



You Don't Own the SaaS Stack. You Still Need to Own Its History.

Why Cloud Resilience Is Not Business Recoverability

Commissioned by Afi



Whit Walters, Field CTO, GigaOm





The Recovery Question

A SaaS application can be fully available and still be wrong for the business. A deleted account, changed permission, damaged workflow, or compromised identity can propagate across a healthy platform. The service stays up while the organization loses the state in which it knows how to operate.

For the CIO, the recovery question is then not whether the provider can restore infrastructure. It is this: if the tenant were wrong tomorrow morning, could the organization determine what right looked like, prove what changed, and reconstruct the dependent state?

This is the boundary the business must govern. **Cloud resilience is provider-owned. Business recoverability remains customer-owned.** The enterprise therefore needs an independently controlled, searchable, historically accurate representation of the cloud operating state on which it depends.

What the CIO Actually Owns

The enterprise may not own the SaaS stack, but it still owns the consequences of its state. That state is broader than files. It includes content, identities, memberships, permissions, policies, schemas, versions, ownership, and the relationships that make an application usable. Recovering content without those relationships can leave every file present and the business unable to operate.



The recovery window is often shorter than the detection window. Native recovery and audit periods vary by service and license. Microsoft, for example, retains core Entra audit and sign-in data for seven or 30 days by license unless the customer exports it. Google permits restoration of a recently deleted Workspace user for 20 days and administrator recovery of deleted Drive data for 25 days across a selected date range. Those are operational windows, not a business retention strategy. A quiet permissions or policy change can remain undiscovered longer than either.



Application semantics determine whether recovery is useful. A directory is not a list of users; its security posture lives in memberships, roles, assignments, and policy. A collaboration site is not a folder; it is content plus structure, ownership, permissions, and workflow. A recovery product that counts objects without preserving their relationships is measuring coverage rather than recoverability.



Change velocity is now a board-level resilience issue. Automation, low-code applications, machine identities, and AI assistants create and modify business state faster than governance teams can document it. The result is a widening gap between what the environment is, what the organization believes it is, and what it can prove it was.



THE CIO'S RECOVERY TEST

A recovery control should answer four questions against the production estate, not against a demonstration tenant.

- **Independence.** If the production tenant or its identity plane were compromised, could the team still reach its recovery evidence and controls?
- **Fidelity.** For each material workload, which objects and relationships are restorable, backup-only, or export-only?
- **Sufficiency.** Does retained and searchable history outlast the time required to detect and investigate a quiet change?
- **Proof.** Can the team identify a known-good state, restore a dependent object set, and produce an honest exception report?

The Control the Business Actually Needs

Traditional backup starts with a copy. SaaS recovery must start with evidence. Before a team can restore anything, it must determine what changed, when it changed, which prior state was trustworthy, and which objects must return together. The useful control is therefore not merely another copy of cloud data. It is an independent operational memory of the enterprise.

Five criteria distinguish that control from a conventional backup claim.

- 1 **Independence: history must survive the system it records.** The protected record, its administrative controls, and the means of authentication should remain reachable when the production tenant is unavailable or hostile. Native retention can be valuable, but it shares more of the production control plane's identity, administrative boundary, and failure modes. That architectural dependency should be explicit in the risk decision.
- 2 **Fidelity: the data model must match the application.** SaaS state cannot be reduced to independent files. Buyers should require a workload-by-workload boundary that distinguishes restorable objects and relationships from data that is backup-only or export-only. No vendor can restore what a provider does not expose through its APIs; the honest response is to document the boundary, not bury it beneath a workload logo.
- 3 **Sufficiency: retained history must outlast realistic detection.** A policy expressed in months is not enough. The relevant measure is the shortest searchable history across the material state, compared with the organization's time to detect and investigate a non-obvious change. If evidence expires before the problem is understood, the nominal retention period has no operational value.
- 4 **Verifiability: evidence must be inspectable before restoration.** A green job status proves that a process ran. It does not prove that the business can recover. Operators need to search, preview, compare versions, and identify a known-good point without first staging a restore. The resulting report should distinguish complete, partial, unsupported, skipped, and failed outcomes.
- 5 **Operability: the control must work at the estate's real scale.** Multi-tenant administration, delegated roles, residency, key control, legal holds, and user lifecycle are not peripheral enterprise features. They determine whether the recovery system can be governed without creating standing privilege or an operating burden the organization cannot staff.

Why Now?

SaaS has crossed a threshold. It is no longer primarily a delivery model for productivity software; it is where the enterprise's operating context accumulates. Identity establishes who can act. Permissions determine what they can reach. Collaboration structures encode how work moves. Low-code applications and AI systems add another layer of state that changes continuously and is difficult to reconstruct from documentation.

At the same time, concentration has increased. A smaller number of cloud control planes mediate a larger share of business activity, while the organization has fewer independent artifacts from which to reconstruct that activity. The operational benefit of consolidation is real. So is the need for a recovery control outside the consolidated failure domain.

The arrival of native SaaS backup offerings, alongside published partner architectures for third-party protection, confirms the market requirement. It does not settle the architecture decision. A CIO still has to decide how much independence is required, which state must be recoverable rather than merely retained, how long evidence must remain searchable, and whether recovery access should survive loss of the production identity plane. The buying question is no longer whether SaaS data deserves protection. It is what kind of recoverability the business is prepared to rely on.


Why Afi?

Afi is a cloud-native data-protection platform for Microsoft 365, Google Workspace, and cloud infrastructure workloads, built for the gap between provider durability and customer recoverability. The platform is designed for cloud data sources rather than adapted from appliance-era backup. That distinction matters for reliable recovery at enterprise scale because cloud recovery must model provider APIs, object relationships, identity, cloud security, search, tenancy, and rate limits rather than treat the environment as a collection of files and volumes.

Afi's primary differentiator is its ability to protect large, complex environments at enterprise scale. Afi protects not only files and other basic content, but also the structures, relationships, permissions, metadata, policies, and application-specific objects required to reconstruct a usable cloud workload. This high-fidelity, application-aware model supports rapid search and access across protected data and enables granular, high-speed recovery of complete operating structures rather than merely restoring disconnected files. For organizations whose business state is embedded in Microsoft 365, Google Workspace, and other cloud applications, that depth can materially reduce recovery time and complexity.

Security is another core part of Afi's differentiation. Afi treats cloud data protection as a broader security problem encompassing identities, permissions, APIs, metadata, administrative controls, and recovery paths—not simply the storage of a secondary data copy. This approach is designed to preserve the integrity of both protected data and the mechanisms used to access and recover it.

Its commercial model is aligned with that origin. Afi sells a subscription based on protected active users, permits customers to protect a subset of users, and retains archived accounts without requiring the user to remain licensed in the source SaaS platform. The business model monetizes the proposition that a business should maintain an independently controlled, searchable, historically accurate representation of its cloud operating environment. The unit of value is not storage alone; it is durable access to business history.



Where Afi Fits

Afi should be evaluated as a specialist, not as a smaller vendor attempting to reproduce every incumbent backup feature. Its strongest fit is an organization that values independence, application-aware recovery, cloud native protection, enterprise-grade security, long-term searchable history, delegated administration, and control over residency and encryption. This assessment is based on vendor briefing and product documentation review rather than a scored product evaluation or independent performance benchmark.

The architecture provides credible answers to the independence test. Customers can configure third-party SAML authentication so recovery access does not depend on Microsoft 365 or Google Workspace identity. Lockdown mode restricts changes to access, backup policy, secrets, and deletion behind a separately stored code and an audited 15-minute unlock window. Afi also offers an optional object-storage retention lock that periodically advances an active lock period of approximately one month. While active, protected objects cannot be changed or deleted by any party, including Afi. Afi states that audit events for data access operations are retained for three years. These are meaningful controls, but they are configurable. Buyers should verify that the proposed operating model enables and governs them.

The product also shows evidence of application-aware recovery. Afi documents recovery of selected Entra ID users, groups, roles, conditional access policies, administrative units, and associated memberships and assignments, while identifying other directory objects as exportable rather than restorable. For SharePoint in-place recovery, Afi documents preservation of original IDs for items still present, together with recovery of libraries, lists, subsites, schema, metadata, and permissions. It also provides a changing view across a selected snapshot range. Those distinctions are more useful than an undifferentiated coverage claim because they let the buyer test the actual recovery boundary.

Its delegated operating model has been exercised at material scale. The company reports that its largest enterprise deployment protects more than 500,000 users. In Afi's published French Red Cross case study, the organization protected roughly 90 TB, archived 60,000 users under residency and GDPR requirements, and scoped administrative access across a distributed operating structure. A vendor case study is not independent proof, but it identifies the right reference questions for a buyer with multiple regions, administrators, or legal jurisdictions.

The diligence obligations remain. Afi is a specialist competing with much larger platform and data-protection vendors. CIOs should examine financial durability, support capacity, recovery throughput at their own scale, API-dependent gaps, data portability, and exit terms. They should also insist on a written matrix showing, for every material workload, what is restorable, backup-only, export-only, or unsupported. Afi belongs on the shortlist when those boundaries meet the business requirement, not simply because its architectural thesis is attractive.

Analyst's Take

The category error in SaaS protection is treating it as a storage problem. The CIO's exposure is not limited to missing files; it is the loss of a trustworthy operating state. Once identity, permissions, policy, structure, and workflow are included, backup becomes a business-control decision shared by the CIO and CISO.

Afi merits consideration because its architecture and commercial model are organized around that broader object of recovery. It is neither unique in protecting SaaS nor complete; no vendor is. Its differentiation lies in the combination of deep, application-aware support for complex cloud workloads; rapid access, search, and granular recovery; operational independence; and enterprise-grade security delivered through a cloud-native architecture. Searchable history, application semantics, and archived-state economics are central to the proposition rather than secondary features on a storage platform.

For the CIO, the next step is not to accept the thesis or reject it. It is to test it. Compromise the production identity assumption in the proof of concept. Ask the team to locate the last known-good state without restoring it. Recover a dependent object set, not a single file. Read the exception report. Confirm the exit path. If a product cannot pass those tests, it is a copy of data, not a control over business recoverability.

The principle stands regardless of vendor: **the enterprise may choose not to own the SaaS stack, but it should still own the history required to understand and reconstruct the business.**

Report Methodology



THE GIGAOM PERSPECTIVE BRIEF provides an independent analyst view of an emerging technology decision. This brief was developed from vendor briefings, Afi product documentation, public pricing and customer material, platform-provider documentation, and GigaOm market coverage, including the Radar for Cloud Data Protection Solutions v1 and the Radar for Managed Backup as a Service. Product claims attributed to Afi have not been independently benchmarked. Buyers should validate recovery fidelity, performance, and operational controls in their own tenant.



ABOUT THE AUTHOR

Whit Walters is a Field CTO at GigaOm with 30 years of enterprise technology experience, including CTO roles at multiple startups, leadership positions within the Google Cloud partner ecosystem, and domain expertise in AI, machine learning, data architecture, and cloud security.



GIGAOM



GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.