


Disaster Recovery, Business Continuity and Backup for OpenShift with NetApp and ACM

How banks can protect containerised applications and virtual machines on Red Hat OpenShift with NetApp ONTAP, Trident and Red Hat Advanced Cluster Management, and produce clear evidence of recovery readiness for risk, compliance and audit.



Why it matters

Recovery targets set the design

DR VERSUS BACKUP

Disaster recovery keeps critical banking services available from a second site after a site-level outage, using replicated data and pre-staged applications.

Backup keeps independent, point-in-time copies to recover from data corruption, accidental deletion or ransomware. A resilient bank needs both, each tested regularly and evidenced for audit.

Business continuity relies on both working together

RPO

Recovery Point Objective: the maximum amount of data, measured in time, that a service can afford to lose.

It is driven by how often, or how synchronously, data is copied to the recovery site.

RTO

Recovery Time Objective: the maximum time a service may be unavailable before it is restored.

It is driven by automation, tested runbooks and the readiness of the recovery site and its staff.

Targets are set per service through business impact analysis

The building blocks

Three layers work together

STORAGE

NetApp ONTAP holds application and VM data. SnapMirror mirrors volumes to a secondary ONTAP system at a remote site, from which data can be served after a failover.

Layer
01

ONTAP, SnapMirror

PROTECTION

Trident is the CSI storage provisioner that connects OpenShift to ONTAP. Trident Protect protects and moves stateful applications. OADP backs up applications, volumes and VMs.

Layer
02

*Trident, Trident
Protect, OADP*

MANAGEMENT


Red Hat Advanced Cluster Management runs a hub that manages clusters, policies and applications across sites. GitOps keeps application definitions consistent between clusters.

Layer
03

*ACM, OpenShift
GitOps*

Reference architecture

Two sites, one management plane



Protection tiers


Match each service to a tier

RPO and RTO values are indicative




Snapshots

Point-in-time ONTAP copies on the same system. RPO: snapshot interval. Fast restore, but no protection from site loss.




Backup

Copies in object storage with Trident Protect or OADP. RPO: backup schedule. RTO is longest, as data must be restored.




Async DR

SnapMirror asynchronous replication to the DR site on a schedule. RPO: replication interval. RTO: set by failover automation.



Sync BC

SnapMirror synchronous or active sync. NetApp states RPO is zero in StrictSync mode, RTO very low. Needs low-latency links.



AppMirrorRelationship: app metadata and volumes replicated with SnapMirror, DC to DR

Containers with ACM

One hub, two clusters, one application

HOW IT WORKS

ACM with OpenShift GitOps deploys the same application definition from Git to the primary and DR clusters, selected through Placement.

Trident Protect replicates the application and its volumes to the DR cluster with SnapMirror, using an AppMirrorRelationship.

At failover, the DR copy is promoted and runs on the DR cluster; ACM keeps the configuration consistent and visible from one console.

Containers with ACM

Three steps to recover a service

ACM keeps the application definition identical on both clusters through GitOps, while Trident Protect replicates its data with SnapMirror. Recovery is a controlled change of state, not a rebuild.

Each step is a declared change in Git or in a Trident Protect resource, with status visible on both clusters, giving the bank a clear trail for audit and post-incident review.

01

Set up replication

Install Trident and Trident Protect on both clusters. Create AppVaults, the application and a replication schedule, then the AppMirrorRelationship on the DR cluster. ACM deploys the app definition.

02

Fail over to DR

On the DR cluster, set the AppMirrorRelationship to Promoted. Replication stops and the application comes online with the replicated data. Point ACM placement at the DR cluster.

03


Fail back to primary

Reverse resync to copy changes made at the DR site back to the primary, then reverse the replication direction. A plain resync would discard data written during failover.

VMs with ACM

Protecting virtual machines

OpenShift
Virtualization on
NetApp ONTAP



REPLICATION

Each VM and its disks are defined as a Trident Protect application. An AppMirrorRelationship replicates the VM disks and definition to a DR namespace with SnapMirror on a schedule.

At the DR site the PVCs exist, but the VM is not created until the relationship is promoted.

ROLE OF ACM

One console for both clusters: the same policies and GitOps configuration on each site, and a single view of where each VM is running.

VMs with ACM

Failover and failback for virtual machines

NetApp documents the same Trident Protect objects for VMs as for containers, so one runbook pattern covers both.

01 Set up


Create an ONTAP S3 AppVault, a Trident Protect app for the VM, a snapshot schedule and an AppMirrorRelationship in the DR namespace.

02 Fail over

Set the relationship to Promoted. The VM is created at the DR site from its replicated disks and can be started there.

03 Fail back

Reverse resync to keep DR changes, then reverse direction to the primary. Setting Established again discards DR writes.




Containers without ACM

Each cluster stands on its own

One Git source of truth


Argo CD on each cluster

Trident Protect replicates



Without a central hub, each OpenShift cluster runs its own OpenShift GitOps (Argo CD) instance, synchronising from the same Git repository.

Trident Protect replicates application data with SnapMirror. Operators run failover on each cluster and check status separately.



VMs without ACM

Operator-led recovery for virtual machines

HOW IT WORKS

Trident Protect replicates each VM's disks and definition to the DR cluster with SnapMirror.

OADP with its kubevirt plug-in keeps independent VM backups in object storage, for recovery from corruption or ransomware.

Without a hub, operators run failover with scripts or runbooks on each cluster and check status cluster by cluster.

Without ACM

Operator-run recovery steps

The same Trident Protect resources are applied directly on each cluster with `oc`, `kubectl` or `tridentctl-protect`, ideally from scripts kept under version control.

Retain each script run and its output as evidence of every DR test and real event, for review by risk, compliance and audit teams.

01

Set up on both clusters

Create AppVaults, the application and a replication schedule with `tridentctl-protect`, then the AppMirrorRelationship on the DR cluster. Configure OADP backups for VMs.

02

Fail over by command

On the DR cluster, patch the AppMirrorRelationship to desiredState Promoted. Confirm apps and VMs are running, then redirect users through DNS or the load balancer.

03

Fail back with care

Reverse resync: delete the DR relationship and create it in the opposite direction, then reverse replication to the primary. A plain resync discards DR writes.

Backup strategy

Backups that survive ransomware

1 Production data on ONTAP

2 Locked local snapshots (tamperproof)

3 Off-site copy: SnapMirror to DR, backups in object storage

3-2-1

3 copies, 2 media, 1 off-site

Immutable

copies that cannot be deleted

Four layers of protection

Replication protects against site loss, but it also copies corruption. Independent, immutable backups are the last line of defence.

- ONTAP snapshots for fast local restore points
- Trident Protect or OADP backups to S3-compatible object storage
- Off-site copy with SnapMirror to the DR site
- Tamperproof snapshots that cannot be deleted before expiry

Readiness checks

An immutable copy exists for each critical service


Restores are tested on a schedule

Backup credentials are kept separate

Backups are encrypted at the storage location

ACM hub resilience

A standby hub ready to take over



With or without ACM

Choosing the operating model

Aspect	With ACM	Without ACM
Management view	One hub console for all clusters	Each cluster managed on its own
Application deployment	GitOps from the hub, targeted by Placement	Argo CD on each cluster from the same Git
Policy and compliance	Central governance policies across clusters	Checks configured cluster by cluster
Failover action	Promote replication; move placement to DR	Operator scripts on each cluster
Data replication	Trident Protect with SnapMirror	Trident Protect with SnapMirror (same)
Management resilience	Standby hub with backup and restore	No hub to protect
Audit evidence	Hub view plus per-cluster records	Collected from each cluster
Best fit	Several clusters, central oversight	Small estates or a first phase

Testing and runbooks

Prove recovery before it is needed

A repeatable drill cycle, for example each quarter

Plan	scope, tier, success criteria
Test	clone, fail over, run the service
Verify	data integrity, RPO and RTO met
Record	evidence, findings, actions

Scheduled DR drills

Test on a clone of the SnapMirror copy while replication continues. From ONTAP 9.14.1, System Manager offers Test Failover; FlexClone volumes serve earlier releases. Verify the application and data, then clean up.

Monitoring


ACM observability brings metrics and alerts from every managed cluster to the hub. ONTAP reports the health and lag of each SnapMirror relationship. Alert on breaches of the target RPO.

Runbooks and evidence

A written runbook for each service tier. Each drill records timings, results, issues and sign-off, which gives audit and risk teams consistent evidence.

Comptech lab

The demo environment



Recommendations and next steps

- 01 Tier every service by its RPO and RTO
- 02 Adopt Trident Protect with SnapMirror for application and VM replication
- 03 Use ACM for multi-cluster governance and a single recovery view
- 04 Keep immutable, tested backups apart from replication
- 05 Run DR drills each quarter and keep the evidence

Comptech Network Systems | www.comptech.com.bd

Next steps

1 Workshop

Business impact analysis and service tiering

2 Proof of concept

On the Comptech DR lab

3 Design and build

Design, implementation and drills