Lucidity

Buyer's Guide to Cloud Storage Optimization

A Comprehensive Framework for Evaluating a Modern Solution

Introduction:
The Cloud Storage Surge

Cloud storage is no longer just an IT necessity — it's a strategic enabler for digital transformation, AI, analytics, and compliance initiatives. Data volumes are exploding, and storage costs are rising faster than overall cloud spend. At the same time, most enterprises face a hidden efficiency problem: over-provisioned disks, under-utilized storage resources, and the wrong disk tier for the workload's needs.

Native cloud provider tools offer basic visibility, but they rarely surface actionable insights or enable intelligent automation or right disk sizing at scale. The result: growing waste, degraded performance, and increased operational complexity.

The good news? Advances in observability, analytics, and autonomous optimization now allow organizations to match storage resources precisely to workload needs — without downtime — delivering both immediate cost savings and long-term performance improvements.

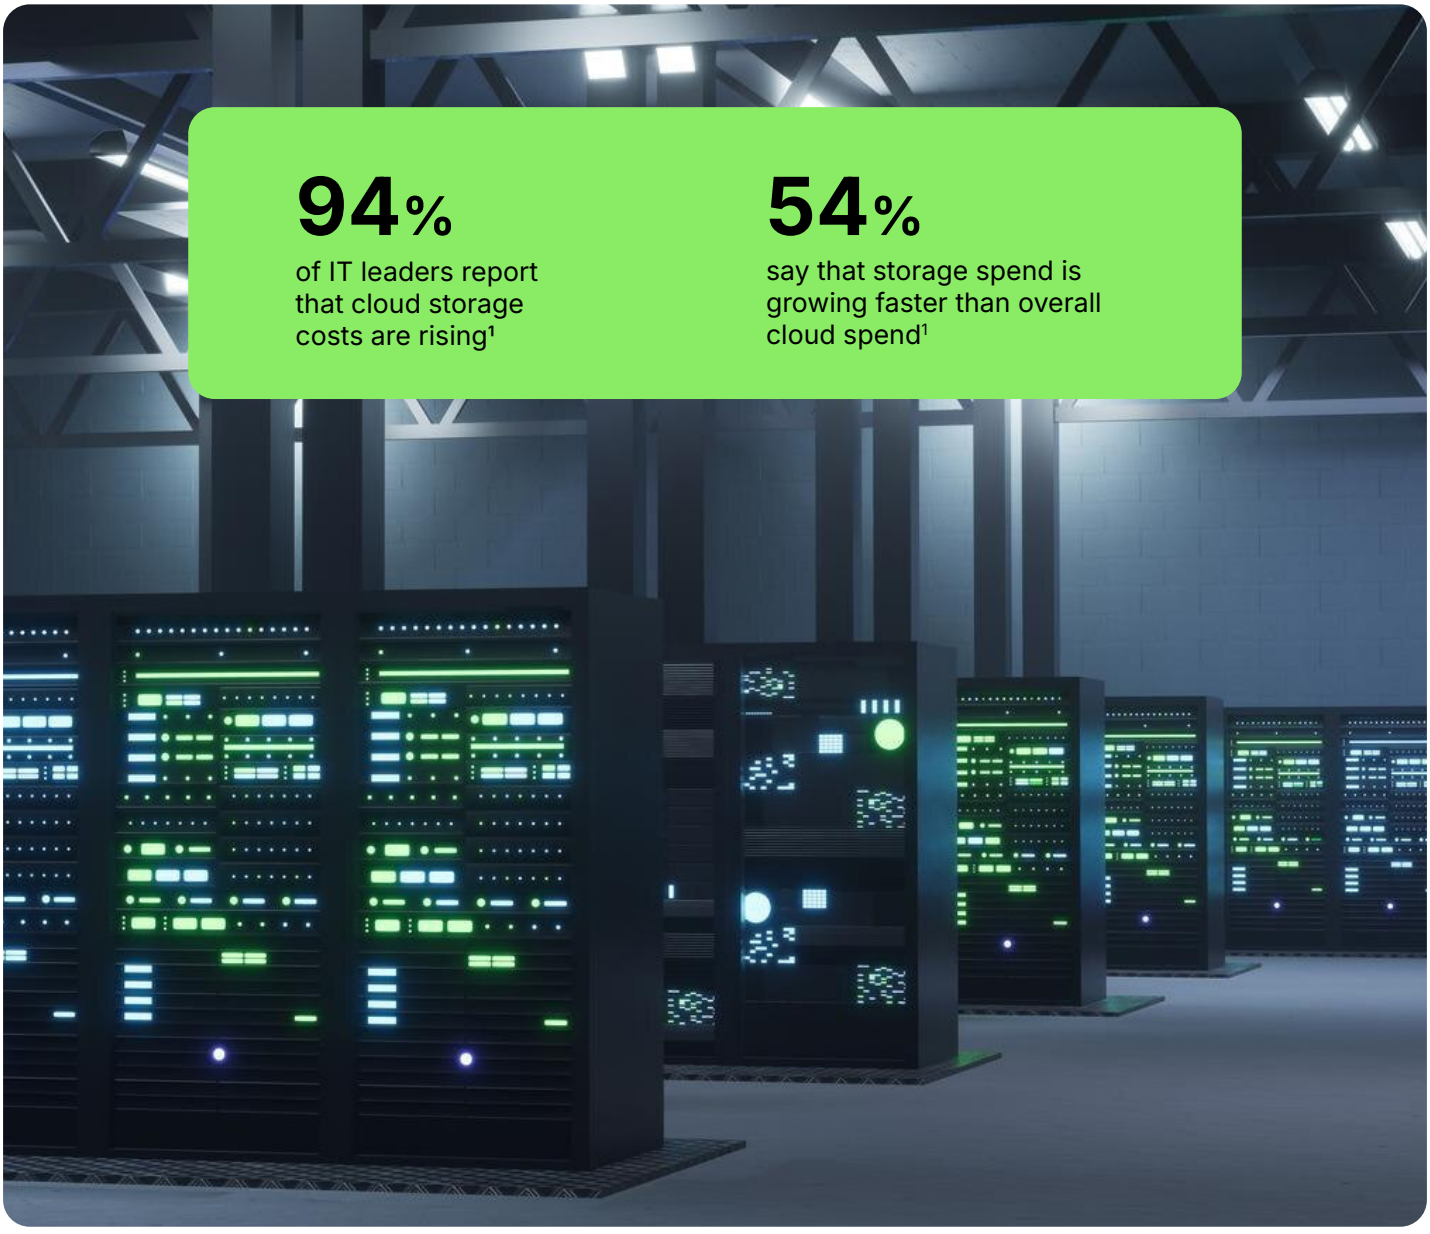

The Case for Change: Current Challenges in Cloud Storage

Despite significant investment in cloud infrastructure, many organizations are still flying blind when it comes to storage utilization and optimization. Common pain points include:

Limited visibility into actual utilization

Without a real-time view into how storage is consumed, organizations are left guessing which resources are delivering value and which are wasted. This blind spot allows inefficiencies to grow unchecked, leading to inflated costs.

Lack of automation forcing manual intervention

Manual resizing and tiering are labor-intensive and prone to delays. This reactive model prevents teams from optimizing quickly enough to capture cost savings in real time.

Over-provisioned block storage leading to stranded capacity

Teams often allocate more storage than applications truly require, just to avoid performance risks or due to commercial terms offered by their cloud service providers. Over time, this "just in case" approach strands vast amounts of unused capacity.

Fragmented management across CSPs and storage types

With separate consoles, policies, and reporting for each provider, managing multi-cloud storage becomes a complex and error-prone process. These silos make it difficult to apply consistent optimization practices.

Idle and orphaned volumes consuming budget

Disks that are no longer attached to active workloads — or remain idle for extended periods — can silently drain budget. In many cases, these resources persist unnoticed for months.

Inability to scale optimization efforts

As storage footprints grow, manual processes and disconnected tools simply cannot keep pace. The result is that optimization lags behind consumption, compounding inefficiencies over time.

Wrong disk tier selection causing waste and performance issues

Selecting a higher-cost tier for a workload that doesn't require it wastes budget, while under-tiering can slow application performance. Without intelligent mapping of workload needs to disk type, both risks are common.

The majority of the organizations utilize ***only 30%*** of their cloud block storage[2]

Traditional approaches — periodic manual reviews, spreadsheets, and ad hoc resizing — cannot keep up with the speed and complexity of modern, multi-cloud workloads. The result is an urgent need for a unified, automated, and intelligent approach to cloud storage optimization.

[1] *Virtana Research - The State of Hybrid Cloud Storage 2023.*
[2] *Lucidity assessment data*

Evaluation Criteria for a Modern Cloud Storage Optimization Solution

Unified Multi-Cloud Visibility and Management

Managing cloud storage across AWS, Azure, and Google Cloud through separate consoles creates inefficiencies and inconsistencies. A unified management platform consolidates visibility and policies across all providers in one interface, simplifying operations. This ensures ITOps, FinOps, and Platform Engineering teams have a single source of truth for cost, performance, and compliance.

Questions to ask your vendors:

- Can your platform manage block storage across AWS, Azure, and GCP from a single interface?
- Does your solution apply optimization policies consistently across providers and storage types/tiers?
- How do you consolidate reporting for multi-cloud costs, utilization, and performance?

Real-Time Utilization Analytics

Static or delayed reporting often leaves teams reacting to problems after waste has already accumulated. Modern solutions provide real-time visibility into IOPS, throughput, and storage utilization. This allows organizations to proactively detect under-utilized resources and make timely adjustments before costs escalate

Questions to ask your vendors:

- How frequently is utilization data refreshed, and at what level of granularity (VM, volume, IOPS)?
- Does your solution correlate storage consumption with application performance requirements?
- Can you provide anomaly detection or trend analysis to identify potential waste before it gets out of control?

Autonomous Scaling Without Downtime

Traditional resizing requires manual intervention and often leads to application disruption. A modern optimization solution must support autonomous scaling — up or down — without downtime. This ensures applications always have the right amount of storage capacity, balancing performance and cost seamlessly.

Questions to ask your vendors:

- Can your platform resize live storage volumes dynamically without disrupting workloads?
- What safeguards are in place to prevent scaling actions from impacting application performance?
- Do scaling operations require manual approval, or can they be fully policy-driven?

Intelligent Tiering

Workloads don't all require the same performance characteristics, but many organizations overpay by running everything on premium storage due to a lack of visibilty, insights, or available options at their cloud service provider. Intelligent tiering matches workloads to the most cost-effective disk type, ensuring optimal performance at the lowest cost. Automated tier migration keeps applications aligned with changing demands over time.

Questions to ask your vendors:

- How does your solution determine the optimal tier for a workload?
- Can your system migrate workloads between tiers without downtime?
- Are recommendations based on historical data, predictive analytics, or both?

Orphaned and Idle Volume Detection

Unmounted disks or idle volumes can silently drain thousands of dollars each month if left unchecked. Modern platforms should automatically identify these resources and recommend their decommissioning or repurposing. Eliminating stranded storage is one of the fastest ways to reduce waste and reclaim budget

Questions to ask your vendors:

- How does your platform identify idle or orphaned volumes across cloud environments?
- Can the system distinguish between temporary low-use volumes and truly unused resources?
- Does the solution provide automated remediation, or does it only generate alerts?

API-First Cloud-Native Integration

Bolt-on solutions or connectors add latency and complexity, undermining efficiency. A truly cloud-native solution integrates directly with CSP APIs to manage native storage services across regions and disk types. This approach reduces risk, accelerates operations, and ensures future compatibility as cloud providers evolve their offerings.

Questions to ask your vendors:

- Does your solution integrate natively with CSP APIs, or does it require third-party gateways?
- How do you ensure compatibility when cloud providers release new storage types or regions?
- Can your API layer be used to extend automation into our CI/CD or FinOps workflows?

Petabyte-Scale Support

As enterprises grow, storage environments can reach petabyte scale, making manual or fragmented tools untenable. Modern solutions must scale linearly with no performance degradation, ensuring optimization is as effective at 10 PB as it is at 10 TB. This scalability prevents bottlenecks and future-proofs the storage optimization strategy.

Questions to ask your vendors:

- How does your platform scale as storage environments grow from terabytes to petabytes?
- Are there any performance trade-offs at large scale, such as longer response times for analytics?
- Can you provide customer examples where your platform manages multi-petabyte environments?

Compliance Enforcement and Auditability

Data residency, sovereignty, and regulatory requirements like GDPR and CCPA demand precise control over where data lives. Optimization tools must not only enforce these rules but also provide audit-ready logs of every action taken. This transparency builds confidence with auditors, regulators, and customers alike.

Questions to ask your vendors:

- Can your solution enforce data residency or sovereignty rules across multiple regions?
- How are audit logs generated, stored, and exported for compliance reporting?
- Can the system provide alerts if a workload is deployed in violation of compliance policies?

What Modern Cloud Storage Optimization Looks Like

In a modern approach, optimization is continuous rather than periodic. Intelligent policies monitor workloads in real time, automatically adjusting capacity and tier to match evolving needs. The result is a unified view across clouds and a storage environment that adapts instantly to business demands — without downtime or disruption.

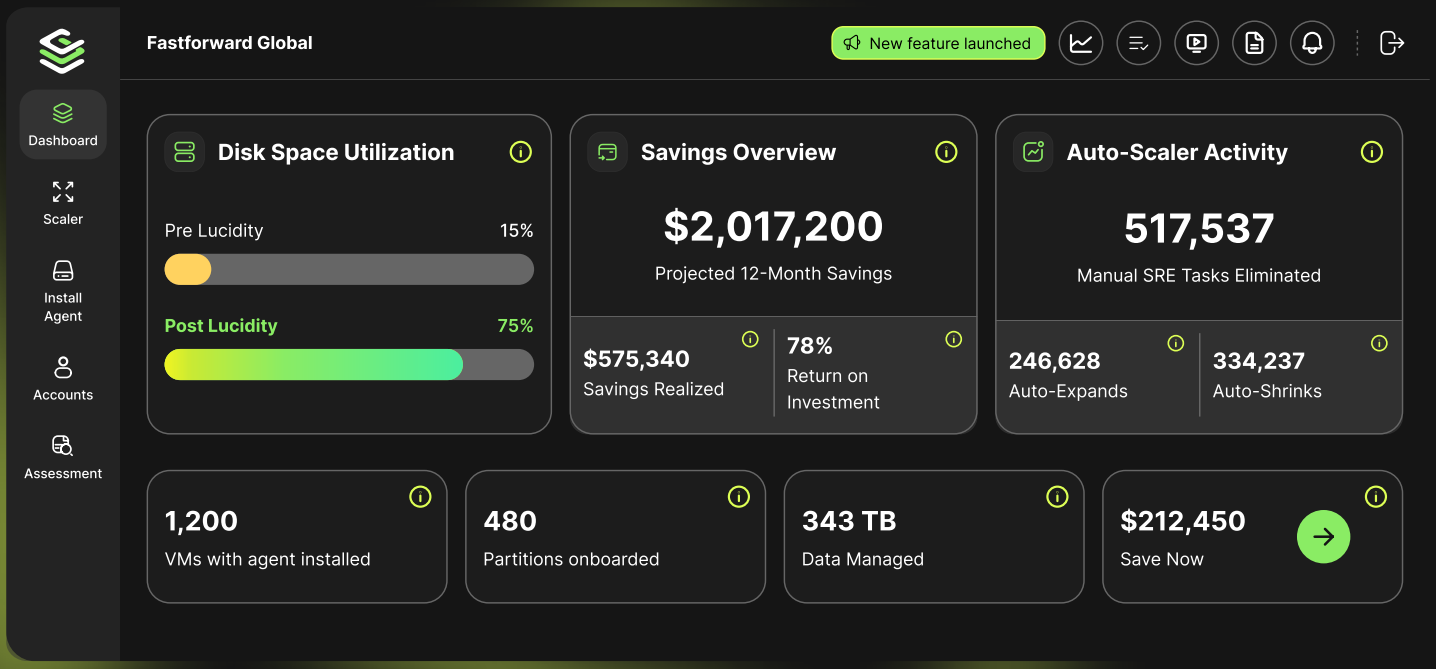

To learn more about Lucidity's Cloud Storage Optimization solutions, visit Lucidity at **www.lucidity.cloud**

Buyer’s Checklist for Cloud Storage Optimization

Requirement	Why It Matters	Vendor 1	Vendor 2	Vendor 3
Unified multi-cloud visibility	Reduces management overhead and enforces consistent optimization policies			
Real-time utilization analytics	Enables proactive optimization before waste builds up			
Autonomous scaling without downtime	Maintains performance and controls costs dynamically			
Intelligent tiering	Matches storage class to workload needs			
Orphaned/idle volume detection	Eliminates waste from unused resources			
API-first CSP integration	Avoids complexity of gateways and bolt-ons			
Petabyte-scale support	Ensures scalability without performance loss			
Compliance enforcement	Avoids penalties and meets regulatory needs			