EON

Enterprise Cloud Backup Buyers Guide

If you're comparing enterprise cloud backup tools, this guide helps you pick the right approach, test it on real workloads, and avoid the surprises that show up at scale.

What does "enterprise-grade" cloud backup mean?

Enterprise-grade backup is not "we can store copies." It's about recovering quickly and cleanly at scale, without turning cost, compliance, and operations into a constant fire drill.

An enterprise-grade cloud backup solution should be able to:

1. Centrally manage and automate backup across multiple clouds, accounts/projects/subscriptions, and regions.
2. Restore reliably under pressure, not just in a demo.
3. Provide immutability, isolation, and auditability as table stakes.
4. Keep storage growth and operational overhead predictable over time.
5. Integrate with existing tools (identity, ticketing, SIEM, reporting) so backup ops fit your workflows.

If it can't do this in your real environment, it's not enterprise-grade.

What types of enterprise cloud backup solutions exist?

Most vendors say they're "cloud-ready." What matters is what you get by default, and what turns into extra tools, extra work, or extra cost later.

In practice, those options fall into four categories:

1. Cloud-native backup platforms (cloud infrastructure + data)
2. Hyperscaler-native tools
3. Orchestration layers for native backups
4. Legacy and hybrid platforms

Why do modern teams prefer cloud-native backup platforms for cloud infrastructure + data?

Cloud-native backup platforms emerged because they fit how cloud teams actually work: API-driven estates, distributed ownership, fast change, and lots of environments.

- They reduce operational drag: fewer (or no) appliances, proxies, or backup servers to size, patch, upgrade, and secure.
- They make governance possible at scale: one place to see coverage, enforce policy, and catch drift across orgs, regions, and services.

Within this category, the real differences tend to show up in three places:

- **Automated posture management:** Cloud Backup Posture Management (CBPM) how well the platform discovers new resources, enforces policy, and flags drift without relying on manual tagging.
- **Granular recovery in real life:** whether file/object/database recovery is practical day-to-day, not just "supported" in edge cases, so incidents aren't all-or-nothing.
- **Data usability without restores:** whether teams can search and query protected data for audits, investigations, analytics, and AI without spinning up full restores first.

What are hyperscaler-native tools best for?

Hyperscaler-native tools are built into a single cloud (AWS, Azure, or Google Cloud) and are typically snapshot-first for backup and restore.

Where enterprises outgrow them:

- Governance gets messy across many accounts/projects/subscriptions and regions.
- Most rely on snapshots, which can be durable, but often aren't optimized for long-term retention economics or everyday incident recovery.
- Recovery and security workflows are split across services, vault types, and settings.
- It's easy to turn on, but at enterprise scale, you often discover extra work: additional components, the "right" vault or storage setup to unlock key restore options, and extra policy setup to roll backups out consistently across accounts/subscriptions/projects.

What are orchestration layers for native backups best for?

These platforms standardize policies, scheduling, and reporting on top of hyperscaler snapshots. Examples include N2WS; the key point is you're still running snapshot-based backups, just with centralized scheduling and reporting.

Where they tend to break down at scale:

- Still snapshot-based: to get specific data, teams often end up restoring whole snapshots/volumes first.
- Limitations show up service-by-service (coverage, recovery granularity, restore paths).
- Cost visibility can stay fragmented because you're still paying through underlying cloud mechanics.
- "Centralized control" can still mean stitching together multiple backup behaviors under one UI.

When do legacy and hybrid platforms still fit?

Legacy and hybrid suites can be a good fit when you have a meaningful on-prem footprint and legacy workloads that need broad coverage.

Where cloud complexity exposes its limits:

- Customer-managed infrastructure (nodes, proxies, appliances) that must be sized, patched, upgraded, and secured.
- Recovery requires restoring to find data and is infrastructure-heavy, which is a mismatch for cloud incidents where teams need targeted fixes fast.
- Cloud scale turns the backup platform itself into another distributed system that you have to operate.
- Layered licensing and add-ons can balloon spend as coverage, retention, and security needs expand.

High-level comparison (by approach)

APPROACH	BEST FIT	WHERE IT TENDS TO BREAK DOWN
Cloud-native backup platforms (cloud infrastructure + data)	Cloud-first or multi-cloud, faster recovery needs	It's not right for enterprises with mostly on-prem environments
Hyperscaler-native tools	Single-cloud, smaller estates, basic restore needs	Governance at scale, fragmented workflows, snapshot-heavy recovery, cost visibility complexity
Orchestration layers for native backups	Standardizing policy/scheduling/reporting across native backups	Still snapshot-based underneath: costs can spike (copies/indexing/scanning), API throttling shows up at scale, and restores stay heavier than needed for everyday incidents
Legacy/hybrid suites	Significant on-prem footprint, broad legacy coverage	Cloud overhead, infrastructure management, restore-first workflows

What should you test when evaluating enterprise cloud backup solutions?

Use these as hands-on evaluation tests. The goal is not "does it exist," it's "does it hold up under real conditions."

1 Can it meet your RTO/RPO targets on realistic workloads?

Validate recovery steps, time-to-first-data, and failure modes on representative workloads.

2 Can you get a cost model you can actually sanity-check?

Get a clear breakdown of what drives cost (storage growth, copies, indexing/search, scanning, cross-region/cross-account requirements, and required infrastructure). Costs will still be estimates, so make sure the vendor explains what inputs they used and how sensitive the model is to retention, change rate, and restore/testing patterns.

3 Does it support multi-cloud operations without adding more moving parts?

This is especially important if you require multi-cloud backups and a single operating model for policy, recovery, and reporting.

4 Can it keep you continuously compliance-ready?

Test retention enforcement, isolation boundaries, auditability, and whether you can prove coverage without a quarterly scramble.

5 Can it restore meaningful workloads under pressure?

In a PoC or lab test, run restores that reflect your real world (encryption, size class, service type, cross-account/region patterns). Validate how many steps it takes, what needs to be pre-provisioned, and whether recovery stays predictable as the environment grows.

6 Does it support granular recovery for real incidents?

Granular recovery is what saves you during partial data loss, accidental deletes, and corruption. File-, object-, and database-level recovery should be practical, not a special project.

7 Is ransomware resilience built in, or assembled?

As part of its ransomware package, backup solutions should have immutability and isolation as the baseline, plus detection signals and the ability to recover from known-clean points without resorting to heroics.

Which operating models do today's vendors use?

Instead of a checkbox grid that goes stale, use a shortlist lens based on what the platform is designed to protect first:

- **Cloud-native backup (cloud infrastructure + data):** Eon
- **Hyperscaler-native:** AWS Backup, Azure Backup, Google Cloud native mechanisms
- **Orchestration layers:** Snapshot policy/scheduling/reporting tools (examples: N2WS and similar products)
- **Legacy/hybrid suites:** Veeam, Commvault, Rubrik, Cohesity

Cloud-native backup for cloud infrastructure + data

EON

Best for: Cloud-first enterprises protecting production infrastructure across AWS, Azure, and Google Cloud that want fast, precise recovery, strong ransomware resilience, predictable cost behavior at scale, and direct access to backup data without making full restores the default.

Eon is a cloud-native backup platform for cloud infrastructure + data built to remove customer-run backup infrastructure while improving day-to-day recovery, governance, and cost control.

It's built for protecting cloud infrastructure and designed to make protected data useful after it's backed up, not just stored.

Unlike restore-first tools, Eon is built so teams can find, inspect, and reuse protected data for audits, investigations, analytics, and AI workflows, without spinning up full restores first.

What to validate in a real-world test:

- **Recovery under pressure:** targeted recovery for real incidents (file/object/database-level), not just full restores
- **Cost behavior:** how storage grows over time and whether the platform reduces long-term overhead
- **Governance:** continuous discovery and policy enforcement without relying on manual tagging
- **Ransomware resilience:** immutability and logical isolation as baseline, plus clean recovery workflows
- **Data usability:** ability to search and query protected data directly for audits, investigations, analytics, and AI (turn backups into a live data lake)

Table stakes, baseline, and built in: compliance-grade retention beyond 35 days, immutability, logical air-gapped backups, cross-region/account recovery patterns, RBAC, and audit logs.

Legacy/hybrid backup platforms

These can be a fit for hybrid estates with significant on-prem and legacy workloads. For cloud-first teams, the key question is what you're signing up to operate.

What to validate (Veeam, Commvault, Rubrik, Cohesity):

- What customer-managed infrastructure is required, and how it scales.
- Whether granular recovery is consistent across cloud-native workloads.
- How costs accumulate across licenses, infrastructure, and cloud consumption.
- Whether policy and reporting keep up with constant cloud change.

Hyperscaler-native cloud backup

These tools are often the default starting point in a single cloud. At enterprise scale, friction usually shows up in governance, operating model complexity, and cost transparency.

Azure native backup: what should you check?
Validate:

- Which vault type you're using, and which resources it actually governs.
- How policies roll out across many subscriptions and regions, and how you keep them consistent.
- How cross-region recovery behaves in practice, and what configuration choices enable or restrict it.
- Whether any workloads require agents/extensions, and what that means operationally.

Google Cloud native mechanisms: what should you check?
Validate:

- Whether centralized backup requires additional operational components (appliances/connectors).
- Which services share the same backup operating model vs. require separate approaches.
- How you handle analytics and object storage protection, where "versioning/replication" patterns may not behave like true backup workflows in incidents.

AWS Backup: what should you check?
Validate:

- Whether you're relying on multiple backup mechanisms across services, and how you prove coverage.
- How isolation and recovery behave across accounts and regions (including copy requirements for restore).
- How optional capabilities like indexing, scanning, or specialized vault workflows affect cost and operations.
- Whether your database model is snapshot-first, PITR-first, or layered (many teams use native PITR for short windows, then another platform for longer retention and cyber recovery).

Case studies: what "right fit" looks like

These can be a fit for hybrid estates with significant on-prem and legacy workloads. For cloud-first teams, the key question is what you're signing up to operate.

What to validate (Veeam, Commvault, Rubrik, Cohesity):

- What customer-managed infrastructure is required, and how it scales.
- Whether granular recovery is consistent across cloud-native workloads.
- How costs accumulate across licenses, infrastructure, and cloud consumption.
- Whether policy and reporting keep up with constant cloud change.

NETGEAR

NETGEAR moved from an appliance-heavy legacy platform to Eon's SaaS-managed approach, cutting operational overhead and improving recovery speed while lowering backup storage costs by **35%** and improving restore speed by **88%**

SoFi

SoFi standardized backup operations across regions and improved visibility into posture and spend, reporting **100% & ROI** and improved resilience.

What's next?

If you're actively evaluating platforms, the fastest next step is a hands-on evaluation of real recovery, real governance, and real cost drivers.

See how Eon performs.

VISIT US AT WWW.EON.IO
INFO@EON.IO