THE RESILIENCE GAP IN REGIONAL BANKING: WHY AUDIT-READY IS NOT THE SAME AS OPERATIONALLY RESILIENT

This guide explores where the compliance-resilience gap comes from, and how regional banks can evaluate whether their infrastructure decisions support real operational continuity, not just documented recoverability.

Most institutions have the right documentation in place. They have recovery plans, vendor reviews, business continuity procedures, and governance frameworks. Those things matter. But they do not always prove that the operating environment can hold up when a real disruption is unfolding.

That gap usually stays hidden until a trusted dependency fails and the institution has to keep serving customers while systems, vendors, or recovery paths are under stress.

The July 2024 CrowdStrike outage made that distinction hard to ignore. Institutions with mature, audit-ready programs still lost access to critical systems when a trusted third-party dependency failed. The lesson was not simply that cyber vendors can create risk.

It was that documented preparedness does not always translate into operational continuity.

That is the resilience gap.

Many financial institutions are well prepared to demonstrate compliance, but less prepared to prove that critical banking services can keep running when infrastructure, vendors, or recovery paths fail under pressure.

This article looks at where that gap comes from and how regional banks can evaluate whether their infrastructure decisions support real operational continuity, not just documented recoverability.

The Difference Between Recoverability and Resilience

Traditional disaster recovery (DR) planning in banking has been built around a deceptively simple premise: if systems fail, they can be restored.

That model shaped RTOs, failover procedures, backup protocols, and recovery runbooks. It works best when outages are contained, measurable, and recoverable in sequence. But that model no longer reflects the way disruptions unfold in reality.

1 The Shift from Recovery to Continuity

Today’s disruptions rarely stay contained to a single system. They can cascade across cloud platforms, core providers, payment networks, telecommunications carriers, authentication tools, and branch operations. The question is no longer only whether systems can be restored after failure. It’s whether critical banking services can continue operating while the failure is still unfolding.

Regulatory guidance increasingly reflects this distinction. The Federal Financial Institutions Examination Council’s Business Continuity Management guidelines direct examiners to assess whether institutions can maintain critical operations during a disruption, not merely restore them afterward.

Even where specific mandates shift, the operational risk remains: waiting for resilience expectations to become mandatory does not reduce the exposure. Institutions that wait for resilience expectations to become mandatory may end up managing to the exam cycle instead of managing to their actual operational risk.

2 The Credit Union Challenge Runs Even Deeper

Credit unions face an added challenge because many critical service providers sit outside direct supervisory reach. That can limit visibility into the very dependencies that determine whether member-facing services remain available during a disruption.

As a result, examinations may focus heavily on the institution itself while offering less visibility into the third-party dependencies that increasingly shape operational resilience.

3 Tabletop Exercises Do Not Replicate Operational Failure

Tabletop exercises are useful, but they can create a false sense of readiness when they become the primary proof of resilience. They validate escalation paths and decision-making roles. They do not prove that systems, vendors, facilities, and business teams can operate together under real disruption conditions.

“A Jones Walker survey found that 76% of bank executives conduct tabletop exercises at least annually, which helps explain why they remain the primary resilience-testing mechanism across the sector.”

Tabletops are useful for validating escalation paths, incident response roles, and leadership communication. But they are not designed to prove whether:

Recovery runbooks reflect current system configurations

Secondary data sites can be brought online within defined tolerances

Critical workflows can continue when multiple dependencies fail at the same time

Facility access, carrier support, and vendor escalation paths work when the institution needs hands-on support

Real disruptions are messy. Recovery sequencing is uneven, information is incomplete, vendors may be responding to many customers at once, and the people needed to execute recovery may not all be available. A bank can satisfy annual testing expectations while still carrying assumptions that have never been validated under real operating pressure.

4 Recovery Is Not a Technical Challenge Alone

Large migration failures, such as the 2018 TSB Bank incident in the UK, show that recovery failures are rarely technical alone. Technology issues matter, but the larger challenge is often coordinating people, vendors, processes, communications, and decision-making while customers are already affected.

The underlying technology issues were significant, but the larger failure was the bank's inability to manage the disruption as it unfolded.

“The incident ultimately took more than 18 months to fully remediate and resulted in regulatory penalties exceeding £48 million.”

The lesson for regional institutions is clear: recovery outcomes depend on the operating model around the technology, not just the technology itself.

5 The Systemic Risks of Third-Party Dependencies

Third-party vendors, in particular, carry a significant risk. Regional banks and credit unions have long relied on outside providers for core systems, payments, telecommunications, cloud platforms, data center facilities, security tools, and recovery services. A disruption in any one of those dependencies can affect the institution’s ability to serve customers..

At the center of this concentration are a handful of dominant technology providers. For many mid-market institutions, critical infrastructure depends heavily on core banking platforms such as Fiserv, Jack Henry, and FIS, as well as payment processors, connectivity providers, cloud platforms, and infrastructure partners. In many cases, developing proprietary alternatives is not economically realistic, leaving institutions with limited opportunities to diversify their operational risk.

The result is a form of exposure that traditional governance frameworks were never designed to address. Vendor management can confirm that a provider was reviewed before engagement.

Operational resilience asks a harder question:

What happens when a critical provider fails while other parts of the environment are already under stress?

Most mid-market institutions have mature processes for the first scenario. Far fewer have tested their ability to operate through the second.

The Regulatory Response

Regulators have begun to recognize this structural gap. Many regional and community banks depend on a small group of core providers for critical infrastructure, but often have limited leverage to negotiate stronger resilience terms or switch vendors when service quality falls short.

The regulatory direction is clear: resilience now requires a deeper understanding of concentration risk and third-party recovery pathways.

The Perils of Systemic Failure in Practice

The Ongoing Operations ransomware attack in 2023 showed how concentrated dependency risk can affect many institutions at once. The attack reportedly took approximately 60 credit unions offline simultaneously, even though they were not the direct target.

Ironically, Ongoing Operations was itself a disaster recovery provider and a firm many institutions relied on to maintain operational continuity during disruptions. Making matters even worse, the vulnerability exploited by attackers had been publicly known and patchable for months before the attack.

> **"The 2022 Rogers telecommunications outage in Canada showed the same dynamic at the infrastructure layer, disrupting ATM access, payment systems, and mobile banking across a broad ecosystem."**

These incidents highlight a fundamental challenge in modern operational resilience. Critical dependencies often extend beyond an institution's direct vendor relationships and into shared technology and infrastructure ecosystems. These types of operational risks cannot be fully understood through vendor assessments, contract reviews, or due diligence questionnaires alone.

Resilience Is a Governance Decision, Not an IT Expense

Managing these risks requires decisions that extend well beyond infrastructure and operations teams.

One of the biggest challenges facing regional banks is that resilience is often funded by one group and experienced by another. Executive teams see recurring infrastructure costs, while IT teams see the systems that keep branches operating, transactions processing, employees connected, and customers served. When everything works, resilience is largely invisible. When it fails, however, it immediately becomes a business problem.

This disconnect can create difficult tradeoffs. Projects tied to growth, efficiency, or new products produce visible returns, but resilience investments are different. Their value is measured by outages that never happen and customers who never experience an interruption. As a result, resilience discussions often begin as budget conversations even though the consequences of failure extend far beyond IT.

1 Resilience Is Now a Board-Level Responsibility

Today, the most important resilience decisions are risk tolerance decisions.

Infrastructure teams understand the tradeoffs behind provider selection, recovery strategies, and redundancy investments. Boards and executives, meanwhile, often encounter those same decisions as cost, risk, and accountability decisions. The challenge is connecting the cost of a resilience investment to the operational risk it is designed to reduce. Every decision to defer an upgrade, consolidate providers, or reduce redundancy reflects a choice about how much disruption the institution is prepared to absorb.

In this context, effective governance starts with a clear understanding of which services need to be available, how long the bank can operate without them, and what level of risk leadership is willing to accept.

2 The True Cost of Operational Failure

Industry research consistently shows that financial services carry some of the highest downtime exposure of any sector. But for regional banks and credit unions, the consequences of operational disruption extend far beyond lost transactions or recovery costs. An outage can prevent customers from accessing accounts, interrupt branch operations, delay payments, or create a flood of support requests at the very moment confidence is most important.

For relationship-driven regional banks and credit unions, the consequences can outlast the outage itself. Customers may forgive a temporary disruption, but repeated failures can undermine their trust and increase the likelihood of attrition.

This is why operational resilience is ultimately a business issue, not just a technology concern. Every disruption carries financial consequences, but it can also affect customer trust, employee productivity, regulatory relationships, and the institution's reputation in the communities it serves.

3 Stable Operations Can Mask Fragility

Examinations, vendor reviews, and uptime metrics may validate normal operations, but they do not always reveal whether critical services can withstand cascading failure conditions, especially when recovery depends on multiple teams, facilities, carriers, vendors, or cloud platforms.

When Operational Failure Becomes a Financial Event

Operational disruption can quickly become financial stress. When a payment provider or infrastructure dependency fails, the issue is no longer limited to IT. In a recent outage involving a third-party payment provider, multiple banks were unable to process Fedwire transactions, and smaller institutions were forced to access same-day liquidity through the Federal Reserve's discount window. The disruption affected liquidity, transaction processing, customer access, and the institutions' ability to move funds.

A dependency failure had disrupted the infrastructure required to move funds, turning an operational outage into a liquidity event. The same principle applies to third-party risk management.

Contracts and SLAs are important, but they mostly define responsibility after something goes wrong. They do not guarantee that a provider can keep critical services operating during the disruption.

Common sources of untested exposure include:

- Fragmented failover ownership across IT, business teams, vendors, and facilities
- Recovery runbooks that no longer reflect current infrastructure
- Untested assumptions about carrier diversity, facility access, and remote hands support
- Recovery paths that depend on the same cloud region, carrier, or provider ecosystem
- SLAs that define obligations but do not prove execution under widespread disruption
- Key-person dependencies during recovery
- Provider responsiveness gaps that only become visible during urgent support needs

None of these vulnerabilities appear in policies or audit reports, but they are immediately felt when something actually fails.

From Documented to Demonstrated Resilience

Closing the resilience gap does not always require a full-scale modernization program. It does require banks to evaluate whether current infrastructure decisions still align with operational continuity requirements, rather than relying on legacy arrangements that have never been seriously tested.

"Governance, testing, vendor management, and infrastructure decisions need to align to continuity requirements, not just compliance documentation. Structurally, governance, testing, vendor management, and infrastructure decisions need to align to continuity requirements, not only compliance documentation."

Regulatory guidance is moving in the same direction, with increasing attention on third-party concentration risk, lifecycle vendor oversight, and validated recovery pathways.

1 What Resilience Testing Needs to Prove

The FFIEC says that tabletop exercises are not enough on their own. Instead, resilience testing needs to show whether technology, business functions, vendors, and escalation paths work together under realistic disruption scenarios.

Most institutions have not made that shift. Integrated testing means running the full operational chain under disruption conditions rather than just the IT recovery sequence.

In practice, that means validating whether:

- Payment processing continues when authentication systems fail
- Branch staff can keep serving customers when core banking access is interrupted
- ATM and digital banking access remain available during provider disruption
- Technical staff can access the facility or remote support needed to restore service
- Treasury and back-office functions can continue when a key third-party integration drops mid-session

A realistic test forces simultaneous failures rather than isolating one system at a time.

2 Recovery Speed Is Not Service Continuity

The FFIEC framework also asks institutions to define maximum tolerable downtime for each critical business service, separately from their Recovery Time Objectives. While an RTO measures how quickly a system can be restored, maximum tolerable downtime measures how long a critical service can remain unavailable before the disruption becomes unacceptable to customers, counterparties, or regulators.

Many DR programs measure recovery time. Fewer test whether critical services can remain available within acceptable disruption limits, particularly among institutions below the $100 billion asset threshold where such testing is not typically required.

3 Infrastructure Decisions Shape Recovery Outcomes

For regional institutions operating across multiple states, infrastructure choices directly shape resilience. Facility location, power availability, carrier access, physical security, remote hands, workload placement, and recovery architecture all influence whether critical services can remain available during disruption.

Cloud can play an important role in modernization and recovery, but single-region architectures can create constraints when production and recovery paths depend on the same environment. Workload placement needs to be evaluated against continuity requirements, not simply platform preference.

The Case for Distributed Infrastructure

Regional colocation and private hosted cloud can help banks distribute critical workloads across geographically separate, independently connected facilities. For mid-market financial institutions, this can provide proximity, operational accountability, carrier flexibility, physical access, and a clearer recovery path than architectures that concentrate too much dependency in one place.

Evaluating those decisions workload by workload is where many regional banks can make practical progress without launching a full modernization program. The question is not simply where the workload runs. The question is what the bank needs that workload to do during disruption and what infrastructure conditions are required to support it.

Focus Area	Question To Ask	What It's Really Testing
Critical services	Which services must continue operating during a disruption?	Business continuity threshold
Proximity	Which workloads require proximity to branches, ops teams, or technical staff?	Operational responsiveness
Dependency overlap	Where do recovery paths depend on the same provider, carrier, cloud region, or facility?	Hidden concentration risk
Recovery execution	Can technical teams access the facility or support resources when something fails?	Practical recovery readiness
Governance	Are resilience investments being evaluated as IT costs or enterprise risk decisions?	Decision ownership and risk tolerance

"For regional financial institutions, infrastructure decisions are not just architecture decisions. They shape continuity, recovery confidence, and operational control across the markets the bank serves."

Compliance Can Support Resilience, but It Cannot Prove It

Regional banks and credit unions have invested heavily in the governance, recovery, and risk management frameworks regulators expect. While these investments are important, compliance and resilience are not the same thing, and the gap often becomes visible only when a critical service, provider, facility, or recovery path is under stress.

The institutions best positioned for the next decade will be those that continuously test assumptions, understand their operational dependencies, and treat resilience as an enterprise capability, not just a compliance exercise.

The remaining question is whether institutions will find the gap through deliberate testing, infrastructure review, and governance, or whether the next disruption will find it for them.

Published by DartPoints

DartPoints helps regional enterprises align infrastructure with resiliency, continuity, and modernization requirements.

For questions or follow-up:

Visit dartpoints.com or send an email to info@dartpoints.com