



ThinkAgile HX Solution for AI with Nutanix Enterprise AI

Solution Brief

AI adoption is rapidly accelerating across enterprises, creating a demand for scalable, GPU-optimized infrastructure that can support increasingly complex and data-intensive workloads. Organizations need datacenter solutions capable of running generative AI, machine learning, virtualized environments, and general inference workloads while maintaining operational simplicity, efficiency, and predictable performance.

Lenovo ThinkAgile HX accelerates time to outcome by delivering a fully engineered, factory-preloaded AI-ready solution that eliminates integration delays and operational guesswork, creating a faster time to production for AI workloads. With validated Best Recipes and automated lifecycle management, customers move from infrastructure to inference faster, with predictable performance and zero-error updates that keep AI pipelines continuously productive.

The [Lenovo ThinkAgile HX solution for AI](#) combines advanced compute, storage, and software components into a unified, hyperconverged architecture in a single, easy to stand up solution within the ThinkAgile HX650a V4. Lenovo and NVIDIA previously worked together to create an optimized hardware stack designed to run single node inference workloads called the [Hybrid AI 221 platform](#), by leveraging this architecture and combining it with ThinkAgile HX validation and support, Lenovo is able to deliver a validated solution for AI inference workloads on the Nutanix AI stack. The solution can be configured [here](#).

The [Lenovo ThinkAgile HX solution for AI](#) is powered by Intel Xeon processors, high-bandwidth NVMe storage, and NVIDIA GPUs including the RTX PRO 6000 Blackwell Server Edition or H200 NVL. It provides a foundation for AI workloads, including inference and agentic AI workloads. Integrated with NAI, NKP and NVIDIA AI Enterprise, the platform offers a secure, production-ready path to enterprise AI deployment. Nutanix Enterprise AI acts as the AI-ready software for your AI Factory, providing a unified, high-performance environment for model deployment, inference services, and end-to-end lifecycle management.

By combining Lenovo’s infrastructure, Nutanix’s AI software, and NVIDIA’s acceleration expertise, Lenovo enables IT teams to deploy, manage, and scale AI workloads confidently, making advanced generative AI and inferencing capabilities accessible to existing staff. These integrated technologies form a comprehensive, enterprise-ready solution that accelerates AI innovation while simplifying operations across hybrid and multi-cloud ecosystems.

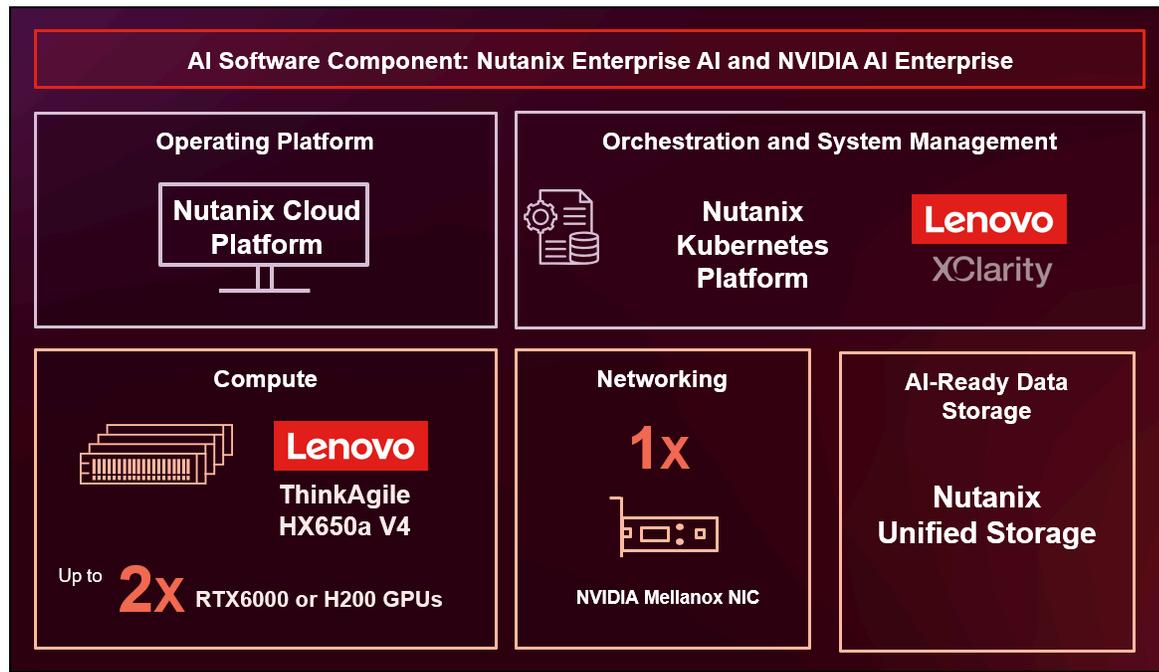


Figure 1. ThinkAgile HX solution for AI overview

Components

The main components of the solution include the AI compute node, GPUs, networking, and the software stack. The software stack consists of the operating platform Nutanix Cloud Platform, AI software components Nutanix Enterprise AI and NVIDIA AI Enterprise, orchestration and system management through Lenovo XClarity and Nutanix Kubernetes Platform, and AI-ready data storage provided by Nutanix Unified Storage

- [AI Compute Node](#)
- [GPU Selection](#)
- [AI Software Stack](#)
- [NVIDIA AI Enterprise](#)

AI Compute Node

Lenovo ThinkAgile HX is a fully engineered, appliance-grade hyperconverged solution designed to simplify infrastructure operations from day-0 deployment through day-N lifecycle management. Built as a complete, pre-validated system, ThinkAgile HX arrives factory preloaded and ready to run, removing integration complexity and accelerating time to value. At the core of the solution are ThinkAgile Best Recipes which are Lenovo's validated configuration blueprints that unify hardware, firmware, drivers, and software into a single, tested package. These Best Recipes eliminate guesswork by defining the correct update sequencing, enabling faster, zero-error updates with predictable outcomes. The result is simplified lifecycle management, consistent compliance, and predictable performance, allowing IT teams to operate with confidence while reducing operational overhead and risk.

The Lenovo ThinkAgile HX650a V4 is a 2-socket 2U hyperconverged solution for customers who want to maximize GPU compute power while deploying a fully integrated platform with Nutanix HCI software in a traditional 2U rack form factor. Built on the ThinkSystem SR650a V4 platform, the HX650a V4 supports two Intel Xeon 6700-series or 6500-series processors, along with support for two NVIDIA H200 NVL, or two RTX PRO 6000 Blackwell Server Edition GPUs in the 221 platform configuration. The HX650a V4 is designed for high-density, scale-out hyperconverged workloads, where capacity and performance scale horizontally by adding additional nodes to the Nutanix cluster.

The Lenovo ThinkAgile HX650a V4 is engineered to accelerate GPU-intensive workloads such as AI inferencing, Retrieval Augmented Generation (RAG), machine learning etc. Powered by Intel Xeon 6 processors, it delivers high GPU density and advanced internal NVMe storage for optimal performance in HCI environments. The platform supports up to 2 double-width or 8 single-width GPUs and up to 8 NVMe drives, enabling fast, low-latency data access for both compute and storage services.

A typical AI-enabled HX650a V4 node can be configured with two Intel Xeon 6530P 32-core 225W 2.3 GHz processors. With 8 memory channels per processor socket, the system delivers exceptional memory bandwidth to support GPU acceleration and virtualization overhead. A total of 512 GB of system memory is allocated to meet the requirements of the GPUs, the Nutanix OS, the hypervisor layer, and customer applications. The system can be configured with a ThinkSystem Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet adapter for or ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port PCIe Ethernet adapter for networking. The 221 platform is designed to use internal NVMe storage by default for hyperconverged solutions.

In the 221 platform configuration of the ThinkAgile HX650a V4, GPUs are installed at the front of the server on Riser 7 as full-height, full-length (FHFL) GPUs.

GPU Selection

The Hybrid AI 221 platform on the ThinkAgile HX650a V4 is engineered to support the full range of NVIDIA double-width PCIe GPUs, including the NVIDIA RTX PRO 6000 Blackwell Server Edition and NVIDIA H200 NVL, enabling flexibility across diverse AI and accelerated compute workloads.

- **NVIDIA RTX PRO 6000 Blackwell Server Edition**

Powered by the next-generation NVIDIA Blackwell architecture, the RTX PRO 6000 Blackwell Server Edition combines advanced AI acceleration with high-end visual computing capabilities for modern data center environments. Featuring 96 GB of high-speed GDDR7 memory, it delivers exceptional performance for a wide variety of workloads, including agentic AI, physical AI, scientific simulation, rendering, 3D visualization, and video processing. This GPU offers strong versatility across both AI and graphics-intensive use cases. When deployed in platform 221, a minimum of 290 GB of system memory is recommended to ensure optimal performance.

- **NVIDIA H200 NVL**

The NVIDIA H200 NVL is purpose-built for demanding generative AI and high-performance computing (HPC) applications. It features an impressive 141 GB of HBM3e memory, nearly doubling the capacity of the H100, along with up to 4.8 TB/s of memory bandwidth. This combination allows the H200 NVL to efficiently process large and complex AI models, including large language models (LLMs), delivering significant performance gains. Designed with power efficiency in mind, it provides higher performance within a similar power envelope as the previous generation. NVIDIA also includes a 5-year NVIDIA AI Enterprise license at no additional cost with H200 NVL. For the 221 platform, a minimum system memory of 453 GB is recommended.

AI Software Stack

The AI Software Stack provides the platform for deploying and managing AI workloads across the infrastructure, with Nutanix Cloud Platform (NCP) forming the core operating platform. It combines Nutanix Enterprise AI (NAI) and Nutanix Kubernetes Platform (NKP) for AI model deployment and orchestration, with NVIDIA AI Enterprise delivering optimized AI frameworks and microservices. Nutanix Unified Storage (NUS) provides high-performance data services for AI workloads, while Lenovo XClarity enables system monitoring and hardware management.

- [Enterprise AI Software](#)
- [Kubernetes Layer](#)
- [AI Ready Data Storage](#)
- [Lenovo XClarity](#)

Enterprise AI Software

Nutanix Enterprise AI (NAI) is a Kubernetes-based application that forms the AI platform layer of the Enterprise AI stack, giving IT teams the ability to deploy, manage, and monitor large language models (LLMs) and inference endpoints. Leveraging the capabilities of NKP, NAI provides the higher-level services required to operationalize generative AI, acting as a centralized inferencing control plane. With support for endpoint APIs from leading LLM providers, including NVIDIA NIM and Hugging Face, organizations can securely run a wide range of generative AI models on-premises or in the public cloud.

NAI includes a streamlined, UI-driven interface, role-based access controls (RBAC), and untethered deployment options for dark-site or air-gapped environments, simplifying Day 2 operations, monitoring, and adaptation of AI models with enterprise-grade resilience and compliance. This approach allows teams to quickly deploy, monitor, and manage AI models and secure endpoints, providing flexibility in model selection and making AI tools accessible across the enterprise, empowering every team to leverage AI effectively.

NAI Unified Endpoints provides consistent authentication, rate-limiting, and monitoring, with fallback and load balancing for inference endpoints, across self-run local models as well as provider models like Anthropic or OpenAI.

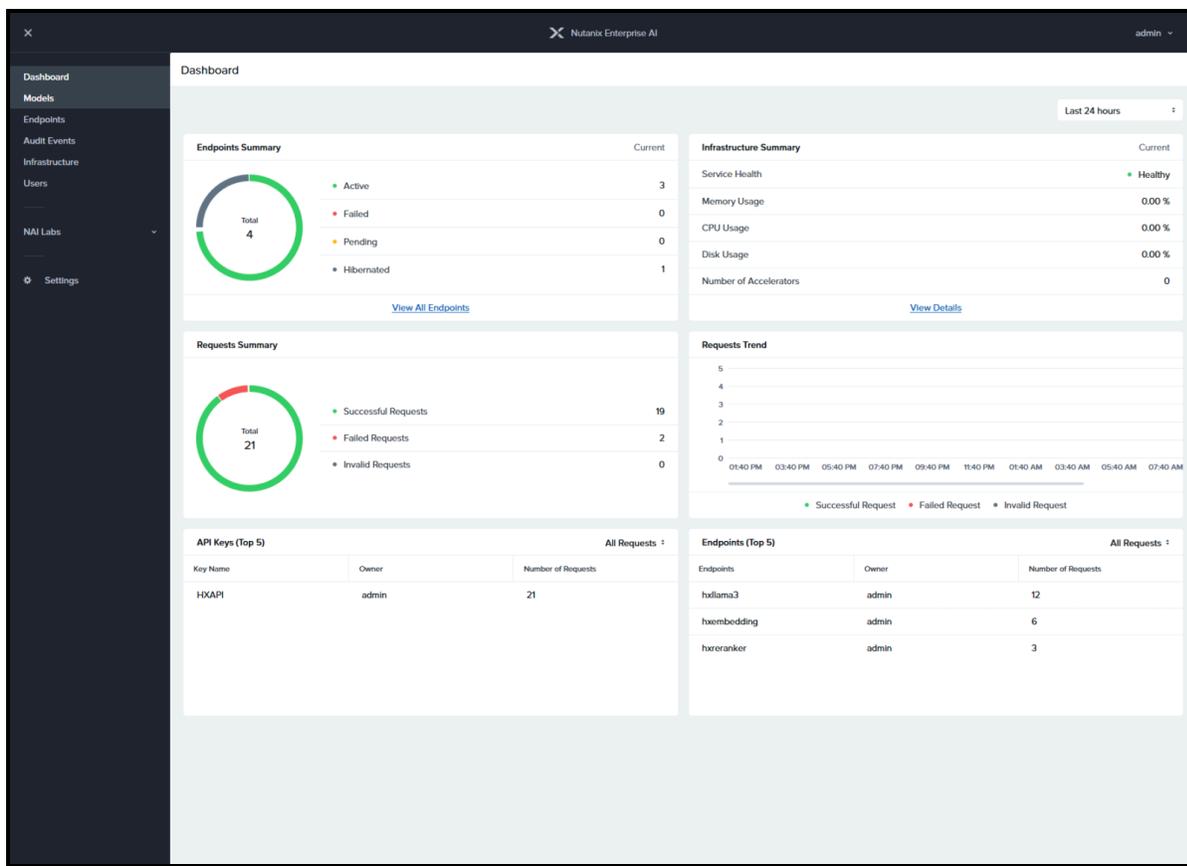


Figure 2. Nutanix Enterprise AI

Kubernetes Layer

Nutanix Kubernetes Platform (NKP) provides the enterprise-grade Kubernetes foundation for Nutanix Enterprise AI on ThinkAgile HX Series, enabling organizations to deploy and operate AI workloads with consistency, security, and scale. As the Kubernetes layer, NKP delivers a CNCF-conformant platform that simplifies cluster lifecycle management and supports production-ready deployment of AI services such as large language models (LLMs), inference APIs, and retrieval-augmented generation (RAG) pipelines.

Tight integration with software-defined compute, storage, and networking ensures efficient resource utilization, high availability, and support for GPU-accelerated workloads, making it well suited for data-intensive AI training and low-latency inference in on-premises and hybrid environments.

Together, ThinkAgile HX Series, NKP and NAI enable a secure, governed, and cloud-consistent AI platform for the enterprise. Built-in security, multi-tenancy, and policy controls allow teams to safely scale AI across business units, while Kubernetes-native portability ensures AI applications can be deployed consistently across on-prem and hybrid cloud environments. This combination provides a trusted foundation for operationalizing enterprise AI at scale.

AI Ready Data Storage

Nutanix Unified Storage (NUS) is a software-defined storage platform that brings file and object storage together into a single, integrated solution. By unifying these storage services on one platform, NUS simplifies operations while delivering high performance, dense capacity, and cost efficiency. This makes it well suited for data-heavy workloads, including AI/ML pipelines and big data analytics, where fast and scalable access to data is critical. Part of the NVIDIA STX, a specialized modular reference architecture for AI-native storage, Nutanix Unified Storage enables continuous, GPU-accelerated data transformation and vectorization directly within the storage cluster. This provides the high-performance data fabric required to

bridge training and inference workloads without bottlenecks.

NUS works seamlessly with the Nutanix Cloud Platform (NCP) to deliver a consistent, high-performance storage foundation across hybrid and multi-cloud environments. By consolidating management into a single interface, NCP enables organizations to optimize performance, reduce costs, and strengthen security, while simplifying operations and allowing teams to focus on innovation. In addition, NUS is fully compatible with NKP and NAI. This integration streamlines the deployment and management of Kubernetes clusters, making it easier to run and scale AI/ML workloads. Together, NUS and Nutanix's AI and Kubernetes tools provide a unified infrastructure solution for modern, data-intensive applications.

Nutanix Unified Storage delivers several advantages for AI and machine learning use cases:

- **High Performance and Scalability** - NUS is designed to meet the intensive I/O demands of AI/ML workloads, providing high-speed read/write performance using GPU-Direct protocols across thousands of GPU clients, ensuring that data availability scales just as fast as your compute, and GPUs are never starved for data.
- **Flexible Deployment and Integration** - The platform supports seamless integration with on-premises systems and public cloud environments, enabling smooth data movement and management across hybrid and multi-cloud architectures.
- **Cost Efficiency** - By consolidating multiple storage services into a single platform and providing a high-capacity tier for KV cache offloading that frees up critical GPU memory, NUS enables the systems to handle larger context windows and more users. as well as reduces operational complexity. This directly lowers overall storage costs, making it economical for managing large-scale datasets.
- **End-to-End AI Enablement** - From data ingestion and preparation to model training and inference, NUS is built to support the full AI/ML lifecycle with scalable and reliable storage services.

Lenovo XClarity

Lenovo XClarity Management on ThinkAgile HX Series provides comprehensive platform-level monitoring and alerting across all hardware subsystems, ensuring issues are detected and addressed before they impact workloads. Hardware events, errors, and health status are reported directly to Lenovo XClarity, giving administrators full visibility beyond the operating system layer. XClarity continues to monitor the underlying hardware independently of the OS, enabling proactive issue detection even during system outages.

Lenovo-specific alerts leverage Lenovo intellectual property and firmware capabilities to help prevent unplanned downtime. This includes live PCIe error recovery, proactive alerts for memory failures and predicted memory faults, and drive health and RAID-related alerts reported through the Lenovo XClarity Controller (XCC).

NVIDIA AI Enterprise

To deliver enterprise-grade AI on ThinkAgile HX Series, Nutanix Enterprise AI (NAI) is deeply integrated with NVIDIA AI Enterprise software and AI microservices, bringing production-ready model deployment, inference performance, and secure AI operations to your hybrid infrastructure.

- **Enterprise-Ready AI Software Stack**
NVIDIA AI Enterprise provides a validated, optimized software layer for AI workloads. When combined with NAI on ThinkAgile HX Series, this enables customers to deploy and manage AI models and inference endpoints with confidence, covering use cases from simple model serving to complex agentic workflows.
- **Pre-Integrated AI Microservices (NVIDIA NIM & NeMo)**
NAI on ThinkAgile HX Series provides direct access to NVIDIA NIM microservices and NeMo model services to simplify AI delivery. These pre-validated components accelerate GenAI deployment, reduce operational overhead, and ensure consistent performance across environments, whether at the edge, in your datacenter, or in cloud extensions.
- **Accelerated Performance & Scalability**
NVIDIA's software stack is engineered to take full advantage of GPU capabilities for inference and reasoning workloads. On ThinkAgile HX Series with NVIDIA GPUs, this means predictable performance, efficient resource utilization, and lower latency for AI applications compared to general-purpose deployments.
- **Unified & Secure AI Operations**
Combined with NAI's centralized RBAC-based user-level access to models and secure endpoint abstraction, enterprises gain unified control over large language model (LLM) endpoints, secure APIs, and governance policies. This integration supports robust Day-2 operations, essential for scaling AI across business units.
- **Simplified Enterprise AI Adoption**
By embedding NVIDIA AI Enterprise and Nutanix Enterprise AI on ThinkAgile HX Series, organizations can reduce time-to-value for AI initiatives. This combination gives IT teams a streamlined path from infrastructure provisioning to model deployment and operationalization with familiar tools and workflows.

RAG Use Case: AI-Powered Manufacturing Operations Assistant

Open-source large language models such as Meta's Llama are trained on vast amounts of public internet data. While they understand general engineering concepts, safety standards, and industrial terminology, they do not have built-in knowledge of your organization's specific machinery, maintenance schedules, operating procedures, or plant-level documentation.

For example, if a plant technician asks:

"What is the approved maintenance interval for CNC Machine 12 in Plant B?"

A general LLM might provide manufacturer-recommended service intervals, but it won't know your company's customized maintenance plan or historical service adjustments.

To bridge that gap, organizations implement Retrieval-Augmented Generation (RAG), a framework that connects large language models with internal operational knowledge.

End-User Query Workflow (RAG in Action)

1. Ask Question

A technician, engineer, or plant manager submits a question to the chatbot interface.

Example: "What troubleshooting steps should I follow for recurring overheating on Press Line 3?"

2. Create Embedding of Query

Instead of immediately sending the question to a text generation model, the application first converts the user's query into an embedding using an embedding model hosted on Nutanix Enterprise AI.

3. Search / Retrieval of Similar Content

The generated query embedding is compared against stored document embeddings in the vector database.

The system retrieves the most semantically similar document chunks, such as:

- Maintenance history records
- SOP troubleshooting sections
- Engineering notes from prior incidents

4. Send Prompt to Inference API

The application augments the original user prompt with the retrieved contextual content. This enriched prompt is then sent to a text generation model hosted on Nutanix Enterprise AI via the inference API.

5. Get Answer

The chatbot generates and returns a response grounded in the organization's actual operational data.

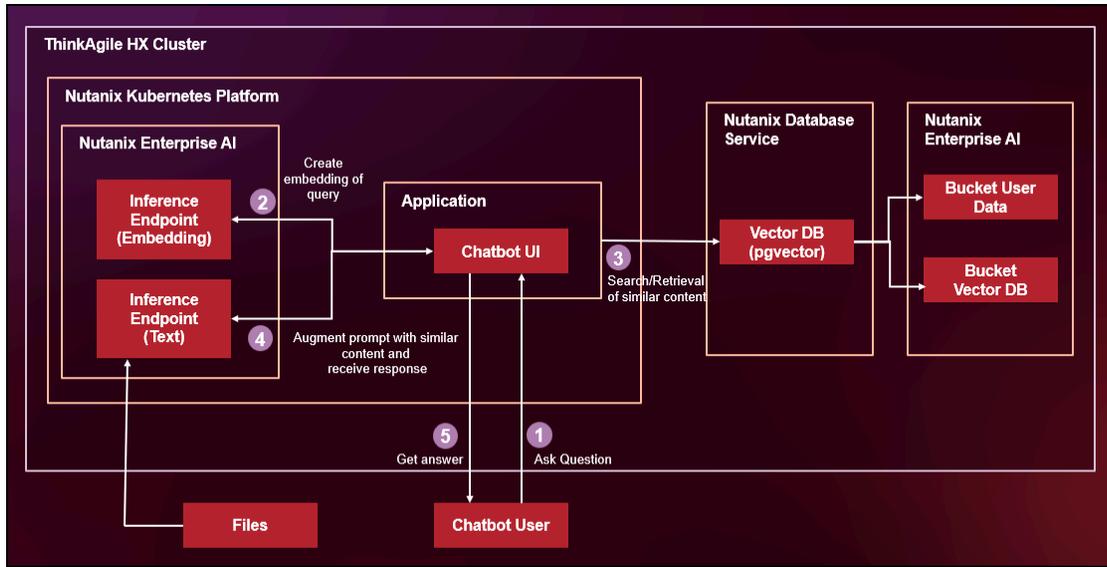


Figure 3. End-User Query Workflow

Bill of Materials - ThinkAgile HX650a V4

Configuration tips:

- NVIDIA H200 GPU can be replaced with NVIDIA RTX 6000 Server Edition GPU
- Based on the workload/use case, the Nutanix SW license quantity can be increased.

Table 1. Bill of Materials - ThinkAgile HX650a

Part number	Description	Quantity
7DG4CTO3WW	Server : ThinkAgile HX650a V4	1
CARX	ThinkAgile HX650a V4 Base	1
BVGL	Data Center Environment 30 Degree Celsius / 86 Degree Fahrenheit	1
B0W3	XClarity Pro	1
B15S	Nutanix SW Stack on Nutanix AHV	1
B0W1	3 Years	1
BM84	ThinkAgile HX Remote Deployment	1
BVKV	Nutanix Cloud Platform (NCP) Pro Software License with Mission Critical Support	1
BSQA	Nutanix Unified Storage (NUS) Pro Software License & Mission Critical Software Support per Tib Data stored, select total for a cluster	1
C84Q	Nutanix Kubernetes Platform (NKP) Pro Software License for NCI with Production Support	1
CC4R	Nutanix Enterprise AI Pro, Mission Critical Support Per GB of GPU RAM	1
C5QT	Intel Xeon 6530P 32C 225W 2.3GHz Processor	2
C3QR	ThinkSystem 2U V4 Performance Heatsink	2
C0TQ	ThinkSystem 64GB TruDDR5 6400MHz (2Rx4) RDIMM	8
B0SW	Nutanix Flash Node Config	1
C3Q9	ThinkAgile HX 2.5" U.2 VA 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	6
C46P	ThinkSystem 2U V4 8x2.5" NVMe Backplane	1
C26V	ThinkSystem M.2 RAID B545i-2i SATA/NVMe Adapter	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BE4T	ThinkSystem Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
B8PP	ThinkSystem Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet Adapter	1
C3V3	ThinkSystem NVIDIA H200 NVL 141GB PCIe GPU Gen5 Passive GPU	2
C62D	ThinkSystem SR650/a V4 x16 Rear Direct Riser Slot 5	1
C9R5	ThinkSystem SR650a V4 x16 600W Front Riser Slot 23	1
CATS	ThinkSystem SR650a V4 x16 600W Front Riser Slot 21	1
C0UD	ThinkSystem 3200W 230V Titanium CRPS Premium Hot-Swap Power Supply	2
B4L0	1.0m, 16A/100-250V, C19 to IEC 320-C20 Rack Power Cable	2
CC8X	ThinkSystem 2U 6056 24K Ultra Fan Module for 600W PCIe Adapter	6
C3UG	ThinkSystem Long Travel Toolless Slide Rail Kit V4 with 2U CMA	1
C5MU	ThinkSystem SR650a V4 Standard Left Rack Latch	1
BPKR	TPM 2.0	1
5977	Select Storage devices - no configured RAID required	1
B7Y0	Enable IPMI-over-LAN	1

Part number	Description	Quantity
C3K9	XClarity Platinum Upgrade v3	1
C4S2	ThinkSystem SR650 V4 Processor Board	1
CB2P	SR650 V4 Laser service indicator	1
B0ML	Feature Enable TPM on MB	1
B6C1	Node Cores	64
B6C2	Node Tebibytes	21
9220	Preload by Hardware Feature Specify	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	2
C3RM	ThinkSystem 2U Air duct Filler for 1P	2
AURS	Lenovo ThinkSystem Memory Dummy	24
BPP5	OCP3.0 Filler with screw	1
BSJ8	ThinkAgile EIA Plate logo Label	1
BKD9	HX All-NVMe configuration	1
CBKJ	ThinkAgile HX650a V4 - Nutanix IP	1
C9R6	ThinkSystem 2U FGPU Riser Cage for 600w	2
C5MT	ThinkSystem SR650 V4 Main air duct mylar for Front AC GPU	1
C6RV	Nutanix Software Selection	1
C4S6	HV 2U V4 long Chassis L1 PKG BOM	1
C7Y8	ThinkSystem SR650 V4 System I/O Board	1
C3S5	ThinkSystem 2U V4 3FH Riser Cage	1
C26Y	ThinkSystem V4 CPU HS Clip	2
B0SQ	HX Badge 1	1
C3RN	ThinkSystem 2U Main Air Duct	1
C1QN	Warranty Services Upgrade	1
CBMX	ThinkSystem 2U Front Double Width Air Duct for 600w	1
C3RJ	ThinkSystem 2U 2LP Riser Cage Filler	2
BHSS	MI for PXE with RJ45 Network port	1
C3RH	ThinkSystem 2U 3FH Riser Cage Filler	1
C3SG	ThinkSystem SR650a V4 2U F GPU 8x25 HDD Cage	1
ATSB	Nutanix Solution Code MFG Instruction	1
BENN	HX First Power-On Label	1
BDYY	Mellanox Low-Profile Dual-Port QSFP56 PCIe Bracket	1
C3QW	ThinkSystem M.2 Signal & Power Cable, ULP 82P-SLX4/2X10 SB, 400/400mm	1
C3QU	ThinkSystem GPU Power Cable, PCIe16-PCIe16, 400mm	2
C6R2	Think System, PCIe Gen5 Cable, MTK 74-MCIOx8, 800mm	2
C4UT	Power Cable, 2x3+6 P-MTK PWR, 900mm	1
C3UD	ThinkSystem BHS 2U 25 HDD Backplane Installation Label	1
C3U0	ThinkSystem 3200W Ti Power rating Label WW	1
C89N	ThinkAgile SR650 V4 Agency Label - Blank	1
C706	ThinkAgile SR650 V4 Service Label - WW	1
AWF9	ThinkSystem Response time Service Label LI	1
C3SZ	ThinkSystem BHS 2U Front GPU 2.5" NVMe HDD Label (on pull tag)	1

Part number	Description	Quantity
C3T7	ThinkSystem BHS 2U Front GPU Slot 20-23	1
B97B	XCC Label	1
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
BZ7F	ThinkSystem WW Lenovo LPK, Birch Stream	1
BE0E	N+N Redundancy With Over-Subscription	1
BK15	High voltage (200V+)	1
BTTY	M.2 NVMe	1
CA4D	MI-RS21, PCIe3&4/ PWR3	1
CA4C	MI-RS23, PCIe1&2/ PWR4	1
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
3444	Registration only	1
7S0PCTO3WW	Nutanix P&P Software for ThinkAgile HX	1
SDKM	Nutanix Kubernetes Platform Pro for NCI, Production Support Per vCPU, 3Yr	1
7S0PCTO3WW	Nutanix P&P Software for ThinkAgile HX	1
SALS	Nutanix Unified Storage Pro 1 To 50 TiB, Mission Critical Support For 1 TiB Data Stored, 3Yr	1
7S0PCTO3WW	Nutanix P&P Software for ThinkAgile HX	1
SE6W	Nutanix Enterprise AI Pro, Mission Critical Support Per GB of GPU RAM, 3Yr	282
7S0PCTO3WW	Nutanix P&P Software for ThinkAgile HX	1
SAPU	Nutanix Cloud Platform Pro, Mission Critical Support Per Core, 3Yr	64
7S0XCTO8WW	XClarity Controller Prem-FOD	1
SCY0	Lenovo XClarity XCC3 premier - FOD	1
5MS7B00045	ThinkAgile HX Remote Deployment (up to 3 node cluster)	1
7Q04CTSAWW	SW DEFINED KEEP YOUR DRIVE ADD-ON	1
QAK6	KYD	1
QAPB	HX650a V4	1
QA0Y	Months	36
7Q04CTS4WW	SW DEFINED PREMIER 24X7 4HR RESP	1
QAPB	HX650a V4	1
QA0Y	Months	36
QA12	24x7 4hr Resp	1
QA18	Premier	1

Author

Amalu Susan Santhosh is the Worldwide Technical Product Manager for Lenovo's ThinkAgile HX and MX/SXM Series of Hyperconverged Infrastructure (HCI) solutions. Amalu is responsible for showcasing the business value and differentiation of Lenovo's hybrid cloud solutions and contributing to the product lifecycle process.

Related product families

Product families related to this document are the following:

- [Artificial Intelligence](#)
- [ThinkAgile HX Series for Nutanix](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP2390, was created or updated on March 16, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2390>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2390>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkAgile®

ThinkSystem®

XClarity®

The following terms are trademarks of other companies:

Intel®, the Intel logo and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Other company, product, or service names may be trademarks or service marks of others.