S&P Global

451 Research Pathfinder Paper

May 2026

Modernizing recovery

Strategic recommendations for managing risk and cost

Commissioned by BROADCOM®

Executive summary

Recovery infrastructure is now a strategic resilience issue, not just a backup storage decision. Outages remain common, and security incidents have made recovery harder, more visible and more costly. In this environment, organizations need to move beyond architectures built mainly for low-cost backup retention and modernize their resiliency infrastructure for fast, clean and cost-aware recovery.

Recovery speed must become a design priority — both technically and operationally. This modernization requires organizations to keep recent backups locally on higher-performance infrastructure and move cold data to lower-cost tiers while improving coordination across backup, infrastructure, security and application teams under a shared operational model. The goal is not just to store more backup data at lower cost; it is to restore business operations faster and reduce the total cost of disruption.

Modernizing recovery infrastructure will not prevent every outage or attack, but it can significantly improve an organization's ability to respond when disruption occurs. In a period of constant data growth, tighter budgets and greater cyber risk, that is becoming one of the most important measures of resilience.

Key findings

Recovery operations remain challenging for most. About a third of respondents (35%) say it takes significant effort to synchronize and restore primary systems after a storage failure, up from 27% in the previous year. An additional 45% say it takes moderate effort to complete recovery.

Outages remain costly. The mean cost of outages came in at just over $2 million in the most recent study, marking a slight decrease from $2.3 million the previous year. According to the survey, 35% of outages cost over $1 million, while 3% of respondents say their most recent outage cost more than $10 million.

Security issues, software failures and cloud service failures are the top causes of recent outages. In the study, 23% of respondents say security issues caused their most recent outage, just ahead of software failures at 22%.

Hybrid cloud-based data protection continues to rise. Hybrid cloud deployments are the most common choice for data protection (37% of organizations), while cloud-based services such as online backup and DRaaS are used by 33% of organizations, and on-premises deployments are chosen by 30% of organizations.

Resilience issues facing organizations

Figure 1: Data resiliency dashboard

Voice of the Enterprise: Storage, Disaster Recovery 2025.
Source: 451 Research from S&P Global Energy Horizons.

Backup costs are growing — and so is the pressure to justify them

Data is among the most valuable assets for any organization, particularly as AI workloads make large, high-quality datasets central to business operations. But this value comes with a cost. Average annual data growth is approaching 30%, and organizations are retaining more than ever before — not just for operational needs, but also for compliance, analytics and AI model training.

This growth is driving greater capital and operating expenditures for backup and recovery. More data means more capacity, more replication, more retention and more operational overhead — including spending on power, cooling, floor space, administration and support. In a period of rapidly escalating hardware costs, organizations cannot afford to maintain existing backup strategies without thoughtfully considering alternatives. The question is not simply how to store more data — it is whether the current architecture is delivering recovery value commensurate with the cost of maintaining it.

Equally important: IT leaders must determine whether existing backup infrastructure is meeting service-level agreements for its core task — recovering from an outage. In many cases, the answer is no.

Outages are expensive, and they are becoming more frequent and longer-lasting

Outages remain a regular problem. About a third of organizations have had an outage in the last two years. In other words, disruption is not rare, and the frequency is increasing, driven by cyberattacks and growing infrastructure complexity. These outage incidents cost organizations millions of dollars, and four out of five report moderate to significant effort required to recover (see Figure 1). For many IT teams, this has become part of normal operations. A key question is whether the organization can recover fast enough to limit the damage.

Downtime disrupts revenue, delays customer service, interrupts employee work and creates uncertainty — and the longer it lasts, the worse these effects become. This is why recovery speed matters more than backup speed alone. An organization can complete backups on time and still fail during recovery if restore processes are too slow or too manual. Recovery is where resilience is truly tested.

Data growth is making recovery harder and more expensive

According to 451 Research's Voice of the Enterprise: Storage, Budgets & Optimization 2026 survey, the average organization expects its data volume to grow by nearly 30% in the next year, and this rapid growth is making recovery more difficult. As organizations create and keep more data, they also increase the amount of data that may need to be scanned, validated, moved and restored during an outage.

This creates both technical and cost challenges. On the technical side, larger datasets put more pressure on storage, networks and restore workflows. Recovery windows grow with the amount of data under management and with increased complexity as workloads evolve and interact with more applications and end users — making performance bottlenecks more likely while making dependencies harder to track and manage.

In many cases, storage systems built mainly for low-cost retention are now being asked to support fast recovery. That is often a poor fit. As environments scale, the cost of this mismatch compounds. Organizations end up paying not just for more storage, but for the operational burden of managing infrastructure that was never designed for the recovery performance they now need.

As a result, scale is no longer just an issue of storage capacity; it is also — and perhaps more importantly — one of recovery cost. The bigger the environment becomes, the more expensive it is to rely on architectures that were not designed for fast recovery. Addressing this requires rethinking how storage tiers are used and how spending is aligned to recovery value rather than retention volume alone.

Security incidents have changed recovery requirements

Security issues are now a leading cause of outages, and they make recovery more complex. In a non-security-related failure, the goal is to restore data and applications. During a cyber incident, the organization must also avoid reinfection, verify that backups are clean and make sure restored systems are safe to bring back into production. Cyberattacks continue to become more complex, making recovery more time-consuming and costlier for organizations, in turn forcing vendors to stay focused on improving the speed and effectiveness of their tools.

This fundamentally changes the operating model. **Recovery can no longer be the responsibility of only the backup or infrastructure team. It requires close, real-time coordination across security, infrastructure and application teams, and this coordination challenge is one of the most critical issues in modern resilience.** In many organizations, backup, infrastructure and security teams still use different tools and follow different processes. That creates significant delays during an incident. In a ransomware event, time is often lost not only because recovery infrastructure is slow, but because teams lack a shared view of the problem or a common response process. Breaking down these organizational silos is as important as any technology upgrade, and this is required to provide real-time coordination across security, infrastructure and application teams.

Though many organizations leverage the assistance of their vendors during incidents, they do not have the luxury to sit idly, waiting for the vendor to send in specialists and fresh hardware for a recovery, as they risk losing revenue and customers.

Recovery costs rise sharply as outages last longer

Recovery costs go far beyond lost data or short-term productivity. The average outage can cost more than $2 million, and the total impact is often broader. Lost revenue, damaged reputation, weakened customer loyalty, compliance penalties and even job loss can follow a major incident.

Critically, most of these costs are not fixed but escalate over time. A short outage may be manageable, but a long outage is far more damaging. Customers may leave. Employees may be unable to work. Backlogs grow. Regulators may take notice. And the organization's reputation can suffer long after systems return. Security incidents intensify this pressure further: They draw negative public, regulatory and executive attention to the organization in ways that a routine outage does not. That scrutiny raises the stakes for every hour of delay.

This is precisely why rapid recovery is not merely a technical aspiration but one of the most important levers an organization can use to control the total cost of a disruption. Every improvement in recovery speed translates directly into reduced business loss, while extensive downtime on the scale of weeks or months can have dire, existential-level consequences for an organization.

Required capabilities for resiliency modernization

Make recovery performance a design priority

The first step in modernization is straightforward: Recovery performance must be a core design requirement, not an afterthought. Many backup environments were built mainly for retention efficiency. That led to hard drive-based systems, data deduplication and architectures focused on storing large backup volumes at low cost. These systems may work well for retention, but they often perform poorly when recovery is needed.

Legacy backup storage typically depends on hard drives and deduplicated backup repositories. These reduce storage costs and expand effective capacity but also dramatically slow down recovery speeds. Hard drives are slower than flash storage, especially for random access, and deduplicated data often must be rehydrated before it can be used. That increases delay, especially when large amounts of data or multiple applications must be restored simultaneously.

Modern backup storage takes a different approach. It keeps recent backups on higher-performance local storage — which is often flash-based and not deduplicated — so they can be restored more quickly. **Beyond raw speed, flash-based infrastructure offers additional operational benefits: lower failure rates relative to spinning disk, reduced physical footprint and meaningfully lower energy consumption. As power and cooling costs rise sharply across data center environments, the energy efficiency of flash-based recovery infrastructure becomes an increasingly important factor in total cost of ownership calculations. This alone does not eliminate every recovery challenge, but it alleviates one major bottleneck — slow access to backup data — while reducing operational overhead in the process.**

But recovery speed is not only a technology problem. It is also an operational one. Even the fastest infrastructure will underperform if teams lack a unified model for responding to incidents. Backup, infrastructure and security teams must be able to act from a shared view of the environment — with clear roles, common tooling and coordinated processes. Modernizing the technology without modernizing the operating model leaves significant recovery performance on the table.

Replace siloed storage and backups with hybrid cloud architecture

Figure 2: Organizations favor cloud over on-premises for data protection and long-term archiving

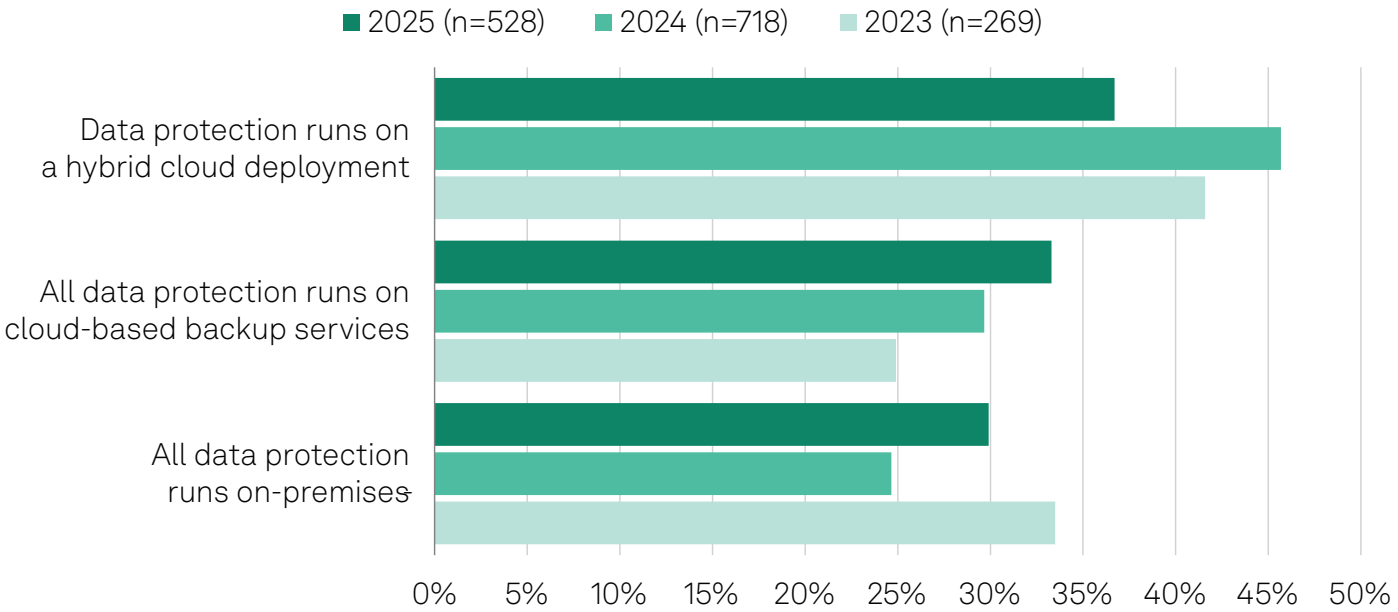

Q. Which of the following best describes your organization's current use of data protection (e.g., backup, disaster recovery)?
Base: All respondents (n=528).
Voice of the Enterprise: Storage, Disaster Recovery 2025.
Source: 451 Research from S&P Global Energy Horizons.

Figure 3: Long-term archiving data is headed for public cloud

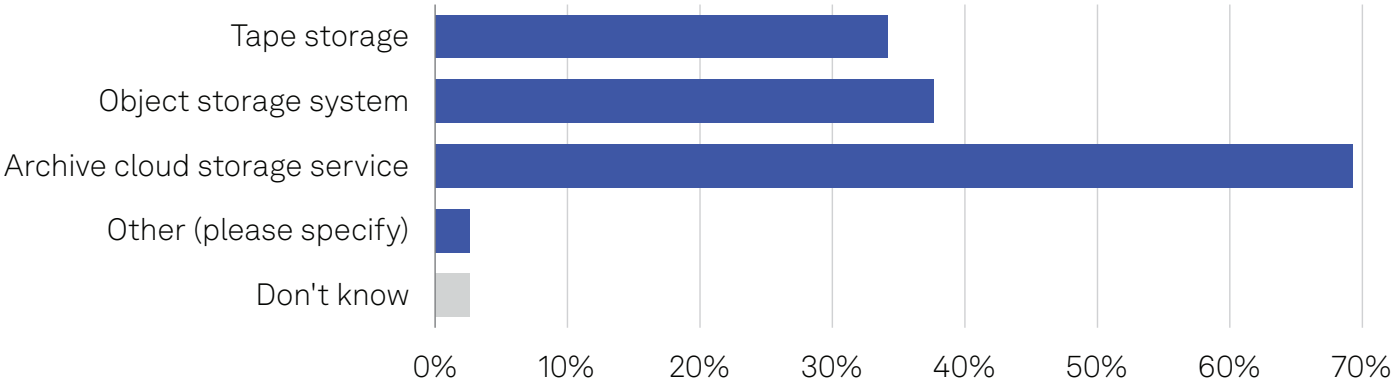

Q. What storage technology or service does your organization use for long-term data archiving? Please select all that apply.
Base: Respondents whose organizations archive data in a lower cost storage system (n=231).
Voice of the Enterprise: Storage, ESG 2025.
Source: 451 Research from S&P Global Energy Horizons.

Modernization also requires a broader architectural shift. Many organizations still manage backup storage, archive storage and cloud storage as separate silos. This increases complexity, raises costs and makes it harder to match infrastructure spending with business value.

Legacy media environments carry more than hardware costs. Tape and spinning-disk systems require ongoing media failure management, physical maintenance and significant space, power and cooling resources. Some purpose-built backup appliances can be expensive relative to the recovery performance they deliver, particularly when compared with commodity hardware or cloud-based approaches. These operational costs accumulate over time, and they often crowd out investment in more capable recovery infrastructure.

In our Voice of the Enterprise: Storage, Disaster Recovery 2025 study, 30% of organizations reported keeping all their data protection resources on-premises, while 37% cited a hybrid cloud deployment, and 33% chose cloud-only services such as disaster recovery as a service (DRaaS) and online backup (see Figure 2). With a hybrid cloud model, recent backups and high-priority recovery data can remain on faster local infrastructure, while long-term retention moves to public cloud storage or other lower-cost tiers. This aligns spending with recovery value: Data needed for rapid recovery stays on fast infrastructure, while data kept mainly for compliance or long-term retention moves to cheaper tiers. Critically, offloading cold data to low-cost storage tiers can help fund the modernization of primary recovery infrastructure, freeing budget that would otherwise be absorbed by legacy operational costs.

A unified operational model across on-premises environments, cloud environments and storage use cases also accelerates incident response. During a ransomware event, when speed and clarity are paramount, a unified hybrid architecture reduces manual effort, simplifies policy management and shortens the time needed to identify clean recovery points and initiate restoration workflows. Equally important, it enables backup, infrastructure and security teams to operate from a shared context during a crisis, rather than each team navigating its own tools and data views in parallel. The operational benefits of this integration are most visible precisely when organizations need them most.

Align recovery with cyber resilience

Immutable snapshots and replication to a DR site are table stakes, but infrastructure improvements alone are not enough. Recovery modernization must also support cyber resilience — because in a cyber event, the goal is to restore clean operations while also eliminating reinfection risks. This modernization must occur across both technology and operations.

On the technology side, organizations should consider new innovations such as isolated clean rooms. The industry has delivered solutions that simplify the provisioning, managing and securing of these environments and end-to-end workflows that integrate all the critical components and steps of cyber recovery.

On the operations side, organizations should also develop stronger coordination across backup, infrastructure, security and application teams, as well as better visibility into the environment, including infrastructure assets, security status, application dependencies and, increasingly, the AI-related systems and agents tied to business operations.

Organizations should choose integrated data protection offerings with unified visibility and coordination and cyber recovery capabilities for a best-in-class solution with superior speed of recovery. Applications depend on identity systems, APIs, cloud services and data pipelines. A recovery process that restores storage but ignores these dependencies may bring workloads back in an incomplete or unstable state. Solutions should align across backup, infrastructure, security and application teams to reduce confusion during an incident because this informs teams on which backups can be trusted, which systems should come back first, how to validate cleanliness and how to restore operations without introducing new risks. That is both a technology challenge and an operating model challenge.

Future considerations

AI will add more pressure to infrastructure and data protection environments

As generative and agentic AI use grows, organizations will create and manage even more data. They will also need faster access to that data and stronger governance around how it is used. From a recovery standpoint, AI creates four main pressures: It increases data volumes; it strains infrastructure budgets; it increases demand for fast data access, which makes slow recovery systems less practical; and it raises the importance of governance, lineage and policy control.

AI may also make recovery more complex. As AI services become part of normal business operations, organizations may need to recover more than traditional applications and databases. They may also need to recover vector stores, orchestration layers, pipelines and policy controls. This broadens the scope of resilience planning.

Recovery modernization should not be treated as a narrow storage upgrade — it is part of a larger architecture shift.

Content provided by BROADCOM®

Recovery modernization is urgent. VMware vSAN delivers on all four dimensions that matter:

Cost Savings: vSAN delivers a recovery TCO 27% below comparable all-flash alternatives with industry-standard servers and a unique licensing model

Speed of Recovery: vSAN's all-flash, NVMe-based architecture delivers dramatically faster restores than legacy spinning-disk systems — ideal for mission-critical applications

Cyber Resilience: Replicate any VCF VM to a dedicated on-premises cyber recovery cluster, protect it with immutable snapshots, and accelerate recovery with VMware Advanced Cyber Compliance (sold as an add-on)

Operational Simplicity: A unified operational model and high degree of automation meaningfully lowers time to complete storage operations

Visit vmware.com/products/cloud-infrastructure/vsan or speak with a Broadcom sales representative to learn more.

About the author

Henry Baltazar

Research Director, Storage

Henry Baltazar is research director of the Storage channel at 451 Research from S&P Global Energy Horizons with a focus on data storage. In his current role, Henry analyzes the market trends around environmental, social and governance (ESG) storage challenges, infrastructure modernization and resiliency. He publishes reports on trends in data storage, disaster recovery and hybrid cloud. He is often cited as a subject expert by publications such as MIT Technology Review, Forbes and TechTarget.

Henry arrived at S&P Global through its 2019 acquisition of 451 Research, where he began working as an analyst in August 2006. After spending three years running the storage research practice at Forrester, he returned to 451 Research in 2015 to fill the research director role and lead the storage practice.

Henry graduated from the University of California, Berkeley with a bachelor's degree in environmental sciences.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About S&P Global Energy

At S&P Global Energy, our comprehensive view of global energy and commodities markets enables our customers to make superior decisions and create long-term, sustainable value. Our four core capabilities are: Platts for pricing and news; CERA for research and advisory; Horizons for energy expansion and sustainability solutions; and Events for industry collaboration.

S&P Global Energy is a division of S&P Global (NYSE: SPGI). S&P Global enables businesses, governments, and individuals with trusted data, expertise, and technology to make decisions with conviction. We are Advancing Essential Intelligence through world-leading benchmarks, data, and insights that customers need in order to plan confidently, act decisively, and thrive economically in a rapidly changing global landscape.

Learn more at www.spglobal.com/energy.

CONTACTS

www.spglobal.com
www.spglobal.com/en/enterprise/about/contact-us.html