citrix

The CIO's business resiliency playbook

Keeping the business running through disruption

White paper

Executive overview

On July 19, 2024, a single misconfigured file in a routine software update took down 8.5 million Windows devices worldwide. Airlines grounded flights. Hospitals turned away patients. Banks went offline. The fix arrived in 80 minutes. The damage lasted for days. Delta Airlines alone reported $500 million in losses.

The organizations that recovered fastest were not the ones with the best IT teams. They were the ones whose employees could keep working while the crisis was still active. Their critical work ran on platforms that were not tied to the affected systems. They did not spend the first day figuring out who was impacted or what the workarounds were.

Recovery gets systems back. Continuity keeps the business moving. Recovery is table stakes. Continuity is the real test.

Business resiliency is no longer a technical concern - it is a strategic priority. Regulatory frameworks like FFIEC, DORA, and UK PRA now require board level accountability for keeping critical business functions running, not just restoring IT systems. Organizations that stay operational during disruption earn customer trust and pull ahead of competitors. Those that go dark trade one crisis for two.

The problem with most resilience strategies

Most continuity planning was built for one scenario: systems go down; IT brings them back. That problem is mostly solved. The harder problem is partial failure - an identity provider slows down, a cloud region degrades, or a SaaS platform goes unresponsive. The infrastructure looks fine on the dashboard - but no one can do their job.

When that happens, employees improvise. Access expands beyond what's governed; controls get bypassed. By the time the incident resolves, the organization is managing two crises: the outage and its aftermath. Audit trails are incomplete; access exceptions weren't closed; a compliance problem quietly accumulated while everyone focused on the operational one.

This is the exception spiral, and most organizations are not built to prevent it. The ones that avoid it are the ones that decided in advance what tightens, what gets restricted, and who can approve exceptions. That is a governance decision, and it must be made before the next incident starts.

If everything is critical, nothing is recoverable. Service tiering is the foundation everything else is built on.

The state of disruption

- **Universal impact:** 100% of large enterprises suffered revenue loss from outages in the past 12 months.
- **Frequency of disruption:** 55% of organizations experience service disruptions at least once per week.
- **Outage costs are material:** High-impact outages average 77 hours of downtime annually, with hourly costs reaching $1.9 million.
- **Breach impact is business impact:** The average data breach costs $4.88 million. Ransomware appears in 44% of breaches.
- **Third-party risk is rising:** Third-party involvement is a leading driver of incident scale. 72% of organizations report rising cyber risk.
- **AI is accelerating the threat curve:** Half of the organizations cite adversarial advances powered by generative AI as a primary concern.
- **Shadow AI is a blind spot:** 78% of AI users are bringing their own tools to work, creating invisible workflow dependencies IT may not provision, govern, or even know about.
- **Customer trust is fragile:** 24% of customers abandon a brand after one bad experience. 70% leave after two.
- **Regulators are raising the bar:** DORA became enforceable across the EU in January 2025. UK PRA required firms to demonstrate impact tolerance compliance by March 2025.

Compliance and regulation

Operational resilience is now audited, not just expected. Regulators across industries want evidence of tested recovery, third party oversight, and controls that hold during disruption, not documentation of plans that were never tested.

What regulators converge on

Despite differences in terminology, FFIEC, DORA, PRA, HIPAA, PCI DSS, ISO 22301, and NIST SP 800-34 converge on the same proof requirements: defined critical services with limits for how long each service can be down; mapped dependencies, including third parties; tested continuity and recovery with evidence; controls that remained effective during disruption; and demonstrated ability to restore safely after change events.

DORA became enforceable across the EU on January 17, 2025. UK PRA required firms to demonstrate compliance with limits for how long each service can be down by March 31, 2025. U.S. FFIEC examiners are increasingly asking not just whether recovery plans exist, but whether they have been exercised and whether controls held.

What regulators expect

At the governance level, regulators want to see that an organization's most critical services are formally ranked; that each one has a named owner and a clear limit for how long it can be down; and that the board has approved a continuity plan that spells out what gets restricted during an incident and who can authorize exceptions.

On execution, they want evidence that drills have been run; lessons addressed; and services can be restored in a controlled way. A plan that has never been tested does not satisfy DORA, UK PRA, or FFIEC. Documented drill results and repeatable recovery steps do.

On controls, they want proof that the organization's security posture held during the incident itself; vendor access stayed limited to what was needed; audit trails were intact throughout; and employees working from personal devices did not create gaps in your data governance. The organizations that can show this leave audits in a stronger position than where they entered them.

If you can't prove it, you can't defend it.

Stage 1: Strategy and service tiering

Resiliency starts as a leadership decision, not a technical one: define which services are non-negotiable, what acceptable degradation looks like, and what outcomes matter most under pressure.

This is where business service tiers and limits for how long each service can be down are established so that incident decisions are defensible to the board, customers, and regulators before the next disruption happens.

If you only do one thing: Tier your services. If everything is critical, nothing is recoverable.

Service tiering is a leadership decision

Without an executive-approved view of what matters most, every incident becomes a negotiation. The outcome of that negotiation is measured in customer trust, regulatory confidence, and reputational impact. Leaders who have tiered their services in advance arrive at every incident with clear answers: what must stay available, what can degrade, and who owns the decision.

Tier 0 services cannot go down. Tier 1 services can be degraded but must remain functional. Everything below can fail in sequence. Once that map exists, everything else -application strategy, vendor management, dependency mapping, and infrastructure investment -flows from it.

Enabling resilient access

When disruption starts, employees need one reliable place to go. A unified workspace delivers exactly that — keeping critical applications accessible through a single, consistent entry point regardless of what is failing behind the scenes. Security policies apply automatically, so the right level of access is enforced without anyone having to make manual decisions under pressure.

The decisions made at this stage eliminate the negotiation that consumes the first hour of every incident. When everyone already knows what must stay available, what can degrade, and who owns the call, the response is faster, calmer, and far less costly.

Stage 2: Dependency and third-party risk mapping

Once Tier 0 to 1 services are defined, the next question is: what can stop them? Cloud outages most often start as partial failures -a specific region, identity provider, API, or third-party service degrades rather than disappearing entirely.

Mapping these dependencies by workflow is how organizations understand where their real exposure lies.

If you only do one thing: Map Tier 0 to 1 dependencies and third-party access by workflow before the next failure.

Partial failure is the common case

In a partial failure, the infrastructure dashboard is often green. The problem is upstream in a service the organization relies on but does not control. Without a dependency map, the first hour of every incident is spent establishing scope. This time translates directly into potential revenue impact, customer friction, and escalation.

Regulators increasingly hold organizations accountable for third party failures. DORA explicitly requires that critical service dependencies, including third parties, be

mapped, governed, and tested. The expectation is not just that you have a list or a plan; it is that you have tested it.

Reducing the scope of damage through access governance

How access is granted determines how far a failure can spread. When employees and vendors connect through broad network access, a single compromised account or failed service can take down everything that account touches. When access is controlled at the application level instead, the damage stays contained. One application goes down, not everything.

Third party access is often where the biggest unmanaged exposure lives. Contractors and partners tend to have more access than they need, through paths that no one audits regularly. Locking that down by workflow, so each vendor can only reach what their work actually requires, is the organizational equivalent of closing the fire doors before a crisis starts rather than scrambling to close them during one.

Stage 3: Continuity design and control

Resiliency is not about restoring systems after a disruption. It is about maintaining safe productivity while disruption is ongoing with tighter controls, auditable processes, and clear rules for exceptions that were agreed upon in advance.

If you only do one thing: Define continuity modes so controls tighten and evidence is captured while work continues.

The exception spiral and how to prevent it

The default response to disruption is to loosen controls to keep work moving. Access exceptions get granted, governance gets deferred, and it works - at least for a while. But by the time systems stabilize, data has moved in ungoverned ways; audit trails are incomplete; and the organization faces compliance consequences that persist long after the operational incident is resolved.

Preventing the exception spiral requires defining in advance what tightens, what is restricted, who can authorize exceptions, and what evidence is captured. This is not an IT decision; it is a governance decision that requires executive approval before the next incident.

Continuity modes that preserve control

A good continuity plan does not just keep the lights on. It keeps the right controls in place while the lights are on. That means access is granted only where it is needed; every action is logged; and anyone who needs emergency access gets it through a process that is tracked and auditable. When the incident ends, there is no clean-up required because governance never lapsed.

For web and SaaS tools accessed during an incident, policy-controlled browsing wraps those sessions in controls that travel with the user across any device. Files cannot be downloaded without authorization. Sensitive data cannot be copied. The organization keeps working and the audit trail stays clean.

Proof in practice

Mass General Brigham (U.S.) kept clinical workflows accessible across diverse endpoints during COVID surge. Scaled from 9,000 to 250,000+ virtual clinical visits per month.

Continuity should tighten controls, not loosen them.

Stage 4: Incident execution and recovery speed

When incidents occur, speed matters. Some disruptions are resolved more quickly than a full disaster recovery cutover can be completed safely. The key factor is executing recovery from a pre-approved playbook, not just assembling the war room quickly.

Make recovery executable: pre-approved recovery playbooks plus fast rollback to known good states.

Time-to-detect is where the actual cost accumulates

Most of the financial impact of a disruption occurs not during the incident itself, but in the hours before the scope is understood. Teams are comparing notes, checking dashboards, and escalating. By the time the war room has a clear picture, hours have passed, and customer impact is already underway.

Real-time observability that correlates user experience with upstream dependency signals compresses the gap between “something is wrong” and “we understand what and how many.” That compression is where recovery time is actually reduced.

Recovery as a practiced capability

The fastest recovery is the one you practiced. Pre-approved recovery playbooks with clear triggers, owners, communications, and rollback steps mean that when the incident happens, execution follows a plan rather than a conversation. Fast rollback to known good states after

bad updates, configuration drift, or compromised systems eliminates the need for manual rebuilding from scratch.

Complete audit trails and comprehensive activity logging mean that recovery can be demonstrated as well as executed. Leadership and regulators expect proof that recovery was controlled, not just that systems are back up.

Stage 5: Value realization and continuous improvement

Resilience is not a project with a completion date. It is an operating rhythm that compounds over time: fewer repeated failures, lower disruption cost, faster audits, and more predictable operations. That compounding effect is what transforms resilience from a risk management concern into a competitive advantage.

If you only do one thing: Drill and measure. Resilience is a muscle, not a memo.

The scorecard that funds the program

Resilience programs lose momentum when they are treated as isolated initiatives. The discipline that sustains them is regular measurement tied to business outcomes: service level objectives for Tier 0 to 1 services, mean time to recovery, disruption cost per quarter, audit findings closed, and drill results. A scorecard that connects these metrics to revenue protection and regulatory posture is what keeps resilience funded.

The governance model does not need to be complex - it needs to be consistent. Review outcomes for the most critical services every quarter. When failures repeat, fix the root cause rather than patching the symptom. Use the observability and performance data already available from your monitoring systems to show leadership that the program is working. That documented track record is what holds up when a regulator, auditor, or executive asks for proof.

If it can't be measured quarterly, it won't be funded annually.

Business resiliency outcomes and how Citrix enables them

The best-prepared organizations have already decided what tightens, what stays available, and how to recover before the next incident. Next is how the Citrix platform enables these outcomes.

Outcomes and what Citrix enables

- **Continuity of access**
  Keeps Tier 0 to 1 work reachable through a consistent workspace even when identity, cloud, or SaaS dependencies are degraded. Traffic steers around failing services automatically.
- **Control under pressure**
  Enforces zero trust access at the application layer. Tightens session controls and governs third party access during disruption instead of expanding exceptions.
- **Visibility to act fast**
  Correlates user experience with upstream dependency signals to shorten time to detect. Preserves full audit trails for compliance and incident review.
- **Repeatable recovery**
  Enables fast rollback to known good states after bad updates, drift, or compromise. Recovery is controlled, rehearsed, and fast, not manual and heroic.
- **Hybrid and multi-cloud flexibility**
  Reduces concentration risk. Supports data residency, latency, cost, and M&A constraints across on-premises, hybrid, and multiple cloud environments.
- **Secure enterprise browsing**
  Creates a governed safe lane for web and SaaS access during disruption. Reduces data leakage even on unmanaged or BYO devices.
- **Least privilege and third-party access**
  Shrinks the scope of damage with application level least privilege access. Provides auditable, governed access paths for contractors and partners.

Delivering resilience through a unified platform

Citrix's unified platform approach addresses the full range of challenges CIOs face during an incident: keeping people working, keeping controls tight, and getting back to normal faster.

Maintaining access under pressure starts with a stable entry point. The Citrix platform gives employees one place to reach their most important apps, desktops, and web-based workflows, whether workloads run on-premises, in the cloud, or split across both. The platform automatically routes work around failing services, so a degraded cloud region or struggling identity provider doesn't have to stop the business.

Keeping work accessible is only half the job. Under pressure, organizations inadvertently open access broadly just to keep things moving; that's where compliance problems start. The Citrix platform gives employees, contractors, and partners

access only to the specific applications they need, without putting them on the broader network. Policy-controlled environments wrap web and SaaS work, so data stays contained even on personal or unmanaged devices. Complete activity recording maintains audit trails throughout the incident, preserving evidence regardless of what else unfolds.

A U.S. investment firm used this approach during a live network issue and maintained a 94 percent work continuity rate without opening up broad access to compensate.

When something does go wrong, speed of detection and recovery determines cost. The observability tools within the Citrix platform give IT teams a real-time view of how users are being affected before the help desk is flooded with calls. That visibility compresses the time between "something is wrong" and "we know exactly what and how many people are impacted." Also, recovery becomes repeatable by rolling back systems to a known working state, rather than rebuilding from scratch.

A Big 5 accounting firm used this approach to eliminate restore windows that had previously taken hours, bringing entire teams back online in a fraction of the time.

Summary

Disruption is a given. How an organization responds to it determines whether customers stay; regulators are satisfied; and the business comes out stronger. The organizations referenced in this playbook share a common trait: they treated resilience as an operating discipline, not a crisis response. They ranked their most critical services; mapped the dependencies that could take them down; designed continuity plans that kept controls tight; and practiced recovery until it became routine.

The Citrix platform supports every stage of that journey by making resilience measurable, repeatable, and defensible.

Your next disruption is already on the calendar; you just don't know the date. The organizations that handle it best will be the ones that are already ready.

Citations

1. Cockroach Labs, "The State of Resilience 2025," Oct 2024. cockroachlabs.com
2. Business Continuity Institute (BCI), "Growing global momentum behind operational resilience," May 2025. thebci.org
3. Uptime Institute Intelligence, "Annual Outage Analysis 2025." uptimeinstitute.com
4. New Relic, "IT Outages Cost Businesses Up to $1.9M Per Hour," Oct 2024. newrelic.com
5. IBM Security, "Cost of a Data Breach Report 2024." ibm.com/reports/data-breach
6. Verizon, "2025 Data Breach Investigations Report (DBIR)." verizon.com
7. World Economic Forum, "Global Cybersecurity Outlook 2025." weforum.org
8. Emplifi, "State of Consumer-Brand Engagement 2025." emplifi.io
9. EBA/ESMA/EIOPA, DORA applicability from January 17, 2025. eba.europa.eu
10. UK FCA, "Operational resilience". impact tolerance deadline March 31, 2025. fca.org.uk
11. FFIEC IT Examination Handbook, "Business Continuity Management." ithandbook.ffiec.gov
12. NIST SP 800-34 Rev. 1, "Contingency Planning Guide for Federal Information Systems." csrc.nist.gov
13. ISO 22301:2019, "Security and resilience. Business continuity management systems." iso.org
14. PCI Security Standards Council, "PCI DSS v4.0." pcisecuritystandards.org

Enterprise Sales
North America | 800-424-8749
Worldwide | +1 408-790-8000

Headquarters
851 Cypress Creek Road, Fort Lauderdale, FL 33309, United States

20260710