Portfolio Cyber Resilience: A Sponsor Control Framework across Critical Services

A Board and Sponsor Operating Model for Critical-Service Continuity, Controlled Cyber Risk and Exit-Ready Evidence

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Cyber risk becomes a portfolio value question when disruption can interrupt critical services, impair customer trust, restrict cash generation, trigger notification duties or complicate financing and exit. A sponsor can set common expectations and mobilise shared capability, while each portfolio-company board remains responsible for its own operations, law, customers and risk. The practical challenge is to create enough consistency for control without replacing local accountability or reducing diverse businesses to a single maturity score.

This paper develops a sponsor control framework built around critical services and their dependencies. It connects acquisition diligence, board governance, identity, information technology, operational technology, cloud, suppliers, incident command, recovery, insurance, investment decisions, portfolio intelligence and exit evidence. Six original figures and six tables translate the framework into operating tools. The research draws on current official and authoritative material from the United Arab Emirates, Saudi Arabia, European Union, United Kingdom, United States, Singapore, Australia and Canada, together with international cybersecurity, operational-resilience and continuity standards.

All scores, thresholds, financial amounts, time horizons, allocations and worked cases are hypothetical management assumptions created to demonstrate the method. They do not describe a client, market benchmark, forecast or guaranteed outcome. Actual requirements and results depend on the company's services, architecture, threat environment, regulation, contracts, people, suppliers, insurance and resources. The paper provides strategic and operational research. It does not replace legal, regulatory, cybersecurity, engineering, accounting, insurance, valuation or investment advice.

JEL Classification: G24, G32, G34, L86, M15, M14

Keywords: private equity, portfolio cybersecurity, critical services, cyber resilience, operational resilience, incident response, supply chain, OT security, cloud risk, exit readiness

1. Define portfolio cyber resilience through critical services

Portfolio cyber resilience is the governed ability of each company to keep its most important services within approved disruption tolerances when cyber events affect people, facilities, technology, data or suppliers. It combines prevention, detection, response, recovery and adaptation. The unit of analysis is the service delivered to a customer, market or community, rather than the number of tools installed or policies issued.

A sponsor framework should begin with a limited set of questions. Which services can create material harm if interrupted or corrupted? What assets, identities, data, sites, operational technologies, cloud platforms and external providers support them? How much disruption can the company tolerate? Which decisions require board authority? What evidence shows that recovery

works? The answers create a control perimeter that management can operate and the sponsor can challenge.

NIST Cybersecurity Framework 2.0 organises outcomes across Govern, Identify, Protect, Detect, Respond and Recover [1]. Its addition of Govern makes accountability, policy, supply-chain oversight and risk appetite explicit. The United Kingdom's Cyber Assessment Framework connects governance, protection, detection and impact reduction to essential functions [6]. CISA's Cross-Sector Cybersecurity Performance Goals provide a prioritised baseline for critical infrastructure [2]. Together, these sources support a service-led framework without prescribing one technology stack.

Portfolio coordination should preserve legal and operating accountability. The sponsor can approve minimum expectations, aggregate material risk, provide expert resources, compare evidence and coordinate mutual aid. Each portfolio-company board should approve service tolerances, material exceptions, investment, incident authority and disclosure decisions. Local management owns implementation and performance. This division enables scale while keeping decisions close to the business facts.

The framework should distinguish resilience from compliance. Regulatory obligations form part of the minimum control environment, and contractual duties may add further requirements. Resilience asks whether the company can sustain or restore the service when controls fail, an attacker changes tactics or a shared supplier is unavailable. A well-governed portfolio therefore reviews control operation, dependency concentration, exercised recovery and management readiness together.

Figure 1. Sponsor cyber-resilience control architecture

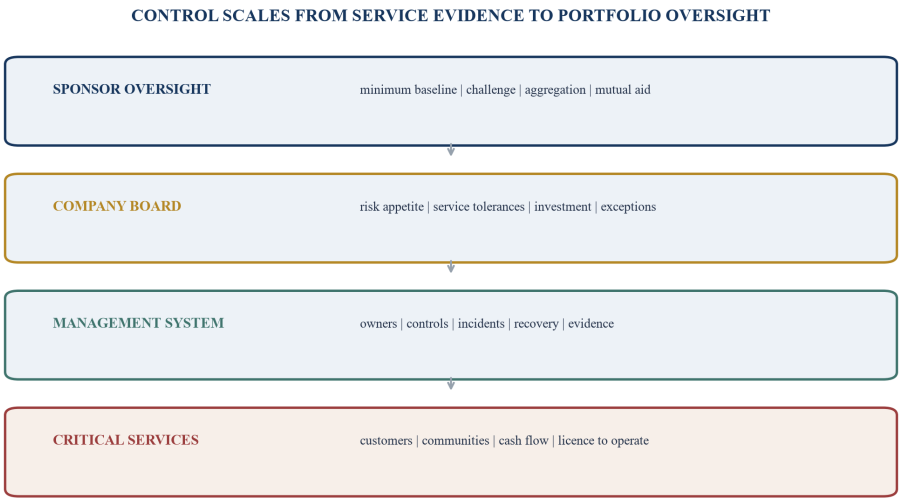

The allocation of responsibilities is illustrative and requires company-specific board approval.

2. Map critical services, consequences and tolerances

A service map translates a complex technology estate into a manageable business view. Management should name the service in customer language, identify the people or institutions that depend on it, describe the harm caused by failure and state the point at which disruption becomes intolerable. The map should cover loss of availability, integrity, confidentiality and control because a service can remain online while producing unsafe, fraudulent or unreliable outcomes.

Impact tolerance is a board decision informed by customer commitments, safety, regulation, liquidity and realistic response capability. A payment service may be measured in minutes, a hospital workflow in clinical consequences, an industrial process in safe-state requirements and a data service in recoverable transactions. The tolerance should state the maximum acceptable impact under a severe but plausible scenario. Recovery objectives, system service levels and supplier contracts then support the tolerance; they do not replace it.

The service owner should map end-to-end dependencies: business process, critical roles, privileged identities, applications, infrastructure, cloud regions, network routes, operational technology, data, facilities, utilities and third parties. The map should show shared components. A single identity provider, remote-access gateway, domain-name service, managed-service provider or power source can create portfolio concentration even when companies operate in different industries.

APRA's CPS 230 requires regulated entities to identify critical operations, set tolerance levels and manage service-provider risks [18]. The UK's operational-resilience approach similarly focuses firms on important business services, impact tolerances and mapping [24]. The regulatory scope differs, while the operating logic is useful across critical-service businesses: define the outcome, understand dependencies, test the limit and remediate material gaps.

Management should keep the register decision-sized. A list of hundreds of processes reduces attention. The first pass can identify the small number of services whose failure can materially affect life, safety, essential supply, customers, cash, regulation or enterprise value. Supporting services can then be mapped where they create a material dependency. The register should be refreshed after acquisitions, architecture changes, outsourcing, major incidents and changes to customer commitments.

Table 1. Critical-service register and tolerance design

Critical service	Material consequence	Minimum dependency evidence	Tolerance expression	Recovery evidence	Board question
customer payments	liquidity, customer harm and regulatory breach	payment rails, identity, ledger, fraud controls and banks	maximum unavailable minutes and unreconciled transactions	failover, reconciliation and controlled backlog test	can the service recover without creating hidden financial loss?
clinical workflow	patient safety and care interruption	devices, applications, data, sites, clinicians and suppliers	defined safe operating period and manual capacity	clinical downtime exercise and record reconciliation	which care pathway fails first and who can authorise alternatives?
energy dispatch	safety, supply and contractual exposure	control systems, telemetry, communications, operators and grid	safe-state conditions and restoration sequence	isolated recovery and control-room exercise	can operations remain safe when enterprise IT is unavailable?
logistics fulfilment	customer loss, penalties and working-capital strain	orders, routing, warehouses, carriers, customs and payments	backlog volume, delay and priority rules	alternate-routing and backlog-clearance test	how long before disruption becomes a cash and customer event?
regulated data access	confidentiality, integrity and legal duties	identity, keys, repositories, processors and audit logs	maximum data loss and access interruption	immutable backup and access-restoration test	can management prove which records were affected?

The entries are hypothetical. Actual tolerances require company evidence, testing and board approval.

3. Set a portfolio minimum baseline and preserve local overlays

A portfolio baseline should contain a small number of control outcomes that materially reduce common paths to disruption. It should avoid a tool catalogue. A useful baseline covers governance, asset and service visibility, identity, secure configuration, vulnerability and exposure management, logging and detection, backups, incident response, supplier access, recovery testing and evidence retention. Each outcome needs an owner, operating frequency, evidence standard and exception route.

CISA describes its cross-sector goals as a voluntary baseline with high-impact actions for critical infrastructure [2]. Australia's Essential Eight provides a prioritised set of mitigation strategies and maturity guidance [19]. Canada's baseline controls provide practical direction for smaller organisations [23]. These sources can inform the common floor. NIST CSF 2.0 profiles and tiers can help each company state current and target outcomes without presenting certification or numerical maturity as a substitute for risk judgement [1].

Local overlays should add what the business actually requires. Operational-technology companies need safety, segmentation, engineering access and safe-state controls. Financial services may require transaction integrity, third-party oversight and regulatory reporting. Healthcare needs clinical continuity and medical-device governance. Data-centre and cloud-dependent services need capacity, region, control-plane and tenant-separation evidence. Consumer platforms need fraud, privacy and customer communication. Geographic overlays incorporate applicable law, regulator expectations, notification timeframes and government reporting.

The baseline should support three evidence states. Designed means the policy, architecture or procedure exists and has an accountable owner. Operating means current evidence shows the

control performed as intended. Resilient means the company has tested failure, recovery and decision-making under a relevant scenario. Reporting these states separately prevents a policy document from being mistaken for exercised capability.

Exceptions should be visible and time-bound. An exception record should describe the affected service, exposure, reason, compensating control, accountable executive, target treatment, funding and expiry. A company may accept an exposure when the business case, alternatives and residual consequence are understood. Portfolio committees should focus on material exceptions and repeated delays rather than demanding uniform closure dates without regard to operating reality.

Table 2. Portfolio minimum-control baseline

Outcome	Minimum operating evidence	Service-level test	Escalation trigger	Local overlay examples
accountable governance	board-approved authority, risk appetite and exception log	material decisions reached within the crisis timetable	unclear authority or expired exception	regulator, licence and disclosure duties
known attack surface	current critical-service, asset, identity and external-exposure records	unknown dependency discovered during exercise	unsupported or ownerless critical asset	OT, medical device, cloud and data classifications
controlled identity	strong authentication, privileged-access review and joiner-mover-leaver evidence	emergency access and revocation test	shared privileged account or dormant access	break-glass, contractor and machine identity
protected recovery	segregated, monitored and tested backups with restoration priorities	restore service and reconcile data within tolerance	untested backup or shared control-plane dependency	safe-state, immutable copy and alternate site
detectable compromise	relevant logs, triage ownership and response coverage	detect and escalate scenario indicators	critical log gap or unsupported monitoring	fraud, safety, privacy and managed detection
governed suppliers	tiering, access control, contract rights and continuity evidence	supplier outage or compromise exercise	concentration without credible alternative	fourth parties, cloud regions and remote engineering

This baseline is a management design aid. Applicable law, sector obligations and architecture determine the final controls.

4. Convert acquisition diligence into ownership controls

Cyber diligence should answer a transaction question: can the target's critical services enter the ownership period without transferring an unmanaged compromise, continuity gap or unfunded obligation? The work should connect the investment thesis, transaction perimeter and integration plan. A vulnerability list without service context can overstate minor findings and understate structural exposure.

The diligence team should request service and architecture maps, material incidents, regulatory correspondence, identity and privileged-access records, external exposures, vulnerability management, endpoint and network coverage, cloud configuration, development practices, operational-technology interfaces, supplier access, backups, recovery tests, cyber insurance and security expenditure. Evidence quality should be recorded. Management statements can guide

inquiry, while decisions need underlying support or an explicit uncertainty treatment.

Findings should move to one of several destinations: valuation or investment case, transaction protection, closing condition, Day One containment, one-hundred-day remediation, longer-term transformation or accepted risk. A critical unmanaged remote-access path may require immediate containment. Unsupported systems may require funded replacement. Weak evidence can justify staged access and additional testing. A historical incident may require legal, regulatory, forensic and insurance review by qualified specialists.

The United States Department of Justice asks whether companies conduct timely risk review of acquisition targets and integrate acquired entities into compliance programmes [27]. NIST supply-chain guidance addresses cybersecurity risk across products and services [4]. The same continuity principle applies to ownership transition: preserve the evidence chain, assign each material finding, fund the response and verify closure.

Pre-close planning must respect merger-control, confidentiality, privacy and securities restrictions. Clean teams may be required for competitively sensitive or personal data. The acquirer should not connect networks, share credentials or exercise control before permitted. The integration plan should state what can be assessed, what can be prepared and what must wait for clearance or close.

Figure 2. Cyber diligence to ownership-control gate

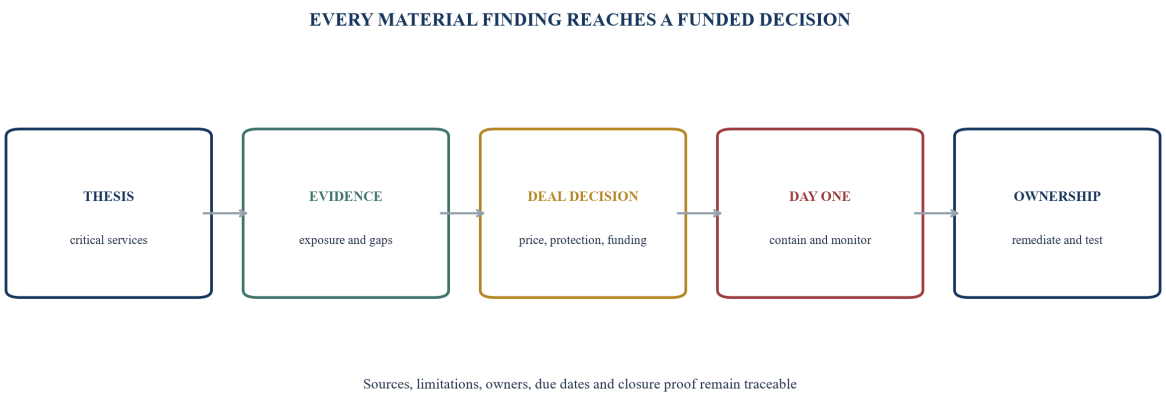

The gate sequence is illustrative. Transaction restrictions and legal advice determine permitted access and timing.

5. Establish board accountability and the three lines

The board should govern cyber resilience as an enterprise risk that can affect strategy, customers, operations, financing and reputation. It should approve risk appetite, identify critical services, review tolerances, understand material dependencies, decide investment and monitor significant exceptions. The board also needs a tested route for urgent decisions when facts are incomplete and time is limited.

Management owns service performance and control operation. The chief executive should assign one executive owner to each critical service. Technology and security leaders operate specialist

controls, but business owners decide priorities and continuity trade-offs. Risk and compliance provide challenge within their mandates. Internal audit provides independent assurance over selected evidence and governance. External specialists can add technical depth; their work should have defined scope, limitations and management ownership.

The SEC's cyber disclosure rules require covered public companies to disclose material incidents and describe processes for assessing and managing material cyber risk, together with board oversight and management roles [11]. The rule applies within its legal scope, while its emphasis on materiality and governance offers a useful reminder: board reporting should explain consequence, decisions and oversight instead of only technical activity.

Board capability can be developed through briefings, scenario exercises and decision rehearsals. The aim is informed challenge, not transforming directors into incident handlers. Directors should understand the company's service architecture, material threats, major suppliers, residual risks, insurance boundaries and notification governance. They should be able to ask what evidence supports management's confidence and what would invalidate it.

Minutes and decision records matter. They should capture the information available, limitations, alternatives, advice, authority, decision and follow-up. During an incident, the board may need to decide on customer communication, market disclosure, regulator engagement, ransom-related governance, service prioritisation and major expenditure. A pre-agreed authority map reduces avoidable delay while preserving escalation for material consequences.

6. Control identity, privilege and remote administration

Identity is a common control plane across employees, contractors, applications, machines, customers, suppliers and administrators. A sponsor baseline should require an authoritative identity lifecycle, strong authentication where appropriate, controlled privileged access, separation of administrative and everyday accounts, rapid revocation, logging and periodic review. The design should include cloud, software-as-a-service and operational-technology access rather than stopping at the corporate directory.

Privileged-access reviews should be service-led. Management should know who can change production, disable monitoring, alter backups, create payment instructions, access sensitive records or administer industrial systems. Each privilege needs a named owner, approved purpose, authentication method, duration and audit trail. Emergency access should be available under controlled conditions and tested before a crisis.

Acquisitions create acute identity risk. Dormant accounts, shared credentials, unknown service accounts, unmanaged remote tools and former supplier access can survive into the ownership period. Day One controls should inventory high-impact access, remove unsupported paths, apply monitoring and establish a controlled process for necessary exceptions. Network connection should follow evidence of containment rather than the closing timetable.

Machine identities and secrets deserve the same discipline. Certificates, API keys, tokens and embedded credentials can outnumber people and remain difficult to rotate. The company should identify which critical services depend on them, where secrets are stored, how rotation occurs and what happens when a certificate or provider fails. Development and infrastructure automation should avoid persistent, broadly privileged credentials.

Metrics should distinguish coverage from control. The percentage of users enrolled in multifactor authentication is useful, while the unresolved privileged accounts, bypass paths and high-impact applications outside coverage determine residual exposure. The board should see critical exceptions, revocation performance, privileged-session evidence and the result of recovery tests for identity services.

7. Govern IT, operational technology, cloud and interconnected systems

Critical services increasingly span enterprise IT, operational technology, cloud services, connected devices and specialised platforms. These environments have different availability, safety, maintenance and vendor constraints. A single patching or segmentation rule can be impractical or unsafe. The portfolio baseline should define outcomes and evidence, while engineering and service owners select controls suited to the environment.

Saudi Arabia's Operational Technology Cybersecurity Controls define requirements for governance, defence, resilience and third-party cybersecurity in OT environments [13]. Saudi Arabia's Essential Cybersecurity Controls provide a broader national control baseline [14]. Singapore's 2026 Code of Practice for critical information infrastructure strengthens senior-management accountability, resilience frameworks, interconnected-system oversight, exercises and threat detection [20]. These sources demonstrate the growing expectation that operational resilience includes connected and externally dependent systems.

Architecture review should identify trust boundaries, management paths, data flows, remote access, internet exposure, safety dependencies, legacy protocols and failure domains. The objective is controlled communication and recoverable operation. Segmentation should be tested for the paths that matter, including supplier maintenance and emergency operations. A diagram without configuration evidence or traffic testing is an incomplete control.

Cloud resilience requires visibility beyond provider certification. Management should know the services, regions, control planes, identity dependencies, encryption keys, data-residency requirements, quotas, backup arrangements, exit mechanisms and provider concentrations supporting each critical service. The UAE National Cloud Security Policy establishes national objectives for secure cloud adoption and data protection [8]. Contract and architecture decisions should reflect the company's applicable obligations and service tolerances.

Technical debt belongs in the investment case. Unsupported systems, scarce engineering skills, opaque vendor appliances and tightly coupled applications can increase downtime and constrain growth. Remediation choices include replacement, isolation, monitoring, compensating controls, service redesign and managed retirement. The board should see cost, risk reduction, operational dependency and delivery capacity rather than a single backlog total.

Table 3. Architecture control decisions across technology environments

Environment	Material dependency	Required evidence	Resilience decision	Typical failure to challenge	Ownership outcome
enterprise IT	identity, endpoint, network and business applications	asset ownership, configuration, logging and recovery	standardise control plane and isolate exceptions	coverage percentage hides unsupported assets	verified baseline and funded exception plan
operational technology	safe operation, engineering station and remote maintenance	zones, conduits, safety dependencies and vendor access	protect safe state and govern approved pathways	enterprise control applied without process-safety review	tested isolation and recovery sequence
public cloud	region, identity, control plane, keys and managed services	architecture, contracts, quotas, backups and exit path	design for failure domains within approved economics	provider certification treated as service recovery proof	exercised restoration and concentration decision
software as a service	provider availability, data access and identity federation	service tier, export, logs, sub-processors and continuity	tier suppliers by critical-service consequence	contract renewal without dependency reassessment	documented alternate process or accepted tolerance
connected product	firmware, update path, customer environment and telemetry	bill of materials, support life, signing and vulnerability process	align lifecycle support with product obligation	sales growth outpaces support and disclosure capability	funded secure lifecycle and retirement plan

The control choices are illustrative and require engineering, safety and legal assessment.

8. Manage suppliers, concentration and fourth parties

The portfolio should tier suppliers by consequence to critical services rather than by annual spend. A low-spend identity provider, remote-maintenance vendor, domain registrar or specialist component can create a larger resilience exposure than a high-spend commodity supplier. Tiering should consider access, substitutability, geographic concentration, financial health, data, operational dependency and the time required to move or recover.

NIST SP 800-161 provides guidance for cybersecurity supply-chain risk management across products and services [4]. ENISA's supply-chain guidance emphasises governance, supplier relationships, vulnerability handling and assurance [5]. The EU's NIS2 framework and related implementing requirements place additional attention on supply-chain security for entities within scope [15][16]. Contracts should reflect applicable obligations, while operational evidence should test whether the supplier can support the company's tolerance.

Supplier diligence should address architecture, ownership, control evidence, incidents, vulnerabilities, subcontractors, data, access, continuity, recovery, insurance, notification, audit rights and exit. Certifications can narrow the inquiry; they do not show how the supplier supports the particular service or how current exceptions affect it. Management should record reliance on shared platforms and fourth parties where the supplier's own concentration can transmit disruption.

Contract design should define service levels, security responsibilities, incident notification, cooperation, evidence access, change control, data handling, subcontracting, business continuity, termination and transition. A contractual right has value when it can be exercised in time. The company should know who invokes the right, what information is needed, which alternative exists and how service continuity is maintained during transition.

At portfolio level, aggregation can reveal hidden concentration across companies. Several businesses may depend on one cloud region, managed-security provider, payroll platform, software library or telecommunications carrier. The sponsor can coordinate review and contingency planning. Procurement leverage should support stronger evidence and resilience; it should not force uniform suppliers where concentration outweighs the economic benefit.

9. Build an incident command system before the incident

An incident plan should define authority, information flow and decision thresholds. It should identify the incident commander, technical response lead, service owners, legal counsel, communications, privacy, regulatory, finance, insurance and board contacts. The plan should state how authority changes outside working hours, when senior leaders are unavailable or when normal collaboration tools cannot be trusted.

The command system needs a common operating picture. It should track affected services, known facts, evidence quality, threat activity, containment, customer consequence, financial exposure, legal and contractual deadlines, recovery dependencies and decisions. Technical severity and business materiality should remain separate because a technically contained event can still require disclosure, while a widespread alert may have little service consequence.

Notification obligations differ by jurisdiction, sector, contract and fact pattern. The EU General Data Protection Regulation, NIS2, SEC rules, UAE requirements and sector regulators can create distinct assessments and timeframes [11][15][17]. Qualified counsel and responsible officers should maintain the applicable matrix. The operational design should collect decision-grade facts early enough for them to act.

Communications should be accurate, timely and coordinated. Customers, employees, regulators, lenders, insurers, law enforcement and the market may require different information. Speculative attribution, premature certainty and inconsistent statements can impair trust and response. The plan should assign approval, preserve versions and state what is known, what remains under investigation, what action recipients should take and when the next update is expected.

The sponsor can support coordination when an incident spans companies or a shared provider. A portfolio escalation protocol should define when a company alerts the sponsor, what assistance is available, how privileged information is handled and who can contact other portfolio companies. The company incident commander retains operational control unless a legally approved arrangement states otherwise.

Figure 3. Incident command and notification clock

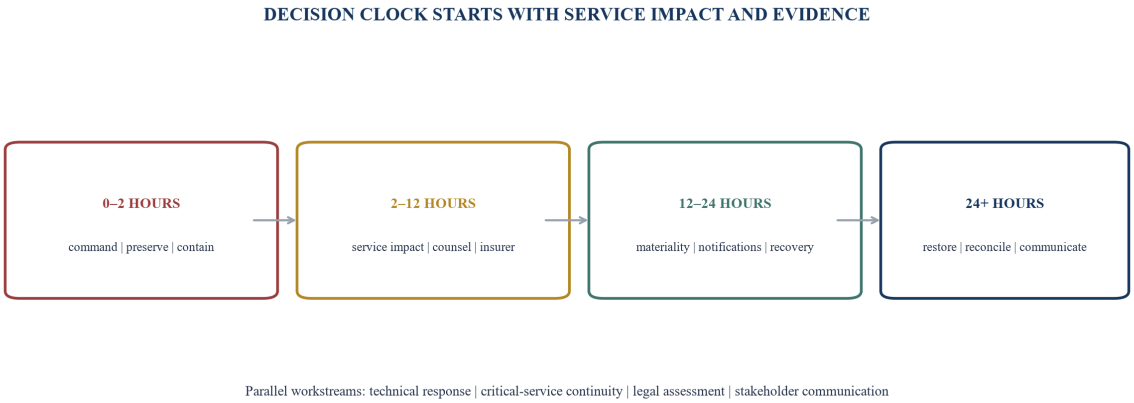

The timing bands are hypothetical management design windows, not legal deadlines. Applicable obligations require qualified assessment.

10. Design recovery around service outcomes and clean evidence

Recovery is a controlled business process. It should restore the service to a known state, reconcile data and transactions, validate security, manage backlog and communicate with affected stakeholders. A system can be online while the service remains unsafe or financially unreliable. Service owners therefore need acceptance criteria beyond technical availability.

Backup design should address segregation, immutability where appropriate, access, monitoring, retention, restoration order and the dependencies needed to use restored data. A successful backup job is evidence of copying; a controlled restore proves whether the copy, keys, applications, identities and procedures work together. The company should test restoration under conditions that remove or distrust parts of the normal environment.

Recovery sequencing should follow critical-service priorities. Identity, communications, network, data, core applications, operational systems and external connections may have circular dependencies. The plan should identify minimum viable service, manual alternatives, capacity constraints and decision points for returning to normal operation. Backlog clearance should be modelled because delayed orders, claims, payments or clinical records can create a second operational crisis.

NIST contingency-planning guidance describes business-impact analysis, recovery strategies, testing and plan maintenance [25]. ISO 22301 provides an international management-system standard for business continuity [22]. CISA's ransomware guidance addresses preparation, response and recovery practices [3]. These sources support repeatable planning, while the company's architecture and consequences determine the actual recovery design.

After restoration, management should reconcile what changed. This may include financial transactions, inventory, orders, access, configurations, customer records and operational commands. Evidence should show the recovery source, restoration time, validation, exceptions,

approvals and outstanding monitoring. The board should receive the service outcome, retained exposure, customer and regulatory consequences, cost and actions needed to prevent recurrence.

11. Quantify exposure and investment without false precision

Cyber investment decisions need financial discipline, yet the evidence rarely supports a single precise loss estimate. A decision framework can identify exposure pathways, plausible consequence ranges, control effects, implementation cost and residual uncertainty. It should distinguish expected operating cost, avoided loss, revenue protection, regulatory exposure, insurance recovery and strategic enablement.

The starting point is the critical service. Management maps scenarios such as service interruption, data corruption, fraud, unsafe operation, intellectual-property loss or supplier compromise. Each scenario identifies affected customers, duration, volume, margin, working capital, remediation, legal costs, contractual remedies, insurance and longer-term commercial effects. The model should avoid adding overlapping consequences or treating gross revenue as lost profit.

Probability estimates require evidence. Internal incidents, control tests, sector data and expert judgement can inform ranges. Where probability cannot be supported, management can use threshold analysis: how frequent or severe would the event need to be for the investment to meet the approved hurdle? Scenario ranges and break-even points are more transparent than a precise figure presented without a defensible basis.

The investment portfolio should separate mandatory remediation, resilience maintenance, growth enablement and discretionary improvement. Some expenditure protects the licence to operate. Some removes a deal-specific exposure. Some enables a product, customer or acquisition. Some reduces operating cost. Each initiative should state the service affected, current evidence, target outcome, cost, dependency, owner, delivery risk and verification method.

The hypothetical bridge below shows the method. It does not estimate market loss or return. The board should review downside concentration and reversibility: whether the investment removes a single failure point, improves detection, shortens recovery, increases strategic flexibility or creates permanent operating cost. Finance should validate claimed cash and earnings effects; risk reduction should remain separately identified unless a recognised financial effect exists.

Figure 4. Hypothetical cyber investment decision bridge

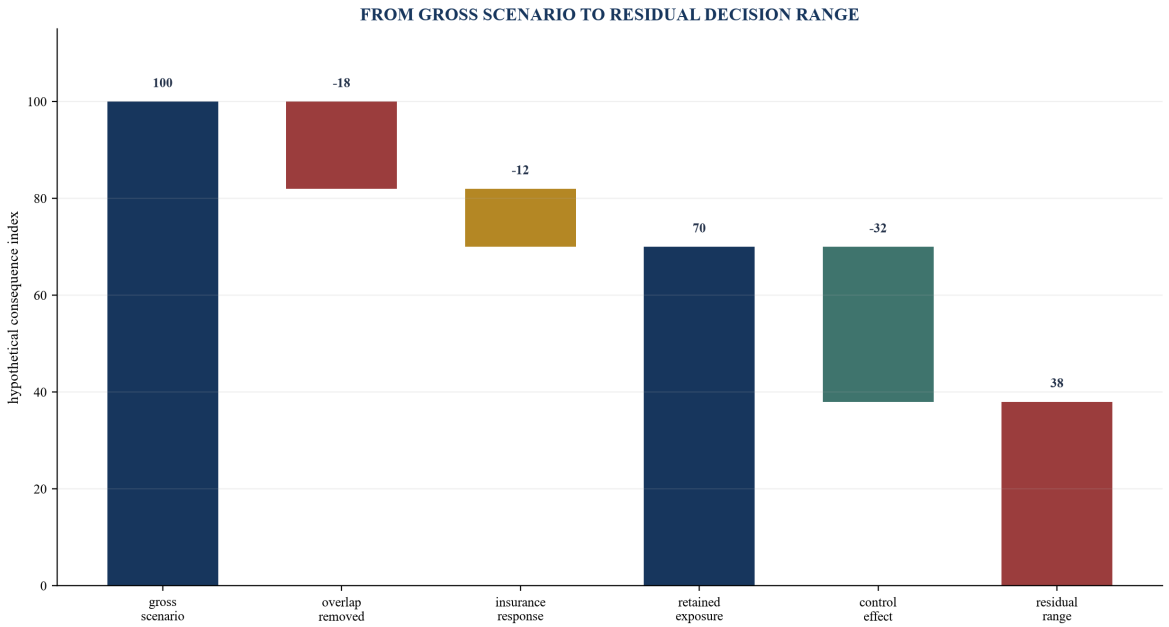

Values are hypothetical management assumptions in index points. They are not market estimates, forecasts or promised benefits.

12. Use cyber insurance as bounded risk transfer

Cyber insurance can transfer defined costs and provide response resources, subject to policy terms. It does not replace resilience or remove operational consequence. The company should understand coverage, limits, sublimits, deductibles, waiting periods, exclusions, territorial scope, notification, consent, panel providers, aggregation and the treatment of systemic events. The policy should be assessed against the critical-service scenarios rather than headline limit alone.

The placement process should reconcile the proposal, underwriting representations and current control evidence. Inaccurate or stale answers can create dispute when the company most needs coverage. Security, legal, finance and insurance advisers should agree who owns each representation and how material changes are reported. Acquisitions, divestments, new cloud dependencies and major control exceptions may affect the risk presented to insurers.

Scenario analysis should identify insured and uninsured loss. Business interruption may depend on waiting periods, measurement methods and whether the trigger occurs within the insured environment or at a dependent provider. Incident response, forensic, legal, notification, restoration, fraud, extortion and liability cover can have different terms. Qualified advisers should assess application to the company's facts and law.

Response plans should integrate the insurer without surrendering management control. The company should know how to notify, obtain consent, appoint specialists, preserve privilege and document cost. A delay in finding policy details or approved providers can impair both response and recovery. Exercises should include the broker or insurer where practical and test a scenario that raises a genuine coverage boundary.

At portfolio level, the sponsor can compare retained exposures, policy structures and concentration. A common programme may create scale, but a uniform limit or wording may not suit different service consequences and jurisdictions. The decision should consider company

autonomy, claims control, aggregation, programme dependency and how a portfolio-wide event would allocate limits.

13. Apply jurisdiction and sector overlays

The portfolio framework needs a controlled overlay process because duties differ across countries and sectors. The overlay should identify covered entities and services, governance duties, required controls, incident reporting, regulator access, record retention, localisation, supplier requirements and personal liability where applicable. It should link each duty to an operating owner and evidence rather than sit as a legal summary.

Dubai Law No. 15 of 2024 establishes the Dubai Electronic Security Centre's role and addresses critical non-government entities [7]. Dubai's published standards and policies provide further local security instruments [9]. The UAE National Cloud Security Policy establishes national direction for secure adoption [8]. Saudi Arabia publishes Essential Cybersecurity Controls and separate Operational Technology Cybersecurity Controls [13][14]. Companies should obtain current qualified advice on scope and implementation.

Within the European Union, NIS2 creates cybersecurity risk-management and incident-reporting requirements for entities in scope, and DORA establishes digital operational-resilience requirements for financial entities [15][17]. The Cyber Resilience Act addresses products with digital elements [28]. The Critical Entities Resilience Directive addresses resilience of critical entities beyond cyber alone [29]. These regimes can overlap with privacy, sector regulation and contractual requirements.

Singapore's 2026 critical-infrastructure code emphasises board and senior-management accountability, resilience frameworks, interconnected systems, exercise planning and evolving threats [20][21]. Australia applies the Security of Critical Infrastructure framework and sector obligations alongside the Essential Eight guidance, while APRA CPS 230 applies to regulated financial entities [18][19][30]. The UK combines sector regulation, the NCSC framework and data-protection obligations [6][24][26].

The overlay should identify common control evidence once and map it to multiple obligations. One tested incident plan may support several requirements, while notification decisions remain specific. This approach reduces duplicate administration and makes gaps visible. The compliance map should retain source, version, effective date, interpretation owner and last review date so the board can distinguish current requirements from planned change.

Table 4. Geographic and sector overlay register

Overlay	Governance focus	Operating evidence	Decision owner	Refresh trigger
UAE and Dubai	critical entities, government direction, cloud and local standards	entity scope, control mapping, notification path and approved contacts	local board, security and counsel	legal update, service classification or cloud change
Saudi Arabia	essential controls, OT controls and national requirements	applicable control assessment, OT boundary and remediation record	local board, security and engineering	control update, acquisition or system change
European Union	NIS2, DORA, privacy, product and critical-entity rules	scope memo, risk measures, supplier register, incidents and testing	entity board, compliance and counsel	new activity, member-state implementation or product launch
United Kingdom	essential functions, sector regulation and data security	service map, impact tolerance, CAF assessment and notification matrix	board, risk and responsible officers	regulatory update, material outsourcing or incident
United States	sector rules, material disclosure and contractual duties	materiality process, board oversight, incident record and disclosure control	board, legal and finance	listing change, acquisition or material incident
Singapore and Australia	CII or critical-infrastructure duties, resilience and service providers	entity classification, exercises, interconnected-system and provider evidence	local board and accountable executives	designation, code update or outsourcing change

The table identifies research anchors, not a legal determination. Companies require current advice on scope and duties.

14. Integrate acquisitions without transferring compromise

Technology integration can transmit risk from target to platform or from platform to target. The ownership plan should therefore separate legal control from technical trust. Closing does not justify immediate network connection, credential federation, data migration or shared administration. Each connection should pass an evidence gate based on service need, exposure, containment, monitoring and rollback.

Day One priorities include privileged access, external exposure, endpoint coverage, active incidents, backup viability, payment and treasury controls, critical supplier access, communication and escalation. High-risk pathways can be disabled or routed through monitored interfaces. Necessary business access can use controlled virtual environments, temporary accounts and segmented collaboration while the architecture is assessed.

The integration pattern should reflect the thesis. Full integration may suit a standardised platform. Federated operation can preserve local systems with common governance and reporting. Shared services can centralise selected capabilities under service levels. Ring-fenced operation can protect regulated, safety-critical or technically fragile environments. The pattern should state controlled interfaces, ownership, funding and conditions for later change.

Migration requires a service acceptance plan. It should include data reconciliation, identity, security testing, capacity, performance, recovery, user readiness, supplier support, rollback and a defined stabilisation period. The team should avoid measuring completion by system cutover

alone. Customer service, transaction integrity, finance close, incident detection and operating workload should remain within approved bounds.

An acquisition closure gate should confirm that material diligence findings have a destination, temporary controls are owned, integration access is approved, service tolerances remain achievable and the first recovery tests are complete. Open strategic transformation can continue beyond the gate. The governance objective is to close unmanaged transition risk, not to declare that cyber work is finished.

15. Operate portfolio detection, intelligence and mutual aid

Portfolio-level capability can improve situational awareness when it aggregates signals that matter across companies. Useful shared intelligence includes exploited vulnerabilities, supplier incidents, credential exposure, domain abuse, attack patterns and government alerts. The sponsor can coordinate distribution and ask companies to confirm exposure and action. The process should protect sensitive company information and follow applicable law.

A common event taxonomy helps compare without forcing identical technology. Companies can report whether a critical service, regulated record, material supplier or privileged control was affected; whether the event is contained; and whether external notification is under assessment. Aggregation should preserve evidence quality. An unconfirmed alert should not be presented as an incident, and an incident count should not be treated as a performance ranking without context.

Shared specialist resources can include incident response retainers, forensic readiness, legal panels, crisis communications, threat intelligence, exercise design and recovery engineering. Each arrangement needs activation authority, scope, confidentiality, privilege, cost and priority rules. A portfolio-wide event can exhaust a shared provider, so contracts and exercises should test concurrent demand.

Mutual aid can provide experienced personnel, alternative facilities, supplier contacts or temporary operating support. It should be pre-authorised and technically safe. A company should not import personnel or equipment into a compromised environment without identity, access and evidence controls. Assistance should operate under the incident commander's priorities and preserve the recipient company's legal and regulatory responsibilities.

The sponsor should maintain its own continuity plan for portfolio communication. It needs current contacts, secure alternate channels, decision authority and the ability to operate when normal collaboration, identity or document systems are unavailable. The plan should be exercised with at least one company and one critical supplier scenario rather than reviewed only as a document.

16. Govern metrics, evidence quality and exceptions

Metrics should help the board decide. A compact scorecard can show critical-service coverage, material exceptions, control operation, detection, recovery, supplier concentration, incidents, investment and closure. Each measure needs a definition, source, owner, frequency and scope. Changes in data quality should be visible so an improving number is not caused by excluding difficult assets or redefining the denominator.

Leading measures include service maps completed, privileged access reviewed, critical vulnerabilities within treatment, protected backups tested, supplier plans current and exercises completed. Lagging measures include disruption time, customer impact, financial loss, notification and recovery performance. Both need consequence context. A high alert count may reflect better detection; a low incident count may reflect weak visibility.

Evidence should be graded by source and recency. A policy is designed evidence. A configuration export or completed access review is operating evidence. A controlled exercise, restoration or independent test is resilience evidence. Management self-assessment can support prioritisation, while material board confidence should identify where it relies on independent assurance or tested outcomes.

Exceptions should be reported by critical service, consequence and age. A backlog count hides whether one expired exception can interrupt an essential service. The board should focus on risk accepted above appetite, repeated missed dates, concentration, weak compensating controls and initiatives that lack funding or accountable delivery capacity. The sponsor can compare patterns and mobilise help where several companies face the same constraint.

The hypothetical heat map below illustrates aggregation. Scores show a management prioritisation method, not measured risk or company performance. A final score should never erase the underlying consequence, evidence and decision. The committee should be able to move from a cell to the service, dependency, exception, owner and treatment plan.

Figure 5. Hypothetical portfolio cyber-resilience heat map

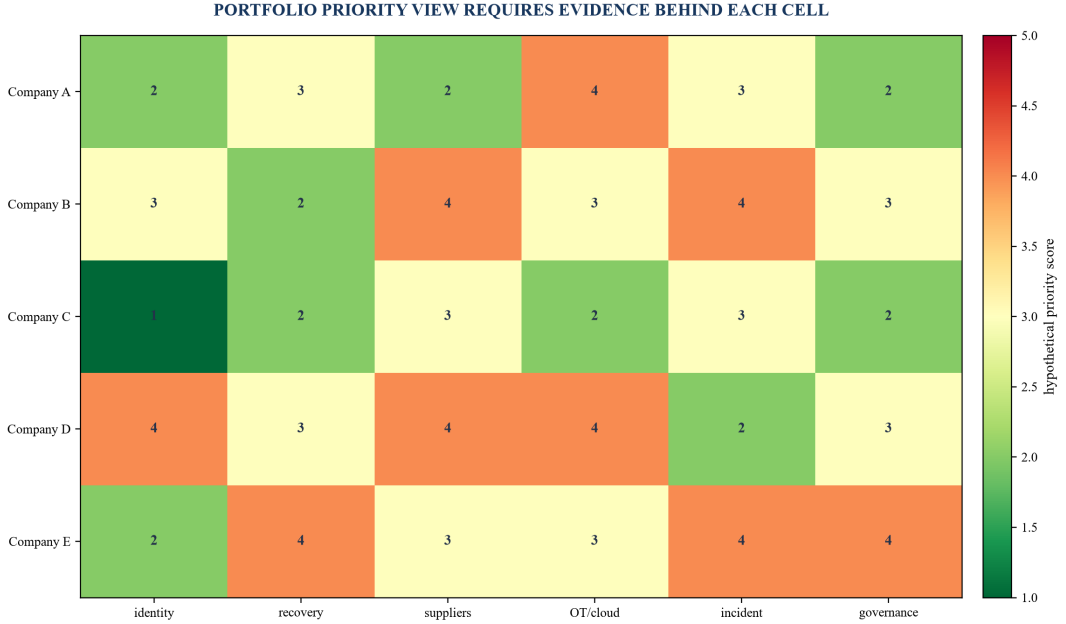

Scores are hypothetical management assumptions for method illustration. They are not assessments of any company or sector.

17. Exercise severe but plausible scenarios

Exercises test decisions, dependencies and evidence that ordinary control reporting cannot reveal. The scenario should threaten a critical service, remove one or more normal capabilities and create uncertainty. Useful cases include identity-provider compromise, destructive ransomware, cloud control-plane failure, manipulated industrial data, supplier remote-access compromise, payment fraud and concurrent disruption across portfolio companies.

The exercise should have explicit objectives. A tabletop can test governance and notification. A technical simulation can test detection and containment. A recovery exercise can test restoration and reconciliation. A supplier exercise can test escalation and alternate operation. A full operational exercise can test customer service and manual workarounds. Combining every objective into one event can create theatre without reliable evidence.

Participants should receive information in realistic stages. The board needs enough uncertainty to practise materiality and communications decisions. Service owners need capacity and backlog data. Technical teams need indicators and access constraints. Legal, regulatory, privacy, insurance and communications teams need deadlines and conflicting facts. Controllers should record decisions, evidence requested, delays and assumptions.

The result should be a prioritised improvement plan. Findings may concern authority, contacts, supplier response, backup access, restoration order, data reconciliation, customer communication or board information. Each action needs a service consequence, owner, funding, due date and validation method. Repeating an exercise without closing material findings can produce false reassurance.

Singapore's updated critical-infrastructure code includes exercise planning and evolving threat considerations [20][21]. NIST, CISA and NCSC guidance also support testing response and recovery [1][2][6]. The portfolio schedule should respond to service importance and material change. A company that migrated its core platform or changed a critical provider may need an exercise before the annual calendar would otherwise require one.

18. Run a sponsor operating cadence

The sponsor cadence should connect company action to portfolio decision-making. A monthly operating review can address material incidents, critical exceptions, investment delivery and shared threats. A quarterly board or portfolio review can assess service tolerances, recovery evidence, suppliers, scenario results and risk appetite. An annual review can reset the baseline, insurance, assurance plan and transformation priorities.

Cadence should be event-driven as well as calendar-driven. Acquisitions, divestments, refinancing, major outsourcing, product launches, cloud migration, material incidents, regulatory change and leadership turnover should trigger targeted review. The sponsor can define trigger criteria so management knows when an ordinary company decision has portfolio relevance.

Portfolio reporting should use a stable data dictionary. Each company reports the same outcome definitions with an allowed local explanation. The sponsor should avoid ranking companies through a composite score that rewards easy evidence or penalises necessary complexity. Comparative views are useful for identifying shared gaps, inconsistent interpretation and opportunities for central capability.

The cadence also needs decision rights. A portfolio committee may challenge an exception, recommend funding, commission assurance or coordinate shared response. The company board approves risk acceptance and expenditure within its authority. Reserved matters and lender or regulator requirements should be mapped. Escalation should state what happens when the company and sponsor disagree about consequence, timing or resources.

The operating model needs capacity. A small sponsor team cannot manually inspect every control across a large portfolio. It should prioritise critical services, material exceptions, recent change and weak evidence. Common tooling can support collection, but management conversations and tested outcomes remain essential. The cadence should be designed around decisions that can be made, not data that can be harvested.

Table 5. Sponsor cyber-resilience operating cadence

Forum	Frequency or trigger	Decision inputs	Primary decisions	Evidence retained
company service review	monthly and after material change	service performance, exceptions, incidents and dependencies	treatment, resources, acceptance and escalation	minutes, owners, due dates and closure evidence
portfolio resilience review	quarterly	company scorecards, concentrations, exercises and investment	shared capability, assurance, challenge and priority	portfolio decision log and company responses
incident coordination	material event or shared exposure	service impact, facts, legal assessment and response capacity	mutual aid, specialist activation and portfolio communication	activation record, permissions and action trail
acquisition gate	diligence, signing, close and connection	critical services, exposure, deal protections and Day One plan	price input, containment, funding and connection approval	finding-to-action register and gate approval
annual reset	strategy, insurance and budget cycle	threat change, incidents, tests, regulation and architecture	baseline, risk appetite, assurance and capital plan	approved framework, budget and assurance plan

The cadence is illustrative. Frequency and authority depend on portfolio risk, company governance and applicable obligations.

19. Build exit-ready resilience evidence

Cyber resilience can influence buyer confidence, diligence scope, transaction protection, insurance, financing and integration planning. Exit readiness begins during ownership. The company should preserve a coherent record of critical services, architecture, governance, material risks, incidents, investment, supplier dependencies, control operation, recovery tests and improvement. The record should show current state and the route by which material gaps were closed.

The evidence room should be structured for decision use. Governance includes board oversight, authority, risk appetite and exceptions. Service evidence includes tolerances, dependency maps and exercised recovery. Control evidence includes identity, exposure, logging, vulnerability, backup and incident operation. Supplier evidence includes tiering, contracts, concentration and tests. Incident evidence should be legally reviewed and disclose facts appropriately.

Metrics should reconcile over time. If a company reports improved coverage, the denominator, scope and source should be available. If an investment claim includes avoided loss, the method and uncertainty should remain visible. If certifications or independent reports are used, scope, date, exclusions and management responses should be retained. Buyers may discount polished dashboards that cannot be traced to operating evidence.

Known gaps should have a credible treatment. Concealing them can damage trust and delay diligence. A buyer can often evaluate a documented legacy system with isolation, monitoring, replacement funding and delivery progress more effectively than an unexplained exception discovered late. Management should distinguish accepted residual risk from overdue action and identify consequences that can transfer to the new owner.

The exit narrative should connect resilience to the business. It explains how controls protect customer commitments, regulatory standing, revenue, cash, growth and integration. This makes cyber evidence relevant to valuation without promising a premium. Actual transaction effects depend on buyer requirements, market conditions, incidents, architecture, sector, financing and negotiation.

20. Use a board scorecard that preserves the underlying decisions

A board scorecard should fit on one page and open into evidence. It can show critical-service coverage, top residual exposures, recovery performance, material suppliers, incidents, investment, assurance and decisions required. Each item should identify trend, evidence date, accountable owner and the service affected. The board should receive narrative where numbers cannot responsibly capture uncertainty.

Red, amber and green can support attention, but colour needs a defined basis. Green should mean the approved outcome is supported by current evidence, not that cyber risk is absent. Amber should identify a bounded exposure with owned treatment or a material evidence limitation. Red should identify consequence above appetite, failed control or recovery, unclear authority, or an overdue material decision.

The scorecard should include decisions rather than observations. Examples include approving a resilience investment, accepting a time-bound exception, requiring independent assurance, changing a critical provider, revising the service tolerance or activating a portfolio response. Management should state recommendation, alternatives, consequence, cost, timing and evidence. The board record should capture the decision and follow-up.

Assurance should be risk-based. Internal audit, penetration testing, red teaming, architecture review, restoration tests and supplier assessment answer different questions. Scope should follow critical services and material change. Findings should reconcile to management's register so the board can see where independent work confirms, contradicts or adds to existing evidence.

The table below provides a practical design. It is a hypothetical template. Boards should adapt measures and thresholds to the business, applicable law and risk appetite.

Table 6. Board cyber-resilience scorecard

Board lens	Decision-grade measure	Required explanation	Evidence source	Example board action
critical services	services with approved tolerance and tested dependency map	gaps, recent change and customer consequence	service register and exercise	revise tolerance or fund dependency remediation
material exposure	top exceptions above or near appetite	consequence, compensating control and expiry	exception register and assurance	accept, accelerate, isolate or redesign
recovery	services restored and reconciled within approved objective	failed steps, backlog and dependency	restoration record and business validation	require repeat test or alternate recovery path
suppliers	critical providers with current evidence and exercised continuity	concentration, fourth parties and exit	supplier register, contract and exercise	renegotiate, diversify or accept concentration
incidents	events by service consequence and notification status	facts, uncertainty, response and lessons	incident record and legal review	approve disclosure, remediation or independent review
investment	delivery, cost and verified service outcome	dependency, delay, benefit and residual risk	programme record and finance validation	continue, resequence, stop or add capacity

Measures and decision thresholds are hypothetical. Company boards should approve definitions and evidence standards.

21. Sequence a one-hundred-day implementation

The first thirty days establish accountability and visibility. The sponsor approves the framework and reporting definitions. Each company confirms board and executive owners, names critical services, identifies urgent authority gaps and records material incidents, suppliers and exceptions. Acquisition targets or recently acquired businesses receive a controlled exposure and privileged-access review. The portfolio identifies common providers and shared response resources.

Days thirty-one to sixty establish the operating baseline. Companies map dependencies for priority services, confirm tolerances, validate identity and backup evidence, tier suppliers, test incident contacts and build funded treatment plans. The sponsor reconciles common gaps, assesses concurrent-response capacity and defines the quarterly portfolio review. Qualified advisers confirm jurisdiction and sector overlays.

Days sixty-one to one hundred test the system. Each company exercises one severe but plausible scenario linked to a critical service. At least one exercise should involve a supplier or recovery process. Management closes urgent findings, takes material exceptions to the appropriate board and updates the investment case. The sponsor runs a portfolio coordination exercise and records lessons for acquisitions, insurance and shared capability.

The sequence should adapt to consequence. An active compromise, unsupported critical system, failed recovery or ungoverned supplier access requires immediate action. A company with stable evidence may focus on deeper exercises and architecture. The framework should not create a uniform transformation programme detached from operating facts.

Success at day one hundred means the portfolio has a working decision system: accountable boards, named critical services, visible material exposure, tested incident authority, recovery evidence, supplier understanding, approved investment and a sponsor cadence. It does not mean cyber risk has been removed. Continued value comes from maintaining the evidence, testing change and using the framework during acquisitions, strategy, financing and exit.

Figure 6. Hypothetical one-hundred-day implementation sequence

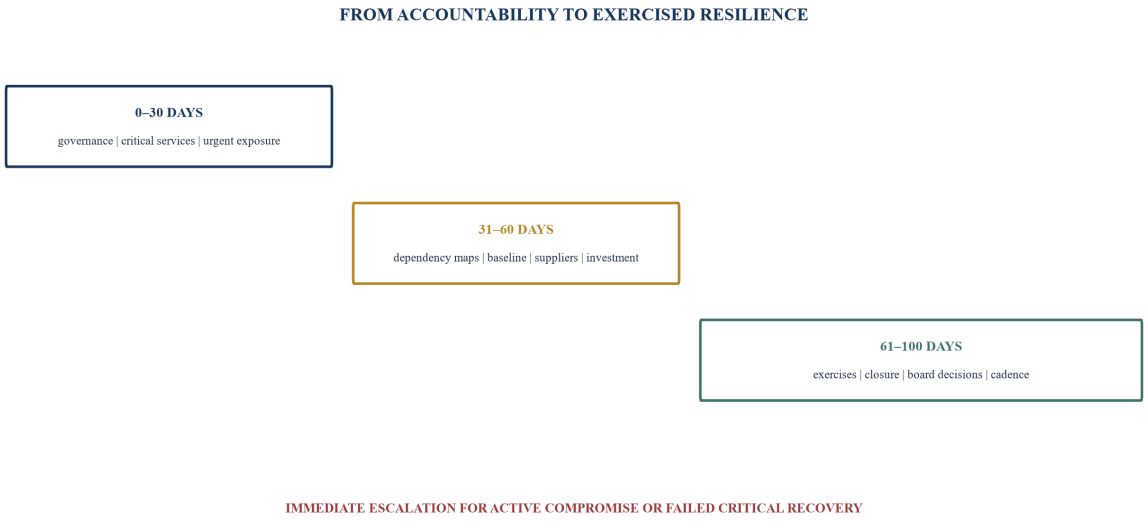

The timing is an illustrative management sequence. Active incidents and material exposures require immediate company-specific action.

References

[1] National Institute of Standards and Technology, The Cybersecurity Framework 2.0 https://www.nist.gov/publications/nist-cybersecurity-framework-csf-20

[2] Cybersecurity and Infrastructure Security Agency, Cross-Sector Cybersecurity Performance Goals https://www.cisa.gov/cybersecurity-performance-goals

[3] Cybersecurity and Infrastructure Security Agency, StopRansomware Guide https://www.cisa.gov/stopransomware/ransomware-guide

[4] National Institute of Standards and Technology, SP 800-161 Rev. 1 Cybersecurity Supply Chain Risk Management Practices https://csrc.nist.gov/pubs/sp/800/161/r1/final

[5] European Union Agency for Cybersecurity, Good Practices for Supply Chain Cybersecurity https://www.enisa.europa.eu/sites/default/files/publications/Good%20Practices%20for%20Supply%20Chain%20Cybersecurity.pdf

[6] United Kingdom National Cyber Security Centre, Cyber Assessment Framework 4.0 https://www.ncsc.gov.uk/files/NCSC-Cyber-Assessment-Framework-4.0.pdf

[7] Government of Dubai, Law No. 15 of 2024 Concerning the Dubai Electronic Security Centre https://dlp.dubai.gov.ae/Legislation%20Reference/2024/Law%20No.%20%2815%29%20of%202024%20Concerning%20the%20Dubai%20Electronic%20Security%20Centre.html

[8] Government of the United Arab Emirates, National Cloud Security Policy https://uaelegislation.gov.ae/en/policy/details/lsy-s-lotny-llamn-lsh-by

[9] Dubai Electronic Security Centre, Standards and Policies https://www.desc.gov.ae/regulations/standards-policies/

[10] European Union, Commission Delegated Regulation 2024/1366 https://eur-lex.europa.eu/eli/reg_del/2024/1366/oj/eng

[11] United States Securities and Exchange Commission, Cybersecurity Risk Management Strategy Governance and Incident Disclosure https://www.sec.gov/rules-regulations/2023/07/s7-09-22

[12] United States Securities and Exchange Commission, Final Rule 33-11216 https://www.sec.gov/files/rules/final/2023/33-11216.pdf

[13] Saudi National Cybersecurity Authority, Operational Technology Cybersecurity Controls https://nca.gov.sa/en/regulatory-documents/controls-list/otcc/

[14] Saudi National Cybersecurity Authority, Essential Cybersecurity Controls ECC 2-2024 https://nca.gov.sa/en/regulatory-documents/controls-list/ecc/

[15] European Union, Directive 2022/2555 on Measures for a High Common Level of Cybersecurity across the Union https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[16] European Union Agency for Cybersecurity, NIS Directive 2 https://www.enisa.europa.eu/topics/state-of-cybersecurity-in-the-eu/cybersecurity-policies/nis-directive-2

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

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

[19] Australian Signals Directorate, Essential Eight Maturity Model https://www.cyber.gov.au/sites/default/files/2025-03/PROTECT%20-%20Essential%20Eight%20maturity%20model%20%28November%202023%29.pdf

[20] Cyber Security Agency of Singapore, Codes of Practice https://www.csa.gov.sg/legislation/codes-of-practice/

[21] Cyber Security Agency of Singapore, Updated Cybersecurity Code of Practice for Critical Information Infrastructure https://www.csa.gov.sg/news-events/press-releases/cybersecurity-code-of-practice-for-critical-information-infrastructure-to-be-updated-to-address-apt-and-ai-enabled-threats/

[22] International Organization for Standardization, ISO 22301 Security and Resilience Business Continuity Management Systems https://www.iso.org/standard/75106.html

[23] Canadian Centre for Cyber Security, Baseline Cyber Security Controls for Small and Medium Organizations https://www.cyber.gc.ca/en/guidance/baseline-cyber-security-controls-small-and-medium-organizations

[24] Bank of England, SS1/21 Operational Resilience: Impact Tolerances for Important Business Services https://www.bankofengland.co.uk/prudential-regulation/publication/2021/march/operational-resilience-impact-tolerances-for-important-business-services-ss

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

[26] United Kingdom Information Commissioner's Office, Security Outcomes https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/security/a-guide-to-data-security/security-outcomes/

[27] United States Department of Justice, Evaluation of Corporate Compliance Programs https://www.justice.gov/criminal/criminal-fraud/page/file/937501/dl

[28] European Union, Regulation 2024/2847 on Horizontal Cybersecurity Requirements for Products with Digital Elements https://eur-lex.europa.eu/eli/reg/2024/2847/oj

[29] European Union, Directive 2022/2557 on the Resilience of Critical Entities https://eur-lex.europa.eu/eli/dir/2022/2557/oj

[30] Australian Government, Security of Critical Infrastructure Act 2018
https://www.legislation.gov.au/C2018A00029/latest/text

About the Author

Chennakeshav Adya is an independent researcher and Managing Partner at Matchpoint Partners. His work focuses on corporate finance, M&A, private capital, strategy and transaction execution across developed and emerging markets. This research helps boards, sponsors and operating teams convert complex control and resilience questions into evidence-led decisions.