Operational Resilience: The Preparedness Paradox

You stress-test for the last crisis – but the next one is always different

Digital Finance

BSc Digital Finance Course

Why Did the 2023 Banking Crisis Hit Exactly the Risks That Stress Tests Ignored?

The Preparedness Paradox

Banks spend billions on stress testing every year. They model credit shocks, market crashes, and liquidity squeezes with extraordinary precision. Yet every major crisis exploits precisely the risks that were not in the scenario library.

What stress tests covered in 2022:

- Credit losses from recession (GDP −4%, unemployment 10%)
- Equity market crash (−40% broad decline)
- Interest rate shock (+300bp parallel shift)
- Sovereign default in peripheral economies

What actually happened in 2023:

- Uninsured deposit concentration triggered bank runs via social media
- Interest rate risk in held-to-maturity bonds – not in trading books
- Speed of contagion: 48 hours from rumour to insolvency
- The risk was operational, not credit – *who* held the deposits and *how fast* they could leave

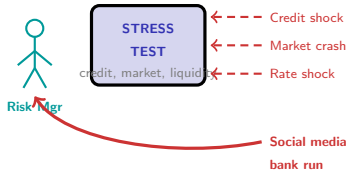

The shield blocks what you planned for.
The crisis comes from where you did not look.

The preparedness paradox: the more precisely you model known risks, the more vulnerable you become to the unknown risks you left out.

What Happens to Your Savings When Your Bank's IT System Crashes for a Week?

Thought Experiment

Your bank's mobile app, website, and ATMs all go offline at 9am on a Monday. You cannot check your balance, pay rent, or transfer money. Your employer confirms your salary was sent – but it has not arrived.

Day 1: You call customer service. Busy signal. You visit a branch – the tellers cannot access any accounts either.

Day 3: Your rent payment bounces. Your landlord charges a late fee. You borrow cash from a friend to buy groceries.

Day 5: The bank issues a press release: "We are working to restore services." No timeline. No explanation.

Day 7: Systems come back. Your balance is correct – but the late fees, the stress, and the broken trust remain.

This is not hypothetical. TSB Bank (UK, 2018) locked 1.9 million customers out of their accounts for weeks after a botched IT migration. Operational resilience is not about preventing every failure – it is about ensuring critical services continue even when systems fail.

Operational resilience is personal: when your bank's systems fail, you cannot pay rent, buy food, or access your own money.

What Is the Difference Between Business Continuity, Disaster Recovery, and Operational Resilience?

Dimension	Business Continuity	Disaster Recovery	Operational Resilience
Focus	Keep running during disruption	Restore systems after failure	Absorb, adapt, recover from any shock
Scope	Single process	IT infrastructure	End-to-end services
Question	"Can we keep going?"	"How fast to restore?"	"Can customers be served?"
Metric	Max tolerable downtime	RTO & RPO	Impact tolerance
Mindset	Plan for known events	Recover from known failures	Prepare for the unknown

The evolution: Business continuity asks "can we survive?" Disaster recovery asks "can we rebuild?" Operational resilience asks "can the customer still get served even when things break?" The shift is from internal systems to external outcomes.

Why the distinction matters

- **BC plans fail** when the disruption is not in the playbook – they assume you know what will go wrong
- **DR plans fail** when the backup site has the same vulnerability as the primary
- **Resilience succeeds** by focusing on outcomes: "customers can make payments within 4 hours" regardless of what broke

The DORA mandate (EU, 2025):

- Define important business services
- Set impact tolerances for each service
- Map dependencies (people, processes, technology, vendors)
- Test with realistic scenarios

Business continuity and disaster recovery are necessary but not sufficient. Operational resilience asks: can the customer still be served when the unexpected happens?

Follow One Bank Through a Coordinated Cyber-Attack – Minute by Minute

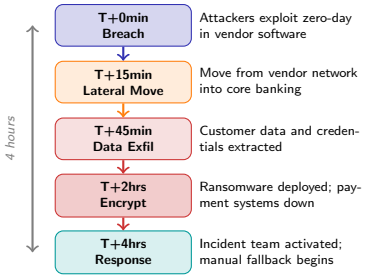

The resilience test

- **Detection:** SIEM flagged anomalous traffic at T+8min – but alert triaged as low priority
- **Containment:** Network segmentation slowed but did not stop lateral movement – vendor had privileged access
- **Communication:** Board notified T+3hrs; regulator T+4hrs (DORA: "without undue delay")
- **Fallback:** Manual payments resumed at T+6hrs – at 5% capacity
- **Recovery:** Full systems restored after 11 days

What resilience changed:

- Pre-agreed impact tolerance: payments within 4 hours
- Tested manual fallback reduced customer impact
- Vendor access controls tightened post-incident

The attack encrypted core systems in 2 hours. The resilience framework ensured customers could still make critical payments within 4 hours – the impact tolerance held.

How Do You Automate Stress Testing When the Scenarios Are Infinite?

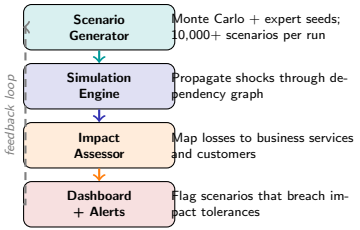

The automation pipeline

- **Scenario generator:** Combines historical events, expert hypotheticals, and random perturbations via Cholesky-correlated sampling
- **Simulation engine:** Models shock propagation – e.g., cloud outage disables payments, triggering liquidity stress
- **Impact assessor:** Translates losses into service outcomes – "customers cannot pay for 6 hours"
- **Dashboard:** Flags scenarios breaching impact tolerances; board sees a heat map, not a spreadsheet

Key insight: You cannot test infinite scenarios, but you can automate the pipeline so *new* scenarios run weekly, not annually.

Automated stress testing does not replace judgment – it amplifies it. Experts design seed scenarios; the machine explores the space around them.

What Happens When the Stress Test Itself Becomes a Source of Systemic Risk?

The stress test paradox
Regulators publish stress test scenarios so banks can prepare. But when every bank prepares for the *same* scenario, they all build the *same* defences – and develop the *same* blind spots.

How standardisation creates fragility:

- All banks hedge the same risks the same way – crowded trades amplify the shock
- All banks pass the test – creating false confidence
- Banks optimise for the test, not resilience – "teaching to the exam"
- Scenarios become public – adversaries can design attacks around them

The monoculture problem:

- Same cloud, same model, same scenarios – one failure mode takes down the system
- Diversity of approaches is itself a form of systemic resilience
- Reverse stress tests help: "what kills *this* bank?" has a different answer for each institution

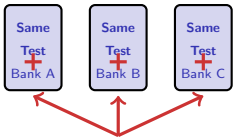

Same test, same blind spot, same systemic failure.

Standardised stress tests make individual banks safer but the system more fragile – every bank develops the same blind spots.

Where Are the Gaps in Current Stress Test Coverage?

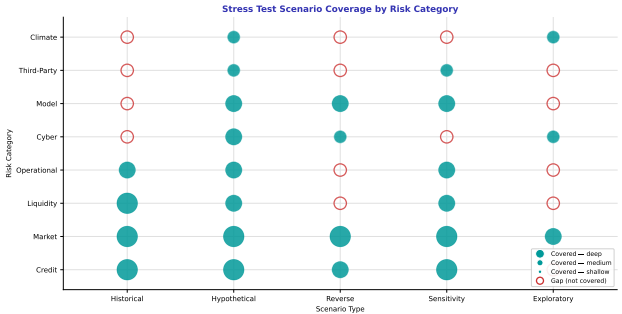

Illustrative coverage assessment. Circle size indicates depth of scenario analysis.

Reading the dot matrix

- **Filled circles:** scenarios actively tested; larger dots mean deeper analysis
- **Hollow red circles:** gaps – risk-scenario combinations not covered
- **Credit and market** are thoroughly tested – decades of regulatory focus
- **Cyber, third-party, and climate** show the most gaps – newer, harder to model
- **Reverse stress testing** is sparse outside credit and market

The pattern: Banks test what regulators require. Gaps cluster where regulations have not yet reached.

The dot matrix reveals a coverage paradox: the risks we understand best are tested most; the risks growing fastest are tested least.

Who Owns Operational Risk – the CRO, the CIO, or the Board?

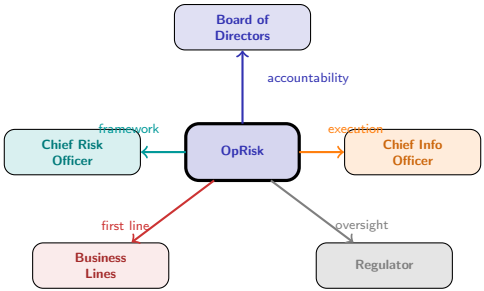

The ownership problem

- **Board:** Ultimately accountable under DORA – must approve impact tolerances and resilience strategy
- **CRO:** Owns the risk framework, sets appetite – but may lack technical depth
- **CIO:** Owns technology stack and DR plans – but may lack risk perspective
- **Business lines:** First line of defence – own processes but often underinvest
- **Regulator:** Sets rules but cannot monitor daily operations

The honest answer: Everyone owns a piece; nobody owns the whole. Resilience works only when accountability is explicit, tested, and enforced.

The governance gap in operational risk is not about who has the title – it is about who makes the call at 3am when systems are down.

The Resilience Assessment Framework: Beyond Check-the-Box Compliance

When evaluating any institution's operational resilience – as an investor, regulator, or employee – use these five tests:

1. Impact tolerances defined?
Has the institution set maximum tolerable disruption for each critical service – in customer outcomes, not system uptime?

2. Dependencies mapped?
Does it know every vendor, system, and process underpinning each critical service? Can it trace a payment end-to-end?

3. Scenarios tested – not just planned?
Has it run a realistic disruption exercise in the past year? Not a tabletop, but an actual test with systems offline.

4. Lessons embedded?
After the last incident or test, what changed? If the answer is "we updated the documentation," that is not resilience.

5. Board engaged?

Can the board articulate the top three resilience risks?

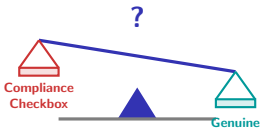

Most institutions tip toward compliance.
The five tests reveal which side you are on.

Compliance asks "did we follow the rules?" Resilience asks "can we still serve customers when things break?" The five tests separate the two.

Your Challenge: Design a Reverse Stress Test for a Digital Bank

Your Challenge

The scenario: A digital-only bank with 2 million customers, no branches, all infrastructure on a single cloud provider. Critical services: (1) payments, (2) account access, (3) customer support chat.

Your task: Design a reverse stress test – work backward from failure:

Step	Your Analysis
1. Define "failure" for each service	.....
2. What single event could cause all three to fail?	.....
3. Most plausible path to that event?	.....
4. What control, if it failed, makes it inevitable?	.....

Discuss with your neighbour:

- Shortest path from "normal" to "all three services down"?
- Which dependency is the single point of failure?
- How does your answer change with two cloud providers?

Reverse stress testing inverts the question: instead of "how bad could this be?" it asks "what would kill us?" – and works backward to the most plausible path.