

Operational Resilience in Banking: From Disaster Recovery to Minimum Viable Bank

WRITTEN BY:

Sam Abadir

Research Director; Risk, Compliance, and Financial Crime


PUBLISHED:

August 2026

The rise in disruption risk is forcing financial services organizations to confront gaps in cyber-recovery, with solution shortfalls, regulatory pressure, and direct loss experience driving 2026 investment.

Figure 1
Why financial services organizations are increasing cyber-recovery investment in 2026: Insufficiency and compliance top the list, but one in four has already paid the price

Q What is driving increased cyber-recovery investment?



Note: n = 387
Source: IDC's *Future Enterprise Resiliency and Spending (FERS) Survey*, December 2025, Wave 9

Introduction

Every chief information security officer (CISO) in financial services is managing a version of the same uncomfortable truth: the organization will face an event that puts operations and data at risk. Figure 1 shows that one in four financial services organizations suffered losses from a successful cyberattack in the past 12 months alone. More telling, nearly 70% that are planning to increase cyber-recovery spending in 2026 cited their current solution's inability to prevent or mitigate losses as a top 2 driver for that investment. The problem is not awareness; it is architecture.

Underlying every resilience investment decision is a question institutions rarely make explicit: what does the organization need to keep running? The concept of a minimum viable bank (MVB) provides the analytical frame. An MVB is the minimum set of systems, data flows, and operational processes required to sustain critical banking functions during and immediately after a disruptive event, including payments, settlements, core deposit access, and regulatory reporting. It is not a degraded state to be avoided but a deliberate, pre-engineered operating posture that the institution can reach and hold while full recovery proceeds. Without an MVB definition, recovery planning lacks a target. Recovery point objective (RPO) and recovery time objective (RTO) commitments float against an undefined endpoint, and incident response teams make scope decisions under pressure that should have been made in advance. For CISOs, defining the MVB is the prerequisite to everything else in this brief. It determines which systems require the highest-integrity recovery architecture, which workloads justify flash-based restore speeds, and what operational actually means when the work recovery time (WRT) clock stops.

Ransomware dominates the threat narrative, but it shares a single critical characteristic with several other threats that can render an institution operationally blind at exactly the moment recovery is most urgent:

- Destructive malware
- Insider sabotage
- Supply chain compromise
- Catastrophic infrastructure failure

At a glance

Key stats

- One in four financial services organizations suffered losses from a successful cyberattack in the past 12 months.
- Nearly 70% of the surveyed organizations said their current recovery solution is insufficient, making it the top driver for increased investment in 2026.
- Around 30% of financial services organizations named compliance and cyber-insurance requirements as their primary driver for increased cyber-recovery investment.

Key takeaways

Legacy backup was not built for modern disruption. Chief information security officers (CISOs) that define recovery time objective (RTO) and recovery point objective (RPO) targets, verify backup isolation, and test recovery realistically will be better positioned operationally and with regulators when an event occurs.

- Extended cloud provider outages

The common thread across these scenarios is not the attack vector. The question is whether a verified, clean, and operationally viable data state exists at the moment recovery is needed. Modern ransomware operators answer that question for their victims by targeting backup infrastructure during extended dwell periods, sometimes spanning weeks, before detonating. Campaigns, including NotPetya and Shamoon, demonstrated that destructive attacks targeting financial sector infrastructure can be deployed as instruments of geopolitical aggression, with recovery dependent entirely on what survives outside the blast radius. A misconfigured mass deletion by a privileged insider or a cascading failure in a cloud provider environment can produce identical consequences without a threat actor involved at all. In every case, the architecture either holds or does not.

For the CISO, this reframes the problem at a fundamental level. RPOs and RTOs, metrics that have historically defined disaster recovery planning, were designed around hardware failure and natural disaster scenarios. They assume a clean restore point exists. Against any event that corrupts, encrypts, or destroys the backup environment alongside production, this assumption breaks down. Achieving a viable RPO requires not only backups but also a guarantee of its integrity. Meeting an aggressive RTO requires that clean data be accessible at recovery speeds commensurate with the operational urgency, not constrained by the throughput of legacy tape or spinning-disk infrastructure.

Cyber-resilience platforms built around immutable, flash-based storage and isolated recovery architecture address this gap directly. Financial institutions that engineer for clean, rapid recovery as a distinct capability, separate from perimeter defense, will measurably reduce the operational, financial, and regulatory severity of any event that threatens data integrity. Those that do not will continue to negotiate from a position of weakness when it matters most.

Trends

Regulators are broadening the definition of resilience

The Digital Operational Resilience Act (DORA), in force and enforceable since January 2025, codifies ICT operational resilience as a standalone regulatory obligation, requiring in-scope European Union (EU) financial entities to manage and withstand ICT-related disruption well beyond malicious cyberattacks, including system failures, outages, and third-party provider failures. Its scope is bounded to ICT risk rather than all-hazards operational risk. The genuinely all-hazards regimes sit elsewhere:

- United Kingdom under the Prudential Regulation Authority's (PRA) and Financial Conduct Authority's (FCA) operational resilience rules
- Singapore under the Monetary Authority of Singapore's (MAS) business continuity expectations

- Australia under the Australian Prudential Regulation Authority's (APRA) Cross-Industry Prudential Standards (CPS) 230.

These require firms to maintain critical operations through any severe but plausible disruption, whatever its source. The direction of travel is consistent even as the hazard scope diverges. The question regulators are asking is no longer only whether an institution can survive a cyberattack but whether it can keep critical operations running under disruption, with ICT treated as one of the most material sources of that disruption rather than a separate regime.

Attackers have adapted to traditional backup architectures

Current threat intelligence indicates that sophisticated ransomware operators and state-sponsored threat actors conduct deliberate reconnaissance of backup infrastructure during dwell periods, targeting backup agents, replication jobs, and cloud-connected storage before triggering their payload. Institutions still relying on backup environments connected to production networks or on cloud backup solutions accessible via the same identity infrastructure as production are operating against adversaries that have already mapped those environments.

Recovery is no longer a disaster recovery exercise

Framing recovery as a disaster recovery (DR) function has become a structural liability. DR programs were designed for infrastructure failure: a failed array, a corrupted database, a datacenter outage. They optimized RTO and RPO against accidental loss and assumed a clean restore point existed. Ransomware and wiper attacks do not share this assumption. They are deliberate operations, and the recovery problem they create is categorically different.

In a DR scenario, the organization controls the timeline. It declares a disaster, activates a runbook, and restores from backup. In a ransomware or wiper scenario, the institution discovers the event after detonation, with material uncertainty about when the environment was first compromised. The question is not how quickly recovery can proceed but what can safely be restored.

That distinction dismantles the DR playbook. Identifying a clean restore point requires knowing the infection boundary: when the attacker entered, what it touched, and whether backup copies were corrupted during the dwell period. Most institutions cannot answer those questions within the time frames required by regulators and business continuity requirements. Recovery is now a security function with infrastructure execution. Institutions that have not reorganized their recovery programs around that reality are testing plans that do not match the threat.

Benefits

Integrity-guaranteed recovery across threat scenarios

Immutable storage architectures prevent the modification, overwrite, or deletion of stored data by any actor, including privileged administrators, after the write window closes. For CISOs, this means the recovery environment cannot be compromised regardless of how an event originates, whether through external attack, insider action, or infrastructure failure. Air-gapped or logically isolated recovery vaults, combined with immutability controls, eliminate the scenario in which an attacker with domain administrator credentials or an administrator making a catastrophic error reaches and corrupts the last line of defense. The clean state is architecturally guaranteed, not operationally assumed.

Aggressive RPO and RTO performance through flash

Flash-based recovery infrastructure fundamentally changes the performance parameters of operational resilience. Tape restoration is measured in hours (or days) and subject to sequential read constraints. Cloud-based recovery introduces bandwidth dependencies that collapse during incidents. Flash delivers recovery speeds measured in minutes. For tier-1 systems, core banking, payment processing, and settlement infrastructure, the difference between a 2-hour RTO and 14-hour RTO is not a technical footnote; it is the difference between a contained incident and systemic operational failure with regulatory and customer-facing consequences.

WRT

RTO measures how long restoration takes. It does not measure how long recovery takes. WRT is the interval between system availability and actual business operability: the time required to validate data integrity, reconcile transactions, restore application dependencies, clear exception queues, and confirm that staff and processes can resume normal operations on the recovered environment.

For financial services institutions, WRT is frequently the longer interval. A core banking system restored in two hours may require an additional four to six hours of reconciliation before payment processing can resume without risk of error or regulatory exposure. Settlement systems require sequenced validation across counterparty records before they are operationally trustworthy. The RTO clock stops when the system comes back. The WRT clock stops when the business comes back.

This distinction has direct implications for how resilience commitments are structured and tested. An institution that meets its RTO against regulators but has not accounted for WRT in its business continuity planning is presenting an incomplete picture of its actual recovery posture. Flash-based recovery infrastructure compresses RTO materially, but the institutions that also engineer for WRT, through pre-validated recovery runbooks, automated reconciliation tooling,

and tested operational hand-off procedures, are the ones that translate technical recovery into business continuity within the windows that matter.

Flash also supports granular recovery, enabling institutions to restore specific workloads or data sets rather than performing a full environment restore, thereby compressing recovery windows and reducing the operational complexity of incident response. Because flash operates without the read/write degradation that affects spinning disks over time, recovery performance is consistent regardless of how long data has been stored. RTO commitments become defensible to regulators and leadership rather than theoretical targets that erode under pressure.

Regulatory defensibility

The compliance driver is not abstract for financial services organizations. As shown in Figure 1, about 30% of financial services organizations cited compliance and cyber-insurance requirements as their primary driver in increasing cyber-recovery investment in 2026, with more than half naming it as a top 2 driver. The DORA, Federal Financial Institutions Examination Council (FFIEC) guidance, Office of the Comptroller of the Currency (OCC) expectations, and Securities and Exchange Commission (SEC) incident disclosure rules converge on one requirement: institutions must demonstrate, not only document, their recovery capability. Critically, these frameworks do not scope their requirements to malicious cyberevents alone. DORA's mandate reaches the full range of ICT-related disruption, including technology failures, third-party provider dependencies, and human error in critical systems, not attacks alone. A tested, immutable recovery architecture with defined, achievable RPO and RTO metrics gives CISOs the evidence base to satisfy supervisory expectations across that scope. Post-incident examinations increasingly scrutinize whether recovery procedures were tested realistically and whether claimed recovery timelines were met in practice.

Enterprise impact and the business case foundation

The financial exposure from an uncontrolled resilience event extends well beyond immediate remediation costs. As Figure 1 illustrates, one in four financial services organizations is already spending from a position of loss, after suffering a successful attack in the prior year. This number reflects only reported and acknowledged losses, not near misses or events attributed to other causes.

Downtime in financial services can have a direct revenue impact of hundreds of thousands to billions of dollars per day for midsize to global institutions, depending on the affected systems:

- Payment processing outages, trading halts, and customer-facing application failures each carry measurable revenue and relationship costs.
- Regulatory fines and supervisory actions following inadequate incident response have reached the nine-figure mark in documented cases.
- Litigation exposure from customers, counterparties, and shareholders adds further tail risk. For wiper attacks and destructive malware events, in which no ransom negotiation option

exists, the recovery timeline is the only variable the institution controls, making the speed and reliability of that capability the single largest determinant of total financial impact.

- Reputational damage affects deposit retention, institutional counterparty relationships, and talent acquisition long after technical recovery is complete. For CISOs building internal support for resilience investment, these enterprise-level consequences are the business case.

Considering Everpure

Everpure delivers flash-based, immutable storage in an air-gapped environment built for enterprise workloads, now positioned within the company's broader data management platform following its 2026 rebrand from Pure Storage. Immutability closes the write window against privileged-credential compromise and insider error, whereas the flash substrate supports the minutes-scale restore speeds tier 1 systems require. The clean state becomes architecturally guaranteed rather than operationally assumed, which is the standard regulators and recovery testing increasingly demand.

Immutability is not a rigid switch with no governed exception path. Early deletion of a locked snapshot, before its retention timer expires, requires two named, authorized contacts working with Everpure support at the same time, outside the institution's own login system. No single administrator, regardless of privilege level, can unilaterally delete protected data ahead of schedule. Every authorization request is logged, creating a full audit trail of who approved what and when.

For CISOs, this matters for two reasons beyond ransomware defense. First, it gives the institution a defensible, documented path to honor a legitimate early-deletion obligation, such as a GDPR or CCPA erasure request, without weakening the immutability guarantee for every other record on the system. Second, it closes the insider threat gap that a single compromised or malicious administrator credential would otherwise represent.

Where CISOs should apply scrutiny is the integration layer because the stakes there sit on the recovery path itself. Everpure controls both the OS release cycle and array plug-in so keeping them in sync is the vendor's job, not the customer's. When they fall out of sync during a firmware upgrade, the system can throw misreported multipath warnings. The danger is not the false alarms but a real storage path failure hides among them, gets tuned out by an operations team used to ignoring the noise, and surfaces only at the moment of recovery. For a capability whose whole purpose is to work when nothing else does, this is the failure mode that matters.

Everpure has moved to narrow this window by publishing compatibility matrices and upgrade guidance ahead of firmware releases, but the responsibility distribution still matters for procurement. CISOs should confirm before deployment that firmware release schedules are communicated proactively and that plug-in compatibility validation is part of standard customer success engagement, not a gap the customer discovers under pressure. For a platform whose value proposition is guaranteed recovery, the integration layer is not a secondary consideration. It is the place where that guarantee either holds or it does not.

Conclusion

For financial services CISOs, the operational risk posed by an underbuilt recovery capability is no longer confined to ransomware threats. Destructive attacks, insider events, infrastructure failures, and cloud provider outages all lead to the same outcome when recovery architecture is inadequate: extended downtime, regulatory exposure, and a loss of control over the institution's narrative.

Three questions are worth answering honestly before any further investment in prevention or detection:

- Are existing backup environments isolated from production in a way that survives both adversarial compromise of privileged credentials and catastrophic infrastructure failure?
- Have RPO, RTO, and WRT commitments for tier 1 systems been tested under realistic conditions, not tabletop exercises but actual recovery from backup under simulated incident pressure, to the point in which business operations have verifiably resumed?
- Can the organization identify, right now, a verified clean restore point for its most critical systems?

For most institutions, the honest answer to at least one of those questions is no. Flash-based, immutable recovery architecture closes this gap in a way that legacy disaster recovery frameworks were never designed to. The institutions that treat resilience as a first-class architectural investment will be the ones that survive disruption events intact and emerge with their regulatory standing and customer trust preserved.

Financial services institutions are not underinvesting in recovery because they do not understand the risk but because their legacy architecture has obscured how wide the gap actually is.

About the IDC analyst(s)



Sam Abadir

Research Director, Risk, Financial Crime, and Compliance

Sam Abadir is the research director for IDC Financial Insights responsible for risk, financial crime, and compliance practice. His core research coverage includes compliance regulators around "bad actors," including anti-money laundering, enhanced customer due diligence (eCDD), and know your customer (KYC), and other regulatory programs, such as the General Data Protection Regulation (GDPR) and compliance watchlists (e.g., Office of Foreign Assets Control [OFAC] and politically exposed person [PEP]). His research covers different aspects of risk and compliance in financial services and how technology can support mandated and best practices in the Digital Operational Resilience Act (DORA), GDPR, and other regulations that impact financial services.

Message from the sponsor

Banks are entering a new era in which AI, cyber-resilience, and operational resilience have become interconnected priorities. Success depends on more than protecting data. It requires a modern data platform that keeps critical information secure, available, governed, and ready to support real-time business decisions, even during disruptions.

Everpure helps financial institutions modernize their data infrastructure with a unified platform designed for mission-critical banking workloads. From ransomware recovery and cyber-resilience to AI-ready data and regulatory compliance, Everpure enables banks to simplify operations, reduce risk, and accelerate innovation while maintaining the performance and availability of today's financial services organizations demand.

Learn how Everpure is helping banks build resilient, AI-ready data platforms by visiting its banking solutions page:

<https://www.everpuredata.com/solutions/industries/financial-services/banking.html>

For practical guidance on strengthening cyber and operational resilience, download the solution brief Strengthening Cyber and Operational Resilience in Banking:

<https://www.everpuredata.com/docs.html?item=/type/pdf/subtype/doc/path/content/dam/pdf/en/solution-briefs/sb-strengthening-cyber-operational-resilience-in-banking.pdf>

IDC Custom Solutions

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for [external use](#), and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.



IDC Research, Inc.

One Beacon Street, Suite 33100, Boston, MA 02108, USA

T +1 508 872 8200

blogs.idc.com | LinkedIn [@IDC](#) | www.idc.com

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events 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.

©2026 IDC. Reproduction is forbidden unless authorized. All rights reserved. [CCPA](#)