Reference Architecture for SAP Applications on ThinkAgile HX Series

Last update: 17 April 2025
Version 2.2

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

Mark Kelly
Andrei Stanescu
Patrick Hartman

LENOVO PRESS

Table of Contents

1 Introduction

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

ThinkAgile HX Series provides a hyper-converged infrastructure which incorporates 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 virtual workloads, while providing flexibility to combine 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.

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 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 Integrated Systems (IS) and Certified Nodes (CN) are designed to help 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.

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

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

- Lenovo ThinkAgile HX630 V3: optimized for ROBO and SMB environments including SAP HANA
- Lenovo ThinkAgile HX650 V3: optimized for Big Data, high capacity and mission critical workloads including SAP HANA
- Lenovo ThinkAgile HX645 V3: optimized for general virtualization, VDI and ROBO
- Lenovo ThinkAgile HX665 V3: optimized for database & enterprise applications, AI/ML, data protection

Table 1: Comparison of ThinkAgile HX Series systems

Feature	**HX Series Integrated Systems (IS)**	**HX Series Certified Nodes (CN)**
Validated and integrated hardware and firmware	Yes	Yes
Pre-installed hypervisor: AHV or ESXi	Yes	Yes
Includes Nutanix licenses	Yes, Nutanix PnP Software	Optional
Lenovo Premier Support for problem reporting and resolution	Yes	Recommended but optional
Includes deployment services	Yes	Optional

For more information, specifications and supported configurations, refer to the Lenovopress product guides:

- ThinkAgile HX650 V3 2U Integrated System and Certified Node
- ThinkAgile HX630 V3 2U Integrated System and Certified Node

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 for SAP HANA support the following hypervisors:

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

2.2.2 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 datacentre workflows.

Nutanix Prism is a part of the Nutanix software preloaded on the HX Integrated Systems 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
    - Analyses trending patterns for accurate capacity planning

2.2.3 Nutanix Foundation

Nutanix Foundation is a separate utility used 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.

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

2.2.4 Nutanix Controller VM (CVM)

The Nutanix Controller VM (CVM) is key to hyper-converged capability and each node in a cluster has its own instance.

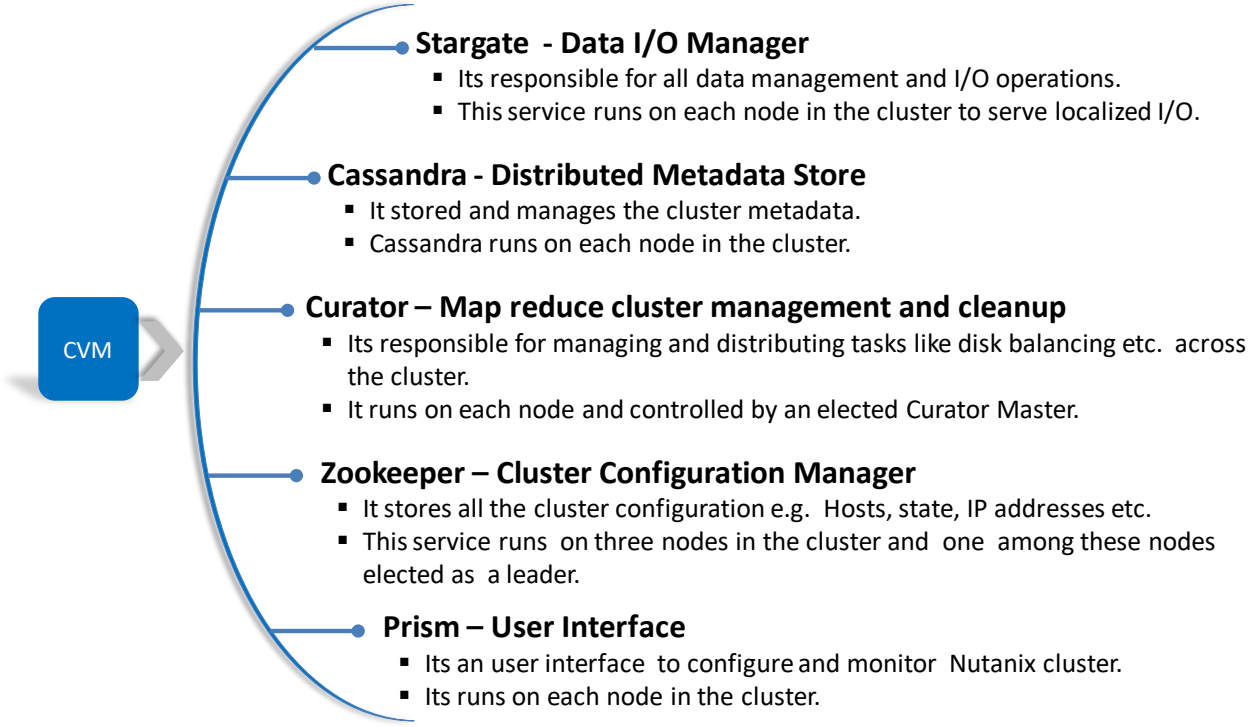

Figure 1 CVM components

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

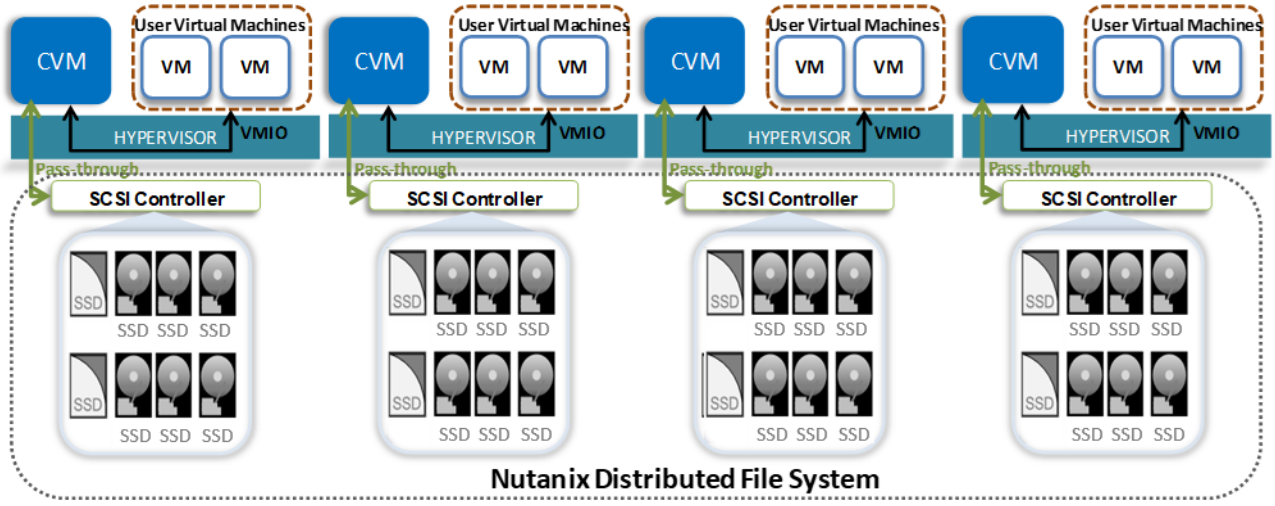

Figure 2: CVM interaction with Hypervisor and User VMs

The CVM virtualizes all the local storage attached to each node in a cluster and presents it as a centralized storage array using Nutanix Distributed File System (NDFS). All I/O operations are handled locally to provide the highest performance. See section 2.4 for more details on the performance features of NDFS.

2.3 Data network components

The data network is the fabric that carries all inter-node storage I/O traffic for the shared Lenovo HX distributed file system, in addition to the user data traffic via the vNICs exposed through the hypervisor to the virtual machines.

Each HX Series node contains at least two dual-port 10GbE network adapters as well as 4 on-board 1GbE ports. The hypervisors are configured by the Nutanix software so that the fastest network ports are pooled for the data network. The hypervisor VM management network should use the same network. Because all of the network ports are pooled, each node only needs two network IP addresses; one for the hypervisor and one for the Nutanix CVM. These IP addresses should all be on the same subnet.

All storage I/O for virtual machines (VMs) running on an HX Series node is handled by the hypervisor on a dedicated private network. The I/O request is handled by the hypervisor which then forwards the request to the private IP on the local CVM. The CVM then performs the remote data replication with other nodes in the cluster using its external IP address. In most cases, read request traffic is served locally and does not enter the data network. This means that the only traffic in the public data network is remote replication traffic and VM network I/O (i.e. user data). In some cases, the CVM will forward requests to other CVMs in the cluster, such as when a CVM is down or data is remote. Also, cluster-wide tasks such as disk balancing temporarily generate I/O traffic on the data network.

For more information on the network architecture see nutanixbible.com.

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 thoroughly tested firmware . 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.

ThinkAgile HX Series nodes have redundant hardware components including dual power supplies, multiple chassis fans, two Intel CPUs, multiple memory DIMMs, multiple SSDs, and dual-port network interface cards.

Hardware performance

The ThinkAgile HX Series nodes have been carefully designed for performance. In addition to all 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 SSD storage on all nodes and presents it as a single storage entity to the cluster.

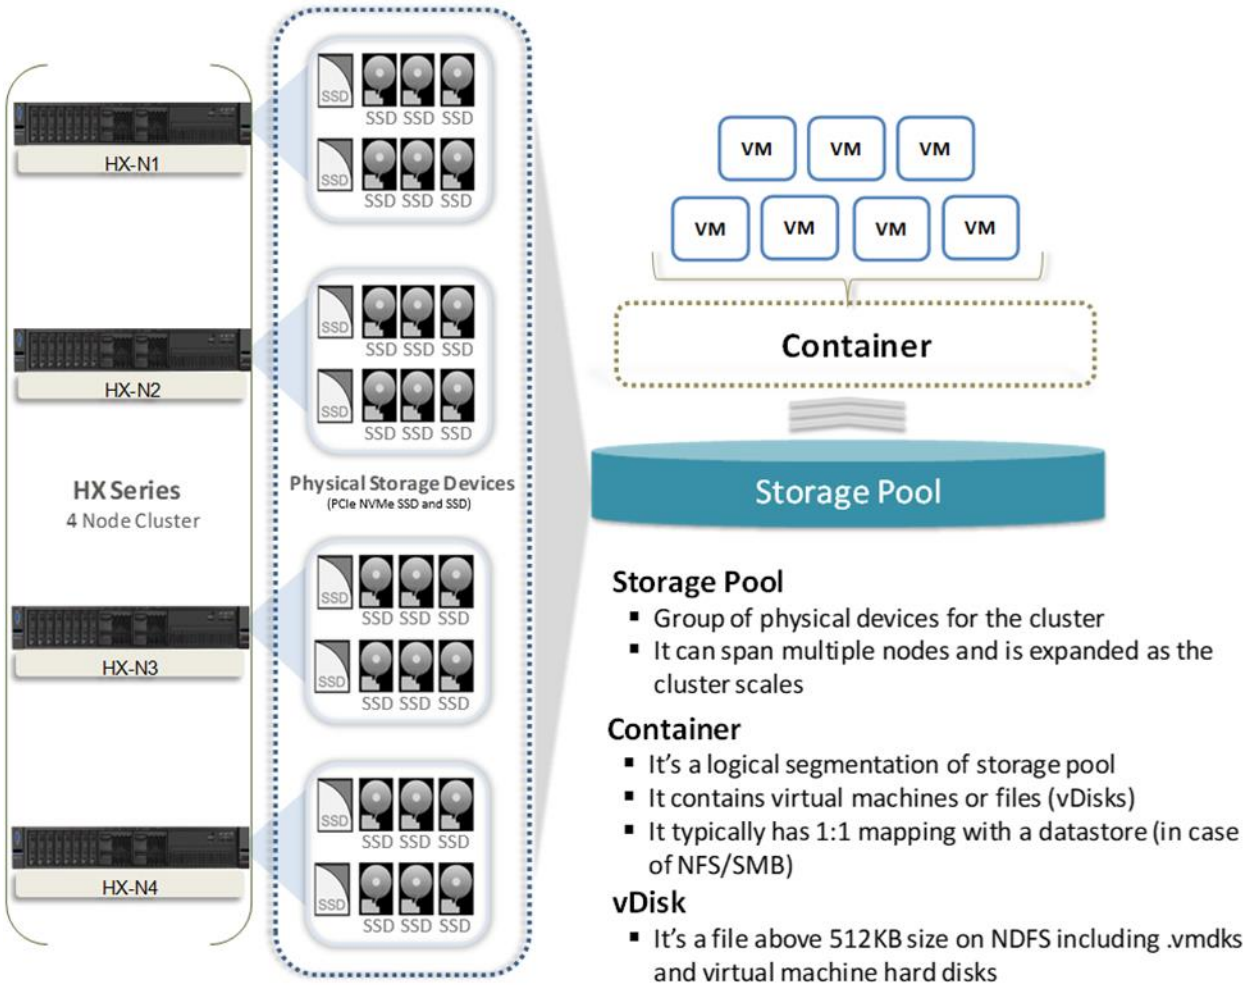

Figure 3: Nutanix Distributed File System

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 is 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 Disk Devices

Performance devices are the highest performance components in the node. These can be NVMe or a mix of NVMe and SSD devices. They store the following items:

- Nutanix Home (CVM core)
- Metadata (Cassandra / AES storage)
- OpLog (persistent write buffer)
- Extent Store (persistent storage)

The following figure shows an example of the storage breakdown for a Nutanix node's performance device:

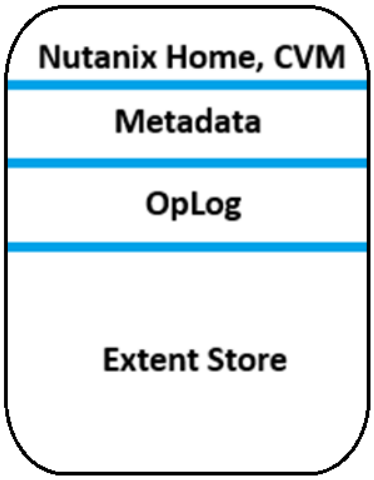

Figure 4: Performance Disk Devices - ***see the Nutanix Bible*** for more details

Performance with data tiering

Nutanix uses a disk tiering concept in which disk resources 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:

- SSDs on a local node always have the highest tier priority for write I/O.
- If the SSDs on the local node SSD are full then other SSDs in the cluster are used for I/O.
- The NDFS Information Lifecycle Management (ILM) component migrates cold data away from local SSDs to free up SSD space. Conversely it also moves heavily accessed data to local SSDs 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 a 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

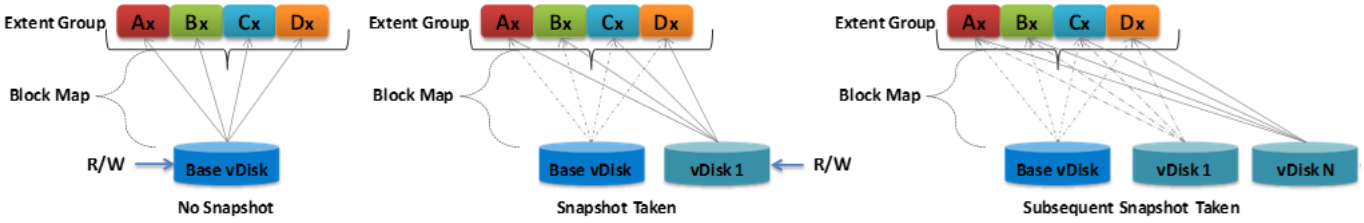

Figure 6: VM snapshots

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.

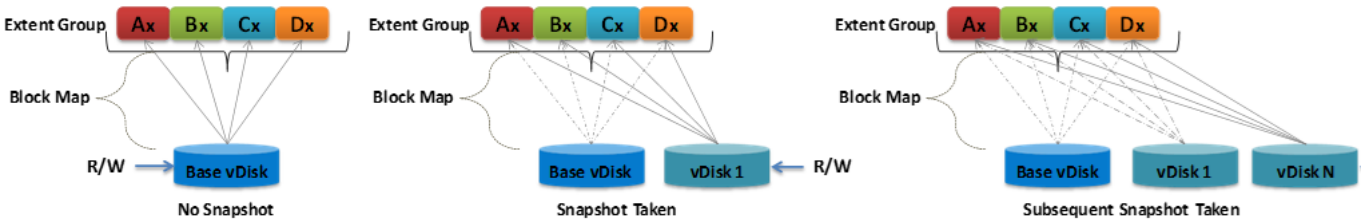

Figure 6: VM snapshots

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.

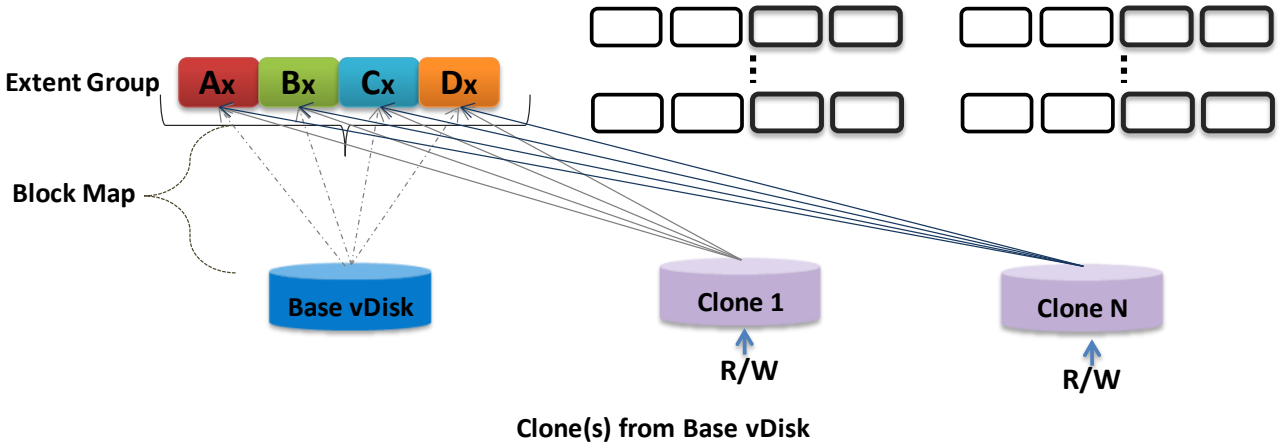

Figure 7: VM clones

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 the following actions to increase storage efficiency:

- Sequential streams of data fingerprinted at 4K granularity
- A 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 the 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 it as offline and in need of replacement. The idea is to avoid expensive data transfers, to maintain data redundancy and avoid possible loss of data.

3 SAP Application Workloads

SAP 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

SAP Applications running on ThinkAgile HX 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

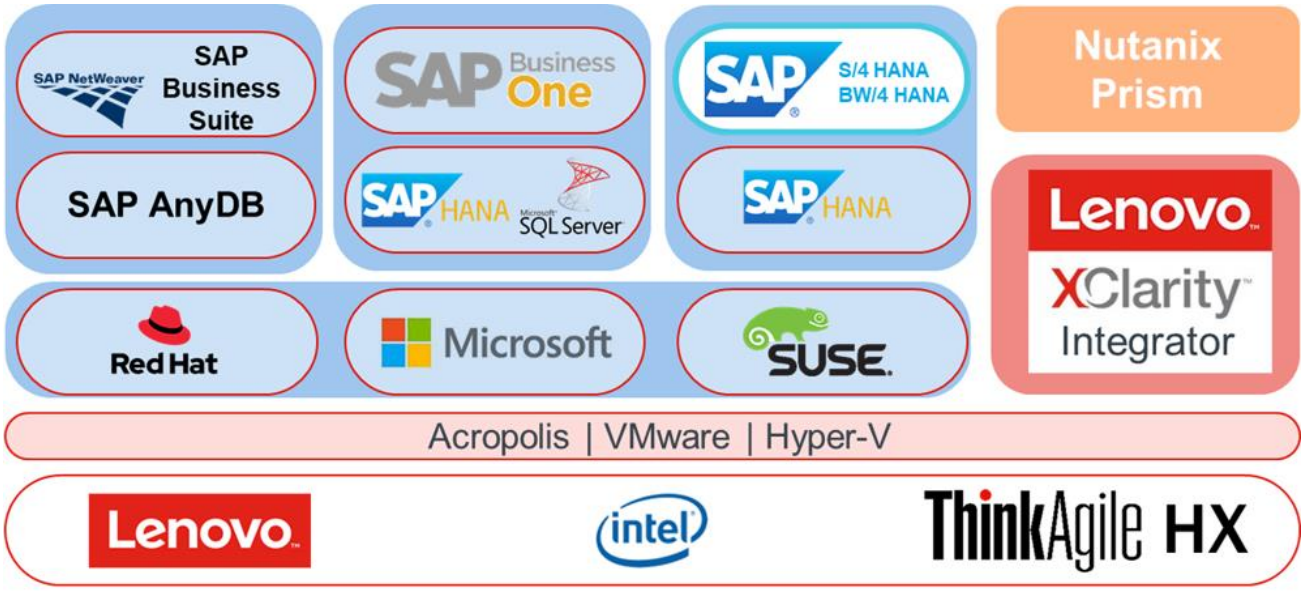

Figure 8: SAP Applications on ThinkAgile HX

3.2.1 Virtualization of SAP landscapes

Traditionally the SAP application tier is executed on physical servers which can be wasteful of resources whereas virtualization can be applied to consolidate servers and reduce CAPEX and OPEX. Virtualization of SAP application tiers provides a number of benefits as standard:

- 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. The Nutanix Distributed file system (NDFS) controlled by the CVM requires only up to four physical processor cores and up to 64GB of physical memory per node. This can be configured in Nutanix Prism.

3.3 SAP HANA deployment

SAP HANA is certified to run on Lenovo ThinkAgile HX Series systems which support the following hypervisors:

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

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

- SAP Note 2686722 - SAP HANA virtualized on Nutanix Acropolis Hypervisor

In order to deploy an instance of SAP HANA the hypervisor, guest operating system, file system and SAP HANA database requirements must be sized.

Table 2 is derived directly from the SAP Note 2686722 - SAP HANA virtualized on Nutanix Acropolis Hypervisor and the Nutanix Best Practices Guide for SAP HANA.

Table 2: SAP HANA on Nutanix Acropolis OS (with ThinkAgile HX630 V3, HX650 V3)

Description	Minimum	Maximum
Virtual Memory (vMEM) allowed	128 GiB	2048 GiB
Virtual CPUs (vCPUs) allowed	8 (16 with hyper-threading)	Intel Sapphire Rapids - 60 (120 with hyper-threading) Intel Emerald Rapids - 64 (128 with hyper-threading)
Sockets / Productive SAP HANA use	1	1.5 (using SNC-2)
VMs / Productive SAP HANA use	1	3 (using SNC-2)
Number of ThinkAgile HX Nodes	3 (4 recommended)	64 with AHV 48 with vSphere ESXi

3.3.1 Intel Sub-NUMA Clustering (SNC-2)

Starting with AOS version 6.8 and Intel Sapphire Rapids (Gen 4) and Emerald Rapids (Gen 5) systems it is supported to configure “half-socket” VMs for SAP HANA using SNC-2 technology. The Intel feature Sub-NUMA Clustering (SNC-2) splits the last level cache (LLC) into two localization domains based on address range, while binding each cluster to a subset of the memory attached to it via the controllers in the system. SNC-2 improves average latency by ensuring the data addresses are mapped into the local memory controller and the local LLC slices. It has the added advantage that for SAP HANA, a sub-socket VM can be configured both for non-productive and productive instances. This increases the density of SAP HANA instances on smaller nodes.
Note that only full-socket VMs or half-socket VMs are supported. 1.5 socket VMs are not supported.

It is also important that VMs on SNC-2 enabled hosts should not be combined with VMs on SNC-2 disabled hosts due to risks associated to of HA takeover or migration (see SAP Note 2686722).

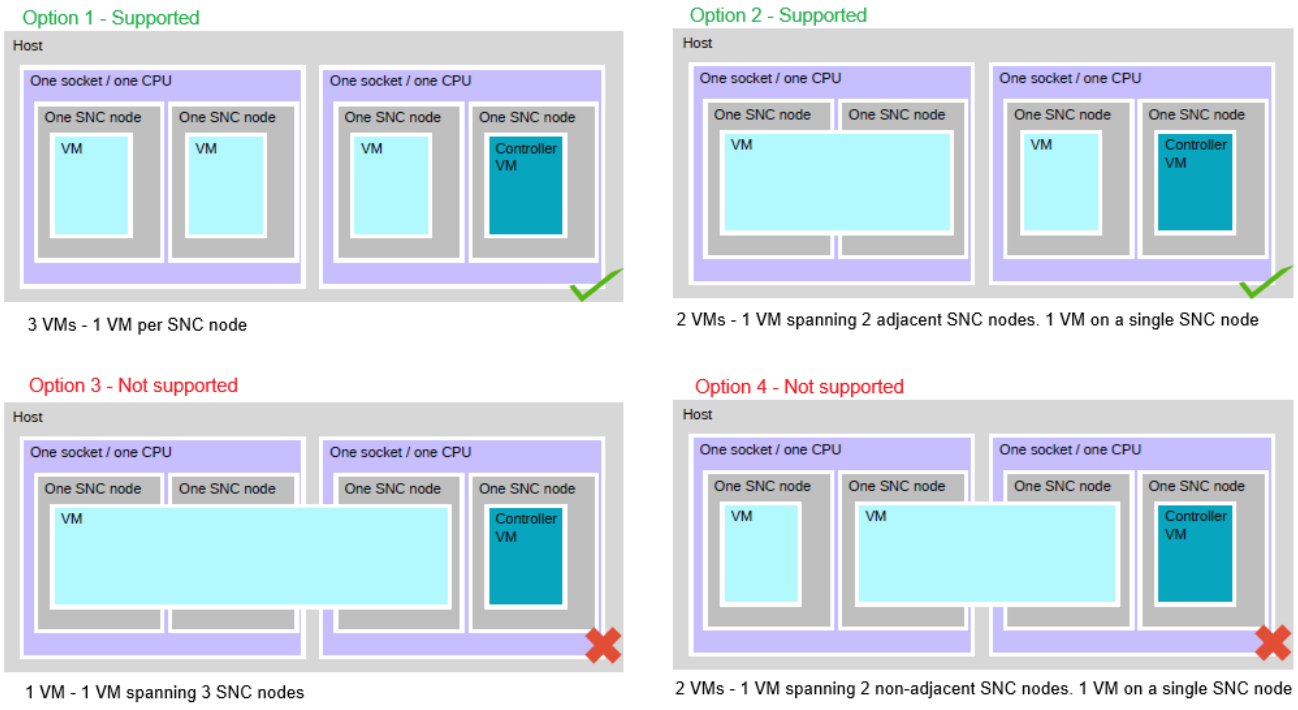

Figure 9:Supported and unsupported SNC-2 configurations *(taken from Note 2686722)*

3.3.2 Nutanix Cluster Guidance with ThinkAgile HX for SAP HANA

- The minimum requirement is for two ThinkAgile HX Integrated Systems or Certified Nodes which have been certified for SAP HANA. The 2nd node will be required for High Availability failover for the primary SAP HANA production instance.
- One (minimum) or two (recommended) additional ThinkAgile HX630 V3 or HX650 V3 nodes to complete the cluster. These may be used for non-production SAP HANA or for other application workloads
- All nodes in a Nutanix cluster must contain a similar storage layer
    - The physical number of disks must be equivalent in each node
    - The first 2 nodes much be configured with identical CPU and memory. Subsequent nodes from node 3 onwards may be different HX models or the same HX model configured with different CPU & RAM.
- For production SAP HANA workloads
    - Full-socket VMs can occupy a maximum of one socket on HX630 V3 and HX650 V3 nodes
    - Half-socket VMs are supported on HX630 V3 and HX650 V3 nodes with AOS 6.8 and later with Sub-NUMA Clustering (SNC-2) enabled. See the Intel Technical Overview Of The 4th Gen Intel® Xeon® Scalable processor family for details of SNC-2 feature
    - Half-socket VM resources using SNC-2 are fixed at 50% of the cores on the socket and 50% of the memory on the DDR channels associated to the socket
    - When SNC-2 is enabled for SAP HANA VMs on HX630 V3 and HX650 V3 systems, then Lenovo advises that all hosts in the cluster are SNC-2 enabled. It is not supported to migrate SNC-2 enabled VMs to hosts which have SNC-2 disabled
    - A half-socket VM is supported on the socket hosting the Nutanix Controller VM (CVM) with 50% of the socket resources
    - 1.5 socket VMs are not supported

- For non-production SAP HANA workloads
    - The same configurations for productive SAP HANA workloads are valid
    - Additionally, any available resources on the socket hosting the CVM may be assigned to non-productive workloads
- For all other workloads
    - Any other workload can run on any available socket, including the socket hosting the CVM
- Turn off any Nutanix Distributed file system (NDFS) features that impact storage performance such as: compression, deduplication or erasure coding (EC-X)

3.3.3 SAP HANA configuration guidance

SAP Note 2686722 provides some basic rules to set up a ThinkAgile HX node for use with SAP HANA.

Processor Rules:

- A quantity of two Intel Xeon Scalable Processors Gen 4 (Sapphire Rapids) or Gen 5 (Emerald Rapids) are required
- Any Silver, Gold or Platinum processor with a minimum of at least eight cores which is configurable in the Lenovo DCSC configurator is supported.

Memory Rules:

- For Intel Xeon Scalable Processors, either one or two DIMM channels per socket may be populated for a total of 16 or 32 physical DIMMS populated in the system
- SAP HANA requires having a homogeneous placement of memory across each populated DIMM channel
- Mixed memory sizes are not permitted on ThinkAgile HX systems

Networking Rules:

- At least 2 pairs of dual-port 10GbE NICs are necessary. 25GbE is recommended. 1 pair will be used for SAP HANA traffic and 1 pair for Nutanix Files

- Two Mellanox ConnectX-6 Lx 2x25GbE SFP28 adapters are recommended.
  Additional adapters may be added up to the maximum supported by the ThinkAgile HX system
  Alternatively, any other ethernet adapter selectable in the DCSC configurator is valid.
  Alternative adapters supporting RoCE include Broadcom 57414, 57504, 57416, 57454, Mellanox CX-6 Dx 100GbE
- If NVMe storage devices are used, then RDMA enabled adapters must be used, like the Mellanox ConnectX-6 series.

GPU Rules:

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

Storage Rules:

- The minimum number of disks is dependent on the storage layout:
  - Systems configured with SATA/SAS SSDs require a minimum of four devices per node on smaller installations. The Nutanix recommendation is for 8 drives or more to ensure optimal performance
  - Mixed NVMe and SAS/SATA drives:
    HX650 V3: the number of NVMe drives configurable in a mixed environment is 4. The minimum number of SAS/SATA SSDs is 4.
    HX630 V3: the number of NVMe drives can be between 2-4. The number of

SAS/SATA SSDs can be between 2-8. 4 drives in total is the minimum while 8 is recommended for optimal performance.

- NVMe SSD drives only: With NVMe-only configurations the minimum supported quantity of drives for smaller environments is 4. The recommended minimum for optimal performance is 8 NVMe drives.

- Any NVMe or SAS/SATA drives configurable in the DCSC configurator are supported, including Read Intensive drives.

Nutanix recommends using a NVMe based approach for SAP HANA databases which are highly transactional. If applications are not highly transactional a configuration with only SAS/SATA SSD devices is acceptable.

Software Rules:

- Nutanix Acropolis OS (AOS) 6.5.4 LTS, 6.8 LTS or 6.10 LTS should be used with SAP HANA applications and either:
  - Acropolis Hypervisor (AHV) 20230302.100173 or 20230302.101026 major versions
  - VMware vSphere ESXi 8.0 U1 hypervisor
- All system management software options are allowed
- Lenovo Professional Services is recommended for installation and deployment of SAP HANA

3.3.4 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 1TB of memory
- S/4HANA System Recover (HA) instance with 1TB of memory
- S/4HANA QA (QAS) instance with 1TB of memory
- Four SAP HANA (Dev) development instances with 256GB of memory each
- BW/4HANA Production (PRD) instance with 500GB of memory
- BW/4HANA System Recovery (HA) instance with 500GB of memory
- BW/4HANA QA (QAS) instance with 500GB of memory

The customer wishes to purchase as few servers as possible and place their S/4HANA and BW/4HANA instances on two socket ThinkAgile HX systems. Following the above requirements it can be shown (see Figure 10) that a minimum of four servers need to be configured.

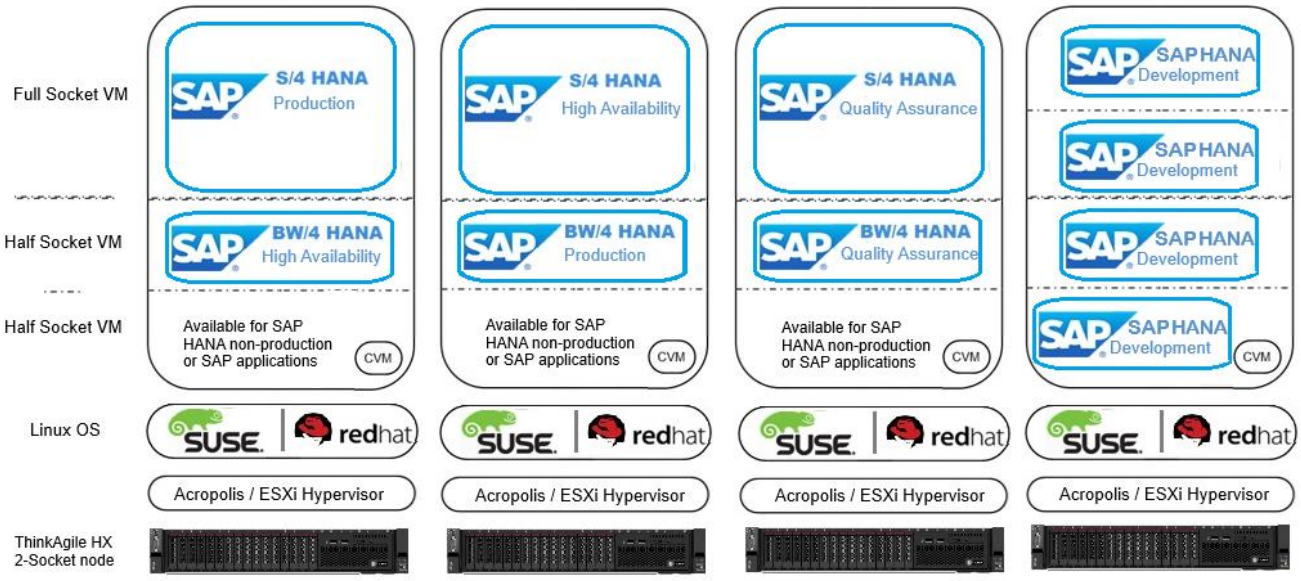

Figure 9: SAP HANA Sizing example showing placement of Production and Non-production workloads

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 a homogeneous placement of memory across all populated DIMM channels

For 4th Generation Intel Xeon Scalable processors, there can be either one or two DIMMs per channel (8 DIMMs per socket channel) for a total of 16 or 32 DIMMs on a 2-socket server.

Table 3: SAP HANA virtual machine memory sizing provides the recommended system memory depending on how one or more SAP HANA VMS are distributed across each of the processors.

Table 4: SAP HANA virtual machine memory sizing

1 Processor VM memory size	Total System Memory for ThinkAgile HX630 V3 / HX650 V3
128GB	256 GB (16 x 16GB)
256GB	512GB (16 x 32GB)
384GB	768GB (16 x 48GB)
512GB	1024GB (16x32 or 16x64GB
768GB	1536GB (16x96GB)
1024GB	2048GB (32 x 64GB)
1536GB	3072GB (32 x 96GB)
2048GB	4096GB (32 x 128GB)

Processor sizing

SAP recommends a minimum of eight cores per socket.

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 SAPS are not available, it is also possible to use 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 recommendations of 50GB and 2GB respectively but has extra capacity for temporary file growth and certain root file system features such as btrfs snapshots.

In general, any solid-state device may be used which is supported by the ThinkAgile HX SAP HANA models. 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.

The amount of storage required is calculated as a ratio of memory.
It is recommended to size all SAP HANA VMs first using SAP HANA TDI Storage Guidelines (rule of thumb is roughly 3x memory, but it can be less or more).

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 Capacity Calculator and follow the below rules:

- Storage efficiency set to none (1:1)
- Storage deduplication disabled (ECX)
- Set failover plan to N+1
- Only SSD are allowed for SAP HANA
- Spinning disks (HDD) are not allowed

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.

Network sizing

For Hyper-converged SAP HANA environments the use of network cards and switches that support Remote Direct Memory Access (RDMA) technology is recommended when using NVMe devices. Any switch described in “Data network components” on page 5 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 and RoCE. In this scenario any standard network card and switch may be used with a speed of 10GbE or higher.

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 for your chosen Linux distribution chosen.

4.3 Guest File system

Supported file systems are documented in 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 Best Practices Guide for SAP HANA

Nutanix Storage Capacity Calculator

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 3372365 - SAP HANA on VMware vSphere 8

SAP Note 2161991 – VMware vSphere configuration guidelines

SAP Note 2000003 – FAQ: SAP HANA

SAP Note 1552925 – Linux: High Availability Cluster Solutions

SAP Note 1612283 – Hardware Configuration Standards and Guidance

SAP Customer Network (SCN)

SAP HANA TDI FAQ

SAP HANA TDI Storage Guidelines

Community Blog: SAP HANA on VMware vSphere

Document History

Version 2.2	11 April 2025	Include ThinkAgile Generation 5 (Emerald Rapids) systems
Version 2.1	20 December 2024	Include latest AOS, AHV and ESXi supported software. Include Sub-Numa Clustering support
Version 2.0	1 March 2024	Include ThinkAgile Generation 4 (Sapphire Rapids) systems. 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