

# Trilio for Kubernetes

Cloud-native data protection for container-based applications

## Highlights

Choose an integrated backup and recovery solution with full Red Hat OpenShift operator certification and Red Hat OpenShift Data Foundation interface compatibility.

Protect both application data and metadata for Red Hat OpenShift applications using container storage interface (CSI) snapshots.

Deploy all operators as a unified solution from the OperatorHub interface in OpenShift Container Platform.

Run container-based applications, provision OpenShift Data Foundation, and protect and manage data operations with Trilio—all through the Red Hat OpenShift console.

Administer with Kubectl commands or the custom Trilio user interface (UI) with support for Red Hat OpenShift dynamic forms.

## The importance of cloud-native data protection and management

Cloud-native platforms are increasingly vital, giving organizations new ways to develop, deliver, and manage applications and microservices. Kubernetes and the stateful container-based applications at the core of new hybrid cloud-based services have recently become critical, driving new needs for data protection and disaster recovery strategies in enterprise operations. Whether on-premise or in a cloud environment, the combination of [Trilio for Kubernetes](#), [Red Hat® OpenShift®](#), and [Red Hat OpenShift Data Foundation](#) provides data protection for container-based applications.

## Trilio for Kubernetes and Red Hat OpenShift

Organizations need flexibility and scalability to innovate rapidly in response to emerging business challenges and opportunities—while lowering their risk of service downtime or data loss. Trilio for Kubernetes is an intelligent recovery software that supports the scale, performance, and mobility requirements of container environments across any public, hybrid, or multicloud deployment. Coupled with Red Hat OpenShift Platform Plus, which includes OpenShift Data Foundation, Trilio for Kubernetes protects the entire application, including data, metadata, and all other associated Kubernetes resources. Once the application is protected, restoration can occur from any point-in-time snapshot with a single click of a button.

Trilio for Kubernetes protects application data and metadata—including pods, persistent volumes (PVs), secrets, configuration maps, and other Kubernetes resources. With Red Hat OpenShift operator certification, Trilio for Kubernetes is available in the OperatorHub, allowing simple installation and configuration. It uses the Red Hat OpenShift custom resource definition application programming interface (API) to provide a custom Trilio user interface (UI) for click-driven workflows. [Trilio for Kubernetes](#) is compatible with simple storage service (S3) or network file system (NFS) compatible storage like Amazon S3 object storage bucket, such as IBM Storage Ceph®, for use as a backup target.

IT managers, backup administrators, and application developers can use Trilio for Kubernetes with OpenShift Platform Plus to support a wide range of use cases:

- ▶ **Backup and recovery.** Schedule point-in-time backup and recovery and automate incremental backups with predefined policies, while selectively restoring containers and applications to the same or a new namespace.
- ▶ **Disaster recovery.** Trilio provides advanced DR capabilities that allows users to restore multiple applications or namespaces as part of a single workflow to recover business operations. The DR configuration is created and saved at the backup target level—which means the plan can be executed even when the source cluster is not available. There is no need to Backup the Backup Solution—everything you need to recover is on the backup target outside of the cluster.
- ▶ **Migration.** Trilio's Architecture automates and simplifies the process of migration from any OpenStack® or Kubernetes distribution to Red Hat OpenStack or Red Hat OpenShift by being able to run on all Kubernetes and OpenStack environments. E.g., between distributions, across clouds, from on-premise to a cloud environment. No more vendor lock in or fearing upgrades.

With Trilio for Kubernetes, namespace users can restore applications on demand. Control for Trilio operations can be delegated based on Kubernetes role based access control. The intuitive click-driven UI provides discovery, monitoring, and operations visibility and management capability. Monitoring and metering is supported with [Prometheus](#) and [Grafana](#) with logging and tracing supported by [Fluentd](#).

Trilio supports Red Hat OpenShift dynamic forms, providing an alternative to YAML files to execute Trilio operations and workflows through simple form-fill capabilities.

- ▶ **Application mobility.** Build multiple testing environments using Trilio backups from a single target location to develop and validate code, and push changes to production.

## Resilient data protection with OpenShift Data Foundation

Trilio natively provides application backup using container storage interface (CSI) snapshots to capture the state of a storage volume at a particular point in time. Cluster administrators use CSI snapshots to back up and restore PVs. App developers can use volume snapshots in a dev/test context and rapidly roll back to previous development versions when necessary. Together, Trilio and OpenShift Platform Plus provide:

- ▶ **An enterprise-ready platform with self-service management.** Trilio features native integration into Kubernetes and Red Hat OpenShift Platform Plus with OpenShift Data Foundation interface compatibility. Organizations can backup to NFS volumes or use any S3-based storage. Support is included for Kubernetes namespaces, operators, Helm, and label backup and restore. Trilio can be deployed via OperatorHub or Helm charts.
- ▶ **Non-disruptive backup.** Lightweight Kubernetes API and CSI snapshots capture and back up the entire application, including data, metadata, and other associated Kubernetes objects. The application supports linear, infinite-scale forever incremental backups with application-aware hooks for backup and restore.
- ▶ **One-click restore.** Organizations can recover entire applications from a particular point in time. Restore can take place to a new Red Hat OpenShift cluster or namespace. Trilio includes policy-based global job scheduling and data retention, with robust restoration policies.
- ▶ **Simplified provisioning, management, and maintenance.** Trilio works with OpenShift Data Foundation, with no extra configuration steps required. Both technologies are deployed as Kubernetes operators within Red Hat OpenShift, providing a simple way to package, deploy, and manage environments. Administrators can automate container backup and storage tasks to enhance efficiency and resiliency.




## Conclusion

Together, Trilio for Kubernetes on Red Hat OpenShift and Red Hat OpenShift Data Foundation provide an integrated platform that supports running container based workloads, supporting those workloads with enterprise-grade persistent storage, and providing those workloads with sophisticated backup and recovery functionality for PVs and complete applications. Organizations can run container-based applications, provision OpenShift Data Foundation, and protect and manage data protection operations with Trilio.



### About Red Hat

Red Hat helps customers standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with [award-winning](#) support, training, and consulting services.

 facebook.com/redhatinc  
 @RedHat  
 linkedin.com/company/red-hat

**North America**  
1 888 REDHAT1  
www.redhat.com

**Europe, Middle East,  
and Africa**  
00800 7334 2835  
europe@redhat.com

**Asia Pacific**  
+65 6490 4200  
apac@redhat.com

**Latin America**  
+54 11 4329 7300  
info-latam@redhat.com