThinkAgile VX compute servers based on the Intel Scalable processor Gen 3 (Whitley Ice Lake). Table 10 lists the VSI max performance results for Login VSI 4.1 workloads with ESXi 7.0 U3 and Windows 11 21H2.The Gold 6338N processors nodes were tested with VSAN All Flash configuration. The Windows 11 requires minimum 4GB memory and 64 GB disk than Windows 10. The Windows 11 performance is 5-10% lower than Windows 10 for the same knowledge worker workload. Also reducing the memory less than 4GB results drop in performance.

Table 10: Windows 11 performance with Intel Xeon SP 3

Server	Processor	Workload	Users (@100% CPU)	Users (@75% CPU
ThinkAgile VX7530	2x Gold 6338N	Knowledge worker (2 vcpu+4GB+80GB disk)	255 users	165 users

4.4.4 ThinkAgile VX V3 with 4th Generation AMD EPYC Processors (Genoa) with vSAN 8 ESA

The benchmark was performed with 3 node ThinkAgile VX665 V3 cluster using vSAN 8 Express Storage Architecture configuration. The testing was done using Horizon 8 instant clones. The server is configured with 2x AMD EPYC 9354 @ 3.25 Ghz 32C, 32x 32GB DDR5 4800 MHz, 6x 3.20TB NVMe SSD and Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapter. The VMware vSAN ESA supports 500 virtual machines per node and AMD EPYC high core processors can accommodate large number of virtual desktops with a smaller number of nodes.

Table 11: Login Enterprise Knowledge Worker Results with AMD EPYC 9004

Server	Processor	RDMA	Workload	Total Desktops	VSIMax	EUX Score	Users @75% CPU
ThinkAgile VX665 V3	2x AMD EPYC 9354 @ 3.25 Ghz 32C	Enabled	Login Enterprise Knowledge worker (3 vcpu+2GB+40GB disk)	950	867	6.2	215
ThinkAgile VX665 V3	2x AMD EPYC 9474F @ 3.6 Ghz 48C	Enabled	Login Enterprise Knowledge worker (3 vcpu+2GB+40GB disk)	1800	1761	6.4	325

Figure 8 shows VSIMax and EUX score pattern for knowledge worker test with ThinkAgile VX 665 V3 with AMD EPYC 9474F. The response time for applications is very good until total logins reach to 1200 on a 4-node cluster and the CPU utilization is below 85% at that point of time.

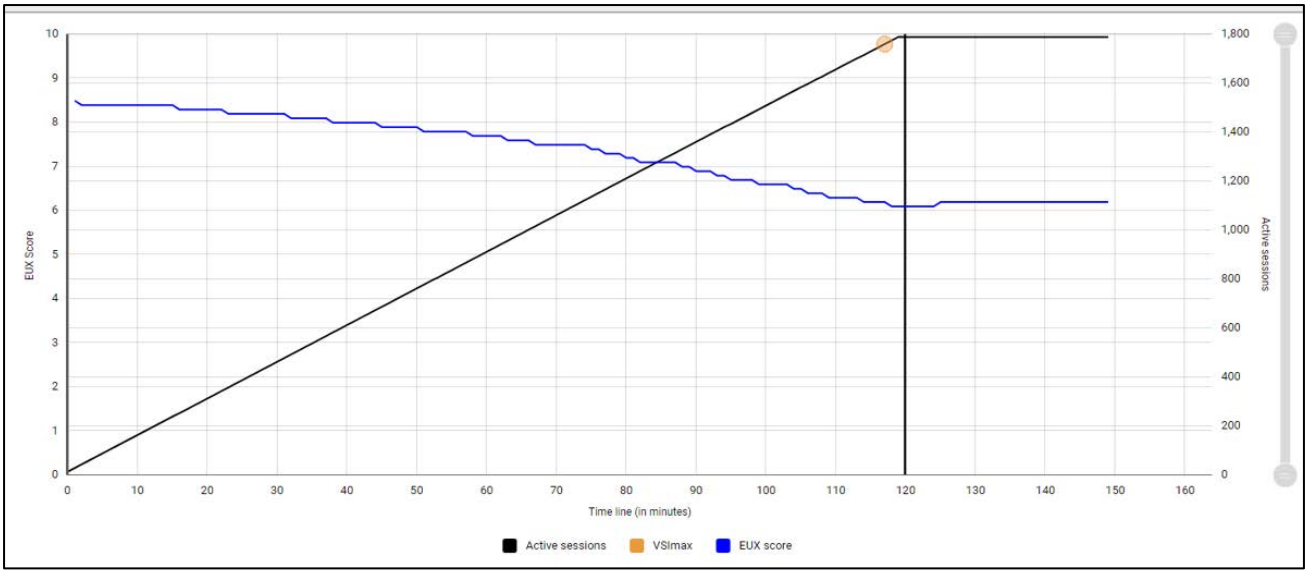

Figure 8: Login Enterprise VSIMax on Intel 4th Gen AMD EPYC vSAN ESA

4.5 Graphics acceleration

The VMware ESXi hypervisor supports the following options for graphics acceleration:

- Dedicated GPU with one GPU per user, which is called virtual dedicated graphics acceleration (vDGA) mode.
- GPU hardware virtualization (vGPU) that partitions each GPU for 1 – 8 users.
- Shared GPU with users sharing a GPU, which is called virtual shared graphics acceleration (vSGA) mode and is not recommended because of user contention for shared use of the GPU.

VMware also provides software emulation of a GPU, which can be processor-intensive and disruptive to other users who have a choppy experience because of reduced processor performance. Software emulation is not recommended for any user who requires graphics acceleration.

The vDGA option has a low user density as it restricts a single user to access each very powerful GPU. This option is not flexible and is no longer cost effective even for high-end power users. Therefore vDGA is no longer recommended especially given that the performance of the equivalent vGPU mode is similar.

When using the vGPU option with ESXi 6.5 and the latest drivers from Nvidia, it is necessary to change the default GPU mode from “Shared” (vSGA) to “Shared Direct” (vGPU) for each GPU using VMware vCenter. This enables the correct GPU support for the VMs which would otherwise result in the VM not powering on correctly and the standard “graphics resources not available” error message. The host needs to be rebooted for the changes to take effect.

There is also a known problem with VMware vSphere that may prevent the VM from booting correctly. Lenovo recommends always setting “pciPassthru.use64bitMMIO” option to TRUE regardless of the memory size of the VM by choosing “VM Options > Advanced”. See also the Nvidia release notes for VMware vSphere: docs.nvidia.com/grid/latest/grid-vgpu-release-notes-VMware-vsphere/index.html.

The performance of graphics acceleration was tested using the Lenovo ThinkSystem SR650 servers. Each server supports up to two GPU adapters. The Heaven benchmark is used to measure the per user frame rate for different GPUs, resolutions, and image quality. This benchmark is graphics-heavy and is fairly realistic for designers and engineers. Power users or knowledge workers usually have less intense graphics workloads and can achieve higher frame rates.

Table 12 lists the results of the Heaven benchmark as FPS that are available to each user with the Nvidia M10 GPU using different vGPU profiles with VMware Horizon 7.1, DirectX 11, and ESXi 6.5U1 on Xeon SP1 CPU 6130.

Table 12: Performance of Nvidia M10 vGPU profiles

Quality	Tessellation	Anti-Aliasing	Resolution	M10-8Q	M10-4Q	M10-2Q	M10-4A	M10-2A
High	Normal	0	1280x1024	Untested	Untested	13.0	25.3	13.4
High	Normal	0	1680x1050	Untested	Untested	10.3	N/A	N/A
High	Normal	0	1920x1200	Untested	16.7	Untested	N/A	N/A
Ultra	Extreme	8	1280x1024	Untested	Untested	6.4	12.9	6.3
Ultra	Extreme	8	1680x1050	21.1	10.6	Untested	N/A	N/A

Ultra	Extreme	8	1920x1080	18.8	9.3	Untested	N/A	N/A
Ultra	Extreme	8	1920x1200	17.2	8.5	Untested	N/A	N/A
Ultra	Extreme	8	2560x1600	11.3	Untested	Untested	N/A	N/A

Table 13 lists the results of the Heaven benchmark as FPS that are available to each user with the Nvidia P40 GPU using different vGPU profiles with VMware Horizon 7.4, DirectX 11, and ESXi 6.7.

Table 13: Performance of Nvidia P40 vGPU profiles

Quality	Tessellation	Anti-Aliasing	Resolution	P40-24Q	P40-12Q	P40-6Q	P40-4Q
High	Normal	0	1280x1024	Untested	Untested	57.7	42.6
High	Normal	0	1680x1050	Untested	65.0	49.5	34.7
High	Normal	0	1920x1200	Untested	64.8	44.7	29.7
Ultra	Extreme	8	1280x1024	Untested	61.3	37.9	24.8
Ultra	Extreme	8	1680x1050	64.8	54.2	31.1	20.9
Ultra	Extreme	8	1920x1080	64.8	50.2	26.5	17.7
Ultra	Extreme	8	1920x1200	65.0	47.5	25.1	16.8
Ultra	Extreme	8	2560x1600	61.0	32.4	16.8	Untested
Ultra	Extreme	8	3840x2160	34.9	18.5	Untested	Untested
Ultra	Extreme	8	4096x2160	32.5	16.5	Untested	Untested

Lenovo recommends that a medium to high powered CPU, such as the Xeon SP1 6130 or Xeon SP2 6230, are used for accelerated graphics applications tend to also require extra load on the processor. For vGPU mode, Lenovo recommends at least 384GB of server memory. Because there are many variables when graphics acceleration is used, Lenovo recommends that testing is done in the customer environment to verify the performance for the required user workloads.

For more information about the bill of materials (BOM) for Nvidia GPUs for Lenovo ThinkSystem SR650 servers, see the following corresponding BOMs:

- “BOM for enterprise and SMB compute servers” section on page 30.
- “BOM for hyper-converged compute servers” on page 32

4.6 Management servers

Management servers should have the same hardware specification as compute servers so that they can be used interchangeably in a worst-case scenario. The VMware Horizon management servers also use the same ESXi hypervisor, but have management VMs instead of user desktops. Table 14 lists the VM requirements and performance characteristics of each management service.

Table 14: Characteristics of VMware Horizon and ESXi management services

Management service VM	Virtual processors	System memory	Storage	Windows OS	HA needed	Performance characteristic
vCenter Server Appliance	8	12 GB	60 GB	2016	No	Up to 2000 VMs.
vCenter SQL Server	4	8 GB	200 GB	2016	Yes	Double the virtual processors and memory for more than 2500 users.
View Connection Server	4	16 GB	60 GB	2016	Yes	Up to 2000 connections.

Table 15 lists the number of management VMs for each size of users following the high-availability and performance characteristics. 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 15: Management VMs needed

Horizon management service VM	300 users	600 users	1200 users	3000 users
vCenter servers	1	1	2	2
vCenter SQL servers	2 (1+1)	2 (1+1)	2 (1+1)	2 (1+1)
View Connection Server	2 (1+1)	2 (1+1)	2 (1+1)	2 (1+1)

Each management VM requires a certain amount of virtual processors, memory, and disk. There is enough capacity in the management servers for all of these VMs. Table 16 lists an example mapping of the management VMs to the two physical management servers for 3000 stateless users.

Table 16: Management server VM mapping (3000 stateless users)

Management service for 3000 stateless users	Management server 1	Management server 2
vCenter servers (2)	1	1
vCenter SQL servers (2)	1	1
View Connection Server (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.

For shared storage systems that support block data transfers only, it is also necessary to provide some file I/O servers that support CIFS or NFS shares and translate file requests to the block storage system. For high availability, two or more Windows storage servers are clustered.

Based on the number and type of desktops, Table 17 lists the recommended number of physical management servers. In all cases, there is redundancy in the physical management servers and the management VMs.

Note for a small number of users the management VMs could be executed on the same servers as the virtual desktop VMs.

Table 17: Management servers needed

Management servers	300 users	600 users	1200 users	3000 users
Stateless desktop model	0	0	2	2
Dedicated desktop model	0	0	2	2
Windows Storage Server 2012	2	2	2	2

For more information, see “BOM for enterprise and SMB management servers on page 39

4.7 Systems management

Lenovo XClarity™ Administrator is a centralized resource management solution that reduces complexity, speeds up response, and enhances the availability of Lenovo® server systems and solutions.

The Lenovo XClarity Administrator provides agent-free hardware management for Lenovo’s ThinkSystem, System x® rack servers and Flex System™ compute nodes and components, including the Chassis Management Module (CMM) and Flex System I/O modules. Figure 9 shows the Lenovo XClarity administrator interface, where Flex System components and rack servers are managed and are seen on the dashboard. Lenovo XClarity Administrator is a virtual appliance that is quickly imported into a virtualized environment server configuration.

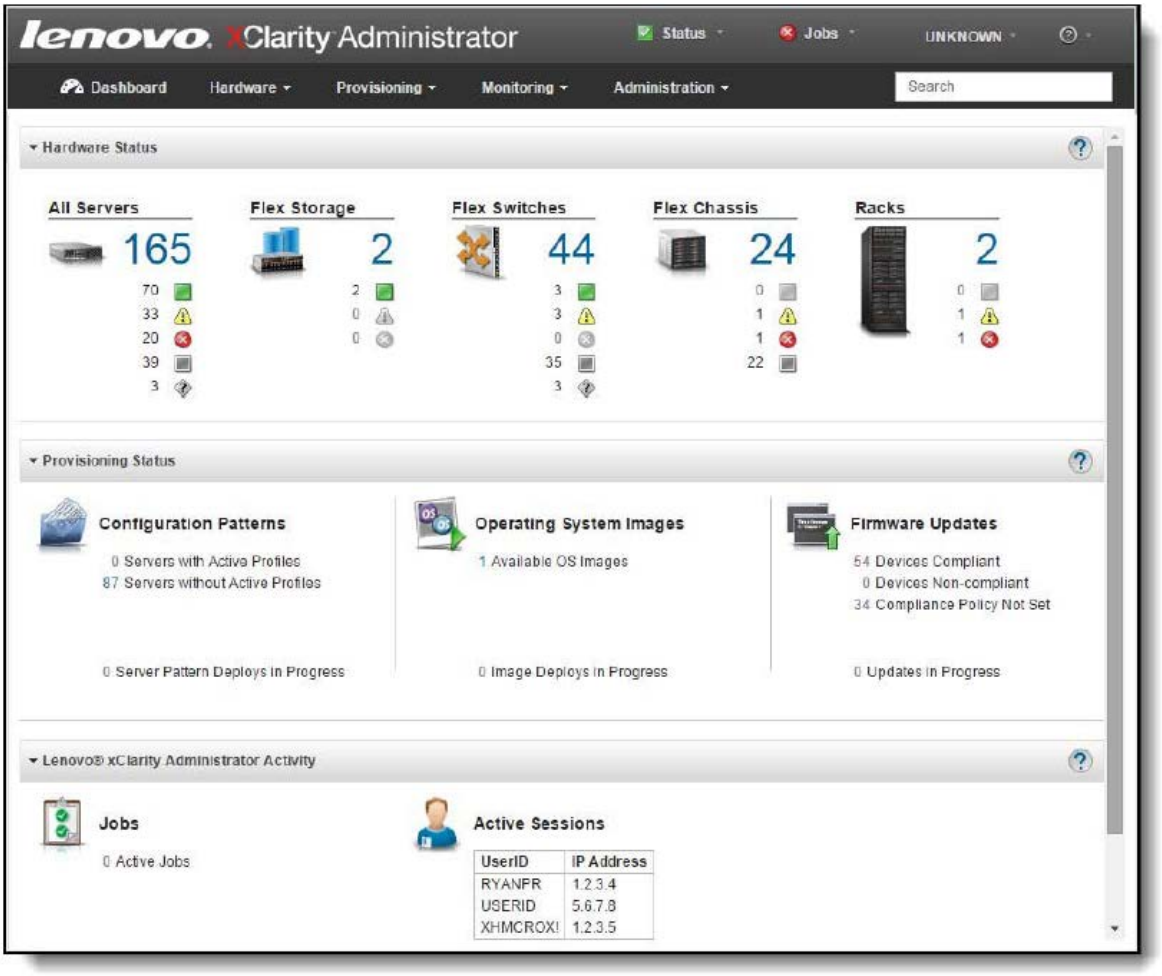

Figure 9: XClarity Administrator interface

4.8 Shared storage

VDI workloads, such as virtual desktop provisioning, VM loading across the network, and access to user profiles and data files place huge demands on network shared storage.

Experimentation with VDI infrastructures shows that the input/output operation per second (IOPS) performance takes precedence over storage capacity. This precedence means that more slower speed drives are needed to match the performance of fewer higher speed drives. Even with the fastest HDDs available today (15k rpm), there can still be excess capacity in the storage system because extra spindles are needed to provide the IOPS performance. From experience, this extra storage is more than sufficient for the other types of data such as SQL databases and transaction logs.

The large rate of IOPS, and therefore, large number of drives needed for dedicated virtual desktops can be ameliorated to some extent by caching data in flash memory or SSD drives. The storage configurations are based on the peak performance requirement, which usually occurs during the so-called “logon storm.” This is when all workers at a company arrive in the morning and try to start their virtual desktops, all at the same time.

It is always recommended that user data files (shared folders) and user profile data are stored separately from the user image. By default, this has to be done for stateless virtual desktops and should also be done for dedicated virtual desktops. It is assumed that 100% of the users at peak load times require concurrent access to user data and profiles.

In View 5.1, VMware introduced the View Storage Accelerator (VSA) feature that is based on the ESXi Content-Based Read Cache (CBRC). VSA provides a per-host RAM-based solution for VMs, which considerably reduces the read I/O requests that are issued to the shared storage. It is recommended to use the maximum VSA cache size of 2GB. Performance measurements by Lenovo show that VSA has a negligible effect on the number of virtual desktops that can be used on a compute server while it can reduce the read requests to storage by one-fifth.

Table 18 summarizes the peak IOPS and disk space requirements for virtual desktops on a per-user basis. Dedicated virtual desktops require a high number of IOPS and a large amount of disk space for the linked clones. Note that the linked clones also can grow in size over time.

Table 18: Shared virtual desktop shared storage performance requirements

Shared virtual desktops	Protocol	Size	IOPS	Write %
Replica	Block/NFS	30 GB	18	85%
Linked clones	Block/NFS	10 GB		
User “AppData” folder				
vSwap (recommended to be disabled)	NFS or Block	0	0	0
User data files	CIFS/NFS	5 GB	1	75%
User profile (through MSRP)	CIFS	100 MB	0.8	75%

The sizes and IOPS for user data files and user profiles that are listed in Table 18 can vary depending on the customer environment. For example, power users might require 10 GB and five IOPS for user files because of

the applications they use. It is assumed that 100% of the users at peak load times require concurrent access to user data files and profiles.

Many customers need a hybrid environment of stateless and dedicated desktops for their users. The IOPS for dedicated users outweigh those for stateless users; therefore, it is best to bias towards dedicated users in any storage controller configuration.

The storage configurations that are presented in this section include conservative assumptions about the VM size, changes to the VM, and user data sizes to ensure that the configurations can cope with the most demanding user scenarios.

4.9 Lenovo ThinkSystem DM and DE Series Storage Arrays

In this section we provide an overview of the newly introduced Lenovo DE and DM series storage arrays. For VDI workloads, the DE/DM series storage provides a wide range of options, supporting high-performance and scale-out options as well as rich set of software capabilities such as storage management, block/file level support, encryption, deduplication, and so forth.

4.9.1 Lenovo ThinkSystem DM Storage

The Lenovo ThinkSystem DM storage arrays provide a unified storage solution to manage all your block-and-file workloads on one array. DM series comes in both Hybrid Flash and All Flash configurations.

DM Series Hybrid Flash systems simplify the task of managing growth and complexity by delivering high performance, supporting a broad range of unified workloads, and seamlessly scaling of performance and capacity. For growing organizations that are concerned about budgets and meeting challenging IT needs, ThinkSystem DM Series Hybrid Flash systems are the perfect choice.

Figure 10: Lenovo ThinkSystem DM Storage

4.9.2 Lenovo ThinkSystem DE Storage

Lenovo ThinkSystem DE Series Storage Arrays are SAN storage systems that are designed to provide performance, simplicity, capacity, security, and high availability for medium to large businesses. ThinkSystem DE Series Storage Arrays deliver hybrid and all flash storage with enterprise-class storage management capabilities and a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features.

Figure 11: Lenovo ThinkSystem DE Storage

Table 19 lists a comparison of the key features of the DM and DE series and suggested use cases.

For more information and specifications of the Lenovo ThinkSystem DE and DM series storage arrays, see the following Lenovo Press publications:

Lenovo ThinkSystem DM Series Unified Storage Arrays

Lenovo ThinkSystem DE Series Storage Arrays

Table 19: Comparison and applications of DM Vs DE series

Property	DM series	DE Series
General **Application**	Robust data management and shared general IT	High performance, streamlined architecture that is optimized for low-latency workloads
Protocol	Unified and file	Block
Data Management	Needed in the storage layer	Advanced management that is provided up stack (app, OS, or file system)
Workload Characteristics	Shared, multi-tenancy, concurrent Benefits from data reduction Transactional	Dedicated, application focused Does not benefit from data reduction Extreme low latency
Performance Requirements	Mixed I/O Scale-out IOPS	High bandwidth Scale-up IOPS Consistent ultralow latency
Data Protection	Integrated	Application-driven
Efficiency Features	Deduplication Compression Compaction	High capacity per unit High performance per unit High performance per unit
Flash	Hybrid cache tier (read or read/write) Persistent storage	Predictive cache tier (read only) Persistent storage

4.9.3 Storage recommendations

Lenovo recommends using all SSDs to provide the best performance both in terms of general usage and for boot and login storms. Note that the sizing does not assume de-duplication and compression features. Therefore the required disk capacity can be determined from the storage capacity of the linked clone or full clone VMs, master images, and user storage.

Table 20 lists the storage configuration that is needed for each of the different numbers of stateless users. It is assumed that the master images are 40 GB (actual) with linked clones of 3GB each. Each user requires up to 5GB of persistent storage. The total storage includes 30% capacity for growth.

Table 20: Storage configuration for stateless users

Stateless storage	300 users	600 users	1200 users	3000 users
Master image count	1	1	2	4
Master image per controller with backup	160 GB	160 GB	320 GB	640 GB
Stateless linked clones (no growth)	1800 GB	3600 GB	7200 GB	18000 GB
User storage	1500 GB	3000 GB	3600 GB	15000 GB
Total storage (with 30% growth)	4.4 TB	8.6 TB	14 TB	33 TB
Number of 800 GB SSDs (RAID 5)	8	14	24	0
Number of 3.84 GB SSDs (RAID 5)	0	0	0	12

Table 21 lists the storage configuration that is needed for each of the different numbers of dedicated users. It is assumed that the master images are 40 GB (actual) with linked clone growth of up to 50%. Each user requires up to 5GB of persistent storage.

Table 21: Storage configuration for dedicated users

Stateless storage	300 users	600 users	1200 users	3000 users
Master image count	1	1	2	4
Master image per controller with backup	160 GB	160 GB	320 GB	640 GB
Stateless linked clones (50% growth)	6900 GB	13800 GB	27600 GB	69000 GB
User storage	1500 GB	3000 GB	3600 GB	15000 GB
Total storage (no growth)	9 TB	16.6 TB	31 TB	83 TB
Number of 800 GB SSDs (RAID 5)	16	24	0	0
Number of 3.84 GB SSDs (RAID 5)	0	0	10	24

4.10 Networking

The main driver for the type of networking that is needed for VDI is the connection to storage. If the shared storage is block-based it is likely that a SAN that is based on 8 or 16 Gbps FC, or 10 GbE iSCSI connection is needed. Other types of storage can be network attached by using 1 Gb or 10 Gb Ethernet. For vSAN 10 GbE

networking should always be used. Also, there is user and management virtual local area networks (VLANs) available that require 1 Gb or 10 Gb Ethernet. Automated failover and redundancy of the entire network infrastructure and shared storage is important. This failover and redundancy is achieved by having at least two of everything and ensuring that there are dual paths between the compute servers, management servers, and shared storage.

4.11 Racks

The number of racks and chassis for Flex System compute nodes depends upon the precise configuration that is supported and the total height of all of the component parts: servers, storage, networking switches, and Flex System Enterprise Chassis (if applicable). The number of racks for ThinkSystem servers is also dependent on the total height of all of the components.

4.12 Deployment Model with ThinkAgile VX

This deployment example supports 1200 dedicated virtual desktops with VMware vSAN in All Flash mode Each VM needs 3 GB of RAM. Six servers are needed assuming 200 users per server using the Office worker profile. There can be up to 1 complete server failure and there is still capacity to run the 1200 users on five servers (ratio of 5:1).

Each virtual desktop with Windows 10 takes 40 GB. The total capacity would be 126 TB, assuming full clones of 1200 VMs, one complete mirror of the data (FTT=1), space for swapping, and 30% extra capacity. However, the use of linked clones for vSAN or de-duplication for all flash means that a substantially smaller capacity is sufficient.

This example uses six ThinkAgile VX3330 1U servers with two Intel Silver Xeon SP3 4314 processors, 768 GB of memory, two network cards, two 800 GB SSDs, and four 3.84TB SSDs configured with two disk groups.

Figure 12 shows the deployment configuration for this example.

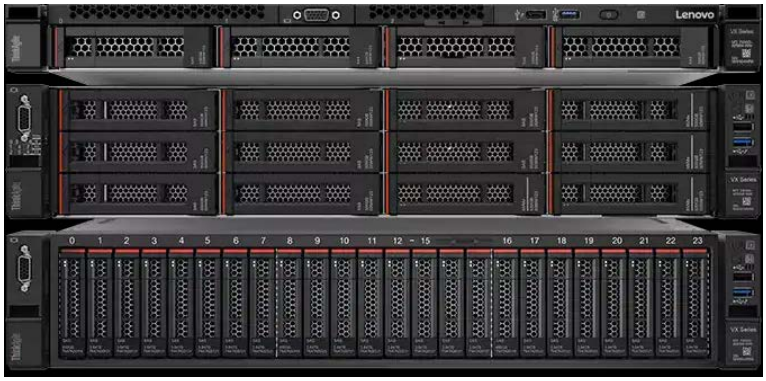

Figure 12: Deployment configuration for 1200 dedicated users using vSAN on ThinkAgile VX servers

5 Appendix: Bill of materials

This appendix contains the bill of materials (BOMs) for different configurations of hardware. There are sections for user servers, management servers, storage, networking switches, chassis, and racks that are orderable from Lenovo. The last section is for hardware orderable from an OEM.

The BOM lists in this appendix are not meant to be exhaustive and must always be double-checked with the configuration tools. Any discussion of pricing, support, and maintenance options is outside the scope of this document.

For connections between TOR switches and devices (servers, storage, and chassis), the connector cables are configured with the device. The TOR switch configuration includes only transceivers or other cabling that is needed for failover or redundancy.

5.1 BOM for enterprise and SMB compute servers

This section contains the bill of materials for enterprise and SMB compute servers. Table 22 and

Table 23 list the BOM for a server (a minimum of two servers per cluster is needed).

Table 22: ThinkSystem SR630V2 (3rd Generation Intel Xeon SP)

Part number	Product Description	Qty
7Z71A04WNA	Server: ThinkSystem SR630 V2, 1xIntel Xeon Silver 4314 16C 2.4GHz 135W, 1x32GB 2Rx8, SW RD, 1x750W, XCC Enterprise,	1
4XG7A63411	ThinkSystem SR630 V2 Intel Xeon Silver 4314 16C 135W 2.4GHz Processor Option Kit w/o Fan	1
4X77A08633	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	23
4XC7A08237	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
4F17A14487	ThinkSystem SR630 V2/SR645 Performance Fan Option Kit	2
5PS7A67541	Premier Essential - 3Yr 24x7 4Hr Resp + YDYD SR630 V2	1
7S06073EWW	VMware HCI 6 Kit Essentials for 3 Nodes (Max 2 processors per node) w/Lenovo 3Yr S&S	1
00MT202	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

Table 23: ThinkSystem SR650V2 (3rd Generation Intel Xeon SP)

Part number	Product Description	Qty
7Z73A06FNA	Server: ThinkSystem SR650 V2, 1xIntel Xeon Silver 4314 16C 2.4GHz 135W, 1x32GB 2Rx8, SW RD, 1x750W, XCC Enterprise,	1
4XG7A63455	ThinkSystem SR650 V2 Intel Xeon Silver 4314 16C 135W 2.4GHz Processor Option Kit w/o Fan	1
4X77A08634	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx8 1.2V) RDIMM	23
4Y37A09739	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
4XB7A17073	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	1
4XC7A08237	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
4F17A14496	ThinkSystem SR650 V2 Performance Fan Option Kit	1

4X97A80440	ThinkSystem SR650 V2/SR665 M.2 Cable Kit v2	1
5WS7A67902	Premier Foundation - 3Yr NBD Resp SR650 V2	1
00MT202	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

5.2 BOM for hyper-converged compute servers

This section contains the bill of materials for hyper-converged compute servers. Table 24, Table 25 and Table 26 list the BOM for a vSAN cluster(a minimum of four servers per cluster is needed).

Table 24: ThinkAgile VX3330 Appliance (3rd Generation Intel Xeon SP)

Part number	Product Description	Qty
7Z62CTO2WW	Server : Lenovo ThinkAgile VX3330 Appliance	1
BJLH	ThinkAgile VX 1U 2.5" Chassis with 8 or 10 Bays	1
B0W3	XClarity Pro	1
BPZQ	VMware HCI Kit 6 Enterprise (per CPU)	1
BN8K	ThinkAgile VX Remote Deployment	1
BB2W	Intel Xeon Gold 6346 16C 205W 3.1GHz Processor	2
B964	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	16
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
B5MC	vSAN All Flash Config	1
B5XH	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2
BMEY	VMware ESXi 7.0 U3 (Factory Installed)	1
B8MX	ThinkSystem 1U 10x2.5" (6x SAS/SATA 4x AnyBay) Backplane	1
B8N2	ThinkSystem 1U PCIe Gen4 x16/x16 Riser 1	1
B8NC	ThinkSystem 1U LP+LP BF Riser Cage Riser 1	1
B5SZ	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
BHS9	ThinkSystem 1100W (230V/115V) v2 Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
BH9M	ThinkSystem 1U Performance Fan Option Kit	8
B8LA	ThinkSystem Toolless Slide Rail Kit v2	1
B97M	ThinkSystem SR630 V2 MB	1
B0ML	Feature Enable TPM on MB	1
B8GY	M.2 Module Cable	1
BMJC	ThinkSystem 8x2.5" BP and 6+4 x2.5" BP Power Cable v2	1
BA22	ThinkSystem SR630 V2 MB to 6SATA/SAS and 4 AnyBay BP Cable 440mm	1
BC4T	ThinkSystem 1U SFF Front 2.5" Backplane SAS/SATA Gen3 Cable	1
BA1W	ThinkSystem SR630 V2 SFF Front BP SATA/SAS Cable	1
BJDQ	ThinkAgile VX3330 Appliance	1
B173	Companion Part for XClarity Controller Standard to Enterprise Upgrade in Factory	1
BJDJ	ThinkAgile SR630 v2 Service Label LI	1
B978	ThinkSystem SR630/SR850/SR860 V2 Standard Heatsink	2
B49C	ThinkSystem 2.5" S4510 3.84TB Read Intensive SATA 6Gb HS SSD	8
B8HU	ThinkSystem 2.5" PM1645a 800GB Mainstream SAS 12Gb Hot Swap SSD	2

7S06CTOEWW	VMware vSAN	1
S6AR	VMware HCI Kit 6 Enterprise (per CPU) w/Lenovo 3Yr S&S	2
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
5WS7A94898	Foundation ThinkAgile APP - 3Yr NBD Resp VX3330	1
5MS7A87711	ThinkAgile VX Remote Deployment (up to 4 node cluster)	1
7S06CTOSWW	VMware vSAN for VX	1
S6J0	VMware vSAN 7 Advanced for Desktop 10 Pack (CCU) w/Lenovo 3Yr S&S	1

Table 25: BOM for ThinkAgile VX3375 Integrated System (3rd Generation AMD EPYC processor) based compute nodes

Part number	Product Description	Qty
7D82CTO2WW	Server : Lenovo ThinkAgile VX3375 Integrated System	1
BN50	ThinkAgile VX 1U 2.5" Chassis with 8 or 10 Bays	1
B0W3	XClarity Pro	1
BPZQ	VMware HCI Kit 6 Enterprise (per CPU)	1
BN8K	ThinkAgile VX Remote Deployment	1
BF7E	ThinkSystem AMD EPYC 7343 16C 190W 3.2GHz Processor	2
B5XE	ThinkSystem 32GB TruDDR4 3200MHz (2Rx4 1.2V) RDIMM-A	16
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
BC4V	Non RAID NVMe	1
B5MC	vSAN All Flash Config	1
B8JP	ThinkSystem 2.5" 5300 3.84TB Entry SATA 6Gb Hot Swap SSD	8
B5XH	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2
BMEY	VMware ESXi 7.0 U3 (Factory Installed)	1
BB3T	ThinkSystem 1U 10x2.5" AnyBay Backplane	1
B5SV	ThinkSystem Broadcom 57454 10/25GbE SFP28 4-port OCP Ethernet Adapter	1
B8N2	ThinkSystem 1U PCIe Gen4 x16/x16 Riser 1	1
B8NC	ThinkSystem 1U LP+LP BF Riser Cage Riser 1	1
B8QC	ThinkSystem 1100W (230V/115V) V2 Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
B8N4	ThinkSystem 1U Performance Fan Option Kit	8
BMPK	1	
B8GM	ThinkSystem SR645 on-board NVMe Cable Kit to Front 10x2.5" AnyBay 1	1
BC4T	ThinkSystem 1U SFF Front 2.5" Backplane SAS/SATA Gen3 Cable	1
BC4S	ThinkSystem SR645 SFF G3 Cable to Front 10x 2.5" AnyBay 6+4 AnyBay	1
B8H2	ThinkSystem SR645 on-board NVMe Cable Kit to Front 10x2.5" AnyBay 2	1
BE0E	N+N Redundancy With Over-Subscription	1
BNHS	ThinkAgile VX3375 IS	1

B13M	ThinkAgile EIA Plate	1
B1DS	VX Badge 1	1
AURR	ThinkSystem M3.5 Screw for Riser 2x2pcs and Planar 5pcs	2
AUWG	Lenovo ThinkSystem 1U VGA Filler	1
AWF9	ThinkSystem Response time Service Label LI	1
B8HU	ThinkSystem 2.5" PM1645a 800GB Mainstream SAS 12Gb Hot Swap SSD	2
7S06CTOEWW	VMware vSAN	1
S6AR	VMware HCI Kit 6 Enterprise (per CPU) w/Lenovo 3Yr S&S	2
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
AVE5	ThinkSystem 430-16i SAS/SATA 12Gb HBA placement	1
5MS7A87711	ThinkAgile VX Remote Deployment (up to 4 node cluster)	1
7S06CTOSWW	VMware vSAN for VX	1
S6J0	VMware vSAN 7 Advanced for Desktop 10 Pack (CCU) w/Lenovo 3Yr S&S	1

Table 26: ThinkAgile VX650 V3 Integrated System (4th Generation Intel Xeon SP)

Part number	Product Description	Qty
7D6WCTO1WW	Server : Lenovo ThinkAgile VX650 V3 Integrated System	1
BRY9	ThinkAgile VX V3 2U 24x2.5" Chassis	1
B0W3	XClarity Pro	1
BUE2	VMware vSAN 8 Enterprise	1
BPQF	Intel Xeon Gold 6426Y 16C 185W 2.5GHz Processor	2
BKTM	ThinkSystem 32GB TruDDR5 4800MHz (2Rx8) RDIMM	16
5977	Select Storage devices - no configured RAID required	1
B8P1	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA	1
B5MC	vSAN All Flash Config	1
B8LU	ThinkSystem 2U 8x2.5" SAS/SATA Backplane	1
BH8B	ThinkSystem 2U/4U 8x2.5" AnyBay Backplane	1
BM8X	ThinkSystem M.2 SATA/x4 NVMe 2-Bay Enablement Kit	1
BTTY	M.2 NVMe	1
BS7M	Intel VROC (VMD NVMe RAID) Standard for M.2	1
BS7F	M.2 NVMe Array 1 RAID 1	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2

BPPW	ThinkSystem Broadcom 57504 10/25GbE SFP28 4-Port OCP Ethernet Adapter	1
BPK9	ThinkSystem 1800W 230V Titanium Hot-Swap Gen2 Power Supply	2
BLL6	ThinkSystem 2U V3 Performance Fan Module	6
BQQ2	ThinkSystem 2U V3 EIA Latch Standard	1
BVMC	Trigger MFG to scan the SN from the CPU Board via this MI	1
BS7J	M.2 NVMe Array 1 HDDs	2
BTSL	ThinkAgile VX650 V3 IS	1
AVEQ	ThinkSystem 8x1 2.5" HDD Filler	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	8
BLL3	ThinkSystem SR650 V3 PSU Duct	1
BPDQ	ThinkSystem SR650 V3 AL Extrusion Entry Heatsink	2
BMPF	ThinkSystem V3 2U Power Cable from MB to Front 2.5" BP v2	2
BRQ3	ThinkSystem V3 2U WH CBL, 20Pin, 320mm,Tin-plated	1
BRPW	ThinkSystem SR650 V3,PCIe Gen4 CBL,SLx8-Swift,M.2-PCIe3	1
BACB	ThinkSystem V3 2U SAS/SATA Y Cable from CFF C0,C1/ C2,C3 to Front 8x2.5" BP	2
BSYM	ThinkSystem SR650 V3,PCIe4 Cable,Swift8x-SL8x,2in1,PCIe 6/5(MB) to BP1/BP2	1
BRPV	ThinkSystem SR650 V3,PCIe Gen4 CBL,SLx8-Swift,CFF IN-PCIe4	1
BPE3	ThinkSystem SR650 V3 MCIO8x to SL8x CBL, PCIe4, 8x2.5AnyBay, 200mm	2
BPK3	ThinkSystem WW Lenovo LPK	1
BQ11	G4 x16/x8/x8 PCIe Riser BLKL for Riser 1 Placement	1
BLKL	ThinkSystem V3 2U x16/x8/x8 PCIe Gen4 Riser1 or 2	1
BNW9	ThinkSystem 2.5" PM1655 1.6TB Mixed Use SAS 24Gb HS SSD	2
BNWF	ThinkSystem 2.5" PM1653 3.84TB Read Intensive SAS 24Gb HS SSD	6
7S06CTOSWW	VMware vSAN for VX	1
SAQT	VMware vSAN 8 Enterprise for 1 processor w/Lenovo 3Yr S&S	2
5PS7B14942	Essential ThinkAgile IS - 3Yr 24x7 4Hr Resp + YDYD VX650 V3	1

Table 27: ThinkAgile VX655 V3 Integrated System (4th Generation AMD EPYC)

Part number	Product Description	Qty
7D9WCTO2WW	Server : Lenovo ThinkAgile VX655 V3 Certified Node	1
BRY9	ThinkAgile VX V3 2U 24x2.5" Chassis	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
B0W3	XClarity Pro	1
BN8K	ThinkAgile VX Remote Deployment	1
BPVJ	ThinkSystem AMD EPYC 9554 64C 360W 3.1GHz Processor	1
BQ3D	ThinkSystem 64GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM-A	12
5977	Select Storage devices - no configured RAID required	1
BC4V	Non RAID NVMe	1
BT2G	vSAN ESA	1
BNF5	ThinkSystem 2.5" U.3 7450 PRO 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	4
BS7Y	ThinkSystem V3 2U 8x2.5" NVMe Gen5 Backplane	1
B8P9	ThinkSystem M.2 NVMe 2-Bay RAID Enablement Kit	1
BTTY	M.2 NVMe	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BQ8S	VMware ESXi 8.0 U1 (Factory Installed)	1
BPPW	ThinkSystem Broadcom 57504 10/25GbE SFP28 4-Port OCP Ethernet Adapter	1
BLKM	ThinkSystem V3 2U x16/x16/E PCIe Gen4 Riser1 or 2	1
BPK9	ThinkSystem 1800W 230V Titanium Hot-Swap Gen2 Power Supply	2
B8LA	ThinkSystem Toolless Slide Rail Kit v2	1
BQQ3	ThinkSystem 2U V3 EIA Latch with VGA Port	1
BLL6	ThinkSystem 2U V3 Performance Fan Module	6
BPKR	TPM 2.0	1
BQ80	ThinkSystem SR655 V3 MB	1
BS6Y	ThinkSystem 2U V3 M.2 Signal & Power Cable, SLx4 with 2X10/1X6 Sideband, 330/267/267mm	1
BSHZ	PCIe MB-BP Cable MCIO 1 2(MB) 12x3.5/8x2.5" 430mm to 8x2.5	1
BSJ2	PCIe MB-BP Cable MCIO 3 4(MB) 12x3.5/8x2.5" 340/390mm to 8x2.5	1
BMPF	ThinkSystem V3 2U Power Cable from MB to Front 2.5" BP v2	1
BQ12	G4 x16/x16/E PCIe Riser BLKM for Riser 1 Placement	1
B8MM	ThinkSystem 2U MS 3FH Riser Filler	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	4
BSR6	ThinkSystem SR635 V3/SR655 V3 RoT Module LV-RoW	1
BQ81	ThinkSystem SR655 V3 Main Airduct	1
BTT4	ThinkAgile SR655 V3 Service Label - WW	1
AWF9	ThinkSystem Response time Service Label LI	1
BTSW	ThinkAgile SR655 V3 Agency Label - No CCC	1
B986	ThinkSystem HV 2U WW General PKG BOM	1

BXGY	Right EIA with FIO assembly	1
BQQ6	ThinkSystem 2U V3 EIA right with FIO	1
AVEQ	ThinkSystem 8x1 2.5" HDD Filler	2
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
BQF9	ThinkSystem SR655 V3 2U High Performance Heatsink	1
B265	ThinkAgile VX Pubkit	1
BQ31	ThinkSystem MS 2U common Airductr Filler	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
5PS7B12166	Premier Essential ThinkAgile CN - 3Yr 24x7 4Hr Resp + YDYD VX655 V3	1
5AS7B12280	Hardware Installation (Business Hours) for VX655 V3	1
5MS7A87711	ThinkAgile VX Remote Deployment (up to 4 node cluster)	1
7S0XCTO5WW	XClarity Controller	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

Table 28: ThinkAgile VX650 V3 Certified Node (5th Gen Intel Xeon Scalable Processors)

Part number	Product Description	Qty
7D6XCTO2WW	Server : Lenovo ThinkAgile VX630 V3 Certified Node	1
BRYA	ThinkAgile VX V3 1U 10x2.5" Chassis	1
B0W3	XClarity Pro	1
B3XQ	3 Year	1
BN8K	ThinkAgile VX Remote Deployment	1
BYW8	Intel Xeon Gold 6538Y+ 32C 225W 2.2GHz Processor	2
BZ4V	ThinkSystem 64GB TruDDR5 Performance+ 5600MHz (2Rx4) 10x4 RDIMM	8
5977	Select Storage devices - no configured RAID required	1
B8P1	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA	1
BT2G	vSAN ESA	1
BYRM	AF-0	1
BNEY	ThinkSystem 2.5" U.3 7450 MAX 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	4
BRQY	ThinkSystem 1U 2.5" 6 SAS/SATA 4 AnyBay Gen5 Backplane	1
B8P9	ThinkSystem M.2 NVMe 2-Bay RAID Adapter	1
BTTY	M.2 NVMe	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BQ8S	VMware ESXi 8.0 U1 (Factory Installed)	1
BPPW	ThinkSystem Broadcom 57504 10/25GbE SFP28 4-Port OCP Ethernet Adapter	1
BLKB	ThinkSystem V3 1U x16/x16 BF PCIe Gen5 Riser1	1
BLK9	ThinkSystem V3 1U MS LP+LP BF Riser Cage	1
BPK9	ThinkSystem 1800W 230V Titanium Hot-Swap Gen2 Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
BH9M	ThinkSystem V3 1U Performance Fan Option Kit v2	8

BLKD	ThinkSystem 1U V3 10x2.5" Media Bay w/ Ext. Diagnostics Port	1
B8LA	ThinkSystem Toolless Slide Rail Kit v2	1
BPKR	TPM 2.0	1
BLK5	ThinkSystem SR630 V3 MB	1
BVMC	Trigger MFG to scan the SN from the CPU Board via this MI	1
9207	VMWare Specify	1
9220	Preload by Hardware Feature Specify	1
BRPJ	XCC Platinum	1
B7XZ	Disable IPMI-over-LAN	1
BK15	High voltage (200V+)	1
B8KM	ThinkSystem 1U 10x2.5" 6 SAS/SATA+4 AnyBay HDD Type Label	1
BPK3	ThinkSystem WW Lenovo LPK	1
B265	ThinkAgile VX Pubkit	1
BVMB	ThinkSystem SR630 V3 CFF G4 Cables, C0/C1 to SAS0/SAS1,10x 2.5" anyBay BP,175/150mm	1
BA20	ThinkSystem 1U CFF RAID to 10x2.5" Backplane SAS/SATA G4 Cable 1	1
BPCP	ThinkSystem SR630 V3 MCIO8x to SL8x PCIe4 Cable from PCIe3(MB) to CFF RAID, 400mm	1
BRLT	ThinkSystem MCIO8x flat CBL, PCIe5, NY PCIe1 to 10AnyBay NVMe 8-9, 360mm	1
BRLW	ThinkSystem MCIO8x CBL, PCIe5, NY PCIe2/3 to 10AnyBay NVMe 6-7/4-5,340mm	1
BRF3	ThinkSystem Cable132	1
BMJF	ThinkSystem MB to CFF RAID v2	1
BS3A	ThinkSystem SR645 V3 M.2 Signal & Power Cable, SLx4 with 2X10/1X6 sideband, 730/300/250mm	1
B8NB	ThinkSystem 1U MS LP Riser Filler	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	6
BTSK	ThinkAgile VX630 V3 CN	1
BQBP	ThinkSystem MCC CPU Clip	2
BP50	ThinkSystem SR630 V3 Performance Heatsink	2
BYQN	ThinkSystem SR630 V3 Firmware and Root of Trust Security Module v2	1
B984	ThinkSystem 1U PLV Top Cover Sponge	1
B989	ThinkSystem V2 1U Package	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
B8Q8	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA Placement	1
5PS7B14796	Premier Essential ThinkAgile CN - 3Yr 24x7 4Hr Resp + YDYD VX630 V3	1
7S0XCTO5WW	XClarity Controller	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

5.3 BOM for enterprise and SMB management servers

Table 17 on page 23 lists the number of management servers that are needed for the different numbers of users. To help with redundancy, the bill of materials for management servers must be the same as compute servers. For more information, see “BOM for enterprise and SMB compute servers” on page 30.

Table 29: ThinkSystem SR630V2 (3rd Generation Intel Xeon SP)

Part number	Product Description	Qty
7Z71A04WNA	Server: ThinkSystem SR630 V2, 1xIntel Xeon Silver 4314 16C 2.4GHz 135W, 1x32GB 2Rx8, SW RD, 1x750W, XCC Enterprise,	1
4XG7A63411	ThinkSystem SR630 V2 Intel Xeon Silver 4314 16C 135W 2.4GHz Processor Option Kit w/o Fan	1
4X77A08633	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	7
4XC7A08237	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
4F17A14487	ThinkSystem SR630 V2/SR645 Performance Fan Option Kit	2
5PS7A67541	Premier Essential - 3Yr 24x7 4Hr Resp + YDYD SR630 V2	1
7S06073EWW	VMware HCI 6 Kit Essentials for 3 Nodes (Max 2 processors per node) w/Lenovo 3Yr S&S	1
00MT202	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

Table 30: ThinkSystem SR650V2 (3rd Generation Intel Xeon SP)

Part number	**Product Description**	**Qty**
7Z73A06FNA	Server : ThinkSystem SR650 V2, 1xIntel Xeon Silver 4314 16C 2.4GHz 135W, 1x32GB 2Rx8, SW RD, 1x750W, XCC Enterprise,	1
4XG7A63455	ThinkSystem SR650 V2 Intel Xeon Silver 4314 16C 135W 2.4GHz Processor Option Kit w/o Fan	1
4X77A08634	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx8 1.2V) RDIMM	7
4Y37A09739	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
4XB7A17073	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	1
4XC7A08237	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
4F17A14496	ThinkSystem SR650 V2 Performance Fan Option Kit	1
4X97A80440	ThinkSystem SR650 V2/SR665 M.2 Cable Kit v2	1
5WS7A67902	Premier Foundation - 3Yr NBD Resp SR650 V2	1
00MT202	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1

5.4 BOM for shared storage

The following two tables list the BOMs for Lenovo ThinkSystem DE6000F storage array.

Table 31: Lenovo ThinkSystem DE6000F control enclosure

Code	Product Description	Qty
7Y79CTO1WW	Storage : Lenovo ThinkSystem DE6000F All Flash Array SFF	1
B38L	Lenovo ThinkSystem Storage 2U24 Chassis	1
B4J9	Lenovo ThinkSystem DE6000 12Gb SAS 4-ports HIC	2
B4JP	Lenovo ThinkSystem DE6000 Controller 64GB	2
B4BV	Lenovo ThinkSystem DE Series 1.6TB 3DWD 2.5” SSD SED FIPS 140-2 2U24	24
B4BP	Lenovo ThinkSystem Storage USB Cable, Micro-USB	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
B4JD	Lenovo ThinkSystem DE6000F Premium Bundle	1
B38Y	Lenovo ThinkSystem Storage Rack Mount Kit 2U24/4U60	1
B4AR	Lenovo ThinkSystem DE Series Ship Kit (RoW), 2U	1
B4M0	Lenovo ThinkSystem DE6000H SMID Controller Base Setting	1
B4AW	Lenovo ThinkSystem Storage Packaging 2U	1

Table 32: Lenovo ThinkSystem DE6000F optional software licenses

Code	**Description**	**Qty**
B4JL	Lenovo ThinkSystem DE6000 Add Snapshot 2048 PFK	1
B4JH	Lenovo ThinkSystem DE6000H Add Synch Mirroring PFK	1
B4JG	Lenovo ThinkSystem DE6000H Add Asynch Mirroring PFK	1
B4JJ	Lenovo ThinkSystem DE6000H Controller Upgrade Key FC to DE6000F	1

Resources

For more information, see the following resources:

- Lenovo Client Virtualization reference architecture

  lenovopress.com/lp0756

- VMware vSphere

  VMware.com/products/datacenter-virtualization/vsphere

- VMware Horizon (with View)

  https://www.VMware.com/products/horizon.html

- VMware vSAN

  VMware.com/products/vsan.html

- VMware vSAN design and sizing guide

  https://core.VMware.com/resource/VMware-vsan-design-guide

- Flex System Interoperability Guide

  lenovopress.com/redpfsig-flex-system-interoperability-guide

- Lenovo Storage Interoperability Links

  lenovopress.com/lp0584-lenovo-storage-interoperability-links

Document History

Version 1.0	18 September 2017	• Initial version
Version 1.1	8 January 2018	• Added section on GPU graphics acceleration • Replaced G8124E RackSwitch with ThinkSystem NE1032 switch • Added ThinkSystem SD530 server
Version 1.2	22 June 2018	• Added NVIDIA M10 and P40 GPUs
Version 1.3	5 November 2018	• Added results for ESXi 6.7 using vSAN 6.7
Version 1.4	30 March 2020	• Added Windows 10 version 1809 and Horizon 7.11 results for ThinkAgile VX Appliances with 2nd Generation Xeon SP processors using ESXi 6.7 U3
Version 1.5	10 June 2020	• Removed Windows 10 version 1803 results for 1st Generation Xeon SP processors • Added Windows 10 version 1909 and Horizon 7.11 results with 2nd Generation Xeon SP processors for ThinkAgile VX Appliances, ThinkSystem SR850 and SR850P servers • Added Windows 10 version 1909 and Horizon 7.11 results with 1st Generation Xeon SP processors for ThinkSystem SR630 • Removed DS6200 storage (withdrawn) and added information about DE and DM series storage arrays • Added ThinkSystem SR635/SR655 AMD EPYC processor systems
Version 1.6	16 Nov 2020	• Added section 4.3.3 Lenovo ThinkSystem with AMD EPYC (Rome) single and dual node VSImax performance
Version 1.7	27 Sep 2022	• Section 4.3.5 Added Login Enterprise benchmark results for ThinkSystem with Intel Xeon SP3 • Section 4.4.2 Added LoginVSI results for Lenovo ThinkAgile VX with Intel 3rd Generation Xeon Scalable Processor (Ice Lake) • Section 4.4.3 Added Windows 11 Performance data with Lenovo ThinkAgile VX • Section 4.4.5 Added LoginVSI results for Lenovo ThinkSystem with 3rd Generation AMD EPYC (Milan) • Section 4.10 Removed information on Lenovo TOR Switches and DS storage

		• Section 4.12 Updated deployment example with Xeon SP3 processors and ThinkAgile VX • Section 5.1, 5.2, 5.3 – Updated with 3rd Generation Intel Xeon SP and 3rd Generation AMD EPYC processors. • Section 5.5 Removed BOM for networking & Rack
Version 1.8	20 December 2022	• Section 4.3.5 – Updated with Login Enterprise 4.9.10 results
Version 1.9	28 June 2023	• Section 3.3.1 – Added vSAN Express Storage Architecture • Section 4.4.3 – Added Login Enterprise results for Intel 4th Gen Scalable Architecture and vSAN Express Storage Architecture • Section 5.2 - Added BOM for ThinkAgile VX V3 systems with Intel 4th Gen Xeon SP
Version 1.10	25 September 2023	• Removed LoginVSI Pro results for sections 4.3.1, 4.3.2, 4.3.3, 4.3.4, 4.4.1, 4.4.4 • Removed Boot Storm Performance section 4.4.7 • Section 4.3.3 Added Login Enterprise results SR850 V3 with Intel 4th Gen SP • Section 4.4.4 Added Login Enterprise results for AMD Genoa • Added BOM for ThinkAgile VX655 V3
Version 1.11	05 March 2024	• Section 4.3.5 added Login Enterprise results for 5th Gen Intel Scalable Processors • Section 4.4.4 Updated Login Enterprise results for AMD Genoa • Section 5.2 BOM is added for ThinkAgile VX630 V3 with 5th Gen Intel Scalable Processors

Trademarks and special notices

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

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

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

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

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

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.