White Paper

The Private Cloud Imperative: Operational Resilience and Regulatory Accountability in Financial Services

Sponsored by: Broadcom

Sam Abadir
April 2026

IDC OPINION

IDC believes that financial services firms are reassessing cloud strategy, not because cloud adoption has slowed, but because accountability has become unavoidable. Regulators, boards, and customers now expect institutions to retain control over operational outcomes, including uptime, recovery timelines, and data integrity, regardless of where technology is hosted. This accountability increasingly sits with compliance, risk, and technology leaders, who must explain, defend, and evidence architectural decisions during examinations, post-incident reviews, and ongoing supervisory engagement. The rapid expansion of generative AI (GenAI) into decisioning, analytics, and customer-facing processes accelerates this shift by compressing the time between system behavior and real-world impact, leaving less tolerance for opaque failures or delayed recovery.

How enterprises currently articulate their priorities reflects this shift. Capitalizing on GenAI ranks as the leading business priority, while operational resilience, risk management, cybersecurity, and regulatory compliance remain prominent alongside digital transformation initiatives. This combination reflects a shift toward pursuing innovation while strengthening accountability and resilience (see Figure 1). Rather than displacing resilience concerns, GenAI investment heightens the need to ensure that underlying architectures can sustain faster decision cycles, increased dependency, and stricter expectations for control and recoverability.

April 2026, IDC #US54456626

FIGURE 1

Top business priorities for financial services organizations

Q. In the next two years, what will your organization's top business priorities be?

n = 917

Source: IDC's *Industry Insights Survey: Financial Services,* 2025

Private cloud (as a cloud operating model, not as a synonym for on-premises infrastructure; refer to the definition in the Situation Overview section) is reemerging, not as a rejection of cloud computing, but as a response to regulatory realism that the expansion of GenAI into core operating processes is shaping. As GenAI increases, system interdependence, decision velocity, expectations around uptime, incident disclosure, and resilience testing become more stringent. In this context, architectural decisions become risk decisions, and direct authority over infrastructure behavior, security posture, and recovery actions takes precedence over abstract assumptions about shared responsibility.

This shift reflects a broader realization across financial services. Institutions can consume innovation from many sources, including public cloud ecosystems, but they cannot outsource responsibility for outcomes. As a result, cloud strategy is increasingly defined by how control is distributed, proven, and exercised across environments, rather than by a single hosting preference.

SITUATION OVERVIEW

Cloud adoption across banking, insurance, capital markets, and payments continues to expand, and cloud operating models are increasingly well understood across financial services. Institutions are broadly comfortable with automation, software-defined infrastructure, and cloud-native design patterns across development, analytics, and customer-facing workloads. Cloud acceptance itself is less frequently a primary constraint in architectural decision-making.

What has changed is the regulatory and operational context in which cloud is consumed. Institutions are increasingly required to demonstrate control over uptime, recovery objectives, data location, execution environments, and third-party dependencies as part of routine supervisory engagement. As AI-driven systems accelerate execution speed and decision cycles, regulators are placing greater emphasis on continuous evidence that controls operate consistently across environments, not just on documented intent or contractual assurances.

For financial services institutions, regulatory compliance is increasingly not a question of awareness or the interpretation of standards. Risk, compliance, and technology leaders generally understand the requirements across resilience, governance, data protection, and third-party oversight. The challenge increasingly lies in how enterprises operationalize those expectations at scale as regulatory regimes evolve and converge across regions. Supervisors are placing greater emphasis on continuous evidence, integrated controls, and demonstrable operational resilience, shifting compliance from policy alignment toward system-level execution.

Operational resilience regulations formalize this expectation. Whether through DORA, U.K. operational resilience frameworks, U.S. cybersecurity disclosure rules, or sector-specific guidance across insurance and capital markets, institutions must identify critical services and demonstrate that they can remain available or recover within defined tolerances. This includes cyberevents, infrastructure failures, and third-party disruptions.

Investment behavior reflects this shift. Regulatory requirements, international regulatory exposure, and the risk of regulatory fines are the primary drivers of operational resilience investment, often outweighing internal modernization or innovation initiatives. Cyberthreats and emerging risk trends remain material, but regulatory obligation is the dominant catalyst

for action (see Figure 2). As a result, architectural decisions increasingly move from discretionary technology optimization into the domain of risk ownership and regulatory accountability.

FIGURE 2

Primary drivers of investment in operational resilience

Q. What drives your investment in operational resilience?

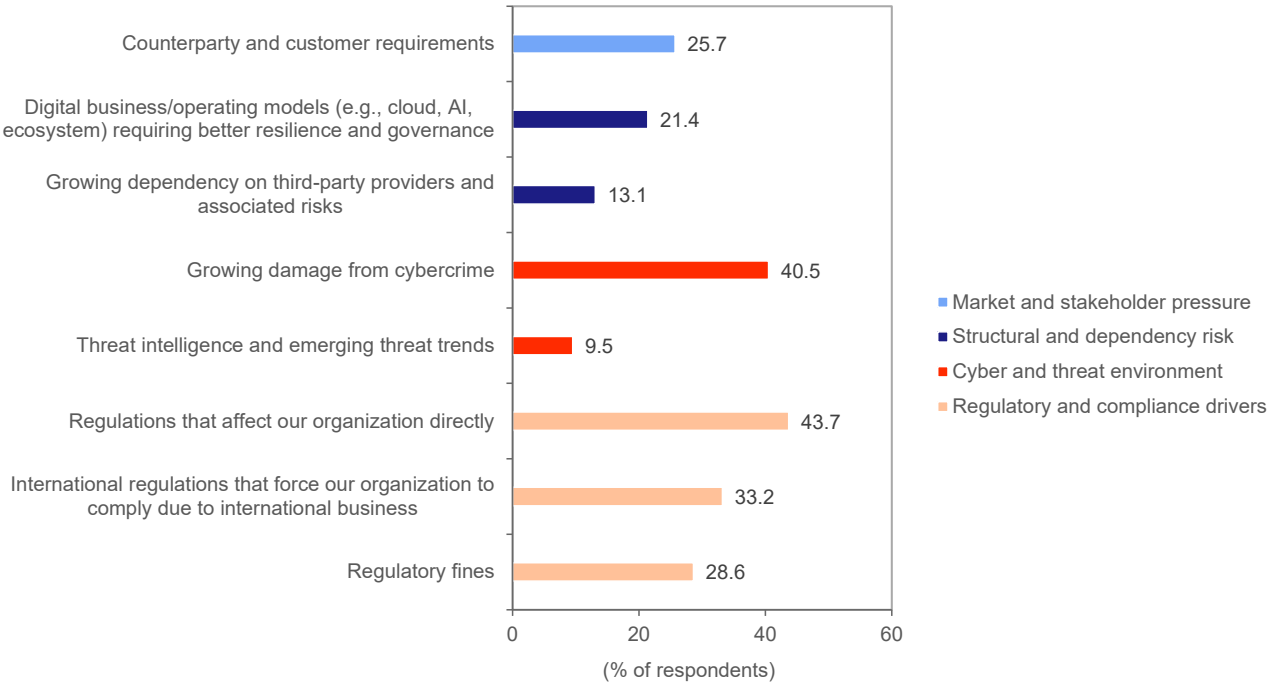

n = 650

Source: IDC's *Financial Insights Survey,* 2024

Public cloud operating models have become a central driver of modernization and remain a critical component of most banking, financial services, and insurance (BFSI) architectures. However, under heightened regulatory scrutiny, accountability for operational outcomes remains unchanged regardless of delivery model. This dynamic has been evident across multiple regulatory actions where institutions experienced service outages or control failures within internally operated or vendor-delivered environments, yet they were held fully accountable for remediation, customer impact, and supervisory response. In these cases, regulators focused on governance, recovery execution, and evidentiary control rather than the specific technology provider involved. When incidents occur, accountability remains with the institution, even when remediation actions depend on third-party systems, vendors, or

external infrastructure. As a result, institutions are placing greater emphasis on governing outages rather than explaining them after the fact. Predictability, transparency, and recovery authority are becoming more important than marginal improvements in scale or resource flexibility. Security concerns have not diminished, but they have become more focused. Institutions are evaluating concentration risk, geopolitical exposure, visibility gaps, and their ability to execute recovery actions independently when needed.

This recalibration exposes a structural gap. While dependency risk is widely recognized, a recent IDC survey shows that almost 75% of institutions report that they do not yet have a holistic and executable cloud exit strategy in place. This gap is material in the context of regulatory expectations that emphasize the recoverability and substitutability of critical services under stress (see Figure 3). Historically, businesses addressed exit planning through contracts and contingency documentation. Emerging resilience expectations now require architectural and operational readiness that many legacy models were not designed to provide.

FIGURE 3

Readiness of financial institutions to execute a cloud exit strategy

Q. Do you have a holistic and executable cloud exit strategy in place today?

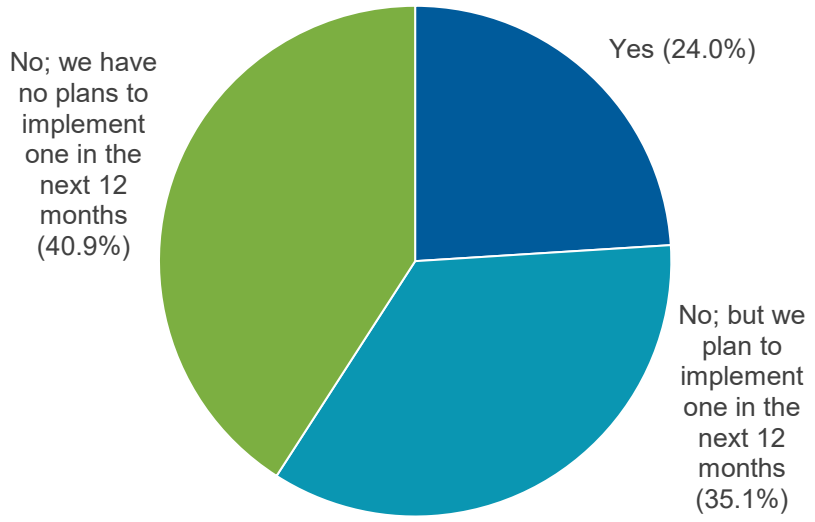

n = 650

Source: IDC's *Financial Insights Survey,* 2024

This has driven a recalibration of cloud strategy. Rather than choosing between public and private cloud, many BFSI organizations are deliberately designing dual and hybrid architectures. In these models, private cloud environments often anchor regulated and systemically important workloads, while public cloud environments are used to access

scalable services, geographic distribution, and specialized capabilities. However, this distribution is increasingly shaped by cost considerations, data locality constraints, and control requirements. Data movement, including ingress and egress patterns, can introduce material cost and latency considerations, particularly for data-intensive workloads such as AI. As a result, institutions are placing greater emphasis on workload placement strategies that account for data gravity, access control boundaries, and the ability to evidence governance across both environments.

In this paper, dual architecture refers to the use of two independently governable execution environments designed to support resilience and recovery. These environments may span private and public clouds, multiple private environments, or other segmented models, provided they support independent control, recovery execution, and regulatory accountability. Throughout this paper, private cloud refers to a cloud operating model defined by dedicated governance and institutional control over execution environments, not to a specific infrastructure location. However, the resilience benefits discussed here assume infrastructure-level independence from public cloud providers, meaning the private cloud environment does not share underlying infrastructure with the secondary cloud environment it is designed to back up or complement.

FUTURE OUTLOOK

How effectively institutions internalize operational accountability across complex, multi-environment architectures will shape the future of regulatory compliance and cyber-risk management in financial services. Cloud adoption will continue, but tolerance for unbounded dependency in critical services will decline as regulators place greater emphasis on uptime, recovery authority, and transparency.

For large global banks, supervisory focus will increasingly center on systemic resilience and cross-environment recovery. Regulators will assess whether institutions can sustain core banking, payment, and liquidity functions during widespread disruption, including scenarios where a single cloud environment becomes unavailable. This will push banks to adopt architectures that distribute workloads and recovery capabilities across private and public environments while maintaining clear command and control over execution.

For regional and midtier banks, the challenge will be to achieve regulatory parity without proportional resources. These institutions face the same resilience expectations as larger peers but with tighter operational margins. Architectures that simplify control, reduce opaque dependencies, and support clear recovery execution across environments will be more favorable than complex designs that are difficult to evidence during examinations.

For insurance carriers, service continuity during catastrophic events and cyberincidents will define future risk exposure. Claims processing, policy administration, and customer

communications must function during periods of extreme stress. The use of dual architectures will increase to balance regulatory expectations with surge demand.

For capital markets firms, tolerance for latency disruption and operational opacity will continue to shrink. Trading, clearing, and settlement systems will face heightened scrutiny around availability and incident disclosure. Firms will need architectures that allow them to directly observe, control, and recover execution environments, including coordinated recovery across dual infrastructure during market volatility.

For payments providers and financial market infrastructure operators, availability remains existential. Even brief outages trigger regulatory intervention and market impact. As these firms are systemically important, expectations around independence, recovery authority, and cross-environment resilience will continue to rise.

Across all BFSI segments, this evolution does not signal a retreat from public cloud adoption. It reflects a more deliberate alignment between regulatory exposure, workload criticality, and architectural control. Institutions that invest in governed, dual architecture resilience will be in a better position to meet supervisory expectations than those that rely on single-environment assumptions.

In this environment, technology leaders are increasingly evaluated on recoverability and control rather than performance alone, while risk and compliance leaders are judged on their ability to evidence resilience outcomes across complex architectures.

Broadcom and the role of private cloud

Broadcom's private cloud strategy is aligned to regulated operating models that require explicit, testable control over execution, recovery, and audit evidence. This focus aligns with an environment in which regulatory pressure is the primary driver of investment in resilience and where the ability to execute orderly cloud exits remains inconsistent across the industry.

In dual-architecture environments, private cloud often serves as the control anchor for regulated workloads. It provides deterministic infrastructure behavior, clear recovery authority, and consistent governance across environments. Broadcom's approach emphasizes standardization across infrastructure layers, helping institutions reduce variability that complicates compliance and operational oversight.

Rather than positioning private cloud as a replacement for public cloud, Broadcom supports segmented and complementary architectures. Institutions can leverage public cloud for elasticity, rapid provisioning, and geographic diversification while anchoring critical services in private environments designed for regulatory control and predictable recovery. Innovation is executed across both environments, with private cloud supporting governed development and production alignment.

Scenarios

Financial services represent a key use case for Broadcom's private cloud strategy:

- **For large global banks,** Broadcom's private cloud supports standardized platforms across regions while enabling coordinated recovery with public cloud environments. This allows banks to manage concentration risk, enforce data sovereignty, and execute cross-environment resilience strategies under regulatory scrutiny.
- **For regional and midtier banks,** Broadcom's private cloud provides a controlled foundation for regulated workloads while allowing selective use of public cloud for development, analytics, and resilience. This balance helps institutions modernize without introducing unmanaged dependency risk.

However, there are strong use cases as well for the following industries:

- **For insurance carriers,** Broadcom's private cloud supports stable execution environments for policy, underwriting, and claims systems, while public cloud resources can provide surge capacity during catastrophic events. This dual approach supports regulatory expectations and operational flexibility.
- **For capital markets firms,** Broadcom's private cloud enables tight control over performance-sensitive systems, with the ability to integrate public cloud for testing, analytics, or contingency scenarios. Direct authority over recovery processes remains essential for meeting disclosure and resilience requirements.
- **For payments providers and financial market infrastructure operators,** Broadcom's private cloud supports high-availability transaction processing, while complementary architectures can provide geographic and operational diversification. This supports increasingly stringent expectations around uptime and independence.

CHALLENGES/OPPORTUNITIES

Broadcom must continue to address several challenges to fully realize this opportunity. In particular, private cloud platforms must meet financial services expectations for automation, scalability, and developer experience. Organizations will not accept controlled environments that materially slow change or increase operational burden, regardless of their governance benefits.

Private cloud platforms must also support high levels of automation and integration without introducing control fragmentation or evidence gaps. Environments that inhibit change or obscure operational behavior increase regulatory risk, even when governance objectives are well intentioned.

Integration remains critical because regulatory evidence is produced across interconnected systems. Financial institutions rely on security, observability, governance, and recovery

tooling that spans environments. Private cloud platforms must integrate in ways that preserve traceability, control consistency, and evidentiary continuity rather than creating isolated control domains.

Broadcom has navigated similar transitions in the past by standardizing complex infrastructure layers, increasing automation, and aligning platform capabilities to enterprise control requirements. If it continues to apply this approach to regulated, dual-architecture environments, these challenges represent a credible path to differentiation rather than structural barriers.

CONCLUSION

Financial services institutions are entering a phase where operational resilience and cyber-risk management are no longer abstract design goals. They are enforceable expectations. Regulators, boards, and customers increasingly judge institutions by their ability to sustain critical services, recover rapidly from disruption, and explain outcomes with clarity and evidence.

Cloud adoption will continue, but it will not be simplistic. Institutions are moving toward deliberate, dual-architecture strategies that align workload criticality with regulatory exposure and recovery requirements. Private cloud plays a central role in this model, not by excluding public cloud, but by anchoring control where responsibility cannot be shared.

Institutions that succeed will be those that align technology execution with governance obligations. They will treat uptime, recoverability, and data control as strategic capabilities exercised across environments. This requires disciplined design and investment, but the alternative is prolonged regulatory friction and weakened supervisory confidence.

For vendors serving financial services institutions, the implication is straightforward. Platforms must support governance, evidentiary requirements, and recovery execution across increasingly complex architectures. In this context, private cloud, when integrated into a governed dual architecture, functions as a structural enabler of sustained compliance and operational resilience rather than a short-term infrastructure decision.

In an environment defined by accountability and disruption, controlled cloud operating models are no longer optional. They are becoming prerequisites for sustainable participation in regulated financial markets.

MESSAGE FROM THE SPONSOR

VMware Cloud Foundation is a unified private cloud platform that combines the scale and agility of public cloud with the security and performance of private cloud, delivering increased productivity and lower TCO. The platform modernizes infrastructure with integrated, enterprise-class compute, networking, storage, management, and security across all endpoints. Automated infrastructure and intelligent operations optimize performance, lower costs, and reduce operational overhead.

Organizations can accelerate innovation with a self-service IaaS platform that delivers a modern cloud interface to run VMs, containers, and AI workloads. Built-in security and resiliency safeguard businesses, ensure business continuity, and free up teams to focus on innovation rather than responding to security threats.

How VCF addresses the needs of financial services:

- Cyber Resilience - VCF delivers robust ransomware and data recovery processes to recover securely and with minimal downtime.
- Compliance - VCF ensures data sovereignty and meets DORA's audit and reporting mandates.
- Private AI-ready - VCF provides a standardized infrastructure stack (vSphere, vSAN, NSX, VMware vSphere Kubernetes Service, etc.,) to deploy localized AI workloads securely on-premises.
- Cost Optimization - VCF automates lifecycle management and rightsizing, leading to a 20-25% decrease in IT labor costs.
- Agility - VCF enables license portability to move workloads between private and public clouds (AWS, Azure, Google) as costs change.

Learn more at https://www.vmware.com/products/cloud-infrastructure/vmware-cloud-foundation.

ABOUT IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com