Backup and disaster recovery

Backup policies and dispatch, S3-compatible immutable repository, verification of each backup by isolated restoration, multi-engine PITR, semi-automatic DR plans and federation between sites.

General model

Status: Available

The backup module resides in the agent installed on each server. The console schedules policies, dispatches jobs and shows results; database credentials are resolved by the agent on the server itself and never pass through the console. Each agent's scope is its own host: it protects only the server where it resides and everything that server contains.

Each backup:

- Travels in stream, without intermediate temporary files, encrypted with **AES-256-GCM**.
- Is accompanied by a manifest with the engine and its version, start and end times, sizes, the SHA-256 fingerprint of the artifact, the row count per table and, after verification, the measured restore time.
- Is governed by a generational retention policy: daily, weekly, monthly and annual.

Destinations and immutable repository

Status: Available

Status by destination:

- **Available:** the S3 object repository — the S3 layer is always provided by **MinIO**, deployed by the vendor —, including **NAS storage**, which is served over S3 and already operates that way in production, and the storage of another federated site.
- **Commitment (committed roadmap): HPE 3PAR**, served likewise over S3 (MinIO layer).
- **Roadmap:** tape and the other destination types.

Available destinations are interchangeable with each other.

- **S3-compatible object repository with immutable object lock.** A retention is applied to the sealed point that prevents its deletion during its period, in compliance mode when the policy protects critical points. The usage pattern requires a dedicated container per agent or organisation, with least-privilege credentials specific to that agent and without administration or retention-bypass permissions.
- **Encryption per agent.** Each backup is encrypted per agent before leaving toward its destination; compromising a destination does not expose backups in cleartext.

- **Second copy.** The policy may define a second destination, immediate or scheduled. The copy verifies the fingerprint during transfer: if it does not match, the copy is aborted and the partial is discarded.
- **Disconnected copy (roadmap).** A disconnectable local medium provides the isolated copy against ransomware.

Verification of each backup

Status: Available

A backup that is never restored is a promise, not a backup. That is why **each backup is verified through a restoration in an isolated environment (sandbox)**, with no network toward production, following this order:

1. Antivirus analysis of the decrypted content. If the analysis is not available on that server, the report declares it as omitted, without pretending it was performed.
2. Startup of the isolated environment and wait for the engine to respond healthy.
3. Restoration of the artifact and row count per table, which must match the manifest exactly.
4. Measurement of the restore time, which is recorded in the manifest.
5. In extended verification, tests on the declared portals and real capture of the check.
6. Destruction of the isolated environment in all cases.

Verification produces a **report** that accompanies the backup. A backup marked as compromised is marked and isolated, and is never deleted automatically; restoring a marked point requires explicit, recorded risk approval.

In addition to restoration, a periodic integrity check re-reads and recomputes the fingerprint of each artifact and contrasts it with its manifest, to detect silent corruption.

Engines and point-in-time recovery (PITR)

Status: Available

Engine	Full backup	Incremental	Point-in-time recovery
PostgreSQL	Dump in custom format	Base copy with archived transaction log segments	Restoration of the base and replay of the log up to the requested date and time or position
MySQL and MariaDB	Transactional dump	Binary log segments from the position captured in the base	Restoration of the base and replay of the log up to the requested date and time or position
SQL Server	Full database copy	Transaction log copies	Restoration without recovery, application of logs up to the requested instant and close with recovery
SQLite	Full copy, always	Not applicable	To the last backup

Point-in-time recovery is **verified end to end on the three engines with transaction logs**: PostgreSQL, MariaDB and SQL Server. Before touching production, restoration is first executed in the isolated environment and can be checked with a read-only query on the chosen instant.

Retention pruning never leaves a broken chain: if an incremental is retained, its base and all intermediate incrementals are retained.

Volumes and file paths are also backed up, with a per-file index that allows restoring a single file without restoring the full volume, as well as virtual machines and hypervisor containers on Proxmox VE, and the console itself.

Continuous backup

Status: Available, for workloads whose data-loss objective is minutes.

Continuous backup sends log segments as they close, in order, verifies each upload by re-reading the object and comparing its fingerprint, and only then confirms the position. Its operational measure is the lag between the current moment and the last confirmed segment, which represents the real potential loss and must be monitored and alerted. The user manual declares the verified limit of this function, which must be consulted before presenting it as verified against production engines.

Disaster recovery (DR) plans

Status: Available

A recovery plan is a versioned plan that rebuilds in order: network, volumes, data and applications, waiting for each layer to respond healthy before continuing.

- **Automatic preparation.** Failure detection and plan preparation — most recent verified point, destination checks and execution order — are automatic.
- **Execution with human approval.** Execution is **semi-automatic**: it requires the approval of an identified person, distinct from whoever configured the plan, through the console or through the foreseen on-call channels. If nobody authorises, nothing affecting data is executed.
- **Drills.** Plans are exercised with periodic drills that truly restore in an isolated environment and measure times against the objectives of the contracted tier, with a signed report.
- **Federation between sites.** Backup destinations may be another site of the federation, and a plan may rebuild at the alternate site from that federated backup. Each site continues to operate autonomously if the federation is not available.

Recovery time and recovery point objectives

The public sources of the dossier set the following **design objectives per tier**, and expressly declare that zero loss is not promised: the design is based on frequent snapshots and continuous backup, not on synchronous replication.

Tier	Recovery point objective (RPO)	Recovery time objective (RTO)
Commercial Essential	24 hours	4 hours

Tier	Recovery point objective (RPO)	Recovery time objective (RTO)
Commercial Professional	1 hour	1 hour
Government	Contractual; design: 15 minutes	Contractual; design: 30 minutes

The time that counts before an auditor is the **measured** time in each verification and in each drill, which is recorded in the manifest and in the corresponding report.

Tape

Status: Committed roadmap. Tape forms part of the committed destinations programme (October 2026); this section records its target design, with the limitation declared in the user manual.

Backup to tape with a generic driver and the complete tape infrastructure on that basis — media sets by generations, off-site media custody with movement logging, several parallel streams and sets with absolute retention — form part of the product. Validation against a specific physical library is subject to the site's compatibility matrix; the validation performed is against a virtual library.