



Reference Architecture for SAP applications on ThinkAgile VX

Last update: 1 March 2024
Version: 2.0

Provides a technical overview of
ThinkAgile VX integrated
solutions and certified nodes

Contains SAP HANA best
practices and sizing
recommendations

Shows how SAP applications
can be used in a hyper-
converged environment

Configure your system for best
use with SAP HANA

Patrick Hartman



Table of Contents

1	Introduction	1
1.1	SAP SE solutions overview	1
2	Technical overview of ThinkAgile VX	2
2.1	ThinkAgile VX	2
2.2	Configurations	2
2.3	Software components	3
2.4	SAP application landscapes	4
2.5	Networking	6
2.6	Reliability and performance features	8
3	SAP Business Platform sizing	9
4	SAP HANA system sizing	10
4.1	ThinkAgile VX650 V3 Certified Nodes for SAP HANA	10
4.2	ThinkAgile VX 4U SAP HANA Nodes	11
4.3	SAP HANA sizing example	13
5	Virtualized SAP HANA Recommendations	21
5.1	SAP HANA architecture	21
5.2	ThinkAgile System Configuration	23
	Resources	25

1 Introduction

The intended audience for this document is technical IT architects, system administrators, and managers who are interested in executing SAP application workloads on the ThinkAgile VX Integrated Systems and Certified Nodes.

ThinkAgile VX provides a hyper-converged infrastructure. Hyper-converged means incorporating multiple components like compute and storage into a single entity through software. A hyper-converged infrastructure seamlessly pools compute and storage to deliver high performance for the virtual workloads and provides flexibility to combine the local storage using a distributed file system to eliminate shared storage such as SAN or NAS. These factors make the solution cost-effective without compromising performance.

The ThinkAgile VX Integrated Systems can be used for all SAP applications except for the SAP HANA®, along with ThinkAgile VX Certified Nodes. ThinkAgile VX Certified Nodes are available in 2U and 4U form factors and utilize the dramatic efficiency and cost savings of industry-leading VMware virtualization technology to handle the high demands of SAP applications such as the SAP HANA database.

Chapter 2 brings a technical overview of ThinkAgile VX and explains why the combination of Lenovo servers with VMware software gives the best system performance and reliability. Chapters 3 and 4 present sizing examples for SAP business applications and the SAP HANA database. Then, best practices using the SAP HANA database are given in chapter 5.

1.1 SAP SE solutions overview

SAP Societas Europaea (SE) is a software vendor of mission critical business software. Their comprehensive solutions for business processes across all industries offer the solutions needed to run a company.

SAP SE has over 1600 products and 330 solutions which may run on ThinkAgile VX. Those can be categorized, in general, as follows:

- Enterprise Resource Planning
 - SAP S/4HANA®
 - SAP Business Suite® (ERP, SCM, PLM, CRM, SRM)
- Digital Platform
 - SAP HANA database
 - SAP Sybase ASE® (and other databases, like Oracle®, IBM DB2®, or Microsoft SQL Server®)
 - SAP NetWeaver® platform (Portal, KM, BI, PI, ABAP, JAVA)
- Business Analytics
 - Business Intelligence with SAP BW/4HANA®

2 Technical overview of ThinkAgile VX

This chapter provides an overview of the ThinkAgile VX Integrated Systems (IS) and Certified Nodes (CN) including the associated software, systems management, and networking.

2.1 ThinkAgile VX

ThinkAgile VX IS and CN systems are designed to help you simplify the IT infrastructure, reduce costs, and accelerate time to value. These hyper-converged appliances from Lenovo combine industry-leading hyper-convergence software from VMware with Lenovo enterprise platforms. Several common uses are:

- Enterprise workloads, like SAP S/4HANA
- Server virtualization
- Virtual desktop infrastructure (VDI)

Starting with as few as three nodes to keep your acquisition costs down, the ThinkAgile VX systems are capable of immense scalability as your needs grow.

2.2 Configurations

ThinkAgile VX IS & CN systems are available in 2U and 4U models that can be tailored to your SAP application needs:

- ThinkAgile VX V3 1U IS and CN based on Intel Eagle Stream Platform – Intel 4.Gen Xeon Processor
- ThinkAgile VX V3 2U IS and CN based on Intel Eagle Stream Platform – Intel 4.Gen Xeon Processor
 - ThinkAgile VX650 V3 supports the new VMware vSphere Distributed Services Engine (DSE) using the NVIDIA Bluefield SmartNIC Data Processing Unit (DPU). This solution improves the performance of VMware vSphere by offloading tasks from server CPUs to the DPU thereby making data center applications more efficient
- ThinkAgile VX 4U IS and CN based on Intel Purley Platform – Intel 2.Gen Xeon Processor

For more information about the system specifications and supported configurations, refer to the product guides for the ThinkAgile VX IS & CN systems see the ThinkAgile VX Series [data sheet](#) (lenovopress.com/ds0104).

In addition, the ThinkAgile VX Solution for SAP HANA [data sheet](#) (lenovopress.com/ds0118) lists the Certified Nodes that have been certified for use with the SAP HANA Platform edition database. The SAP HANA data sheet also highlights limitations to the type and number of options which can be configured for use with the high performance in-memory database.

2.3 Software components

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

2.3.1 Hypervisor

The ThinkAgile VX IS and standard CN models come with VMware ESXi hypervisors preloaded in the factory, which is optionally preloaded in the factory for the ThinkAgile VX server per customer request.

The ThinkAgile VX CN for SAP HANA models, while similar, have special restrictions to the versions of VMware ESXi hypervisors allowed to be used to run SAP HANA dependent on the processor generation of the series chosen.

Table 1: Versions of VMware vSphere ESXi for SAP HANA usage

System Type	VMware vSphere ESXi version
ThinkAgile VX 4U for SAP HANA	VMware ESXi 7.0 U1, U2, U3
ThinkAgile VX650 V3 CN for SAP HANA	VMware ESXi 8.0 U1

2.3.2 Software Defined Storage

Using VMware vSAN, the ThinkAgile VX systems come with a world class software defined storage application. While standard systems support of VMware vSAN software is listed in the best recipe listings online, those for SAP HANA have specific versions which are validated by VMware exclusively at SAP SE and Lenovo provided below

Table 2: Versions of VMware vSAN for SAP HANA usage

System Type	VMware vSphere vSAN version
ThinkAgile VX 4U for SAP HANA	VMware ESXi 7.0 U1, U2, U3
ThinkAgile VX V3 CN for SAP HANA	VMware vSAN 8.0 U1

2.3.3 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized system management solution that helps administrators deliver infrastructure faster. This solution integrates easily with Lenovo servers, ThinkAgile VX IS and CN nodes providing automated agent-less discovery, monitoring, firmware updates, and configuration management.

Lenovo XClarity Pro goes one step further and provides entitlement to additional functions such as XClarity Integrators for Microsoft System Center and VMware vCenter, XClarity Administrator Configuration Patterns and Service and Support. Lenovo XClarity Administrator is an optional software component and can be also used to manage firmware upgrades.

Lenovo XClarity Administrator is provided as a virtual appliance that can be quickly imported into a virtualized environment. XClarity can either be installed on a separate server or a server within a VMware vSAN cluster, providing that the hardware management network with the server xClarity Console (XCC) is routable from the server hosting the XClarity virtual machine.

2.4 SAP application landscapes

Logically, an SAP application landscape has 3 tiers: a presentation layer, a business logic or application layer, and a database layer. Scalability is key for a successful implementation of an SAP solution. As a business grows, SAP provides scaling using 2 and 3-tier deployment architectures as follows:

- A 2-tier (also known as vertical scaling) architecture uses a single server for all SAP applications, except the presentation layer. This means that the presentation layer is one tier and the application and database layer is mapped to the second tier.
- A 3-tier (also known as horizontal scaling) architecture means that each of three logical tiers maps to an individual set of one or more physical servers.

A 2-tier architecture has the following key advantages:

- Because the application and database layers are on the same host, this architecture provides the best performance, since there is no overhead for database connections, nor network traffic, and no shadow processes.
- The 2-tier approach makes it easier to administer small installations or installations with performance demands that are not dynamic. Larger installations can also use the extra performance that can be achieved from this type of installation.

A 3-tier architecture offers the following key advantages:

- The application tier can be scaled separately from the database tier to provide both, load sharing and high availability in each tier. Short-time load peaks can be easily managed by temporarily adding more application servers. Larger SAP application implementations typically use a 3-tier or multi-tier architecture to allow easier implementation of high availability for each component or tier.
- Separating the database layer from the application layer gives an organization the flexibility to choose the database of choice, such as Microsoft SQL Server, Oracle, IBM DB2, SAP MaxDB, SAP Sybase ASE or SAP HANA and keep it separate from SAP application servers.

The flexibility of the SAP software enables one to decide the best approach whether that is a two, three or multi-tiered architecture.

An SAP application landscape is a group of two or more SAP systems. An SAP system consists of one or more SAP instances across one or more servers. A traditional SAP landscape may consist of one or more of the following types of SAP instances:

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

Figure 1 shows an example of a typical landscape for an SAP customer.




Figure 1: SAP customer landscape

Virtualization of SAP landscapes

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

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

Lenovo and their partners have conducted benchmarks with virtualization enabled to show the performance differences between a virtualized and non-virtualized SAP application system. According to the SAP Notes for VMware ESXi, the performance overhead has been measured to be 10% of the CPU and Memory resources. These overheads should always be included when sizing a new hardware platform. VMware vSAN has its own memory overhead that depends not only on the processing power, but also the storage configurations. This overhead calculation can be determined using a formula in the VMware KB [2113954](#) Understanding vSAN memory consumption in ESXi 7.0 as an example.

The storage of VMware vSAN monitors for any disk to be above an 80% threshold. When this value is reached, VMware vSAN starts moving the data onto other disks. If all disks are above 80%, then VMware vSAN will continuously try to swap data in and out of every drive, reducing thus the storage bandwidth to nothing. The recommendation from VMware is to stay safely aware from that condition. This “slack” space is a comfortable value set around 30% of the whole physical space of the cluster to be left aside. A cluster-wide slack space of 30% additional cluster space must be added with a factor of 1.42 multiplied to the total raw

space of the entire cluster. In general, assuming a replication factor value of 2 (or failures to tolerate, FTT, equal to 1), we create a formula for the storage overhead as such:

$$Raw\ Capacity = \sum_{nodes=1}^{\#Nodes} \left(\sum_{VMs=1} (HANA_{shared} + HANA_{data} + HANA_{log} + OS_{FS} + OS_{swap}) \right) \times 2_{FTT=1} \times 1.42_{slack}$$

2.5 Networking

There are three logical networks defined in this reference architecture:

- **External:** The external network is used for the internet access to the clusters, ingress to the exposed applications (services and routes).
- **Internal:** This is the non-routable network used for vMotion, and vSAN data networks.
- **Out-of-band network:** This is a secured and isolated network used for switch and server hardware management, such as access to the xClarity Controller (XCC) module on the servers and SoL (Serial-over-LAN).

Network redundancy

A deployment of SAP applications on the ThinkAgile VX vSAN platform uses 10GbE or higher network as the primary fabric for inter-node communication. This example uses two switches to provide redundant data layer communication and deliver maximum network availability. The typical deployment architecture for this setup is shown in Figure 2.

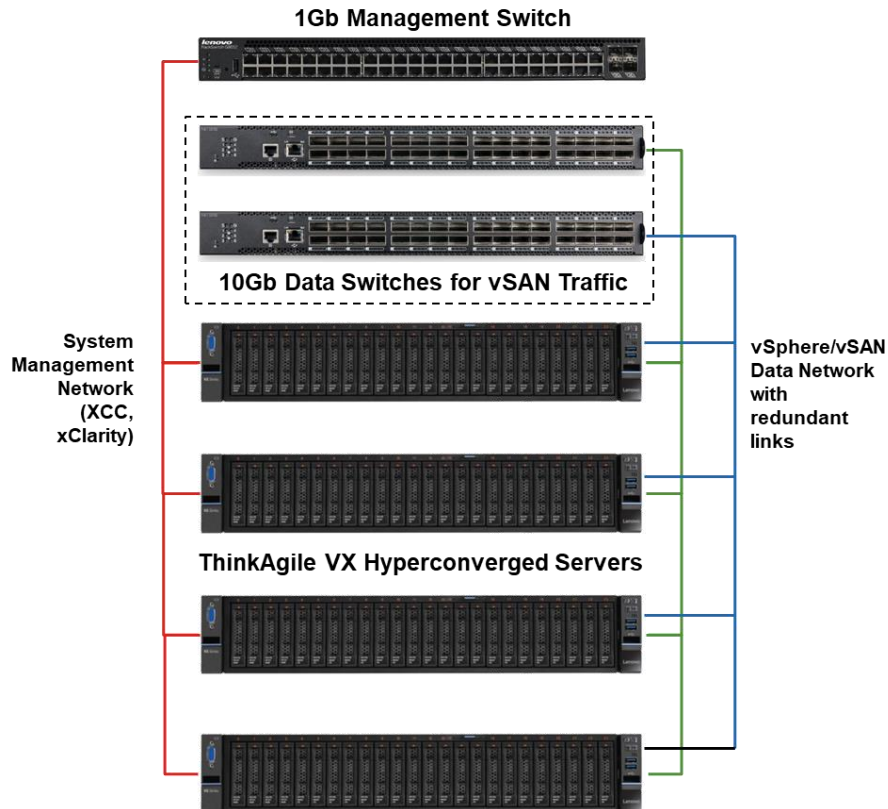


Figure 2: ThinkAgile VX network connectivity

The two primary network fabrics shown in the diagram are the systems management network and the internal data/user network. Typically, 1Gbps Ethernet is enough for the systems management network, which provides out-of-band access to the on-board management processors on the servers and network switches. The data/cluster internal fabric is recommended to be 10Gbps ethernet. This fabric is also recommended to have redundant switches for high availability of the network fabric.

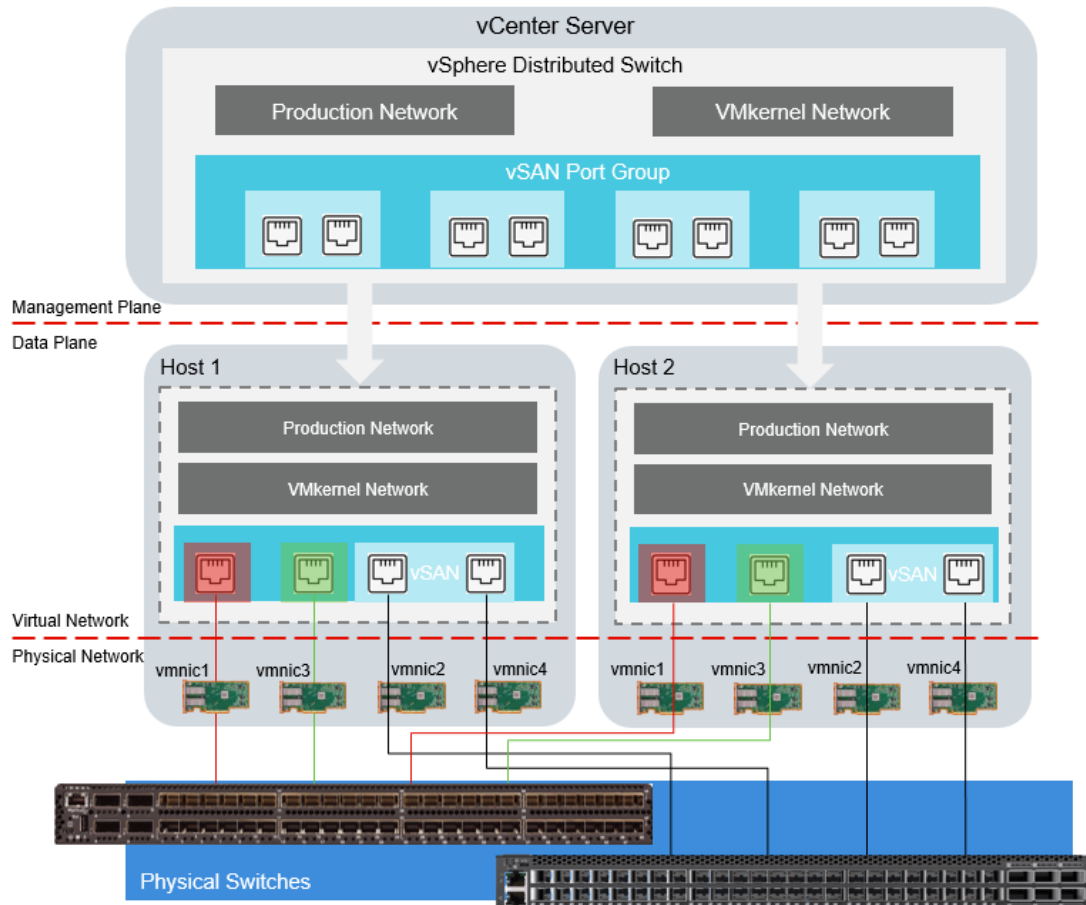


Figure 3: Networking

Figure 4 shows the redundant network architecture and the MCLAG configuration.

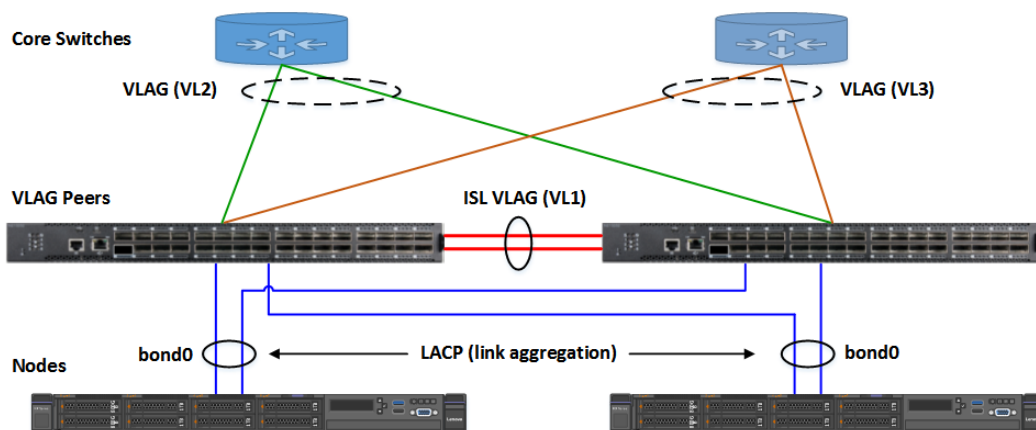


Figure 4: Network fabric redundancy and MCLAG

Multi-chassis link aggregation group (MCLAG) is a feature of the switch operating system that allows a pair of switches to work as a single virtual switch. Its implementation may vary by vendor; most notably that the protocol between each chassis is proprietary. Each of the cluster nodes has a link to each MCLAG peer switch for redundancy. This provides high availability (HA) for the nodes using the link aggregation control protocol (LACP) for aggregated bandwidth capacity. Connection to the uplink core network is facilitated by the MCLAG peers, which present a logical switch to the uplink network, enabling connectivity with all links active and without a hard requirement for spanning-tree protocol (STP). The link between the two MCLAG peers is an inter-switch link (ISL) and provides excellent support of east-west cluster traffic between the nodes. The MCLAG presents a flexible basis for interconnecting to the uplink/core network, ensures the active usage of all available links, and provides high availability in case of a switch failure or a required maintenance outage needs to be connected to a 1GbE TOR switch.




Figure 5: System Management Network

2.6 Reliability and performance features

Reliability and excellent performance are important for any workload but particularly for hyper-converged infrastructures like ThinkAgile VX. These requirements are met through the following design features of VMware vSAN software combined with Lenovo Servers.

Hardware reliability

A ThinkAgile VX IS or CN has redundant hardware components by including two power supplies, multiple chassis fans, two Intel CPUs, multiple memory DIMMs, multiple SSDs and HDDs, and optionally up to two dual-port network interface cards. Lenovo uses the highest quality hardware components combined with firmware that is thoroughly tested. Consequently, Lenovo servers have been rated #1 in hardware reliability for the last 3 years.

Hardware performance

The ThinkAgile VX IS or CN have been carefully designed for performance. In addition to all the usual attributes like processors and memory; for example, the 32 drive ThinkAgile VX650 V3 CN uses three HBA controllers instead of only one. Therefore, the latency is halved for some workloads that heavily utilize the cold tier. This allows higher throughput and improved transaction rates.

3 SAP Business Platform sizing

Classical applications based on SAP NetWeaver Technology like the SAP Business Suite may run on any of the ThinkAgile VX IS and CN systems based on the SAP NetWeaver validation produced by VMware. The Lenovo and SAP sizing methodology is based on SAP benchmarks, information from SAP SE, and actual customer experiences. Lenovo uses sizing tools and customer input to approximate the system resource requirements; however, actual customer results can always vary. A sizing questionnaire is used to gather the requirements that are used to estimate the hardware resources to run SAP Business Suite on Lenovo servers.

This sizing estimate approximates the hardware resources (processor, memory, and disk) that are required to support an SAP application or component implementation. This estimate provides an introductory understanding of the customer's initial hardware requirements based on the information given. Your actual experiences vary from the sizing estimate for many reasons, including batch and reporting workloads and custom code.

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

VMware has created SAP community documents such as the *SAP on VMware vSphere* [FAQ](#), [SAP HANA on VMware vSphere](#) and on [SAP HANA on VMware vSAN](#) which describe the best practices and configuration guidance when deploying the VMware vSphere hypervisor for SAP applications. In addition, the Resources section at the end has more links and SAP Notes which should be helpful:

4 SAP HANA system sizing

VMware has created an SAP community document entitled “SAP HANA on VMware vSphere” [FAQ](#), which describes the best practices and configuration guidance when deploying the VMware vSphere hypervisor for SAP applications. In addition, the following SAP notes should be helpful:

- SAP Note [2937606](#): SAP HANA on VMware vSphere 7.0 in production
- SAP note [3406060](#): SAP HANA on VMware vSphere 8 and vSAN 8
- SAP Note [2718982](#): SAP HANA on VMware vSphere and vSAN

Lenovo provides two ThinkAgile products for use with SAP HANA. These are the ThinkAgile VX 4U certified node for SAP HANA based on Intel 2nd Generation Xeon processors and ThinkAgile VX V3 CN for SAP HANA based on Intel 4th Generation Xeon Processors. These servers have been certified by Lenovo and SAP for use of SAP HANA with VMware vSAN software defined storage.

4.1 ThinkAgile VX650 V3 Certified Nodes for SAP HANA

4.1.1 Supported SAP HANA Configurations using 2 Sockets:

Model Name	ThinkAgile VX650 V3 CN	
Machine Type/Model	7D6WCTO3WW	
Processor	1x or 2x 4th Gen Intel® Xeon® Scalable processors	
Memory	Up to 8TiB TruDDR5 Memory	
Virtual	S4/HANA (OLTP)	BW/4HANA (OLAP)
Memory Sizes	128GiB – 6 TiB vMemory / VM	128GiB – 2 TiB vMemory / VM
Drive Bays	32*x 2.5" (HS) NVMe/SAS/SATA	
Disk Groups per node* (min. 2)	2 / 4	
vSAN Nodes / Cluster	3-64 nodes	
OSA Disks allowed	Cache: Any NVMe/SAS SSD (max 4)	
	Capacity: Any NVMe/SAS/SATA (max 28)	
ESA Disks allowed	Any NVMe SSD (max 24)	
PCIe Networking	Minimum 10GbE networking required	
	Recommended 4 x 25GbE ports or higher	

*Each Productive SAP HANA VM requires a Disk Group using vSAN Original Storage Architecture.

4.2 ThinkAgile VX 4U SAP HANA Nodes

4.2.1 Supported SAP HANA Configurations using 4 Sockets:

Model Name	ThinkAgile VX 4U Node for SAP HANA	
Machine Type/Model	7Z12CTO2WW	
Processor	2x, 4x 16 Core second-generation Intel Xeon Scalable Gold or Platinum Processors or higher	
Memory	Up to 6TiB TruDDR4 Memory	
Virtual	S4/HANA (OLTP)	BW/4HANA (OLAP)
Memory Sizes	128GiB – 6 TiB vMemory / VM	128GiB – 3 TiB vMemory / VM
Persistent Memory	Up to 6TiB Intel Optane Persistent Memory	
Drive Bays	Up to 24x 2.5-inch Disks Total	
Disk Groups per node* (min. 2)	2 / 4	
vSAN Nodes / Cluster	3-64 nodes	
Disks allowed	Cache: Any NVMe/SAS (max 4)	
	Capacity: Any NVMe/SAS/SATA (max 20)	
Network	Minimum 10GbE networking required	
	Recommended 4 × 25GbE ports or higher	

*Each Productive SAP HANA VM requires a Disk Group.

4.2.2 Supported SAP HANA Configurations using 8 Sockets:

Model Name		ThinkAgile VX 4U Node for SAP HANA	
Machine Type/Model		7Z12CTO2WW	
Processor	8x 16 Core second-generation Intel Xeon Platinum processors	8x 24 Core second-generation Intel Xeon Platinum processors	8x 28 Core second-generation Intel Xeon Platinum processors
Memory	Up to 12TiB TruDDR4 Memory		
Drive Bays	Up to 18x 2.5” Disks Total	Up to 14x 2.5” Disks Total	Up to 12x 2.5” Disks Total
Virtual	S4/HANA (OLTP)		BW/4HANA (OLAP)
Memory Sizes	128GiB – 12 TiB vMemory / VM		128GiB – 3 TiB vMemory / VM
Persistent Memory	Up to 12TiB Intel Optane Persistent Memory		
Disk Groups per node* (min. 2)	2 / 4	2	2
vSAN Nodes / Cluster	3-64 nodes		
Disks allowed	Cache: Any NVMe/SAS SSD (max 4)	Cache: Any NVMe/SAS SSD (max 2)	Cache: Any NVMe/SAS SSD (max 2)
	Capacity: Any NVME/SAS SSD (max 14)	Capacity: Any NVME/SAS SSD (max 12)	Capacity: Any NVME/SAS SSD (max 10)
Network	Minimum 4 × 10GbE ports required		
	Recommended 4 × 25GbE ports or higher		

*Each Productive SAP HANA VM requires a Disk Group.

The **ThinkAgile VX CN for SAP HANA** have special hardware configuration rules that must be considered when configuring a server for use with SAP HANA. These rules may have been described elsewhere in this guide but are collected here for completeness' sake.

Memory Rules:

- At least one (1) fully populated memory channel with:
 - (VX V3) eight (8) similar TruDDR5 DIMMs
 - (VX 4U) twelve (12) similar TruDDR4 DIMMs
- Memory must be added in groups of either eight (8) or twelve (12) DIMMS from Table 4: ThinkAgile VX 4U DRAM Memory Configurations:
 - 1 DIMM Per Channel (DPC) × any feature code

- 1 DIMM Per Channel (DPC) x any feature code of **N** GB & 1 DIMM Per Channel (DPC) x any feature code of **N*2** GB
- 2 DIMMs Per Channel (DPC) x any feature code
- 1 DIMM Per Channel (DPC) of 12 DIMMS x any feature code and 1 DIMM Per Channel (DPC) of 12 Intel Optane PMem for the ThinkAgile VX 4U models
- Intel Optane Persistent Memory must fully populate one memory channel across all sockets, and fully populate the second memory channel with TruDDR4 DIMMS as shown in Table 5: ThinkAgile VX 4U Persistent Memory in App-Direct Mode Configurations

Networking Rules:

- A minimum of two (2) 10/25GbE Adapters are required for vSAN SDS traffic.
- A minimum of two (2) 10/25GbE Adapters are required for the ESXi hypervisor traffic.
- Other adapters may be added up to maximum supported by the equivalent base model

GPU Rules:

- VX nodes for SAP HANA do not support GPUs.

Software Rules:

- For VX nodes based on Intel Xeon SP 2. Generation processors
 - VMware ESXi 7.0 U1, U2, and U3 may be used with SAP HANA applications
- For VX nodes based on Intel Xeon SP 4. Generation processors
 - VMware ESXi 8.0 U1 may be used with SAP HANA applications
- All system management software FC and Option part number are allowed
- An additional Professional Service Unit exists for installation and deployment of SAP HANA

4.3 SAP HANA sizing example

In order to size a ThinkAgile VX system correctly with SAP HANA, one needs to consider various system requirements. We configure a system here to illustrate the concepts using an example customer. The customer wants to begin her journey with S/4HANA. She already uses VMware ESXi for her company's virtual infrastructure. Her company already uses SAP ERP on SAP HANA on a bare metal server from Lenovo. She now wants to move this physical server into a VMware vSAN hyper-converged infrastructure.

The customer's current SAP HANA system uses 128GB of main memory and runs on a Lenovo ThinkSystem SR950 server based on Intel Xeon SP family of processors. Using the sizing reports, we determine that she requires a workload of 25,000 *SAPS* and a maximum of 150GB of memory to handle the current system. Her team expects a growth of 50% over the next 3 years for a total of 225GB of main memory.

Immediately, it can be determined that she will require a new cluster of at least four (4) ThinkAgile VX 2U certified nodes. SAP HANA will run on two (2) of the nodes as a production VM and one for SAP HANA system replication. The fourth node will be used for her SAP ERP application servers. The following guidelines describe only the sizing of the productive SAP HANA database.

4.3.1 SAP based expert sizing

The size of an SAP HANA VM is the result of estimating the memory, processor, storage, and network components for each instance. Additionally, any non-productive SAP HANA workload and other SAP applications need to be calculated into the total to find the appropriate ThinkAgile VX solution size.

If possible, the easiest approach to determine the amount of processing power required is to use the *SAPS* value. *SAPS*, standing for SAP Application Performance Standard, describes a workload requirement for the CPU capacity of an SAP application and it can be calculated using the outputs of tools provided by SAP. They enable Lenovo experts to translate the sizing results into a Lenovo configuration with the appropriate Intel Xeon processors and memory. If not, one can use various CPU and memory utilization monitoring to determine the processing power required. For each additional virtual machine (VM), you will need to add the results together to form a complete picture of the resources necessary to run SAP applications.

Memory sizing

We decided at the very beginning that the estimated database size was using 225GB of virtual memory. This will be the memory required used in a single virtual machine instance running an OS, not the entire physical memory on one server. We want to use as many sockets as possible for a single instance. There are resources required for the VMware ESXi hypervisor and they must be also considered.

With 16GB RDIMMs, the ThinkAgile VX can be configured with a total of 384GB system memory (24 total DIMMS). The team determines that SAP HANA will run on two (2) sockets. See the following section on Processor (CPU) Sizing for more details.

VMware states in their VMware KB Article [2113954](#) – Understanding vSAN memory consumption – that vSAN will need system resources assigned to itself and should be calculated into any overall memory considerations per node. For the above example using a four disk groups per host, our all flash configuration uses roughly 42960 MB. VMware describes it, in short, as so:

```
vSANFootprint = HOST_FOOTPRINT + NumDiskGroups * DiskGroupFootprint
DiskGroupFootprint = DISKGROUP_FIXED_FOOTPRINT + DISKGROUP_SCALABLE_FOOTPRINT + CacheSize *
                    CACHE_DISK_FOOTPRINT + NumCapacityDisks * CAPACITY_DISK_FOOTPRINT
```


The ThinkAgile VX V3 CN use TruDDR5 memory with a possible 8 DIMMs per memory channel and the following DRAM memory configurations:

Table 3: ThinkAgile VX V3 DRAM Memory Configurations

Memory DIMM Type	Capacity (GB) per CPU	VX V3 2U
		2
16GB DIMM	128	256
16GB DIMM	256	512
32GB DIMM	256	512
16/32GB DIMM mix	384	768
32GB DIMM	512	1024
64GB DIMM	512	1024
32/64GB DIMM mix	768	1152
64GB DIMM	1024	2048
128GB DIMM	1024	2048
128GB DIMM	2048	4096
256GB DIMM	2048	4096

For the ThinkAgile VX 4U based certified nodes use TruDDR4 memory with a possible 6 DIMMs per memory channel and the following DRAM memory configurations.

Table 4: ThinkAgile VX 4U DRAM Memory Configurations

Memory DIMM Type	Intel Xeon 2. SP Type	Capacity (GB) per CPU	Capacity (GB) with #CPUs		
			VX 2U	VX 4U 4S	VX 4U 8S
			2	4	8
16GB DIMM	base	96	192	384	768
16GB DIMM	base	192	384	768	1536
32GB DIMM	base	192	384	768	1536
16/32GB DIMM mix	base	288	576	1152	2304
32GB DIMM	base	384	768	1536	3072
64GB DIMM	base	384	768	1536	3072
32/64GB DIMM mix	base	576	1152	2304	4608
64GB DIMM	base	768	1536	3072	6144
128GB DIMM	base	768	1536	3072	6144
64/128GB DIMM mix	M	1152	2304	4608	9216
128GB DIMM	M	1536	3072	6144	12288
256GB DIMM	M	1536	3072	6144	12288
128/256GB DIMM mix	L	2304	4608	9216	-
256GB DIMM	L	3072	6144	12288	-

VMware vSAN does not support Intel Optane PMem in App-Direct mode directly, it can use virtual machines which are configured in VMware ESXi to use Intel Optane PMem in App-Direct mode with some limitations. For SAP HANA, there are additionally specific restrictions and ratios allowed between the amount of traditional DRAM and Intel Optane PMem (shown in Table 5: ThinkAgile VX 4U Persistent Memory in App-Direct Mode Configurations). An SAP hardware sizing is always required before considering using Intel Optane PMem with SAP HANA. Limitations and restrictions to the usage of SAP HANA and Intel Optane PMem with VMware vSAN, can be found in the SAP Community Network (SCN) [Wiki](#).

Table 5: ThinkAgile VX 4U Persistent Memory in App-Direct Mode Configurations

Intel Optane PMem + TruDDR4 DRAM	Intel Xeon 2. SP Type	Usable Capacity (GB) per CPU	Usable Capacity (GB)			SAP HANA ratio
			VX 2U	VX 4U 4S	VX 4U 8S	
			2	4	8	
128GB PMem + 32GB DRAM	base	960	1920	3840	-	1:4
128GB PMem + 64GB DRAM	M	1152	2304	4608	-	1:2
128GB PMem + 128GB DRAM	M	1536	3072	6144	-	1:1
256GB PMem + 64GB DRAM	M	1920	3840	7680	-	1:4
128GB PMem + 256GB DRAM	L	2304	4608	9216	-	2:1
256GB PMem + 128GB DRAM	L	2304	4608	9216	-	1:2
256GB PMem + 256GB DRAM	L	3072	6144	12288	-	1:1
512GB PMem + 128GB DRAM	L	3840	7680	15360	-	1:4
512GB PMem + 256GB DRAM	L	4608	9216	18432	-	1:2

Processor (CPU) sizing

While no other workload may run on the socket (or half-socket) running a productive SAP HANA instance according to the SAP documentation, unused sockets may be used for any other type of workload including non-productive instances of SAP HANA. This means that when using both sockets for a single SAP HANA VM, there will be no remaining sockets to be used for other purposes like SAP application servers, and other non-productive SAP HANA instances. Instead, you can add other nodes into the cluster if you so desire for those applications or create a second cluster for explicit use of those applications.

In our example, the local team uses long-term load history to determine that she can use at least 10 cores per socket in total, so we can configure with two (2) Intel Xeon Gold 5218 16C processors.

Using guidelines from SAP, we can begin to define the processor type and sizes required. SAP recommends a minimum of eight (8) cores per socket. Since the team was previously using the Intel Xeon Gold 5218 16C processor in her original ThinkSystem SR650 VMware based machines, she has chosen to stay with the same processor in the ThinkAgile VX 2U node.

With the initial memory sizing, and applying the SAP HANA definitions for configuring an instance, this defines a virtual machine that could run on two (2) physical sockets with an estimated total of 352GB of physical memory. The remaining memory is reserved for the VMware ESXi hypervisor.

Storage sizing

SAP has created several configuration guides for the use of SAP HANA. A Tailored Datacenter Integrated approach (or TDI) of SAP HANA uses specific rules for the storage, network, and CPU to Memory ratios to ensure that enough raw performance is available for the SAP HANA database. Estimating the size of virtual SAN necessary to use with SAP HANA requires [guidance](#) from SAP. For each virtual machine, the individual needs for SAP HANA should be calculated accordingly, then these can be fitted into the VMware vSAN capacity layer appropriately. This document does a very good job describing the various shares necessary for SAP HANA persistence, backup, and operating system requirements for a system. We additionally recommend increasing the operating system to at least 128GB and have a swap space of a maximum of 32GB. While a bit larger than the recommendations from SAP SE, they allow temporary file growth and certain root file system features, like snapshots, when used with various Linux file system types.

VMware vSAN now has two storage architectures. VMware vSAN Extended Storage Architecture (ESA) and Original Storage Architecture (OSA). VMware vSAN OSA allows the installed disk drives to be divided up into cache and capacity drives called Disk Groups (DG), while the newer VMware vSAN ESA uses all NVMe SSD devices removing the need for a caching layer. Depending on the number of VMs per node and the type of capacity drive chosen, several options are available to the customer to configure the cache to capacity drive quantity ratio. We recommend the settings shown in Table 6: Original Storage Architecture Disk Group settings.

Table 6: Original Storage Architecture Disk Group settings

Original Storage Configuration	<p>For All Flash storage configurations, the ThinkAgile VX Certified Nodes supports from 1 to 5 cache drives and from 2 to 21 capacity drives depending on the quantity of the drive groups.</p> <p>For SAP HANA models, the following number of drive groups are limited due to performance of running SAP HANA to:</p> <ul style="list-style-type: none">• 2-4 drive groups: 1 cache drive and 2 to 7 capacity drives in each drive group. <p>The quantity of the capacity drives in each drive group must be the same. All cache drives in the node must be of the same model and capacity. All capacity drives in the node must be of the same model and capacity. All nodes in a cluster must have the same storage configuration.</p> <p>For optimal storage performance, VMware and Lenovo recommends using 1DG per virtual machine running SAP HANA within the cluster of nodes.</p> <p>For SAP HANA Configurations, drive groups of 1 or 5 are not recommended to be used.</p>
Disks allowed	<p>Cache: Any NVMe SSD or SAS SSD</p> <p>Capacity: Any SATA/SAS SSD or NVMe SSD</p>

Table 7: Extended Storage Architecture Disk Group Settings

Extended Storage Configuration	For Extended Storage Architecture configurations, the ThinkAgile VX Certified Nodes supports from 2 to 24 storage drives. The concept of capacity drives is no longer applicable and there are no cache drives or drive groups any longer.
Disks allowed	Any NVMe SSD



Figure 6: ThinkAgile VX650 V3 node with 2 ESA disk group layout using NVMe SSDs




Figure 7: ThinkAgile VX 4U node with 4 OSA disk group layout using NVMe/SATA SSDs

Using the above requirements, we calculate for each SAP HANA VM the sum of the following storage requirements:

- Storage for SAP HANA Shared, Data and Log directories.
 - SAP HANA TDI guidelines require roughly 2.5 * amount of VM memory
 - Example: 3 VMs each with 225GB = 2.5 * 225GB = 562.5GB

- Storage for each VMs Linux OS = 128GB¹
- Storage for each VMs Linux swap = 32GB²

We add all individual SAP HANA VM requirements together for one server and then multiply this by 2 because VMware states to set the failures to tolerate to one (1) which implies the doubling of the storage.

Finally, we add the storage per node required together for the total raw storage required. Then an additional slack space is calculated and added to the overall computation of the needed data space and FTT multiplier. Using the recommendation from VMware of 30% slack space and an FTT=1, we calculate the disk overhead as according to the formula in follows:

$$\text{Raw Capacity} = ((3 * (562.5 + 128 + 32) * 2)) * 4 * 1.42 = (4335 * 4) * 1.42 \cong \text{25TB total cluster capacity}$$

As an example, a ThinkAgile VX 4U node which has 4 disk groups and each drive group having at least four (4) SAS SSD devices can be set up as follows. For a cluster with four (4) nodes and using a capacity drive size of 800GB, this provides 51.2 TiB raw capacity (16×800×4) in total or more than enough for the use of the three (3) SAP HANA virtual machines in this example. Each drive group should also use one (1) ThinkSystem 1.6TB NVMe SSD to complete the VMware vSAN configuration for the cache device.

Network sizing

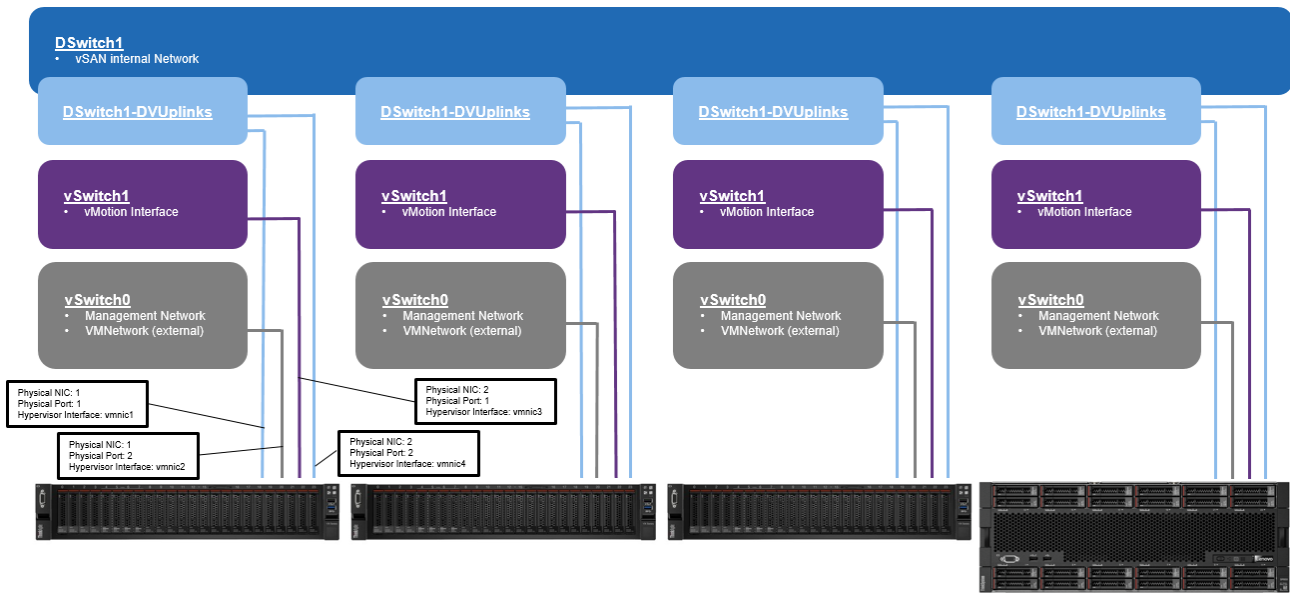
Running SAP HANA on ESXi requires network cards and switches that support Remote Direct Memory Access (RDMA) technology when using NVMe devices. Any switch described in Section 2.3 may be used for SAP HANA as they all support the RDMA over Converged Ethernet (RoCE) standard, but some additional configuration on the switch will be required (e.g. PFC parameters).

ThinkAgile VX for SAP HANA certified nodes require at least two (2) Mellanox ConnectX-4 Lx 2x25GbE SFP28, or equivalent adapters that are running in Ethernet mode to achieve a redundant connection. These adapters should always be added and connected to a 10GbE or faster enabled switch. Networking guidelines for SAP HANA are also in the SAP HANA [Network Guidelines](#). Setting up your switch should be done according to the switch vendors best practices guides.

The VLAN policies must be created in VMware vCenter and pushed over to the switch CLI according to the switch vendor's instructions.

¹ SAP [recommends](#) only 50GB for the Operating System, we do not recommend such a low value in hyperconverged scenarios.

² SAP [recommends](#) only 2GB for Linux Swap, we do not recommend such a low value in hyperconverged scenarios.



When using virtual machines larger than 224 vCPUs and 6TiB of virtual memory, consider the use of pass-through network configurations instead of the VMware virtual VMXNET3 driver. If undesired network latencies between the application servers and the SAP HANA database appear, please refer to SAP Note [3102813](#) for further options.

5 Virtualized SAP HANA Recommendations

This section describes at a high level the installation and configuration of supported Linux operating systems running SAP HANA. For more details on the configuration and operation of VMware vSAN as well as file layouts, please refer to the *Best Practices Guide for SAP HANA solutions using ThinkAgile VX* on Lenovo.com.

5.1 SAP HANA architecture

In order to deploy an instance of SAP HANA, you need to size and configure VMware ESXi hypervisor, guest, operating system, file system, and SAP HANA itself correctly. The following are the best practices and sizing guidance to create the most efficient use of the ThinkAgile VX system chosen.

5.1.1 ThinkAgile VX system specification

For installations of SAP HANA, certain minimum and maximum values need to be considered before starting to size a VM and the entire landscape. Table 9 is derived directly from the SAP Notes [2718982](#), [2393917](#) and the Lenovo certification.

Table 8: ThinkAgile VX 2.Gen SAP HANA limitations on VMware ESXi for vSAN

Description	Minimum	Maximum
Virtual Memory (vMEM) allowed	128GiB	4TiB
Virtual CPUs (vCPUs) allowed	8	up to 240
Sockets / Productive SAP HANA use	1	2 / 4 (using SNC ³)
Sockets / Non-SAP HANA Productive	1	2 / 4 (using SNC)
Number of ThinkAgile VX Nodes	3	64

Table 9: ThinkAgile VX 4U SAP HANA limitations on VMware ESXi for vSAN

Description	Minimum	Maximum
Virtual Memory (vMEM) allowed	128GiB	12TiB
Virtual CPUs (vCPUs) allowed	8	up to 448
Sockets / Productive SAP HANA use	1	2, 4, 8
Sockets / Non-SAP HANA Productive	.5	2 / 4, 4 / 8, 8 / 16
Number of ThinkAgile VX Nodes	3	64

³ Intel Sub-NUMA Clustering

5.1.2 VMware vSAN Cluster Best Practices

To best handle the nature of software defined infrastructure's inherent need for a cluster of servers, The VMware vSAN servers require a minimum of three (3) servers to run a vSAN cluster, while for SAP HANA, we highly recommend at least four (4) servers to ensure that VMware can handle not only a system fail-over, but also allows for maintenance of the cluster while maintaining resilience to any high availability outages.

If the customer wishes to use workloads other than productive SAP HANA DBs, extra ThinkAgile VX Integrated Systems or Certified Nodes may be used, but VMware vSphere and vSAN must be configured so that no SAP HANA workload is run on these extra nodes. All nodes within the cluster where SAP HANA would be run, must conform to the storage restrictions of SAP HANA TDI Storage Guidelines:

- For production SAP HANA VMs
 - Do not place any other VM on socket(s) running the productive SAP HANA VM
 - Do not share resources across any VMs and the productive VM(s)
 - Allow only whole multiples of processor sockets (1,2,3 or 4) per productive VM
 - Using Intel Sub-NUMA Clustering (starting with Intel Xeon 4th Generation SP)
- For non-production SAP HANA VMs
 - The same configurations for productive VMs are valid
 - Half socket non-productive VMs are allowed (but may not run on the same socket as a productive SAP HANA VM)
- For all other VM workloads
 - Any socket not running an instance of a productive SAP HANA VM can run productive or non-productive workloads, such as SAP ERP, BW or NetWeaver.
- Turn off any VMware vSAN features that save storage impacts, such as compression, deduplication or erasure coding (EC-X), on any containers holding the SAP HANA productive data files

Intel Sub-NUMA Clustering (SNC)

According to Intel, they have a new feature starting with newer Intel Xeon Scalable Processors which provides localization benefits similar to Cluster-On-Die (COD), a feature found in previous processor families, without some of COD's downsides. This feature is called Sub-NUMA Clustering or SNC and splits the last level cache (LLC) into disjoint clusters based on address range, while binding each cluster to a subset of the memory attached to it via the controllers in the system. SNC improves average latency and can replace the COD feature found in previous processor families. It has the added advantage that for SAP HANA, a sub-socket VM can be configured both non-productive and productive instances. This increases the density of SAP HANA instances on smaller nodes.

5.1.3 SAP HANA requirements

SAP Note [2718982](#) provides some basic rules to set up VMware vSAN with a ThinkAgile VX node to use with SAP HANA. In addition, Lenovo provides the following best practices:

- SAP HANA demands a minimum of 128GB RAM and 8 physical cores per VM for any SAP HANA workload. This correlates to a minimum of 16 virtual CPUs (vCPU) with hyperthreading turned on.
- The VM containing the productive SAP HANA appliance instance may use either half or all the physical cores contained on a set of processor sockets. For example, on a ThinkAgile VX 2U system, if two Intel

Xeon SP 8276 processors are used, then each of the two possible VMs will contain multiples of 28 physical cores (CPUs) and 28 vCPUs, or one large VM containing 56 vCPUs (2 × 28 physical cores).

- The SAP HANA Storage TDI guidance and calculations for SAP HANA Data, Log, and Shared directories should be used as described in this document.
- The SAP HANA Network TDI guide gives guidance on how to calculate the quantity of SAP HANA VMs on an HCI network with the following formula:

Network bandwidth / highest throughput KPI = number of VMs on the cluster possible.

5.2 ThinkAgile System Configuration

5.2.1 ThinkAgile UEFI Configuration

For better results of your VMware vSphere ESXi installation, see the VMware Guide “*SAP HANA on Hyper-converged Infrastructure (HCI) Solutions based on VMware vSAN [Guide](#)*” for the specific configuration changes necessary for SAP HANA. Those settings recommended by VMware are valid for ThinkAgile VX systems.

5.2.2 ThinkAgile Storage Configuration

A ThinkAgile VX node for SAP HANA must use an All-Flash (AF) configuration with a combination of NVMe and SAS/SATA SSD devices for the cache and capacity layers. We highly recommend any Lenovo Mainstream or Performance NVMe or Lenovo Performance SSDs as the cache layer while any Lenovo Entry, Mainstream or Performance SSDs may be used for the capacity layer. We recommend using VMware vSAN's own configuration rules for use with SAP HANA as a best practice.

5.2.3 ThinkAgile Sub-NUMA Configuration

Using Intel SNC to increase the VM density for SAP HANA on the ThinkAgile VX systems, follow the practices discussed in the VMware Paper, “THIS PAPER HERE” found on the VMware web [link](#).

5.2.4 VMware vSphere Configuration

For better results of your VMware vSphere ESXi installation, see the VMware Guide “*SAP HANA on Hyper-converged Infrastructure (HCI) Solutions based on VMware vSAN [Guide](#)*” for the specific configuration changes necessary for SAP HANA.

5.2.5 Operating System Configuration

Install the desired Linux distribution according to their installation guidelines for SAP. Apply operating system settings for SAP HANA inside the VM as recommended in SAP Note [2235581](#) - SAP HANA: Supported Operating Systems (SAP account required) for the Linux distribution chosen.

SAP requires that method to detect whether an SAP HANA VM runs on a certified SAP HANA HCI solution. VMware and SAP have created a SAP HANA on vSAN Health Check vib and Installation [script](#) that can be run on the ESXi host which allows during runtime if an SAP HANA instance is running on a vSAN datastore. This script is available from VMware directly.

The script, based on Python, runs every five minutes on the ESXi host. It should be installed by the customer or Lenovo partner into the VMware ESXi before the SAP HANA VMs are created. It adds the following information to the VM configuration file.

- "guestinfo.vsan.enabled": True/False,
- "guestinfo.SDS.solution": esxi_version, this is equivalent to the vSAN release
- "guestinfo.vm_on_vsan": True/False

Prerequisites:

- Enable SSH on VMware ESXi before installing the script. Deactivate SSH after installation.
- We recommend having identical passwords of the super user across all ESXi hosts. This may be changed afterwards.

5.2.6 File System Configuration

Along with the VMware Guide “*SAP HANA on Hyper-converged Infrastructure (HCI) Solutions based on VMware vSAN [Guide](#)*”, we also recommend to read the Lenovo Best Practices for ThinkAgile VX using SAP HANA [paper](#) for the specific configuration changes necessary for SAP HANA. Use the Linux logical volume management (LVM) tools to create your volumes for SAP HANA as described in these guides.

5.2.7 SAP HANA configuration

Apply SAP HANA settings inside the VM as recommended in SAP documentation for the version of SAP HANA you are installing. This can be found on help.sap.com/hana.

5.2.8 Check compliance

It is highly recommended to work with the latest version of the system check script. You can find it in SAP Note [2533844](#) – Lenovo Support Tool for SAP HANA Integrated Systems. Verify all hardware requirements and basic operating system requirements using this script.

Resources

VMware

VMware vSphere

VMware vSphere - www.vmware.com/products/vsphere

SAP and VMware - www.vmware.com/solutions/business-critical-apps/sap-virtualization.html

SAP on VMware Best Practices [Guide](#) (2015)

Blog: SAP HANA on vSphere – Support Status and Best Practices [Summary](#)

VMware vSAN

VMware Virtual SAN - www.vmware.com/products/vsan

SAP HANA on Hyperconverged Infrastructure (HCI) Solutions Powered by VMware vSAN™ [Guide](#)
[VSAN Design and Sizing Guide](#)

[SAP HANA on vSAN Health Check Vib and Installation Script](#)

VMware Knowledge Base Articles

VMware KB Article [2113954](#) - Understanding vSAN memory consumption in vSAN 7.0

VMware KB Article [55636](#) - Overview of 'L1 Terminal Fault' Speculative-Execution vulnerabilities

SAP SE

SAP Notes (Authorization required)

SAP Note [2235581](#) – SAP HANA: Supported Operating Systems

SAP Note [2533844](#) – Lenovo Support Tool for SAP HANA Integrated Systems

SAP Note [1788665](#) – SAP HANA Support for virtualized / partitioned (multi-tenant) environments

SAP Note [2718982](#) – SAP HANA on VMware vSphere and vSAN

SAP Note [2652670](#) – SAP HANA VM on VMware vSphere

SAP Note [2937606](#) – SAP HANA on VMware vSphere 7.0 in production

SAP Note [3102813](#) – SAP HANA on VMware vSphere 7.0 U2/U3 with up to 12 TB DRAM 448 vCPUs VM sizes

SAP Note [3372365](#) - SAP HANA on VMware vSphere 8

SAP Note [2161991](#) – VMware vSphere configuration guidelines

SAP Note [1900823](#) – SAP HANA Storage Requirements Whitepaper

SAP Note [2000003](#) – FAQ: SAP HANA

SAP Note [171356](#) – SAP Software on Linux: General Information

SAP Note [1552925](#) – Linux: High Availability Cluster Solutions

SAP Note [1501701](#) – Single Computing Unit Performance and Sizing

SAP Note [1612283](#) – Hardware Configuration Standards and Guidance

SAP Customer Network (SCN)

SAP HANA TDI [Overview](#)

SAP HANA TDI [FAQ](#)

SAP HANA Network [Guidelines](#):

Community Blog: [SAP on VMware vSphere](#)

Community Blog: [SAP HANA on VMware vSphere](#)

Lenovo

ThinkAgile VX Data Sheets and Product Guides

Intel Xeon 3rd Generation Scalable Processors

- ThinkAgile VX Solution for SAP HANA: lenovopress.lenovo.com/ds0118
- ThinkAgile VX7820 Series: lenovopress.lenovo.com/lp1338
- ThinkAgile VX 4U node: lenovopress.lenovo.com/lp1339
- ThinkAgile VX 4U node for SAP HANA: lenovopress.lenovo.com/lp1341

Intel Xeon 4th Generation Scalable Processors

- ThinkAgile VX Series: lenovopress.lenovo.com/ds0104
- ThinkAgile VX650 V3 2U Integrated System and VX650 V3 2U Certified Node:
lenovopress.lenovo.com/lp1673

Document History

Version 2.0	1 March 2024	Add Intel Xeon Scalable Processor 4th Generation (Sapphire Rapids) on ThinkAgile VX V3 2-Socket systems Add VMware 8.0 support Remove VMware 6.7 support Remove ThinkAgile VX 2U systems based on Intel Xeon 2 nd Gen. SP
Version 1.8	5 May 2021	Update ThinkAgile VX 4U 8-Socket configurations Adding SAP Notes 2718982 & 3102813 regarding VMs larger than 4-Sockets. Included a statement in network sizing of SAP HANA
Version 1.7	5 December 2020	Add ThinkAgile VX 4U 8-Socket configurations
Version 1.6	28 September 2020	Add Intel Optane Persistent Memory App-Direct Mode Clean cache and capacity device SAP HANA recommendations
Version 1.5	14 July 2020	Add Intel Xeon Scalable Processor 2. Generation (Cascade Lake) Add ThinkAgile VX 4U 4-Socket models
Version 1.0	6 February 2019	Initial version

© 2024 Lenovo. All rights reserved.

Availability: Offers, prices, specifications and availability may change without notice. Lenovo is not responsible for photographic or typographical errors. **Warranty:** For a copy of applicable warranties, write to: Lenovo Warranty Information, 1009 Think Place, Morrisville, NC, 27560, Lenovo makes no representation or warranty regarding third party products or services. **Trademarks:** Lenovo, the Lenovo logo, ThinkSystem, ThinkAgile, System x, ThinkServer are trademarks or registered trademarks of Lenovo. Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, the Intel Inside logo, Xeon and Optane are registered trademarks of Intel Corporation in the U.S. and other countries. Other company, product, and service names may be trademarks or service marks of others.

03-2024