

# Red Hat OpenShift Data Foundation

Backup and recovery solutions for Red Hat OpenShift applications

## Highlights

Provide robust data protection for production container-based applications with Red Hat OpenShift and Red Hat OpenShift Data Foundation.

Choose proven data protection management solutions for consistent backup and recovery across diverse deployment environments.

Depend on open and standard technologies to help provide portability and version independence.

## Introduction

DevOps methodologies and emerging microservices architectures are dramatically changing how applications are developed and deployed. Containers and Kubernetes technologies offer dynamic scale and new levels of automation, benefiting both new applications and the migration of existing applications to hybrid cloud environments. These very advantages, however, can present challenges for backup and recovery in production environments. Traditional backup approaches assume that an application state is maintained at the server or virtual machine level. They simply cannot handle highly dynamic and automated Kubernetes environments, much less capture the context and state of cloud-based applications.

As a overall data resiliency, [Red Hat® OpenShift®](#) and [Red Hat OpenShift Data Foundation](#) let organizations extend their current enterprise backup and recovery solutions to production Kubernetes applications. Multiple data protection management partners now offer backup and recovery solutions for Red Hat OpenShift applications. With enhanced API-enabled backup applications, container-based applications gain granular, cluster-consistent backups that include cluster metadata and app data. Even complex Kubernetes applications can be backed up and restored within a cluster, or even moved or cloned to a separate cluster—while retaining all state and resource information.

## Kubernetes-native backup and recovery

Achieving critical business continuity and disaster recovery for production container-based applications requires comprehensive infrastructure support. At the orchestration platform level, Red Hat OpenShift provides application resilience, restarting application pods as necessary.<sup>1</sup> At the data services level, OpenShift Data Foundation adds extensive data protection, including:

- Files striped across objects and objects randomly distributed within a cluster.
- Integral replication or erasure coding that protects against device or storage server failures to automate recovery.
- A shared persistent storage layer that provides failover support across failure zones—either on-premise or in the cloud.

In addition to these data mechanisms, Red Hat OpenShift and OpenShift Data Foundation offer technology features for backup and recovery, including:

- **Snapshots** that provide point-in-time data copies to protect against logical failures.
- **Backups** (based on snapshots) that enable application restoration to a specific pre-failure state.



[www.facebook.com/redhatinc/](https://www.facebook.com/redhatinc/)  
[@RedHat](https://twitter.com/RedHat)  
[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

<sup>1</sup> Additionally, some applications provide their own application and data resilience.

Red Hat OpenShift Data Foundation is persistent software-defined storage optimized for Red Hat OpenShift Container Platform, offering data services for Red Hat OpenShift in hybrid and multicloud environments.

Unlike static environments, capturing the full context of a container-based application running in Kubernetes requires a multifaceted approach. Effective backup and recovery solutions must record more than just persistent volumes (PVs) in use by an application (Figure 1). Backup orchestration must first define all components of the application, including Kubernetes namespaces, resources, internal images, and persistent volume data. Resources must then be extracted and saved to a reliable backup target, either inside or outside the cluster. Restoration must then reliably reverse that process to a cluster of the organization's choosing. To support business continuity and disaster recovery requirements, application state must be preserved throughout the process.

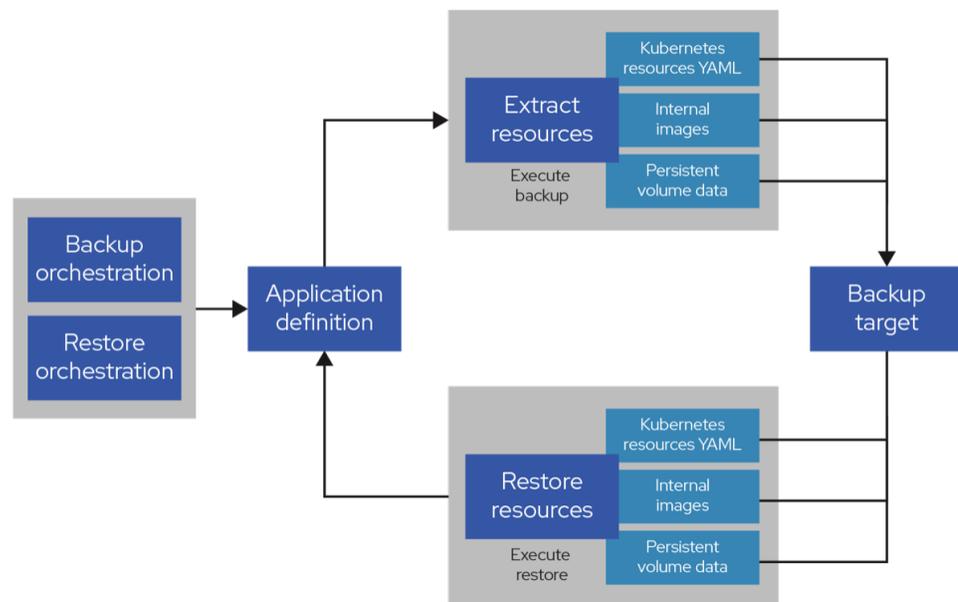


Figure 1. Backup and recovery workflow for production Kubernetes applications

Capturing the context for cluster resources and state requires API-enabled data protection management applications that can work with Red Hat OpenShift Container Platform and OpenShift Data Foundation. These API-enabled backup applications capture complete application data and metadata, providing application-granular and cluster-consistent backup and recovery (Figure 2).

- Data protection management applications handle backup policy management, backup scheduling, retention and restore management, and data movement. These applications provide full stack protection with application awareness.
- APIs collect cluster resources and metadata such as *Backup*, *Restore*, *Scheduling*, *BackupStorageLocation*, and *VolumeSnapshotLocation*. While these APIs are typically built into Kubernetes-native data protection applications, Red Hat provides an optional OpenShift APIs for data protection operator.<sup>2</sup> The data protection operator enables namespace or label-scoped backups with all cluster resources and application data (PVs).
- OpenShift Data Foundation provides efficient incremental PV snapshots via the Container Storage Interface (CSI). Backups of persistent volumes are performed through standard CSI interfaces. Ceph®-CSI snapshots can be used with or without OpenShift APIs for data protection.

<sup>2</sup> The OpenShift APIs for data protection Kubernetes operator supports the upstream [Velero API](#). The OpenShift APIs for data protection operator is represented in the upstream [Konveyor](#) upstream project.

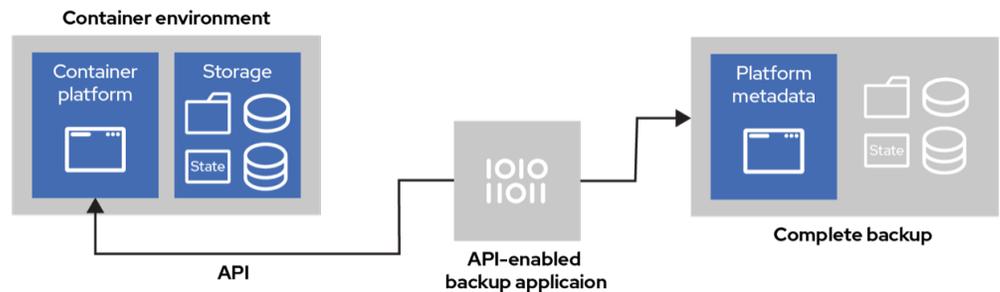


Figure 2. API-enabled backup applications provide application-granular and cluster-consistent backups.

## Data protection management solutions for Kubernetes

Organizations now have a choice of advanced data management solutions for backup and recovery. A number of popular choices are described in the sections that follow.

### IBM Spectrum Protect Plus

[IBM Spectrum Protect Plus](#) is a modern data protection solution that provides recovery, replication, retention, and reuse for virtual machines (VMs), databases, applications, file systems, Software-as-a-Service (SaaS) workloads, and containers in hybrid cloud environments. When coupled with Red Hat OpenShift and OpenShift Data Foundation, [the solution](#) is deployed as a virtual appliance or as a container application. With support for Kubernetes, IBM Spectrum Protect Plus lets organizations protect, back up, and restore complete containerized applications, including persistent volumes, namespaces, resources, and metadata for disaster recovery or development and testing operations.

### Kasten K10 by Veeam

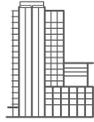
The [Kasten K10 data management platform](#) works with Red Hat OpenShift and OpenShift Data Foundation to offer backup, restore, and disaster recovery support. With a Red Hat-certified K10 operator available in the Red Hat OpenShift OperatorHub, organizations can back up and restore all the components of an application without time-consuming guesswork. [Combined with Red Hat OpenShift Data Foundation](#), Kasten K10 provides snapshots and cloning of persistent volumes and application metadata to one or more namespaces for backup or to multiple Red Hat OpenShift clusters for backup or disaster recovery.

### TrilioVault for Kubernetes

[TrilioVault for Kubernetes](#) protects application data and metadata—including pods, persistent volumes, secrets, configuration maps, and other Kubernetes resources. With Red Hat OpenShift operator certification, [TrilioVault for Kubernetes](#) protects both application data and metadata for Red Hat OpenShift applications using CSI snapshots. Namespace users can restore applications on demand. The intuitive, click-driven user interface provides discovery, monitoring, operations visibility, and management. Monitoring and metering is supported with [Prometheus](#) and [Grafana](#), with logging and tracing supported by [Fluentd](#).

### Bacula Enterprise

[Bacula Enterprise](#) is a broad enterprise backup and recovery solution that offers advanced and automated backup for Red Hat OpenShift applications. Bacula lets organizations protect container deployments and other IT infrastructure with a single backup and recovery solution. Ideal for hybrid cloud environments, Bacula Enterprise can help modernize your enterprise datacenter backup strategy, increase enterprise



data backup software efficiency, and significantly reduce costs. Native hybrid cloud integration includes support for Amazon Web Services (AWS) Simple Storage Service (S3), Microsoft Azure, Google Cloud, Oracle Cloud, and Glacier.

## Building in data resilience

Achieving stringent service-level objectives requires true data resilience. Red Hat realizes that business continuity and effective disaster recovery for Kubernetes applications requires a committed and comprehensive approach. Red Hat innovates and develops technologies that help organizations move toward zero or near-zero recovery time objectives (RTO) and recovery point objectives (RPO), allowing them to recover from:

- **Localized failures.** Backup and recovery using CSI snapshots and API-enabled backup and recovery applications is the first step in recovering from localized failures.
- **Regional disasters.** Asynchronous mirroring and site failover and failback automation can provide resilience from regional disasters.
- **Metro site failures.** Stretch clusters with arbitration can provide resilience and true-zero RPO and RTO in the face of larger-scale failures and disasters.

With many of the foundational elements already in place, these technology approaches and others remain an ongoing focus for providing even greater data resilience in upcoming versions of Red Hat OpenShift Data Foundation.

## Conclusion

Red Hat OpenShift Data Foundation, paired with third-party data protection management solutions, provides robust data services for production container-based applications running in Red Hat OpenShift. Open and standard technologies like Ceph-CSI snapshots help provide portability across clusters and clouds, as well as version independence, while giving developers and administrators tools for cloning, backing up, and restoring their apps and associated data. Multiple third-party backup and recovery solutions now support API-enabled applications running under Kubernetes. Organizations can use their favorite backup and recovery solution as they move to containers, gaining consistency across multiple deployment environments, whether physical, virtual, or in a hybrid cloud.

### About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry leading operating system, and automate, secure, and manage complex environments. Award winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

#### North America

1 888 REDHAT1

[www.redhat.com](http://www.redhat.com)

#### Europe, Middle East, and Africa

00800 7334 2835

[europe@redhat.com](mailto:europe@redhat.com)

#### Asia Pacific

+65 6490 4200

[apac@redhat.com](mailto:apac@redhat.com)

#### Latin America

+54 11 4329 7300

[info-latam@redhat.com](mailto:info-latam@redhat.com)



[facebook.com/redhatinc](https://facebook.com/redhatinc)

@Redhat

[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

redhat.com

O-F28056