KVM / DR

BEST PRACTICE GUIDE

RPO, RTO & DR Testing

Setting recovery objectives you can actually meet — and building the testing discipline that turns a disaster recovery plan from an assumption into a proven capability.

Published by **KVMDR** · kvmdr.ai · 2026

A VENDOR-NEUTRAL GUIDE FOR INFRASTRUCTURE AND PLATFORM TEAMS

EXECUTIVE SUMMARY

Objectives you can't test are objectives you don't have.

Every disaster recovery plan rests on two numbers — Recovery Point Objective (RPO) and Recovery Time Objective (RTO). They define how much data loss and downtime the business can tolerate. Most teams can recite them. Far fewer have ever proven they can meet them.

The gap between a documented objective and a demonstrated one is where DR plans fail. A plan that has never been rehearsed is a hypothesis, and the day it is finally tested is rarely the day you would choose. This guide covers how to set RPO and RTO targets that are realistic per workload, and how to run a testing program that keeps those targets honest.

THE TWO NUMBERS

RPO and RTO, precisely

OBJECTIVE	THE QUESTION IT ANSWERS
RPO Recovery Point Objective	How much data can we afford to lose? Measured as time — an RPO of five minutes means losing up to five minutes of changes is acceptable. It is set by your replication frequency, not your backup schedule.
RTO Recovery Time Objective	How long can we be down? Measured from incident to service restored. It is governed by how quickly workloads can be brought online, in the right order, on a reachable network.

> *RPO is a replication problem. RTO is an orchestration problem. Backup frequency alone determines neither — and a backup-only posture quietly degrades both.*

SETTING OBJECTIVES

Not every workload deserves the same target

A single estate-wide RPO/RTO is almost always wrong — too lax for the systems that matter, too expensive for the ones that don't. The right approach is to tier workloads by business impact and assign objectives per tier.

Group interdependent systems into **consistency groups** that fail over together, so application tiers stay coherent, then size replication and retention to the tier's target rather than to a blanket policy.

TIER	TYPICAL RPO / RTO	EXAMPLE WORKLOADS
Tier 0 — Critical	Seconds / Minutes	Transactional databases, core revenue systems, auth services.
Tier 1 — Important	Minutes / ~1 hour	Line-of-business apps, internal APIs, shared services.
Tier 2 — Standard	Hours / Hours	Reporting, dev/test, batch and back-office systems.

→ Make the targets defensible

Set objectives with the business owners who feel the loss, not in isolation. A target nobody has signed off on is a target nobody will fund — and a target you set without measuring your real change rate and recovery times is a guess. Validate both before you commit the number to a plan.

THE DISCIPLINE

Testing is what makes the objective real

A protection strategy is only as good as its last successful, verified recovery — not its last successful backup. Recovery you have not rehearsed is recovery you do not have.

The single biggest obstacle to regular testing is disruption: if a test means impacting production or pausing replication, it will be deferred until it is skipped entirely. **Non-disruptive testing** removes that excuse. Spinning up workloads from any recovery point in an isolated sandbox — while live replication continues untouched — converts a theoretical plan into a proven one, on a cadence you can actually sustain.

1 What a real test verifies

- ✓ **Recoverability.** The recovery point boots, the filesystem is consistent, and the application actually starts — not just that the data copied.
- ✓ **Ordering.** Consistency groups come up in the correct dependency order, so tiered applications are coherent on recovery.
- ✓ **Network.** Re-IP and network reconfiguration land workloads on a reachable network without manual intervention.
- ✓ **Timing.** Measured RTO against the target — and measured RPO from the recovery point's actual lag.

> *If a test never produces a number you can compare to your objective, it isn't a DR test — it's a demo.*

PUTTING IT TOGETHER

A testing program that scales

Sustainable DR testing is a cadence, not a heroic annual event. Match the frequency to the tier, automate what you can, and record the results so objectives stay honest over time.

CADENCE	WHAT TO RUN
Continuous	Automated boot / recovery-point health checks on protected workloads.
Monthly	Non-disruptive sandbox recovery of Tier 0 consistency groups, with RTO/RPO measured.
Quarterly	Full multi-group failover rehearsal, including re-IP and dependency ordering.
Annually	End-to-end failover and clean failback to source, with stakeholders observing.

DR TESTING CHECKLIST
Confirm each, every cycle

- ✓ Objectives tiered per workload and signed off by business owners
- ✓ Tests run in an isolated sandbox, replication never paused
- ✓ Measured RTO and RPO recorded and compared to target
- ✓ Consistency-group ordering and re-IP verified on recovery
- ✓ Failback to source rehearsed, not just failover
- ✓ Results trended over time, drift investigated

CONCLUSION

Set the number. Then prove it, on a schedule.

RPO and RTO are only as credible as the last test that met them. The teams that recover cleanly are not the ones with the most ambitious targets — they are the ones who rehearse modest, well-understood targets often enough that recovery is routine.

Tier your workloads, set objectives with the people who own the risk, and make testing non-disruptive enough to do regularly. Disaster recovery you have proven is the layer no organization regrets investing in — and the one nobody wants to be testing for the first time during an outage.