

The Lenovo logo is displayed in white text on a black rectangular background.

Reference Architecture for SAP Applications on ThinkAgile HX Series

Last update: 1 March 2024
Version 2.0

Provides a technical overview of
Lenovo ThinkAgile HX Series
systems

Contains SAP HANA
configuration and sizing
recommendations

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

Explains reliability and
performance features of
hyper-converged appliances

Patrick Hartman



Table of Contents

1	Introduction	1
2	Technical overview of appliances	2
2.1	ThinkAgile HX series	2
2.2	Software components	3
2.3	Data network components	6
2.4	Reliability and performance features	7
3	SAP Application Workloads	12
3.1	SAP SE solutions overview	12
3.2	SAP Business Suite and NetWeaver for traditional DBs	12
3.3	Operational model	14
3.4	Sizing SAP applications	16
3.5	SAP Business Suite deployment	16
3.6	SAP HANA deployment	17
4	Other best practices	24
4.1	Nutanix Acropolis	24
4.2	Guest Operating system	24
4.3	Guest File system	24
4.4	SAP HANA parameters	24
4.5	Check compliance	24
	Resources	25

1 Introduction

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

ThinkAgile HX Series 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 the performance.

Chapter 2 provides a technical overview of ThinkAgile HX Series and explains why the combination of Lenovo servers and Nutanix software provides best of breed system performance and reliability. Chapter 3 describes the SAP SE virtualized workloads and provides recommendations on what model to use and how to size the system to that workload. Some best practice recommendations are also listed. ThinkAgile HX Series Integrated Systems (IS) and Certified Nodes (CN) are not limited to just the workloads described in this reference architecture and can execute any virtualized workload on the supported hypervisors.

This Reference Architecture describes the following workloads:

- SAP Business Applications
- SAP HANA Platform Edition

2 Technical overview of ThinkAgile HX Series

This chapter provides an overview of the ThinkAgile HX Series including the associated software, systems management, and networking. The last section provides an overview of the performance and reliability features.

2.1 ThinkAgile HX series

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

- Enterprise workloads, like SAP HANA
- Private and hybrid clouds
- Remote office and branch office (ROBO)
- Server virtualization for SAP NetWeaver
- Virtual desktop infrastructure (VDI)
- Small-medium business (SMB) workloads, like SAP Business One

Starting with as few as three nodes to keep your acquisition costs down, the Lenovo ThinkAgile HX Series IS and CN nodes are capable of immense scalability as your needs grow.

Lenovo ThinkAgile HX Series nodes are available in five families that can be tailored to your SAP workloads:

- Lenovo ThinkAgile HX630 V3 ROBO IS/CN: optimized for ROBO environments
- Lenovo ThinkAgile HX630 V3 IS/CN: optimized for SMB environments including SAP HANA
- Lenovo ThinkAgile HX650 V3 Storage Integrated IS/CN: optimized for storage-heavy workloads
- Lenovo ThinkAgile HX650 V3 IS/CN: optimized for mission critical workloads including SAP HANA
- Lenovo ThinkAgile HX645 V3 IS/CN: optimized for general virtualization environments and ROBO
- Lenovo ThinkAgile HX665 V3 IS/CN: optimized for enterprise applications

Table 1 shows the similarities and differences between ThinkAgile HX Series nodes.

Table 1: Comparison of ThinkAgile HX Series systems

Feature	HX Series Appliances	HX Series certified nodes
Validated and integrated hardware and firmware	Yes	Yes
Certified and preloaded with Nutanix software	Yes	Yes
Includes Nutanix licenses	Yes	No
ThinkAgile Advantage Single Point of Support for quick 24/7 problem reporting and resolution	Yes	Yes
Includes deployment services	Optional	Optional

For more information about the system specifications and supported configurations, refer to the product guides for the Lenovo ThinkAgile HX Series systems based on the Intel Xeon Scalable processor. For appliances see the ThinkAgile HX Series [data sheet](http://lenovopress.lenovo.com/ds0019) (lenovopress.lenovo.com/ds0019). For certified nodes see the ThinkAgile HX Certified Nodes [data sheet](http://lenovopress.lenovo.com/ds0046) (lenovopress.lenovo.com/ds0046).

In addition, the ThinkAgile HX Solution for SAP HANA [data sheet](https://lenovopress.lenovo.com/ds0068) (lenovopress.lenovo.com/ds0068) lists certain ThinkAgile HX Series systems 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 amount of options which can be configured for use with the high performance in memory database.

2.2 Software components

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

2.2.1 Hypervisor

The ThinkAgile HX Series nodes (generally) support the following hypervisors:

- Nutanix Acropolis Hypervisor (AHV)
- VMware ESXi vSphere hypervisor
- Microsoft Hyper-V hypervisor

The ThinkAgile HX Series nodes for SAP HANA support the following hypervisors:

- Nutanix Acropolis Hypervisor (AHV)
- VMware ESXi vSphere hypervisor

2.2.2 Lenovo XClarity Administrator

Lenovo XClarity Administrator is a centralized systems management solution that helps administrators deliver infrastructure faster. This solution integrates easily with Lenovo servers, ThinkAgile HX Series 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 used to manage firmware upgrades outside of the Nutanix Prism web console. Note that XClarity should not be used to install hypervisors and Nutanix Foundation should be used instead.

Lenovo XClarity Administrator is provided as a virtual appliance that can be quickly imported into a virtualized environment. XClarity can either be installed on a separate server or a server within a Nutanix cluster providing that the hardware management network with the server XClarity Controller (XCC) is routable from the server hosting the XClarity VM. Figure 1 shows the Lenovo XClarity administrator interface.

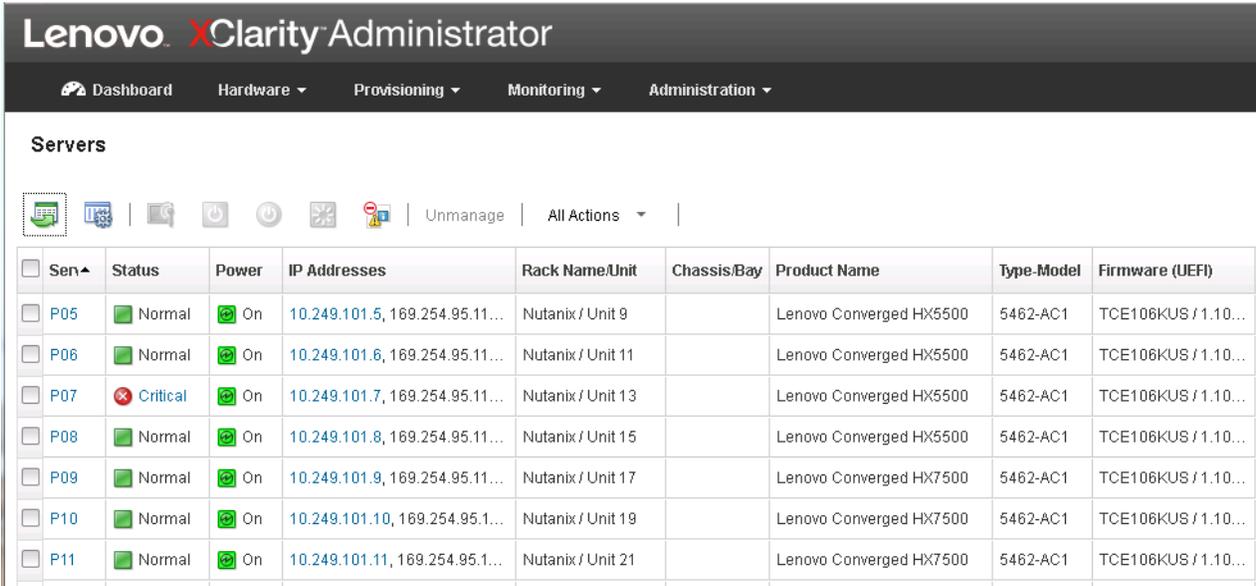


Figure 1: XClarity Administrator interface

2.2.3 Nutanix Prism

Nutanix Prism gives administrators a simple and elegant way to manage virtual environments. Powered by advanced data analytics and heuristics, Prism simplifies and streamlines common workflows within a data center.

Nutanix Prism is a part of the Nutanix software preloaded on the appliances and offers the following features:

- Single point of control
 - Accelerates enterprise-wide deployment
 - Manages capacity centrally
 - Adds nodes in minutes
 - Supports non-disruptive software upgrades with zero downtime
 - Integrates with REST APIs and PowerShell
- Monitoring and alerting
 - Tracks infrastructure utilization (storage, processor, memory)
 - Centrally monitors multiple clusters across multiple sites
 - Monitors per virtual machine (VM) performance and resource usage
 - Checks system health
 - Generates alerts and notifications
- Integrated data protection
 - Offers customizable RPO/RTO and retention policies
 - Supports configurable per-VM replication (1:1, 1:many and many:1)
 - Provides efficient VM recovery

- Deploys affordable data recovery (DR) and backup to the cloud
- Diagnostics and troubleshooting
 - Provides time-based historical views of VM activity
 - Performs proactive alert analysis
 - Correlates alerts and events to quickly diagnose issues
 - Generates actionable alerts and reduces resolution times
 - Analyzes trending patterns for accurate capacity planning

2.2.4 Nutanix Foundation

Nutanix Foundation is a separate utility that you use to orchestrate the installation of hypervisors and Nutanix software on one or more nodes. The maximum number of nodes that can be deployed at one time is 20.

Foundation is available both as a stand-alone VM and also integrated into the CVM. Because CVM is pre-installed in the factory, the CVM integration of Foundation simplifies the deployment and cluster creation of new servers delivered from the factory.

The dual M.2 boot drives must be configured as a RAID 1 mirrored array for installation to be successful.

2.2.5 Nutanix Controller VM

The Nutanix Controller VM (CVM) is the key to hyper-converged capability and each node in a cluster has its own instance. Figure 2 shows the main components of the CVM.

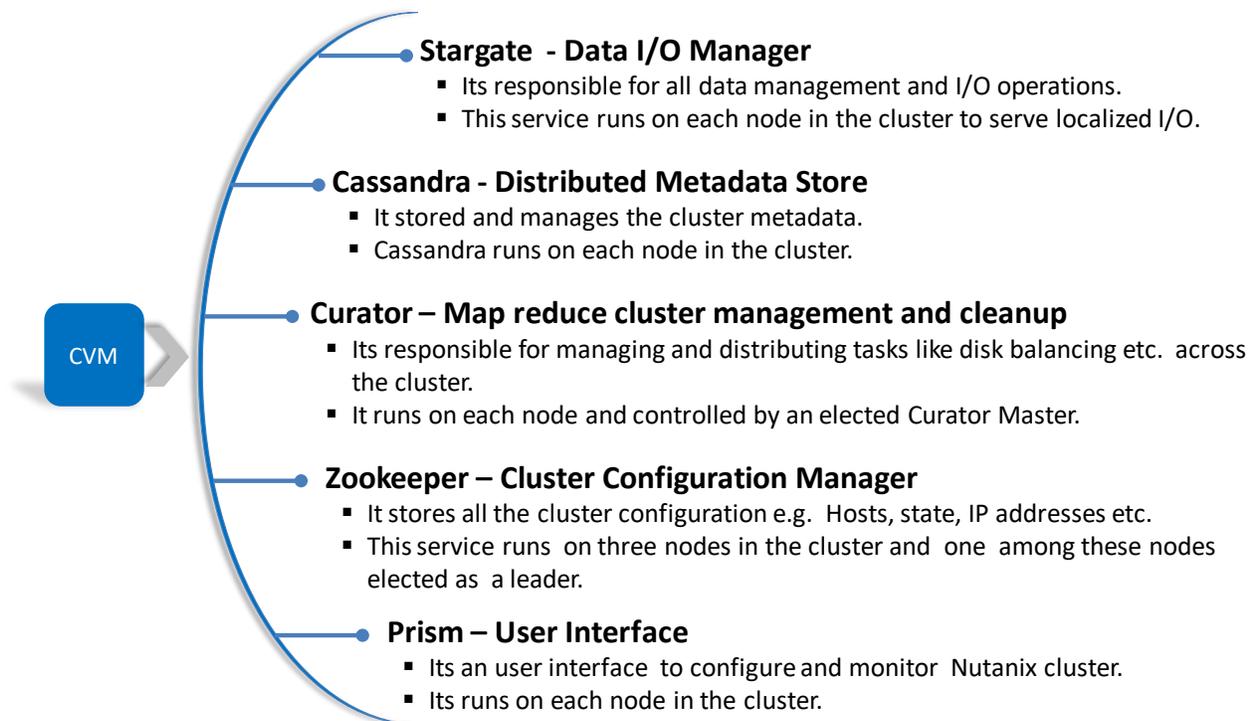


Figure 2: Controller VM components

The CVM works as interface between the storage and hypervisor to manage all I/O operations for the hypervisor and user VMs running on the nodes as shown in Figure 3.

2.4 Reliability and performance features

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

Hardware reliability

Lenovo uses the highest quality hardware components combined with firmware that is thoroughly tested. As a consequence, Lenovo servers have been rated #1 in hardware reliability for the last 9 years (ITIC – lenovopress.lenovo.com/lp1117). This is important as it lowers the frequency of a server failure which in turn lowers OPEX.

A ThinkAgile HX Series has redundant hardware components by including two power supplies, multiple chassis fans, two Intel CPUs, multiple memory DIMMs, multiple SSDs and HDDs, and optionally up to two dual-port network interface cards.

Hardware performance

The ThinkAgile HX Series nodes have been carefully designed for performance. In addition to all of the usual attributes like processors and memory, the 24 drive HX650 V3 uses three HBA controllers instead of the one. As a consequence, the latency is halved for some workloads that heavily utilize the cold tier. This allows a higher throughput and improved transaction rates.

Distributed file system

The Nutanix Distributed file system (NDFS) is an intelligent file system which virtualizes the local attached storage (SSD/HDD) on all the nodes in a cluster and presents it as single storage entity to cluster. Figure 4 shows the high level structure of NDFS:

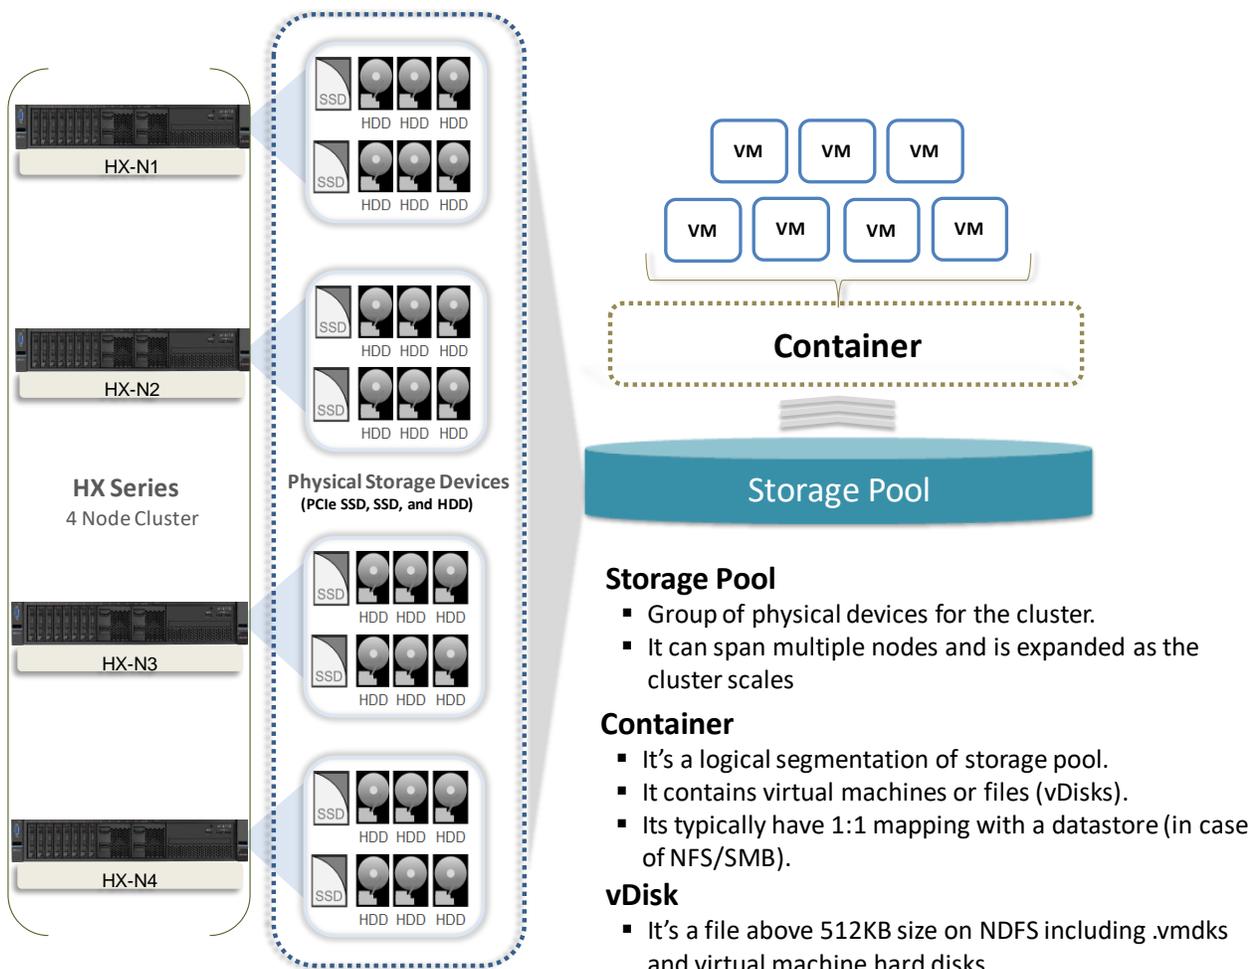


Figure 4: Nutanix Distributed File System

Data protection via replication

The Nutanix platform replication factor (RF) and checksum is used to ensure data redundancy and accessibility in the event of a node or disk failure or corruption. It uses an OpLog which acts as a staging area for incoming writes on low latency SSDs which are then replicated to the OpLogs for one or two other Controller VMs before acknowledging a successful write. This approach ensures that data available in at least two to three different locations and is fault tolerant. While the data is being written a checksum is calculated and stored as part of its metadata.

In the case of a drive or node failure, that data is replicated out to more nodes to maintain the replication factor. A checksum is computed every time the data is read to ensure the data validity. If the checksum and data mismatch, then the data replica is read to replace the invalid copy.

Performance with data tiering

Nutanix uses a disk tiering concept in which disk resources (SSD and HDD) are pooled together to form a cluster wide storage tier. This tier can be accessed by any node within the cluster for data placement and can leverage the full tier capacity. The following data tiering functions are provided:

- The SSD on a local node always has the highest tier priority for write I/O.
- If the local node's SSD is full then the other SSDs in the cluster are used for I/O.
- The NDFS Information Lifecycle Management (ILM) component migrates cold data from the local SSD to HDD to free up SSD space. It also moves heavily accessed data to the local SSD to provide high performance.

Performance by data locality

Data locality is a crucial factor for cluster and VM performance. In order to minimize latency the CVM will work to ensure that all I/O happens locally. This ensures optimal performance and provides very low latencies and high data transfer speeds that cannot be achieved easily with shared storage arrays, even if all-flash.

The following occurs in case of a VM migration or high availability event that moves a VM from Node-A to Node-B:

- The VM's data is provided by the CVM running on Node-B.
- All write I/O requests occur locally i.e. to the local storage of Node-B.
- When a request comes for reading old data, the I/O request is forwarded by Node-B to Node-A. NDFS detects that the I/O request originated from different node and migrates the data locally in the background i.e. from Node-A to Node-B so that all subsequent read I/O operations are served locally. This approach (migration only on a read) helps to avoid network flooding.

Performance of snapshots and clones

NDFS provides support for offloaded snapshots and clones using a redirect-on-write algorithm. When a snapshot or clone is created, the base vDisk is marked as read only and another vDisk is created with read/write permissions as shown in

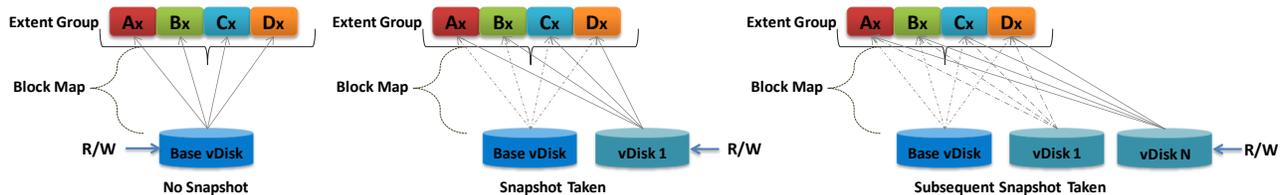


Figure 5 and Figure 6 below.

At this point both vDisks have the same block map - a metadata mapping of the vDisk to its corresponding extents. This approach reduces the overhead of creating snapshots and allows snapshots to be taken very quickly with little performance impact.

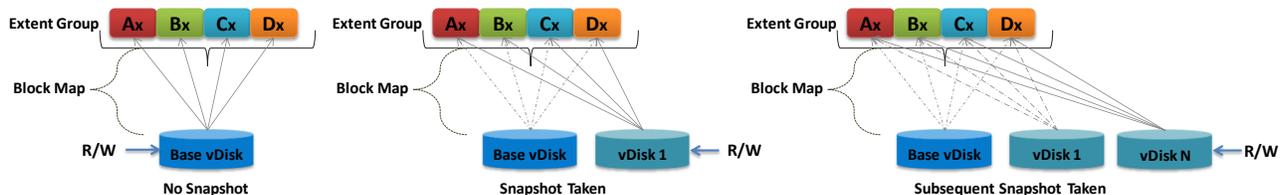


Figure 5: VM snapshots

When a VM is cloned the current block map is locked and then clones are created. These updates are metadata only so again no actual I/O takes place. The logic applies for clones of clones as well where a previously cloned VM acts as a base vDisk. All the clones inherit the prior block map and any new writes take place on the individual block maps.

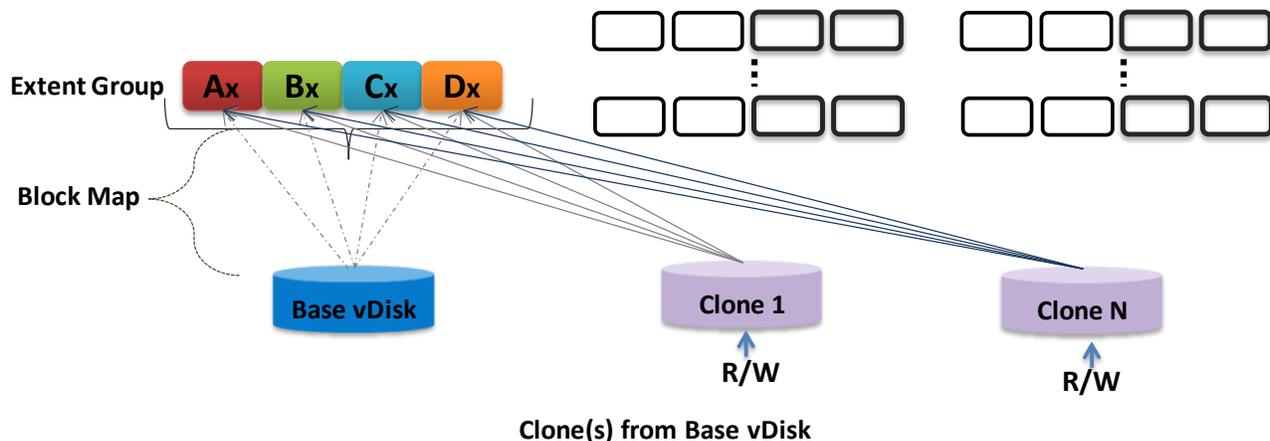


Figure 6: VM clones

Storage reduction via De-duplication and Compression

The Nutanix elastic de-duplication engine increases the effective capacity of a disk, as well as the RAM and cache of the system by removing duplicate data. It's an intelligent technology which performs following actions to increase storage efficiency:

- Sequential streams of data fingerprinted at 4K granularity
- Single instance of the shared VM data is loaded into the cache upon read
- Each node in a cluster performs its own fingerprinting and deduplication

The Nutanix capacity optimization engine is responsible for performing data transformations and compression to achieve data optimization. NDFS provides following compression methods:

- In-line compression sequential streams of data or large I/O sizes are compressed in memory before written to the disk
- Post-process compression whereby data is written in an uncompressed state and the curator framework is used to compress the data in a cluster wide manner

The Nutanix capacity optimization engine uses the Google snappy compression library to deliver good compression ratios with minimal compute overhead and very fast compression or decompression rates.

Elimination of “split-brain” errors

In a distributed system it is possible for one participant to become disconnected which will cause differences in the stored data. NDFS uses the proven “Paxos” algorithm to eliminate these “split-brain” issues by reaching a consensus (quorum) among the participants in a distributed system before the writes are made.

Drive reliability via active monitoring

The CVM actively monitors the performance of every drive in a node. The deterioration of a drive's performance may indicate that the drive is about to fail. The CVM proactively moves data off the drive before it fails and marks the drive offline and in need to replacement. The idea is to avoid the expensive data transfers to maintain data redundancy and possible loss of data.

3 SAP Application Workloads

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

3.1 SAP SE solutions overview

Error! Reference source not found. shows these solutions on the ThinkAgile HX platform, which can be categorized as follows:

- Enterprise Resource Planning
 - SAP S/4HANA®
 - SAP Business Suite®
- Digital Platform
 - SAP HANA® database
 - SAP Sybase ASE® (and other databases, like Oracle®, IBM DB2®, or Microsoft SQL Server®)
 - SAP NetWeaver® platform
- Business Analytics
 - Business Intelligence with SAP BW/4HANA®
- Small & Medium Businesses
 - SAP Business One®

3.2 SAP Business Suite and NetWeaver for traditional DBs

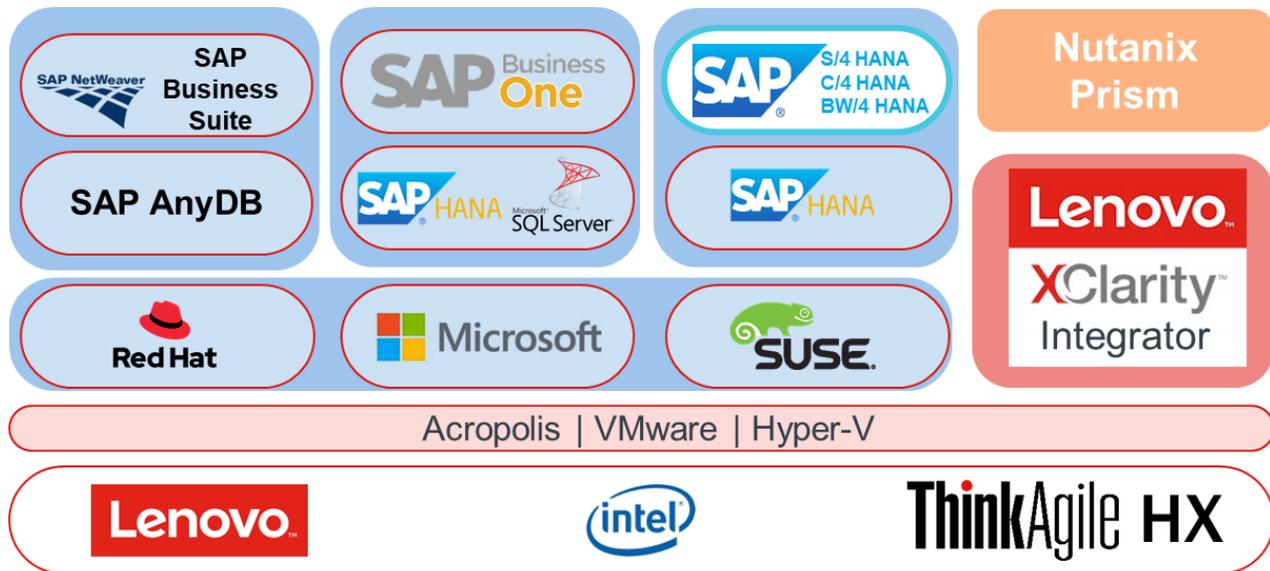


Figure 7: SAP Applications on ThinkAgile HX

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

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

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

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

3.3 SAP Business Platform applications for SAP HANA

SAP SE bases their newest business application offerings on a propriety in-memory database called SAP HANA. These solutions by SAP SE include applications for enterprise resource planning (SAP S/4HANA) and business analytics (SAP BW/4HANA) and small and medium businesses (SAP Business One).

- **SAP S/4HANA** offers a new user experience with SAP Fiori® that is personalized and ready for consumers. Whatever industry or business, SAP S/4HANA will drive instant value for you as a customer.
- **SAP BW/4HANA** is the next-generation data warehouse solution that can help capitalize on the value of your data, whether from SAP applications or third-party solutions. Even unstructured data like geospatial or Hadoop-generated data lakes can be analysed.
- **SAP Business One powered by SAP HANA** is an enterprise resource planning (ERP) system for small and medium sized businesses.

3.3.1 SAP HANA relational database

The in-memory high performance row- and column-orientated database developed by SAP SE primarily performs standard SQL based database transactions to store and retrieve data just like any traditional relational database. Via the innovative column store architecture, it can additionally perform advanced analytical operations which can simplify customer's business processes. Before this technology advance, separate data warehouses were needed using applications such as SAP Business Warehouse®.

3.4 SAP Business One for Microsoft SQL Server and SAP HANA

SAP Business One is an enterprise resource planning (ERP) system for small and medium sized businesses. Its focus is to automate key business processes like financials, operations and human resources and is arranged into 15 functional modules, which cover typical business functions.

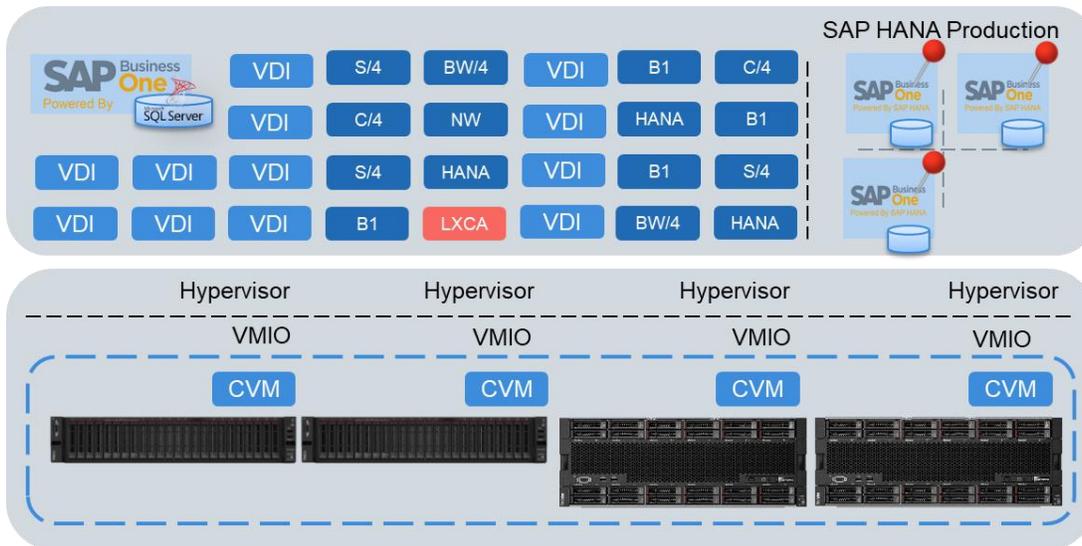


Figure 8: SAP Business One System Diagram

The SAP Business One architecture is a client-server model where the data can be stored traditionally in a Microsoft SQL Server (running under Windows) or more recently in an SAP HANA Platform Edition database (running under Linux). These server components run very well with the ThinkAgile HX family of servers.

3.5 Operational model

This section describes different models for deploying the various components in the SAP Business Application suite.

3.5.1 SAP application landscapes

Logically an application has 3 tiers: presentation layer, business logic or application layer, and database layer. Scalability is key to a successful implementation of an SAP solution.

As a business grows, SAP provides scaling using 2- and 3-tier deployment architectures as follows:

- Scale up (also known as vertical scaling) is provided using a 2-tier architecture. This means that the presentation layer is one tier and the application and database layer is mapped to the second tier.
- Scale out (also known as horizontal scaling) is provided using a 3-tier architecture. This means that each of three logical tiers maps to a physical server tier.

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, no 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.
- Some SAP applications, like SAP Advanced Planner and Optimizer (APO), require a 3-tier architecture.

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. 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 (PPD)
- Production system (PRD)

Figure 9 shows an example of a typical landscape for an SAP customer.

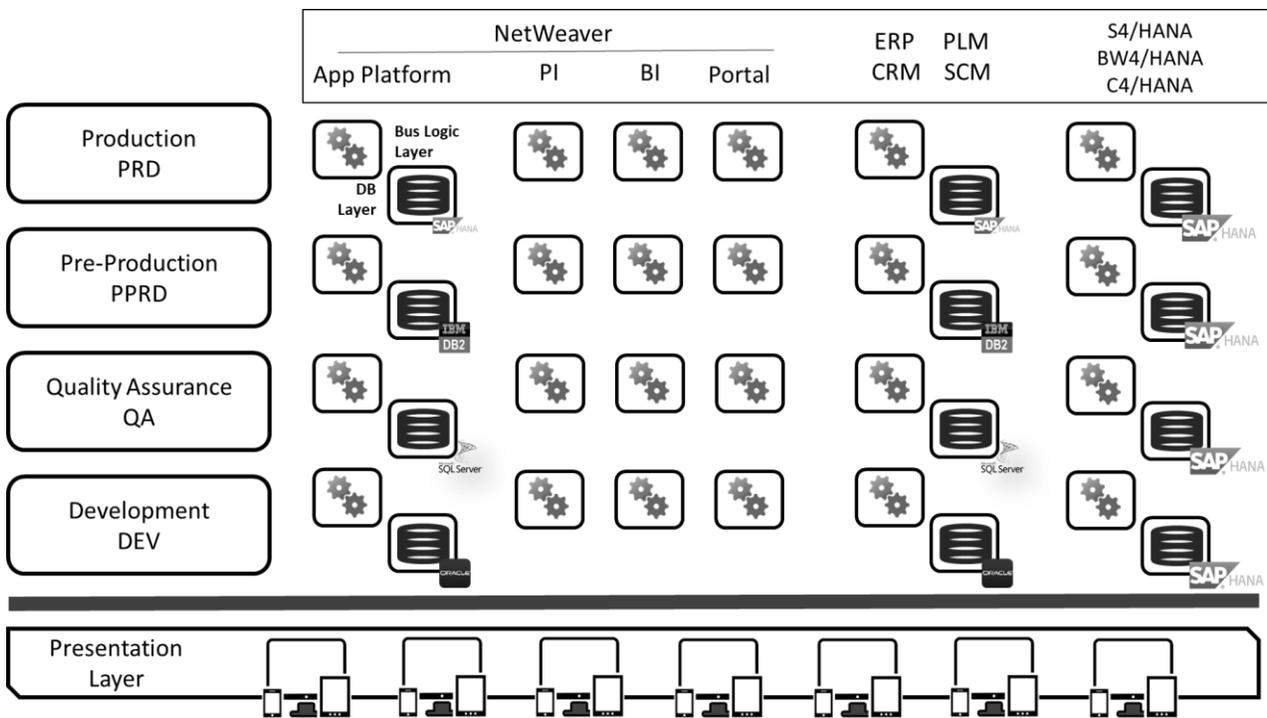


Figure 9: SAP customer landscape

3.5.2 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 the hypervisor vendors, the performance overhead of each hypervisor has been measured as follows:

- Nutanix Acropolis Hypervisor (AHV) – 4%
- VMware ESXi – 10%

These overheads should always be included when sizing a new hardware platform. In addition, Nutanix is a software defined, or hyperconverged, infrastructure. This means there are storage overhead costs that must also be included in the overall calculations. The Nutanix Distributed file system (NDFS) is controlled by a virtual machine called the Control VM, or CVM. This storage controller takes up to four (4) physical processor cores and up to 64GB of physical memory per node additionally. This can be configured in Nutanix Prism.

3.6 Sizing SAP applications

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

A 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 upon standard assumptions and benchmark performance data. It cannot and should not replace capacity planning for installed systems.

3.7 SAP Business Suite deployment

SAP NetWeaver and SAP Business Suite are certified to be run on the Lenovo ThinkAgile HX Series systems. Lenovo ThinkAgile HX Series supports the following hypervisors for use with all SAP applications:

- Microsoft Hyper-V for applications running on Microsoft Windows

- VMware vSphere ESXi for applications running on Microsoft or Linux operating systems
- Nutanix Acropolis Hypervisor (AHV) for applications running on Linux operating systems

See the following SAP notes when deploying SAP NetWeaver applications on Acropolis Hypervisor:

- Note 2656072 - Steps required to run virtualized SAP applications on Nutanix AHV

Additionally, the following SAP notes are useful when deploying SAP applications in a virtualized environment:

- Note 1492000: General Support Statement for Virtual Environments
- Note 1122387: Linux: SAP Support in Virtualized Environments
- Note 171356: SAP Software on Linux: General Information
- Note 1552925: Linux: High Availability Cluster Solutions
- Note 1501701: Single Computing Unit Performance and Sizing
- Note 1612283, Sizing SAP VMs within a NUMA node when configuring SAP application VMs
- Note 1612283: Hardware Configuration Standards and Guidance

3.8 SAP HANA deployment

SAP HANA is certified to run on the Lenovo ThinkAgile HX Series systems. Lenovo ThinkAgile HX supports the following hypervisors for use with all SAP applications:

- VMware vSphere ESXi for SUSE and Red Hat Linux operating systems
- Nutanix Acropolis Hypervisor (AHV) for SUSE and Red Hat Linux operating systems

See the following SAP notes when deploying SAP NetWeaver applications on Acropolis Hypervisor:

- Note 2686722 - SAP HANA virtualized on Nutanix Acropolis Hypervisor

In order to deploy an instance of SAP HANA, you need to correctly size and configure the hypervisor, guest operating system, file system and SAP HANA. The following are sizing guidance for using ThinkAgile HX Series systems with SAP HANA.

3.8.1 ThinkAgile HX configuration

For installations of SAP HANA, certain minimum and maximum configuration values must be considered before sizing a VM and the entire landscape. Table 2 is derived directly from the SAP Note [2686722](#) and the Nutanix [Best Practices Guide for SAP HANA](#).

Table 2: SAP HANA limitations on Nutanix Acropolis OS

Description	Maximum	Minimum
Virtual Memory (vMEM) allowed	2048 GiB	128 GiB
Virtual CPUs (vCPUs) allowed	60 (120 with hyper-threading)	8 (16 with hyper-threading)
Sockets / Productive SAP HANA use	3	1
Number of ThinkAgile HX Nodes	64	3

3.8.2 Nutanix Cluster Guidance

Below is a list of deployment guidance for running SAP HANA in a ThinkAgile HX cluster. We presume here a single instance of SAP HANA is installed on a cluster.

- One ThinkAgile HX appliance or certified node for productive SAP HANA which has been certified for SAP HANA
- Two or more ThinkAgile HX630 V3 or HX635 V3 nodes for other application workloads
- All nodes in a Nutanix cluster must contain similar storage layer
 - The physical number of disks must be equivalent in each node
 - Nodes not running productive SAP HANA workloads
- For production SAP HANA workloads
 - Do not place any other VM on the same socket(s) running the productive VM.
 - Do not share resources across any VMs and the productive VM(s)
 - Allow only whole multiples of processor sockets (1, 2 or 3) per productive VM
- For non-production SAP HANA workloads
 - The same configurations for productive SAP HANA workloads are valid
 - Half socket non-productive workloads are additionally allowed, but may not run on same socket as a productive SAP HANA workload
 - The socket running the Nutanix Control VM (CVM) can run non-productive SAP HANA workloads
- For all other workloads
 - Any other workload can run on any socket not running an instance of a productive SAP HANA workload
- Turn off any Nutanix Distributed file system (NDFS) features that impact storage performance such as: compression, deduplication or erasure coding (EC-X)

3.8.3 SAP HANA configuration guidance

SAP Note [2686722](#) provides some basic rules to set up a ThinkAgile HX node for use with SAP HANA.

- SAP HANA demands a minimum of 128GB RAM and 8 physical cores or 16 threads per VM for any SAP HANA workload.
- The VM containing the productive SAP HANA appliance instance must use all physical cores contained on a whole number of processor sockets (1, 2 or 3). For example, if four Intel Xeon SP 8176 processors are used, then each of the three possible VMs will contain multiples of 28 physical cores (CPUs) or 56 virtual CPUs (vCPUs) with hyperthreading.
- 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 Storage 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.

In addition, the following SAP notes should be helpful:

- SAP Note [2428012](#) – SAP on Nutanix
- SAP Note [2686722](#) – SAP HANA virtualized on Nutanix Acropolis Hypervisor

- SAP Note [2652670](#) – SAP HANA VM on VMware vSphere
- SAP Note [2656072](#) – Steps required to run virtualized SAP applications on Nutanix AHV

There are 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.

Processor Rules:

- A quantity of two (2) or four (4) Intel Xeon Scalable Processors (SP) are required
- Based on the CPU model, a minimum of at least eight (8) cores per processor is necessary

Memory Rules:

- Either one (1) or two (2) fully populated memory channel(s) with twelve (12) similar TruDDR4 DIMMs each is allowed
- SAP HANA requires having a homogeneous placement of memory across each populated DIMM channel.
- For Intel Xeon Scalable Processors, either one (1) or two (2) DIMM channels per socket may be populated for a total of 16 or 32 physical DIMMS populated in the system
- Memory sizes may be mixed if the DIMM type is the same (RDIMM, LRDIMM or 3DS RDIMM) and must be added in pairs. The sizes must be neighbours to each other (8/12, 16/32, or 32/64) and matching technology (RDIMM, LRDIMM, RDIMM 3DS).

Networking Rules:

- At least 2 pairs of a minimum 10GbE NICs are necessary,
 - 1 pair for SAP HANA and 1 pair for Nutanix Files
 - Two (2) Mellanox ConnectX-4 Lx 2x25GbE SFP28 adapters are recommended
- Other adapters may be added up to maximum supported by the ThinkAgile HX system
- If NVMe devices are used, then RDMA enabled adapters must be used, like the Mellanox ConnectX-4 series.

GPU Rules:

- ThinkAgile HX systems for SAP HANA do not support GPUs.

Storage Rules:

- A minimum number of disks dependent storage configured
 - Systems configured with SATA/SAS SSD require at least eight (8) devices per node
 - Systems configured with NVMe & SATA/SAS SSD devices require at least six (6) devices, two (2) NVMe and four (4) SSD
 - For every NVMe cache drive added; there should be at least two (2) SSD drives be configured
- In other words: If using NVMe devices (for highly transactional workloads),
 - The minimum quantity of NVMe is two (2) with a maximum of four (4).

- For configuration with up to six (6) x SSD, two (2) x NVMe are required.
- For configuration with more than six (6) x SSD, four (4) x NVMe are required.
- o If using SSD Devices only (for analytical workloads, or light transactional workloads)
 - The minimum quantity of SSD is eight (8).
 - The maximum depends upon the system used. (10-24)

Nutanix recommends using a NVMe based approach for SAP HANA databases which are highly transactional based. If applications are not highly transactional a configuration with only SSD devices is acceptable. It is best to size all VMs first using SAP HANA TDI Storage Guidelines (rule of thumb is roughly 2.5+mem, but it can be less or more) and then add that together.

Finally go to the Nutanix Storage Calculator online and determine how much effective storage you would have in the cluster you create with the drives configured. To define the storage available on your HX cluster you should use the [Nutanix storage sizing tool](#) and follow the below rules:

- Storage efficiency set to none (1:1).
- Storage deduplication disabled (ECX).
- Set failover plan to N+1.
- Do not include cache drives in the storage space calculation.
- Only SSD used: spinning drives not allowed.
 - o NVMe devices are recommended when highly transactional workloads are to be used

Software Rules:

- o Nutanix Acropolis OS (AOS) 5.20 LTS or 6.5.4 LTS should be used with SAP HANA applications and either:
 - Acropolis Hypervisor (AHV) 20190916 or 20220304.439 major versions
 - VMware vSphere ESXi 6.7 U3 hypervisor
 - VMware vSphere ESXi 7.0 U1, U2 or U3c hypervisor
- o All system management software FC and Option PN are allowed.
- o An additional Professional Service Unit exists for installation and deployment of SAP HANA

3.8.4 SAP HANA deployment example

In order to size a ThinkAgile system correctly with SAP HANA, one needs to consider various system requirements. This section contains an example to illustrate the concepts and best practices.

The size of an SAP HANA VM is the result of sizing the memory, processor, storage and network components for each instance. Additionally, any non-production SAP HANA workload and other SAP applications need to be calculated into the total to find the appropriate Lenovo ThinkAgile HX solution size. An example configuration begins with the customer request to place the following SAP applications in a ThinkAgile HX cluster.

- S/4HANA Production (PRD) with 1500GB of memory
- S/4HANA Pre-Production (PrePRD) with 1500GB of memory

- S/4HANA System Recover (HSR) instance with 1500GB of memory
- S/4HANA QA (QAS) instance with 1500GB of memory
- Four (4) SAP HANA (HANA) development instances with 256GB of memory each
- BW/4HANA Production (PRD) instance with 750GB of memory
- BW/4HANA Pre-Production (Pre) instance with 750GB of memory
- BW/4HANA System Recovery (HSR) instance with 750GB of memory
- BW/4HANA QA (QAS) instance with 750GB of memory

The customer wishes to purchase as few servers as possible. They wish to place their productive S/4HANA instance on a four socket (4S) system, while the productive BW/4HANA system could run on a two socket (2S) system. Following these above restrictions, it can be shown (see Figure 10) that a minimum of five (5) servers need to be configured.

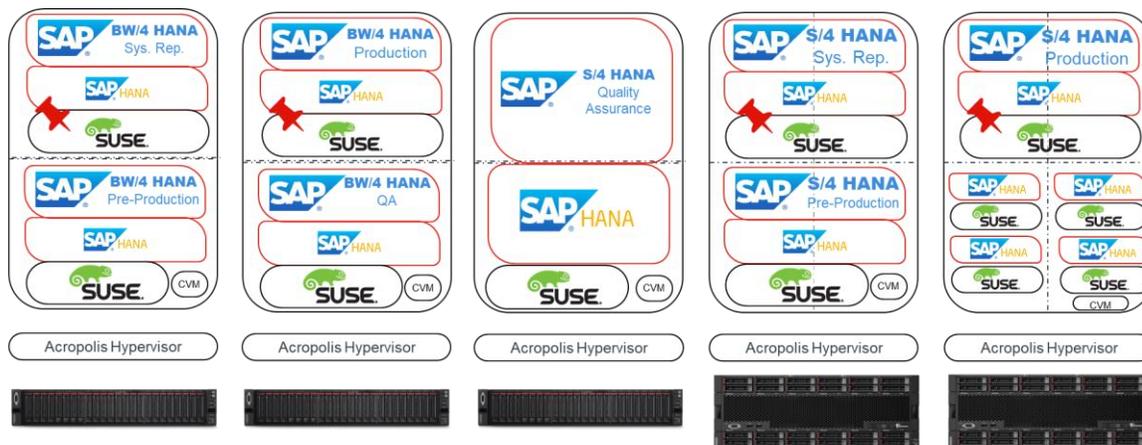


Figure 10: SAP HANA Production VM pinned  to a set of processor sockets

Memory sizing

For SAP HANA the main driver in a configuration is the amount of main memory for the database. The total amount of system memory is the size of the memory for the SAP HANA VMs plus memory for any other VMs plus the hypervisor and a small 4GB overhead for Nutanix Files metadata.

SAP HANA requires having a homogeneous placement of memory across all populated DIMM channels. For 2nd Generation Intel Xeon Scalable processors, there can be either one (1) or two (2) DIMMs per channel (up to 6 DIMMs per sockets channel) for a total of 24 DIMMS on a 2 socket or 48 DIMMS populated on a 4 socket system. For 4th Generation Intel Xeon Scalable processors, there can be either one (1) or two (2) DIMMs per channel (up to 8 DIMMs per sockets channel) for a total of 32 DIMMS on a 2 socket.

Error! Reference source not found. provides the recommended system memory depending on how one or more SAP HANA VMS are distributed across each of the processors.

Table 3: SAP HANA virtual machine memory sizing

1 Processor VM memory size (all)	2 Processor VM memory size (HX782x)	3 Processor VM memory size (HX782x)	Total System Memory for ThinkAgile HX782x
N/A	192GB	288GB	384GB
192GB	384GB	576GB	768GB
288GB	576GB	864GB	1152GB
384GB	768GB	1152GB	1536GB
576GB	1152GB	1728GB	2304GB
768GB	1536GB	2304GB	3072GB
1344GB	2304GB	3456GB	4608GB
1536GB	3072GB	4608GB	6144GB

1 Processor VM memory size (all)	Total System Memory for ThinkAgile HX630v3 / HX650v3
128GB	256 GB
256GB	512 GB
512GB	1024 GB
768GB	1536 GB
1024GB	2048 GB
1536GB	3072 GB
2048GB	4096 GB

Processor sizing

SAP recommends a minimum of eight (8) cores per socket. When using four (4) sockets, we highly recommend using a processor that has three (3) UPI ports and allows a fully meshed UPI architecture.

SAP HANA has defined that only whole number of sockets can be used for productive use, as well as no other workload may run on the same socket as the one executing a productive SAP HANA instance. The remaining sockets can be used for the Nutanix Control VM (CVM), other VMs such as SAP application servers, or other non-production SAP HANA instances.

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 the system performance of an SAP application, and 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 long-term load and usage information 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.

Storage sizing

The storage capacity calculated via the SAP HANA [TDI Guidelines for Storage](#) defines the amount of storage for the SAP HANA database and the operating system for each instance of SAP HANA. The TDI Guidelines also describe the requirements for various volumes needed for SAP HANA persistence, backup and the operating system.

Lenovo recommends increasing the operating system to 128GB and allocating a maximum swap space of 32GB. This is larger than the SAP SE recommendations of 50GB and 2GB respectively but has extra capacity for temporary file growth and certain root file system features such as btrfs snapshots.

The type of storage must be carefully considered along with the requirements of the customer. For example, faster NVMe devices should be used to handle workloads with high transactional properties while a pure SSD configuration can be used for analytical workloads.

In general, any solid-state device may be used which is supported by the ThinkAgile HX SAP HANA models according to the workload sizing from Nutanix. It is recommended to use NVMe devices when the percentage of transactional workload is expected to be very high. To improve performance of the ThinkAgile HX file system, you should also include RMDA enabled adapters and RDMA over converged ethernet enabled switches in order to obtain the best network throughput.

If a mixed NVMe and SAS/SATA configuration is chosen, the number of NVMe devices should be added in a ratio of one (1) NVMe device for every two (2) or more SATA/SAS SSD devices used. For other nodes in the same cluster not running SAP HANA productively, they need to include the same quantity and type of disks as the SAP HANA model. It is not necessary to use NVMe devices in these nodes.

Additional storage capacity may be needed for other VMs that reside on the ThinkAgile HX cluster and this should be factored into the overall cluster sizing calculation. Consult the Nutanix Storage Capacity [Calculator](#) for details on total cluster requirements.

Network sizing

Hyper-converged SAP HANA recommends network cards and switches that support Remote Direct Memory Access (RDMA) technology when using NVMe devices. Any switch described in “Data network components” on page 6 may be used for SAP HANA providing it supports the RDMA over Converged Ethernet (RoCE) standard.

If NVMe devices are not used, it is not necessary to use RDMA an RoCE, and; therefore, any standard network card adapter and switch may be used with at least a speed of 10GbE or higher.

We recommended that the ThinkAgile HX node for SAP HANA is configured with at least two Mellanox ConnectX-4 Lx 2x25GbE SFP28 adapters running in ethernet mode to provide appropriate network redundancy.

4 Other best practices

4.1 Nutanix Acropolis

For best results using Nutanix Acropolis OS and Prism software, see the Nutanix [Best Practices Guide](#) for the specific configuration changes necessary for SAP HANA.

4.2 Guest Operating system

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 your chosen Linux distribution chosen.

4.3 Guest File system

Use a supported file system refer to SAP Note [405827](#). Nutanix describes how to install a file system based on LVM in the instructions in the [Nutanix Best Practices Guide](#). All SAP HANA file systems should be based on the XFS file system, while the file system used for the OS can be chosen for the customer's best interest. More information can be found in the Lenovo ThinkAgile HX [Best Practice](#) guide.

4.4 SAP HANA parameters

Apply SAP HANA settings inside the VM as recommended in SAP documentation for the version of SAP HANA you are installing. Review the SAP HANA [implementation resources](#) and [help](#) documentation.

4.5 Check compliance

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

Resources

Nutanix

Nutanix [Portal](#) (requires registration)

Nutanix [Bible](#)

Nutanix [Tech Note](#): VMware vSphere Networking on Nutanix

Nutanix [Best Practices Guide for SAP HANA](#)

VMware [vSphere](#)

Acropolis Hypervisor

SAP HANA on Nutanix [FAQ](#)

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](#)

SAP SE

SAP Notes (Authorization required)

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

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

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

SAP Note [2393917](#) – SAP HANA on VMware vSphere 6.5 and 6.7 in production

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

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

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 TDI Storage [Guidelines](#)

SAP HANA TDI Networking [Guidelines](#)

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

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

SAP Virtualization Wiki: wiki.scn.sap.com/wiki/display/virtualization

Document History

Version 2.0	1 March 2024	Include latest AOS, AHV and VMware ESXi supported software offerings
Version 1.5	20 June 2023	Include latest AOS, AHV and VMware ESXi supported software offerings
Version 1.0	5 December 2020	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, System x, ThinkServer are trademarks or registered trademarks of Lenovo. Microsoft and Windows are registered trademarks of Microsoft Corporation. Intel, the Intel logo, Xeon and Xeon Inside 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