DataSitr — Disaster Recovery Summary

Generated: 2026-05-17T05:10:49Z

Source document: docs/disaster-recovery-summary.md

Git commit: dee50b6d

Generator: [operator-tooling]

Benchmark artifact: docs/generated/pii_benchmark_latest.json (generated 2026-04-29T14:40:52Z, Arabic NER loaded, 1K p95 47.92 ms)

For: Operational resilience reviewers, IT risk assessors, business continuity diligence

This summary reflects the current documented encrypted-first backup posture and

the still-narrow continuity claim boundary.

Recovery objectives

Metric	Current posture	Notes
RPO (Recovery Point Objective)	Last successful encrypted backup interval	Backup cadence is operator-configured; the live pilot uses encrypted local backup plus verified OSS off-host replication
RTO (Recovery Time Objective)	Operator-directed restore window measured in minutes	Service restart from local backup is fast; remote recovery still depends on transfer and operator execution

These are not contractual SLAs. They reflect observed operational capability on

the current Saudi-hosted pilot stack. They do not imply replication-backed

continuity, automatic failover, or zero-downtime disaster recovery.

What is backed up

Data	Included in backup	Encryption
PostgreSQL runtime dump (vault, compliance, auth, billing, pricing tables)	Yes	Backup encryption

Legacy compliance JSONL snapshots and app data	Yes	Backup encryption
Configuration snapshot (deployment configuration summary, tenant policies)	Yes	Backup encryption
API key store and billing state	Yes	Keys stored as SHA-256 hashes; runtime state captured in PostgreSQL dump
Application code	No	Re-deploy from git / image artifact
Dashboard build	No	Rebuild on deploy
Provider catalogs and pricing	No	Re-synced on startup

What is NOT backed up

Data	Reason	Recovery path
Live deployment configuration, master key material, and backup passphrase file	Kept outside the archive by design	Preserve separately; required before a restored stack can start correctly
Backup encryption passphrase file	Kept outside the archive by design	Preserve separately; required to decrypt `.tar.gz.enc` archives
External provider-console settings (quota changes, console-side metadata)	Not part of local runtime state	Reapply from the provider console if needed
In-flight webhook deliveries	Active/ directory may contain in-progress items	Crash recovery moves active → pending on restart

---

Backup procedures

Local backup

[See source document for diagram/code]

Creates a timestamped encrypted archive containing database state, compliance snapshots, and a redacted config snapshot. It does not include the live deployment configuration, master key material, backup passphrase file, or runnable application code. Archive is stored in the local backup directory.

Off-host replication

Two supported options exist. The live pilot now uses Alibaba Cloud OSS as the

primary off-host path and SSH/SCP as fallback only.

Primary pilot command:

[See source document for diagram/code]

`backup_oss.sh --latest-only --verify` now prefers the newest encrypted `.tar.gz.enc` archive, fails closed if no encrypted archive exists, and records authoritative host-side state in `[internal path]

The current dated pilot evidence also includes a verified encrypted upload -> download -> restore drill at `[internal path]

Method	Script	Destination
Alibaba Cloud OSS	`[operator-tooling]`	Cloud object storage bucket
SSH/SCP	`[operator-tooling]`	Remote server via SSH

Backup verification

[See source document for diagram/code]

The first command verifies that the active off-host replication candidate is an encrypted archive and that upload/download integrity succeeds. The second runs a non-destructive restore verification: extracts the archive, validates file integrity, checks vault database readability, and reports result without modifying the running system.

Restore procedures

From local backup

[See source document for diagram/code]

Restores runtime data from the archive. Service must be stopped before restore and restarted after. If the archive includes a configuration snapshot, the operator still needs to restore or verify the active deployment configuration deliberately before starting the service.

From off-host backup

[See source document for diagram/code]

Downloads the latest backup from OSS, then restores using the standard restore procedure.

Post-restore verification

After restore:

1. Start the service
2. Run the standard health and readiness checks for the restored environment
3. Run the operator smoke test for the restored environment
4. Verify vault integrity: process a test request and confirm rehydration works

---

Current proof boundary

The current live pilot proves:

- encrypted local backup creation as the intended pilot/production path
- verified OSS upload/download integrity for the newest encrypted archive
- isolated restore-recovery from a real encrypted backup
- operator-directed public restored-state cutover from a dated encrypted backup
- operator-directed public Host A ingress independence during a restored-state Host B direct-ingress window
- multi-AZ ACK customer-route cutover with a 4-hour soak (2026-05-04 evidence bundle at `evidence/ha/alibaba-live-2026-05-04T01:17:03Z/`, signed)
- in-Kingdom multi-region warm-standby drill rehearsal covering DNS, GKE Ingress, and TLS routing only (2026-05-16 evidence bundle at `evidence/multi-region-drill/multi-region-warm-standby-20260516T220433Z/`, Ed25519-signed)

The current live pilot does **not** prove:

- full-vault verification after restore
- cross-cloud PostgreSQL / Redis replication-backed continuity
- automatic database-tier or authentication failover
- unplanned full Host A machine-loss tolerance
- unplanned full-region failure tolerance (warm-standby drill is operator-directed routing only; data-tier failover remains manual)
- blanket HA
- zero-downtime disaster recovery

Failure scenarios

Scenario 1: Service crash / restart

Aspect	Behavior
Recovery	Automatic via systemd restart or Docker restart policy
Data loss	None — vault and logs are on disk
Webhook queue	`_recover_active()` moves interrupted deliveries back to pending on startup
Time to recover	Seconds

Scenario 2: Disk corruption on VPS

Aspect	Behavior
Recovery	Restore from most recent backup to new or repaired VPS
Data loss	Changes since last backup
Prerequisite	Encrypted archive and backup passphrase must both be available; required runtime key material and configuration must be restorable
Time to recover	Minutes from local backup; variable from remote

Scenario 3: Full VPS loss

Aspect	Behavior
Recovery	Provision new VPS, deploy application, restore from off-host backup
Data loss	Changes since last off-host backup
Prerequisites	Off-host backup configured and current; encrypted archive and backup passphrase both available
Time to recover	30–60 minutes (provisioning + deploy + restore + verification)
Current live pilot posture	Authoritative off-host backup state and dated encrypted upload -> download -> restore proof exist on the pilot; each additional deployment still needs its own destination, schedule, verification, and monitoring rollout

Scenario 4: Master key loss

Aspect	Behavior

Recovery	**Not possible** — vault data is cryptographically unrecoverable without the master key
Mitigation	Master key backup procedures documented in runbook; operator must secure separately
Recommendation	Stricter external KMS/HSM-backed steady-state custody would narrow this single point of failure

Scenario 5: Provider outage (external AI)

Aspect	Behavior
Detection	Circuit breaker opens after consecutive failures
Behavior during outage	Requests fail-safe — either use fallback provider or return error
Recovery	Circuit breaker transitions to half-open after timeout; auto-recovers when provider returns
Data impact	None — no data loss from provider outage

Scenario 6: Region-level disruption (Alibaba Riyadh)

Aspect	Behavior
Detection	Cloud Monitoring uptime checks alert when the Alibaba ACK customer route fails health checks
Behavior during outage	The GCP Dammam warm-standby (`me-central2`) is reachable at `standby.gcp.datasitr.com` and returns HTTP 200 in degraded mode; the standby validates DNS, GKE Ingress, TLS routing, and Cloud Armor WAF — it does not yet serve customer traffic with cross-cloud database state
Operator-directed cutover	Operator updates the public DNS record to point at the Dammam standby and brings up degraded service; full restoration requires restoring the database tier from off-host backup at the standby region (RTO measured in tens of minutes, not seconds)
Data impact	Up to the last off-host backup until cross-cloud DB replication is implemented
Boundary	The drill rehearsal evidence proves DNS/GKE/TLS routing only; it does not prove cross-cloud DB replication, automatic failover, or unplanned full-region tolerance
Current evidence	`evidence/multi-region-drill/multi-region-warm-standby-20260516T220433Z/`, signed `sha256:158f604a...`, public key fingerprint `ed25519:f4c7e408...`

Deployment architecture for DR

[See source document for diagram/code]

Gaps and next steps

Gap	Impact	Remediation
Backup monitoring rollout beyond the pilot	The pilot now has state-backed off-host backup and restore-drill monitoring, but each additional deployment still needs the same flow wired and verified	Mirror the pilot's state-backed monitor path, alert destination, and restore-drill cadence into every customer-facing deployment
Cross-cloud database replication and auth failover	Multi-region warm-standby covers DNS/GKE/TLS routing only; PostgreSQL/Redis primary failover remains operator-directed; authentication failover not yet automated	GCP Dammam warm-standby (`me-central2`) was rehearsed via signed scoped drill 2026-05-16 (DNS/GKE/TLS only); cross-cloud DB replication is the next major buildout, gated on a low-traffic maintenance window for the RDS replication prerequisite
No replication-backed continuity or automated failover	Recovery remains backup-based and manual	Only widen the claim after dated proofs exist for replication-backed continuity and automatic failover
Backup encryption passphrase handling remains operator-managed	Loss of the passphrase blocks decryption of encrypted archives	Preserve the passphrase separately and document rotation/retention discipline
Restore drills are not yet on a mature periodic cadence	The pilot now has dated restore and restored-state evidence, but recovery confidence will decay without repeats	Schedule and document repeat off-host restore drills and keep the dated summary current
Full-vault verification after restore is still narrow	Restored-state proofs currently sample oldest/newest vault rows only	Keep the current claim narrow and add a full-vault verification pass before widening integrity claims

---

Version: 0.1.1 | **Last updated:** 2026-05-17

See also: Operations Runbook | Pilot Deployment Guide | Production Readiness Checklist