



Red Hat

Day 2 and Beyond

Cloud to Edge. AI to Virtualization.

Running OpenShift with Confidence at Enterprise Scale.

BEFORE DAY 2 ...PREPARATION IS CRITICAL

Veeam delivers value at every stage. From migration planning, through Day 0 landing and Day 1 operations, to Day 2 scale and beyond.

Deployment/Migration planning

Day 0 · Landing

Day 1 · Operate

Day 2 · Scale

Beyond · Evolve

THE GAP WHERE IT BREAKS

Unknown protection scope

Recovery absent from design

Coverage drifts as apps land

Recovery unproven at scale

Audit evidence is manual

Modernization stalls

Not Meeting Regulatory Requirements

Unrecoverable Data Loss

veeam

WHAT VEEAM DELIVERS AT EACH STAGE

Workload discovery & scoping

Protection designed in, not added

Automated policy enrollment

Unified VM + container recovery

Validated recoverability

Security & Resiliency

Mobility for what's next

Every stage makes a decision you inherit. Design In Day 2 Requirements Early.

THE MIGRATION WAVE IS ALREADY HERE

Plan for 'similar to vSphere' Data Protection ahead of Migration



417% more VMs in a year

VMs running on Red Hat OpenShift Virtualization grew 417% during 2025. Accounts running VMs grew 70%. (Red Hat, April 2026)



1.46M VMs assessed

Red Hat Services has assessed more than 1.46 million VMs through virtualization migration engagements. Planning is happening now, at scale. (Red Hat, April 2026)



October 11, 2027

vSphere 8 reaches end of general support. Decisions made in the next eighteen months set the operating model for the decade after.



Confidence outruns proof

90% of organizations are confident they can recover from a cyber incident. Fewer than one in three can actually prove it. (Veeam Data Trust & Resilience Report, 2026)

THE PLANNING WINDOW

AI ADOPTION WITH LIMITED GOVERNANCE / GUARDRAILS

Data Resiliency is more important than ever



1

3X Adoption of AI agents

Number of active AI agents in organizations have tripled from last year. ¹

Large Enterprises have a scaling fleet of 50 to over 200 Agents. ²



2

66% of Orgs Do Not Trust deployed Agents

The study found that 86% of organizations have moved beyond the AI agent pilot stage, yet just 34% say they trust the actions their AI agents are taking. ²



3

6.3M Orders Impacted

Amazon Kiro AI agent coding tool created two separate incidents: Two outages that wiped out 6.3M in orders



4

1000s of High-Severity & Zero-Day Vulnerabilities

Anthropic's Claude Mythos demonstrated autonomous cyber-offense capabilities. Resulted in project Glasswing.

Threat actor groups are undoubtedly developing their own vuln scanning tools to have parity (and surpass.)


THE PLANNING WINDOW

¹ Znet 2026 <https://www.zdnet.com/article/ai-agent-adoption-tripled-measurable-roi/>


² Frontier Enterprise 2026 - <https://www.frontier-enterprise.com/most-firms-have-deployed-ai-agents-few-trust-them/>

OPENSIFT IS BECOMING THE PLATFORM FOR MORE

The platform started stateless and ephemeral. Almost everything landing on it now has state that has to survive.



AND IT IS INCREASINGLY STATEFUL



Source: Dimensional Research survey of 500+ Kubernetes enterprises, published via CNCF, 2025.

ACT ONE



What breaks after migrations or new deployments

PLATFORM COMPLEXITY

VMS + CONTAINERS

DATA SERVICES

AI WORKLOADS

TEAMS & POLICIES

MULTI-CLUSTER / EDGE

MIGRATION HAD A PLAN. DOES DAY 2?

From bound, owned, funded, to unbound, unowned, and always on. Day 2 fails in four recognizable ways.



01

SCALE OUTPACES POLICY

New workloads (and clusters) appear faster than protection policies can be assigned.

AT RISK IF

New namespaces and VMs go live before a protection policy is attached.



02

OWNERSHIP FRAGMENTS

Every team assumes another team owns recovery after the project closes.

AT RISK IF

Nobody can name who owns recovery twelve months after go-live.



03

RECOVERY SPLINTERS BY WORKLOAD

VMs and containers create separate tools, procedures, and expertise requirements.

AT RISK IF

VMs and containers run through two tools, two teams, two runbooks.



04

GOVERNANCE BECOMES MANUAL

Audit evidence depends on spreadsheets, job logs, and institutional memory.

AT RISK IF

Audit evidence takes a week and a spreadsheet to assemble.

COMPLEXITY INCREASES

DAY 2 AND BEYOND • RED HAT SUMMIT CONNECT 2026

WHAT CLOSES EACH GAP

Four challenges, four specific mechanisms.



AUTOMATED ENROLLMENT

Selectors attach protection to new workloads as they appear, GitOps-aligned.

PROOF IT WORKED

New namespaces are covered without anyone filing a request.



ONE POLICY OWNER

Platform team owns policy. App and infra teams recover inside RBAC boundaries.

PROOF IT WORKED

You can name the recovery owner. So can the auditor.



ONE RECOVERY MODEL

Kasten for platform-native workloads, VBR for datacenter-rooted VMs. One vendor.

PROOF IT WORKED

VM and container coverage reported in the same place.



EVIDENCE ON DEMAND

Immutable targets, restore-point validation, evidence generated on request.

PROOF IT WORKED

Audit evidence takes minutes, not a week.

COMPLEXITY ABSORBED

DAY 2 AND BEYOND • RED HAT SUMMIT CONNECT 2026

ACT TWO



The modern resilience operating model

Modern platforms need a modern resilience operating model.

PLATFORM COMPLEXITY

VMS & CONTAINERS

DATA SERVICES

AI WORKLOADS

TEAMS & POLICIES

MULTI-CLUSTER & EDGE

Seamless Veeam | Protecting Every Workload on OpenShift

Veeam protects the entire Red Hat OpenShift estate



CONTAINERS, VMs, AI, and EDGE MOBILITY

VEEAM KASTEN

Kubernetes-native data protection for OpenShift

- Certified Operator & OLM integration
- OpenShift Virtualization protection
- Application-aware Kubernetes recovery
- Multi-cluster mobility & governance
- GitOps, ACM, and OpenShift ecosystem integrations

- Operates like a native OpenShift service
- Protect containers and VMs together
- Recover complete applications, not just data
- Standardize resilience across hybrid OpenShift environments
- Embed resilience into platform engineering workflows



VMS IN THE DATACENTER MOVING TO OPENSIFT VIRT

VDP FOR OPENSIFT VIRTUALIZATION

Veeam Backup & Replication with the Veeam KubeVirt Plugin

- Native OpenShift Virtualization support inside the VBR platform teams already trust
- Protect migrated VMs with the same jobs, immutability, and retention as the rest of the estate
- Enterprise-grade: object-lock storage, tape, backup copy, DR

- Protects OpenShift Virt VMs only
- Native VBR-based support via new on-cluster proxy
- Full VBR repository support (including Direct to Object)
- Cross-hypervisor restore
- File-level Restore

VEEAM KASTEN FOR KUBERNETES

Secure, Kubernetes-native resilience and application mobility for the most complete set of distributions and platforms on an enterprise scale.

BACKUP

RECOVERY

SECURITY

FREEDOM

INTELLIGENCE



SECURE, SCALABLE, RESILIENT MODERNIZATION



Kubernetes-native
data resilience for
modern containers
and KubeVirt VMs



UNIFIED DATA RESILIENCE

- Native protection for Kubernetes applications and OpenShift KubeVirt VMs
- Unified operations reduce complexity and cost



SECURE, AUTOMATED OPERATIONS

- Immutable, encrypted backups with RBAC and policy-driven automation
- Centralized, consistent security and compliance across workloads



ACCELERATED MODERNIZATION

- Protect modernize at your pace without disruption
- Built-in mobility and recovery across hybrid and multi-cloud environments

RELIABLE, SECURE, NATIVE

Veeam Kasten Integrations with Red Hat OpenShift

Unified data protection
for OpenShift everywhere



OpenShift
AI



OpenShift
Container
Platform



OpenShift
Platform
Plus



OpenShift
Virtualization
Engine



Advanced
Cluster
Security



Advanced Cluster
Management



OpenShift Data
Foundation

RELIABLE, SECURE, RESILIENT

WHY VEEAM KASTEN?

Complete data protection, no matter the OpenShift workload



Secure, Trusted Data

Backups are insulated from primary storage and can be used to restore operations quickly.



Isolation from Ransomware Attacks

Isolate production and backup environments to ensure operations are restored from a clean and reliable backup in the event of a security incident.



Point-in-Time Recovery

Recover your data from specific points in time, including situations where rolling back to a known state of operations is necessary.



Application & Data Mobility

Enable graceful migration of applications and data across OpenShift clusters, including to a separate product resource if a live platform is compromised.



Early Threat Detection

Raise an early flag on potentially malicious activity or imminent attacks, with integration into Red Hat Advanced Cluster Security (ACS).



Encrypted & Immutable Backups

Always have a safe and consistent copy of your application and data.

VEEAM KASTEN | SAFEGUARD AI APPLICATIONS

Protecting AI-Native & AI-Enhanced Apps

veeam Kasten



VECTOR DATABASES

And other DBs, underlying
persistent volumes, and
configuration



OBSERVABILITY & AUDIT

Logs required for compliance,
debugging, and rollback



APPLICATION CONFIG

e.g., web frontend and its related
secrets, manifests, and
ConfigMaps



MODEL ARTIFACTS

Including fine-tuned LLMs
and their associated
metadata workflows

Kubernetes increasingly popular for deploying AI apps – stateful by nature

Lenovo Validated Design: AI Resilience and Data Governance with Veeam on Lenovo AI Factory | JH Reviewed
May 2026

Lenovo Validated Design

AI Resilience and Data Governance with Veeam on Lenovo AI Factory

*Kubernetes-native protection, ransomware resilience, and AI governance for Lenovo
Hybrid AI 285*

Version 1.1 | Last updated: May 2026 | JH Reviewed
Authors: Vanita Meyer, Hapsara Sukasadi, Connor Blumsack, Edgar Rodriguez
Editorial review: Joe Home, Director Global Alliances, Veeam Software

JH Review Note: Title updated from 'AI Resilience with Veeam Kasten' to 'AI Resilience and Data Governance with Veeam.' The original title scopes the document to a single product and capability. DAICC is the governance layer that reaches the CDO/CDATO buyer and multiplies deal value. Version bumped to 1.1 to distinguish from the original lp2382. Version 1.0 remains the published validated design.

Confidential — For partner and internal use only. Version 1.1 | 1

Veeam being integrated as the AI Trust and
resilience layer for 'AI Factories'

PROTECTING AI INFERENCE WORKLOADS

App-consistent protection for relational and vector databases via Kanister blueprints in Kasten



THE CHALLENGE

WITH VEEAM KASTEN

RELATIONAL DATABASE

- Inference apps depend on relational state
- Crash-consistent snapshots risk corruption
- Manual dump scripts don't scale on Kubernetes



- App-consistent backup via Kanister blueprints
- Templates for all major relational databases

**CONSISTENT STATE
ON EVERY RESTORE**

VECTOR DATABASE

- Loss or corruption forces full re-embedding
- Hours of GPU compute to rebuild indexes
- Stale answers while pipelines re-run



- PGVector and Qdrant blueprints today, Weaviate coming shortly
- Restore to a point before loss or corruption

**RESTORE
EMBEDDINGS, DON'T
RECOMPUTE**

THE DAY 2 ADVANTAGE

- GPU capacity burned on data re-ingestion
- Customers lack resilient inference services
- Resilience gaps slow enterprise AI adoption



- Policy-driven protection for full inference stacks
- Embedding compute saved is GPU time for inference

**GPUS SERVE
INFERENCE, NOT
REBUILDS**

TWO THINGS BEFORE YOU GO...



Come Find Us

Visit the Veeam kiosk

- Learn more about how Veeam Kasten natively protects OpenShift everywhere it runs
- Chat about operational efficiency and how to simplify and automate Day 2 resilience
- Ask our team your most pressing questions about protecting mixed workloads



Kubernetes Korner

Learning, Tips and Product Info

- Join our world-class practitioner online community with a dedicated Kubernetes sub-section for tips, events and more.



See Kasten in Action

Ungated, Interactive Demos

What you'll find:

- Veeam Kasten on Red Hat OpenShift
- Protecting VMs on OpenShift Virt
- vSphere to OpenShift migration
- Application Mobility made easy



Scan Me!
You'll be glad you did.

And so much more at veeam.com



Thank you