Disaster Recovery Plan

Disaster Recovery Plan

> **Project:** {{PROJECT_NAME}} **Version:** {{VERSION}} **Date:** {{DATE}}
> **Author:** {{AUTHOR}} **Status:** Draft | In Review | Approved **Reviewers:** {{REVIEWERS}}

Document History

Version	Date	Author	Changes
0.1	{{DATE}}	{{AUTHOR}}	Initial draft

---

1. Business Continuity Overview

This plan documents the procedures to recover {{PROJECT_NAME}} services following a disaster event (data center failure, data corruption, security breach, or catastrophic failure).

Plan Owner: {{DR_OWNER}} **Plan Reviewer:** {{DR_REVIEWER}} **Last Tested:** {{LAST_TEST_DATE}} **Next Scheduled Test:** {{NEXT_TEST_DATE}}

Disaster types covered:

- Infrastructure failure (AZ/region outage)
- Data corruption or accidental deletion
- Security incident (ransomware, data breach)
- Vendor/provider outage
- Catastrophic application failure

---

2. RPO / RTO Targets Per Service Tier

Tier	Description	RPO	RTO	Examples
Tier 1 — Critical	Core user-facing services; downtime has direct revenue impact	0 (real-time replication)	< 15 min	Auth, checkout, core API
Tier 2 — Important	Supporting services; degraded experience without them	< 1 hour	< 4 hours	Notifications, reports
Tier 3 — Standard	Background/admin services; business can operate without temporarily	< 24 hours	< 24 hours	Analytics, admin panel

---

3. Service Tier Classification

Service	Tier	Owner	Rationale
{{SERVICE_1}}	Tier 1	{{OWNER}}	Core user journey
{{SERVICE_2}}	Tier 1	{{OWNER}}	Authentication
{{SERVICE_3}}	Tier 2	{{OWNER}}	Supporting
{{SERVICE_4}}	Tier 3	{{OWNER}}	Admin only
Database — Primary	Tier 1	Platform	All services depend on it
Object Storage	Tier 2	Platform	User uploads

---

4. Backup Strategy

4.1 Database Backups

Database	Backup Type	Frequency	Retention	Location	Verified
{{DB_PRIMARY}}	Automated snapshot	Daily	30 days	{{BACKUP_LOCATION}}	Monthly
{{DB_PRIMARY}}	Point-in-time recovery	Continuous	7 days	{{BACKUP_LOCATION}}	Monthly
{{DB_READ_REPLICA}}	Not backed up separately	—	—	Rebuilt from primary	—

Automated backup tool: {{BACKUP_TOOL}} **Backup encryption:** AES-256, key managed in {{KMS_TOOL}} **Cross-region copy:** {{CROSS_REGION}}

4.2 File / Object Storage Backups

Storage	Backup Method	Frequency	Retention	DR Copy
{{S3_BUCKET}}	S3 versioning + replication	Continuous	{{RETENTION}}	{{DR_BUCKET}}
{{FILE_STORE}}	Snapshot	Daily	30 days	Cross-region

4.3 Configuration Backups

Config	Backup Method	Location	Frequency
IaC (Terraform)	Git repository	{{GIT_REPO}}	On change
Application config	Git repository	{{GIT_REPO}}	On change
Secrets	Secrets manager replication	{{SECRETS_BACKUP}}	Real-time
DNS records	Export to Git	{{GIT_REPO}}	Weekly
TLS certificates	Secrets manager	{{CERTS_BACKUP}}	On renewal

4.4 Backup Testing Schedule

Backup Type	Test Frequency	Last Test	Result	Tester
Database full restore	Monthly	{{DATE}}	{{RESULT}}	{{TESTER}}
Point-in-time restore	Quarterly	{{DATE}}	{{RESULT}}	{{TESTER}}
Object storage restore	Quarterly	{{DATE}}	{{RESULT}}	{{TESTER}}
Full DR failover drill	Bi-annually	{{DATE}}	{{RESULT}}	{{TESTER}}

---

5. Failover Procedures

5.1 Automated Failover

Component	Automatic Failover	Mechanism	Failover Time
Database (Multi-AZ)	Yes	RDS automatic failover	60-120 seconds

Component	Automatic Failover	Mechanism	Failover Time
Load balancer	Yes	Health check → route to healthy targets	< 30 seconds
CDN	Yes	Origin health checks	< 60 seconds
Redis (if clustered)	Yes	Redis Sentinel / ElastiCache	< 30 seconds

Monitoring automatic failover:

- Alert fires: `MultiAZFailover` CloudWatch event or equivalent
- On-call notified immediately
- No manual action required, but on-call must confirm recovery

5.2 Manual Failover Steps

Prerequisite: Automatic failover has NOT occurred or has failed.

Database Manual Failover (Tier 1)

1. Confirm primary is unavailable: `ping {{DB_PRIMARY_HOST}}` — should timeout
2. Connect to standby: `psql {{STANDBY_HOST}}`
3. Promote standby to primary: `SELECT pg_promote();`
4. Update DNS record `db.{{INTERNAL_DOMAIN}}` → `{{STANDBY_HOST}}`
5. DNS TTL: Ensure TTL was set to 60s pre-incident (if not, wait {{DNS_TTL}} seconds)
6. Verify applications are reconnecting: Check application logs for successful DB connections
7. Page on-call to verify all services healthy

Regional Failover (Catastrophic)

1. Declare DR event (approval from {{DR_AUTHORITY}})
2. Confirm primary region {{PRIMARY_REGION}} is unreachable
3. Activate standby in {{DR_REGION}}: `terraform apply -var-file=envs/dr.tfvars`
4. Restore database from latest cross-region snapshot
5. Update Route 53 / DNS to point to {{DR_REGION}} endpoints
6. Run smoke tests: `bash scripts/smoke-tests.sh {{DR_REGION}}`
7. Notify stakeholders (see Communication Plan)
8. Monitor enhanced metrics for {{MONITOR_PERIOD}}h

---

6. Recovery Procedures Per Service

Tier 1 Services

Service	Recovery Procedure	Recovery Script	Est. Time
{{SERVICE_1}}	1. Restore from snapshot 2. Verify config 3. Run smoke tests	`scripts/restore-{{SERVICE_1}}.sh`	{{TIME}}min
Authentication	1. Deploy from last known good image 2. Verify JWT keys 3. Test login flow	`scripts/restore-auth.sh`	{{TIME}}min

Tier 2 Services

Tier 3 Services

---

7. DR Drill Schedule & Scenarios

Drill Type	Frequency	Participants	Last Executed	Next Scheduled
Tabletop exercise	Quarterly	On-call team + engineering lead	{{DATE}}	{{DATE}}
Database failover test	Quarterly	DevOps + one developer	{{DATE}}	{{DATE}}
Full DR failover	Bi-annually	Entire engineering team	{{DATE}}	{{DATE}}
Backup restore test	Monthly	DevOps	{{DATE}}	{{DATE}}

Drill Scenarios to Cover:

1. Database primary failure (automatic failover test)
2. Accidental data deletion (point-in-time restore)
3. Single AZ outage (multi-AZ failover)
4. Full region failure (cross-region DR)
5. Ransomware/data corruption (restore from offline backup)
6. CDN outage (origin fallback)
7. Secret store unavailable (cached credentials)

---

8. Communication Plan During DR Event

Internal Communications

Audience	Channel	Frequency	Owner
Engineering team	Slack #incidents + war room call	Real-time	Incident commander
Engineering management	Direct message	At declaration + hourly	Incident commander
Product/Business leadership	Email + Slack	At declaration + hourly	Incident commander
Customer support	Dedicated Slack channel	At declaration + 30 min	Support lead

External Communications

Audience	Channel	Trigger	Message
Customers	Status page ({{STATUS_PAGE}})	Within 15 min of confirmed incident	"We are investigating an issue"
Customers	Status page update	Every 30 min	Progress update
Customers	Email	If impact > {{EMAIL_THRESHOLD}}h	Direct notification
SLA customers	Direct contact	Per SLA contract	As contractually required

Communication templates: See go-live-runbook.md communication section

---

9. War Room Setup

War Room: {{WAR_ROOM_LINK}} **Bridge Line:** {{BRIDGE_NUMBER}} **Document:** Live incident doc created at: {{INCIDENT_DOC_TEMPLATE}}

Roles during DR event:

Role	Responsibility	Primary	Backup
Incident Commander	Coordinates response, final decisions	{{IC}}	{{IC_BACKUP}}
Technical Lead	Leads technical recovery	{{TECH_LEAD}}	{{TECH_BACKUP}}
Communications Lead	Internal/external updates	{{COMMS_LEAD}}	{{COMMS_BACKUP}}
Scribe	Documents timeline, actions taken	{{SCRIBE}}	Rotate

---

10. Post-Recovery Verification Checklist

- ☐ All Tier 1 services healthy (health checks passing)
- ☐ Error rate back to baseline (< {{ERROR_BASELINE}}%)
- ☐ P99 latency back to baseline (< {{P99_BASELINE}}ms)
- ☐ Database connections stable
- ☐ Replication lag < {{REPLICATION_LAG}}s (if applicable)
- ☐ Backup jobs resumed and completed successfully
- ☐ Monitoring and alerting functional
- ☐ No data loss confirmed (or data loss quantified and documented)
- ☐ All Tier 2 services healthy
- ☐ Stakeholders notified of recovery
- ☐ Status page updated to "Resolved"
- ☐ Incident timeline documented
- ☐ Post-mortem scheduled (within {{POSTMORTEM_SLA}}h)

---

11. DR Test Results Log

Date	Test Type	Scenario	RTO Achieved	RPO Achieved	Issues Found	Resolved By
{{DATE}}	{{TYPE}}	{{SCENARIO}}	{{RTO}}	{{RPO}}	{{ISSUES}}	{{RESOLVED}}

Related Documents

- Monitoring & Observability
- Operational Runbook
- Incident Report
- Post-Mortem

Approval

Role	Name	Date	Signature
Author			
Reviewer			
Approver			

Revision #7
Created 2026-02-23 12:05:57 UTC by John
Updated 2026-05-25 07:33:57 UTC by John