

Rancher 2.6

Kasten K10 by Veeam

Kubernetes-native backup/restore, disaster recovery, and application migration

Rancher 2.6 by SUSE

Adam Bergh, Solutions Architect, Cloud Native Technical Partnerships (Kasten by Veeam)

Terry Smith, Director, Global Partner Solutions (SUSE)

Gerson Guevara, IHV Solutions Architect (SUSE)



Kasten K10 by Veeam

Kubernetes-native backup/restore, disaster recovery, and application migration

Date: 2022-12-16

Summary

This document provides a brief introduction and demonstration of Kasten K10 by Veeam with Rancher by SUSE for enterprise cloud native backup/restore, disaster recovery, and app mobility.

SUSE One Partner Solution Stacks are featured, SUSE-confirmed co-innovations, collaboratively developed by SUSE and partners to arm enterprises with tools and agility to help overcome challenges and drive success.

Disclaimer

Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Contents

1	Introduction	4
2	Technical overview	4
3	Prerequisites	5
4	Installing Kasten K10	6
5	Accessing the Kasten K10 dashboard	10
6	Creating a location profile	15
7	Creating a policy	18
8	Working with policies	27
9	Summary	31
10	Legal notice	33
11	GNU Free Documentation License	34

1 Introduction

1.1 Motivation

Organizations are shifting to cloud native, leveraging containerized workloads and Kubernetes management platforms like Rancher by SUSE. The goal is to gain greater flexibility, scale, and resilience to accelerate innovation and quickly adjust to dynamic conditions. In this always-on IT environment, application mobility and data protection are critical considerations.



The Kasten K10 by Veeam® data management platform provides enterprise operations teams with an easy-to-use, scalable and secure system for backup and restore, disaster recovery, and mobility of cloud native applications.

1.2 Scope

This guide provides the steps to install and set up Kasten K10 by Veeam in your Rancher 2.6 by SUSE Kubernetes environment and an overview of application backup and restoration.

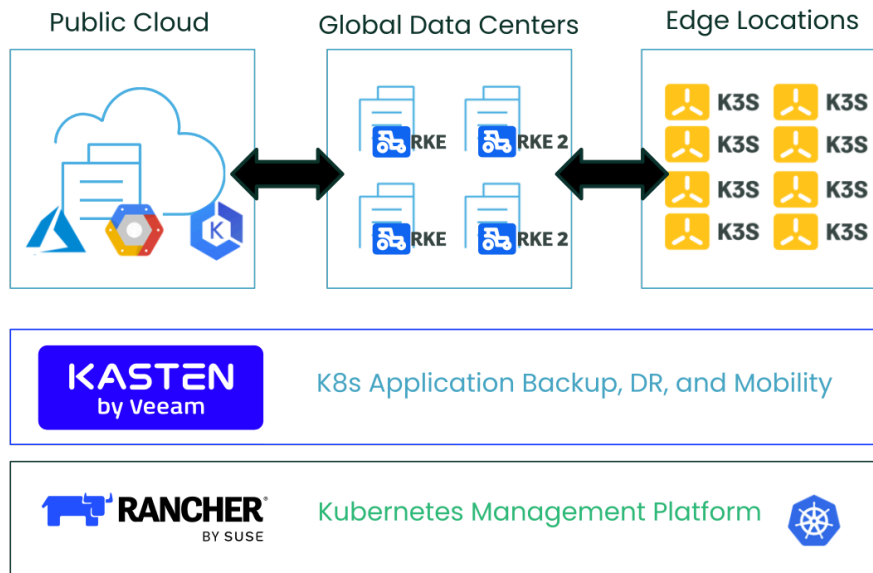
1.3 Audience

This document is intended for IT operations teams, backup administrators, DevOps and DevSecOps teams, and others who are responsible for ensuring business continuity, disaster recovery, ransomware and threat reduction, and application migration for cloud native landscapes.

2 Technical overview

The Kasten K10 by Veeam® data management platform has deep integrations with Rancher by SUSE and comes with an extensive ecosystem of support across Kubernetes distributions and cloud platforms. This gives enterprise operations teams the flexibility to choose the deployment

environments that best meet their needs - on-premises, public cloud, and hybrid. Kasten K10 is policy-driven and extensible. It delivers enterprise features such as full-spectrum consistency, database integration, automatic application discovery, multi-cloud mobility, and a powerful Web-based user interface.



3 Prerequisites

For this guide, you will need the following:

- Rancher by SUSE
In this guide, we use Rancher 2.6.
See [Rancher Installation and Upgrades guide](https://documentation.suse.com/cloudnative/rancher-manager/latest/en/installation-and-upgrade/installation-and-upgrade.html) (<https://documentation.suse.com/cloudnative/rancher-manager/latest/en/installation-and-upgrade/installation-and-upgrade.html>) for more information.
- Kubernetes cluster managed by Rancher
Any CNCF-certified Kubernetes cluster can be used.
See [Rancher Support Matrix](https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/rancher-v2-6-3/) (<https://www.suse.com/suse-rancher/support-matrix/all-supported-versions/rancher-v2-6-3/>)


- Storage for backup target

An external backup storage target, such as an NFS file server or cloud object store. This document uses an external, S3-compatible object storage bucket.

- User application to demonstrate backup and restore capability

For example, WordPress can be easily installed by [Helm chart \(https://bitnami.com/stack/wordpress/helm\)](https://bitnami.com/stack/wordpress/helm) .

Kasten K10 can be installed in a variety of different environments. To ensure a smooth installation experience, you can use the primer tool to perform several pre-flight checks to make sure that the prerequisites are met. This tool runs in a pod in the cluster and does the following:

- Validates that the Kubernetes settings meet the Kasten K10 requirements.
- Catalogs the available StorageClasses.
- If a CSI provisioner exists, it will also perform a basic validation of the cluster's CSI capabilities and any relevant objects that may be required. See [CSI pre-flight checks \(https://docs.kasten.io/latest/install/storage.html#csi-preflight\)](https://docs.kasten.io/latest/install/storage.html#csi-preflight)  in the documentation for more details.

Run the following command to deploy the primer tool:

```
curl https://docs.kasten.io/tools/k10_primer.sh | bash
```




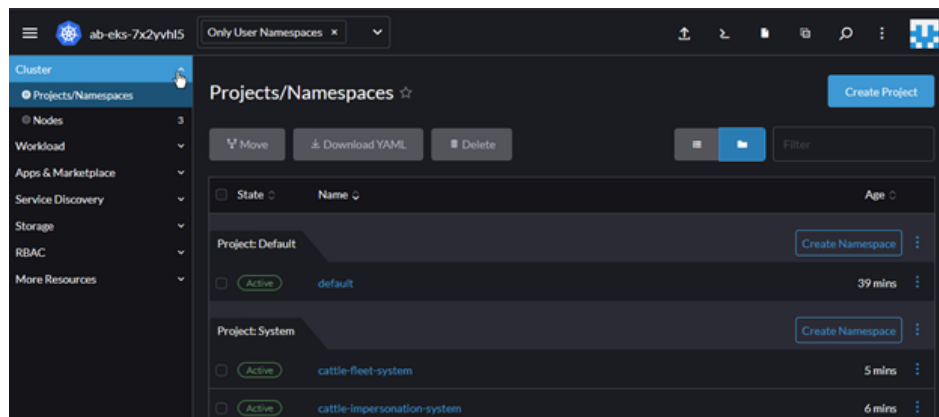
Note

This will create and clean up a ServiceAccount and ClusterRoleBinding to perform sanity checks on your Kubernetes cluster.

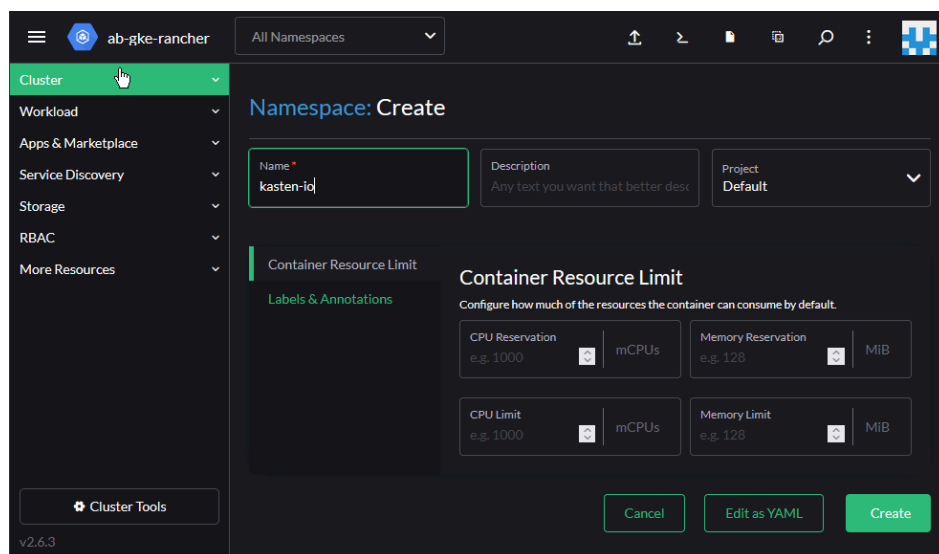
4 Installing Kasten K10

Kasten K10 can be easily deployed from the Rancher Apps Catalog.

1. Create a new [namespace \(https://documentation.suse.com/cloudnative/rancher-manager/latest/en/cluster-admin/namespaces.html\)](https://documentation.suse.com/cloudnative/rancher-manager/latest/en/cluster-admin/namespaces.html)  for the Kasten K10 application.
 - a. In the Rancher user interface (UI), navigate to *Clusters* → *Project/Namespace*.

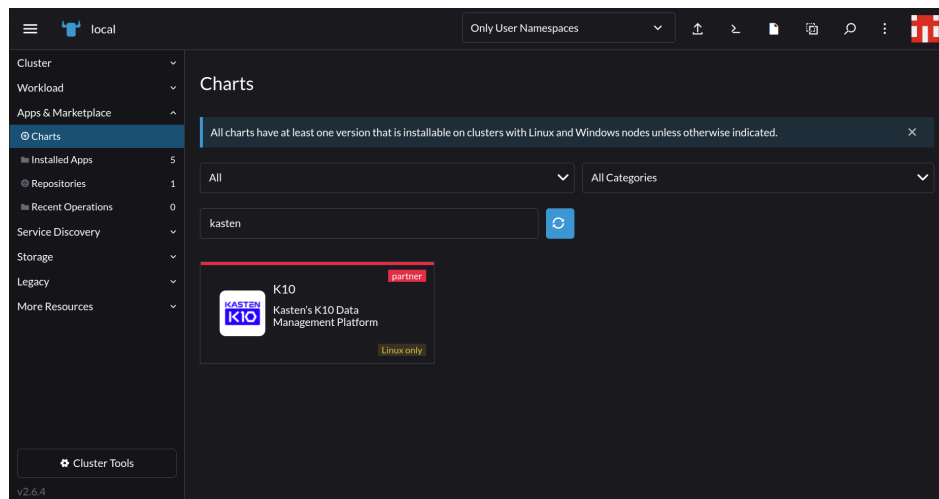


b. Create a "kasten-io" namespace for Kasten K10.

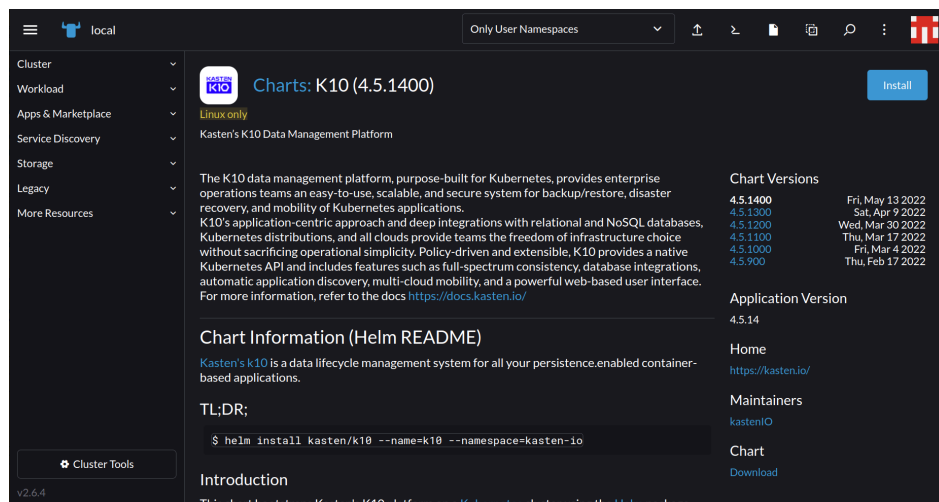


2. Install Kasten K10.

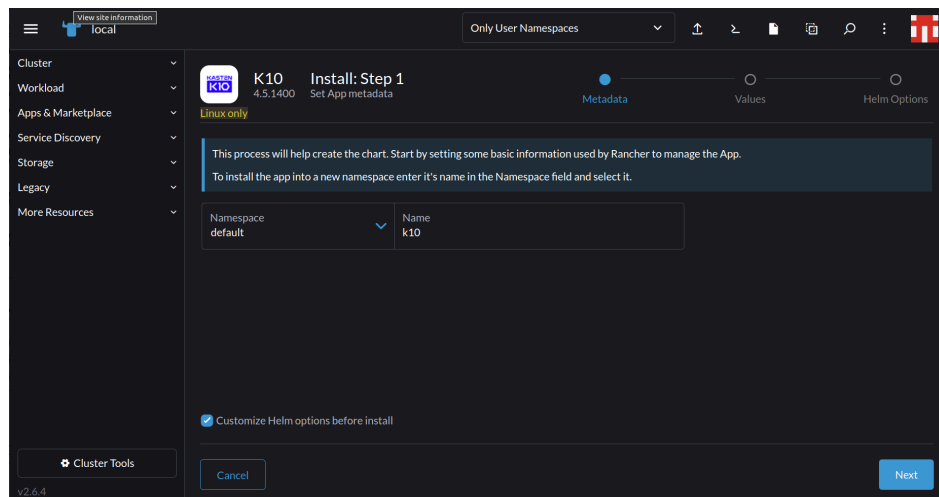
a. Navigate to *Apps & Marketplace* > *Charts* within the Rancher UI and search for “Kasten.”



b. Select the Kasten K10 chart and click *Install*.



c. Select the namespace "kasten-io" from the Namespace drop-down box. Optionally select *Customize Helm options before install* to customize the deployment. See the [Complete list of Helm options \(https://docs.kasten.io/latest/install/advanced.html#complete-list-of-k10-helm-options\)](https://docs.kasten.io/latest/install/advanced.html#complete-list-of-k10-helm-options) for detailed descriptions.



d. After setting your chart values, click *Next*, then click *Install*.

3. Validate the installation.

To validate that Kasten K10 has been properly installed, run the following command in the "kasten-io" namespace and watch for the status of the pods:

```
kubectl get pods --namespace kasten-io --watch
```

It may take a couple of minutes for all pods to come up and display "Running" status.

```
kubectl get pods --namespace kasten-io
```

NAMESPACE	NAME	READY	STATUS	RESTARTS
AGE				
kasten-io 1m26s	aggregatedapis-svc-b45d98bb5-w54pr	1/1	Running	0
kasten-io 1m26s	auth-svc-8549fc9c59-9c9fb	1/1	Running	0
kasten-io 1m26s	catalog-svc-f64666fdf-5t5tv	2/2	Running	0
...				



Note

In the unlikely scenario that pods are stuck in any other state, see the [support documentation \(https://docs.kasten.io/latest/operating/support.html#support\)](https://docs.kasten.io/latest/operating/support.html#support) to debug further.

5 Accessing the Kasten K10 dashboard

1. The Kasten K10 dashboard is not exposed externally by default. To establish a connection, use the following `kubectl` command:

```
kubectl --namespace kasten-io port-forward service/gateway 8080:8000
```

2. Open your Web browser to <http://127.0.0.1:8080/k10/#/>.



Note

If you are running on GKE and want to access the dashboard without local `kubectl` access, see [K10 Dashboard Directly From the Google Cloud Console \(https://docs.kasten.io/latest/access/gcp_details/gcp_console_dashboard.html\)](https://docs.kasten.io/latest/access/gcp_details/gcp_console_dashboard.html).

Direct access (<https://docs.kasten.io/latest/access/authentication.html#id5>) to the Kasten K10 dashboard requires a properly configured authentication method to secure access. For more information, see [Kubernetes authentication \(https://kubernetes.io/docs/reference/access-authn-authz/authentication/\)](https://kubernetes.io/docs/reference/access-authn-authz/authentication/). The next two sections provide an overview of the steps you can follow to configure an authentication method.

Proceed by following the steps to configure either **basic authentication** or **token authentication** in the next sections.

5.1 Basic authentication

Basic authentication (<https://docs.kasten.io/latest/access/authentication.html#id6>) allows you to protect access to the Kasten K10 dashboard with a user name and password.

Enable basic authentication by first generating `htpasswd` (<https://httpd.apache.org/docs/2.4/programs/htpasswd.html>) credentials using either an [online tool \(http://www.htaccesstools.com/htpasswd-generator/\)](http://www.htaccesstools.com/htpasswd-generator/) or via the `htpasswd` command found on most systems. When generated, you need to supply the resulting string with the `helm install` or `helm upgrade` command using the following flags:

```
--set auth.basicAuth.enabled=true \  
--set auth.basicAuth.htpasswd='example:$apr1$qrAVXu.v$Q8YVc50vtiS8KPmiyrkld0'
```

Alternatively, you can use an existing secret contained in a file created with `htpasswd`. The secret must be in the "kasten-io" namespace with the key named "auth" and the value as the password generated using `htpasswd`.

```
--set auth.basicAuth.enabled=true \  
--set auth.basicAuth.secretName=my-basic-auth-secret
```

5.2 Token authentication

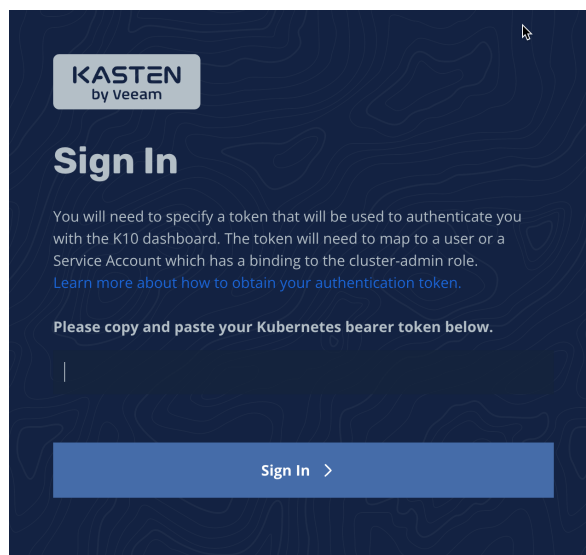
Token authentication (<https://docs.kasten.io/latest/access/authentication.html#id7>) allows the use of any token that can be verified by the Kubernetes server. For more information about token authentication, see:

- Obtaining Tokens (<https://docs.kasten.io/latest/access/authentication.html#id8>)
- Authentication Strategies (<https://kubernetes.io/docs/reference/access-authn-authz/authentication/#authentication-strategies>)

1. Enable token authentication by using the following flag as part of the initial `helm install` or subsequent `helm upgrade` command.

```
--set auth.tokenAuth.enabled=true
```

2. Next, provide a bearer token that will be used when accessing the dashboard.



The most common token type that you can use is a service account bearer token.

1. You can use `kubectl` to extract such a token from a service account that you know has the proper permissions.

- a. Get the SA secret

```
sa_secret=$(kubectl get serviceaccount my-kasten-sa -o  
jsonpath="{.secrets[0].name}" --namespace kasten-io)
```

- b. Extract the token

```
kubectl get secret $sa_secret --namespace kasten-io -ojsonpath="{.data.token}  
{'\n'}" | base64 --decode
```


2. Alternatively, you can create a new service account from which to extract the token.

```
kubectl create serviceaccount my-kasten-sa --namespace kasten-io
```

You can create a role binding or cluster role binding for the account to ensure that it has the appropriate permissions for Kasten K10. To learn more about permissions, see [Authorization \(https://docs.kasten.io/latest/access/authorization.html#authz\)](https://docs.kasten.io/latest/access/authorization.html#authz).

5.3 Kasten K10 dashboard overview

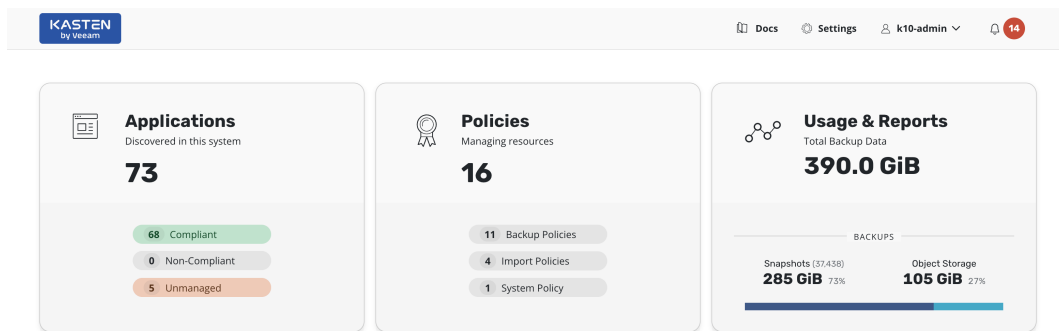
The Kasten K10 dashboard is divided into several different sections, described below.



Tip

A guided tour is available when the Kasten K10 dashboard is accessed for the first time or via the option on the [Settings \(https://docs.kasten.io/latest/usage/overview.html#k10-settings\)](https://docs.kasten.io/latest/usage/overview.html#k10-settings) page.

The top of the Kasten K10 dashboard displays a list of applications currently mapped to namespaces, any policies that might exist in the system, and a summary of the cluster's backup data footprint.



After filtering for applications that have stateful services (defined as containing a persistent volume), this screen further categorizes applications as:

- **Unmanaged:** There are no protection policies that cover this object.
- **Non-compliant:** A policy applies to this object, but the actions associated with the policy are failing (because of underlying storage slowness, configuration problems, etc.) or the actions have not been invoked yet.
- **Compliant:** These objects have policies, and the policy SLAs are being respected.



Tip

You can filter the view by clicking the *Compliant*, *Non-Compliant*, or *Unmanaged* buttons.

The Kasten K10 platform equates namespaces to applications for ease of use and consistency with Kubernetes best practices. This also allows use of role-based authentication controls (RBAC) and mirrors the most common application deployment patterns. As shown later, policies can be defined to operate on more than one namespace or only operate on a subset of an application residing in a single namespace.

If you already have installed applications, clicking the *Applications* card on the dashboard will provide you with details.

Applications

View details or perform actions on applications.

Filter by Status

Filter by Name

2 applications

default

Compliant With Policies

Latest snapshot was Today, 3:28pm

12 GB 4 4 3 7 2

snapshot restore export details

gitlab

Not Protected by Policies

Create a Policy

12 GB 4 4 3 7 2

snapshot restore export details

An application is made up of multiple Kubernetes resources and workloads, including deployments and stateful sets.

KASTEN by Veeam

< Clusters < Dashboard

Applications

View details or perform actions on applications.

Filter by Status Filter by Name

crawltest-8kcqv

Compliant With Policies

Latest snapshot was Today, 10:42am

1 GiB 4 4 5 7 2

snapshot restore export details

worldtest-1bjc9

Not Compliant With Policies

Latest snapshot was Today, 10:42am

Application Details

crawltest-8kcqv

LABELS

app:k10 chart:k10-21470.master.a23f9b727792ef03c331a4e32564ae7... heritage:Tiller release:k10 app:picture-gallery

Show 9 more labels ...

kubectl

\$ kubectl get --raw /apis/apps.k10.kasten.io/v1 copy

Data (4)

PVC

pic-gal-mysql

SIZE

1 MiB

...

PVC

pic-gal-php-cfg

SIZE

1 GiB

...

PVC

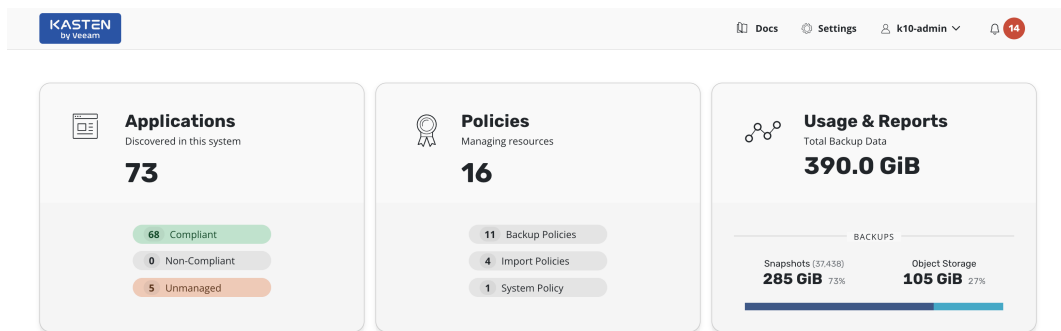
SIZE

...

Kasten K10 policies are used to automate your data management workflows. Policies combine actions you want to take (such as making a snapshot), frequency or schedule for how often you want to perform the action, and selection criteria for the resources you want to manage.

14

Kasten K10 by Veeam



On the *Policies* card, you notice that no default policies are created at install time. A policy can be either created from this page or from the application page shown earlier.

Policies

Policies are used to automate your data management workflows. To achieve this, they combine actions you want to take (e.g., snapshot), a frequency or schedule for how often you want to take that action, and a label-based selection criteria for the resources you want to manage.

[+ Create New Policy](#)

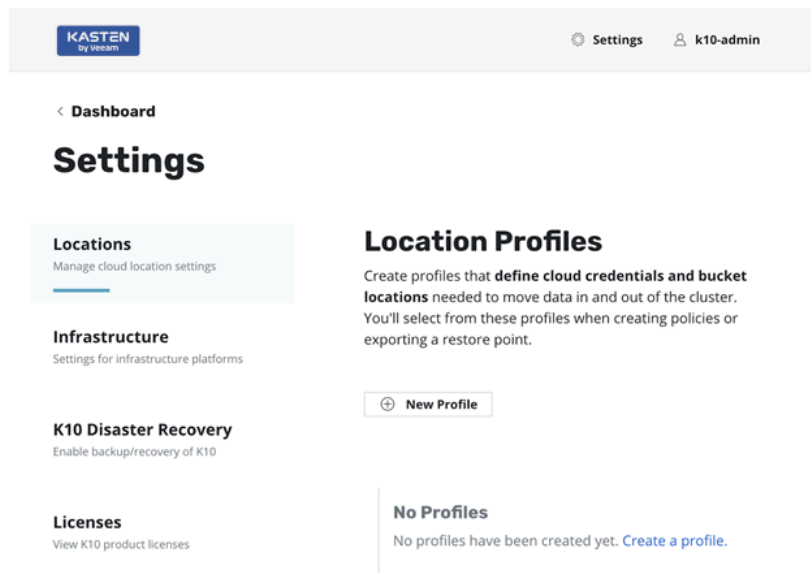
No Policies

No policies have been created yet. [Create your first policy.](#)

6 Creating a location profile

Kasten K10 can invoke protection operations, such as snapshots, within a cluster without requiring additional credentials. This may be sufficient if Kasten K10 is running in the major public clouds and actions are limited to a single cluster. It is not sufficient for most production situations, where performing real backups, enabling cross-cluster and cross-cloud application migration, and enabling disaster recovery are essential.

To enable these actions that span the lifetime of any one cluster, Kasten K10 needs to be configured with access to external object storage or external NFS file storage. This is accomplished with location profiles.



Access profile creation from the *Settings* icon in the top-right corner of the dashboard or via the [CRD-based Profiles API \(https://docs.kasten.io/latest/api/profiles.html#api-profile\)](https://docs.kasten.io/latest/api/profiles.html#api-profile). Location profiles are used to create backups from snapshots, move applications and their data across clusters and potentially across different clouds, and to subsequently import these backups into another cluster.

To create a location profile, click *New Profile* on the profiles page.

×

New Profile

Profile Name
Only lowercase letters, numbers, dash, and dot

Cloud Storage Provider

☐ Google Cloud Storage

☐ Amazon S3

☐ Azure Storage

☒ S3 Compatible

☐ NFS FileStore

☐ Veeam Backup Server

S3 Access Key

S3 Secret

👁

Endpoint
URL or domain of the S3 service API

☐ Skip certificate chain and hostname verification 🗨

Region
The geography in which the bucket is located

Bucket
If the bucket exists, ensure the region matches the bucket.

☐ **Enable Immutable Backups**
The bucket listed above must already exist and it must have *object locking* enabled. [More about Locked Bucket Setup...](#)

When exporting to or importing from an object storage location, you must pick an object storage provider, a region for the bucket if being used in a public cloud, and the bucket name. If a bucket with the given name does not exist, it will be created.

If you use an S3-compatible object storage system that is not hosted by one of the supported cloud providers, an S3 endpoint URL must be specified.



Note

When certain cloud providers (like AWS or Microsoft Azure) are selected, provider-specific options (such as IAM Roles) will appear for configuration.

When you click *Validate and Save*, the configuration profile is created and a profile similar to the following appears:

Settings

Locations

Manage location settings

Infrastructure

Settings for infrastructure platforms

Blueprints

Extensions for application-specific data management tasks

K10 Disaster Recovery

Enable backup/recovery of K10

Licenses

View K10 product licenses

User Roles

Manage User Access Privileges

Support

Cluster and contact info

Dashboard

Light/Dark mode, Guided Tour.

Location Profiles

Create profiles that **define credentials and locations** needed to move data in and out of the cluster. You'll select from these profiles when creating policies or exporting a restore point.

[+ New Profile](#)

LOCATION PROFILE

my-location-profile

revalidate yami edit delete

CLOUD PROVIDER	REGION	BUCKET NAME
AWS S3	US West (Oregon) • us-west-2	test.kasten.io
STATUS		
Valid		

7 Creating a policy

Protecting an application with Kasten K10 is usually accomplished by creating a policy.

In this section, you learn about snapshots and backups, scheduling, and selection in the context of application protection policies.

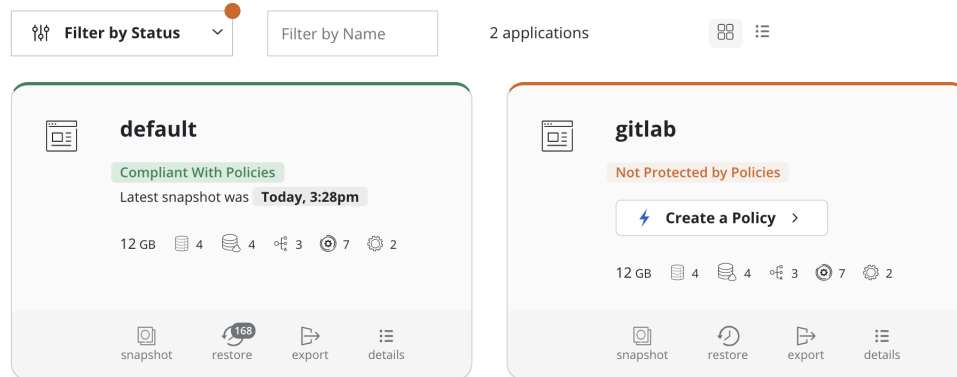
Kasten K10 defines an application as a collection of namespaced Kubernetes resources associated with

- a workload (such as ConfigMaps and Secrets),
- relevant non-namespaced resources used by the application (such as StorageClasses),
- Kubernetes workloads (including Deployments, StatefulSets, standalone pods, etc.),
- deployment and release information available from Helm v3,
- and all persistent storage resources (such as PersistentVolumeClaims and PersistentVolumes).

While you can always create a policy from scratch through the policies page, the easiest way to define policies for unprotected applications is to click the *Applications* card in the main dashboard. This will allow you to see all applications in your Kubernetes cluster.

Applications

View details or perform actions on applications.

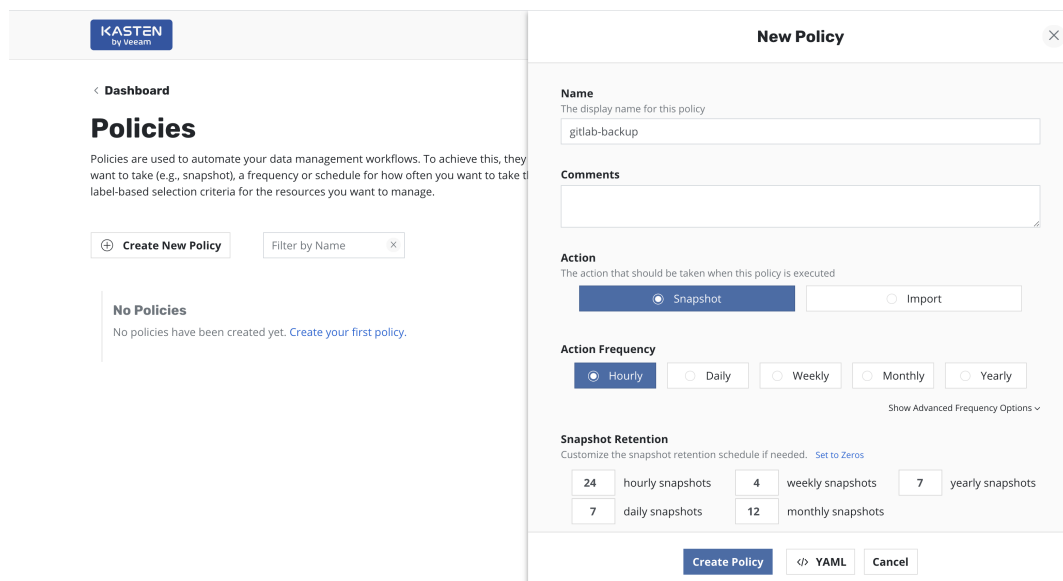


Filter by Status Filter by Name 2 applications

default
Compliant With Policies
Latest snapshot was Today, 3:28pm
12 GB 4 4 3 7 2
snapshot restore export details

gitlab
Not Protected by Policies
Create a Policy >
12 GB 4 4 3 7 2
snapshot restore export details

To protect any unmanaged application, simply click *Create a Policy* to open the *New Policy* dialog.



KASTEN by Veeam

< Dashboard

Policies

Policies are used to automate your data management workflows. To achieve this, they want to take (e.g., snapshot), a frequency or schedule for how often you want to take the label-based selection criteria for the resources you want to manage.

Create New Policy Filter by Name

No Policies
No policies have been created yet. [Create your first policy.](#)

New Policy

Name
The display name for this policy
gitlab-backup

Comments

Action
The action that should be taken when this policy is executed
Snapshot Import

Action Frequency
Hourly Daily Weekly Monthly Yearly
Show Advanced Frequency Options

Snapshot Retention
Customize the snapshot retention schedule if needed. [Set to Zeros](#)

24 hourly snapshots 4 weekly snapshots 7 yearly snapshots
7 daily snapshots 12 monthly snapshots

Create Policy YAML Cancel

7.1 Snapshots and backups

All Kasten K10 policies center around the execution of actions. You start by selecting the snapshot action with an optional backup (called an **export**).

See [Snapshots and Backups \(https://docs.kasten.io/latest/usage/protect.html#snapshots-and-backups\)](https://docs.kasten.io/latest/usage/protect.html#snapshots-and-backups) in the Kasten documentation for more details.

7.1.1 Snapshots

Snapshots form the basis of persistent data capture in Kasten K10. They are typically used in the context of disk volumes (PVC/PVs) used by the application but can also apply to application-level data capture (such as with [Kanister \(https://docs.kasten.io/latest/kanister/kanister.html#kanister\)](https://docs.kasten.io/latest/kanister/kanister.html#kanister)).

New Policy ✕

Name
The display name for this policy

Comments

Action
The action that should be taken when this policy is executed
☒ Snapshot ☐ Import



Note

Several public cloud providers (including AWS, Azure, and Google Cloud) actually store snapshots in object storage, and they are retained independent of the lifecycle of the primary volume. However, this is not true of all public clouds. An independent backup would provide essential safety. Check with your cloud provider's documentation for more information.

Snapshots, in most storage systems, are very efficient and have a very low performance impact on the primary workload, require no downtime, support fast restore times, and enable incremental data capture.

Storage snapshots usually suffer from constraints, such as having relatively low limits on the maximum number of snapshots per volume or per storage array. Most importantly, snapshots are not always durable. A catastrophic storage system failure will destroy your snapshots along with your primary data. Further, in several storage systems, a snapshot's lifecycle is tied to the source volume. So, if the volume is deleted, all related snapshots might automatically be garbage collected at the same time.



Tip

It is highly recommended that you create backups of your application snapshots to ensure durability.

7.1.2 Backups

Backups overcome the limitations of application and volume snapshots by converting them to an infrastructure-independent format, deduplicating, compressing, and encrypting them before they are stored in an external object store or NFS volume.

To convert your snapshots into backups, activate *Enable Backups via Snapshot Exports* during policy creation.

The screenshot shows a configuration panel titled "Enable Backups via Snapshot Exports" with a toggle switch turned on. Below the title is a subtitle: "After snapshot completes, export restore points to enable backups or cross-cluster migration." The panel contains several sections:

- Frequency:** A dropdown menu set to "Every snapshot".
- Export Location Profile:** A section with the subtitle "The cloud location that restore points will be exported to" and a dropdown menu set to "test-profile".
- Retention of Exported Snapshots:** A section with the subtitle "Manage how many exported snapshots to retain" and a dropdown menu set to "Use the same retention schedule as above".
- Snapshot Durability / Portability:** A section with two radio buttons: "Export Snapshot Data" (selected) and "Export Snapshot References Only".
- Checkmark and Note:** A checkmark icon followed by the text: "Exports complete snapshot data for durable and portable backups. Data is compressed, encrypted, and deduplicated, however this will use additional compute resources." with a speech bubble icon.

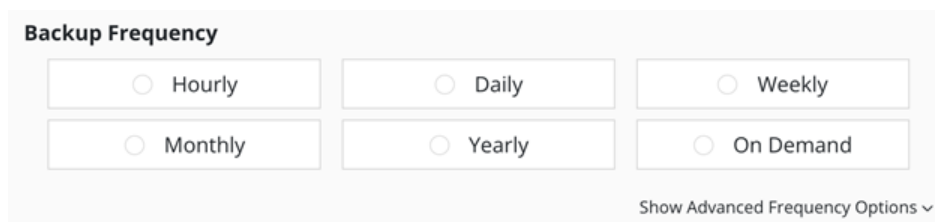
7.2 Scheduling

There are four components to Kasten K10 scheduling:

- Action frequency: how frequently the primary snapshot action should be performed
- Export frequency: how often snapshots should be exported to backups
- Retention schedule: how snapshots and backups are rotated and retained
- Timing: when the primary snapshot action should be performed

7.2.1 Action frequency

Kasten K10 snapshots can be set to execute at an hourly, daily, weekly, monthly, or yearly frequency, or on demand. By default, hourly snapshots execute at the top of the hour, while weekly, monthly, and yearly snapshots execute at midnight UTC.



The screenshot shows a configuration panel titled "Backup Frequency". It contains six radio button options arranged in two rows of three. The first row includes "Hourly", "Daily", and "Weekly". The second row includes "Monthly", "Yearly", and "On Demand". All radio buttons are currently unselected. At the bottom right of the panel, there is a link that says "Show Advanced Frequency Options" followed by a downward-pointing chevron icon.

You can also specify the time at which scheduled actions execute and sub-frequencies that execute multiple actions per frequency. Sub-hourly actions can be useful when you are protecting mostly Kubernetes objects or small data sets. See [Advanced Schedule Options \(https://docs.kasten.io/latest/usage/protect.html#advanced-schedule-options\)](https://docs.kasten.io/latest/usage/protect.html#advanced-schedule-options) in the Kasten documentation for more information.

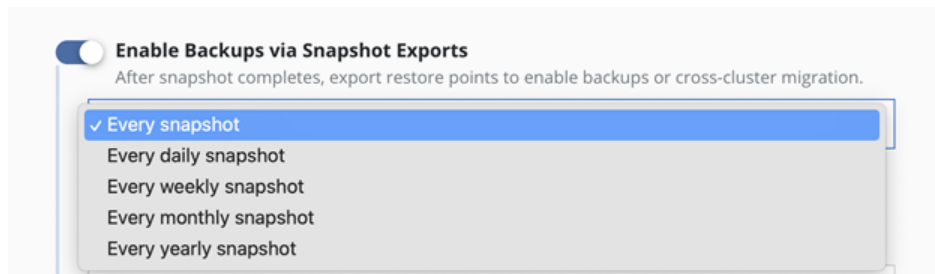


Warning

Care should be taken not to stress the underlying storage infrastructure or running into storage API rate limits. Further, sub-frequencies do also interact with retention (described below). For example, retaining 24 hourly snapshots at 15-minute intervals would only retain 6 hours of snapshots.


7.2.2 Export frequency

When *Enable Backups via Snapshot Exports* is enabled, snapshots are exported as backups. By default, every snapshot is exported, but you can limit this to a subset by selecting a daily, weekly, monthly, or yearly export frequency.



7.2.3 Retention schedule

A powerful scheduling feature in Kasten K10 is the ability to use a https://en.wikipedia.org/wiki/Backup_rotation_scheme#Grandfather-father-son [GFS retention scheme] for cost savings and compliance. With this backup rotation scheme, hourly snapshots and backups are rotated each hour with one graduating to daily every day, daily snapshots and backups are rotated each day with one graduating to weekly, and so on. You can set the number of hourly, daily, weekly, monthly, and yearly copies that need to be retained, and Kasten K10 will take care of cleanup.



Note

It is not possible to set a retention schedule for on-demand policies.

The default backup retention schedule is the same as the snapshot retention schedule. You can make these independent schedules, if needed. This allows you to create policies where a limited number of snapshots are retained for fast recovery from accidental outages while a larger number of backups are stored for long-term recovery. This separate retention schedule is also valuable when a limited number of snapshots are supported on the volume, but a larger backup retention count is needed for compliance.

Retention of Exported Snapshots

Use the same retention schedule as above

✓ Custom retention ...

24 hourly snapshots 4 weekly snapshots 7 yearly snapshots

7 daily snapshots 12 monthly snapshots

Snapshots and backups created by scheduled runs of a policy can be retained and omitted from the retention counts by adding a `k10.kasten.io/doNotRetire: "true"` label to the [RunAction](https://docs.kasten.io/latest/api/actions.html#api-run-action) (<https://docs.kasten.io/latest/api/actions.html#api-run-action>) created for the policy run.



Note

The retention schedule for a policy does not apply to snapshots and backups produced by [manual policy runs](https://docs.kasten.io/latest/usage/protect.html#manual-policy-runs) (<https://docs.kasten.io/latest/usage/protect.html#manual-policy-runs>). And you will need to clean up any artifacts created by manual policy runs.

7.2.4 Timing

By default, actions set to hourly execute at the top of the hour. Other action frequencies execute at midnight UTC.

Unhide *Advanced Options* to select how many times actions are executed within the frequency interval. For example, if the action frequency is daily, you can specify the hour of the day and the minutes after the hour when the action is to start.

Action Frequency

☐ Hourly ☒ Daily ☐ Weekly ☐ Monthly ☐ Yearly

Hide Advanced Options ^

Hour(s) of the Day for Daily Snapshots Local Time ☒ UTC [Reset](#)

Actions can be scheduled for one or more hours each day.

12am	1am	2am	3am	4am	5am	6am	7am	8am	9am	10am	11am
12pm	1pm	2pm	3pm	4pm	5pm	6pm	7pm	8pm	9pm	10pm	11pm

Minutes After the Hour: :00 ▼

> Snapshot at 12:00am UTC (5:00pm local) each day

Note: Times are stored in UTC, which does not change with Daylight Savings Time.

You can also customize the retention schedule by selecting which snapshots and backups will graduate and be retained for longer periods.

Snapshot Retention		
Customize the snapshot retention schedule if needed. Set to Zeros		
7	daily snapshots	
4	weekly snapshots	Sun 5pm used for weekly retention
12	monthly snapshots	01 used for monthly retention
7	yearly snapshots	Jan used for yearly retention



Tip

You can toggle whether to display and enter times in local time or UTC, but all times are converted to UTC and do not change for daylight savings time.

7.3 Selection

You can specify which applications are bound to a policy by name or label.

7.3.1 Application name

The most straightforward way to apply a policy to an application in Kasten K10 is to use its name (derived from the name of the namespace). You can even select multiple application names for the same policy.

If you need a policy to span similar applications, use the asterisk ('*') wild card. For example, if you specify 'mysql-*', Kasten K10 will match all applications whose names start with 'mysql-'.

Select Applications
Choose which application namespaces this policy should target. Select applications by name or by label.

☒ By Name ☐ By Labels ☐ None

Choose one or more applications to target with this policy. ⓘ

mysql-*

mysql-*



Note

For policies that need to span all applications, use the asterisk wild card by itself.

7.3.2 Application label

You can also use labels to bind a policy to multiple applications. For example, you could protect all applications that use MongoDB or applications that have been annotated with, say, the 'gold' label. Matching occurs on labels applied to namespaces, deployments, and StatefulSets. If multiple labels are selected, a union (logical OR) will be performed, matching all applications with at least one of the labels.

Label-based selection can be used to create forward-looking policies, as such policy would automatically apply to any future application with the matching label. For example, if you use a label of 'heritage: Tiller' (for Helm v2) or 'heritage: Helm' (for Helm v3), the selector will apply the policy to any new Helm-deployed applications because the label is applied to any Kubernetes workload created by the Helm package manager.

Select Applications
Choose which application namespaces this policy should target. Select applications by name or by label.

☐ By Name ☒ By Labels ☐ None

Use labels to target applications that you want to protect. Multiple labels will be unioned (OR).

heritage: Tiller x

7.3.3 Other resources

Kasten K10 can also protect cluster-scoped resources without targeting any applications. To specify this, select *None*.

Select Applications
Choose which application namespaces this policy should target. Select applications by name or by label.

☐ By Name ☐ By Labels ☒ None

☒ **Snapshot Cluster-Scoped Resources**
These include non-namespaced resources that are not captured in application snapshots, such as Custom Resource Definitions, ClusterRoles, and ClusterRoleBindings.

☒ All Cluster-Scoped Resources ☐ Filter Cluster-Scoped Resources

For more information about protecting cluster-scoped resources, see [Cluster-Scoped Resources \(https://docs.kasten.io/latest/usage/clusterscoped.html#clusterscoped\)](https://docs.kasten.io/latest/usage/clusterscoped.html#clusterscoped).

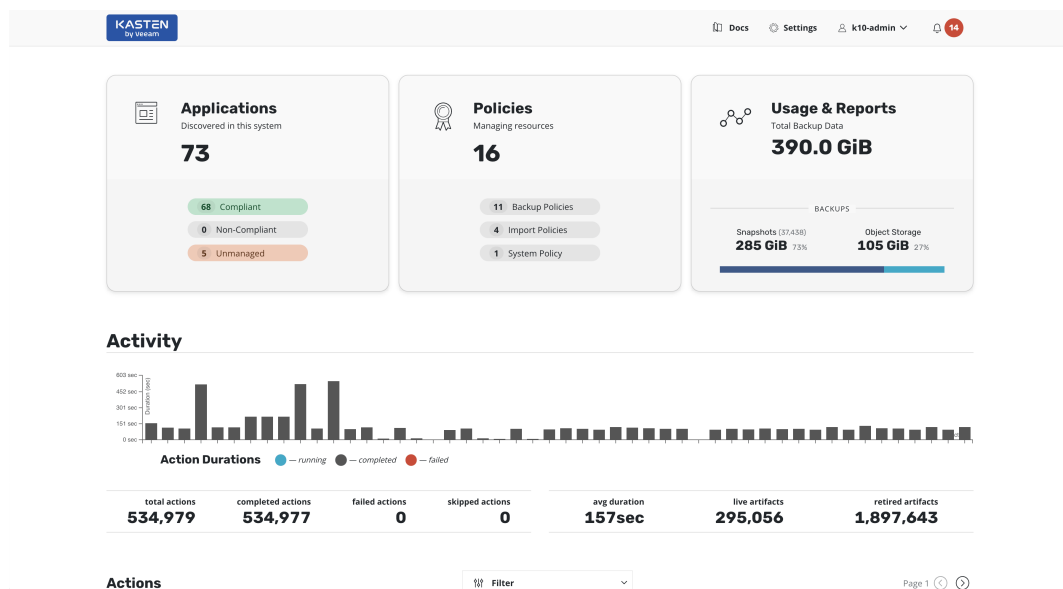
7.3.4 Customization

You can further customize what is and what is not covered under a Kasten K10 application protection policy with:

- namespace exclusions (<https://docs.kasten.io/latest/usage/protect.html#namespace-exclusion>)
- exceptions (<https://docs.kasten.io/latest/usage/protect.html#exceptions>)
- resource filtering (<https://docs.kasten.io/latest/usage/protect.html#resource-filtering>)

8 Working with policies

When you have created a policy and have navigated back to the main Kasten K10 dashboard, you see the selected applications quickly switch from unmanaged to non-compliant. That is, a policy covers the objects, but no action has been taken yet. The applications will switch to compliant as snapshots and backups are run and the application enters a protected state. You can also scroll down the page to see the activity, how long each snapshot took, and the generated artifacts.



Note

More detailed job information can be obtained by clicking the in-progress or completed jobs.

8.1 Manual policy runs

You can manually run a policy by clicking the *run once* button on the desired policy. Any artifacts created by this action will not be eligible for automatic retirement and will need to be manually cleaned up.

POLICY
mysql-backup

Valid

mysql

Snapshot *hourly* and retain
24 hourly snapshots
7 daily snapshots
4 weekly snapshots
12 monthly snapshots
7 yearly snapshots

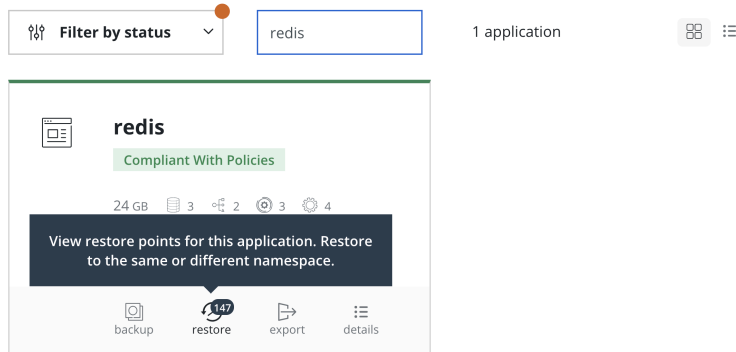
revalidate
edit
yaml
run once
pause
delete

8.2 Restoring existing applications

Restore an application via the *Applications* page in the Kasten K10 dashboard. To restore an application, simply click the *restore* icon in the application's card.

Applications

View details or perform actions on applications.



Note

While Kasten K10 uses the "export" term for backups, no import policy is needed to restore from a backup. Import policies are only needed when you want to restore the application into a different cluster.


Next, you select a restore point. These are distinguished in Kasten K10 as having been generated manually or automatically through a backup policy.

Restore application *redis*

Restore an application to a previous state. Restore points are shown and ordered based on scheduled execution time which may be different from the actual creation time. During a restore, the existing application is deleted and then recreated with the data artifacts restored from backups.

Select a restore point for details.

Past day



A restore point may include snapshots (native to the cluster) and backups (exported outside the cluster) with the same data. When both snapshots and backups are present, the Kasten K10 provides you with the option to select the instance you want to use to restore the application.

Restore application *redis*

Restore an application to a previous state. Restore points are shown and ordered based on scheduled execution time which may be different from the actual creation time. During a restore, the existing application is deleted and then recreated with the data artifacts restored from backups.

Select an instance...

Select a restore point for details.

This restore point has two instances - one that is native to the cluster and another that has been exported outside the cluster.

Past day

Selecting a restore point brings up a side-panel containing more details on the restore point for you to preview before you initiate an application restore.

When you click *Restore*, Kasten K10 automatically re-creates the entire application stack into the selected namespace. This not only includes data associated with the original application but also the versioned container images.



Note

Restored PersistentVolumes may not have the annotations specified in the original PersistentVolume.

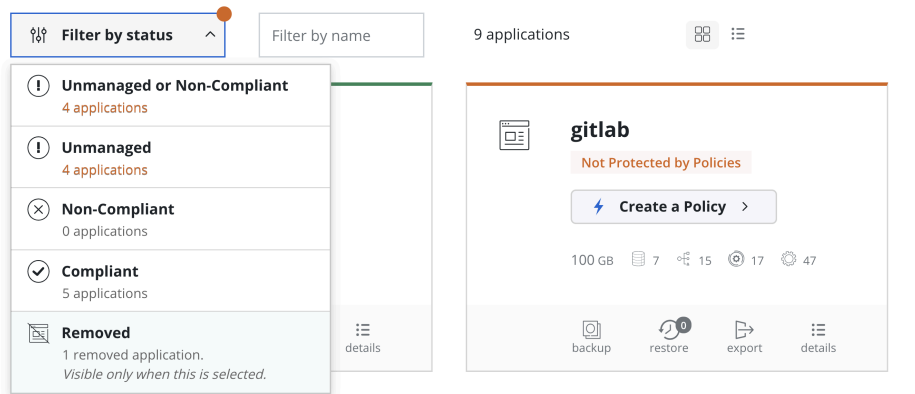
After the restore completes, you can go back to your application and verify that the state was restored to what existed at the time the restore point was obtained.

8.3 Restoring deleted applications

Restoring a deleted application follows nearly the same process, except that removed applications are not shown on the Applications page by default. To discover them, you simply need to filter and select *Removed*.

Applications

View details or perform actions on applications.



When the filter is in effect, you see applications that Kasten K10 had previously protected but which no longer exist. These can now be restored using the normal restore workflow.

9 Summary

Rancher by SUSE enables enterprises to streamline multi-cluster Kubernetes operations everywhere with unified security, policy and user management. Kasten K10 by Veeam delivers easy-to-use, Kubernetes-native application backup and restore, disaster recovery, and application mobility. Together, SUSE and Veeam provide enterprises with the tools they need to reduce risk and accelerate cloud native success.

In this guide, you learned how to seamlessly deploy Kasten K10 into your Rancher 2.6 by SUSE Kubernetes landscape, create policy-driven backups, and restore applications.

Continue your journey with these additional resources:


- "Cloud native workload protection with Kasten K10 by Veeam and SUSE Rancher" - demonstration video (https://youtu.be/c_mSNy6Q9RU) ↗
- "Kasten K10 by Veeam and SUSE Rancher: Enterprise K8s data protection" - blog article (<https://www.suse.com/c/kasten-k10-by-veeam-and-suse-rancher-enterprise-k8s-data-protection/>) ↗

- "Deploying Multicloud Day 2 Operations with SUSE Rancher, Fleet, and Kasten K10"
- blog article (<https://www.suse.com/c/deploying-multicloud-day-2-operations-with-suse-rancher-fleet-and-kasten-k10/>) ↗
- SUSE Rancher Prime product page (<https://www.suse.com/products/rancher/>) ↗
- Download Kasten K10 for free and use it for up to 5 nodes (<https://www.kasten.io/free-kubernetes>) ↗
- Kasten by Veeam Documentation (<https://docs.kasten.io/latest/index.html>) ↗
- SUSE Rancher Prime Documentation (<https://documentation.suse.com/cloudnative/rancher-manager/latest/en/about-rancher/what-is-rancher.html>) ↗

10 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/> .

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the series SUSE Technical Reference Documentation have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

11 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2

```
or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled "GNU  
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.