

Optimizing Backup Strategies

Trilio's Integrated Approach to
OpenShift Virtualization
Environment(s)



**Certified for
Red Hat OpenShift**

This outline provides a framework to better manage your backup strategy and Disaster Recovery strategy of your OpenShift Virtualization clusters with Trilio. It addresses the various deployment and management options, and helps you choose the most suitable approach for different scenarios.

At the end of this document, we also provide “best practices” that will ensure proper backup management, and the ability to leverage Trilio’s advanced features to amplify your experience in protecting OpenShift applications and Virtual Machines.

TABLE OF CONTENTS

Introduction	4
Things to Consider First.....	5
Managed vs. Self-Managed	5
On-Prem vs. Cloud.....	5
Core Components	5
Protection Strategy	5
Advanced Use Cases	5
Operational Options for Your Backup Strategy	6
OpenShift UI Driven Configuration	6
Trilio UI Driven Configuration	7
Command Line Interface Based Configuration and GitOps Integration	8
Ansible Playbook Automation.....	9
Red Hat Advanced Cluster Management Governance and Policy Integration	10
Choosing the Right Option	11
Factors to Consider:.....	11
Decision Matrix.....	11
Best Practices	12
Conclusion.....	12
Protecting Your OpenShift Investment with Confidence	12
Key Takeaways	13
Integrations with the Ecosystem	13
Moving Forward	13
Additional Considerations	14


Disaster Recovery.....14

About Trilio.....16

Introduction

Purpose-built for Red Hat [OpenShift](#) and [OpenShift Virtualization](#), Trilio protects cloud-native applications in [multi-cluster](#) private, public and hybrid-cloud environments. Trilio believes that recovery and application portability should be effortless, automated, and predictable. As a Backup tab in the OpenShift User Interface and coupled with Trilio's Certified [Ansible Collections](#) & [RHACM Policies](#), provide zero-touch operational resilience. Recover or migrate in place, or into other clusters and clouds.

Regardless of how an OpenShift cluster/s is/are deployed (managed service, public cloud, on-prem infrastructure, etc.), Trilio interacts with the cluster's Kubernetes API for backup and restore functionalities and is always installed as an OLM operator. See below:



Trilio for Kubernetes

4.0.2 provided by Trilio

[Install](#)

Latest version

4.0.2

Capability level

- ☒ Basic Install
- ☒ Seamless Upgrades
- ☒ Full Lifecycle
- ☐ Deep Insights
- ☐ Auto Pilot

Source

Certified

Provider

Trilio

Trilio for Red Hat OpenShift provides fast and easy backup of application containers and virtual machines, disaster and ransomware recovery and application mobility across any cluster, anywhere. Trilio for Red Hat OpenShift:

- Red Hat Certified - first backup and recovery solution with OpenShift Operator Certification.
- Application-Centric - protects both data and metadata for Helm or Label based applications.
- Namespace protection: Protect single or multi-namespace protection for complete cluster recovery
- Fully multi-tenant tool, adhering to OpenShift RBAC policies, allowing you to operate hybrid, on-prem, air-gapped, public cloud (ROSA or ARO) or OpenShift on public cloud instances.
- OpenShift Virtualization - protects data and metadata of OpenShift Virtualization Virtual Machines
- Operators - protect and recover Red Hat OpenShift Lifecycle Management (OLM) Operators
- Native to Kubernetes & OpenShift: Packaged and deployed as an Operator, provides new CRDs, integrated via Kubernetes API and with all features that it provides.
- Infrastructure Agnostic: Compatible with any CSI Persistent Volume storage; backup storage based on NFS, S3, Blob/Container; and run on any Cloud (Private or Public).

Note - Trilio backup configurations can be handled differently, depending on the tool/tools you choose/need to manage the backup strategy. It does not depend on the OpenShift deployment model.

Things to Consider First

There might be indirect considerations based on the deployment model:

Managed vs. Self-Managed

- In a managed service, some configuration options *might* be limited or pre-configured by the provider. This could influence your choice between UI or CLI for basic adjustments.
- Self-managed deployments offer more control, making Ansible, Red Hat Advanced Cluster Management (ACM), Ansible Automation Platform (AAP) or even DevOps pipelines potentially more relevant for complex setups.

On-Prem vs. Cloud.

Both on-prem and cloud deployments should prioritize strong access controls.

- On-prem environments might have a perception of greater physical control, but robust access management is still crucial.
- Cloud deployments rely on provider security measures, which usually are outstanding, but you still need to configure access controls within your environment for Trilio. You can't completely rely on the public cloud provider to do everything for you.

Trilio will store and manage the lifecycle of your backups. This makes it a critical component for data protection and disaster recovery. It's essential to control access to Trilio to prevent unauthorized modifications or deletions of backups. You should carefully plan your RBAC configuration to follow the "least privilege" methodology to enhance your security posture.

Core Components

- OpenShift Cluster(s)
- Trilio Deployment
- Backup Target(s) - NFS or S3

Protection Strategy

- Backup Types: Virtual Machine(s), namespace, multi namespace, Operator, Helm Charts, Applications
- Backup Scheduling: Full/Incrementals, Frequency-based, Retention Policies
- Recovery Options: Granular vs. Full Application Restore, Transforms to Migrate/Disaster Recovery applications between dissimilar Environments / Architectures / Distributions / Clouds / Storage

Advanced Use Cases

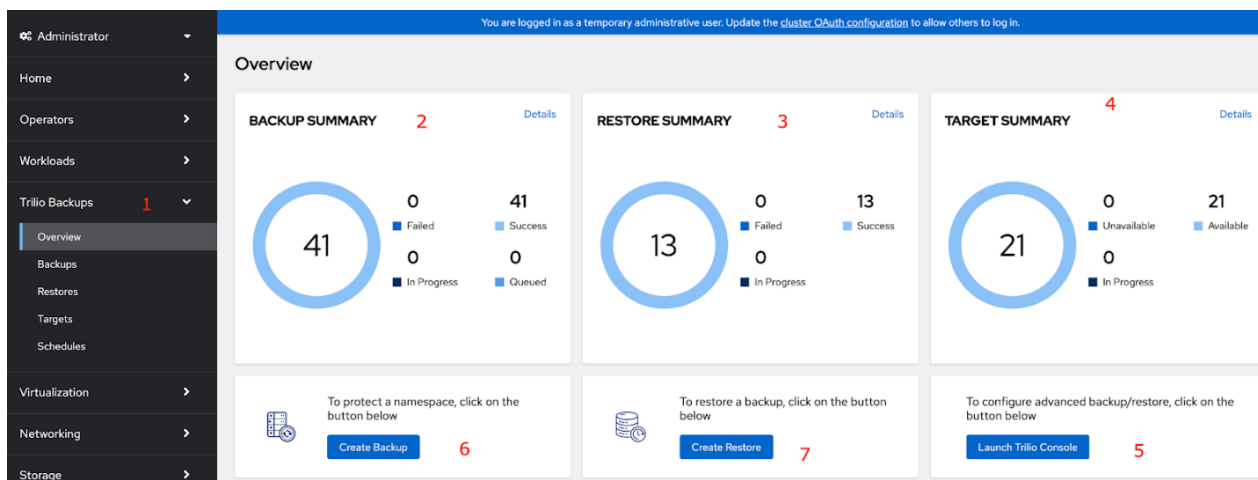
- Encryption
- Ransomware Protection

- Continuous Restore
- Migration
 - Application Mobility
 - Disaster Recovery (including DR periodic testing)
 - Database Consistent Backups

Operational Options for Your Backup Strategy

OpenShift UI Driven Configuration

Use Cases	<ul style="list-style-type: none"> • Initial configuration of Trilio in a single OpenShift cluster. • Simple deployments with basic backup schedules. • User-friendly interface for testing and exploration.
Benefits	<ul style="list-style-type: none"> • OpenShift-Centric View: Tailored specifically for OpenShift environments, potentially offering a more intuitive experience for those focused on OpenShift. • Streamlined Workflow: Simplifies VM protection within OpenShift, reducing context switching and simplifying tasks for OpenShift administrators.
Limitations	<ul style="list-style-type: none"> • Not scalable for managing complex deployments or multi-cluster environments. • Lacks automation capabilities for repetitive tasks (still it can handle scheduling and retention policies).



- 1 Trilio OpenShift Console Plugin
- 2 Backup Summary
- 3 Restores Summary
- 4 Targets Summary

- 5 Open Trilio Advanced Dashboard
- 6 Create backup Wizard
- 7 Create restore Wizard

Trilio UI Driven Configuration

Use Cases	<ul style="list-style-type: none">• Multi and hybrid cluster configuration• Complex configurations
Benefits	<ul style="list-style-type: none">• Centralized Management: Provides a single pane of glass for managing backups and recoveries across all your Kubernetes clusters, including OpenShift, giving you a holistic view.• Advanced Features: May offer additional capabilities not available in the OpenShift Trilio UI, depending on the specific Trilio version.• Easy cloning of objects between clusters.• Easy visual tools for migrations, DR, Continuous Restore.
Limitations	<ul style="list-style-type: none">• Potential Information Overload: If you only manage VMs in OpenShift, the native Trilio UI might present a broader view than you necessarily need.

CREATE NEW BACKUP - APPLICATION

1 Select Application

2 Select Backup Plan

3 Add Backup Name

4 Status Logs

25%

Progress

Type

LabelsHelm & OperatorsObjectsVirtual Machines

Search Virtual Machines

Namespace: All

Choose Virtual Machines to include0 Selected | Select all

ubuntu-demo (demo-ubuntu-migration)

rhel9-rbd (my-vms)

testvm (my-vms)

centos-stream9-nokigrig26qpi455 (rbd-vms)

centos-stream9-w0cfg2lbwpybys5r (rbd-vms)

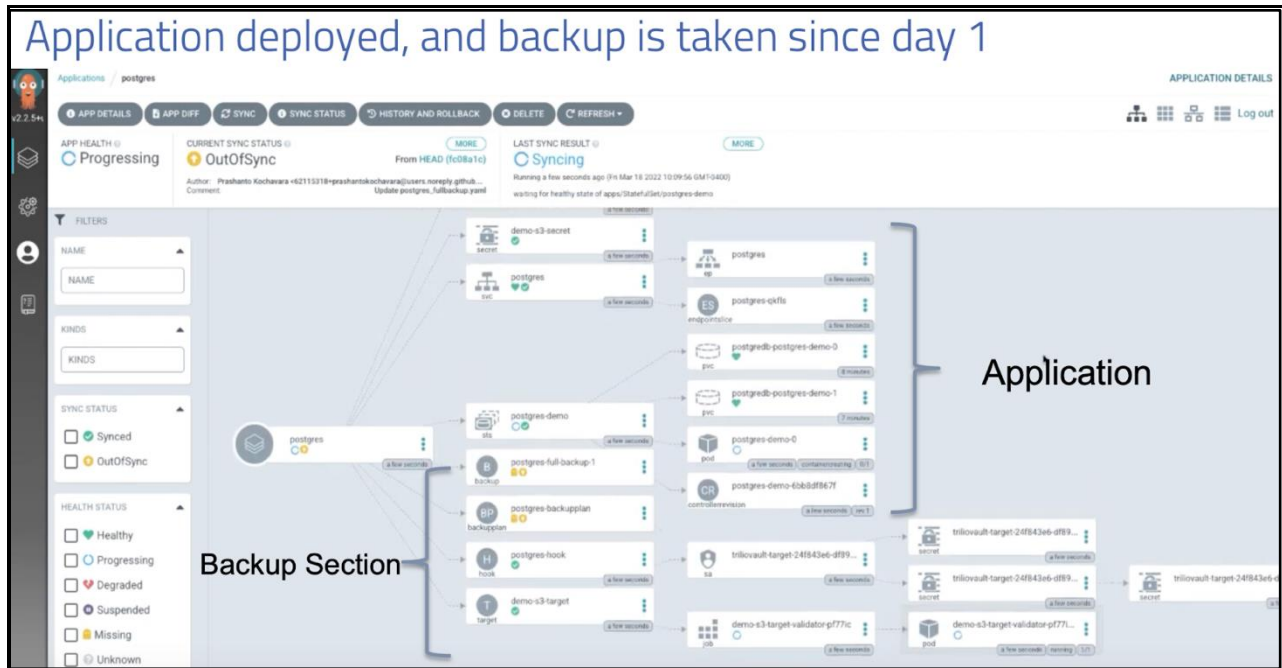
fedora-02 (rodolfo-tyson)

rhel9-rodolfo (rodolfo-vm)

Next

Command Line Interface Based Configuration and GitOps Integration

Use Cases	<ul style="list-style-type: none">• Automation for repeatable backup and restore workflows.• Integration with CI/CD pipelines for automated deployments.• Granular control over Trilio configuration for advanced users.• Within your existing application/Virtual Machine(s) deployment pipeline, add steps to configure a Trilio backup strategy for the new application.
Benefits	<ul style="list-style-type: none">• Reusability and version control for deployment automation.• Declarative approach defines desired state, simplifies configuration changes.• Automated Backups: Backups are automatically triggered as part of the deployment pipeline, eliminating the need for manual intervention. This ensures consistent and timely backups for newly deployed applications or updates.• Reduced Errors: Automating backup tasks minimizes the risk of human errors that can occur during manual backup processes.• Improved Disaster Recovery: Faster recovery times in case of incidents with readily available backups automatically created alongside deployments.• Streamlined Workflows: The entire deployment lifecycle, from code building to application deployment and backup creation, is managed within a single pipeline.
Limitations	<ul style="list-style-type: none">• Requires scripting/devops knowledge and familiarity with Trilio CLI commands.• Can become complex for managing intricate backup strategies.
Considerations	<ul style="list-style-type: none">• Secret Management: Securely store sensitive credentials within your CI/CD pipelines using tools like Hashicorp Vault. This ensures proper access control and prevents unauthorized access to backup functionalities.• Error Handling: Implement proper error handling and retry mechanisms within your pipelines. This ensures that even if a deployment fails, the associated backup creation process is retried to avoid missing backups for critical applications.• Testing: Integrate testing steps within your pipeline to verify successful deployment of the application and the associated Trilio backup strategy.



Ansible Playbook Automation

Trilio Ansible Galaxy and Automation Roles Collection ([Link Here](#))

Use Cases	<ul style="list-style-type: none"> • Complex deployments with multi-cluster support. • Standardized configuration management across multiple OpenShift clusters. • Reusability and version control for deployment automation.
Benefits	<ul style="list-style-type: none"> • Declarative approach defines desired state, simplifies configuration changes. • Integrates well with existing Ansible infrastructure for centralized management.

Launch | ACME | Namespace Backup

Development

1 Credentials
2 **Survey**
3 Preview

Which Namespace to Backup? *

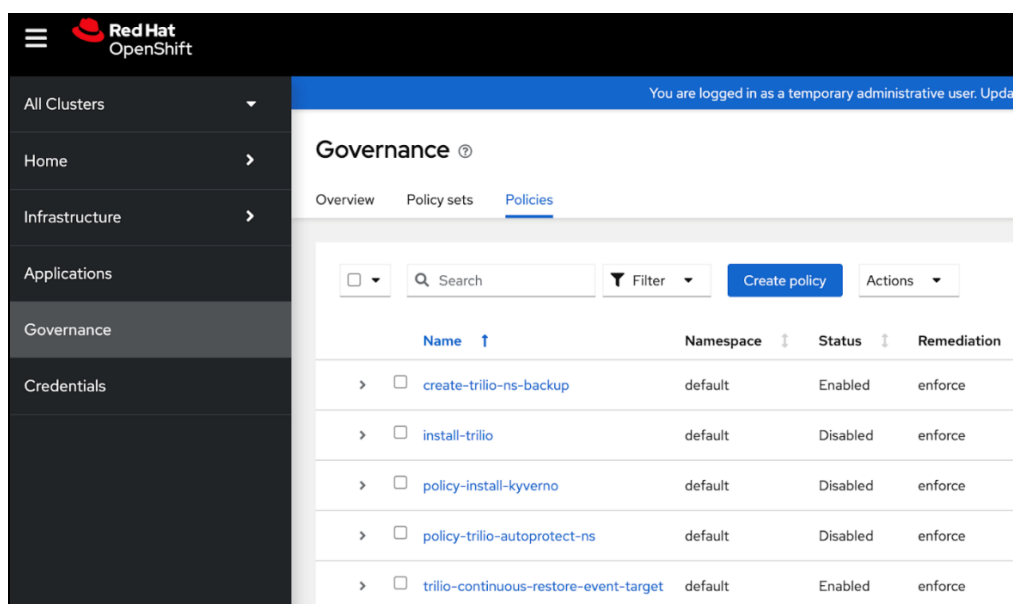
Select an option

- kevin-demo
- rodolfo-demo
- my-vms

Red Hat Advanced Cluster Management Governance and Policy Integration

Trilio Documentation & Resources for Deploying with ACM Policies ([Link Here](#))

Use Cases	<ul style="list-style-type: none">• Large-scale deployments with centralized policy-based management for Trilio.• Enforcing automated and consistent backup strategies across numerous OpenShift Virtualization clusters.• Automating Trilio deployments and configuration updates via ACM policies.
Benefits	<ul style="list-style-type: none">• Single pane of glass for managing backup policies across clusters.• Automated compliance for enforcing security best practices, to streamline the process of adhering to security regulations and internal security policies.• Reduced Workload: Automating compliance tasks frees up IT security teams to focus on more strategic initiatives.• Faster Response: Automated tools can continuously monitor systems for security issues and enable quicker remediation.• Compliance Reporting: ACM could generate reports demonstrating adherence to specific security regulations or internal policies.
Considerations	<ul style="list-style-type: none">• Requires an existing ACM deployment and familiarity with its functionalities.



Choosing the Right Option

Factors to Consider:

- Deployment size and complexity
- Automation requirements
- User skills and preferences

Decision Matrix

Feature	UI-driven Configuration	CLI-based Configuration / GitOps	Ansible	ACM Integration	DevOps Pipelines
Deployment Model	Any	Any	Any	Managed & Self-Managed	Any
Complexity	Simple deployments	Complex deployments	Complex deployments	Large-scale deployments	Complex deployments
Automation	Limited	High	High	High	High
User Skill Set	Beginner	Advanced user	Advanced user	Experienced user	DevOps knowledge
Multi-cluster Support	No	Yes	Yes	Yes	Yes
Centralized Management	No	No	No	Yes	No
Security Best Practices	Limited control	More granular control	More granular control	Enforced centrally	Limited control
Scalability	Limited scalability	High scalability	High scalability	Very High scalability	High scalability
Integration with CI/CD	No direct integration	Can be integrated	Can be integrated	No direct integration	Integrated

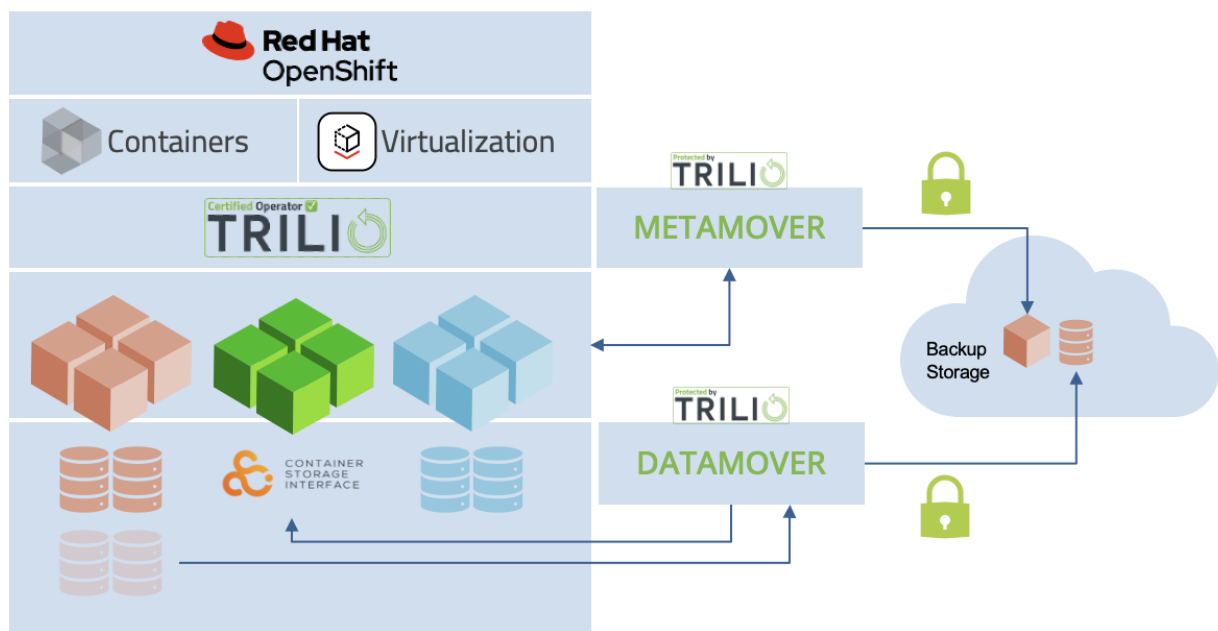
Best Practices

- ✓ Security considerations (access control, encryption)
<https://kubernetes.io/docs/concepts/security/secrets-good-practices/#least-privilege-secrets>
- ✓ Monitoring (Prometheus, Grafana)
 - Configuration Adjustment
 - Resources control (CPU/RAM)
 - Quota Management
- ✓ Alerting for backup jobs
- ✓ Disaster Recovery (DR) strategy using backups
- ✓ Testing and validation of backups (using ansible automation)
- ✓ Maintain proper DR documentation
- ✓ Maintain proper offsite backup copies
- ✓ Unattended recovery with Event Driven Ansible

Conclusion

Protecting Your OpenShift Investment with Confidence

By leveraging Trilio, you can ensure robust data protection, streamlined disaster recovery capabilities, and peace of mind for your critical applications and Virtual Machines.

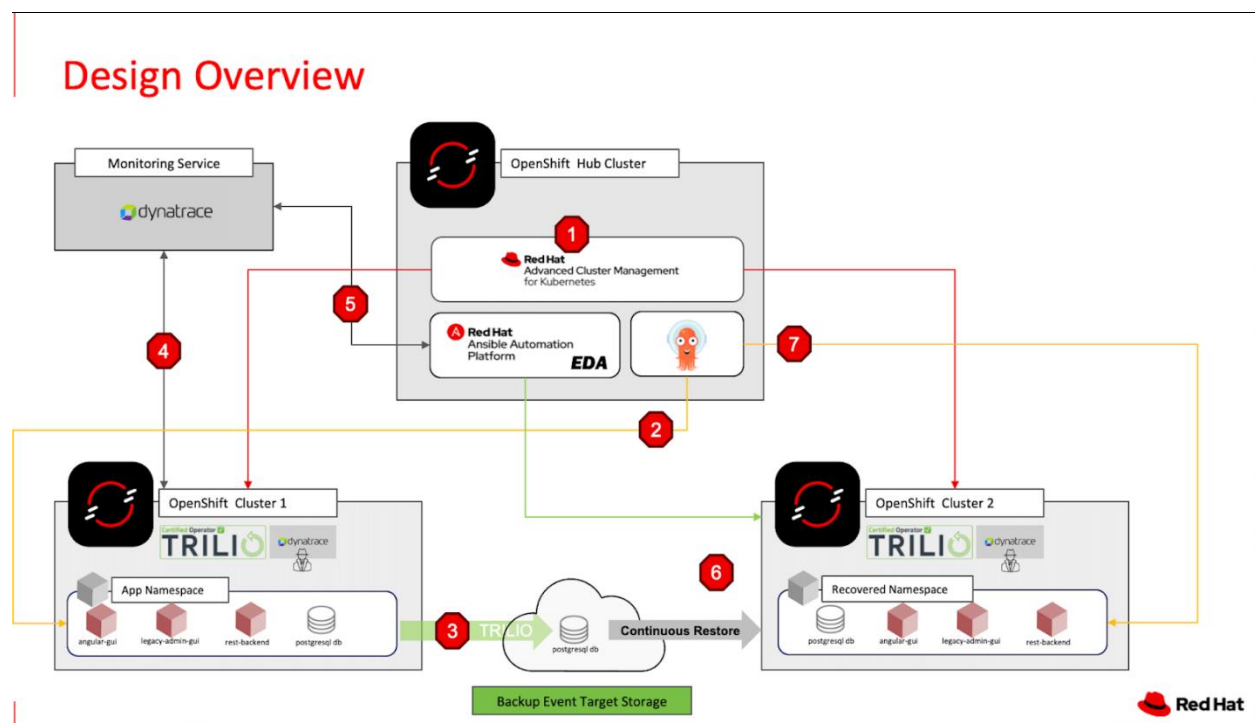


Key Takeaways

- Trilio offers a flexible and scalable solution for backing up your OpenShift clusters' Virtual Machines and Applications, regardless of the deployment model (managed vs. self-managed, on-prem vs. cloud).
- The choice of deployment option (UI, CLI, Ansible, ACM, DevOps Pipelines) depends on factors like deployment complexity, desired automation level, and user skillset.
- Each option provides varying degrees of control, scalability, and integration capabilities to suit your specific needs.
- Strong access control configurations and adherence to security best practices are crucial for protecting your backups.
- Integrating Trilio with your DevOps pipelines can further streamline your application deployment and backup workflows.

Integrations with the Ecosystem

We would like to highlight that as Trilio is a cloud native tool, it can be easily integrated with other tools. Actually, this is going to be presented on the Red Hat Summit in May



Moving Forward

By carefully considering the information presented in this document, you can make an informed decision about the best approach for deploying and managing Trilio within

your OpenShift Virtualization environment. This will ensure a secure and efficient backup strategy, allowing you to focus on innovation and delivering value with confidence.

Additional Considerations

Trilio goes beyond basic backups by offering **advanced features** that further enhance your data protection posture.

- ✓ Encryption ensures your backups remain secure at rest and in transit.
- ✓ Ransomware protection safeguards your backups from malicious attacks, providing an additional layer of security.
- ✓ Trilio's application mobility capabilities empower you to seamlessly migrate or recover applications across diverse environments, simplifying disaster recovery and fostering greater flexibility for your OpenShift deployments.
- ✓ Trilio's Continuous Restore (CR) shines in disaster recovery for Virtual Machines in OpenShift-Virtualization by minimizing downtime and being storage agnostic. CR asynchronously updates and maintains VM data copies readily available across different cloud environments, providing a reliable source for immediate restoration. This eliminates dependence on a single location, ensuring your OpenShift Virtual Machines and applications can recover quickly and remain fault tolerant even during a disaster. This approach saves additional money and resources as the target / DR site does not required pods to be up and running at all times. This could be impactful in hybrid or public cloud cost savings.

NOTE: High Availability is NOT Disaster Recovery

Disaster Recovery

<ul style="list-style-type: none">× Human error: Ansible playbooks executed with wrong scope, misconfiguration of unknown resources× Hardware Failure× Network Outage× Storage Failure× Ransomware attacks & Security Breach× DDOS attack× Power Outage× Natural Disasters× Software bugs× ETCD corrupted, sluggish	<ul style="list-style-type: none">× External service outage, that is available on the secondary site× DNS issues on the primary site× Physical Security Threats× Cloud Provider Outage× Cloud Instance Termination× Accidental Resource Deletion× Extreme Weather Events× Insider Threat× Volume Replication Failure× Poor Cluster Upgrade× Load Balancing Issues
--	---

While a robust backup strategy with Trilio forms the foundation of your disaster recovery plan, preparedness is key. Disaster Recovery planning and testing are no longer afterthoughts.

- ✓ Easily simulate disaster scenarios and meticulously test your restore processes. With Trilio, identify potential bottlenecks and refine your recovery procedures. This proactive approach ensures a smooth and efficient recovery in the event of an actual incident, minimizing downtime and data loss.
- ✓ Regularly testing your disaster recovery plan with restored backups instills confidence in your ability to quickly resume operations and safeguard your critical business data.

To empower your successful Trilio deployment and ensure you leverage its full potential, Trilio provides comprehensive resources at your fingertips. Trilio's **detailed documentation** offers step-by-step guides, configuration best practices, and troubleshooting tips. Additionally, Trilio's **responsive support team** is available to address any questions or challenges you may encounter throughout the implementation process. With these valuable resources readily available, you can confidently navigate your Trilio deployment and establish a robust backup strategy for your OpenShift environment.

About Trilio

Trilio is a leading provider of cloud-native Data Protection software solutions, supporting private, public and hybrid-clouds, engineered from ground up for Kubernetes, KubeVirt and OpenStack environments. At Trilio, we believe that data protection should be effortless, automated, and predictable. Our platforms deliver a modern data protection experience that gives customers more power and control over their applications and data. Cloud Architects, Platform Engineers, ITOps and DevOps departments, rely on Trilio technology for operational resiliency to perform critical tasks such as data backup and recovery, migration, ransomware protection, application mobility and disaster recovery.

Either in place, or into other clusters and clouds, Trilio's software dramatically reduces the amount of time spent on restoration and migration activities empowering customers from diverse sectors, such as telecommunications, financial services, defense, automotive and healthcare with the ability to easily deploy, manage and scale applications with confidence. Trilio has been a Premier Red Hat partner since 2017.

For more information, visit or contact us:

Trilio.io

[linkedin.com/company/Trilio](https://www.linkedin.com/company/Trilio)

Tel: + 1 (508) 233 3912

e-mail: info@trilio.io

Copyright © 2024 by Trilio Data, Inc.

All rights reserved. This document or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher.