



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

Last update: 16 Nov 2020

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, networking, and racks

Chandrakandh Mouleeswaran

Cristian Ghetau

Srihari Angaluri



Table of Contents

1	Introduction	1
2	Architectural overview	2
3	Component model	3
3.1	VMware Horizon provisioning	5
3.2	Storage model	5
3.3	VMware vSAN	6
3.3.1	VMware vSAN storage policies	6
4	Operational model	9
4.1	Operational model scenarios	9
4.1.1	Enterprise operational model	10
4.1.2	Hyper-converged operational model	10
4.2	Hypervisor support	11
4.3	Compute servers for virtual desktops	11
4.3.1	Lenovo ThinkSystem with Intel Xeon Scalable Processor Gen 1 (Skylake)	12
4.3.2	Lenovo ThinkSystem with Intel Xeon Scalable Processor Gen 2 (Cascadelake)	14
4.3.3	Lenovo ThinkSystem with AMD EPYC (Rome)	15
4.4	Compute servers for VMware vSAN	16
4.4.1	ThinkAgile VX3320 with Intel Xeon Scalable Processor Gen 2 (Cascadelake)	17
4.4.2	ThinkSystem SR635/SR655 with AMD EPYC™ Processors (Rome)	19
4.4.3	vSAN Boot Storm Performance	20
4.5	Graphics acceleration	21
4.6	Management servers	23
4.7	Systems management	24
4.8	Shared storage	26
4.9	Lenovo ThinkSystem DM and DE Series Storage Arrays	27
4.9.1	Lenovo ThinkSystem DM Storage	27
4.9.2	Lenovo ThinkSystem DE Storage	27
4.9.3	Storage recommendations	29
4.10	Networking	29
4.10.1	10 GbE networking	30
4.10.2	1 GbE administration networking	30

4.11	Racks.....	31
4.12	Deployment models	31
4.12.1	Deployment example 1: vSAN hyper-converged using ThinkSystem servers.....	31
5	Appendix: Bill of materials.....	33
5.1	BOM for enterprise and SMB compute servers.....	33
5.2	BOM for hyper-converged compute servers	36
5.3	BOM for enterprise and SMB management servers	42
5.4	BOM for shared storage.....	44
5.5	BOM for networking	45
5.6	BOM for Flex System chassis.....	46
5.7	BOM for rack.....	46
	Resources	47

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 describes the reference architecture for VMware Horizon 7.x. This document should be read with the Lenovo Client Virtualization (LCV) base reference architecture document that is available at this website: lenovopress.com/lp0756.

The business problem, business value, requirements, and hardware details are described in the LCV base reference architecture document and are not repeated here for brevity.

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.

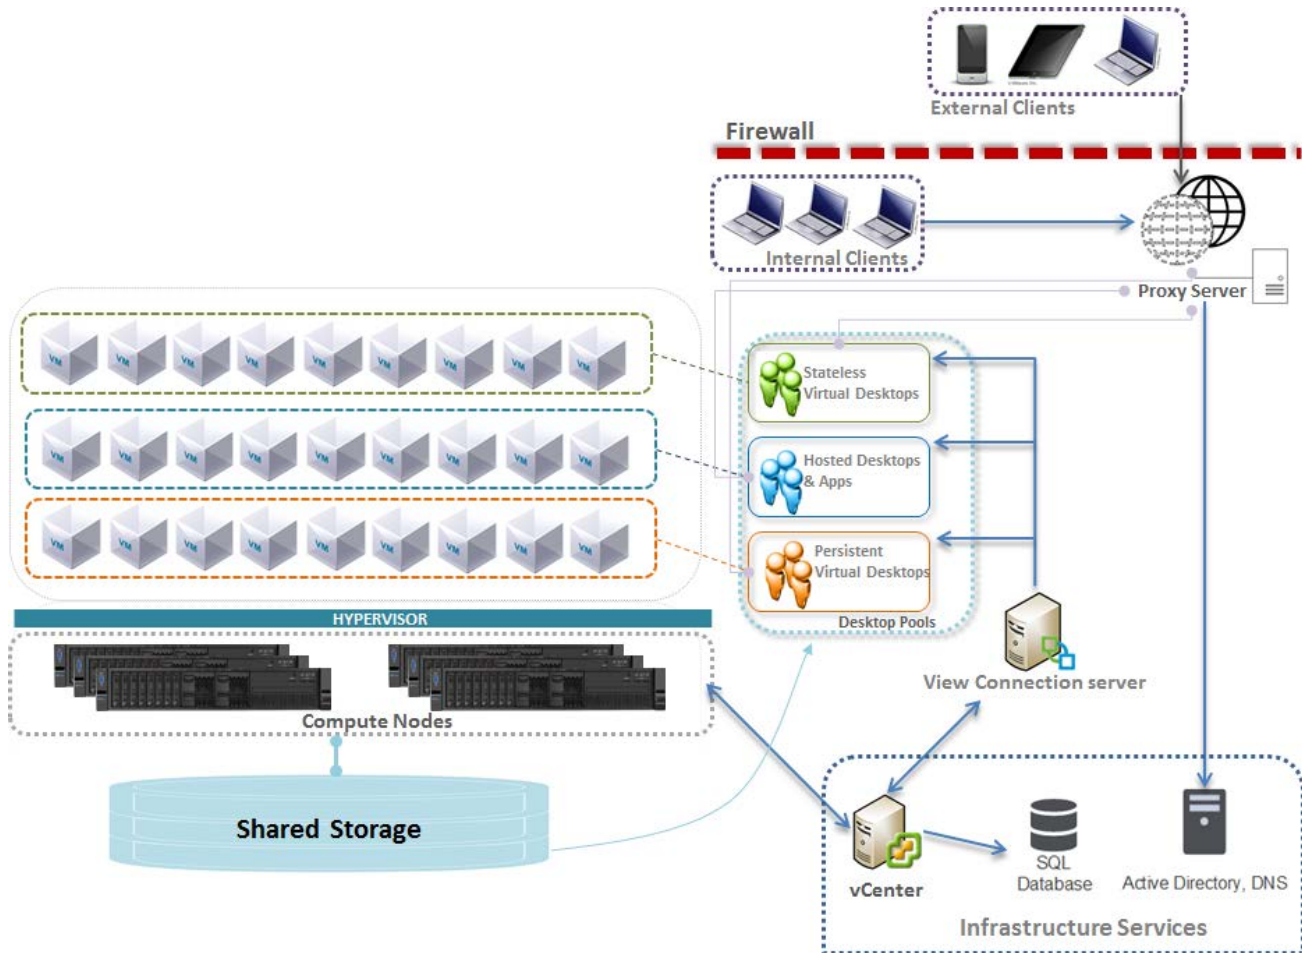


Figure 1: LCV reference architecture with VMware Horizon

3 Component model

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

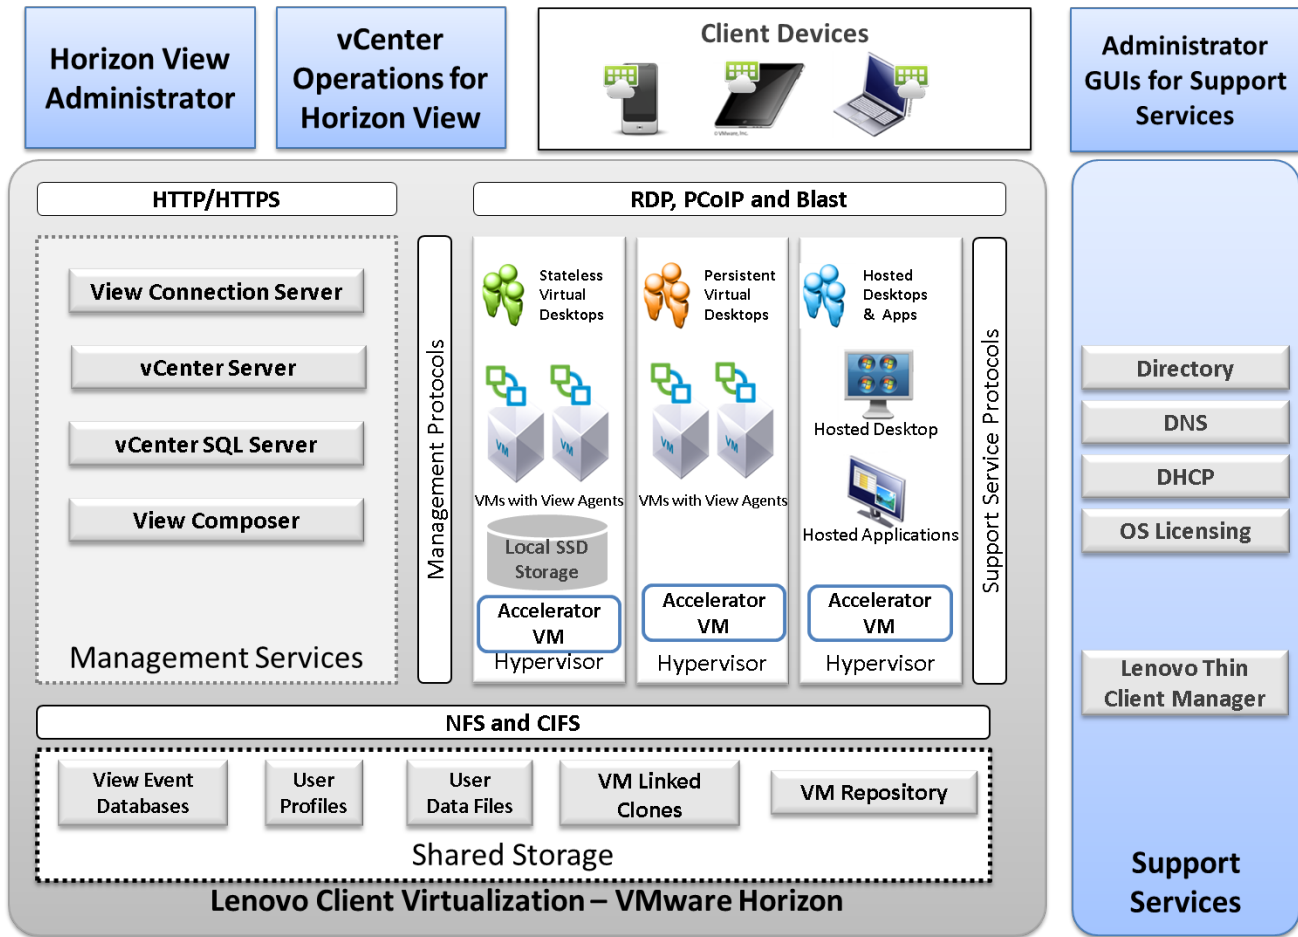


Figure 2: Component model with VMware Horizon

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	<p>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.</p> <p>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.</p>
View Composer	<p>In a VMware vCenter Server instance, View Composer is installed. View Composer is required when linked clones are created from a parent VM.</p>
vCenter Server	<p>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.</p>
vCenter SQL Server	<p>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.</p>
View Event database	<p>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.</p>
Clients	<p>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.</p>
RDP, PCoIP	<p>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).</p>
Hypervisor ESXi	<p>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.</p>
Accelerator VM	<p>The optional accelerator VM.</p>
Shared storage	<p>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.</p>

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](https://www.vmware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf)

3.3.1 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 reserved automatically. When deduplication and compression is enabled, this should be set to either 0% (do not apply) or 100%.	0%	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.	No	Yes

	Object checksums carry a small disk IO, memory and compute overhead and can be disabled on a per object basis.		
--	--	--	--

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

Storage Policy	Dedicated				Stateless		
	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 33. 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.

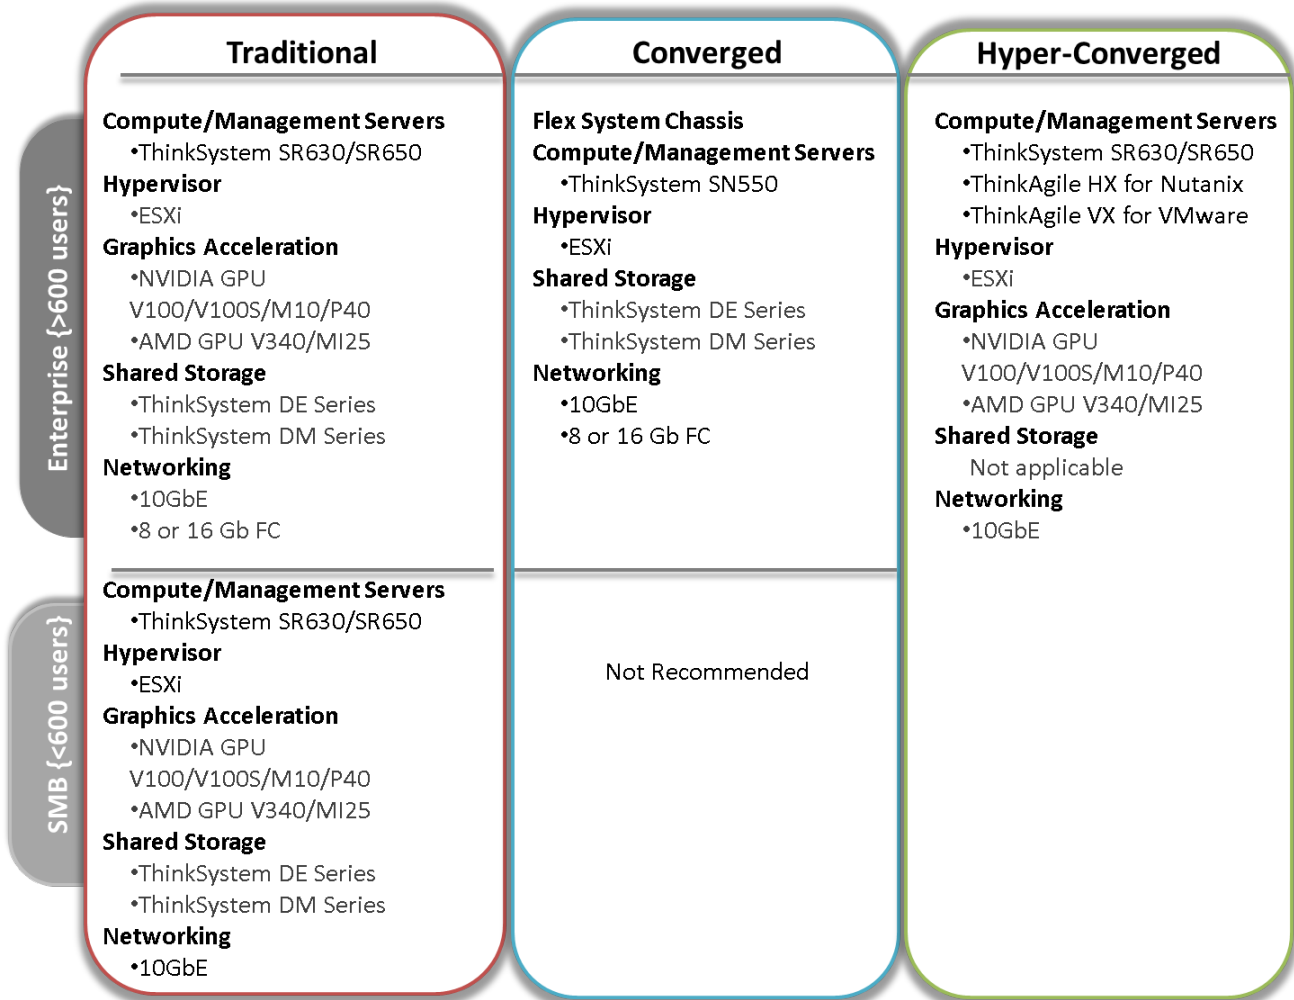


Figure 3: Operational model scenarios

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 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 Hypervisor support

This reference architecture was tested with VMware ESXi 6.7 U3 hypervisor with Intel Spectre and Meltdown security patches.

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. For the VMware ESXi hypervisor, the number of vCenter clusters depends on the number of compute servers.

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

https://my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_7#custom_iso
my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5#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 33.

4.3.1 Lenovo ThinkSystem with Intel Xeon Scalable Processor Gen 1 (Skylake)

This section shows the performance results and sizing guidance for Lenovo ThinkSystem compute servers based on the Intel Scalable processor Gen 1 (Purley Skylake). The testing was done with single ESXi 6.7u3 node using 2x 800GB Local SSDs. Table 3 lists the VSI max performance results for different Login VSI 4.1 workloads with ESXi 6.7 U3 and Windows 10 version 1909.

Table 3: Login VSI Performance for Intel Xeon SP1

Processor	Workload	Stateless	Dedicated
Two Scalable 6130 processors 2.10 GHz, 16C 125W	Office worker	203 users	204 users
Two Scalable 6130 processors 2.10 GHz, 16C 125W	Knowledge worker	169 users	176 users
Two Scalable 8176 processors 2.10 GHz, 28C 165W	Knowledge worker	260 users	268 users
Two Scalable 6130 processors 2.10 GHz, 16C 125W	Power worker	160 users	160 users
Two Scalable 8176 processors 2.10 GHz, 28C 165W	Power worker	241 users	250 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 Scalable 8176 processor has significantly better performance than the Scalable 6130 processor and is directly proportional to the additional 12 cores but at an added cost.
- The LoginVSI benchmark is CPU intensive and the local SSDs used for the testing was enough for linked clones and no storage latency issues noticed. Please refer the shared storage section “4.8 Shared storage for sizing IOPS and capacity.

The Xeon Scalable 6130 processor has a good cost/user density ratio. Many configurations are bound by memory; therefore, a faster processor might not provide any added value. Some users require a very fast processor and for those users, the Xeon Scalable 8176 processor is a good choice.

The default recommendation is two Xeon Scalable 6130 processors and 512 GB of system memory. The memory can be increased based on the user requirements, but it cannot increase user density as the Windows 10 results show that it is compute bound.

For an office worker, Lenovo testing shows that 136 users per server is a good baseline and has an average of 75% usage of the processors in the server. If a server goes down, users on that server must be transferred to the remaining servers. For this degraded failover case, Lenovo testing shows that 166 users per server has an average of 84% usage of the processors. It is important to keep a 25% or more headroom on servers to cope with possible failover scenarios. Lenovo recommends a general failover ratio of 5:1.

For a knowledge worker, Lenovo testing shows that 110 users per server is a good baseline and has an average of 75% usage of the processors in the server. For the degraded failover case, Lenovo testing shows that 132 users per server have an average of 73% usage of the processors.

For a power user, Lenovo testing shows that 104 users per server is a good baseline and has an average of 75% usage of the processors in the server. For the degraded failover case, Lenovo testing shows that 126 users per server have an average of 87% usage of the processors. Moving up to two Xeon Scalable 8176

processors shows those densities increasing to 170 users and 204 users respectively with lower CPU utilizations.

Table 4 summarizes the processor usage with ESXi for the recommended user counts for normal mode and failover mode.

Table 4: Processor usage

Processor	Workload	Normal Mode Users per Server (CPU Utilization%)	Failover Mode (5+1) Users per Server (CPU Utilization %)
Two Scalable 6130	Office worker	136 Users (75%)	166 Users (84%)
Two Scalable 6130	Knowledge worker	110 Users (75%)	132 Users (82%)
Two Scalable 6130	Power worker	104 Users (75%)	126 Users (87%)
Two Scalable 8176	Knowledge worker	176 Users (75%)	212 Users (86%)
Two Scalable 8176	Power worker	170 Users (75%)	204 Users (85%)

Table 5 lists the recommended number of virtual desktops per server with 2x6130 processors 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 5: Recommended number of virtual desktops per server

Workload	Office worker	Knowledge worker	Power worker
Processor	Two 6130	Two 6130	Two 6130
VM memory size	3 GB	4 GB	5 GB
System memory	512 GB	512 GB	512 GB
Desktops per server (normal mode)	136	110	100
Desktops per server (failover mode)	166	132	120

Table 6 lists the approximate number of compute servers that are needed for different numbers of users and Office worker workloads.

Table 6: Compute servers needed for Office workers and different numbers of users

Office workers	300 users	600 users	1200 users	3000 users
Compute servers @136 users (normal)	3	5	9	22
Compute servers @166 users (failover)	2	4	8	18

Table 7 lists the approximate number of compute servers that are needed for different numbers of users and Knowledge worker workloads.

Table 7: Compute servers needed for Knowledge workers and different numbers of users

Knowledge workers	300 users	600 users	1200 users	3000 users
Compute servers @110 users (normal)	4	7	11	27

Compute servers @132 users (failover)	3	5	10	23
---------------------------------------	---	---	----	----

Table 8 lists the approximate number of compute servers that are needed for different numbers of users and Power worker workloads.

Table 8: Compute servers needed for Power workers and different numbers of users

Power workers	300 users	600 users	1200 users	3000 users
Compute servers @104 users (normal)	4	6	12	29
Compute servers @126 users (failover)	3	5	10	24

4.3.2 Lenovo ThinkSystem with Intel Xeon Scalable Processor Gen 2 (Cascadelake)

This section shows the performance results and sizing guidance for Lenovo ThinkSystem compute servers based on the Intel Scalable processor Gen 2 (Purley Cascadelake). Table 9 lists the VSI max performance results for different Login VSI 4.1 workloads with ESXi 6.7 U3 and Windows 10 version 1909. The Gold 6230 processors node was tested with ThinkSystem DE6000 All Flash storage and Platinum 8260M nodes were tested with Local SSDs configured with RAID5.

Table 9: Login VSI Performance for Intel SP 2

Server	Processor	Workload	Stateless Users	Dedicated Users
ThinkSystem SR850	4x Platinum 8260M processors 2.40 GHz, 24C 125W	Office worker	650	671
ThinkSystem SR850	4x Platinum 8260M processors 2.40 GHz, 24C 125W	Knowledge worker	570	585
ThinkSystem SR850P	2x Platinum 8260M processors 2.40 GHz, 24C 125W	Knowledge worker	618	639
ThinkSystem SR630	2x Gold 6230 processors 2.10 GHz, 20C 125W	Knowledge worker	-	238
ThinkSystem SR850	4x Platinum 8260M processors 2.40 GHz, 24C 125W	Power worker	543	551
ThinkSystem SR850P	4x Platinum 8260M processors 2.40 GHz, 24C 125W	Power worker	564	590

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

- The performance for dedicated is little higher than stateless and it could be related to caching of user profile and application data for dedicated desktops.
- The Platinum 8260M processor has significantly better performance than the Gold 6230 processor and is directly proportional to the additional 4 cores and higher processor speed but at an added cost.
- The ThinkSystem SR850P results are 5-10% higher than ThinkSystem SR850 and it can be correlated to the mesh topology design followed in SR850P.

- The ThinkSystem SR630 with 6230 CPU LoginVSI result on Shared Storage is similar to ThinkAgile VX results on vSAN with same CPU. It shows compute performance does not change as long as storage is not a bottleneck. The compute sizing would not change for vSAN and Shared Storage, so please refer the section “4.4.1 ThinkAgile VX3320 with Intel Xeon Scalable Processor Gen 2 (Cascadelake)” sizing compute nodes with Gold 6230 Processor. For sizing shared storage, please refer “4.8 Shared storage”.

The Xeon SP2 Gold 6230 processor has a good cost/user density ratio. Some users require a very fast processor and high density and for those users, the 4 socket ThinkSystem SR850/SR850P with Xeon SP2 Platinum 8260 processor is a good choice. The default recommendation is 2x Gold 6230 processors and 768GB of system memory which provides configuration to meet different virtual desktop users workloads.

Table 10 summarizes the processor usage with ESXi for the recommended user counts for normal mode and failover mode.

Table 10: Processor usage

Server	Processor	Workload	Normal Mode Users per Server (CPU Utilization%)	Failover Mode (5+1) Users per Server (CPU Utilization %)
ThinkSystem SR850	4x Platinum 8260M	Office worker	360 (75%)	432 (85%)
ThinkSystem SR850	4x Platinum 8260M	Knowledge worker	328 (75%)	396 (86%)
ThinkSystem SR850P	4x Platinum 8260M	Knowledge worker	344 (75%)	414 (84%)
ThinkSystem SR630	2x Gold 6230	Knowledge worker	140 (75%)	168 (83%)
ThinkSystem SR850	4x Platinum 8260M	Power worker	300 (74%)	360 (86%)
ThinkSystem SR850P	4x Platinum 8260M	Power worker	324 (74%)	390 (84%)

4.3.3 Lenovo ThinkSystem with AMD EPYC (Rome)

This section presents the single node and dual node performance for Lenovo ThinkSystem hyper-converged servers based on AMD EPYC processors in conjunction with VMWare ESXi 6.7 U3 and VMWare Horizon 7.11. The tests were done using single and dual socket nodes Lenovo ThinkSystem SR655 and ThinkSystem SR665 server with 32 cores AMD EPYC 7452 CPU and 64 core AMD EPYC 7742 CPU with 1024GB RAM. The local storage consists of 2xNVMe drives (745GB per drive) and 6xSSD SAS drives (1.75TB per drive).

Table 11 presents the single node VSImax performance for Lenovo ThinkSystem SR655 (2U single socket server) with 32-core AMD EPYC 7452 CPU.

Table 11: Login VSI Performance with ESXi 6.7 U3 (AMD EPYC 7452 32c Single Node)

Server	Processor	Workload	Stateless	Dedicated
1xThinkSystem SR655	1x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Knowledge worker	239 users	237 users
1xThinkSystem SR655	1x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Power worker	205 users	193 users

Table 12 presents the single node VSI_{max} performance for Lenovo ThinkSystem SR655 (2U single socket server) with 1x64-core AMD EPYC 7742 CPU.

Table 12: Login VSI Performance with ESXi 6.7 U3 (AMD EPYC 7742 64c Single Node)

Server	Processor	Workload	Stateless	Dedicated
1xThinkSystem SR655	1x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Knowledge worker	395 users	380 users
1xThinkSystem SR655	1x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Power worker	316 users	312 users

Table 13 lists the VSI_{max} performance for Lenovo ThinkSystem SR665 (2U Dual-socket server) with 2x 32-core AMD 7452 CPUs.

Table 13: Login VSI Performance with ESXi 6.7 U3 (AMD EPYC 7452 32c Dual Socket Node)

Server	Processor	Workload	Stateless	Dedicated
1xThinkSystem SR665	2x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Knowledge worker	496 users	499 users
1xThinkSystem SR665	2x AMD EPYC 7452 (125W) 32C (64 threads) @ 2.35GHz	Power worker	421 users	428 users

4.4 Compute servers for VMware vSAN

This section presents the compute servers for the hyper-converged system using VMware vSAN 6.7 with ESXi 6.7 U3 for Intel Xeon Scalable Processor Gen 2 (Cascade lake). Additional processing and memory are required to support vSAN. Systems 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.

For more information, see “BOM for hyper-converged compute servers” on page 36.

4.4.1 ThinkAgile VX3320 with Intel Xeon Scalable Processor Gen 2 (Cascadelake)

VMware vSAN 6.7 U3 All Flash is tested by using Horizon 7.11 and Login VSI 4.1.28. Four Lenovo ThinkAgile VX3320 Appliances with two Xeon Scalable Gold 6230 processors and 1536 GB (24x64GB DIMM) of system memory were networked together by using a 10 GbE TOR switch with two 10 GbE connections per server.

Each server was configured with two disk groups because a single disk group does not provide the necessary resiliency. Each disk group had one 800 GB SSD for cache tier and three 1.84 GB SSDs for capacity tier. This provided more than enough capacity for linked clone or full clone VMs. Additional 1.84 GB HDDs could be added for more capacity.

Table 11 lists the per server VSI max performance results for different Login VSI 4.1 workloads with ESXi 6.7 U3 and Windows 10 version 1909. Linked clones were used with the VMware vSAN default storage policy of number of failures to tolerate (FTT) of 0 and stripe of 1. No de-duplication or compression was used.

Table 11: Login VSI Performance with ESXi 6.7 U3 (4 Node vSAN Cluster)

Server	Processor	Workload	Stateless	Dedicated
4xThinkAgile VX3320	2x Gold 6230 processors 2.10 GHz, 20C 125W	Office worker	1005 users (251 per node)	1040 users (260 per Node)
4xThinkAgile VX3320	2x Gold 6230 processors 2.10 GHz, 20C 125W	Knowledge worker	890 users (222 per node)	935 users 234 per node
4xThinkAgile VX3320	2x Gold 6230 processors 2.10 GHz, 20C 125W	Power worker	827 users (206 per node)	857 users (214 per node)

The Xeon Scalable Gold 6230 processor has a good cost/user density ratio. The persistent desktop provides around 5% better performance than stateless on vSAN and it could be related to profile and application data caching.

The default recommendation is two Xeon Gold 6230 processors and 768 GB of system memory because this configuration provides the best coverage and density for a range of users.

For an office worker, Lenovo testing shows that 170 persistent users per server is a good baseline and has an average of 75% usage of the processors in the server. If a server goes down, users on that server must be transferred to the remaining servers. For this degraded failover case, Lenovo testing shows that 204 users per server have an average of 87% usage of the processors. It is important to keep this 25% headroom on servers to cope with possible failover scenarios. Lenovo recommends a general failover ratio of 5:1. By using a target of 170 users per server, the theoretical maximum number of office workers is 10,880 in a 64 node cluster.

For a knowledge worker, Lenovo testing shows that 146 persistent users per server is a good baseline and has an average of 76% usage of the processors in the server. For the degraded failover case, Lenovo testing shows that 175 users per server have an average of 85% usage of the processors. By using a target of 146 users per server, the theoretical maximum number of knowledge workers is 9344 in a 64 node cluster.

For a power user, Lenovo testing shows that 125 persistent users per server is a good baseline and has an average of 75% usage of the processors in the server. For the degraded failover case, Lenovo testing shows that 150 users per server have an average of 84% usage of the processors. By using a target of 125 users per server, the theoretical maximum number of knowledge workers is 8,000 in a 64 node cluster.

Table 12 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 12: Recommended number of virtual desktops per server

Workload	Office worker	Knowledge worker	Power worker
Processor	Two 6230	Two 6230	Two 6230
VM memory size	3 GB	4 GB	5 GB
System memory	768 GB	768 GB	768 GB
Memory overhead of vSAN	32 GB	32 GB	32 GB
Desktops per server (normal mode)	170	146	125
Desktops per server (failover mode)	204	175	150

Table 13 lists the approximate number of compute servers that are needed for different numbers of users and Office worker workloads.

Table 13: Compute servers needed for Office workers and different numbers of users

Office workers	300 users	600 users	1200 users	3000 users
Compute servers @170 users (normal)	3	4	7	18
Compute servers @204 users (failover)	2	3	6	15

Table 14 lists the approximate number of compute servers that are needed for different numbers of users and Knowledge worker workloads.

Table 14: Compute servers needed for Knowledge workers and different numbers of users

Knowledge workers	300 users	600 users	1200 users	3000 users
Compute servers @146 users (normal)	3	5	9	21
Compute servers @175 users (failover)	2	4	7	18

Table 15 lists the approximate number of compute servers that are needed for different numbers of users and Power worker workloads.

Table 15: Compute servers needed for Power workers and different numbers of users

Power workers	300 users	600 users	1200 users	3000 users
Compute servers @125 users (normal)	3	6	10	24
Compute servers @150 users (failover)	2	4	8	20

4.4.2 ThinkSystem SR635/SR655 with AMD EPYC™ Processors (Rome)

The new Lenovo ThinkSystem SR635 and ThinkSystem SR655 introduce the AMD EPYC 2nd generation 7002 series “Rome” processors. The EPYC processor is designed for high performing enterprise workloads which benefit from a high number of processor cores, high-speed memory, and scalable I/O via PCIe Gen 4 interconnect. Because of the EPYC processor’s high core density, systems with a single processor socket provide excellent compute capability, which supports a variety of workloads that typically require 2-socket servers. More information on the Lenovo ThinkSystem SR655 and SR635 can be found at the following locations:

<https://lenovopress.com/lp1161-thinksystem-sr655-server>

<https://lenovopress.com/lp1160-thinksystem-sr635-server>

This section shows the performance results and sizing guidance for Lenovo ThinkSystem compute servers based on the AMD EPYC processor (Rome).

VMware ESXi 6.7 U3 performance results

VMware vSAN 6.7 U3 All Flash is tested by using Horizon 7.11 and Login VSI 4.1.28. Four Lenovo ThinkSystem SR655 servers with AMD EPYC 7452 (125W) 32 – Core processors (64 threads) and 1024GB (16x64GB DIMM) of system memory were networked together by using a 10 GbE TOR switch with two 10 GbE connections per server.

Each server was configured with two disk groups because a single disk group does not provide the necessary resiliency. Each disk group had one 750 GB NVMe SSD for cache tier and three 1.84 GB SSDs for capacity tier. This provided more than enough capacity for linked clone or full clone VMs. Additional 1.84 GB HDDs could be added for more capacity.

Table 16 lists the per server VSI max performance results for different Login VSI 4.1 workloads with ESXi 6.7 U3 and Windows 10 version 1909. Linked clones were used with the VMware vSAN default storage policy of number of failures to tolerate (FTT) of 0 and stripe of 1. No de-duplication or compression was used.

Table 16: Login VSI Performance with ESXi 6.7 U3 (4 Node vSAN Cluster)

Server	Processor	Workload	Stateless	Dedicated
4xThinkSystem SR655	1x AMD EPYC 7452 (155W) 32 - Core (64 threads) @ 2.35GHz	Knowledge worker	879 users (220 per node)	880 users (220 per node)
4xThinkSystem SR655	1x AMD EPYC 7452 (155W) 32 - Core (64 threads) @ 2.35GHz	Power worker	710 users (178 per node)	717 users (180 per node)

The AMD EPYC 7452 processor has a good cost/user density ratio. The persistent desktop provides around the same performance as stateless on vSAN.

The default recommendation is one AMD EPYC 7452 processors and 1024 GB of system memory because this configuration provides the best coverage and density for a range of users.

Table 17 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 17: Recommended number of virtual desktops per server

Workload	Knowledge worker	Power worker
Processor	One EPYC 7452	One EPYC 7452
VM memory size	4 GB	5 GB
System memory	1024 GB	1024 GB
Memory overhead of vSAN	32 GB	32 GB
Desktops per server (normal mode)	160	135
Desktops per server (failover mode)	198	160

Table 18 lists the approximate number of compute servers that are needed for different numbers of users and Knowledge worker workloads.

Table 18: Compute servers needed for Knowledge workers and different numbers of users

Knowledge workers	300 users	600 users	1200 users	3000 users
Compute servers @160 users (normal)	2	4	8	19
Compute servers @198 users (failover)	2	3	6	15

Table 19 lists the approximate number of compute servers that are needed for different numbers of users and Power worker workloads.

Table 19: Compute servers needed for Power workers and different numbers of users

Power workers	300 users	600 users	1200 users	3000 users
Compute servers @135 users (normal)	3	5	9	22
Compute servers @160 users (failover)	2	4	8	19

4.4.3 vSAN Boot Storm Performance

A boot storm occurs when a substantial number of VMs are all started within a short period of time. Booting a large number of VMs simultaneously requires large IOPS otherwise the VMs become slow and unresponsive.

Different numbers of VMs were booted on a cluster of 4 vSAN all flash servers. The VMs were unpowered in vCenter and the boot storm created by powering on all of the VMs simultaneously. The time for all of the VMs to become visible in VMware Horizon was measured.

Figure 4 shows the boot times for different numbers of VMs in the 4 node cluster with vSAN 6.6 and ESXi 6.5 U1. With an even spread of VMs on each node, the boot time for the VMs on each node was similar to the overall cluster boot time. These results show that even a large number of VMs can be very quickly restarted using a vSAN all flash cluster.




Figure 4: Boot storm performance

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 20 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 7.2, DirectX 11, and ESXi 6.5U1 on Xeon SP1 CPU 6130.

Table 20: 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 21 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 21: 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 33.
- “BOM for hyper-converged compute servers” on page 36.

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 22 lists the VM requirements and performance characteristics of each management service.

Table 22: Characteristics of VMware Horizon and ESXi management services

Management service VM	Virtual processors	System memory	Storage	Windows OS	HA needed	Performance characteristic
vCenter Server	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 23 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 23: 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 24 lists an example mapping of the management VMs to the two physical management servers for 3000 stateless users.

Table 24: 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 25 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 25: 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 42

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 5 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.

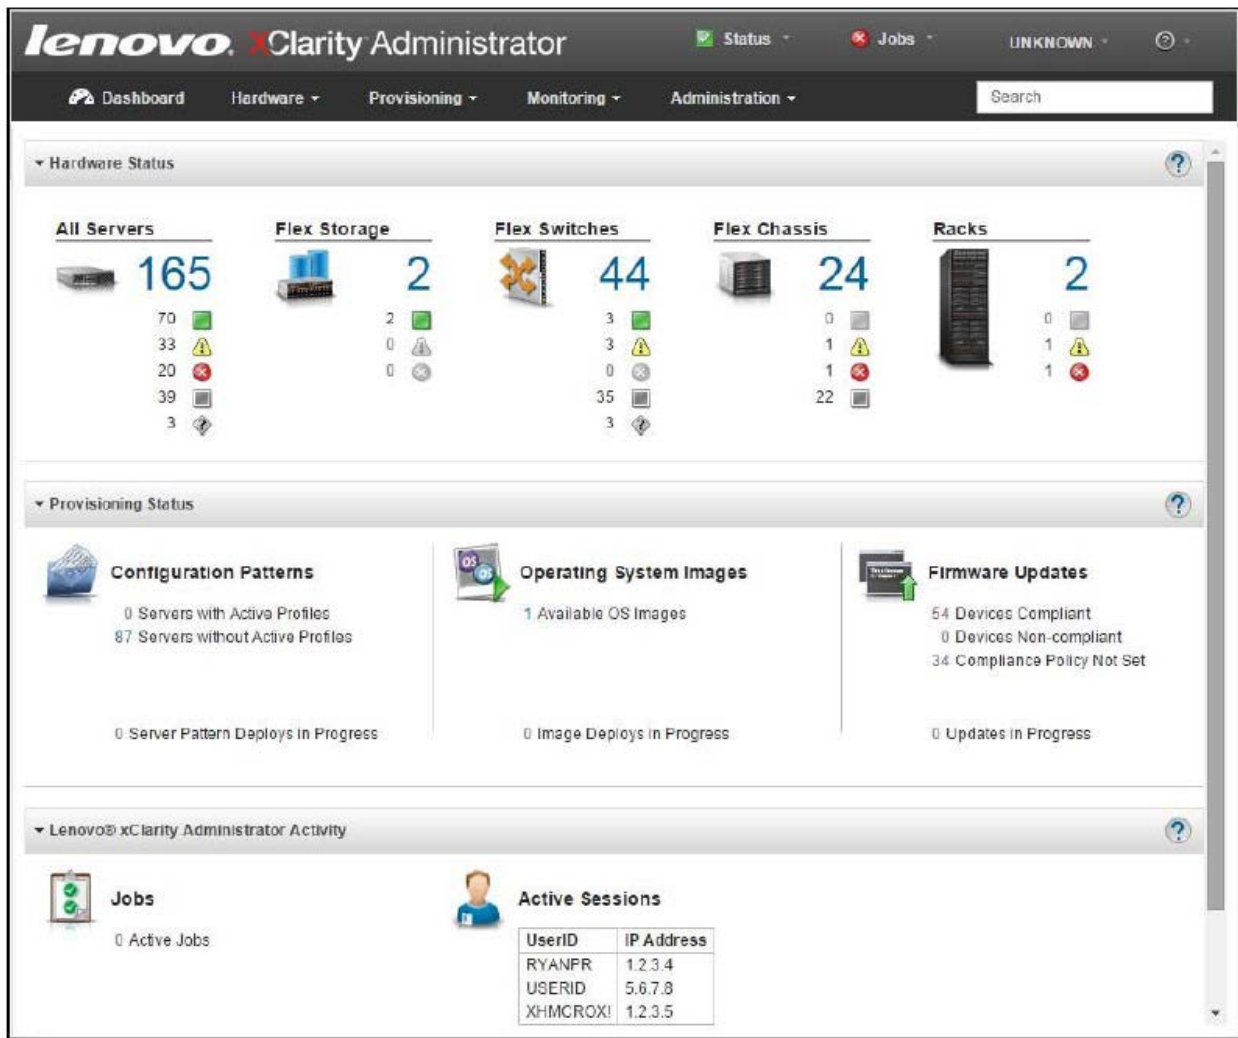


Figure 5: XClarity Administrator interface

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 26 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 26: 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 26 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 6: 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 7: Lenovo ThinkSystem DE Storage

Table 27 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 27: 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)	Predictive cache tier (read only) Persistent storage

	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 28 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 28: 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 29 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 29: 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 (such as the Lenovo ThinkSystem DS6200 storage array), 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 as described in the Lenovo Client Virtualization reference architecture, which is available at this website: lenovopress.com/lp0756.

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.

If only a single Flex System Enterprise Chassis is used, the chassis switches are sufficient and no other TOR switch is needed. For rack servers, more than one Flex System Enterprise Chassis TOR switches are required.

For more information, see “BOM for networking” on page 44.

4.10.110 GbE networking

For 10 GbE networking, the use of CIFS, NFS, or iSCSI, the Lenovo RackSwitch™ G8124E, and G8272 TOR switches are recommended because they support VLANs by using Virtual Fabric. Redundancy and automated failover is available by using link aggregation, such as Link Aggregation Control Protocol (LACP) and two of everything. For the Flex System chassis, pairs of the EN4093R switch should be used and connected to a G8124 or G8272 TOR switch. The TOR 10GbE switches are needed for multiple Flex chassis or external connectivity. iSCSI also requires converged network adapters (CNAs) that have the LOM extension. Table 30 lists the TOR 10 GbE network switches for each user size.

Table 30: TOR 10 GbE network switches needed

10 GbE TOR network switch	300 users	600 users	1200 users	3000 users
G8124E – 24-port switch	2	2	2	0
G8272 – 64-port switch	0	0	0	2

4.10.21 GbE administration networking

A 1 GbE network should be used to administer all of the other devices in the system. Separate 1 GbE switches are used for the IT administration network. Lenovo recommends that redundancy is also built into this network at the switch level. At minimum, a second switch should be available in case a switch goes down. Table 31 lists the number of 1 GbE switches for each user size.

Table 31: TOR 1 GbE network switches needed

1 GbE TOR network switch	300 users	600 users	1200 users	3000 users
G8052 – 48 port switch	2	2	2	2

Table 32 shows the number of 1 GbE connections that are needed for the administration network and switches for each type of device. The total number of connections is the sum of the device counts multiplied by the number of each device.

Table 32: 1 GbE connections needed

Device	Number of 1 GbE connections for administration
ThinkSystem rack server	1
Flex System Enterprise Chassis CMM	2
Flex System Enterprise Chassis switches	1 per switch (optional)
Lenovo ThinkSystem DS6200 storage array	4
TOR switches	1

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. For more information, see the “BOM” section on page 46.

4.12 Deployment models

This section contains the following example deployment models:

- vSAN hyper-converged using ThinkSystem servers

4.12.1 Deployment example 1: vSAN hyper-converged using ThinkSystem servers

This deployment example supports 1200 dedicated virtual desktops with VMware vSAN in hybrid mode (SSDs and HDDs). 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 ThinkSystem SR630 1U servers and two G8124E RackSwitch TOR switches. Each server has two Xeon SP2 6230 processors, 768 GB of memory, two network cards, two 800 GB SSDs, and six 1.2TB HDDs. Two disk groups of one SSD and three HDDs are used for vSAN.

Figure 8 shows the deployment configuration for this example.



Figure 8: Deployment configuration for 1200 dedicated users using vSAN on ThinkSystem 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 33, Table 34, Table 35, and Table 36 list the BOM for a server (a minimum of two servers per cluster is needed). See Table 52 on page 46 for the Lenovo Flex System chassis BOM.

Table 33: ThinkSystem SR630 (Xeon SP Gen 1)

Code	Description	Quantity
7X02CTO1WW	ThinkSystem SR630 - 3yr Warranty	1
AUW3	Lenovo ThinkSystem Mainstream MB - 1U	1
AUW0	ThinkSystem MS 1U 8x2.5" Chassis	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	24
AUWB	Lenovo ThinkSystem SR630 8x2.5" SATA/SAS Backplane	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
AUWQ	Lenovo ThinkSystem 1U LP+LP BF Riser BKT	1
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCB	ThinkSystem Toolless Slide Rail Kit with 1U CMA	1
AVW9	ThinkSystem 750W (230V) Titanium Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1

Table 34: ThinkSystem SR650 (Xeon SP Gen 1)

Code	Description	Quantity
7X06CTO1WW	ThinkSystem SR650 - 3yr Warranty	1
AUQB	Lenovo ThinkSystem Mainstream MB - 2U	1
AUVX	ThinkSystem MS 2U 2.5" Chassis, up to 16	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2

AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	24
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCH	ThinkSystem Toolless Slide Rail Kit with 2U CMA	1
AUS8	ThinkSystem SR550/SR590/SR650 EIA Latch w/ VGA Upgrade Kit	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1
Select optional GPU adapters (with both risers)		
AUR3	ThinkSystem SR550/SR590/SR650 x16/x8 PCIe FH Riser 1 Kit	1
AURC	ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit	1
B15U	ThinkSystem NVIDIA Tesla P40 24GB PCIe Passive GPU	2
B15V	ThinkSystem NVIDIA Tesla M10 32GB PCIe Passive GPU	2

Table 35: ThinkSystem SN550 (Xeon SP Gen 1)

Code	Description	Quantity
7X16CTO1WW	ThinkSystem SN550 - 3yr Warranty	1
5977	Select Storage devices - no configured RAID required	1
AXHY	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	24
AUXP	Lenovo ThinkSystem SN550 Server	1
ATBS	Flex System CN4052S 4-port 10Gb Virtual Fabric Adapter	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1

Table 36: ThinkSystem SD530 (Xeon SP Gen 1)

Code	Description	Quantity
7X21CTO1WW	SD530 : ThinkSystem SD530 - 3yr Warranty	1
AUXN	ThinkSystem SD530 Computing Node	1
AX6D	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUNE	ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM	12
AUYG	ThinkSystem SD530 3x2 SAS/SATA BP	1
B0SS	ThinkSystem 430-8i SAS/SATA 12Gb Dense HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AUUL	ThinkSystem M.2 CV1 32GB SATA 6Gbps Non-Hot Swap SSD	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1

Note that 64GB LRDIMMs are slower than 32GB RDIMMs and this should be taken into account when sizing a solution with the ThinkSystem SD530 that is limited to 12 DIMMs and may require the larger 64GB LRDIMMs for the same amount of system memory as the other servers that can support 24 DIMMs.

Table 37: ThinkSystem D2 Enclosure (for SD530)

Code	Description	Quantity
7X20CTO1WW	Chassis1 : ThinkSystem D2 Enclosure -3yr Warranty	1
AUXM	ThinkSystem D2 Enclosure	1
AUY9	ThinkSystem D2 10Gb 8 port E10M SFP+	1
AUYC	ThinkSystem D2 Slide Rail	1
AUYD	ThinkSystem D2 CMA (Cable Management Arm)	1
AUZ2	ThinkSystem D2 2000W Platinum PSU	2
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
AVR1	ThinkSystem Single Ethernet Port SMM	1
AUY7	ThinkSystem D2 8-slot x8 Shuttle ASM	1

5.2 BOM for hyper-converged compute servers

This section contains the bill of materials for hyper-converged compute servers. Table 38 and Table 39 list the BOM for a vSAN server (a minimum of four servers per cluster is needed).

Table 38: ThinkSystem SR630 (Xeon SP Gen 1)

Code	Description	Quantity
7X02CTO1WW	ThinkSystem SR630 - 3yr Warranty	1
AUW3	Lenovo ThinkSystem Mainstream MB - 1U	1
AUW1	ThinkSystem MS 1U 10x2.5" Chassis	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUW9	Lenovo ThinkSystem SR630 2.5" Anybay 10-Bay Backplane Kit	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
AUWQ	Lenovo ThinkSystem 1U LP+LP BF Riser BKT	1
AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCB	ThinkSystem Toolless Slide Rail Kit with 1U CMA	1
AVW9	ThinkSystem 750W (230V) Titanium Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
AUMG	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12Gb Hot Swap SSD	2
AUMH	ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	8
AUMK	ThinkSystem 2.5" PM1633a 3.84TB Capacity SAS 12Gb Hot Swap SSD	6

Table 39: ThinkSystem SR650 (Xeon SP Gen 1)

Code	Description	Quantity
7X06CTO1WW	ThinkSystem SR650 - 3yr Warranty	1
AUQB	Lenovo ThinkSystem Mainstream MB - 2U	1
AUVV	ThinkSystem MS 2U 24x2.5" Chassis	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	2
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	2
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCH	ThinkSystem Toolless Slide Rail Kit with 2U CMA	1

AUS8	ThinkSystem SR550/SR590/SR650 EIA Latch w/ VGA Upgrade Kit	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
AUMG	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12Gb Hot Swap SSD	2
AUMH	ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	8
AUMK	ThinkSystem 2.5" PM1633a 3.84TB Capacity SAS 12Gb Hot Swap SSD	6
Select optional GPU adapters (with both risers)		
AUR3	ThinkSystem SR550/SR590/SR650 x16/x8 PCIe FH Riser 1 Kit	1
AURC	ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit	1
B15U	ThinkSystem NVIDIA Tesla P40 24GB PCIe Passive GPU	2
B15V	ThinkSystem NVIDIA Tesla M10 32GB PCIe Passive GPU	2

Table 40: ThinkSystem SD530 (Xeon SP Gen 1)

Code	Description	Quantity
7X21CTO1WW	SD530 : ThinkSystem SD530 - 3yr Warranty	1
AUXN	ThinkSystem SD530 Computing Node	1
AX6D	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUNE	ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM	12
AUYG	ThinkSystem SD530 3x2 SAS/SATA BP	1
B0SS	ThinkSystem 430-8i SAS/SATA 12Gb Dense HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AUUL	ThinkSystem M.2 CV1 32GB SATA 6Gbps Non-Hot Swap SSD	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
AUMG	ThinkSystem 2.5" HUSMM32 400GB Performance SAS 12Gb Hot Swap SSD	2
AUMH	ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	4
AUMK	ThinkSystem 2.5" PM1633a 3.84TB Capacity SAS 12Gb Hot Swap SSD	4

Note that 64GB LRDIMMs are slower than 32GB RDIMMs and this should be taken into account when sizing a solution with the ThinkSystem SD530 that is limited to 12 DIMMs and may require the larger 64GB LRDIMMs for the same amount of system memory as the other servers that can support 24 DIMMs.

Table 41: ThinkSystem D2 Enclosure (for SD530)

Code	Description	Quantity
7X20CTO1WW	Chassis1 : ThinkSystem D2 Enclosure -3yr Warranty	1
AUXM	ThinkSystem D2 Enclosure	1
AUY9	ThinkSystem D2 10Gb 8 port E10M SFP+	1
AUYC	ThinkSystem D2 Slide Rail	1
AUYD	ThinkSystem D2 CMA (Cable Management Arm)	1

Code	Description	Quantity
AUZ2	ThinkSystem D2 2000W Platinum PSU	2
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
AVR1	ThinkSystem Single Ethernet Port SMM	1
AUY7	ThinkSystem D2 8-slot x8 Shuttle ASM	1

Table 42: ThinkAgile VX3320 Appliance (Xeon SP Gen 2)

Part number	Product Description	Qty
7Y93CTO1WW	Server : Lenovo ThinkAgile VX3320 Appliance	1
B1DK	ThinkAgile VX 1U 2.5" 10 Bay Chassis	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	1
AVWB	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
B88T	VMware ESXi 6.7 U3 (factory installed)	1
AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
AUWQ	Lenovo ThinkSystem 1U LP+LP BF Riser Bracket	1
AXCA	ThinkSystem Toolless Slide Rail	1
AUW9	ThinkSystem 1U 2.5" 4 AnyBay 10-Bay Backplane	1
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 128GB SATA 6Gbps Non-Hot Swap SSD	2
B7Y0	Enable IPMI-over-LAN	1
AVKG	ThinkSystem SR630 MB to 10x2.5" HDD BP NVME cable	1
AUWV	10x2.5"Cable Kit (1U)	1
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
B5MD	vSAN Hybrid Config	1
B4Y5	ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM2	ThinkSystem 2.5" 1.8TB 10K SAS 12Gb Hot Swap 512e HDD	6
B5MC	vSAN All Flash Config	1
B4Y7	ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD	6

Table 43: ThinkAgile VX 1U Certified Node (Xeon SP Gen 2)

Part number	Product Description	Qty
7Y93CTO4WW	Server : Lenovo ThinkAgile VX 1U Certified Node	1
B1DK	ThinkAgile VX 1U 2.5" 10 Bay Chassis	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUW9	ThinkSystem 1U 2.5" 4 AnyBay 10-Bay Backplane	1
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1

AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 128GB SATA 6Gbps Non-Hot Swap SSD	2
B88T	VMware ESXi 6.7 U3 (factory installed)	1
AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	1
AVWB	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
B5MD	vSAN Hybrid Config	1
B4Y5	ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM2	ThinkSystem 2.5" 1.8TB 10K SAS 12Gb Hot Swap 512e HDD	6
B5MC	vSAN All Flash Config	1
B4Y7	ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD	6

Table 44: ThinkAgile VX 2U Certified Node (Xeon SP Gen 2)

Part number	Product Description	Qty
7Y94CTO5WW	Server : ThinkAgile VX 2U Node	1
B1DH	ThinkAgile VX 2U 2.5" 8/16/24 Bay Chassis	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUR5	ThinkSystem 2U/Twr 2.5" AnyBay 8-Bay Backplane	1
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	2
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 128GB SATA 6Gbps Non-Hot Swap SSD	2
B88T	VMware ESXi 6.7 U3 (factory installed)	1
AURC	ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit	1
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
AUSM	MS 2nd 2U 8X2.5" Cable Kit	2
AUSQ	On Board to 2U 8x2.5" HDD BP NVME Cable	1
Select drive configuration for hyper-converged system (all flash or SDD/HDD combination)		
B5MD	vSAN Hybrid Config	1
B4Y5	ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	2
AUM2	ThinkSystem 2.5" 1.8TB 10K SAS 12Gb Hot Swap 512e HDD	6
B5MC	vSAN All Flash Config	1
B4Y7	ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD	6

Table 45: BOM for ThinkSystem SR635 and SR655 (AMD EPYC processor) based compute nodes

Part number	Product Description	Qty
7Y99CTO1WW	ThinkSystem-SR635 : ThinkSystem SR635 - 3yr Warranty	1

B5VL	ThinkSystem SR635 10x2.5" Chassis	1
B6TU	ThinkSystem AMD EPYC 7452 32C 155W 2.35 GHz Processor	1
B5XC	ThinkSystem 64GB TruDDR4 3200MHz (2Rx4 1.2V) RDIMM-A	8
B5VS	ThinkSystem 1U 10x2.5" PCIe Gen4 AnyBay Backplane	1
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	4
B5XH	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2
B88T	VMware ESXi 6.7 U3 (factory installed)	1
B5VV	ThinkSystem SR635/SR655 x16 PCIe Riser1	1
B5W0	ThinkSystem SR635/SR655 x16/x16 PCIe Riser2 (BF)	1
B5WZ	SR635/SR655 MS FH Riser BKT	1
B715	ThinkSystem SR635 MS LP Riser BKT Slot3	1
AUKP	ThinkSystem Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	1
B6XQ	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
AXCA	ThinkSystem Toolless Slide Rail	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

7Z01CTO1WW	ThinkSystem-SR655 : ThinkSystem SR655 - 3yr Warranty	1
B5VJ	ThinkSystem SR655 24x2.5" Chassis	1
B6TU	ThinkSystem AMD EPYC 7452 32C 155W 2.35 GHz Processor	1
B5XC	ThinkSystem 64GB TruDDR4 3200MHz (2Rx4 1.2V) RDIMM-A	8
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	3
5977	Select Storage devices - no configured RAID required	1
AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1
AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	4

B5XH	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2
B88T	VMware ESXi 6.7 U3 (factory installed)	1
B5VW	ThinkSystem SR635/SR655 x8 PCIe Internal Riser	1
B69R	ThinkSystem SR655 x16 PCIe 2U Riser 1	1
B5VY	ThinkSystem SR655 x16/x16 PCIe Riser2	1
AUKP	ThinkSystem Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	1
B6XQ	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	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

5.3 BOM for enterprise and SMB management servers

Table 25 on page 24 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 33.

Because the Windows storage servers use a bare-metal operating system (OS) installation, they require much less memory and can have a reduced configuration as listed below.

Table 46: ThinkSystem SR630 (Xeon SP Gen 1)

Code	Description	Quantity
7X02CTO1WW	ThinkSystem SR630 - 3yr Warranty	1
AUW3	Lenovo ThinkSystem Mainstream MB - 1U	1
AUW0	ThinkSystem MS 1U 8x2.5" Chassis	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUU1	ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	12
AUWB	Lenovo ThinkSystem SR630 8x2.5" SATA/SAS Backplane	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
AUWQ	Lenovo ThinkSystem 1U LP+LP BF Riser BKT	1
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCB	ThinkSystem Toolless Slide Rail Kit with 1U CMA	1
AVW9	ThinkSystem 750W (230V) Titanium Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1

Table 47: ThinkSystem SR650 (Xeon SP Gen 1)

Code	Description	Quantity
7X06CTO1WW	ThinkSystem SR650 - 3yr Warranty	1
AUQB	Lenovo ThinkSystem Mainstream MB - 2U	1
AUVX	ThinkSystem MS 2U 2.5" Chassis, up to 16	1
AWEN	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUU1	ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	12
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCH	ThinkSystem Toolless Slide Rail Kit with 2U CMA	1
AUS8	ThinkSystem SR550/SR590/SR650 EIA Latch w/ VGA Upgrade Kit	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2

6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1

Table 48: ThinkSystem SN550 (Xeon SP Gen 1)

Code	Description	Quantity
7X16CTO1WW	ThinkSystem SN550 - 3yr Warranty	1
5977	Select Storage devices - no configured RAID required	1
AXHY	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUU1	ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	12
AUXP	Lenovo ThinkSystem SN550 Server	1
ATBS	Flex System CN4052S 4-port 10Gb Virtual Fabric Adapter	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	1

Table 49: ThinkSystem SD530 (Xeon SP Gen 1)

Code	Description	Quantity
7X21CTO1WW	SD530 : ThinkSystem SD530 - 3yr Warranty	1
AUXN	ThinkSystem SD530 Computing Node	1
AX6D	Intel Xeon Gold 6130 16C 125W 2.1GHz Processor	2
AUU1	ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	12
AUYG	ThinkSystem SD530 3x2 SAS/SATA BP	1
B0SS	ThinkSystem 430-8i SAS/SATA 12Gb Dense HBA	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AUUL	ThinkSystem M.2 CV1 32GB SATA 6Gbps Non-Hot Swap SSD	2
5977	Select Storage devices - no configured RAID required	1
B88T	VMware ESXi 6.7 U3 (factory installed)	1

5.4 BOM for shared storage

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

Table 50: Lenovo ThinkSystem DE6000F control enclosure

Code	Product Description	Qty
7Y79CTO1W W	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 51: 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

5.5 BOM for networking

For more information about the number and type of TOR network switches that are needed for different user counts, see the “Networking” section on page 29.

RackSwitch G8052

Code	Description	Quantity
7159G52	Lenovo System Networking RackSwitch G8052 (Rear to Front)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
A3KP	Lenovo System Networking Adjustable 19" 4 Post Rail Kit	1

ThinkSystem NE1032

Code	Description	Quantity
7159A1X	Lenovo ThinkSystem NE1032 RackSwitch (Rear to Front)(Standard)	2
39Y7938	2.8m, 10A/100-250V, C13 to IEC 320-C20 Rack Power Cable	4

RackSwitch G8272

Code	Description	Quantity
7159CRW	Lenovo System Networking RackSwitch G8272 (Rear to Front)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
A3KP	Lenovo System Networking Adjustable 19" 4 Post Rail Kit	1
A1DP	1m QSFP+-to-QSFP+ cable	1
A1DM	3m QSFP+ DAC Break Out Cable	0

5.6 BOM for Flex System chassis

Table 52 lists the BOM for a Lenovo Flex System chassis.

Table 52: Chassis BOM

Code	Description	Quantity
8721HC2	Lenovo Flex System Enterprise Chassis w/ CMM2	1
ASUS	Lenovo Flex System Enterprise Chassis w/ CMM2	1
ASUV	Lenovo Flex System Fabric SI4093 System Interconnect Module	2
6292	2m, 16A/100-250V, C19 to IEC 320-C20 Rack Power Cable	6
A0UC	Flex System Enterprise Chassis 2500W Power Module Standard	2
A0UD	Flex System Enterprise Chassis 2500W Power Module	4
A0UA	Flex System Enterprise Chassis 80mm Fan Module	4

5.7 BOM for rack

Table 53 lists the BOM for a rack.

Table 53: Rack BOM

Code	Description	Quantity
93634PX	42U 1100mm Enterprise V2 Dynamic Rack	1
39Y8941	DPI Single Phase C19 Enterprise PDU (without line cord)	6
40K9614	DPI 30a Line Cord (NEMA L6-30P)	6

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 6.6 design and sizing guide
storagehub.vmware.com/#!/vmware-vsan/vmware-r-vsan-tm-design-and-sizing-guide
- Flex System Interoperability Guide
lenovopress.com/redpfsig-flex-system-interoperability-guide
- Lenovo Storage Interoperability Links
lenovopress.com/lp0584-lenovo-storage-interoperability-links
- View Architecture Planning – VMware Horizon 6.0
docs.vmware.com/en/VMware-Horizon-7/7.2/view-72-architecture-planning.pdf

Document History

Version 1.0	18 September 2017	<ul style="list-style-type: none"> Initial version
Version 1.1	8 January 2018	<ul style="list-style-type: none"> Added section on GPU graphics acceleration Replaced G8124E RackSwitch with ThinkSystem NE1032 switch Added ThinkSystem SD530 server
Version 1.2	22 June 2018	<ul style="list-style-type: none"> Added NVIDIA M10 and P40 GPUs
Version 1.3	5 November 2018	<ul style="list-style-type: none"> Added results for ESXi 6.7 using vSAN 6.7
Version 1.4	30 March 2020	<ul style="list-style-type: none"> Added Windows 10 version 1809 and Horizon 7.11 results for ThinkAgile VX Appliances with Xeon SP Gen 2 processors using ESXi 6.7 U3
Version 1.5	10 June 2020	<ul style="list-style-type: none"> Removed Windows 10 version 1803 results for Xeon SP Gen1 processors Added Windows 10 version 1909 and Horizon 7.11 results with Xeon SP Gen 2 processors for ThinkAgile VX Appliances, ThinkSystem SR850 and SR850P servers Added Windows 10 version 1909 and Horizon 7.11 results with Xeon SP Gen 1 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	<ul style="list-style-type: none"> Added section 4.3.3 Lenovo ThinkSystem with AMD EPYC (Rome) single and dual node VSI_{max} performance

Trademarks and special notices

© Copyright Lenovo 2020

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.