Your Path to Resilience with Veeam Kasten

A roadmap for your first 100 days and beyond.

This guide walks IT administrators through deploying, configuring, and operationalizing Veeam Kasten *for Kubernetes*. Whether installed through Helm or Marketplace, the deployment path is designed to be fast, secure, and accessible. By Day 100, you will have a Kubernetes-native backup environment with immutable offsite exports that's integrated with your toolsets and applications, as well as a documented recovery runbook.

Your 100-Day Journey

PHASE 1 Foundation Days 1–14	PHASE 2 Harden the Basics Days 15–45	PHASE 3 Optimize & Harden Days 46–75	PHASE 4 Prove Value Days 76–100
M1: Size and plan	M4: Centralized management	M7: Observability integration	M10: Recovery testing
M2: Deploy Veeam Kasten	M5: App-aware processing	M8: Enhanced resiliency	M11: Documentation and reporting
M3: First backup jobs	M6: Policy best practices	M9: Close coverage gaps	M12: Ongoing hygiene

Roadmap at a Glance

Phase	Name	Timeline	Focus
PHASE 1	Foundation	Days 1–14	Plan environment, deploy Veeam Kasten, and run first backup jobs
PHASE 2	Harden the basics	Days 15–45	Centralized management, app-aware processing, and policy best practices.
PHASE 3	Optimize and harden	Days 46–75	Observability, resiliency and DR, and close coverage gaps.
PHASE 4	Prove value	Days 76–100	Recovery testing, documentation, and ongoing hygiene.

How to Use This Guide

Days are guidelines, not hard deadlines. Small environments may compress Phase 1 to a week, and complex ones may extend Phase 3.

Each milestone builds on the previous one, so it's recommended to follow the steps in order.

Callout boxes throughout highlight decision points, architecture options, and common gotchas worth taking into account.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Contents

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

PHASE 1 • Days 1–14

Foundation

Goal: Infrastructure sized, deployed, and first workloads protected.

Milestone 1: Size and Plan

Before deploying anything, invest time in planning. Most challenges with deploying a stable and well-protected environment can be avoided with proper attention to this step.

Workload Inventory

- Document total workload and virtual machine (VM) count, data footprint (e.g., volume size/file count), and estimated daily change rate.
- Identify your most critical workloads. These drive your recovery point objective and recovery time objective (RPO and RTO) targets.
- Note any databases and workloads that could benefit from application-aware processing.

Veeam Kasten Sizing and Deployment Planning

- Veeam Kasten has minimal system requirements and default values are usually sufficient.
- Choose your deployment method, either Marketplace/Operator (recommended when possible) or Helm. It is possible to switch methods post-install if needed.
- Verify your version of Kubernetes is compatible with the latest version of Veeam Kasten. If not, you will either have to upgrade first, or use an older version.
- Determine if this environment is air-gapped. If so, the Veeam Kasten images may need to be pushed to a local registry that's available offline through your vendor's marketplace, or have our registry whitelisted.
- Planning your backup network by isolating backup traffic onto a dedicated VLAN is a security best practice, and possible with Kubernetes depending on your CNI.

Helm VS Marketplace

Helm: This is the most flexible and compatible with all Kubernetes distributions.

Marketplace/Operator: The deployment and upgrade lifecycle is simpler and supports the same parameters for customization as Helm. It's also commonly used with Red Hat OpenShift.

The trade-off: There is a potential delay with receiving new releases through the marketplace. Updates are available immediately on Helm, but marketplace vendors often delay new releases for one or more days for validation testing.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Workload Storage Architecture

- **Snapshot-capable storage:** Veeam Kasten is compatible with virtually any CSI storage driver that supports snapshots, as well as direct storage integrations.
- **Other storage (e.g., generic NFS, local disk, etc.):** Veeam Kasten can still protect these workloads, but the lack of true snapshot capability affects the reliability and performance of backups. We strongly recommend you upgrade to using volumes with snapshot capabilities, but GSB or NFS tar could be used if you accept their inherent limitations.

Backup Storage Architecture: Overview of Options

Your storage choice shapes the resilience and availability of your backup data. Here is a general overview of the pro and cons of each type of backup storage that you can configure as a location profile.

Veeam Kasten + Veeam Vault

Snapshots are exported to off-site and logically air-gapped immutable object storage via Veeam Vault. No advanced configuration or maintenance is needed.

Veeam Kasten + Veeam Backup & Replication Repository

Snapshots are exported to Veeam Backup & Replication managed infrastructure, including Veeam Hardened Repositories and Scale-out Backup Repositories (SOBRs). Block-mode is required. **Note:** Pre-9.x versions also require a small additional Location Profile (Vault/Object/NFS/CIFS) for metadata.

Veeam Kasten + Object Storage

Snapshots are exported to object storage (e.g., S3, cloud), which should be configured to be offsite and immutable.

Veeam Kasten + NFS/CIFS

Snapshots exported directly to NFS/CIFS are generally incapable of immutability, even in WORM-capable devices, since versioning would be needed to function. This option is fully supported, but not advised unless you also have offsite immutable duplication or copies.

Veeam Kasten + Multiple Location Profiles

Snapshots are exported to 2+ locations. If designed correctly, using at least one immutable option, this can be the most resilient path. **Note:** Pre-9.x versions require multiple Policies, using 1 Location Profile per Policy.

Storage Modes Determine Capabilities

Depending on your storage type, you may have the following choices available:

Volume Mode: File or Block. Both are compatible with Veeam Kasten (e.g. Ceph FS or RBD).

Export Mode: File or Block. Both are possible with Veeam Kasten.

If a volume is in block mode, you may be able to export your snapshots in block mode. This offers special features, like CBT (when available) and exports to a Veeam Backup & Replication repository.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Using Veeam Backup & Replication + Veeam Infrastructure Appliance (VIA) as a Hardened Repository

For those without existing immutable storage, integrating Veeam Kasten with Veeam Backup & Replication and deploying VIA could be a reasonable path to achieving local immutability.

A hardened repository provides immutable local backups, which means ransomware cannot encrypt or delete them during the retention period. The OS is hardened against unauthorised access and should ideally run on a protected physical server instead of a VM.

Traditionally, a hardened repository requires manually installing and hardening a Linux distribution. VIA removes this barrier entirely since it is pre-hardened and deploys via bootable ISO. It can also be self-deployed on compatible hardware or purchased from a vendor as a certified appliance.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Milestone 2: Deploy Veeam Kasten

Deploy Veeam Kasten

- Use pre-flight tool to verify requirements are met (optional).
- Marketplace/operator method: See instruction details here.
- Helm method: See instruction details here.

Dashboard Access and Authentication Integration

- To access the dashboard externally, add a ingress or route to Veeam Kasten.
- It is best practice to integrate with your distribution's authentication provider (e.g. OpenShift Auth) too, but other options like direct LDAP are also possible.
- Once the Veeam Kasten dashboard is visible, you can continue to the steps below.

Connect Infrastructure to Veeam Kasten

- If applicable, (vSphere/cloud/others) add an Infrastructure Profile for direct storage integration.
- Add a Location Profile to define your backup destination.

Milestone 3: First Backup Jobs

Backups in Kubernetes Environments

Storage snapshots are a powerful feature that Kubernetes natively understands and can manage, if they're supported by your storage provisioner.

Veeam Kasten makes use of this capability or direct integration to trigger snapshots that live locally on the storage device. Veeam Kasten can then use these snapshots to reliably export data to an external Location Profile (e.g., NFS, S3, VBR repo, Vault, etc.).

This method allows for maximum speed and consistency, while still using Kubernetes-native features and APIs for maximum compatibility.

What Is Needed To Get Started

Veeam Kasten, being a Kubernetes-native application, needs the installing user to have kubectl access to the target environment, cluster admin privileges, and ideally access to your distribution's management interface too, if any.

Tip: If you encounter any errors, solutions can be found on our Community and KB sites.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

- Learn more about protecting KubeVirt VMs if applicable to your environment.
- Create your first backup Policy and target your Location Profile.
- Set a sensible retention policy to start: 1 local snapshot daily, with exports to the Location Profile set for 1 daily, 4 weekly, 3 monthly retention.
- Schedule the job to run during off-peak hours and verify it does not conflict with other maintenance windows.
- Run the job manually on first execution and monitor it through to completion.
- Confirm the job completes without warnings or errors before proceeding to Phase 2.

Your First Restore Test: Do Not Skip This

Before moving to Phase 2, perform a restore on a non-critical application to confirm recoverability.

You are not protected until you have verified you can restore. This takes minutes and can prevent days of pain later.

Local VS Export

Remember that snapshots are local only to your storage device. This means:

- They are not a true backup.
- More local snapshots retained means more primary storage consumed.

Thus, it is important to export the snapshots to a Location Profile, which can have an independently configured retention policy.

When exporting snapshots, all data is automatically deduplicated, compressed, encrypted, and saved incrementally.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

PHASE 2 • Days 15–45

Harden the Basics

Goal: Consistent protection, offsite copy, and visibility across your environment.

Milestone 4: Centralized Management (MCM)

Multi-Cluster Manager (MCM) enables centralized visibility and management of multiple clusters. Each cluster will have Veeam Kasten installed to protect workloads and the instances can be joined together.

- Once another cluster running Veeam Kasten has been deployed, promote one Kasten instance to Primary, and use a Join Token to link them (instructions). Be sure self-signed certs are trusted to prevent errors.
- View all joined clusters via the MCM dashboard.
- Define Global Policies and Global Profiles if you wish to have a consistent configuration distributed throughout your environment.
- Add your enterprise license key to the Primary MCM cluster. Licenses will seamlessly be distributed to all other clusters with no other configuration needed.

Milestone 5: Application-Aware Processing

Application-aware processing using Kanister Blueprints ensures crash-consistent backups become application-consistent or logically exported. This is critical for databases like PostgresSQL, MongoDB, MySQL, Elastic, and others.

- Discuss application requirements and access with your app team.
- Find pre-made Blueprint examples you can customize, or create your own.
- Add the Blueprint to Veeam Kasten.
- Associate the Blueprint with workloads manually or by creating a Blueprint Binding.
- Test a backup and restore using your Blueprint to verify end-to-end application recovery.

Types of Backups

Crash-consistent (default): This type of backup captures the entire filesystem at a precise moment, typically via storage provider snapshots. This is more data-consistent than basic file copying, but in-flight writes or unflushed transactions could still be missed.

Application-consistent: This type of backup typically quiesces the application briefly by using the application's native tools to ensure everything is written to the volume before the snapshot is created.

Logical: This type of backup uses the application's native tools to create a logical dump or export of its data as a file. This is often pushed to an external location like S3 instead of the volume itself.

Blueprints are powerful and infinitely flexible. Plus, with its open-source nature, you can build one for any application or workflow if one doesn't exist.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Milestone 6: Policy Best Practices

In this milestone you will create additional jobs to protect all your workloads and use advanced settings to ensure their reliability.

Best Practices (for Most Scenarios)

- Use multiple jobs with labels to select applications instead of one wild card policy protecting everything.
- Use VM-based policies instead of namespace-based when protecting KubeVirt VMs.
- If you wish to back up cluster-scoped resources, create a separate policy for only those backups.
- Define a backup window to ensure production workloads are not impacted during critical hours.
- Consider scheduling jobs with staggering to improve performance.
- Configure immutable backup protection.
- Exclusions / Resource filters if needed.
- (Optional) Global Resources, Policy Presets

Other Tips

For manually run jobs, unless an expiration time is specified, any artifacts will need to be manually removed.

Even if the job was configured using local time, all times are converted to UTC and policy schedules do not change for daylight savings time.

Check your math! Retention calculation is not always obvious. For example, retaining 24 (total) hourly snapshots at 15-minute intervals would only retain 6 hours of snapshots, not 24 hours.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

PHASE 3 • Days 46–75

Optimize and Harden

Goal: Close protection gaps, improve RTO/RPO, and enhance visibility.

Milestone 7: Observability Integration

- Enable Red Hat ACM Observability integration if applicable.
- Export metrics to external monitoring systems (e.g., Grafana, Datadog, Thanos, etc.).
- Create alerts and dashboards within your monitoring systems (Grafana example).
- Enable Reporting.

Milestone 8: Enhanced Resilience and DR

Having a backup of all workloads should also include having a backup of Veeam Kasten itself, which can be accomplished by Kasten Disaster Recovery (KDR). This should not be confused with having a DR environment, or a DR plan, which should all be considered based on the resilience required.

KDR (Backup of Your Veeam Kasten Catalog)

- Enable Kasten DR: This saves the Veeam Kasten backup catalog so that restores are possible even if the Kasten installation is no longer accessible. This is important, as Veeam Kasten backup files are not standalone and require the catalog and encryption passkey to restore.
- Set the Kasten DR to "Export local catalog snapshots" and save the following information in a secure location: Passphrase, cluster ID, KDR Location Profile details, and credentials.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

DR for Your Workloads

- Determine the level of resilience needed based on your desired RTO/RPOs.
- If a RTO/RPO of 1+ days is acceptable, only having backups and rebuilding the environment as needed could be a reasonable solution.
- If RTO/RPOs need to be measured in hours or minutes, careful attention must be made to design a plan. Some examples are given below.

DR Example Configurations (Depending on RTO/RPO needed)

- **Unprotected application:** No restore points (i.e. no local snapshots or exports).

  Result: No recovery possible.

- **Export only:** Exports of KDR and exported restore points exist in an offsite Location Profile.

  Result: To recover, first rebuild a Kubernetes environment, install Veeam Kasten, restore the Kasten catalog, then restore applications from exported restore points.

- **DR environment + Import:** Veeam Kasten is installed and an Import Policy runs regularly to import the restore points into a local catalog.

  Result: Restore points are visible at a secondary location and are ready to be restored. Minimal resources are consumed until they're needed.

- **DR environment + Import + Restore, scaled to zero:** Veeam Kasten is installed, Import + Restore Policy runs regularly to import the restore points into local catalog and then restore applications to Kubernetes environment. Applications are then restored but with Transforms setting replicas to 0 on all workloads.

  Result: Applications are restored into Kubernetes (via pods, PVs, everything selected). However, replicas equal 0, so minimal resources are consumed until they're needed aside from storage.

- **DR environment + Import + Restore, full:** Veeam Kasten is installed, Import + Restore Policy runs regularly to import the restore points into a local catalog then restore applications to Kubernetes environment.

  Result: Applications are restored into Kubernetes (via pods, PVs, everything selected) running exactly as they were in production and ready to be fully utilized.

Disaster Planning

Veeam Kasten is distribution agnostic, meaning it can restore to and from any Kubernetes distribution, even between completely different ones like Tanzu to Azure. However, restores require that a working Kubernetes environment exists to run; Kasten cannot bootstrap itself. Be sure to create a proper disaster plan so that you are ready in case there is an incident and regularly test that your plan is meeting your objectives.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Milestone 9: Fine Tuning

- **Review the Veeam Kasten dashboard:** Address all unprotected workloads before doing any other tuning.
- **Review job schedules for conflicts:** Stagger start times to prevent resource contention.
- Confirm all jobs are completing within your defined backup window.
- **Validate RPO compliance:** Are all critical workloads generating restore points within your target recovery window?
- Confirm all workloads are exporting to at least one immutable location.
- Enabling Garbage collection to remove old actions can aid in catalog performance.
- Test your upgrade process! New Veeam Kasten versions are released roughly every two weeks.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

PHASE 4 • Days 76–100

Prove Value and Operationalize

Goal: Verify recoverability, establish ongoing hygiene, and demonstrate ROI.

Milestone 10: Recovery Testing

The only backup that matters is one you can restore from. Phase 4 is where you prove — with documented evidence — that your environment meets its RTO and RPO commitments.

Granular and Full Restore Tests

- Test a full namespace or VM restore from an export. Simulate total on-premises loss to validate your offsite copy.
- Test application item recovery: Restore a database and run database queries to test that your data is intact.
- Test file-level restore (if desired): Recover individual files from an export to a test location.
- Test your full DR plan end-to-end, note any issues to improve, and extensively document the process in a manner that makes sure your plan and access to it are available
- Record actual recovery times and compare against your RTO targets. Document the results.

Recovery Testing Best Practice

Always restore to a non-production target and never overwrite live workloads during a test.

Document what was restored, from which restore point, to which target, and how long it took.

These results are your proof of recoverability. Retain them for compliance reviews, audits, and management reporting.

Additionally, if you're restoring within the same environment, even to different namespaces, be aware that certain resources (especially networking services) may conflict. Be sure to review your configuration before restoring, restore to a different environment, and/or use Transforms to prevent conflicts.

Restoring to Different Eviroments

If restoring to/from different sites, Kubernetes distributions, platforms, or otherwise, Transforms can be useful to seamlessly modify your configuration upon restore. For example, if you're changing network paths from production to DR, or change storage classes. No manual changes are needed to have your restores running immediately with a modified configuration or in a new environment.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Milestone 11: Documentation and Reporting

- Enable Reporting, if not already done.
- Document your final backup architecture, including policy list, export locations, schedules, and retention policies.
- Review your Veeam Vault storage consumption and confirm usage aligns with your expected budget.
- Archive recovery test results alongside architecture documentation.

Next Steps

Ready for more?

- Visit Veeam University
- Explore Success Stories from Your Peers
- Connect with Customer Success
- Share Your Data Resiliency Story with Veeam

Milestone 12: Establish an Ongoing Hygiene Cadence

By Day 100, your environment should be stable and fully documented. These habits keep it that way:

- **Weekly:** Review the Veeam Kasten dashboard and investigate any failed or warned jobs before they accumulate.
- **Monthly:** Review recent Kasten releases and update regularly.
- **Quarterly:** Perform a documented recovery test and rotate through different workload types each quarter.
- **Annually:** Review your backup architecture against current business requirements.
- **Before renewal:** Review licensing usage in your Kasten Dashboard (e.g., total via MCM, or in each cluster individually) to confirm you are within entitlement.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.

Appendix: Quick Reference

Key Components: Veeam Integration

- **Veeam Software Appliance (VSA):** Pre-hardened appliance with Veeam Backup & Replication pre-installed. Deploy as OVA (VM) or ISO (physical). Can act as a limited immutable repository itself or works with other appliances or backup storage including dedicated VIAs.
- **Veeam Infrastructure Appliance (VIA):** Pre-hardened appliance deployed as a dedicated Hardened Repository or other role. Provides local immutability without Linux expertise. One role per appliance.
- **Veeam Backup & Replication:** Core backup engine, hosted on the VSA. Manages jobs, repositories, proxies, and recovery operations.
- **Scale-Out Backup Repository:** Logical construct combining a local Performance Tier (e.g. VIA) and Capacity Tier (e.g. Vault/S3/NFS) into a single backup target.
- **Veeam Vault:** Immutable cloud object storage for offsite copies. It's Veeam-managed, so no separate cloud account required.

Key Terms

- **RPO:** Maximum acceptable data loss, measured in time. Drives backup schedule frequency.
- **RTO:** Maximum acceptable downtime before a workload must be recovered.
- **Immutability:** Backup data that cannot be modified or deleted for a defined retention period. Protects against ransomware encryption of backup files.
- **Veeam Hardened Repository:** A Linux-based backup repository with immutability enforced at the OS level. Provided ready-to-use by VIA with no Linux administration required.
- **KubeVirt:** An open-source project allowing Kubernetes to run and manage VMs alongside containerized workloads. This enables VMs to run on Kubernetes platforms like OpenShift Virtualization and SUSE Harvester, which can be protected by Veeam Kasten.
- **Application-aware processing:** Processing that creates application-consistent backup points, which Veeam Kasten can do by utilizing Blueprints.
- **Blueprint:** Kanister Blueprints are an open-source mechanism to describe custom tasks for application-level data management on Kubernetes.

Useful Resources

- Veeam Help Center
- Veeam Customer Portal (downloads)
- Veeam Community Forums
- Veeam KB Articles
- Veeam Kasten Docs
- Kanister Blueprints Docs

About Veeam Software

Veeam is the Data and AI Trust Company, specializing in helping organizations ensure their data and AI are fully understood, secured, and resilient to enable the acceleration of safe AI at scale. As the market leader in both data resilience and data security posture management, Veeam is built for the convergence of identity, data, security, and AI risk.

Headquartered in Seattle, with offices in more than 30 countries, Veeam protects over 550,000 customers worldwide, including 82% of the Fortune 500,

Learn more at www.veeam.com or follow Veeam on LinkedIn @veeam-software and X @veeam.

© 2026 Veeam Software. All rights reserved. All trademarks are the property of their respective owners.