Disaster Recovery Planning Playbook for SaaS Companies

Jared Knedler & Erik Jones March 17, 2026

Executive Summary

Disaster recovery is the discipline most often deferred until it is too late to plan calmly. This guide explains how growing SaaS companies can build defensible recovery objectives, automate failover where it matters, and pass **SOC 2 and HIPAA** audits without theatrical "DR tabletops" that never get exercised. Numbers throughout reflect outcomes from infrastructure programs run for venture-backed and PE-backed SaaS customers.

Why DR Matters Now

For most SaaS companies, an unplanned outage is not catastrophic for the business — it is catastrophic for the renewal cycle. Enterprise customers increasingly write recovery objectives into MSAs, with contractual penalties for breaches. **SOC 2 (CC9.1, A1.2)** and **HIPAA**

Security Rule (45 CFR 164.308(a)(7)) both require documented, tested recovery plans. **ISO 27001:2022 Annex A.5.29 and A.5.30** codify the same controls.

Cloud-native architectures change the economics of DR. With AWS, Azure, and GCP all offering managed multi-AZ and multi-region primitives, achieving a one-hour RTO no longer requires hot-standby infrastructure that doubles your bill. As the AWS Well-Architected Framework's Reliability Pillar observes, *"the goal is not to eliminate failure but to reduce its blast radius and accelerate recovery."* The discipline is now about choosing the right pattern per workload, codifying it in Terraform, and exercising it regularly enough that the runbook reflects reality.

Core Framework: RTO, RPO, and Workload Classification

Two numbers anchor every DR conversation: **Recovery Time Objective (RTO)** — how long you can tolerate being down — and **Recovery Point Objective (RPO)** — how much data you can tolerate losing. Setting these per workload class, not per server, is the inflection point between mature and immature programs.

The Tier Model

> **Tier 0 - Customer-facing API + auth (Stateless):** RTO 15-60 min, RPO 0-5 min. Multi-AZ active-active behind a load balancer. Failover is automatic.

Tier 1 - Primary database (Stateful): RTO 1-4 hr, RPO 1-15 min. RDS Multi-AZ for AZ-level failure; cross-region read replicas with promotion runbook for region-level failure.

Tier 2 - Analytics, batch, internal tools: RTO 4-24 hr, RPO 1-24 hr. Restore-from-backup is acceptable; no hot standby.

Tier 3 - Archives, compliance evidence: RTO over 24 hr, RPO 24 hr. S3 Glacier Deep Archive with cross-region replication; cost dominates over speed.

The Tier-0 Anti-Pattern

The most common failure mode is treating every workload as Tier 0. Hot-standby infrastructure for analytics and internal tools wastes 50-70% of DR budget on workloads that could safely run as Tier 2 with 1-day RPO. The discipline is to tier honestly.

Practical Implementation on AWS

The implementation pattern we deploy on most engagements maps cleanly to AWS-native primitives. The same shape applies on Azure and GCP with vendor-equivalent services.

Database Tier

RDS Multi-AZ with a 7-day automated backup window, plus cross-region read replica for Tier 1 workloads. Aurora customers get **global database** when budget allows. Point-in-time recovery (PITR) tested monthly, not just configured.

Replica Promotion Runbook

Cross-region read replicas are useful only with a tested promotion runbook. The runbook covers DNS cutover via Route 53 health-check-based failover, secret rotation if credentials are region-specific, application-side connection-string updates, and the rollback path.

Object Storage

S3 Cross-Region Replication on every bucket holding customer data, with lifecycle policies that age objects to Glacier Deep Archive at 90 days. Versioning enabled to defend against ransomware-style overwrite attacks. **S3 Object Lock** for buckets that hold immutable audit evidence.

Compute Tier

Auto Scaling Groups span at least three Availability Zones with launch-template-based AMIs. For Kubernetes, **Velero** handles cross-cluster restore with snapshots stored in a separate region's S3.

Backup Centralization

AWS Backup centralizes RDS, EBS, and EFS backups under a single audit trail. For on-prem segments, **Veeam** writes to cross-region S3 with WORM-locked immutable backups.

The DR Drill: What Real Testing Looks Like

Most companies "test" DR by reading the runbook in a conference room. That is not a test. A real DR drill, run quarterly, produces measurable evidence.

The Drill Cadence

Tabletop (every quarter): Simulated incident; on-call walks the runbook; gaps flagged; runbook updated within 7 days.

Functional test (semi-annually): Failover one specific component to the DR region. Measure actual RTO/RPO. Roll back. Document.

Full DR exercise (annually): Failover Tier 0 and Tier 1 workloads to DR region. Run synthetic load for 1-4 hours. Validate every dependency. Roll back during a maintenance window. Audit-grade evidence retained.

Chaos engineering (continuous, optional): AWS Fault Injection Simulator or Gremlin run weekly against staging.

Audit-Grade Evidence

The auditor wants dated artifacts: who ran the drill, what was tested, what failed, what was fixed, when. A runbook from 2023 is not evidence of a working program in 2026 — it is the most-cited finding in mid-stage SOC 2 audits.

Common Pitfalls

One-size-fits-all RTO/RPO: forces over-engineering and inflates infrastructure spend. Tier the workloads first.

Single-region "multi-AZ": multi-AZ protects against AZ failure, not region-wide events. The November 2020 us-east-1 outage took out everything from EC2 to IAM.

Backups but no restore practice: restoring 50 TB from S3 Glacier takes 12+ hours; discovering this during an incident means RTO is 12 hours regardless of what the runbook says.

Runbook drift: infrastructure changes in Terraform; the runbook does not. Generate runbook content from infrastructure metadata where possible.

Compliance theater: a Disaster Recovery Plan PDF in SharePoint is not a DR program. The auditor wants test records, post-incident reviews, and linked tickets.

Measured Outcomes

Tier 0 RTO under 60 minutes measured across the last four DR drills, with most landing 15-30 minutes

Tier 1 RPO 1-15 minutes via RDS cross-region replicas plus PITR

Annual full DR exercise completed with documented evidence, dated runbooks, and a post-mortem within 30 days

SOC 2 CC9.1 and A1.2 evidence package auto-generated from Terraform plan plus Datadog uptime data plus drill records

35-60% infrastructure cost savings versus a uniform hot-standby model, achieved by tiering honestly

How Jacobian Helps

Most SaaS teams do not need a dedicated DR engineer; they need a partner who has done DR programs across dozens of cloud-native customers and can codify the right pattern per workload. Our SREs have audit and compliance experience — runbooks, evidence, and SOC 2 / HIPAA controls are designed in from week one rather than retrofitted before an audit. Every DR resource ships through Terraform; every drill produces auditor-grade evidence; every runbook is reviewed quarterly against actual infrastructure state. When you eventually hire your own SRE team, you inherit a working program with documented runbooks — not a black box.

Resource Details

Author: Jared Knedler & Erik Jones

Published: March 17, 2026

Categories:

disaster recovery | business continuity | SaaS | compliance

Download

Full Document

About This Resource

A comprehensive playbook for SaaS disaster recovery planning covering RTO/RPO definition, cloud-native strategies, and compliance-aligned testing.

Categories:

disaster recovery **business continuity** **SaaS** **compliance**