CLOUDFLARE

WHITEPAPER

Resilience at scale:

A strategic blueprint for financial services

Content

A foundational shift in resilience thinking

In the financial services industry, resilience is not a feature — it is a foundation. It must be treated as a systemic, enterprise-wide capability rather than a technology fix. Disruptions in this area can have outsized ripple effects: customer trust erodes in seconds, regulatory scrutiny intensifies, and systemic risks can propagate across markets. With increasing digital dependence, heightened threat activity, and evolving regulations, such as the Digital Operational Resilience Act (DORA), Paperwork Reduction Act (PRA) guidance, and US Federal mandates, financial institutions must rethink how resilience is designed, governed, and operationalized. It is no longer enough to recover. Institutions must demonstrate continuous availability while maintaining business continuity under stress.

This foundational shift is critical, yet true progress means confronting outdated assumptions. Some common resilience myths continue to hinder strategic action.

Debunking myths about resilience

Myth #1: Resilience equals redundancy

In financial services, resilience is too often equated with redundancy: backup data centers, high-availability clusters, and best-of-breed tools stitched together across environments. While these solutions address isolated risks, they often create fragmented architectures, expensive duplicative infrastructure, and operational silos that struggle under real-world pressure.

But resilience is not just about uptime; it is about preserving trust and systemic stability in the face of volatility. The ability to deliver core financial functions during a crisis — payments clearing, fraud detection, real-time trading — is far more valuable than keeping every system online. Leaders must shift from technology-based redundancy to outcome-based continuity, ensuring the services that matter most stay available when the stakes are highest.

Myth #2: 100% availability is the goal

Achieving 100% availability across all systems, in all regions, at all times, is neither practical nor necessary, particularly in a global bank or financial institution with hundreds of applications and interdependencies. Instead, firms must embrace graceful degradation: the ability to maintain critical operations while allowing non-essential functions to temporarily degrade or pause.

Think of a digital bank during a disruption: perhaps real-time balance checks are unavailable, but customers can still complete fund transfers. Or a trading platform slows portfolio performance refreshes, but core execution remains functional. The goal is not perfection, but functional continuity, aligned to customer and market expectations.

Myth #3: All channels deserve equal resilience

Not all user journeys are created equal — especially in financial services, where different channels serve distinct purposes. A mobile banking application used by millions of retail customers carries a different risk profile than an internal reporting dashboard. Similarly, APIs used for interbank settlement or fraud scoring warrant higher assurance than rarely accessed back-office portals.

A modern resilience strategy requires channel-aware prioritization. Financial institutions must map criticality not just by system, but by customer engagement, transaction volume, and time sensitivity. Overinvesting in low-impact systems and underprotecting key customer channels is a recipe for service failure and reputational damage.

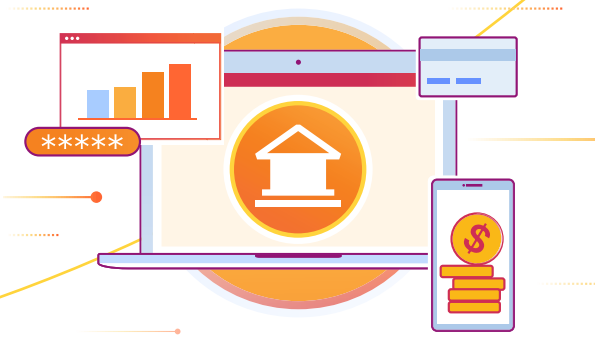

Myth #4: Everything is mission-critical

One of the most common missteps in the financial sector is treating every application or process as "critical." This bloats resilience plans, drains resources, and obscures what truly matters. Instead, institutions must take a page from compliance disciplines like the Payment Card Industry Data Security Standard (PCI DSS) and shrink the scope.

Clarifying and shrinking the scope means identifying which systems directly impact core operations, regulatory reporting, liquidity management, and customer experience. Those should be architected for continuous availability. Non-critical systems can tolerate slower recovery times or alternate access paths. A more focused scope reduces cost, simplifies governance, and improves audit readiness.

Myth #5: Best-of-breed automatically equals best resilience

Financial institutions often lean heavily on best-of-breed tooling to meet specific control requirements: web application firewalls (WAFs) for web protection, separate fraud analytics engines, independent distributed denial-of-service (DDoS) services, and standalone observability platforms. While this may optimize individual layers, it introduces integration risk, inconsistent telemetry, and brittle failover logic.

Moving beyond these persistent myths, real-world resilience is an engineering discipline. It demands a principles-based approach, embedding resilience into the very fabric of modern financial institutions.

Architecture principles for real-world resilience

For financial services organizations, resilience must be engineered, not assumed. Reactive recovery plans or fragmented failover strategies are not sufficient. From targeted cyber attacks to infrastructure-wide outages, disruptions are increasingly complex and consequential. They do not just cause downtime — they erode trust, trigger regulatory violations, and threaten financial stability. Architecture is the first line of defense against these risks. What follows are five foundational principles for building resilience into the architecture of modern financial institutions.

1. Design for architectural independence

Modern disruptions span both malicious attacks — such as DDoS, credential abuse, and web application exploits — and systemic infrastructure failures, including cloud outages, Border Gateway Protocol (BGP) leaks, and DNS breakdowns. Relying on a single cloud provider or tightly coupled technology stack introduces an unacceptable concentration of risk, making organizations vulnerable to service

disruptions, data loss, and reputational damage. To mitigate this, consider adopting a distributed, multi-cloud, or edge-centric architecture that eliminates single points of failure and allows for dynamic rerouting of traffic and workloads across diverse environments.

Architectural independence is now a regulatory and operational requirement. Regulations like DORA, PRA guidance, and US Federal mandates emphasize the need for multi-vendor strategies and the ability to maintain operational continuity regardless of where a failure occurs. Infrastructure must be able to fail over seamlessly and reroute traffic dynamically, enabling organizations to adapt in real time without compromising service delivery.

2. Optimize spend by moving beyond redundancy

Legacy architectures may appear resilient, but they often conceal significant operational costs:

- Data egress fees from multi-cloud environments
- Cold start latency from idle failover systems
- Overprovisioned standby zones that inflate cost without delivering proportional value

These approaches rest on a flawed premise: that redundancy equals resilience. But duplicating infrastructure without intelligence adds complexity and cost without improving outcomes. For financial institutions, the economic and operational burden of legacy recovery models is unsustainable. Resilience should be architected for efficiency — not excess.

3. Simplify through cohesive architecture, not tool sprawl

While a best-of-breed approach offers specialized capabilities, it often results in siloed tools, fragmented visibility, and inconsistent failover behavior. True resilience requires architectural coherence, where components work together seamlessly and policies are enforced consistently acrossenvironments. Instead of layering more disparate tools on top of varied infrastructure, firms should focus on establishing a unified control plane. This control plane, agnostic to the underlying cloud or network providers, enables real-time observability and integrated logic across all layers. This strategic unification simplifies complexity, enabling true agility and consistent defense, particularly when faced with critical incidents or the need to dynamically leverage multiple infrastructure vendors for maximum resilience.

4. Implement layered protection across the stack

Resilience is not achieved at a single layer; it must span the full technology stack. A layered model provides defense-in-depth and enables intelligent failover when disruptions occur:

- At the network layer, global Anycast routing absorbs large-scale attacks and maintains availability.
- At the application layer, edge-based protection and failover reduce dependency on centralized data centers.
- Across all layers, smart routing and real-time telemetry enable automated responses and rapid mitigation.

This approach is not about duplicating entire platforms. It is about adaptive, context-aware continuity, where the system can recover intelligently and prioritize the services that matter most.

5. Decouple policy logic from infrastructure enforcement

One of the most overlooked inhibitors to resilience is operational entanglement. In many institutions, business logic — such as WAF rules, access policies, or fraud controls — is hardcoded into vendor-specific stacks. This makes policy changes slow, brittle, and high risk during a crisis.

Resilience requires logical abstraction. By decoupling rules from tools and managing them through a portable policy layer, organizations can:

- Test and iterate on WAF rules in safe environments
- Maintain version control over bot mitigation logic
- Ensure access policies persist across identity providers

This Rules-as-a-Service approach empowers faster response and consistent enforcement, regardless of where or how systems are deployed.

While these principles form the bedrock of resilience efforts within organizations, the economic, geopolitical, threat, and regulatory environments are evolving rapidly. From the operational disruptions caused by geopolitical instability and sanctions to the rise of AI-enabled attacks that dramatically increase the scale and complexity of threats, financial institutions are facing pressure on multiple fronts.

At the same time, regulatory frameworks like DORA, PRA, and emerging US standards are shifting from static controls to expectations of continuous availability, third-party diversification, and built-in operational independence. In this environment, organizations can no longer rely on fixed, one-time strategies. They must embrace resilience as a dynamic capability — anchored in continuous compliance, proactive threat management, and enterprise-wide agility. Today, the ability to adapt quickly to shifting risks, technologies, and regulatory expectations is a requirement. Now is the time to architect for agility, act with intent, and embed resilience into the core of how your organization operates.

While robust architectural principles form the bedrock, the dynamic threat and regulatory landscape demands continuous adaptation. Savvy leaders are already taking decisive actions to evolve their resilience programs and proactively address emerging challenges.

Strategic actions for evolving resilience programs

1. **Incorporate crisis simulations into strategic planning and budget cycles.** Annual disaster recovery tests are no longer enough. Leading institutions are integrating scenario outcomes into board-level investment decisions and product roadmaps — ensuring that resilience lessons translate into concrete action. After a major ransomware attack, one global financial services organization began using crisis simulations to shape mergers and acquisitions assessments and supplier diversification strategies. Doing so enabled them to mitigate further attacks.
2. **Manage third-party dependencies as systemic exposure, not isolated contracts.** With over 73% of UK financial institutions dependent on just three cloud providers,1 concentration risk has become a regulatory concern. Institutions must run failure simulations for top-tier vendors, map dependency chains across critical services, and build viable exit strategies — just as they would diversify financial risk in asset portfolios.
3. **Assign product-level resilience ownership in high-impact business journeys.** Service continuity is a customer experience imperative. One global retail bank embedded "resilience product owners" into core workflows, like real-time payments and onboarding, empowering those teams to design for graceful degradation, fallback paths, and failover thresholds directly aligned with business priorities.
4. **Develop an AI-aware defense layer to counter emerging threats.** As generative AI fuels highly personalized phishing and polymorphic attacks — 82% of all phishing emails utilized AI[2] — traditional controls are falling behind. Institutions must build AI-native detection layers that can analyze behavioral anomalies in real time, flag prompt injection patterns, and mitigate threats before they escalate.
5. **Address observability debt as a leading indicator of resilience failure.** The longer it takes to see what is broken, the longer it takes to recover. In a 2022 incident, a Tier 1 bank faced a prolonged API outage largely due to a lack of granular telemetry, forcing reliance on manual reports and guesswork. Strengthening observability is not just about performance — it is foundational to timely incident response.
6. **Shift from static recovery targets to adaptive resilience thresholds.** Traditional RTOs often fail during real-world incidents because they are decoupled from business context. According to Cohesity's 2024 Global Cyber Resilience Report,[2] a striking 98% of organizations aim to recover from cyber incidents within one day — but in practice, just 2% actually succeed in meeting that goal. Institutions should develop tiered recovery thresholds that adapt based on service criticality, customer impact, and operational context — enabling more precise, outcome-aligned responses.
7. **Treat agility as a core control, not just a cultural trait.** In a crisis, the ability to pivot quickly is often more valuable than the ability to withstand stress. Following 2022 sanctions, several Eastern European banks rerouted SWIFT transactions and onboarded new providers within weeks — thanks to pre-approved workflows and dynamic governance models. Formalizing agility through emergency overrides, pre-configured fallback routing, and empowered cross-functional teams is now essential for resilience.

1. Mohammed, Gharbawi, Ewa, Ward, Emelie, Bratt, Laurence, Diver, Henrike, Mueller, Rocco Quartu, Haydn, Robinson. "Artificial intelligence in UK financial services – 2024." Bank of England and Financial Conduct Authority. November 21, 2024. https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024
2. "Global cyber resilience report 2024." Cohesity. https://www.cohesity.com/resource-assets/research-reports/cyber-resilience-global-survey-report-en.pdf

1 888 99 FLARE | enterprise@cloudflare.com | Cloudflare.com

REV:BDES-7816.2025AGO12