Operational Resilience for Portfolio Companies: Critical Services, Tolerances and Tests

A Board Framework for Protecting Service, Cash Flow and Enterprise Value through Disruption

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Operational resilience is the ability to keep delivering the services that matter when people, technology, facilities, data, suppliers or infrastructure fail. It is broader than disaster recovery and more decision-oriented than a register of operational risks. For an investor-owned business, resilience affects customer retention, cash conversion, covenant headroom, regulatory standing, insurance recovery, management credibility and the range of financing or exit options available to the board.

This paper develops a board framework for identifying critical services, setting measurable impact tolerances, mapping end-to-end dependencies and testing severe but plausible disruption scenarios. It links service performance to customer harm, contractual obligations, liquidity, EBITDA, enterprise risk and equity value. It also distinguishes service tolerances from system recovery objectives, separates tested capability from documented intention and connects remediation to the ownership agenda.

The analysis draws on current governance, operational-resilience, business-continuity, cyber, third-party and testing sources from international bodies and authorities in the United Kingdom, European Union, United States, United Arab Emirates, Saudi Arabia, India, Singapore and Australia. Several cited requirements apply only to specified regulated entities. They are used as reference architecture for a proportionate portfolio-company operating model and do not establish legal applicability to another company or jurisdiction.

All values, scores, time periods, tolerances, scenarios and outcomes in this paper are hypothetical management assumptions used to demonstrate the framework. They do not describe a client, forecast, valuation opinion, insurance position, regulatory conclusion or assured result. The paper does not provide legal, regulatory, accounting, tax, investment, insurance, cyber-security, engineering, health-and-safety or technical advice. Responsible executives and qualified advisers should assess the company, contracts, sector, jurisdictions and applicable requirements.

JEL Classification: G24, G32, G34, L21, M11, M15, M21

Keywords: operational resilience, critical services, impact tolerances, scenario testing, business continuity, third-party risk, portfolio operations, private equity, value creation, board governance

1. Treat resilience as a service-and-value system

Operational resilience begins with the service received by a customer, counterparty or other stakeholder. A business may have functioning servers, staffed offices and available cash while still failing to deliver the outcome that matters. An industrial-services company can restore its scheduling platform yet remain unable to dispatch technicians because access permits, specialist tools or subcontractors are unavailable. A healthcare operator can recover clinical systems while a laboratory dependency prevents safe treatment. A subscription business can keep its application online while identity verification, billing or customer support remains impaired.

The board should therefore govern resilience from the outside in. It identifies which services must continue, the harm created by disruption, the point at which that harm becomes unacceptable and the resources required to remain within that boundary. The resulting map joins customers, contracts, revenue, people, process, technology, data, facilities, utilities, equipment, suppliers, logistics and decision rights. It turns a collection of continuity documents into an operating system.

For an investor-owned company, the value chain has a second dimension. A disruption changes service levels, backlog, revenue recognition, refunds, working capital, variable cost, remediation spending, insurance claims and cash. It can also affect licences, customer confidence, supplier terms, covenant forecasts and transaction diligence. Resilience investment should be prioritised through that complete chain rather than through the replacement cost of an isolated asset.

The governing discipline is evidence. A policy demonstrates intention. A dependency map demonstrates understanding. A test demonstrates capability under specified conditions. A live incident demonstrates capability under actual conditions. A closed remediation record demonstrates learning. The board needs each form of evidence, with its scope and limitations clearly stated.

Figure 1. From critical service to equity-value protection

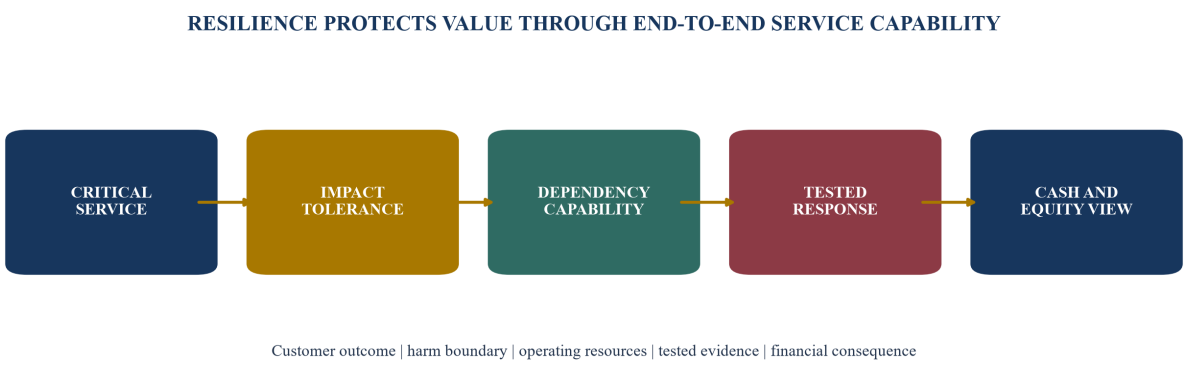

A resilience decision becomes economically relevant when service, harm, operating dependencies and financial consequences are connected.

2. Use a precise operating language

Operational risk, business continuity, disaster recovery, crisis management, cyber security and operational resilience overlap but answer different questions. Operational risk identifies and manages the possibility of loss or disruption arising from people, processes, systems or external events. Business continuity organises how priority activities continue or resume. Disaster recovery focuses on the recovery of technology and data. Crisis management establishes command, escalation and communications. Cyber security protects information and systems. Operational resilience asks whether the organisation can continue delivering critical services through disruption within an approved tolerance.

The distinction matters because a system objective can be met while the service tolerance is breached. A database may recover within two hours, yet the service may require additional time for reconciliation, control checks, queued transactions and customer communication. Conversely, a service can remain within tolerance through manual processing even when the primary system misses its recovery objective. The board should see both measurements and the dependency between them.

The Basel Committee defines operational resilience for banks around the ability to deliver critical operations through disruption and encourages an assumption that disruption will occur [17][18]. The Financial Conduct Authority uses important business services and impact tolerances, with mapping and testing to show whether a firm can remain within them [1][2]. ISO 22301 provides a generic management-system framework for continuity across organisations of different types and sizes [27]. These sources use different scopes and legal effects. Their shared logic is useful: identify the outcome, define the boundary, understand the delivery system, test capability and learn.

A portfolio company should establish a controlled glossary. Each term needs an owner, definition, unit of measure, source system and relationship to other measures. This prevents management from using recovery time objective, maximum tolerable downtime, impact tolerance and service-level agreement as interchangeable labels. It also makes portfolio reporting comparable without forcing different businesses into the same service model.

3. Identify critical services through harm and materiality

A critical service is an end-to-end outcome whose disruption beyond a defined point would create unacceptable harm. Harm may affect customers, patients, employees, communities, markets, contractual counterparties or the company itself. It can arise through safety events, inability to access essential products, lost or corrupted data, financial loss, legal or regulatory breach, environmental damage, liquidity pressure or a threat to the company's viability.

Selection should begin with the service rather than an organisational chart. "Information technology", "operations" and "finance" are functions. "Process a customer withdrawal", "maintain cold-chain delivery", "dispatch emergency field repair", "administer a regulated claim" and "release a safety-critical batch" are services. A service statement should name the recipient, outcome, start and end points, minimum acceptable capacity and the events that count as disruption.

The board can use five lenses. The customer lens asks who loses an essential outcome and how quickly. The obligation lens tests contracts, licences, safety duties and public commitments. The financial lens estimates revenue, cash, margin, remediation and liquidity effects. The concentration lens identifies services whose failure affects many products or entities. The strategic lens asks whether disruption would impair the investment thesis, customer franchise, financing plan or exit route.

The list should remain selective. Labelling every process critical dilutes resources and creates an untestable programme. Management can retain a wider business-impact inventory while designating a smaller board-approved critical-service perimeter. The rationale, exclusions and review triggers should be documented. Acquisitions, new products, site changes, outsourcing, system migrations, major customers and revised regulation should trigger reassessment.

Table 1. Critical-service identification screen

lens	evidence question	example measure	escalation signal	board decision
customer or stakeholder	who loses which outcome?	affected users, vulnerable groups, delayed units	harm accelerates with time	include, narrow or exclude service
contractual and legal	which obligations continue through disruption?	notice time, delivery duty, licence condition	breach can occur before recovery	approve control and advice pathway
safety and environment	can disruption cause injury or environmental harm?	exposure window, safe-state capacity	safe operating boundary may be crossed	define stop authority and fallback
financial	how do service loss and recovery consume value?	revenue, cash, refunds, cost, working capital	liquidity or covenant headroom weakens	fund capacity or change tolerance
concentration	which shared resources support many services?	services per dependency, substitutability	one failure crosses several tolerances	diversify, isolate or create fallback
strategic and transaction	can failure impair the ownership case?	customer retention, licence, diligence finding	exit or financing path is constrained	integrate into value-creation plan

The board can combine harm, obligation, financial and concentration evidence without reducing the decision to one score.

4. Set impact tolerances as decision boundaries

An impact tolerance states the maximum disruption a service can sustain before harm becomes unacceptable. Time is usually necessary and rarely sufficient. The tolerance may also need minimum service capacity, maximum backlog, maximum data loss, customer segment, transaction value, geographic scope, safety state or communication interval. A single four-hour tolerance can conceal a service that must retain ninety per cent capacity for vulnerable customers while lower-priority activity can pause.

The tolerance should be derived from harm and obligations, then tested against capability. Setting it equal to current recovery performance converts a weakness into an approved boundary. Setting it without reference to operations produces an aspiration. The board needs two facts: the tolerance required by the service and the capability demonstrated by evidence. Any gap becomes a funded decision with an owner and target date.

The metric hierarchy should remain explicit. Maximum tolerable downtime describes the outer time boundary for a process or service. Recovery time objective sets the target time for restoring a supporting resource and should allow for service reconciliation and restart. Recovery point objective sets the acceptable data-loss point. Minimum business continuity objective defines the capacity or service level delivered under alternative arrangements. Contractual service levels describe promises to customers or between internal teams. None should silently replace the impact tolerance.

Tolerance calibration requires scenarios. The same service may survive an isolated server failure and fail during a regional power interruption, supplier outage or cyber incident that corrupts identity, communications and backups together. Management should test the tolerance against plausible combinations, downstream queues and recovery dependencies. The board then approves the boundary, capability gap, interim controls and residual exposure.

Table 2. Board-approved service tolerance card

field	decision content	evidence source	example governance question
service and recipient	end-to-end outcome and affected stakeholder	service catalogue, contracts, journeys	is the outcome stated outside-in?
harm boundary	unacceptable customer, safety, legal or financial effect	impact analysis, advice, incident history	what makes further disruption unacceptable?
time tolerance	elapsed time to harm boundary	process timing and scenario analysis	when does the boundary occur?
capacity and backlog	minimum output and maximum queue	operational throughput and demand data	can priority demand be served?
data tolerance	acceptable loss and reconciliation requirement	data flows, backup and control design	what must be accurate before restart?
supporting objectives	system, site, supplier and people recovery targets	architecture and continuity plans	do objectives leave time for service restart?
tested capability	last demonstrated result and conditions	test record and incident evidence	what has actually been shown?
gap and decision	remediation, interim control, owner and date	funded action plan	who accepts exposure until closure?

The card separates the harm boundary from supporting recovery objectives and current tested capability.

5. Map the complete delivery chain

Dependency mapping shows how a critical service is produced. The minimum categories are people, process, technology, data, facilities and third parties. Many companies also need utilities, equipment, inventory, transport, permits, communications, intellectual property and decision authority. The map should extend through material subcontractors and shared services where the company depends on them.

The useful unit is the service step. For each step, management records the input, output, owner, location, system, data, equipment, supplier, control, alternative route and recovery requirement. It identifies single points of failure and common-mode dependencies. Two suppliers provide little resilience if both use the same port, cloud region, power feeder, specialist component or

subprocessor. Two systems provide little resilience if they share identity, network, database or deployment pipelines.

Mapping depth should be proportionate to the tolerance and uncertainty. A high-harm service with a short tolerance needs transaction-level and infrastructure-level detail. A lower-harm service can use a simpler process map. The board does not need every technical object. It needs enough detail to see vulnerabilities, resource conflicts, concentration and the basis of tested capability.

The map should be governed as operational data. New vendors, product changes, role moves, facility closures, cloud migrations, acquisitions and system releases can make it stale. Owners should attest to material changes, and tests should compare the map with what participants actually use. Differences are findings, not administrative errors. They reveal hidden dependencies and informal workarounds.

Figure 2. End-to-end dependency map for a critical service

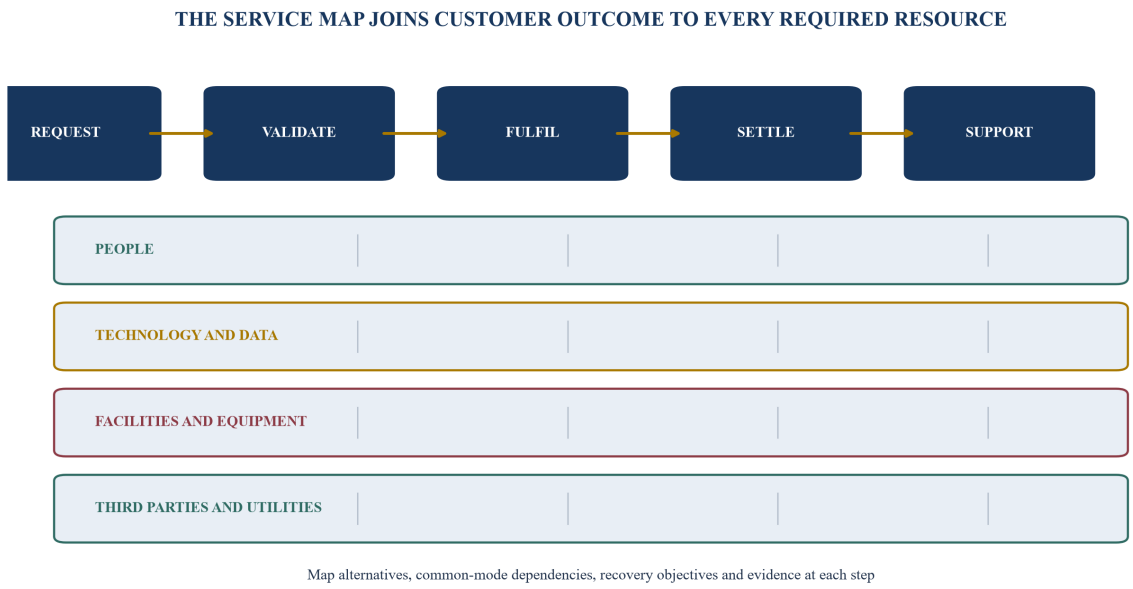

The map exposes shared resources, concentration and recovery requirements across the service chain.

6. Build a baseline of capability and evidence

The first resilience assessment should distinguish design, implementation and performance. Design asks whether the service, tolerance, map, plan and controls exist and are internally coherent. Implementation asks whether named people, suppliers, systems, facilities and alternatives are actually available. Performance asks whether the company has demonstrated the required outcome through tests or incidents.

Evidence should be traceable. Service records need owners and approval dates. Dependency maps need source systems and change dates. Contracts need service, audit, notification, data, continuity, subcontracting and exit terms. Technology evidence needs backup results, restoration logs, configuration, capacity and failover records. People evidence needs rosters, alternates, authority and training. Tests need objectives, conditions, observations, measured results and remediation closure.

The baseline should also record uncertainty. A service may have a documented four-hour recovery objective but no end-to-end test. A supplier may hold certification while the company's specific service and recovery priority remain unclear. A backup may complete daily while restore time under production volume is unknown. These are not equal evidence. Management can use confidence grades such as verified, partly verified, documented only and unknown, with objective definitions.

The portfolio view should aggregate themes without hiding company-specific exposure. Common findings may include untested manual workarounds, identity concentration, unsupported equipment, fragile key-person dependencies, inadequate generator autonomy, unclear data reconciliation or vendor exit plans that assume cooperation during failure. The sponsor can fund shared expertise and establish minimum evidence, while each board retains accountability for its services and risk decisions.

7. Design a test programme that produces decisions

Testing should answer a decision question. A tabletop exercise tests judgement, roles, escalation and communications. A walkthrough tests whether procedures and information are usable. A component test measures a system, backup, call tree or generator. A failover test moves activity to an alternate resource. A simulation combines people, process and systems under controlled conditions. A live exercise introduces operational load and may involve customers, suppliers or authorities. Each format provides different evidence.

The programme should progress from safe and narrow to integrated and demanding. NIST distinguishes tests, training and exercises and recommends objectives, scenarios, conduct and evaluation [20]. The FCA expects testing to mature from judgement-based and desk-based work toward empirical evidence such as failover tests, simulations and lessons from actual incidents [3]. CISA provides scenario packages and resilience assessments that help organisations structure cross-functional exercises [22][23].

Every test should define scope, excluded conditions, initial state, participants, injects, expected decisions, metrics, stop criteria and evidence capture. The team should record elapsed times rather than reconstructing them later. Observers need authority to identify unsafe conditions and preserve records. External providers should participate when their capability is material to remaining within tolerance.

A successful test is not one in which the plan appears to work. It is one that provides reliable evidence and identifies decisions. A service can pass the headline tolerance while revealing a dangerous manual-control weakness. It can fail for a condition that was excluded from the approved tolerance. Results should show service performance, harm, data integrity, decisions, communication, resource consumption, assumptions and findings separately.

Table 3. Evidence ladder for resilience testing

format	primary question	evidence produced	limitation	suitable progression
document review	is the plan internally complete?	controlled plan and gap list	no execution evidence	before every exercise cycle
walkthrough	can owners follow the procedure?	role, contact and sequence findings	low operational realism	after material plan change
tabletop	can leaders make timely decisions?	decision log and communication gaps	systems may remain untested	board and crisis-team rehearsal
component test	does a resource meet its objective?	measured technical or facility result	service outcome may be unknown	backups, generators, call trees
failover or restoration	can activity move and reconcile?	recovery, data and control evidence	scope may exclude peak demand	staged production-equivalent test
integrated simulation	can the service remain within tolerance?	end-to-end performance and harm result	requires strong safety controls	mature annual programme
live incident review	what happened under real conditions?	actual response, recovery and learning	scenario is uncontrolled	feed lessons into next test

The test format should match the decision and the risk created by exercising production capability.

Figure 3. Progressive scenario-testing architecture

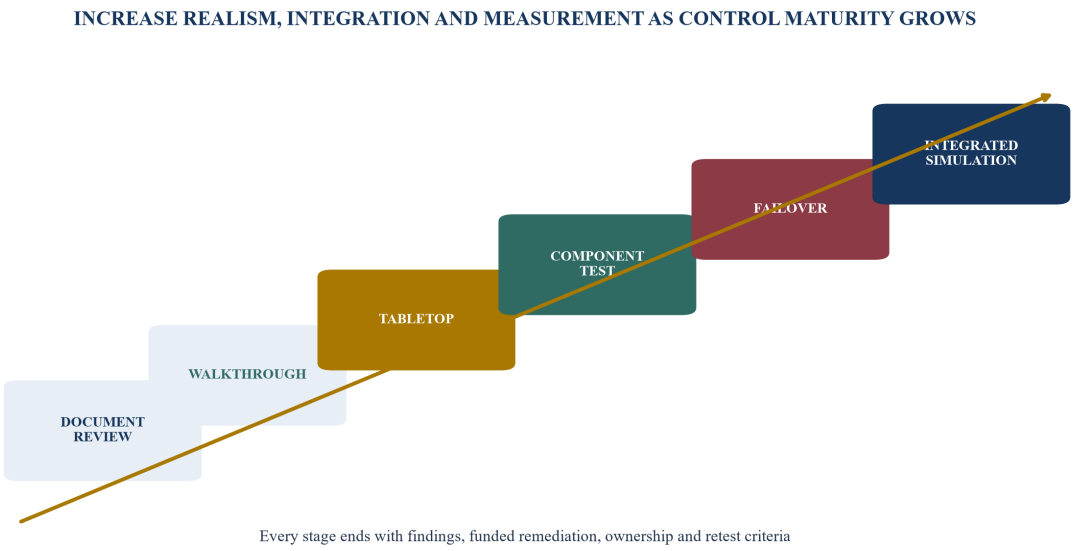

Evidence matures from plan review to integrated service testing and lessons from live incidents.

8. Build severe but plausible scenarios around vulnerabilities

A scenario should challenge the service and the assumptions supporting it. Generic labels such as cyber attack, supplier failure or flood are insufficient. The scenario needs a cause, affected resources, duration, geographic scope, demand condition, information quality, third-party behaviour and recovery constraints. It should be severe enough to reveal the boundary while remaining credible for the company's exposure.

Scenario design begins with vulnerabilities in the map. A single cloud region suggests regional technology failure. One qualified component supplier suggests insolvency, quality hold or trade

disruption. A thin engineering team suggests simultaneous key-person absence during a critical release. Shared identity suggests compromise that affects primary and backup environments. A site dependent on water, cooling or power suggests utility loss with constrained replenishment. An acquisition suggests incompatible escalation, data and authority during integration.

The programme should test combinations and duration. Disruptions rarely remain isolated. A ransomware event can also remove communications, supplier portals, credentials and reliable data. A regional weather event can affect staff travel, power, logistics and customer demand. A product-quality event can create a recall, media pressure, cash refunds and regulator engagement at the same time. The service map helps the team combine factors without inventing an implausible spectacle.

Scenario severity should be escalated through measurable injects. Demand rises, capacity falls, the outage extends, a key supplier becomes unavailable, an alternate site is delayed or information is uncertain. The test continues until the service approaches or crosses tolerance, subject to safety stop criteria. The result shows the vulnerability threshold and the additional capacity required. It should also reveal which management decisions preserve the most service per unit of scarce resource.

9. Establish board ownership and decision rights

The board owns the resilience approach, critical-service perimeter, tolerance framework and acceptance of material gaps. Management owns implementation. Service owners are accountable for end-to-end outcomes even when resources sit in technology, operations, finance or a third party. Risk and compliance provide challenge within their mandates. Internal audit can provide independent assurance over governance, control design and evidence.

Decision rights should cover normal conditions and incidents. The company needs authority to declare a disruption, invoke alternative arrangements, stop unsafe operations, prioritise customers, move cash, engage advisers, notify authorities, communicate externally and accept temporary control changes. Succession must be practical. An alternate decision-maker needs information, access and legal authority before an event.

The ownership model should prevent three gaps. The first is technology ownership of a business-service outcome. Technology can restore systems but cannot decide customer priority, product safety or financial exposure alone. The second is committee ownership without one accountable executive. The third is sponsor intervention without clear company governance. An investor can set expectations, provide resources and monitor exposure, while the portfolio-company board and executives retain their legal and operating responsibilities.

Board reporting should focus on decisions. It should show which services remain outside tolerance, why, for how long, the interim control, the funded remediation, accountable owner, test date and residual exposure. It should identify changes in customers, contracts, sites, suppliers, technology, people and regulation that alter the service map. Detailed control metrics can sit beneath this view.

10. Build an incident control room that protects facts

Incident management converts plans into coordinated action. The control room needs a clear incident commander, service owner, technical and operational leads, finance, legal, risk, communications, people and supplier contacts. The structure can scale with severity. Roles should be named by function and supported by alternates, access, contact routes and delegated authority.

The common operating picture should separate facts, hypotheses, decisions and actions. Facts have time, source and confidence. Hypotheses are tested. Decisions record owner, rationale, assumptions and review time. Actions record owner, deadline and status. This discipline limits confusion when information changes and supports later review. It also gives finance a reliable event chronology for customer credits, lost output, remediation cost, insurance and disclosure.

Communications should be service-led and audience-specific. Customers need the affected outcome, available alternatives and next update. Employees need instructions and authority. Suppliers need priorities and interfaces. Lenders, insurers, authorities and investors may have contractual or legal notification requirements. Messages must remain accurate while investigation proceeds. The company should pre-approve channels, templates and decision thresholds without pre-writing conclusions.

The control room should track tolerance consumption. A service with a six-hour time boundary and a minimum capacity should show elapsed time, delivered capacity, backlog, data state, next recovery milestone and decision triggers. Leaders can then allocate scarce resources based on the approaching harm boundary. The incident closes only after service stability, control reconciliation, communications, evidence preservation and ownership of follow-up actions.

11. Govern third parties, concentration and exit

Outsourcing changes the delivery model and does not transfer the service outcome. The company remains exposed to the provider's people, technology, locations, subcontractors, utilities, security, financial condition and recovery priorities. Concentration can sit below the contracted provider through a common cloud, identity service, logistics node, data processor or equipment manufacturer.

Due diligence should connect the provider to the critical service and tolerance. The company needs to know the provider's service scope, capacity, locations, dependencies, incident process, tested recovery, data arrangements, subcontracting, customer prioritisation and exit capability. Certifications and assurance reports are inputs. They do not prove that the company's service will remain within its own tolerance.

Contract terms should support the operating model. Relevant areas include service and recovery commitments, notification, participation in exercises, evidence access, audit, information security, data location and return, subcontracting, insurance, financial distress, transition assistance and termination. Applicability and enforceability require qualified review. The operational team should understand which rights can be used during an incident and how quickly.

Exit planning should assume impaired cooperation. The company should know the data, configurations, credentials, expertise, equipment, licences and transition capacity required to move or insource the service. A viable exit may require parallel capacity, escrow, modular architecture, inventory, step-in rights or a tested manual fallback. The board should see the cost of

substitutability against the service harm and concentration risk.

Table 4. Third-party resilience and exit evidence

decision area	evidence before reliance	monitoring evidence	exit evidence	common weakness
service and priority	end-to-end scope and customer priority	capacity, incidents and service trend	replacement service design	contract describes system, not outcome
recovery	tested recovery method and objectives	test results and remediation	transition and fallback test	provider result excludes client restart
concentration	locations, platforms and subcontractors	material dependency change	alternate route or provider	apparent suppliers share one dependency
data and access	ownership, integrity, backup and access	restoration and access evidence	export, reconciliation and deletion	data return exists only as a clause
incident response	notice, roles and communication	incident timeliness and lessons	contact continuity during termination	escalation depends on one relationship
financial and legal	viability, insurance and enforceable terms	adverse change indicators	distress and insolvency pathway	exit assumes normal cooperation

Each dependency should connect contractual rights to tested operational capability.

12. Integrate technology, cyber and data recovery

Technology resilience should be designed from the service tolerance. Architecture, identity, network, applications, databases, integrations, end-user devices, monitoring, deployment tools and support processes all contribute. Redundancy at one layer can be defeated by a shared dependency elsewhere. The map should show whether alternatives are independent enough for the scenario being tested.

Recovery objectives need transaction integrity. A near-zero recovery point does not establish that records are complete, ordered and reconciled. The service may require replay, duplicate detection, manual approval, control totals and downstream confirmation. The recovery plan should state how data is validated before the service returns, who can authorise exceptions and how backlog is processed without creating a second incident.

Cyber scenarios add uncertainty and adversarial behaviour. The organisation may need to isolate systems, distrust credentials, preserve evidence and rebuild from known-good states. Restoring compromised infrastructure quickly can extend the event. The response plan should connect security containment to service continuity and define minimum safe capability. CISA encourages corporate leaders to exercise response plans and focus continuity investment on systems supporting critical business functions [24].

Change management is part of resilience. Major incidents often follow releases, migrations, capacity changes, certificate expiry, unsupported components or control drift. The board should see resilience debt alongside technology debt: single points of failure, untested recovery, manual reconciliation, end-of-support exposure and resource constraints. Remediation should be tied to the service and tolerance rather than presented as an undifferentiated infrastructure programme.

13. Design people and facility resilience for actual operating conditions

People resilience is more than a contact list. The company needs sufficient trained staff, alternates, authority, access, physical availability and decision support to deliver the minimum service. Specialist knowledge should be documented and exercised. Key-person exposure should include executives, technical experts, operators, commercial decision-makers and external advisers. Simultaneous absence can matter more than the loss of any one role.

Alternative work arrangements should be tested under realistic constraints. Remote work depends on identity, devices, bandwidth, communications, home conditions and access to records. Alternate sites depend on transport, security, equipment, capacity and local infrastructure. A declared recovery seat is not useful if it lacks the tools, data or authorisation required for the service. Tests should include the actual team and workload where safe.

Facility resilience begins with the service's physical requirements. Power, water, cooling, fuel, ventilation, controlled environments, safety systems, loading access, communications and specialist equipment can be critical. Backup capacity needs runtime, replenishment, maintenance, testing and operating staff. A generator test without a sustained load or fuel-supply scenario provides limited evidence.

The workforce is also an affected stakeholder. Incident decisions should preserve safety, lawful working arrangements, fatigue management and clear communication. Long disruptions create shifts, travel, accommodation and wellbeing demands. Succession and delegation should avoid dependence on exhausted leaders. The recovery plan should define the point at which reduced service must be accepted because continuing would create unsafe or uncontrolled operations.

14. Map supply chains, utilities and infrastructure as operating systems

Supply resilience requires more than counting suppliers. The relevant question is whether the service can obtain the qualified materials, components, logistics, equipment and support it needs within tolerance. Qualification time, regulatory approval, tooling, intellectual property, minimum order quantities, transport routes, customs, storage conditions and working capital can make an apparent alternative unusable.

The company should map nodes and flows. A tier-one supplier may depend on a single tier-two manufacturer. Several distributors may hold the same source. Geographic diversification can still share a port, data network, energy source or commodity. The map should record substitutability, time to qualify, buffer, visibility and the commercial conditions required to activate an alternative.

Utilities and critical infrastructure need explicit service assumptions. CISA's infrastructure-dependency guidance asks how long operations can continue without a service and highlights secondary suppliers, interconnections, backup power and continuity plans [25]. A portfolio company can apply the same questions to electricity, water, telecommunications, transport and waste. It should also test whether priority arrangements or backup supplies remain available during a region-wide event.

Buffers should be economically governed. Inventory, spare equipment, redundant capacity and alternative suppliers consume cash. Their value depends on the loss avoided, activation probability, shelf life, financing, maintenance and opportunity cost. The resilience case should show the service and harm protected, scenario coverage, cost, residual risk and review point. A

buffer without ownership and rotation can become false assurance.

15. Connect resilience to liquidity, financing and insurance

A disruption affects several financial statements and time horizons. Revenue may pause, backlog may rise, refunds and credits may be issued, variable costs may fall or increase, overtime and expedite costs may emerge, inventory may be lost and working capital may absorb cash. Recovery can require capital expenditure, specialists, duplicate capacity and customer remediation. The model should separate event effects, response spending, recovery investment and recurring resilience cost.

The short-term view belongs in the cash forecast. It should include lost or delayed receipts, customer concessions, supplier prepayments, emergency procurement, payroll, tax, debt service, insurance timing and restricted cash. The base, downside and severe scenarios should show covenant headroom and liquidity triggers. A service tolerance can be breached before a financial covenant, and a liquidity constraint can limit the actions required to remain within service tolerance.

Insurance should be modelled through coverage, exclusions, limits, deductibles, waiting periods, evidence and cash timing. A policy does not replace capability, and an expected claim should not be treated as available cash without support. The incident record, decision log and loss evidence can materially affect the company's ability to prepare and support a claim. Qualified insurance and legal advice remains necessary.

Financing documents may contain information, consent, representation, insurance, material-adverse-effect or covenant provisions relevant to disruption. The board should understand who assesses these terms and when lenders are engaged. A credible resilience programme can improve the quality of lender dialogue because the company can show service boundaries, tested capability, scenarios, liquidity and remediation rather than a generic assurance.

16. Integrate resilience into change, growth and the operating model

Growth can consume resilience capacity. Customer volume, new geographies, added products and tighter service promises increase demand on people, systems, facilities and suppliers. A platform may remain technically available while queues exceed tolerance. Capacity planning should therefore include disrupted-mode capacity, backlog clearance and resource competition between critical services.

Change can temporarily weaken controls. System migrations, enterprise-resource-planning releases, site moves, outsourcing, leadership turnover and organisational redesign change dependencies and authority. Major changes should include a resilience impact assessment before approval. The assessment states affected services, tolerances, transition risks, rollback, dual-running, data reconciliation, supplier readiness and the tests required before the old capability is removed.

The target operating model should make resilience part of business-as-usual decisions. Product approval asks whether a service or tolerance changes. Procurement asks how the supplier supports critical services. Architecture review asks whether a new common-mode dependency is created. Capital allocation compares remediation with service harm and value. Workforce planning covers alternates and critical skills. Performance reviews include test and remediation ownership.

The company should avoid creating a parallel resilience bureaucracy. Service owners use existing governance, risk, technology, operations, finance and audit mechanisms with a shared service-and-tolerance record. A small central team can maintain methodology, challenge evidence, coordinate tests and aggregate board reporting. Accountability stays with the executives who control the service.

17. Use resilience throughout the transaction lifecycle

Pre-acquisition diligence should test whether the target can deliver its critical services through disruption and whether the investment case depends on fragile capability. The review can identify critical services, material incidents, tolerance assumptions, dependency concentration, technology recovery, facility continuity, supplier exposure, insurance, contract rights, test evidence and unresolved remediation. Findings should connect to valuation, warranties, indemnities, conditions, financing and the value-creation plan as appropriate.

The first one hundred days should establish the board perimeter, decision rights and highest-risk evidence gaps. Quick actions may include confirming crisis contacts, validating backups, protecting credentials, testing a manual workaround, securing key suppliers, extending fuel autonomy or funding a critical spare. Structural actions may require architecture, supplier qualification, facility investment, process redesign or recruitment and should retain staged gates.

Add-on acquisitions create immediate integration risk. The platform and target may use different incident thresholds, service definitions, identity, data, suppliers and communication routes. Integration plans should preserve safe standalone capability until the combined service is mapped and tested. Synergy actions that remove sites, systems, stock, suppliers or people should state the resilience effect before implementation.

Exit preparation should convert the programme into diligence-ready evidence. The data room can include the service catalogue, board-approved tolerances, dependency maps, incident history, test results, remediation closure, third-party evidence, insurance and financial scenarios. The company should present limitations and open actions. A buyer can then assess capability and residual exposure through evidence rather than management assertion.

18. Apply jurisdictional overlays without creating false equivalence

Operational-resilience rules differ by sector, entity and jurisdiction. The FCA framework applies to specified financial firms and requires important business services, impact tolerances, mapping and testing [1][2]. DORA establishes digital operational-resilience obligations for covered European financial entities and addresses ICT risk, incidents, testing and third parties [5][6][7]. APRA CPS 230 applies to APRA-regulated entities and covers operational risk, critical operations, tolerance levels, business continuity and material service providers [8][9].

The Central Bank of the UAE issued an in-force operational-resilience regulation for licensed financial institutions in 2026, including strategy, threat identification and board escalation [11][12]. The Saudi Central Bank's business-continuity framework uses business-impact analysis, recovery objectives, plans, testing and independent review for member organisations [13]. The Reserve Bank of India's 2023 directions address IT governance, operational resilience, business continuity, disaster recovery, testing and third-party arrangements for covered regulated entities [14][15].

The Monetary Authority of Singapore's revised business-continuity guidance emphasises critical business services and end-to-end dependencies for financial institutions [16]. Australian APRA CPS 230 was updated for commencement on 1 July 2026 and requires covered entities to maintain critical operations within tolerance through severe disruptions [8][9][10]. Public-company and listing frameworks may add disclosure or controls obligations, such as the SEC's cyber rules and the UK Corporate Governance Code [30][31][32].

A portfolio group should maintain an applicability register. It identifies the legal entity, service, sector, jurisdiction, rule owner, obligation, evidence, reporting and change date. Shared methodology can improve consistency, while local requirements and qualified advice control. Regulatory concepts should not be presented as universal legal duties.

Table 5. Illustrative jurisdiction and framework overlay

source	stated scope	resilience concepts used	portfolio-company adaptation	applicability control
FCA, United Kingdom	specified financial firms	important services, impact tolerance, mapping, testing	service-led board framework	confirm entity and handbook scope
European Union DORA	covered financial entities	ICT risk, incidents, testing, third parties	digital dependency and evidence model	confirm entity, service and implementation acts
APRA CPS 230, Australia	APRA-regulated entities	critical operations, tolerances, BCP, providers	multi-dimensional tolerance card	confirm prudential status and transition
CBUAE, United Arab Emirates	licensed financial institutions	critical operations, response, recovery, learning	board escalation and capability gap	confirm licence and rulebook provisions
SAMA, Saudi Arabia	member organisations	BIA, MAO, RTO, RPO, BCP and tests	operating recovery hierarchy	confirm membership and current rule
RBI, India	specified regulated entities	IT governance, BCP, DR, RTO, RPO, provider testing	data reconciliation and technology recovery	confirm entity classification and direction
MAS, Singapore	financial institutions	critical services and end-to-end dependencies	service-centric continuity design	confirm institution and guideline status
ISO 22301 and global principles	generic or standard-setting reference	continuity management and board risk oversight	proportionate management system	distinguish voluntary framework from law

The table summarises reference architecture; responsible advisers must determine applicability and current requirements.

19. Test the framework on a hypothetical portfolio company

Consider a hypothetical multi-country provider of equipment maintenance, remote monitoring and emergency field repair. It has 1,200 employees, four operating sites, a central scheduling platform, specialist spares, mobile technicians, outsourced identity and cloud services, and major customers in utilities and industrial production. All values, timings and scores in this section are illustrative assumptions.

The board identifies three critical services: accept and triage an emergency request, dispatch a qualified technician with required equipment, and restore customer equipment to a safe operating

state. The first service has a one-hour time tolerance and minimum ninety-five per cent triage capacity. Dispatch has a four-hour tolerance for priority cases. Safe restoration varies by contract and asset and is governed through customer-specific escalation.

Mapping reveals five vulnerabilities. Identity is shared across primary and recovery environments. Two spare-parts distributors use one manufacturer. Dispatch supervisors hold undocumented knowledge. The alternate site lacks sufficient radio capacity. The customer-notification process depends on the same platform used for scheduling. Current documents describe a two-hour system recovery objective, but no end-to-end test has demonstrated dispatch within four hours.

Management designs a staged programme. It creates an emergency identity route, qualifies an alternate component, documents supervisor decisions, expands alternate communications and separates customer notifications from scheduling. A tabletop tests decisions and communications. Component tests validate identity and radio. An integrated simulation removes the primary platform, one supervisor and the common distributor during peak demand.

The hypothetical result keeps triage within tolerance, restores alternate identity in forty minutes and reaches eighty per cent dispatch capacity after three hours. Twelve per cent of cases exceed the four-hour dispatch tolerance because spares cannot be allocated accurately from the alternate process. The test therefore fails the complete service objective and identifies a specific inventory-data and allocation decision. Management funds a controlled offline spare register, barcode reconciliation and a retest.

Figure 4. Hypothetical service test against approved tolerances

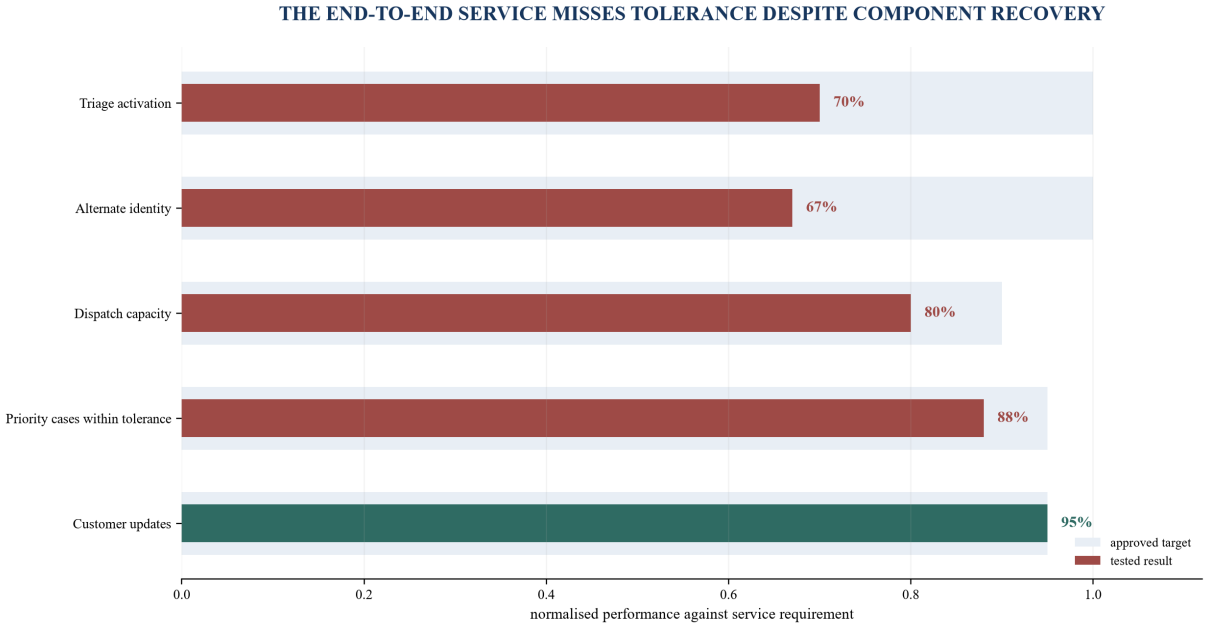

Illustrative results show why component recovery and end-to-end service performance must be reported separately.

20. Translate resilience into an economic and equity bridge

Resilience can protect value through several channels. It can reduce expected lost contribution, customer remediation, emergency procurement, regulatory cost, working-capital absorption and recovery expenditure. It can preserve customer retention, licence continuity, lender confidence and transaction optionality. It can also consume recurring operating cost and capital. The economic case should show gross loss exposure, control effect, full cost, residual uncertainty and timing.

The model should avoid multiplying one event across several labels. Lost revenue, lost contribution, churn and valuation impact can overlap. Insurance may offset part of the same loss after deductible, exclusion, limit and delay. A resilience investment can support several services, creating allocation questions. The model should keep operational results, accounting effects, cash, enterprise value and equity value distinct.

Scenario-weighted estimates can support prioritisation but should not create false precision. The board can compare investments through ranges: loss severity under defined scenarios, capability improvement, service tolerance achieved, cost, implementation risk and strategic option value. Low-frequency, high-severity exposures may still require action because the company cannot absorb the event or because law, safety or contract sets a boundary.

In a transaction, resilience evidence can affect diligence, debt capacity, normalised cost, capital-expenditure expectations, warranties, insurance and the buyer's view of execution risk. It does not mechanically create a valuation multiple. The bridge should show how a tested capability changes the relevant risk or cash-flow scenario and which assumptions remain for the buyer, lender or investor to assess.

Figure 5. Resilience investment to equity-value bridge

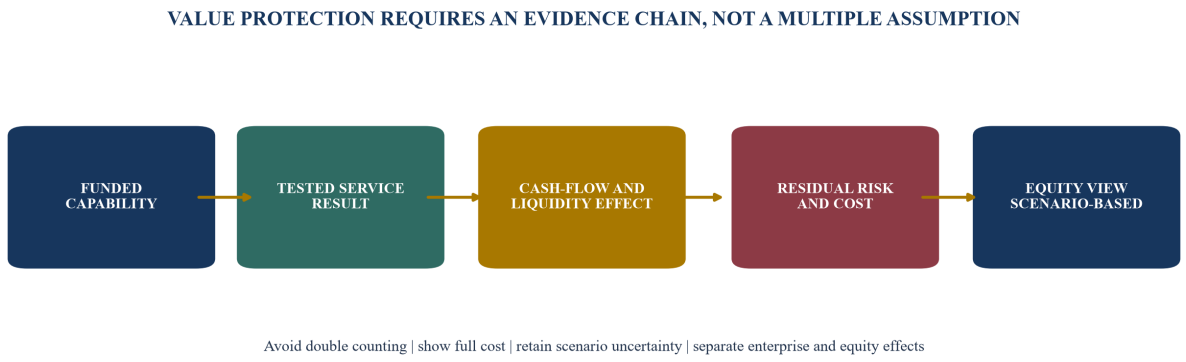

The bridge separates service capability, financial effects, residual risk and transaction judgement.

21. Use a board scorecard that drives decisions

The scorecard should begin with service outcomes. For each critical service it shows the approved tolerance, latest tested capability, evidence date, confidence, incidents, near misses, open vulnerabilities, interim controls, funded remediation and next decision. It should include tolerance consumption during material incidents and the status of customer, regulator, lender, insurer and supplier actions where relevant.

Supporting indicators should be tied to service causality. Examples include backup restoration, identity failover, alternate-site capacity, qualified supplier coverage, key-role alternates, generator autonomy, data reconciliation, call-tree completion and time to customer notification. A metric belongs on the board pack only when management can explain how it changes service capability or harm.

Portfolio reporting can use a small number of comparable fields: services outside tolerance, high-confidence test coverage, overdue critical remediation, concentrated dependencies, material incidents and investment required. The sponsor should preserve each company's service detail and avoid using a single maturity score as a substitute for risk. A business can score well overall while one critical service remains exposed.

The board pack should include decisions requested. Typical decisions are tolerance approval, funding, risk acceptance, supplier concentration, system replacement, test scope, exit investment or change sequencing. Papers should show alternatives, cost, capability effect, implementation risk and residual exposure. The decision log becomes part of the evidence chain for future incidents and transaction diligence.

Table 6. Board operational-resilience scorecard

board field	current view	trend or threshold	required evidence	decision trigger
critical services	approved perimeter and owner	material business change	service record and rationale	add, remove or redefine service
tolerance status	within, gap or unknown	any service outside boundary	approved tolerance and capability	accept, control or fund gap
test coverage	service, scenario and date	overdue or insufficient realism	test plan, measured result, findings	expand scope or accelerate retest
incidents and near misses	harm, duration and service effect	repeated cause or rising severity	chronology, decision and loss record	structural remediation
dependencies	concentration and substitutability	new common-mode exposure	current map and provider evidence	diversify, buffer or redesign
remediation	owner, cost, milestone and residual risk	overdue critical action	closure evidence and retest criteria	intervene or revise plan
liquidity and value	downside cash and ownership effect	covenant or minimum-cash trigger	cash forecast and scenario bridge	fund reserve or change action
assurance	management, control and independent view	evidence confidence deteriorates	attestations, audit and external work	commission deeper assurance

The scorecard prioritises service outcomes, evidence confidence and accountable decisions.

22. Implement the framework through a one-hundred-day sequence

Days 1 to 15 establish governance. The board confirms the executive sponsor, service owners, definitions, escalation, current incidents and immediate safety or continuity concerns. Management identifies a preliminary critical-service perimeter and gathers existing plans, maps, contracts, tests, audits and insurance evidence. Urgent gaps receive interim controls rather than waiting for the full programme.

Days 16 to 35 define services and tolerances. Cross-functional teams validate customer outcomes, harm, obligations, time, capacity, backlog and data boundaries. Finance builds the first disruption cash scenarios. Legal, regulatory, safety, cyber, engineering and insurance specialists assess issues within their mandates. The board approves the initial perimeter, tolerance method and unresolved questions.

Days 36 to 60 map dependencies and test evidence. Service owners document end-to-end delivery, shared resources, third parties, alternatives and recovery objectives. The programme reconciles plans to actual operating data and grades evidence confidence. It identifies common-mode dependencies and capability gaps. Procurement and technology address contractual and architecture evidence; people and facilities validate operational readiness.

Days 61 to 80 run controlled exercises. The sequence starts with walkthroughs and tabletops, then component or failover tests where risk permits. Tests measure service results, decisions, data, communications and resource consumption. Findings are prioritised by tolerance and harm. Each action receives an owner, funding, target date, interim control and retest criterion.

Days 81 to 100 close the first governance cycle. The board reviews service status, test evidence, residual risk, investment and the next twelve-month test plan. The company integrates resilience into change, procurement, capital allocation, incident management and transaction readiness. Completion means that the governance and evidence cycle is operating. It does not mean every risk has been removed.

Figure 6. One-hundred-day operational-resilience programme

THE FIRST CYCLE ESTABLISHES TESTED CAPABILITY AND AN OWNED REMEDIATION AGENDA

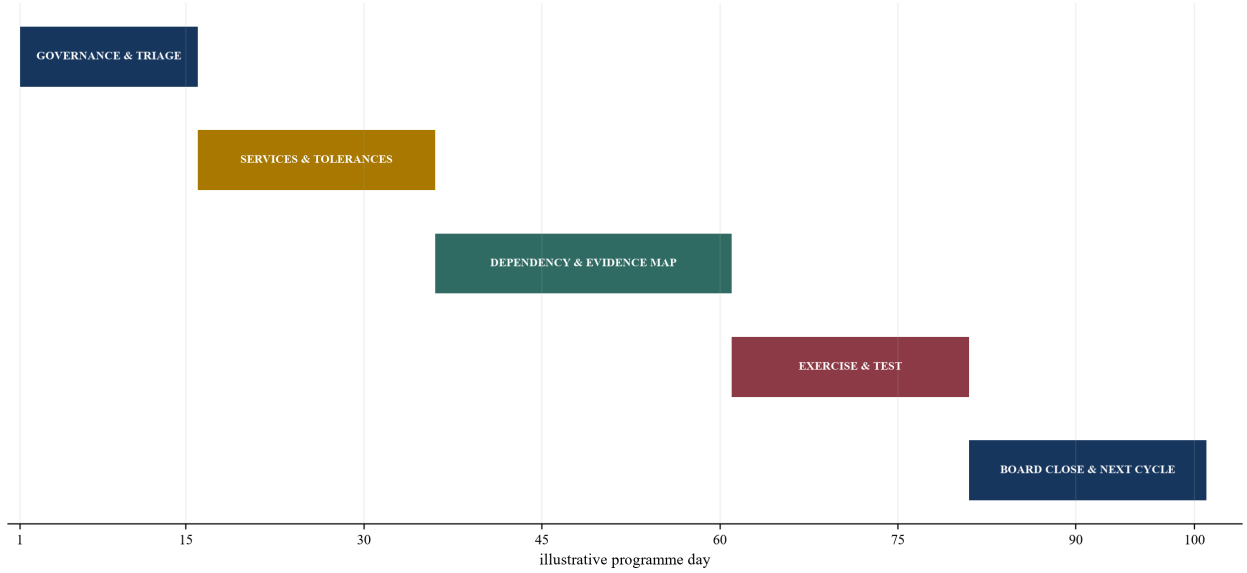

The sequence moves from governance and service definition to tested evidence, funded remediation and recurring board oversight.

23. Recognise the limits and govern the next cycle

The framework does not determine which services are legally critical, what harm is acceptable, which tolerance is correct or which control is sufficient for a particular company. Those decisions depend on customers, safety, contracts, licences, law, regulation, technology, engineering, financial capacity, insurance and professional judgement. Sector-specific sources cited in this paper do not establish requirements for entities outside their scope.

Maps and tests remain incomplete models. A map can omit an informal dependency. A successful test can exclude a condition that later matters. Participants can behave differently in a live incident. A provider can prioritise customers differently during a systemic event. A backup can restore while data integrity fails. Management should state the scope, conditions, exclusions and evidence date behind every capability claim.

Value estimates also remain uncertain. Event frequency, loss severity, customer behaviour, insurance, recovery cost and transaction treatment cannot be known with precision. Resilience can protect service and strategic options without producing an identifiable accounting gain. The board should use scenarios and ranges, avoid double counting and update assumptions after tests, incidents and operating changes.

Operational resilience is therefore a recurring governance cycle. The company identifies the services that matter, sets harm-based tolerances, maps delivery, tests capability, funds gaps, responds to incidents and learns. Under ownership, this cycle connects customer protection and operating control to cash, financing and transaction readiness. Its strongest claim is evidential: the board can see what must continue, what has been demonstrated, what remains exposed and which decision is required next.

References

[1] Financial Conduct Authority, Operational Resilience, https://www.fca.org.uk/firms/operational-resilience

[2] Financial Conduct Authority, PS21/3 Building Operational Resilience, https://www.fca.org.uk/publications/policy-statements/ps21-3-building-operational-resilience

[3] Financial Conduct Authority, Operational Resilience: Insights and Observations One Year On, https://www.fca.org.uk/publications/good-and-poor-practice/operational-resilience-insights-observations-one-year

[4] Financial Conduct Authority, Operational Resilience: Insights and Observations for Firms, https://www.fca.org.uk/firms/operational-resilience/insights-observations

[5] European Union, Regulation (EU) 2022/2554 on Digital Operational Resilience for the Financial Sector, https://eur-lex.europa.eu/eli/reg/2022/2554/oj

[6] European Commission, Cyber Resilience in Financial Services, https://finance.ec.europa.eu/digital-finance/cyber-resilience_en

[7] European Commission, Digital Operational Resilience Regulation Implementing and Delegated Acts, https://finance.ec.europa.eu/regulation-and-supervision/financial-services-legislation/implementing-and-delegated-acts/digital-operational-resilience-regulation_en

[8] Australian Prudential Regulation Authority, CPS 230 Operational Risk Management, https://www.apra.gov.au/standards/cps-230

[9] Australian Prudential Regulation Authority, Prudential Standard CPS 230 Operational Risk Management, https://www.apra.gov.au/sites/default/files/2023-07/Prudential%20Standard%20CPS%20230%20Operational%20Risk%20Management.pdf

[10] Australian Prudential Regulation Authority, Operational Risk Management Consultation and 2026 Materials, https://www.apra.gov.au/consultations/operational-risk-management

[11] Central Bank of the United Arab Emirates, Article 3 Operational Resilience, https://rulebook.centralbank.ae/en/rulebook/article-3-operational-resilience

[12] Central Bank of the United Arab Emirates, Operational Resilience Rulebook Section, https://rulebook.centralbank.ae/en/entiresection/7278

[13] Saudi Central Bank, Business Continuity Management Framework, https://rulebook.sama.gov.sa/en/business-continuity-management-framework-1

[14] Reserve Bank of India, Information Technology Governance Risk Controls and Assurance Practices Directions 2023, https://www.rbi.org.in/scripts/NotificationUser.aspx?Id=12562

[15] Reserve Bank of India, Master Direction on Information Technology Governance Risk Controls and Assurance Practices, https://systemhealth.rbi.org.in/Scripts/BS_ViewMasDirections.aspx_id%3D12562%283%29.html

[16] Monetary Authority of Singapore, Response to Feedback on Revised Business Continuity Management Guidelines, https://www.mas.gov.sg/-/media/mas/regulations-and-financial-stability/regulatory-and-supervisory-framework/risk-management/bcm-guidelines/response-to-feedback-received_bcmg-second-consultation-paper-2.pdf

[17] Bank for International Settlements Basel Committee on Banking Supervision, Principles for Operational Resilience, https://www.bis.org/bcbs/publ/d516.htm

[18] Bank for International Settlements Financial Stability Institute, Principles for Operational Resilience Executive Summary, https://www.bis.org/fsi/fsisummaries/op_resilience.htm

[19] National Institute of Standards and Technology, SP 800-34 Revision 1 Contingency Planning Guide for Federal Information Systems, https://csrc.nist.gov/pubs/sp/800/34/r1/upd1/final

[20] National Institute of Standards and Technology, SP 800-84 Guide to Test Training and Exercise Programs for IT Plans and Capabilities, https://csrc.nist.gov/pubs/sp/800/84/final

[21] National Institute of Standards and Technology, Cybersecurity Framework 2.0, https://www.nist.gov/cyberframework

[22] Cybersecurity and Infrastructure Security Agency, Cyber Resilience Review, https://www.cisa.gov/resources-tools/services/cyber-resilience-review-crr

[23] Cybersecurity and Infrastructure Security Agency, Cybersecurity Scenarios, https://www.cisa.gov/resources-tools/resources/cybersecurity-scenarios

[24] Cybersecurity and Infrastructure Security Agency, Shields Up Guidance for Corporate Leaders and Chief Executives, https://www.cisa.gov/shields-guidance-corporate-leaders-and-ceos

[25] Cybersecurity and Infrastructure Security Agency, Infrastructure Dependency Primer Implement, https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/resilience-services/infrastructure-dependency-primer/implement

[26] United States Federal Emergency Management Agency, Continuity Guidance Circular 2024 Update, https://www.fema.gov/sites/default/files/documents/fema_continuity-guidance-circular_082024.pdf

[27] International Organization for Standardization, ISO 22301:2019 Business Continuity Management Systems, https://www.iso.org/standard/75106.html

[28] Organisation for Economic Co-operation and Development, G20 OECD Principles of Corporate Governance 2023, https://www.oecd.org/en/publications/2023/09/g20-oecd-principles-of-corporate-governance-2023_60836fcb.html

[29] Organisation for Economic Co-operation and Development, The Responsibilities of the Board, https://www.oecd.org/en/publications/g20-oecd-principles-of-corporate-governance-2023_ed750b30-en/full-report/com

ponent-8.html

[30] Financial Reporting Council, UK Corporate Governance Code 2024, https://www.frc.org.uk/library/standards-codes-policy/corporate-governance/uk-corporate-governance-code/

[31] United States Securities and Exchange Commission, Cybersecurity Risk Management Strategy Governance and Incident Disclosure Final Rule, https://www.sec.gov/rule-release/33-11216

[32] United States Securities and Exchange Commission, Cybersecurity Risk Management Strategy Governance and Incident Disclosure Small Entity Compliance Guide, https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/cybersecurity-risk-management-strategy-governance-incident-disclosure

[33] United Kingdom Cabinet Office, The UK Government Resilience Framework, https://www.gov.uk/government/publications/the-uk-government-resilience-framework/the-uk-government-resilience-framework-html

About the Author

Authored by Chennakeshav Adya, Independent Researcher