Lenovo Validated Design: AI Resilience with Veeam Kasten on Lenovo AI Factory

Last update: **02 March 2026**

Version 1.0

Transforming backup from an add-on into a core architectural capability

Enterprise-class protection for mission critical for GenAI and RAG workloads

Kubernetes-native protection for AI models and pipelines

Built-in resilience across the Lenovo AI Factory architecture

Vanita Meyer
Hapsara Sukasdadi
Connor Blumsack
Edgar Rodriguez

Table of Contents

Introduction

Executive Summary

Enterprises are rapidly scaling generative AI, inferencing pipelines, and retrieval-augmented generation (RAG) workloads into production environments. As AI becomes embedded in core business operations, protecting models, vector databases, feature stores, and Kubernetes application state becomes essential to maintaining availability, compliance, and cyber resilience.

This Lenovo Validated Design (LVD) delivers a validated, scalable architecture for AI data resilience using Veeam Kasten on the Lenovo AI Factory (Hybrid AI 285 platform). The solution embeds Kubernetes-native, application-centric protection directly into a GPU-optimized AI infrastructure stack, enabling rapid recovery, immutable backup, and automated, policy-driven protection for production AI workloads. Rather than treating protection as a secondary layer, this design validates how resilience is architected as a core component of the Hybrid AI 285 Reference Architecture, ensuring that high-performance AI pipelines operate with built-in recovery, mobility, and security capabilities.

By combining Lenovo's high-performance compute and enterprise storage platforms with Veeam Kasten's Kubernetes-native protection capabilities, this design enables organizations to:

- Protect AI pipelines end-to-end across scalable units
- Preserve immutable recovery points to strengthen ransomware resilience
- Enable seamless application mobility across clusters and environments
- Minimize downtime for inferencing and mission-critical AI services
- Deploy AI infrastructure with resilience engineered from the start
- Recover from AI Prompt Injection attacks and restore state for forensic analysis

This approach enables organizations to confidently scale AI initiatives while maintaining operational continuity, cyber resilience, and performance integrity across mission-critical workloads.

Intended Audience

This document is intended for:

- Enterprise Architects designing AI platforms
- Infrastructure and Platform Engineering teams operating Kubernetes environments
- AI/ML Platform Owners responsible for model lifecycle management
- DevOps and Cloud-Native Operations teams
- Cybersecurity and Risk Management stakeholders
- System Integrators and Channel Partners deploying AI Factory solutions
- Lenovo and Veeam sales, solution architects, and technical specialists

It is particularly relevant for organizations deploying generative AI, inferencing pipelines, or hybrid AI workloads that require enterprise-grade protection, immutability, and rapid recovery capabilities.

Challenges and Opportunity

Key Challenges

As AI platforms scale into production environments, resilience risk increases due to architectural complexity and workload criticality.

- Kubernetes has become the de facto control plane for AI application deployment
- AI workloads are stateful and distributed across models, datasets, and metadata
- Critical data may exist both in-cluster and externally (e.g., vector DBaaS)
- Legacy VM-centric backup tools lack Kubernetes application awareness
- Inferencing services require low downtime and fast recovery to protect revenue
- Ransomware and malicious deletion risk increase as AI becomes mission-critical
- Rapid data growth from model iterations and embedding drives storage complexity
- Hybrid AI deployments require cross-cluster and cross-environment mobility

Without Kubernetes-native protection, recovery often requires manual reconfiguration of application components, increasing downtime and operational risk.

Strategic Opportunities

Embedding resilience into the Lenovo AI Factory architecture transforms protection from reactive backup into proactive operational enablement.

- Shift from reactive recovery to policy-driven AI protection
- Enable rapid model rollback and version control
- Protect Kubernetes applications holistically — including metadata and state
- Implement immutable recovery tiers for cyber resilience
- Separate high-performance recovery from long-term archival tiers to optimize cost
- Enable secure workload mobility across clusters and hybrid environments
- Deliver AI platforms with built-in governance and risk mitigation

This LVD positions AI resilience as a core architectural layer rather than an add-on, ensuring AI systems remain available, recoverable, and secure at enterprise scale.

Business Impact

Understanding the financial and operational impact of downtime, data loss, cyber-attacks, and compliance failures is essential for organizations evaluating resilience investments. Industry benchmarks estimate typical enterprise downtime at $300K+ per hour, with many organizations reporting costs exceeding $1M per hour[1]; the average total cost of a data breach in 2024 was $4.88M[3], and ransomware incidents have reached average total costs of approximately $5.13M[5]. In parallel, retraining modern AI models, particularly at scale, can cost from hundreds of thousands to tens of millions of dollars in compute and engineering time[6]. Embedding Veeam Kasten within Lenovo's architectures reduces these exposures by accelerating recovery, protecting critical AI assets, and aligning backup and retention with compliance and audit requirements.

Table 1: Enterprise AI Risk & Financial Exposure

Line Item	**Stat**	**Source**
Downtime $/hr – most enterprises[1]	>$300K/hr; 41% see $1M+	ITIC 2024 Global Server Hardware Survey
Downtime cost per minute[2]	~$5,600/min (~$336K/hr)	Gartner (widely cited benchmark)
Avg breach cost (2024)[3]	$4.88M	IBM Cost of a Data Breach Report 2024
Breach cost in high-risk sectors[4]	Up to ~$9.77M (healthcare)	IBM Security industry breakdown
Ransomware avg total cost[5]	~$5.13M	Industry ransomware cost analysis 2025
Large model training cost ranges[6]	Hundreds K – $100M+	Cloud provider AI cost analyses & published LLM estimates

Footnotes:

[1] ITIC 2024 survey: Avg enterprise downtime cost >$300K/hr, many at $1M+

[2] Gartner-based outage cost benchmark

[3] IBM Cost of a Data Breach Report 2024 ($4.88M)

[4] Security Scorecard healthcare breach impact analysis

[5] Ransomware cost ~ $5.13M (PurpleSec)

[6] Cloud AI cost guidance & LLM training cost reporting

Solution Overview

AI on Kubernetes - The Foundation of Modern AI

Kubernetes has become the de facto orchestration platform for enterprise AI and ML workloads. It enables scalable GPU scheduling, autoscaling, and orchestration of complex AI stacks including feature stores, vector databases, fine-tuning pipelines, and real-time inference services.

However, as AI environments transition from pilots to production, the data residing inside Kubernetes clusters, models, embeddings, vector databases, metadata, configuration, and secrets, become mission-critical. Protecting these digital assets against accidental loss, corruption, misconfiguration, or cyberattack is now a business requirement, not an operational afterthought.

Why Traditional Backup is Inadequate

Traditional infrastructure-centric backup tools were designed for monolithic applications and virtual machines. They protect disks, not distributed applications.

AI workloads on Kubernetes are fundamentally different:

- Highly dynamic and distributed
- Composed of persistent volumes, ConfigMaps, Secrets, CRDs, and Operators
- Dependent on stateful services like vector databases
- Spread across clusters and environments

Protecting AI applications requires treating the entire Kubernetes application as the complete unit of protection capturing not only storage volumes, but also configuration, environment state, metadata, and orchestration logic. Without Kubernetes-native protection, recovery often requires manual reconfiguration of components, increasing downtime and operational risk.

Veeam Kasten - The Resilience Layer for Lenovo Hybrid AI 285

The Lenovo Hybrid AI 285 platform enables rapid deployment of enterprise AI Factories supporting LLM fine-tuning, inference, and RAG workloads. Built on Kubernetes distributions such as Canonical Kubernetes and Red Hat OpenShift, it requires container-native resilience architecture. Veeam Kasten provides that resilience layer. Integrating Veeam Kasten into Lenovo’s Hybrid AI stack ensures data protection, recovery, and mobility are strategic capabilities, not peripheral backup add-ons.

As a Lenovo Alliance Technology Partner, Veeam delivers Kubernetes-native data protection purpose-built for cloud-native and AI environments. Unlike VM-centric backup solutions, Kasten is designed specifically to protect distributed, containerized workloads at scale.

Together, Lenovo Hybrid AI 285 + Veeam Kasten deliver:

- Production-grade AI resilience
- Policy-driven Kubernetes protection
- Ransomware-ready immutability
- Cross-cluster mobility and recovery

This transforms AI infrastructure from experimental to enterprise ready.

Protecting the Kubernetes AI Software Stack

The Lenovo Hybrid AI 285 platform supports complex AI software stacks, including:

- Container runtimes (Containerd)
- Helm-managed applications
- Load balancers
- NVIDIA AI Enterprise
- NVIDIA Inference Microservices (NIM)
- Vector databases and distributed storage systems

Kasten protects these environments by capturing a complete application “recipe”:

- Persistent volumes
- Databases
- ConfigMaps
- Secrets
- Kubernetes Operators
- Application metadata

If a critical AI service fails, the entire stack can be restored identically, eliminating configuration drift and accelerating recovery.

Safeguarding Stateful AI Data and RAG Workloads

AI pipelines are inherently stateful. Vector databases, embeddings, training artifacts, and checkpoints represent high-value enterprise assets.

Lenovo Hybrid AI 285 integrates with enterprise storage platforms such as:

- Lenovo ThinkSystem DM/DG Series
- IBM Storage Scale

Kasten leverages Kanister, its open-source data management framework, to:

- Automatically discover stateful services
- Quiesce databases during backup
- Apply policy-driven protection
- Ensure consistent, corruption-free recovery

This ensures vector databases supporting RAG workloads remain consistent and recoverable — protecting both data integrity and AI output reliability.

In terms of model training, Veeam Kasten’s support for AI/ML workloads means it captures model metadata, training checkpoints and datasets that minimizes the need for full retrains.

Security and Ransomware Resilience for High-Value AI

AI models, corporate data, and vector databases are incredibly valuable digital assets, making them prime targets for cyberattacks. Kasten integrates directly into the Kubernetes control plane of the Lenovo platform to provide zero-trust security without breaking native Kubernetes protocols.

- Immutability: Kasten supports immutable backup targets, preventing ransomware from deleting or tampering with AI backups.
- Encryption and RBAC: Kasten enforces strict Role-Based Access Control (RBAC) and ensures that all AI data is encrypted both in transit and at rest using algorithms like AES-256-GCM.
- Isolated Recovery: Kasten provides secure recovery environments to verify clean restores before production reintroduction.

Veeam Kasten embedded at the policy layer enables real-time alignment of backup and retention with service levels. This enables AI environments to meet enterprise security, compliance, and sovereign AI requirements.

Seamless Scaling and Application Mobility

Lenovo Hybrid AI 285 scales from entry Starter Kits to large deployments supporting up to 8 Scalable Units and 256 GPUs. Kasten automatically scales protection policies as workloads grow, eliminating manual intervention. Additionally, Kasten enables Kubernetes Application Mobility:

- Migrate AI workloads across clusters
- Support hybrid deployments (on-prem training, cloud inference)
- Enable DR across data centers
- Facilitate test/dev and model promotion workflows

This flexibility aligns with Lenovo’s Hybrid AI strategy, meeting customers wherever they operate.

Compliance Readiness & Auditability

Veeam Kasten ensures FIPS 140-compliant encryption, full audit logging, and retention policies that are aligned to enterprise and regulated AI environments.

Positioning these as core features of Lenovo’s AI stack satisfies customers governing sensitive workloads (healthcare AI, financial modeling) with audit-ready backups and tamper-proof records.

Solution Components

The Lenovo Hybrid AI 285 platform is a modular, production-ready AI infrastructure stack that integrates high-performance GPU compute, scalable enterprise storage, and advanced networking. This section outlines the validated hardware and software components that collectively form the foundation of the Lenovo AI Factory architecture, including GPU-optimized compute nodes, next-generation NVIDIA acceleration, unified storage tiers, deployment models, and Veeam Kasten for Kubernetes-native protection. Together, these components deliver scalable AI inference and RAG capabilities with resilience, security, and operational continuity directly into the platform.

Hybrid AI 285 AI Platform with Veeam Components Overview

Table 2: Compute

Component	Key Specifications	Business / Technical Value
Lenovo ThinkSystem SR675 V3	• 2U GPU-optimized server • 2x AMD EPYC 9535 (64-core, 2.4GHz, 3.5GHz boost) • Up to 8x PCIe double-wide GPUs • Up to 5x network adapters	Optimized for NVIDIA 2-8-5 configuration. Supports MIG partitioning with sufficient CPU cores per instance. Ideal for scalable AI training and inference workloads.
RTX Pro 6000 Blackwell Server Edition	• 96GB GDDR7 • 4 PFLOPS AI FP4 performance • Up to 5x inference vs L40S	High-performance acceleration for generative AI and visual computing. Enterprise-ready multi-workload AI GPU.
NVIDIA H200 NVL	• 141GB HBM3e memory • 4.8 TB/s memory bandwidth • Energy-efficient architecture • 5-year NVIDIA AI Enterprise license included	Designed for large LLM and HPC workloads. Handles larger models with improved throughput and cost efficiency.

Table 3: Storage

Platform	Core Capabilities	AI / RAG Benefits
DM7200F / DG5200 (All-Flash)	• Unified file, block, object storage • Deduplication & compression • Flexible scale-out architecture • Autonomous ransomware protection	Eliminate data silos. Delivers high-performance data access for training and embeddings. Enterprise-grade security and compliance for GenAI workloads.

Table 4: Data Recovery

Software	Core Capabilities	AI / RAG Benefits
Veeam Kasten v8.5.2	• Kubernetes-native backup and restoration • Automatic workload discovery	Downtime reduction and data protection.

Platform

The Lenovo Hybrid AI 285 platform is a modular, scalable AI infrastructure designed to support enterprise training, fine-tuning, and high-performance inference workloads.

The platform scales from:

- Single-node starter configuration with 4 GPUs
- To a full Scalable Unit (SU) consisting of four servers and 32 GPUs
- Up to 5 Scalable Units, supporting 20 servers and 160 GPUs

This modular architecture enables organizations to begin with a right-sized deployment and expand seamlessly as AI workload demands increase.

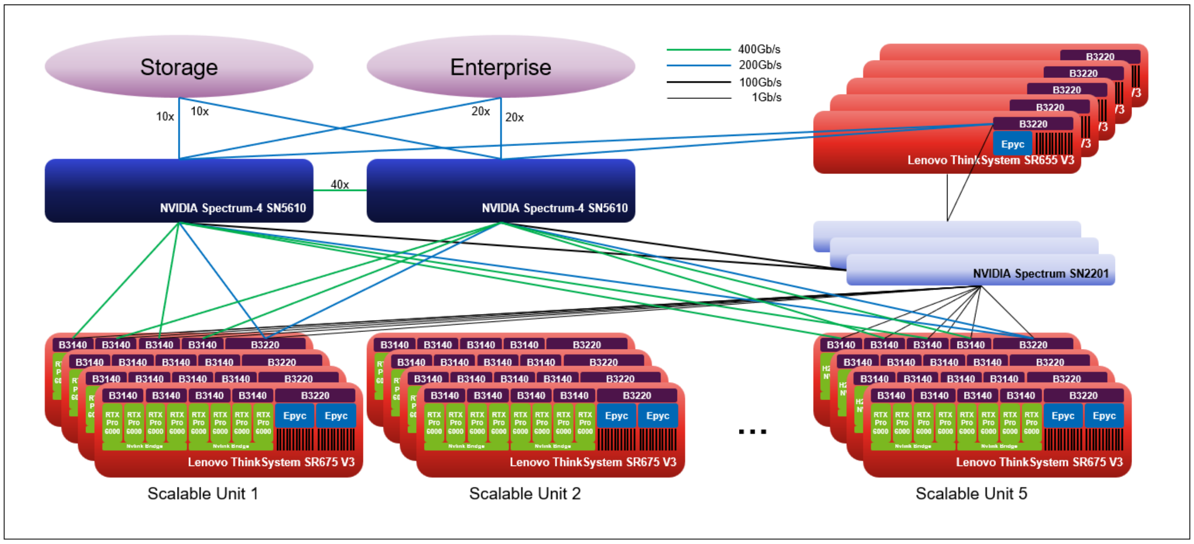

Figure 1: Lenovo Hybrid AI 285 platform with 5 Scalable Units

For larger enterprise deployments, the platform can scale up to **8 Scalable Units**, supporting 32 servers and 256 GPUs. In these configurations, additional NVIDIA SN5600 switches can be introduced to create a dedicated east-west (E/W) fabric optimized for GPU-to-GPU communication.

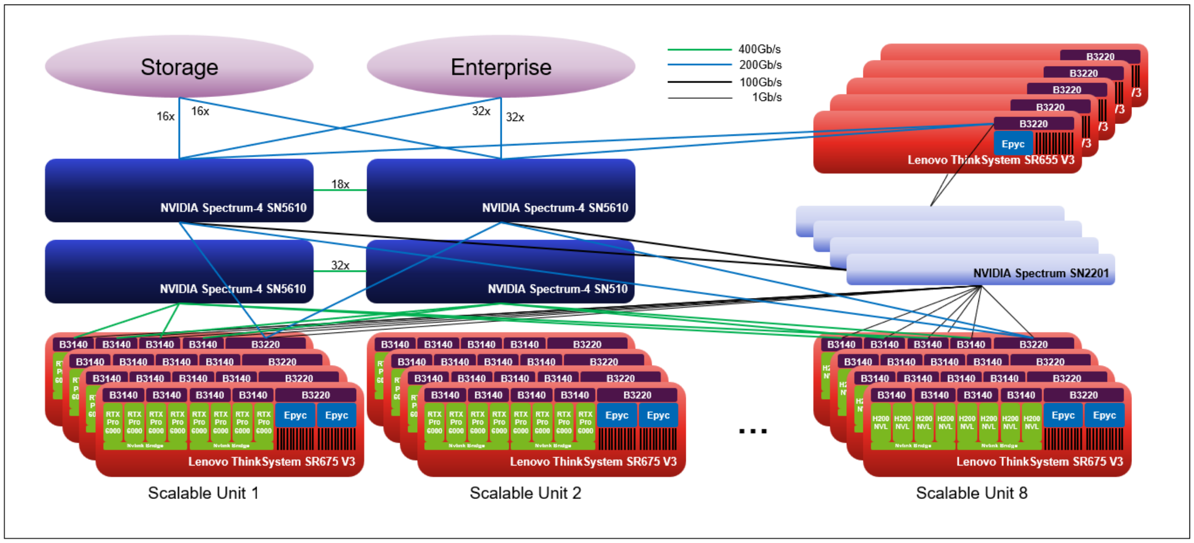

Figure 2: Lenovo Hybrid AI 285 platform with 8 Scalable Units

Enterprise AI Starter Kits

The Lenovo Enterprise AI Starter Kits are pre-validated, modular infrastructure solutions designed to simplify and accelerate the deployment of AI workloads, including Generative AI (GenAI) Inferencing and Retrieval-Augmented Generation (RAG). By integrating compute, storage, and networking into a cohesive, scalable architecture, the Starter Kits provide a turnkey foundation that removes the complexity typically associated with building AI environments from scratch.

Built on pre-engineered and performance-optimized configurations, the Starter Kits reduce deployment time while ensuring architectural consistency and scalability. High-performance NVMe flash storage and NVIDIA GPUDirect technology enable accelerated data movement between storage and GPUs, minimizing bottlenecks and reducing time to insight for mission-critical AI workloads.

Organizations can begin with a single-node configuration and seamlessly scale to full Scalable Unit deployments as performance, concurrency, and capacity requirements grow — allowing AI initiatives to expand without architectural redesign, operational disruption, or infrastructure rework.

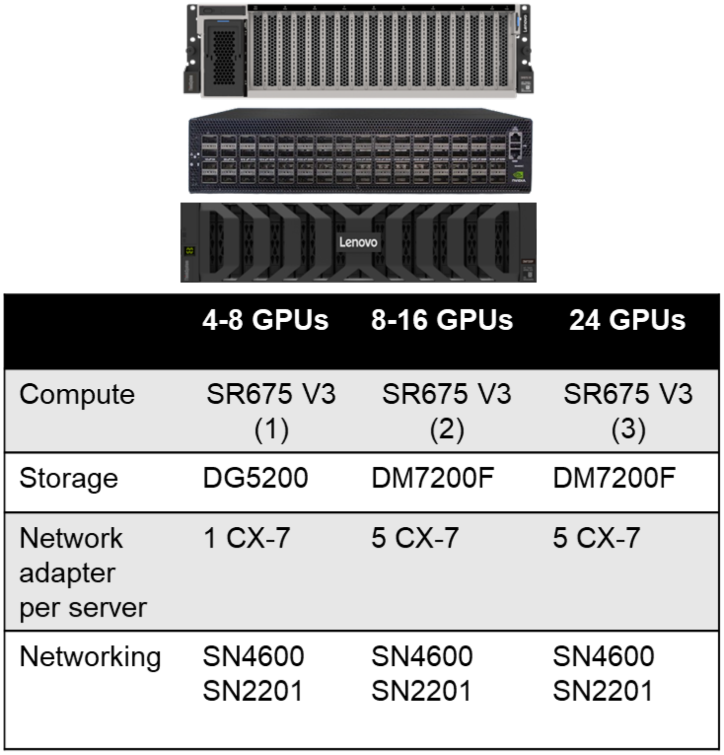

	4-8 GPUs	8-16 GPUs	24 GPUs
Compute	SR675 V3 (1)	SR675 V3 (2)	SR675 V3 (3)
Storage	DG5200	DM7200F	DM7200F
Network adapter per server	1 CX-7	5 CX-7	5 CX-7
Networking	SN4600 SN2201	SN4600 SN2201	SN4600 SN2201

Figure 3: Lenovo Enterprise AI Starter Kits

Scalable Unit Deployment

For deployments exceeding two AI compute nodes, Lenovo recommends implementing a full Scalable Unit (SU) along with the required networking and service infrastructure. A single SU supports up to 32 GPUs establishes a standardized, enterprise-ready foundation that supports predictable performance, high availability, and future expansion.

The first Scalable Unit consists of:

- Up to four AI Compute nodes
- A minimum of five service nodes (management, scheduling, and control-plane functions)
- High-bandwidth networking switches

As AI workload demand increases, additional Scalable Units, each containing four AI Compute nodes, can be added incrementally. This modular expansion model enables linear scaling of GPU capacity while preserving architectural consistency and operational simplicity across the environment.

This design ensures that enterprises can scale from initial production deployments to large, multi-SU AI factories without infrastructure redesign or service disruption.

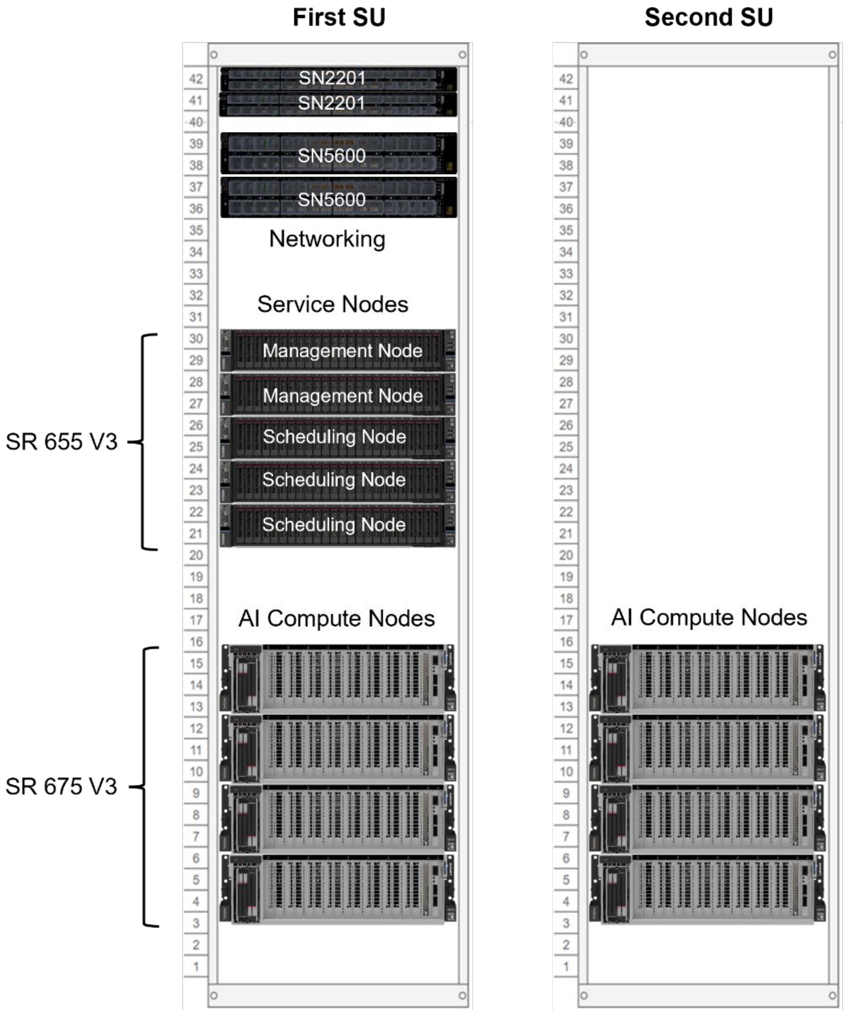

Figure 4: Scalable Unit Deployment

The networking decision depends on whether the platform is designed to support up to five or eight Scalable Units in total, and whether it will handle exclusively inference workloads or also encompass future fine-tuning and re-training activities. Subsequently, the solution can be expanded seamlessly without downtime by incorporating additional Scalable Units, ultimately reaching a total of five or eight as needed.

For more information on the Lenovo Hybrid AI 285 Platform, see the full platform guide.

Software Components

Veeam Kasten

As Kubernetes adoption accelerates, organizations need to address the critical requirement of protecting their Kubernetes applications, KubeVirt virtual machines (VMs), and their data. To keep things running, robust protection and recovery of the entire application, along with data services, must be prioritized to overcome misconfiguration, outage, and security threats that compromise availability.

Kasten is a Kubernetes-native data protection, backup, and restoration product by Veeam. Data is the most valuable resource for AI workloads and ensuring enterprise data is safe and secure is a top priority. With Veeam Kasten, enterprises can have their workloads be automatically discovered, backed up off cluster, and reduce downtime in the case of a cluster failure. Kasten has additional features to protect data, including encryption, granular access management, KubeVirt Virtual Machine protection, and an intuitive GUI.

Veeam Kasten delivers secure, Kubernetes-native data protection and application mobility at scale and across a wide range of distributions and platforms. Proven to recover entire applications simply, quickly, and reliably, Kasten gives operations and application teams the confidence to withstand the unexpected. Unlike traditional backup tools that view data as raw disk volumes, Kasten is application-centric, meaning it understands the entire Kubernetes application stack, including configuration, metadata, and stateful data.

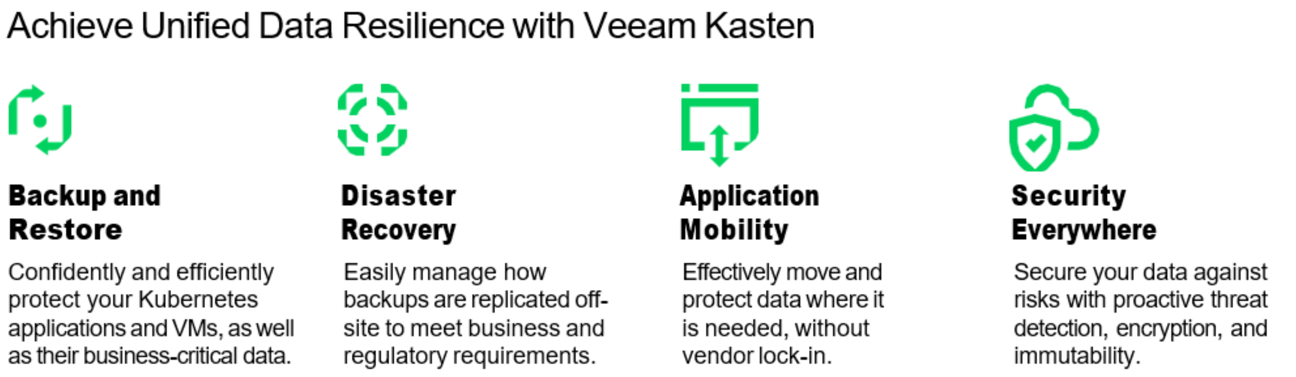

Achieve Unified Data Resilience with Veeam Kasten

Backup and Restore
Confidently and efficiently protect your Kubernetes applications and VMs, as well as their business-critical data.

Disaster Recovery
Easily manage how backups are replicated off-site to meet business and regulatory requirements.

Application Mobility
Effectively move and protect data where it is needed, without vendor lock-in.

Security Everywhere
Secure your data against risks with proactive threat detection, encryption, and immutability.

Figure 5: Veeam Kasten

Core Capabilities

Backup and Restore

Automatically discovers Kubernetes applications and captures both persistent data such as PVCs and Kubernetes resources including ConfigMaps, Secrets, CRDs, and Deployments. This enables complete application recovery.

Disaster Recovery

Orchestrates failover and failback workflows to recreate full application stacks in new namespaces or across clusters to maintain business continuity.

Application Mobility

Enables cross cluster portability through the Transform Engine, supporting migration between Kubernetes distributions such as Amazon EKS and Red Hat OpenShift, as well as across different storage backends including Lenovo ThinkSystem DM and DG Series.

Ransomware Resilience

Supports immutable backups and air gapped storage, including integration with Veeam Vault, to protect AI assets from deletion or encryption.

Optimized Performance

Uses an efficient data mover with built in deduplication, compression, and incremental backups to minimize storage consumption while maintaining fast backup and restore performance.

Protection of On and Off Cluster Data

Through application blueprints and Kanister integration, Kasten protects both in cluster Kubernetes resources and external data services such as DBaaS or vector databases under a unified recovery point.

Key Technical Features

Kanister Integration

Leverages the open source Kanister framework to perform application consistent backups by quiescing databases such as Qdrant, PostgreSQL, MySQL, or MongoDB prior to snapshot creation to ensure data integrity.

Native Kubernetes Integration

Built using Kubernetes Custom Resource Definitions, allowing backup and restore workflows to be managed through kubectl, GitOps pipelines, or standard Kubernetes automation frameworks.

Unified Container and VM Protection

Supports both containerized workloads and KubeVirt based virtual machines, including Red Hat OpenShift Virtualization, enabling consistent protection across Kubernetes environments.

Policy Driven Automation

Allows administrators to define protection policies that automatically secure new workloads based on labels, namespaces, or application metadata.

Highly Configurable Architecture

Advanced Helm configuration options allow tuning for performance, storage targets, and workload specific requirements.

Cloud Native Footprint

Designed with a minimal resource footprint, consuming cluster resources only during active data operations and releasing them once complete

Data Protection Solution Design

The purpose of this section is to provide a blueprint of the Hybrid AI 285 Data Protection Solution to meet the following objectives:

- Enable consistent protection of AI inferencing workloads and associated data
- Provide multiple recovery mechanisms aligned with workload criticality
- Support fast recovery, ransomware resilience, and operational continuity
- Preserve performance characteristics of production inferencing environments

It describes the physical and logical components of the solution and how they interact at a high level. The intent is to establish a common architectural baseline before diving into detailed design decisions.

Core Infrastructure Components

This section details the primary infrastructure elements that form the foundation of the AI 285 Data Protection Solution. Each component is described in terms of its role within the architecture and its relevance to data protection workflows.

Lenovo Hybrid AI 285
GPU-optimized compute platform suitable for inferencing pipelines, model hosting, and vector search.

Lenovo DM5200F (All-Flash)
High-performance primary protection tier used to store **instant copies, snapshots, and fast-recovery backups**. Integrated into the **100 GbE fabric** of the AI cluster.

Lenovo DG Series Storage
Scalable, efficient archival and long-term retention tier Integrated into the **100 GbE fabric** of the AI cluster. Ideal for historical versions of inferencing models, datasets, and application state.

Veeam Kasten K10 for Kubernetes Data Protection
Manages backup and restore, disaster recovery, application mobility, and policy-driven protection for all Kubernetes-based AI workloads. This LVD integrates Veeam Kasten v8.5 to deliver Kubernetes-native data resilience across the AI lifecycle.

Trident CSI Driver
Provides dynamic and consistent persistent volume provisioning for Kubernetes workloads, enabling seamless storage consumption without manual intervention.

Lenovo provides standard support for all systems interfacing with Trident. Direct support for the Trident platform is available through community resources.

Architecture Overview

This design ensures organizations can protect inferencing models, maintain immutable recovery points, migrate AI services across environments, and recover quickly from failures or cyber events — all while supporting the high throughput and low latency required in production AI inferencing environments.

The architecture integrates high-performance compute, multi-tiered storage for SU deployment, and Kubernetes-native data management. The solution is designed for both small-scale and large-scale environments:

- **Starter Kit Configuration:** Supports 1-3 SR675 V3 AI compute nodes. Networking options include high availability (HA) with dual NVIDIA SN4600 switches and dual SN2201 management switches, or a non-HA configuration with single switches.
- **SU (Scalable Unit) Environment:** Scales from a first SU containing service nodes (management and scheduling nodes) and AI compute nodes up to eight SUs.

Networking design follows the standard networking design of Hybrid AI 285. For AI Starter Kit deployment, all networking is based on a single 400G switching fabric. For Scalable Unit deployment, networking is based on 800G switching fabric, with either a single converged fabric or a separate East-West and North-South switching fabric. Please refer to Lenovo Hybrid AI 285 Platform Guide for more details.

AI Data Protection for Lenovo Hybrid AI 285 – Logical Design

This tiered approach allows ONTAP to handle high-performance, short-term snapshots, while DG storage provides cost-efficient, immutable, and portable backups for recovery across clusters or locations.

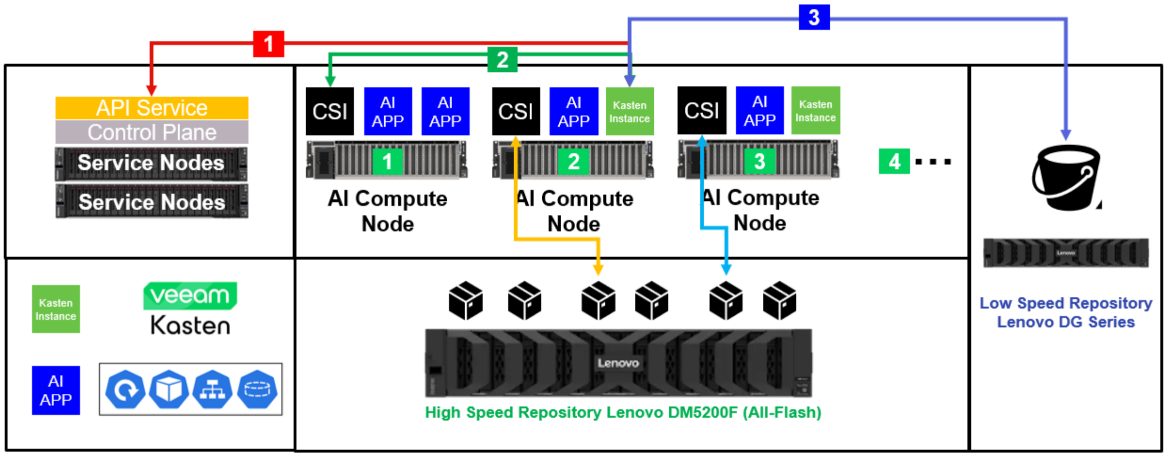

Figure 6: Logical Architecture of AI Data Protection with Veeam Kasten

Cluster Visibility

Veeam Kasten integrates with the Kubernetes API to automatically discover clusters, namespaces, applications, and their associated YAML resources. This provides full awareness of application topology and dependencies. Furthermore, Kasten supports integration with Red Hat OpenShift Advanced Cluster Manager (ACM) allowing for metric aggregation and observability across multiple clusters.

Data Capture (Primary Storage Tier – Lenovo ThinkSystem DM5200F)

Using the Lenovo DM Trident CSI driver, Kasten orchestrates application-consistent snapshots of persistent volumes (PVCs/PVs) stored on ONTAP. These snapshots are taken directly at the storage layer, ensuring fast, space-efficient protection without impacting application performance.

Export & Protection (Secondary Storage Tier – Lenovo ThinkSystem DG)

- Application metadata and Kubernetes configuration are compressed, duplicated, and securely exported to immutable object storage.

- ONTAP snapshots are then converted into durable backup copies and transferred from the primary ONTAP storage tier to object storage, enabling long-term retention, immutability, and disaster recovery.

AI Data Protection for Lenovo Hybrid AI 285 – Starter Kit Design

The Starter Kit Architecture provides baseline architecture for protecting AI workloads, specifically focusing on the integration between AI 285 Starter Kit and Veeam Kasten K10. This configuration is designed to handle the protection of applications while maintaining the high-throughput requirements of an AI pipeline.

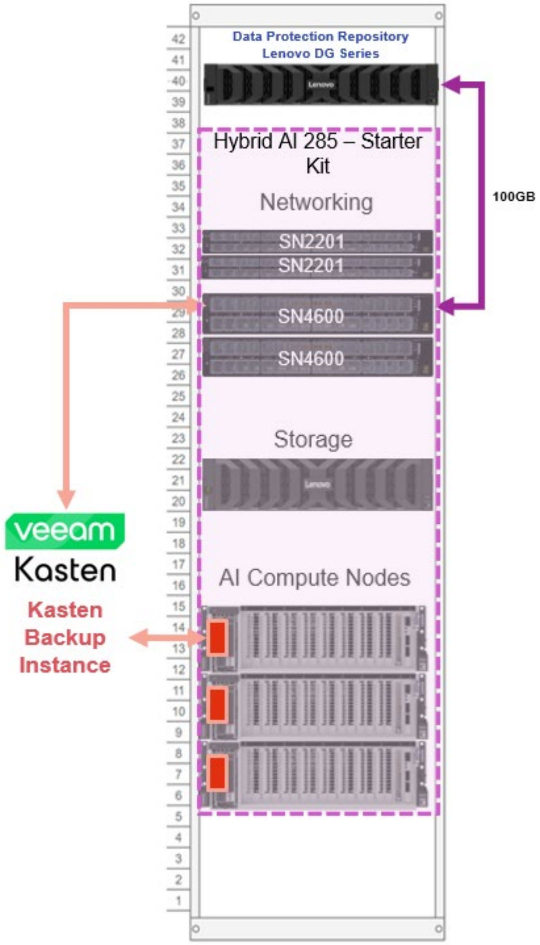

Figure 7: AI Data Protection for Lenovo Hybrid AI 285 – Starter Kit

Data Protection Repository

Unique Tier (Lenovo ThinkSystem DG series)

Purpose: Instant copies, fast recovery, continuous snapshotting and archival storage, long-term retention, compliance.

Network: 200 GbE fabric. Network connections using 200G to 2x100G splitter cables to DG series

Data to be stored: AI Model + Model Instances + Output Archiving

Rationale:

- Ensures model versions can be instantly recovered, maintains low-latency access for fast restore of inferencing workloads. Ideal for recent model iterations, vector DB snapshots, and “hot” datasets. Offloads “cold” data and historical inferencing model versions. Reduces TCO by placing
- non-performance-critical backups on a cost-efficient tier. Separates archival traffic from primary AI traffic to prevent performance degradation.

Considerations:

Depending on the data requirement, we recommend DG5200 (2.2PB Raw Max Raw Max no clustered configs) or DG7200 (7.2PB Raw Max no clustered configs). If the cluster reaches maximum capacity, please consider adding interconnect switching for the DG series.

Veeam Kasten Backup instances

Purpose
Provide automated backup scheduling, rapid application recovery, policy driven archival management, long term data retention, and regulatory compliance for Kubernetes based AI workloads.

Value Delivered

- **Instant Recovery and Local Backups**
  Enables fast restore of AI models, vector databases, and application state directly from the primary protection tier to minimize downtime.
- **Policy Driven Archival Copies**
  Automatically moves backup copies to the Lenovo ThinkSystem DG Series for cost-efficient long-term retention and compliance.
- **Disaster Recovery Workflows**
  Supports replication to additional DG targets or cloud object storage to enable cross site recovery and business continuity.
- **Kubernetes Persistent Volume Snapshots**
  Leverages Trident CSI to orchestrate storage level snapshots and ensure data is placed on the appropriate performance or archival tier.

AI Data Protection for Lenovo Hybrid AI 285 – Scalable Unit Design

The SU architecture provides baseline architecture for protecting AI workloads, specifically focusing on the integration between AI 285 SU and Veeam Kasten K10. This configuration is designed to handle the protection of applications while maintaining the high-throughput requirements of an AI pipeline.

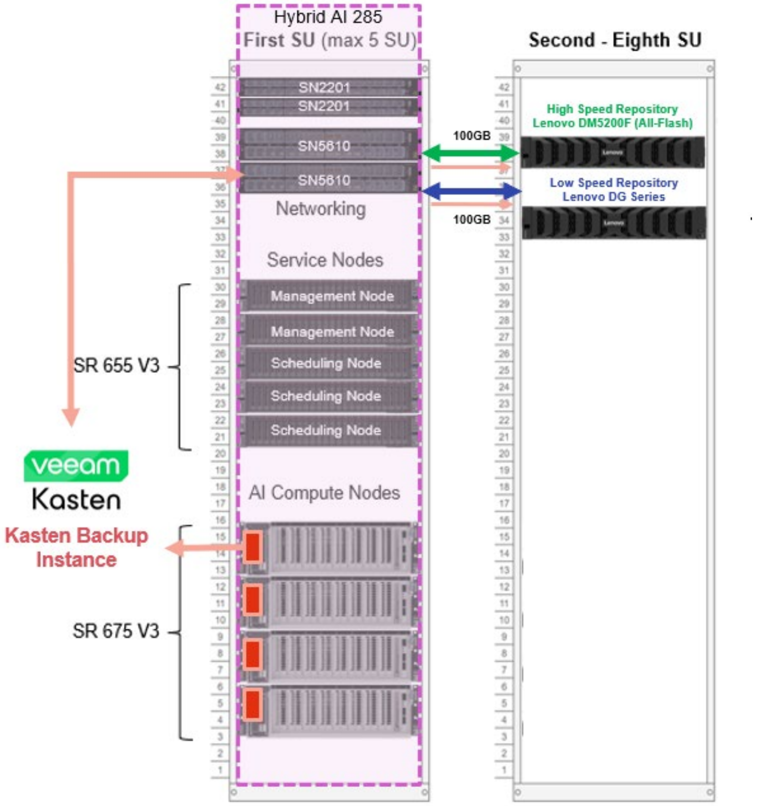

Figure 8: AI Data Protection for Lenovo Hybrid AI 285

High Speed Repository

Tier 1 – High-Performance Protection Tier (Lenovo ThinkSystem DM5200F)

Purpose: Instant copies, fast recovery, continuous snapshotting
Network: 800 GbE fabric. Network connections using 400G to 4x100G splitter cables to DM5200F

Rationale:

Ensure model versions can be instantly recovered. Maintains low-latency access for fast restore of inferencing workloads. Ideal for recent model iterations, vector DB snapshots, and “hot” datasets.

Low Speed Repository

Tier 2 – Archive & Long-Term Retention Tier (Lenovo ThinkSystem DG Series)

Purpose: Archival storage, long-term retention, compliance
Network: 800 GbE fabric. Network connections using 400G to 4x100G splitter cables to DG Series

Rationale:

Offloads “cold” data and historical inferencing model versions from the high-performance tier. Reduces total cost of ownership by placing non-performance-critical backups on a cost-efficient storage layer. Separates archival traffic from primary AI workloads to prevent performance impact on production inference.

If the cluster requires a switched or clustered DG configuration, interconnect switching should be included.

Additional NICs may be required depending on port availability and scale requirements.

Considerations:

Depending on the data requirement, we recommend DG5200 (2.2PB Raw Max Raw Max no clustered configs)) or DG7200 (7.2PB Raw Max no clustered configs). If the cluster reaches maximum capacity, please consider adding interconnect switching for DG series. Please refer to Lenovo ThinkSystem DG5200 and DG7200 Unified Storage Arrays for more details.

Veeam Kasten Backup instances

Enable automated backup scheduling, rapid application recovery, policy driven archival, long term data retention, and regulatory compliance for Kubernetes based AI workloads running on Lenovo Hybrid AI 285.

Data Stored by Protection Tier

Tier 1 High Performance Protection Tier
AI models and active model instances supporting production inference and real time AI services.

Tier 2 Archive and Retention Tier
AI models, model instances, and output archives retained for compliance, historical analysis, and long term recovery**.**

How Kasten Leverages Both Tiers

Instant Recovery and Local Backups
Backup copies are stored on the Lenovo ThinkSystem DM5200F to enable fast restore of production workloads with minimal downtime.

Policy Driven Archival Copies
Backup data is automatically moved to Lenovo ThinkSystem DG Series for cost efficient long term retention.

Disaster Recovery Workflows
Backup copies may be replicated to additional DG targets or supported cloud object storage platforms to support cross site recovery.

Kubernetes Persistent Volume Snapshots
Trident CSI is used to orchestrate storage level snapshots and ensure data is placed on the appropriate performance or archival tier.

Veeam Kasten Licensing and Design

This section outlines the infrastructure resources required to deploy Veeam Kasten and summarizes key licensing considerations for operating within a Lenovo Hybrid AI 285 environment.

Backup Instance Systems Requirements

The following minimum compute configuration supports protection of approximately 100 Kubernetes applications:

Table 5: System Requirements

Component	CPU	RAM

K10 base system	1–2 cores	1–4 GB
K10 DR (optional)	1 core	0.3 GB
Dynamic Backup Processing Pods	1 core	0.4 GB
TOTAL recommended	**~4 cores**	**~4.7 GB**

These resource requirements are lightweight and designed to minimize impact on production AI workloads while ensuring reliable backup and restore operations.

Licensing & Sizing Rules:

Veeam Kasten is licensed on a per node basis**.**

- One Veeam Kasten subscription is required per AI Compute Node per year
- Subscription terms are available from 1-5 years
- Standard Support is included.

This model aligns licensing directly with compute scale, ensuring protection scales consistently was additional Scalable Units are deployed.

Example

A customer deploying a two Scalable Unit Lenovo Hybrid AI 285 configuration with a five-year subscription would require one Kasten license per AI Compute Node across both Scalable Units**.**

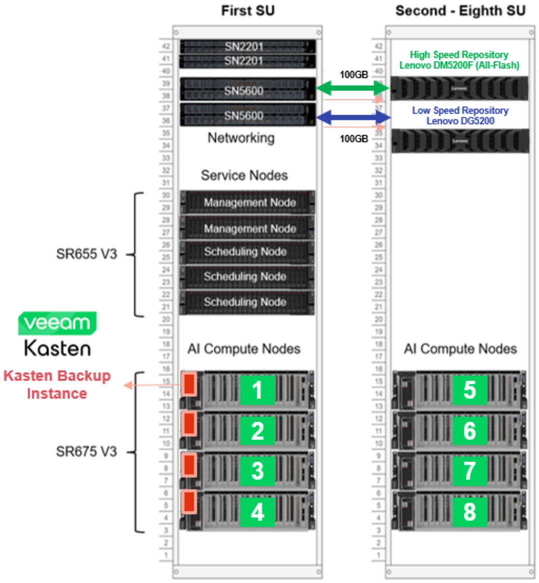

Kasten Backup Instance Licensing

Licensing & Sizing Rules:

1 x Veeam Kasten **subscription per AI Compute Node** node/per year (1-5 years) with Standard Support included

Example: Customer looking for a 5-year Kasten with Standard support for a 2xSU Lenovo ThinkSystems 285

Lenovo SKU	Description	Quantity
7S0L00L7WW	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 5Yr w/Veeam Basic Support	8
7S0L00NKWW	Veeam Support uplift from Kasten Basic support to Kasten Premium Support for Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition. w/Veeam support 5Yr	8

Figure 9: Veeam Kasten Licensing

Veeam Kasten Licensing

Kasten is licensed per-node, with options from 1 to 5 years and including Veeam support. See appendix A for Lenovo SKUs.

Lenovo ONTAP Trident CSI and Veeam Kasten Integration

Lenovo ONTAP Trident CSI enables Kubernetes clusters to consume persistent storage from Lenovo ThinkSystem DM Series systems running ONTAP through standard Container Storage Interface (CSI) mechanisms. Trident acts as the storage orchestration layer, allowing Kubernetes applications to request and manage persistent volumes using native Kubernetes constructs such as StorageClasses and PersistentVolumeClaims (PVCs).

Through this integration, storage provisioning and lifecycle management are automated and aligned with Kubernetes operations, providing consistent and scalable access to enterprise ONTAP storage for containerized workloads.

Integration with Veeam Kasten (K10)

Veeam Kasten integrates with Trident CSI by leveraging CSI-based snapshot capabilities to protect Kubernetes applications and their associated persistent data. When deployed together, Veeam Kasten coordinates application-level protection while Trident facilitates snapshot operations on the underlying ONTAP storage.

This integration enables:

- Protection of Kubernetes application data stored on Lenovo ThinkSystem DM Series.
- Snapshot-based backup and restore workflows aligned with Kubernetes resources.
- Support for recovery and mobility use cases within Kubernetes environments.

Lenovo ThinkSystem DM Series has been validated within the Veeam Ready – Kubernetes program for CSI-based storage integration, confirming compatibility with Veeam Kasten when using Trident CSI.

Assessment Requirements

Before deploying the Lenovo Hybrid AI 285 Data Protection architecture, a structured technical and operational assessment is required to ensure performance integrity, resilience alignment, and policy compliance.

Backup & Data Protection Analysis

- Define and validate RPO/RTO targets for AI inferencing and production workloads.
- Classify AI data by protection tier (high-performance instant recovery vs. archival retention).
- Identify mission-critical models, vector databases, and stateful services requiring priority protection.
- Align retention policies with business, compliance, and sovereign AI requirements.

Connectivity & Infrastructure Assessment

- Validate network throughput and bandwidth allocation for storage fabrics (200 GbE / 800 GbE) and archival paths (10/25 GbE where applicable).
- Confirm jumbo frames, RDMA configuration, and acceptable latency boundaries for GPU and storage traffic.
- Ensure proper NIC redundancy, link aggregation, and switching fabric design (HA vs. non-HA).
- Validate separation of production AI traffic from backup and archival traffic to avoid performance degradation.

Kubernetes & Veeam Kasten Configuration Review

- Verify Kubernetes distribution and version compatibility.
- Confirm ONTAP Trident CSI installation and health.

- Ensure the Trident CSI VolumeSnapshotClass is properly annotated: k10.kasten.io/is-snapshot-class: "true"
- Configure and validate a Kasten Location Profile for off-cluster backup targets (Veeam Vault, S3/S3-compatible object storage, Azure Blob, GCS, NFS, or SMB).
- Audit namespace structure, labeling strategy, stateful workloads, PVC architecture, and operator versions.
- Validate Kasten protection policies, application blueprints, retention logic, immutability settings, and DR orchestration workflows.
- Perform a documented test restore to confirm end-to-end recoverability.

Lenovo Services

Lenovo Solution and Services Group (SSG) provides Hybrid AI Factory Services that delivers end-to-end lifecycle services spanning AI architecture design, deployment, configuration, and ongoing operations management. Services include multi-node GPU cluster implementation, Kubernetes and OpenShift enablement, scalable networking and storage integration (DDN, Lenovo DM/DG, Cloudian), and performance validation aligned to AI Factory reference architectures. The offering also integrates Veeam Kasten for Kubernetes-native backup, disaster recovery, and policy-driven data protection, ensuring resilient AI model training, fine-tuning, and inference environments. Together, these services provide a production-ready, secure, and scalable AI platform with built-in operational governance and business continuity.

Solution Summary

The Lenovo Validated Design (LVD) for AI Resilience with Veeam Kasten on Lenovo Hybrid AI 285 delivers a production-grade, Kubernetes-native data protection architecture purpose-built for enterprise AI, Generative AI, and RAG workloads.

As AI platforms transition from pilot to mission-critical production systems, protecting models, vector databases, training artifacts, configuration state, and application metadata become essential to maintaining availability, compliance, and cyber resilience.

This validated solution integrates Veeam Kasten directly into the Lenovo AI Factory architecture, transforming backup from a secondary add-on into a core architectural layer. Leveraging Kubernetes-native protection, policy-driven automation, immutable backup targets, and application-centric recovery, the design ensures complete protection of AI pipelines. Integration with Lenovo ThinkSystem DM (high-performance recovery tier) and DG Series (archival and long-term retention tier) storage enables a tiered protection model that balances rapid recovery with cost-efficient retention.

The architecture scales from Enterprise AI Starter Kits to multi-Scalable Unit deployments, ensuring resilience scales with compute growth. It supports instant recovery of inference services, cross-cluster mobility, ransomware protection, and compliance-ready auditability, while preserving the high throughput and low latency required for GPU-accelerated AI workloads.

By combining Lenovo Hybrid AI 285 infrastructure with Veeam Kasten's Kubernetes-native protection, organizations can confidently deploy AI factories with built-in operational continuity, reduced downtime risk, accelerated recovery objectives, and enterprise-grade governance aligned to modern AI lifecycle requirements.

Appendix A: Lenovo Bill of materials (BOM)

Lenovo Hybrid AI 285

For Lenovo Hybrid AI 285 Platform BoM please refer to:

Lenovo Hybrid AI 285 Platform Guide

DM7200 BoM

Example Bill of Materials (BoM) for the Lenovo ThinkSystem DM7200F All-Flash array used as the high-speed repository in Scalable Unit (SU) deployments. Final configuration details will vary based on customer capacity, performance, resiliency, and retention requirements.

Table 6: DM7200 BoM

Part number	Product Description	Qty
7DJ3CTO1WW	Controller : Lenovo ThinkSystem DM7200F All Flash Array	1
BF3C	Lenovo ThinkSystem Storage 2U NVMe Chassis	1
BWU8	Storage Complete Bundle Offering	1
C4A4	Lenovo ThinkSystem DM7200 Series Controller, 128GB	2
C3XK	Lenovo ThinkSystem 30.7TB (2x 15.36TB NVMe SED) Drive Pack	9
C4AA	Lenovo ThinkSystem Storage 100Gb 2 port Ethernet (Host/Cluster)	2
C4AA	Lenovo ThinkSystem Storage 100Gb 2 port Ethernet (Host/Cluster)	4
AV1Z	Lenovo 1m Passive 100G QSFP28 DAC Cable	10
BY6K	USB-A to USB-C Cable	1
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
CCWL	Lenovo ThinkSystem Storage ONTAP 9.17 Software Encryption - IPAv2	1
B0W1	3 Years	1
C6S2	Premier 24x7 4hr Response and KYD	1
C48T	Configured with Lenovo ThinkSystem DM7200F 3Yr Warranty	1
BWUC	Storage Complete Bundle License Key	2
BWUE	Storage Encryption Bundle License Key - RoW	2
C49B	Lenovo ThinkSystem DM/DG Series Jupiter All Flash Ship Kit - Multi-Language	1
B6Y6	Lenovo ThinkSystem NVMe Rail Kit 4 post	1
B738	Lenovo ThinkSystem NVMe Accessory	1
C48X	Lenovo ThinkSystem DM/DG/DS Jupiter/Saturn 2U24 NVMe Bezel	1
C8V9	7-segment LED cover Label	1
C9UY	PESS mandatory not be enforced	1
B6Y5	Lenovo ThinkSystem NVMe SFF Filler	6
C48W	I/O Slot Cover	2
C498	Lenovo ThinkSystem Storage Controller 2U24 NVMe Agency Label	1
C492	Lenovo ThinkSystem DM7200F Model Name Label	1
B4E7	EIA NamePlate	1
C3HV	Lenovo Logo nameplate	1

C5H3	Lenovo ThinkSystem NVMe Jupiter/Saturn Packaging	1
7S0SCTOMWW	ThnkSys DM7200F 7DJ3 SW License	1
SDJE	Lenovo ThinkSystem DM7200F NVMe SSD Unified Complete SW License with 3 Years Support, Per 0.1TB	2765
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
5WS7C06619	3Yr Premier 24x7 4Hr Resp DM7200F+KYD	1
5WS7C07259	3Yr Premier 24x7 4Hr Resp+KYD (0.1TB NVMe TLC)	2765

DG5200 BoM

Example Bill of Materials (BoM) for the Lenovo ThinkSystem DG5200 used as the Low-Speed Repository in Scalable Unit deployments and as the Data Protection Repository in Starter Kit configurations. Final BoM specifications will vary based on customer capacity, performance, and retention requirements.

Table 7: DG5200 BoM

Part number	Product Description	Qty
7DHYCTO1WW	Controller : Lenovo ThinkSystem DG5200 QLC All Flash Array	1
BF3C	Lenovo ThinkSystem Storage 2U NVMe Chassis	1
BWU8	Storage Complete Bundle Offering	1
C4A5	Lenovo ThinkSystem DM/DG5200 Series Controller, 64GB	2
C5RG	Lenovo ThinkSystem 30.7TB (2x 15.36TB QLC NVMe SED) Drive Pack	9
C4AA	Lenovo ThinkSystem Storage 100Gb 2 port Ethernet (Host/Cluster)	2
C4AA	Lenovo ThinkSystem Storage 100Gb 2 port Ethernet (Host/Cluster)	2
AV1Z	Lenovo 1m Passive 100G QSFP28 DAC Cable	6
BY6K	USB-A to USB-C Cable	1
6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord	2
CCWL	Lenovo ThinkSystem Storage ONTAP 9.17 Software Encryption - IPAv2	1
B0W1	3 Years	1
C6S2	Premier 24x7 4hr Response and KYD	1
C48U	Configured with Lenovo ThinkSystem DG5200	1
BWUC	Storage Complete Bundle License Key	2
BWUE	Storage Encryption Bundle License Key - RoW	2
C49B	Lenovo ThinkSystem DM/DG Series Jupiter All Flash Ship Kit - Multi-Language	1
B6Y6	Lenovo ThinkSystem NVMe Rail Kit 4 post	1
B738	Lenovo ThinkSystem NVMe Accessory	1
C48X	Lenovo ThinkSystem DM/DG/DS Jupiter/Saturn 2U24 NVMe Bezel	1
C8V9	7-segment LED cover Label	1

C9UY	PESS mandatory not be enforced	1
B6Y5	Lenovo ThinkSystem NVMe SFF Filler	6
C48W	I/O Slot Cover	4
C494	Lenovo ThinkSystem DG5200 Model Name Label	1
B4E7	EIA NamePlate	1
C498	Lenovo ThinkSystem Storage Controller 2U24 NVMe Agency Label	1
C3HV	Lenovo Logo nameplate	1
C5H3	Lenovo ThinkSystem NVMe Jupiter/Saturn Packaging	1
7S0SCTOPWW	ThnkSys DG5200 7DHY SWLicense	1
SDJG	Lenovo ThinkSystem DG5200 NVMe QLC Unified Complete SW License with 3 Years Support, Per 0.1TB	2765
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
5WS7C06770	3Yr Premier 24x7 4Hr Resp DG5200+KYD	1
5WS7C07660	3Yr Premier 24x7 4Hr Resp+KYD (0.1TB NVme QLC)	2765

Kasten Licensing

Table 8: Kasten Licensing

Type	Description	Charge Metric	Lenovo FC	Lenovo PN
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 1Yr w/Veeam Support	Per Node	SD6X	7S0L00KZWW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 1Yr w/Veeam Support Public Sector	Per Node	SD6Y	7S0L00L0WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 2Yr w/Veeam Support	Per Node	SD6Z	7S0L00L1WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 2Yr w/Veeam Support Public Sector	Per Node	SD70	7S0L00L2WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise	Per Node	SD71	7S0L00L3WW

	Edition Subscription. 3Yr w/Veeam Support			
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 3Yr w/Veeam Support Public Sector	Per Node	SD72	7S0L00L4WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 4Yr w/Veeam Support	Per Node	SD73	7S0L00L5WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 4Yr w/Veeam Support Public Sector	Per Node	SD74	7S0L00L6WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 5Yr w/Veeam Support	Per Node	SD75	7S0L00L7WW
Subscription	Veeam Kubernetes Backup, Restore, DR, and application mobility. Veeam Kasten Enterprise Edition Subscription. 5Yr w/Veeam Support Public Sector	Per Node	SD76	7S0L00L8WW

Resources

Resources	**Links**
Kasten Documentation	https://docs.kasten.io https://go.veeam.com/wp-kubernetes-radar https://go.veeam.com/kubernetes-enterprise-trends
Lenovo AI Factory 285	https://lenovopress.lenovo.com/lp2181-lenovo-hybrid-ai-285
Lenovo DM5200	https://lenovopress.lenovo.com/lp2149-lenovo-thinksystem-dm5200h-unified-storage-array
Lenovo DG5200 and DG7200	https://lenovopress.lenovo.com/lp2074-lenovo-thinksystem-dg5200-and-dg7200-unified-storage-arrays
1. Enterprise downtime cost: ITIC 2024 Global Server Hardware, Server OS Reliability Survey	https://itic-corp.com/itic-2024-hourly-cost-of-downtime-report
2. Downtime per minute estimate - Gartner-based analysis commonly cited across industry references	https://netcov.com/resources/average-cost-of-data-loss-data-center-outages
3. Average breach cost (2024): IBM Security Cost of a Data Breach Report	https://www.ibm.com/think/insights/whats-new-2024-cost-of-a-data-breach-report
4. Breach cost in high-risk sectors	https://securityscorecard.com/blog/what-does-a-data-breach-cost-key-insights-for-cyber-leaders
5. The average cost or ransomware attacks	https://purplesec.us/learn/average-cost-of-ransomware-attacks
6. AI training and retraining cost	https://cloud.google.com/blog/topics/cost-management/unlock-the-true-cost-of-enterprise-ai-on-google-cloud
7. Ransomware Trends by Veeam	https://go.veeam.com/ransomware-trends

Document history

Version 1.0	March 2, 2026	Lenovo Hybrid AI Factory 285 and Veeam Kasten v8.5.2

Trademarks and special notices

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®
ThinkEdge®
ThinkShield®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Intel Core® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Veeam® and Veeam Kasten® are registered trademarks of Veeam Software Group GmbH

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

Information is provided "AS IS" without warranty of any kind.
All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.
Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serves as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites are at your own risk