INSIGHT SERIES · GOVERNANCE, RISK & HUMAN FACTORS

When the Ground Shifts

Resilience is not the ability to avoid disruption. It is the ability to absorb it, adapt to it, and keep moving — and the organisations that do this best have spent years building the capability before they needed it.

Author: E van Doorn, Analysis & Opinion · **Date:** June 2026 · **Topic:** Governance / Risk

EXECUTIVE SUMMARY

Operational resilience — the demonstrated capacity to absorb disruption and maintain critical functions under pressure — has moved from a compliance afterthought to the central governance challenge of our time. Drawing on the Fukushima disaster, the British Airways IT collapse of 2017, and the Danish financial sector's contingency model, this article examines the four capacities that resilient organisations build: preparation, absorption, adaptation, and recovery. It argues that the gap between business continuity documentation and actual operational capability is, in most organisations, wider than leadership believes — and that structured supplier oversight, genuine testing, and the quality of human relationships are the factors most likely to determine whether that gap matters when disruption arrives.

1. Introduction

On the morning of March 11, 2011, a magnitude 9.0 earthquake struck off the Pacific coast of Japan's Tōhoku region, triggering a tsunami of devastating scale that killed nearly 20,000 people. At the Fukushima Daiichi nuclear power plant, the loss of external power, the failure of backup generators, and the subsequent meltdowns in three reactor cores produced the worst nuclear disaster since Chernobyl — a crisis that at its most acute threatened the evacuation of the entire Tokyo metropolitan area, home to 37 million people. What the disaster revealed was not that catastrophic events cannot be planned for, but how organisations fail when multiple systems that each work in isolation fail simultaneously, when the dependencies between them were never fully mapped, and when the people responsible lack the shared mental models and clear decision rights to act decisively under extreme uncertainty. It revealed, in other words, that resilience is not a property of individual systems. It is a property of the arrangements — human, structural, and relational — that determine how an organisation holds together when its environment stops cooperating.

> *"Resilience is not a property of individual systems or individual organisations. It is a property of the arrangements — the human, structural, and relational arrangements — that determine how an organisation holds together when its environment stops cooperating."*

That insight has become the animating principle of one of the most rapidly developing areas in organisational risk. Operational resilience — distinct from, and more demanding than, business continuity planning in its traditional form — has moved from the margins of governance thinking to near the centre of regulatory expectation in financial services, critical infrastructure, healthcare, and beyond. The question it asks is deceptively simple: if the worst happens, can you keep going? And if you cannot, do you at least know which of your functions matters most, why it matters, and what it would take to restore it?

2. Resilience and Continuity: An Important Distinction

Business continuity planning, in its classical form, is a documentation exercise. It produces plans — for data backup, for alternate site activation, for crisis communication, for vendor substitution — that

describe what the organisation will do when specific, anticipated scenarios occur. Business continuity plans are valuable. They are also, in many organisations, among the most thoroughly written and least thoroughly tested artefacts in the governance library.

Operational resilience is a different concept. It is not primarily concerned with plans — it is concerned with outcomes: the organisation's actual capacity to maintain the delivery of its most critical functions through disruption, regardless of the specific nature of that disruption. The resilience question is not "do we have a plan for this?" but "can we actually do this, under pressure, with degraded resources and incomplete information?" An organisation with an excellent business continuity plan and no tested resilience capability is not a resilient organisation. It is an organisation that has documented its intentions without verifying its capacity to realise them.

> **Key Point:** The Bank of England's operational resilience policy, in full force from March 2025, requires financial services firms to identify important business services, set explicit impact tolerances, and demonstrate through testing that they can remain within those tolerances through severe but plausible disruptions — deliberately agnostic about the specific form of disruption.

Similar frameworks are taking shape in telecommunications, energy, healthcare, and across the operational technology sectors that underpin critical national infrastructure. The direction of travel is consistent: from planning as a compliance output to resilience as a demonstrated operational capability.

3. The Four Capacities

Resilience, examined across the organisations that display it most consistently, rests on four capacities that operate across different time horizons and that together constitute what might be called the resilience cycle: preparation, absorption, adaptation, and recovery.

3.1 Preparation

Preparation is the work done before disruption arrives: the identification of what matters most and the investment in structural conditions that reduce both the likelihood and the impact of disruption to critical functions. It includes the obvious — backup systems, redundant infrastructure, tested incident response procedures — but also the less obvious: mapping of internal and external dependencies, management of supplier and contractor relationships through which a significant

proportion of operational risk now flows, and maintenance of the institutional knowledge and human relationships that allow effective response when formal procedures prove inadequate. The organisations that respond best to disruption are almost invariably the ones whose people know each other, trust each other, and have practised working together under pressure. A crisis management team that has never run a simulation together is not a crisis management team — it is a collection of individuals who have been assigned roles on a chart.

3.2 Absorption

Absorption is the capacity to take a shock without losing the ability to function — to experience a significant disruption without it cascading into a loss of critical services. It is a function of slack: the resources, redundancy, and excess capacity that most organisations under efficiency pressure are inclined to eliminate. The tension is genuine: redundancy costs money continuously while its benefit is hypothetical and intermittent. The organisations that have maintained genuine absorption capacity tend to be those whose leadership has been persuaded, either by experience or by clear-eyed analysis of their risk environment, that the cost of redundancy is lower than the cost of its absence at the moment it is needed.

3.3 Adaptation

Adaptation is the capacity to reconfigure under pressure — to find new ways of delivering critical functions when the normal ways are unavailable. It is less well-developed than the other three capacities in most resilience frameworks, but in many respects the most important. Real disruptions rarely conform to planned scenarios: a cyberattack does not behave like the one in the business continuity plan, and a supplier failure happens differently from how it was modelled in the risk register. The organisation that can only respond according to its playbook is dependent on the disruption being kind enough to resemble the disruption it prepared for. The organisation that has built genuine adaptive capacity — investing in breadth of skills, flexibility of process, and clarity of authority that allow intelligent improvisation — is resilient in the deeper sense.

3.4 Recovery

Recovery is the capacity to restore full function after a disruption has been absorbed and managed. It is the least glamorous dimension of resilience and the one that most directly affects long-term consequences. An organisation that absorbs a shock well but recovers slowly will lose customers, relationships, and market position in the period between stabilisation and restoration. Recovery is

also where lessons from a disruption are either captured or lost — the organisations that learn most from adverse events are the ones that have built structured post-incident review into their resilience frameworks, treating the period after a crisis as the most information-rich environment they will encounter, where the gaps between plan and reality are most visible and the motivation to close them most acute.

4. Case Study I: The Disruption That Exposed a System — British Airways IT Failure, 2017

On the late May bank holiday weekend of 2017, British Airways suffered a catastrophic IT failure that grounded its entire global operation for the better part of three days. Approximately 75,000 passengers were stranded or severely disrupted, over 700 flights were cancelled, and the eventual cost — in compensation, operational disruption, regulatory fine, and reputational damage — ran to hundreds of millions of pounds. The immediate cause was a power supply issue at a Heathrow data centre. The deeper cause was the state of British Airways' IT resilience.

In the years preceding the incident, the airline had undertaken a significant programme of IT outsourcing and cost reduction. The infrastructure that remained was complex, ageing in parts, and characterised by dependencies and single points of failure that were not fully mapped or understood by those responsible for managing them. When the power event occurred, the safeguards that should have prevented cascading system-wide failure did not perform as intended. Recovery procedures that should have allowed systems restoration within hours took days. And crisis management arrangements were overwhelmed by the scale of disruption and the absence of manual workarounds that a previous generation of airline operations had maintained as a matter of course.

> **Key Point:** The British Airways case is a textbook illustration of the most common resilience failure mode: not the failure to plan, but the failure to maintain the operational capacity to execute when the plan is needed. The airline had continuity documentation. It did not have a continuity capability. The gap between the two — exposed over a bank holiday weekend with a full flight programme — produced one of the most damaging operational failures in the history of British commercial aviation.

5. Case Study II: Resilience by Design — The Danish Financial Sector's Contingency System

In contrast to the British Airways experience, the Danish financial sector offers one of the most instructive examples of deliberately designed operational resilience in the world. Since the 1990s, Denmark's major financial institutions have operated a shared contingency system — known as the Sumclearing — for interbank payments, maintained alongside the primary digital infrastructure as an explicit fallback in the event of systemic failure. The philosophy is unusually frank: the assumption is not that primary infrastructure will always work, but that it will at some point fail, and that the consequences for the financial system and broader economy are severe enough to justify permanent maintenance of a less efficient but more robust alternative. The redundancy is not incidental. It is the point.

Danish financial regulators have, over decades, required institutions to participate in regular, coordinated resilience testing — not of individual firms' systems in isolation, but of the sector's collective capacity to maintain critical payment and settlement functions through scenarios in which primary systems are unavailable. This testing is genuine: it involves the actual activation of fallback systems, actual involvement of operations staff who must demonstrate they can work in a degraded environment, and actual identification of gaps that must then be remediated. The result is a financial system that has, on several occasions when digital infrastructure has been disrupted, maintained critical functions without the cascading failures that similar disruptions have produced in markets where resilience investment has been less systematic.

6. The Supplier Dependency Challenge

No discussion of operational resilience in the modern organisation can avoid, for long, the question of what happens to resilience when a critical supplier fails. This is not a theoretical risk — it is, increasingly, among the most common precipitating causes of operational disruption, and among the least well-managed. Every function that an organisation has outsourced to a supplier is a function whose resilience is, at least in part, a function of the supplier's own resilience. The organisation that has outsourced its payroll processing, cloud hosting, cybersecurity monitoring, or facilities management has not outsourced its resilience obligation. It has created a chain of dependencies through which its resilience obligation now flows — and over which it has, in most cases, considerably less visibility and control than it has over its own operations.

The failure mode is well-documented. It begins with supplier drift: the gradual accumulation of small lapses in supplier oversight that, individually, are unremarkable but that collectively represent a meaningful degradation in the organisation's knowledge of its own dependency base. A renewal missed. A compliance attestation outstanding for longer than it should be. A critical supplier's

business continuity plan unreviewed for twelve months or more. An insurance certificate that lapsed six months ago without anyone noticing. None of these is an emergency in isolation. Together, they constitute an oversight environment in which the organisation does not actually know the resilience posture of the suppliers on whom its critical functions depend.

Structured supplier oversight is designed to prevent this — not by eliminating the dependency, which is often neither possible nor desirable, but by maintaining the visibility that allows the organisation to know, at any given moment, whether its suppliers are current, compliant, and capable of meeting their obligations under stress. Platforms such as RailGuard address this directly: by maintaining a live provider register that surfaces which suppliers are current, overdue, or missing required information, and by automating the reminders that prevent the drift that is otherwise endemic to supplier relationship management conducted across email threads, shared drives, and individual memory.

> **Key Point:** An organisation that discovers a critical supplier's insurance has lapsed or their business continuity documentation has not been reviewed for eighteen months — at the moment a disruption requires reliance on that supplier — is in a fundamentally different position from one that identified and addressed those gaps proactively. The structured decision history that good oversight creates also preserves institutional memory: the record on which sound crisis management depends.

7. Testing: The Gap Between Planning and Capability

Of all the components of an effective resilience programme, testing is the one most consistently underdone and most consistently consequential in its absence. The reason is not mysterious: testing is disruptive, resource-intensive, and revealing of gaps that organisations would often prefer not to have revealed. The incentive to defer it — or to conduct it in a form thorough enough to satisfy a regulatory requirement but not thorough enough to surface genuine weakness — is significant.

The organisations with the strongest resilience track records resist that incentive. They test genuinely — including scenarios they find most uncomfortable, involving functions they are most dependent on and suppliers whose failure would be most damaging. They involve the people who would actually manage a crisis, not just those whose names appear on the crisis management chart. They debrief honestly and act on what the testing reveals. The regulatory frameworks increasingly require this approach: the Bank of England's operational resilience rules require firms to test their ability to remain within impact tolerances; the EU's Digital Operational Resilience Act requires threat-led penetration testing at regular intervals; and CISA's critical infrastructure resilience guidance

emphasises exercise programmes that go beyond tabletop discussions to involve actual activation of fallback systems and manual procedures. What all of these frameworks are reaching for is the same thing: closing the gap between the resilience an organisation believes it has and the resilience it can demonstrate.

8. The Human Foundation

Behind all the frameworks, systems, and regulatory requirements, the most important determinant of organisational resilience is something that cannot be documented, tested, or certified: the quality of the people who will manage a crisis when it arrives, and the quality of the relationships between them. Crisis management at its best is a human performance — a function of clarity of purpose, speed of judgment, quality of communication, and the trust between individuals that allows decisions to be made and executed without the friction that uncertainty and unfamiliarity create.

Leadership behaviour in the period before a crisis determines, to a remarkable degree, the quality of the leadership response when the crisis arrives. Leaders who have made clear, through consistent behaviour over time, that they want accurate information rather than comfortable information — that they reward the identification of problems rather than penalising it, and that they take resilience investment seriously rather than treating it as a compliance cost — will find, when disruption comes, that their organisations are capable of a quality of response that reflects the investment made. The leaders who have not will find that the gap between their business continuity documentation and their actual operational capability is exposed at precisely the moment when exposure is most costly.

Conclusion

The organisations that have genuinely achieved resilience — that have demonstrated, through adverse events and rigorous testing, that they can absorb disruption, adapt under pressure, and recover without catastrophic loss of function or reputation — tend to describe the condition less in terms of systems and documentation than in terms of clarity. They know what matters most. They know their dependencies — including third-party dependencies — well enough to understand where their exposure lies. They have tested their capacity to perform under pressure often enough that the gap between plan and capability is narrow and visible rather than wide and hidden. And they have the leadership culture that ensures honest information travels fast and difficult decisions are made by the right people with the right authority.

None of this is achieved once and then maintained without effort. Resilience is not a project with an end date. It is a continuous organisational practice, requiring sustained investment of attention, resource, and leadership commitment, against the persistent pressure of operational efficiency and the natural human tendency to underweight risks that have not yet materialised. The disruptions will come — in forms that are partly anticipated and partly surprising, at moments that are never convenient, with consequences that depend heavily on the quality of preparation that preceded them. The organisations that meet them with confidence will be the ones that decided, before the disruption arrived, that building the capacity to absorb, adapt, and recover was worth the cost. It always is.

This article draws on the Investigation Committee on the Accident at the Fukushima Nuclear Power Stations report (2012), the UK Civil Aviation Authority's investigation into the British Airways IT disruption of May 2017, the Bank of England's Operational Resilience Policy Statement (2021, effective 2025), the EU Digital Operational Resilience Act (Regulation 2022/2554), and published analysis of the Danish financial sector's contingency payment infrastructure by Danmarks Nationalbank. Information on RailGuard's platform capabilities is drawn from published documentation at assessum.com.