How Financial Institutions Unlock More Value from Data with Object Storage

An FStech and Hitachi Report on Object Storage

FStech In collaboration with HITACHI

Executive Summary

Object storage has moved beyond a low-cost repository into an intelligent engine for AI and lakehouse ecosystems. While adoption is universal, the research reveals a critical gap: 65 per cent of firms still prioritise cost over the performance and intelligence required to scale modern data strategies.

To fully benefit from the technology, organisations need to balance cost with high-performance metadata and automation. Should they treat the technology as a storage expense, infrastructure will become a significant barrier in implementing advanced AI systems.

Introduction

Object storage has quietly evolved within financial services. Once seen mainly as a low-cost repository for unstructured data, it is turning into a foundational element of enterprise data strategies. This reflects a broader shift in how institutions view data not just as something to store and protect, but as an asset to be analysed and monetised.

Adoption is widespread across on-premises, cloud, and hybrid environments, but organisations continue to balance competing priorities, including regulatory compliance, data sovereignty, cost control, and scalability. Approaches vary: some institutions focus on compliance and archiving, others embed object storage into data lakes and AI pipelines, and many pursue hybrid models that offer flexibility while introducing operational complexity. These trade-offs illustrate the tension between maintaining control and unlocking the strategic value of data in a rapidly evolving landscape.

This report explores how financial institutions are using object storage today and how they are positioning it for the future. It reveals a sector that recognises the strategic value of data but is still working to fully operationalise it.

Methodology

FStech surveyed 100 financial services decision-makers across a range of organisations, including retail banks, investment firms and payment providers who are all actively utilising object storage in their business. The research examined how object storage is currently used, the maturity of deployments, key selection criteria, and expected future developments.

The survey also explored how object storage supports AI and analytics initiatives, the role of hybrid and multi-cloud architectures, and the impact of data sovereignty requirements on deployment decisions.

Please note: Percentages may not total 100 per cent due to rounding or multiple-choice questions.

Key Findings

1. Object storage is moving beyond archive into core data infrastructure

44 per cent of organisations use object storage for hybrid or multi-cloud workloads, 35 per cent for cloud-native or API-driven applications, and 27 per cent for data lake or lakehouse architectures. This indicates a clear shift beyond traditional archive and compliance use cases towards active data environments supporting analytics and AI.

2. Hybrid deployment is the dominant model, but adds complexity

38 per cent of organisations use object storage across both on-premises and cloud environments, making hybrid the most common approach. While this supports flexibility and sovereignty requirements, it introduces challenges around governance, visibility, and data consistency.

3. Adoption maturity is split, creating a two-speed market

36 per cent of organisations remain in early-stage adoption, while 35 per cent report enterprise-wide use. This highlights a widening gap between firms building scalable data platforms and those still operating in siloed pilots.

4. Cost dominates, but value-oriented capabilities are rising

65 per cent of organisations cite cost as a primary selection factor. However, resilience, availability, and integration with analytics platforms are gaining importance, signalling a shift towards value-driven decision-making.

5. AI ambition is high, but foundational readiness lags

Object storage supports AI or analytics in all organisations, with 43 per cent placing it at the centre of their strategy. Despite this, relatively low prioritisation of AI-ready platforms and data hubs suggests that many institutions lack the governance and data architecture needed to scale effectively.

Contents

1. Current approaches to object storage

Financial institutions now treat object storage as a strategic platform, underpinning initiatives from regulatory compliance to AI and advanced analytics. The three operating models reflect different approaches to balancing control and scalability, each with distinct advantages and trade-offs.

Organisations using on-premises object storage prioritise control, predictability, and regulatory assurance. This model is well suited to environments where data sovereignty, security, and latency are non-negotiable, enabling direct oversight of sensitive data and close alignment with regulatory expectations.

On-premises deployments prioritise control and regulatory assurance but limit scalability and access to AI-ready compute resources. Overreliance on this model risks a defensive data strategy focused on storage rather than value creation, potentially leaving institutions behind as innovation accelerates.

At the opposite end of the spectrum, organisations using cloud-based object storage exclusively are optimising for scalability, flexibility, and speed of innovation. This model supports rapid experimentation and is well suited to AI, analytics, and cloud-native workloads, while also reducing infrastructure overheads through access to integrated services.

Cloud-only deployments offer scalability and speed but introduce risks, including vendor lock-in, cost volatility, and potential conflicts with data sovereignty. Tightening regulations could challenge purely cloud-based strategies and constrain architectural flexibility.

The largest group of organisations follow a hybrid model, reflecting the operational reality of financial services, where the need for control must be balanced with demand for scalability.

Hybrid architectures balance control with scalability, enabling sensitive data to remain on-premises while leveraging the cloud for analytics, AI, and customer-facing workloads. This model supports nuanced strategies but introduces operational complexity.

However, this model introduces operational complexity. Organisations must manage latency differences across environments, reconcile inconsistent governance policies, and integrate diverse security controls. Data replication and synchronization across sites can increase overhead, while inconsistent monitoring and visibility can obscure performance bottlenecks. Addressing these challenges requires disciplined orchestration and tooling to ensure the hybrid strategy delivers value rather than friction.

Mitigating hybrid complexity requires structured approaches, such as defining standard governance policies, using automated metadata and orchestration tools, and monitoring latency across environments. A simple checklist of controls can help ensure hybrid deployments support performance and compliance objectives.

Across models, adoption choices reflect trade-offs between control, scalability, and strategic value, highlighting that object storage is as much a governance and operational challenge as a technology decision.

Which of the following best describes your organisation's current use of object storage? [select one option]

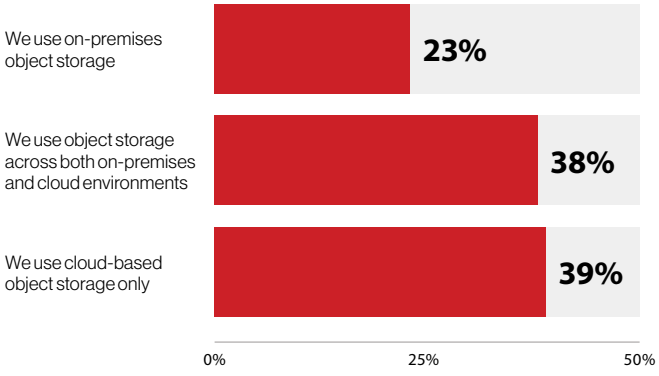

2. Maturity of object storage adoption

The survey shows a near-even split in adoption maturity. Enterprise-wide users have moved beyond experimentation, integrating object storage into core data architectures to enable cross-functional access and advanced analytics.

Early-stage adopters, constrained by legacy systems and silos, illustrate a two-speed market and the uneven pace of digital transformation. Such organisations risk falling behind as competitors build more sophisticated data capabilities. Those further along the maturity curve face different challenges, particularly around governance, scalability, and integration.

Maturity correlates with measurable business outcomes. Enterprise-wide adopters report faster AI deployment cycles, reduced operational costs, and more efficient cross-team data access, whereas early-stage adopters experience slower time-to-insight and higher risk of siloed data. A visual summary of maturity versus use case adoption could illustrate this two-speed market clearly.

Scaling from pilot to enterprise deployment requires organisational change, not just technology. Fragmented data ownership, limited system integration, and weak governance can block enterprise-wide success.

Without standardised access, consistent data management, and alignment across teams, object storage remains a point solution rather than a platform. This creates a risk that organisations overestimate their progress, mistaking isolated success for enterprise readiness.

Bridging this gap requires a more deliberate approach. Object storage must be integrated into end-to-end data pipelines, supporting multiple workloads rather than single use cases. This shift is necessary to establish a cohesive data foundation capable of enabling innovation at scale.

Where object storage is widely deployed, the benefits extend beyond storage efficiency. It becomes a shared data foundation, enabling data to be accessed and reused across teams rather than confined to individual use cases.

This supports faster innovation, as teams can work with existing datasets rather than building new pipelines. It also improves decision-making by enabling insights that span business units, rather than being limited to siloed data.

Centralising data within object storage allows for more consistent governance, strengthening security, compliance, and lifecycle management. In this context, object storage evolves from infrastructure into a core enabler of enterprise-wide data strategy.

Which of the following best describes the maturity of your organisation's use of object storage? [select one option]

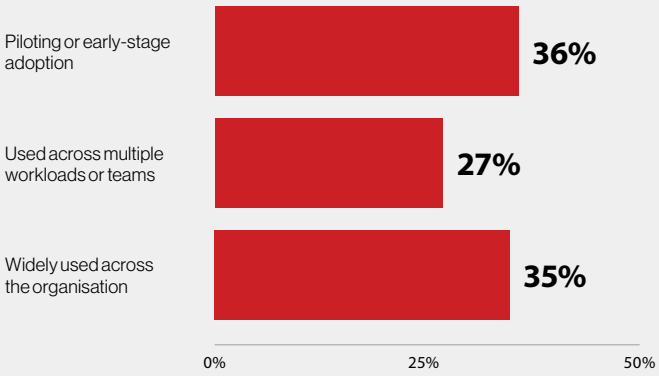

3. Primary uses of object storage

Object storage now supports diverse workloads, spanning traditional archiving, AI, and cloud-native applications.

This breadth reflects the flexibility of the technology, but also highlights the challenge of using a single storage foundation to meet both regulatory obligations and innovation goals.

At the traditional end, 18 per cent of respondents cite regulatory compliance and records retention as a primary use case. Object storage remains well suited to this role, with capabilities such as immutability, long-term retention, and policy-driven governance supporting regulatory requirements.

The limitation is that compliance-led environments can become rigid. Data stored primarily for retention is often difficult to access, integrate, or reuse. As regulatory requirements increasingly intersect with analytics and AI, governance frameworks must enable broader data usage rather than restrict it.

A further 34 per cent use object storage for long-term archive and cold storage, reflecting its cost efficiency and scalability. This supports the retention of large volumes of historical data without the expense of traditional systems.

The risk is that cold storage reinforces a passive approach to data. Information is retained but rarely used, creating a growing pool of underutilised assets and limiting the ability to extract value as analytics capabilities mature.

By contrast, 31 per cent use object storage for active archives or high-value data repositories. This signals a shift towards treating stored data as a source of insight rather than simply a compliance requirement.

IoT data emerges as a major use case, with half of respondents citing it. In financial services, this covers telemetry from payments systems, digital banking channels, transaction monitoring, and real-time fraud detection. Supporting high-volume streaming requires robust metadata, governance, and performance management; without these, institutions risk unstructured, difficult-to-use data estates. The position of IoT as one of the most popular use cases reflects the strategic shift from passive storage to active, metadata-rich environments.

Institutions must embed object storage into daily workflows, balancing compliance and security with accessibility for analytics and AI to unlock value and mitigate governance risk.

What are the primary use cases your organisation uses object storage today [select top three]

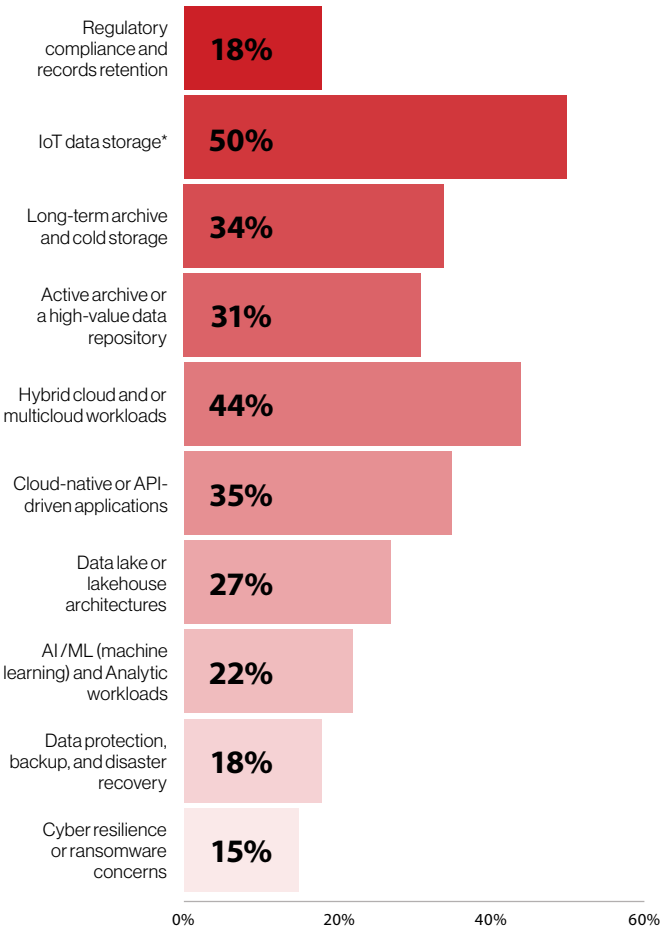

**In financial services, IoT refers to the high-velocity telemetry from digital banking channels and real-time fraud detection.*

4. Selection factors for object storage platforms

Cost dominates platform selection (65 per cent), reflecting object storage's role as a scalable, cost-efficient solution for unstructured data.

This emphasis also exposes a potential blind spot. Prioritising cost can lead to underinvestment in capabilities required to extract value from data. Features such as metadata management, analytics integration, and automation are often treated as secondary considerations despite being critical for enabling active use cases.

This creates a false economy. While storage costs may be reduced, organisations can incur indirect costs through increased complexity, limited data accessibility, and missed opportunities for insight generation. The value of object storage lies in enabling data usage, not simply reducing storage spend.

Beyond cost, 46 per cent of respondents prioritise data resilience and availability. This reflects the role of object storage in supporting business continuity, where durability, protection against data loss, and rapid recovery are essential.

Integration with analytics, AI, and cloud-native platforms is also a key consideration, selected by 35 per cent of respondents. Platforms that integrate effectively enable organisations to move from storing data to using it. Poor integration creates bottlenecks, often requiring data movement or duplication, which increases both cost and complexity.

Only 15 per cent prioritise scalability without sacrificing performance. Given that scalability is a core characteristic of object storage, this suggests that some organisations accept performance trade-offs as data volumes grow.

As object storage assumes a central role, institutions should balance cost with performance, functionality, and integration to maximise strategic value.

Which factors are most important when selecting an object storage platform? [select top three]

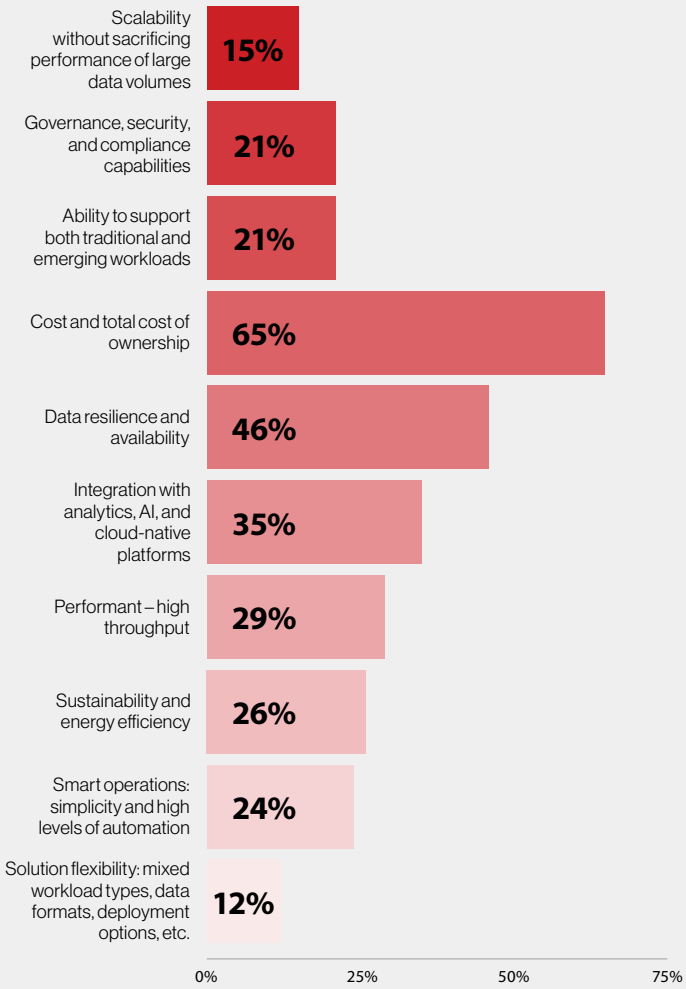

5. Object storage capabilities

The survey indicates that object storage is increasingly supporting analytics, business intelligence, and AI workloads. Integration with analytics tools is identified as a key capability by 39 per cent of respondents, highlighting the growing importance of enabling data-driven decision-making.

Data is now an active asset, valued for insight and business impact rather than mere compliance or retention. This shift places greater emphasis on accessibility, organisation, and governance. Many organisations invest in storage capacity but underinvest in the capabilities required to make data usable.

Advanced metadata, analytics integration, lifecycle management (48 per cent), and policy-driven automation (27 per cent) directly increase operational efficiency and accelerate AI deployment by ensuring data is organised, governed, and accessible.

Advanced metadata and data classification are identified as priorities by 35 per cent of respondents. These capabilities are critical for enabling searchability, governance, and data reuse. However, they are often not natively embedded within object storage platforms and instead rely on additional tools.

The survey highlights that modern object storage delivers value through capability, not just capacity. Organisations gain competitive advantage by investing in metadata, classification, automation, governance, and analytics integration. Without investment in metadata, automation, and integration, simply increasing storage volume will notnot improve intelligence or outcomes.

These gaps have tangible consequences. For example, fragmented governance or insufficient AI-ready platforms can delay AI project completion, while duplicated or inaccessible datasets increase operational overhead and slow model training. Quantifying these impacts helps highlight the cost of underprepared infrastructure.

Organisations need to move beyond scale towards looking at how intelligently data can be organised, governed and activated. This reinforces the idea that the strategic value of object storage lies not just in how much it can hold, but in how well it helps organisations use what they already have.

Which object storage capabilities are most valuable to your organisation today? [select top three]

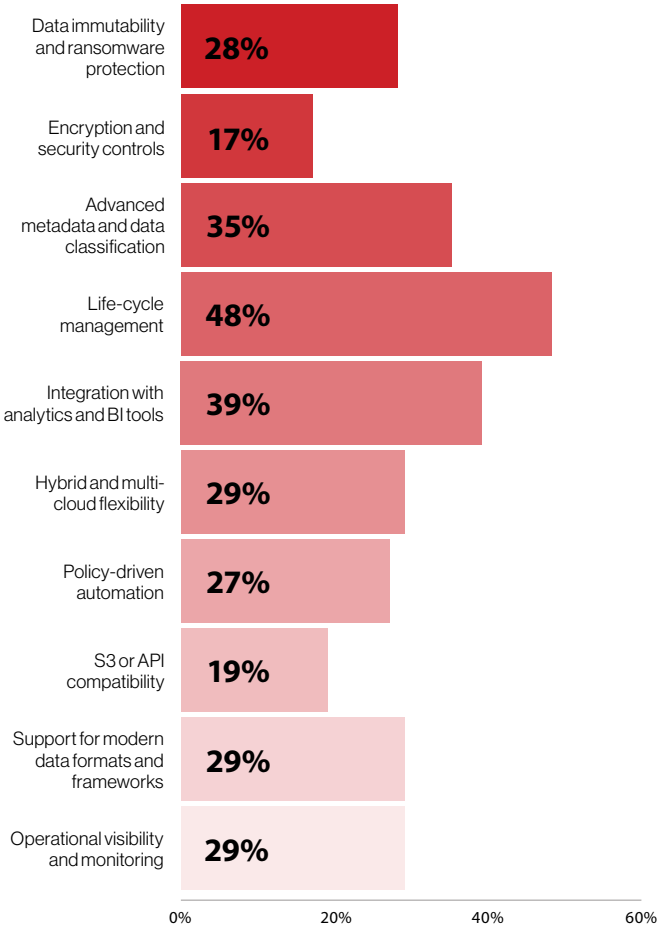

6. The evolving role of object storage

All respondents expect the role of object storage to evolve, indicating sustained investment and growing strategic importance.

Nearly half expect improvements in performance and accessibility—critical for analytics and AI. Without architectural upgrades, scaling for projected data growth may create bottlenecks in model training and real-time analytics.

Forty per cent anticipate broader use across hybrid and multi-cloud environments. This aligns with the need to balance control, cost, and scalability, as institutions continue to distribute workloads across different environments.

One-third anticipate more active, revenue-generating use cases, signalling recognition of data as a value asset – but many still lack the foundations to support this shift.

Only 26 per cent expect increased support for AI, analytics, and data science, indicating a gap between strategic ambition and near-term operational readiness.

As object storage becomes more central to data strategies, expectations are rising. Meeting them will require not only incremental improvements, but also more deliberate architectural alignment with performance, governance, and integration requirements.

How do you expect the role of object storage to evolve in your organisation over the next 1-2 years? [select all that apply]

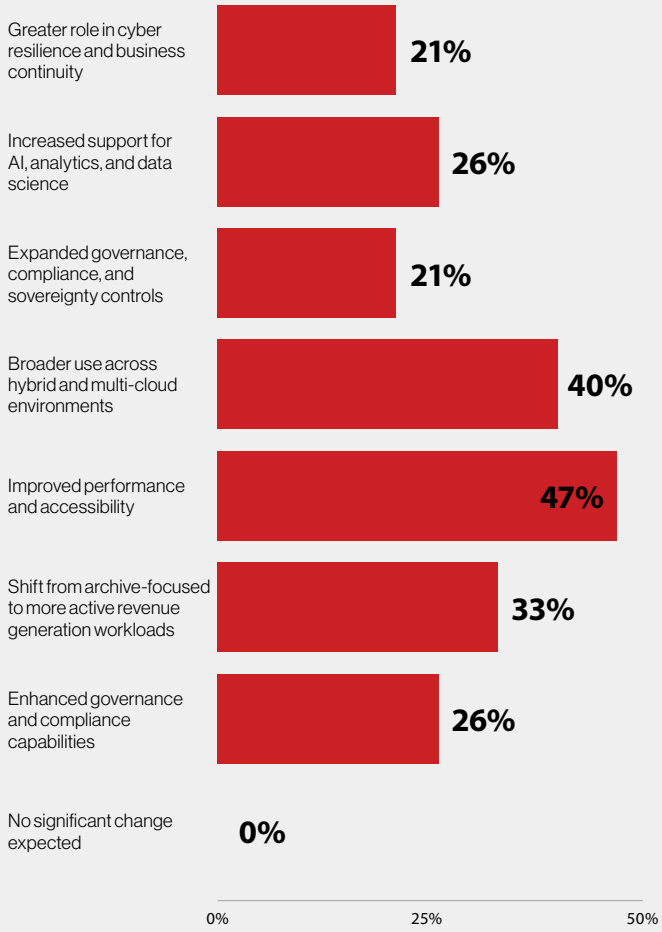

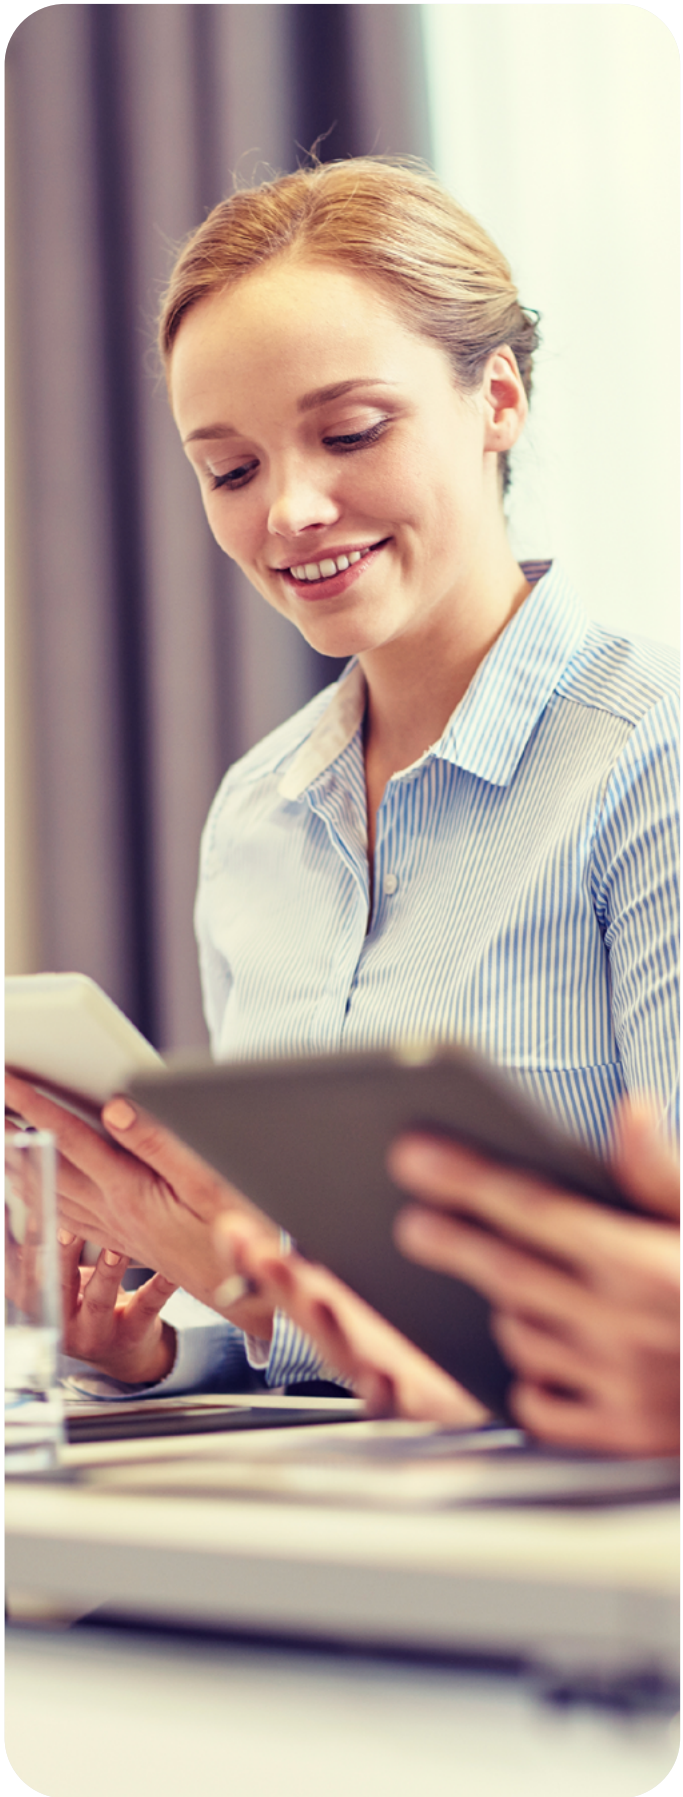

7. Object storage in AI and analytics

Object storage is embedded in AI and analytics strategies, with 43 per cent of organisations treating it as central and the remainder supporting selected workloads.

In many cases, object storage forms part of a broader data lakehouse architecture, supporting how data is stored, accessed, and used within machine learning workflows.

For 57 per cent using it selectively, object storage supports specific AI or analytics workloads without a unified strategy, offering flexibility for experimentation but limiting broader impact.

Over time, however, it can create fragmentation. Data may be duplicated across environments, and AI initiatives can remain siloed, increasing both cost and complexity while slowing progress.

When central to AI strategy, object storage enables shared, reusable datasets, supporting scalability and consistency. This however requires mature governance, metadata management, and analytics integration.

Where object storage is central to AI strategy, the benefits are more significant. Data can be accessed, shared, and reused across models and applications, supporting scalability and consistency.

This model requires greater maturity. It depends on strong governance, effective metadata management, and seamless integration with analytics platforms.

Using object storage only as a support component may suffice short-term, but long-term AI ambitions require it to be fully integrated into governance, data pipelines, and analytics workflows.

What role does object storage currently play in your organisation's AI or advanced analytics initiatives? [select the most appropriate answer]

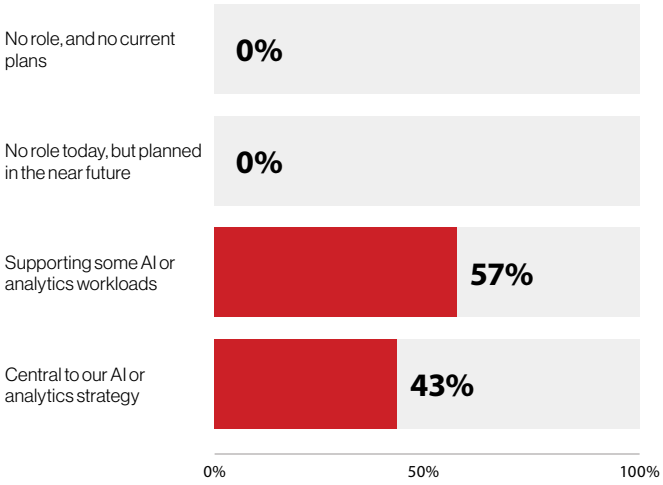

8. Future use cases of object storage

Cyber resilience and long-term retention (44 per cent) remain top priorities, reflecting the increasing importance of protecting data against growing cyber threats.

Interest in AI, machine learning, and generative AI signals a shift toward value-driven use. Organisations aim to leverage data for both protection and insight, maximising its utility.

This dual role introduces complexity. Architectures designed for retention are not always optimised for high-performance analytics, creating potential friction as organisations attempt to support both objectives.

Expanding object storage to support AI, generative AI, and cloud-native workloads further amplifies operational demands. Ensuring consistent performance, maintaining metadata integrity, and coordinating governance across distributed environments are critical to prevent fragmentation and inefficiency.

Only 26 per cent of respondents expect object storage to support data lake or lakehouse platforms. This suggests that some organisations are focusing on advanced AI outcomes without fully establishing the underlying data foundations required to support them.

Continued interest in API-driven applications, data ecosystems, and hybrid cloud enablement indicates that object storage is becoming more integrated into application environments. This supports flexibility and innovation, but increases the importance of effective integration and governance.

Failure to invest in integration, governance, and performance risks fragmentation, duplicated data, and operational inefficiencies, limiting the value of expanding use cases.

Which use cases do you expect object storage to support more actively over the next 1–2 years? [select all that apply]

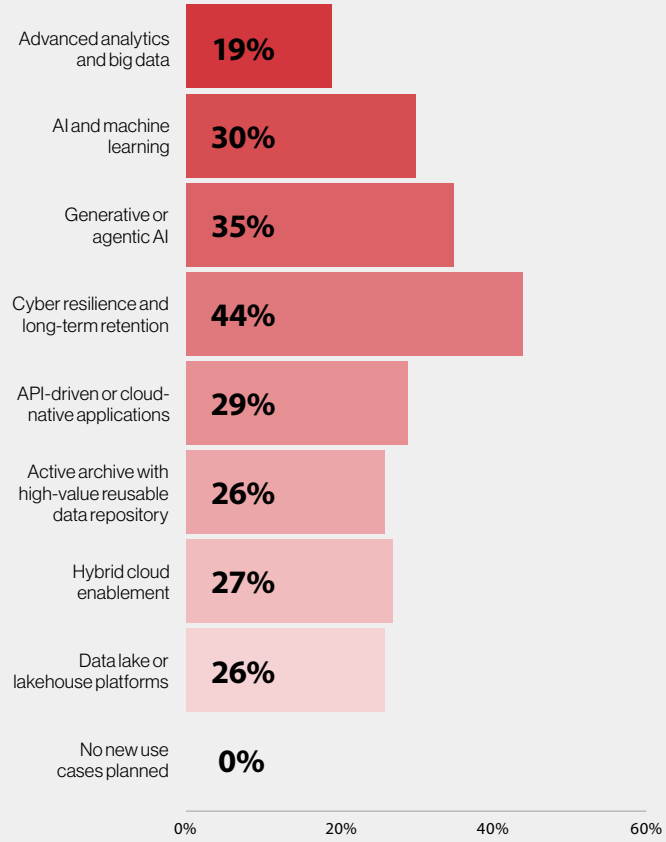

9. Near-term strategic priorities

Respondents prioritise pragmatic operational tasks: managing data growth (35 per cent), improving accessibility, ensuring compliance, and modernising platforms. This is critical in a sector with rising volumes and regulatory demands.

AI-ready platforms (10 per cent) and centralised data hubs (9 per cent) are low priorities, indicating a focus on optimisation over building foundations for future capabilities.

Focusing solely on incremental improvements risks leaving architectural fragmentation unaddressed, potentially exposing inconsistencies rather than improving accessibility.

If near-term priorities focus exclusively on operational tasks like cost reduction and growth management, underinvestment in AI-ready platforms or centralised data hubs could slow future transformation. A simple consequence matrix – mapping operational focus against long-term readiness – can illustrate the risks of an imbalanced approach and guide strategic decision-making.

Performance and responsiveness are low priorities (4 per cent). As workloads become more active, underestimating performance needs could create bottlenecks in analytics and AI.

Managing growth and reducing costs remain essential, but they are not sufficient on their own. Without investment in structural capabilities, organisations risk building more efficient versions of constrained systems.

What are your organisation's top data storage priorities over the next 12 months [select all that apply]

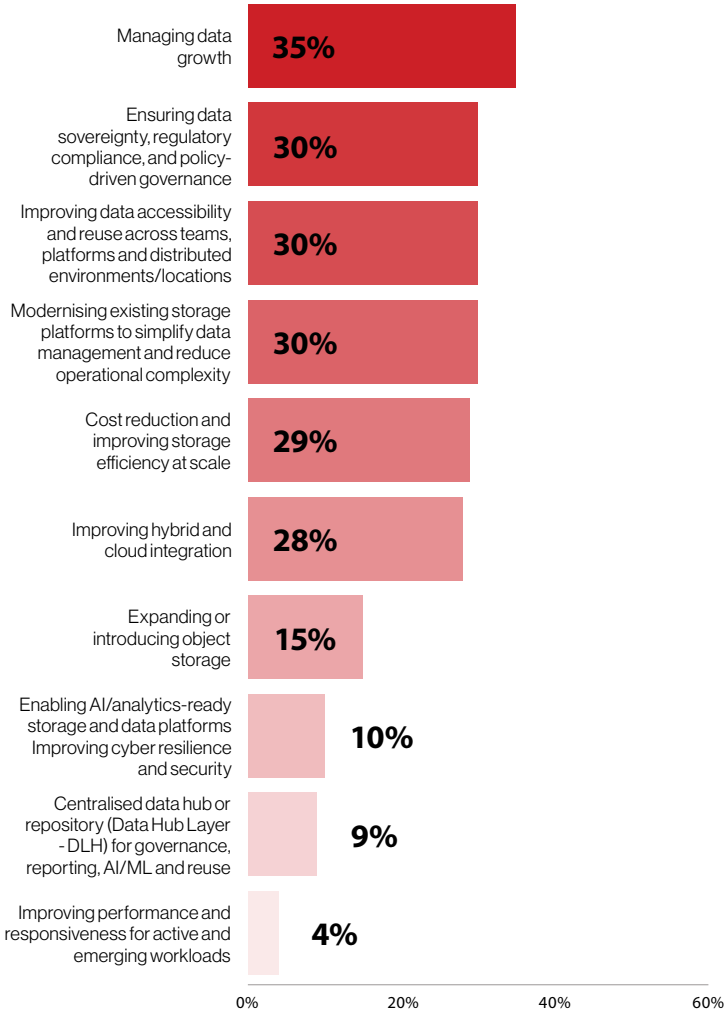

10. Current deployment models

Deployment models vary widely, reflecting the complex balance of control, cost, and scalability across financial institutions.

Each model has trade-offs: managed services reduce overhead, on-premises ensures control, and hybrid approaches offer flexibility.

Managing multiple deployment models increases operational complexity, especially with differing governance, access, and performance requirements, risking fragmentation without a unifying strategy.

Low adoption of multi-cloud (6 per cent) and public cloud-only (4 per cent) signal caution in externalising storage environments due to control, cost, and regulatory concerns.

Hybrid and private cloud models dominate, reflecting efforts to balance control, compliance, and scalability.

How is object storage currently deployed within your organisation? [select all that apply]

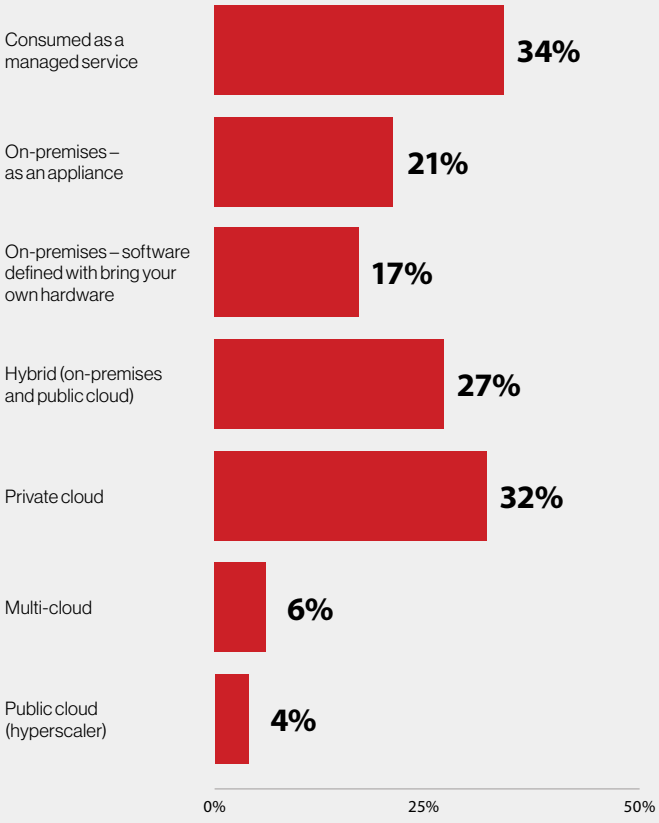

11. Data sovereignty and AI workloads

Data sovereignty significantly shapes how financial institutions deploy object storage for AI, with approaches ranging from regin-specific workloads to splitting training and inference across locations.

In highly regulated environments, data cannot move freely, directly influencing how AI models are developed and deployed.

Keeping data local supports compliance but can restrict access to large-scale compute, while centralising workloads improves efficiency but may conflict with sovereignty requirements.

These tensions result in architectural trade-offs. Models may be trained in one location while data remains in another, or workloads may be distributed across multiple environments. While this enables compliance, it can introduce latency, increase complexity, and reduce overall performance.

Nineteen per cent of respondents report that sovereignty requirements already limit scalability or performance, highlighting a tangible operational impact.

Sovereignty constraints shape AI strategy by limiting where models can be trained and forcing trade-offs between compliance and performance. Mitigation approaches include regional data hubs, federated learning, and edge compute, enabling organisations to balance regulatory requirements with operational efficiency.

Sovereignty must be addressed at the architectural level. Without this, organisations risk building systems that are compliant but inefficient, limiting their ability to scale AI initiatives effectively.

How does data sovereignty influence where and how you run AI workloads using object storage? [select one that best describes your situation]

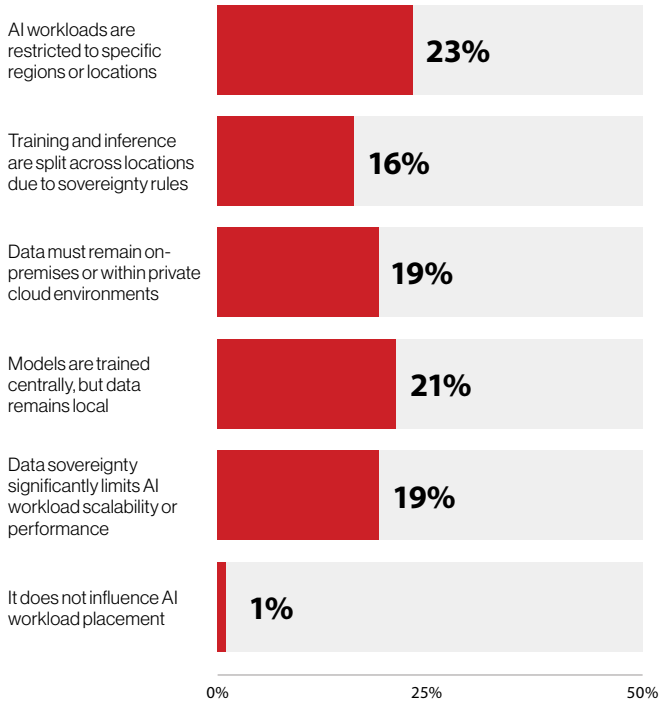

Conclusion

The findings show that the challenge for financial services organisations is no longer adopting object storage, but using it effectively within broader data strategies.

Although data's value for analytics, AI, and differentiation is recognised, uneven execution and incomplete governance, architecture, and operating models limit impact

Misalignment risks wasted investment: passive approaches underutilise data, while premature active use without governance, integration, and performance controls can create operational complexity.

This misalignment risks wasted investment: passive approaches underutilise data, while premature active use without governance, integration and performance controls can create operational complexity.

Success requires addressing both structural and technological challenges – establishing a cohesive data foundation that supports regulatory requirements and innovation at scale. This requires investment in metadata, automation, and integration, alongside alignment across teams and processes.

Object storage is a strategic control point. Organisations that balance control, accessibility, and capability will convert data into competitive advantage; others risk being constrained by their own infrastructure.

Object storage is a strategic control point. Organisations that balance control, accessibility, and capability, will be better positioned to convert data into competitive advantage: others risk being constrained by their own infrastructure.

Hitachi's perspective

At Hitachi, we see object storage evolving through a clear progression. For many years, the foundation was defined by the essentials: scale, durability, and S3 compatibility. Those capabilities remain critical, but they are no longer enough on their own. Today, the market is demanding something more: object storage that is not only resilient and scalable, but also intelligent, composable, and ready to support modern data strategies.

That is the direction we are taking with our platform. We are building a more intelligent object storage layer with native data services such as PII-aware capabilities, support for S3 Tables with Iceberg, and an architecture designed for composable scale. This reflects a broader shift in the industry. Data gravity is real. It is increasingly more efficient to move intelligence to the data than to move data to the intelligence. At the same time, sovereignty has become a baseline requirement. For highly regulated organizations, the ability to govern, audit, and control data within defined boundaries is no longer optional.

Looking ahead, we believe the next phase of object storage will be defined by richer metadata engines that actively enrich and classify data, accelerate RAG workloads, support high-performance interfaces such as S3 over RDMA, and introduce vector-aware capabilities for AI pipelines.

To us, the direction is clear: object storage is moving from a passive bucket to the intelligent platform that AI and lakehouse ecosystems depend on. That is the future we are building for.

Learn more →

FStech In collaboration with HITACHI