CMORG
CROSS MARKET OPERATIONAL RESILIENCE GROUP

Guidance for Firm Cyber Recovery Capabilities

VERSION 1.0 | APRIL 2026 | TLP CLEAR

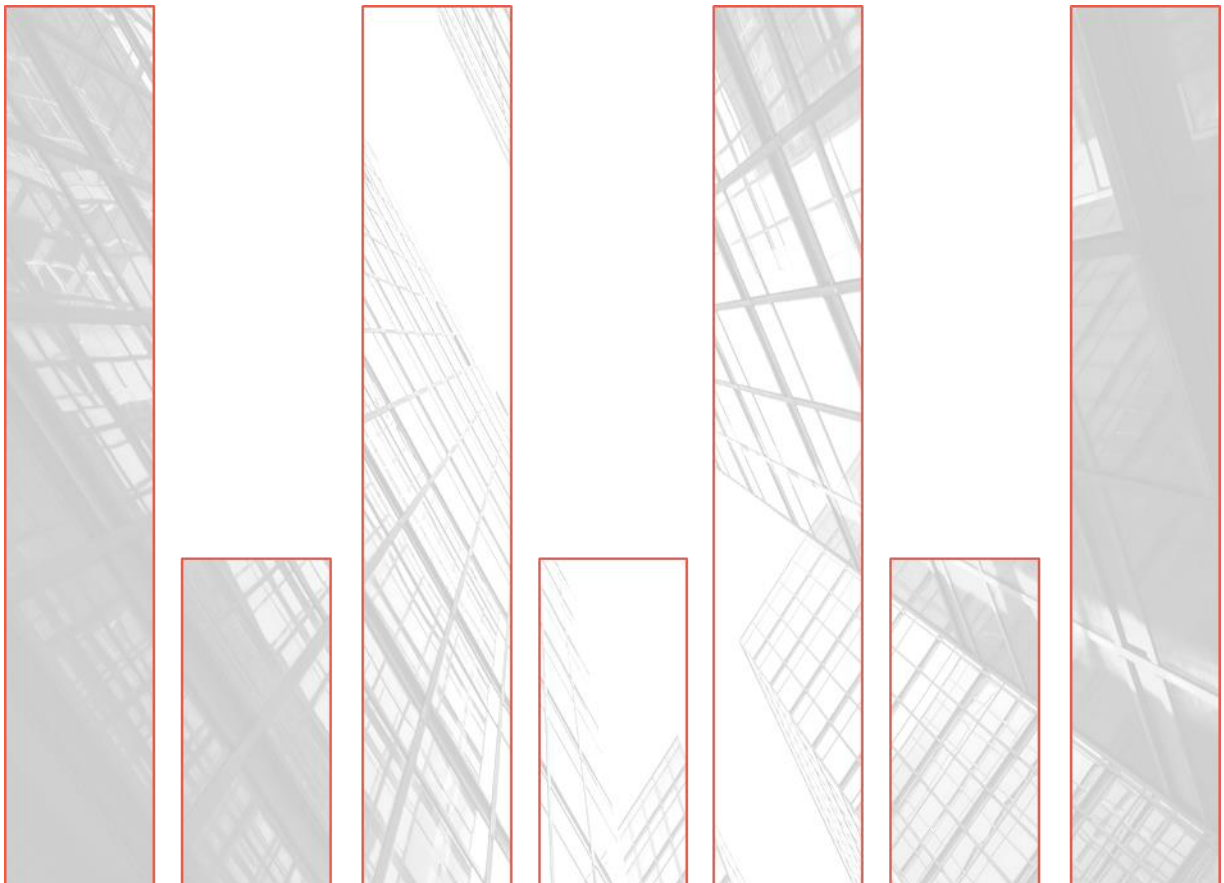

Contents

Introduction **2**

1 Governance and strategy **5**

1.1 Board oversight and risk appetite 5
1.2 Cyber recovery strategy 5
1.3 Prioritisation and scoping 6

2 Cyber recovery capabilities **8**

2.1 Capabilities lifecycle 8
2.2 Operating model and processes 9
2.3 Data protection 9
2.4 Isolation 12
2.5 Trusted recoverability 13
2.6 Interim contingency measures 14

3 Recovery sequencing **16**

3.1 Invocation and orchestration 16
3.2 Tiered recovery 17
3.3 Disconnection and reconnection 19
3.4 Data repair and reconciliation 20
3.5 Service resumption 21

4 Assurance and continuous improvement **22**

4.1 Pre-incident testing 22
4.2 During incident 25
4.3 Continuous improvement 27

Next steps **28**

Appendix A: List of stressed impacts **29**

Appendix B: Key reference material **30**

CMORG-endorsed capabilities (including good practice guidance, response frameworks and contingency tools) have been developed collectively by industry to support the operational resilience of the UK financial sector. The financial authorities support the development of these capabilities and collective efforts to improve sector resilience. However, their use is voluntary, and they do not constitute regulatory rules or supervisory expectations; as such, they may not necessarily represent formal endorsement by the authorities.

Introduction

No-one is fully immune to a sophisticated, targeted cyber attack, even financial institutions with mature prevention, detection and response capabilities. Therefore, supporting firms' ability to recover effectively from, or in some way withstand, a cyber incident that threatens their viability remains a strategic priority for the Cross Market Operational Resilience Group (CMORG) and the financial sector as a whole. Recovery from such an incident is a stretching goal for any organisation and there is currently no guaranteed or proven capability to ensure this outcome. Developing and strengthening this capability will continue to be a priority for the foreseeable future and **requires long-term effort**.

The Cyber Recovery Steering Group (CRSG) was formed under CMORG in May 2025 to support sector coordination through creation and issuance of guidance to support firms in developing capabilities and approaches that can (i) accelerate recovery to a pre-determined scope, and (ii) lessen the impacts of prolonged disruption. This paper represents a first contribution to this objective, but we expect the sector's thinking to continue to evolve and mature.

Purpose and audience

This guidance is intended to support a **shared understanding of cyber recovery** across financial institutions by establishing common language and outcomes. The guidance draws from subject matter experts and relevant third parties, and is principles-based and non-prescriptive, enabling firms to consider the practices proportionately in line with their risk appetite, cybersecurity strategy, size, complexity, and operating model. The guidance brings together current thinking and emerging approaches to cyber recovery, with coverage of:

- **Governance and Strategy (Section 1)**: includes leadership oversight, scoping and definition, developing requirements and establishing a strategy that can be linked to each organisation's risk appetite.
- **Cyber Recovery Capabilities (Section 2)**: introduces currently available technological building blocks which can be combined to support recovery outcomes.
- **Recovery Sequencing (Section 3)**: explains the considerations for restoring services safely and in a controlled order, alongside parallel activities such as incident response.
- **Assurance and Continuous Improvement (Section 4)**: provides a high-level view of how organisations can test and validate aspects of their recovery capabilities and maintain them over time.

This guidance is **aimed at senior decision makers and those with operational accountability** and/or responsibility for their organisation's approach to cyber recovery.

Defining cyber recovery

This guidance is particularly focused on approaches and capabilities necessary to **recover from a cyber incident that could jeopardise a firm's viability**.

The scenario envisaged is a cyber incident where the production and disaster recovery environments are known or suspected to be compromised. An investigation is triggered, however, options such as failing over to disaster recovery arrangements or restoring backups to production may be ruled out due to the nature of the ongoing incident and associated threat. In this scenario, a cyber recovery process is initiated as part of the firm's overall incident response with a central goal of providing (at least) minimum levels of service until normal operating

conditions are restored. It is assumed that some level of breach of Impact Tolerance (IToL)[1] is likely in this scenario and therefore the objective is to minimise impacts that could contribute to further intolerable harm.

Cyber recovery under such circumstances has a greater emphasis on implementation of capabilities that represent a '*last resort*' for recovery; in less severe scenarios where disaster recovery and other resilience measures are available, these will enable a more efficient and rapid recovery that minimises operational disruption more effectively. As the cyber threat environment worsens and examples of incidents that threaten a firm's viability continue to increase in frequency, CMORG will continue to explore how to extend the severity of scenarios under which firms can recover, as well as corresponding capabilities, with the aim of supporting recovery and reducing impacts.

Limitations and key considerations

What this is: Voluntary, principles-based guidance describing outcomes and enabling capabilities that could contribute to cyber recovery in the most severe cyber scenarios, including those that threaten firm viability.

What this isn't: Regulatory rules, supervisory expectations, or a prescriptive technical architecture. It does not replace existing regulatory requirements, nor does it provide step-by-step recovery documentation. Related disciplines of detection, response and incident management are well documented in existing guidance and are explicitly excluded from the scope of this document.

The discipline of cyber recovery is still **relatively immature** and few case studies exist to demonstrate best practice. This is evidenced by firms that suffer a breach and, despite the relative maturity of their cybersecurity programmes, experience a prolonged recovery period (including to re-establish trust).

There are many inherent complexities and logistical challenges that combine to present a significant challenge for firms to achieve. Risks include:

- being over ambitious in the scope of designing a cyber recovery strategy that then becomes too complex to feasibly deliver or too brittle to rely on;
- inherent limitations in testing capabilities such as the inability to fully spin up and test a recovery site with live data mean it may not be possible to complete end-to-end testing; and
- technologies to support areas such as data integrity checking continue to evolve in maturity.

While it is in the interests of all firms to work towards the target of strengthening their ability to recover from a cyber incident that could threaten a firm's viability, this should be considered a continuous process of evolution, and not a project with a known and clearly defined end state that can be progressed towards in a linear fashion, fully measured, tested and completed.

Baseline principles

This guidance sets out **17 baseline principles** to help financial institutions to best prepare for and recover from the most disruptive of cyber incidents. For ease of reference and consistent cross-firm discussion, the principles are organised into four themes, aligned to the structure of this document, and phrased as **outcome statements** that are elaborated further in subsequent sections. Supplemental appendices are included with example scenarios and reference material to support shared understanding. The core objective is to support financial institutions as they explore how to restore critical services in a timely, safe and controlled manner

[1] For Impact Tolerance (IToL) definition, see CMORG Guidance for Firm Operational Resilience v3.0 (Apr 2025)

following such incidents, thereby limiting risk to the firm's viability, harm to its customers and the wider financial system. Each principle addresses a key consideration for this objective, from board oversight and cyber recovery solution(s) to recovery orchestration and continuous improvement.

Section	#	Principle	Outcome
Governance and Strategy	1	**Board Oversight and Risk Appetite**	Boards oversee and senior management own cyber recovery preparedness, understand its limits, and align it with risk appetite.
	2	**Cyber Recovery Strategy**	A cyber recovery strategy sets a clear, documented direction for how a firm will withstand and recover from cyber incidents, defining the scenarios it prepares for, the recovery objectives it intends to achieve, and requirements across services, data and supporting capabilities to meet these objectives.
	3	**Prioritisation and Scoping**	A scope for recovery is established that focuses on firm survivability and prioritises key services (or parts of services) for restoration following any disruption.
Cyber Recovery Capabilities	4	**Capabilities Lifecycle**	Cyber recovery capabilities are designed, implemented, tested, operated, and maintained as an end-to-end lifecycle, ensuring they remain effective.
	5	**Operating Model and Processes**	A formalised operating model and operating processes enable cyber recovery capabilities to be sustained in steady state during normal operations and executed effectively under disrupted conditions.
	6	**Data Protection**	Recovery data is protected to the level necessary to prevent interference, unauthorised access, and loss of integrity.
	7	**Isolation**	Exposure to disruptive events is minimised through approaches that distance or separate critical recovery assets from resultant impacts.
	8	**Trusted Recoverability**	Data used for recovery is subject to controls that detect tampering and confirm its correctness, accuracy, and readiness for use.
	9	**Interim Contingency Measures**	Pre-planned contingency arrangements are established and maintained that allow critical functions and obligations to be fulfilled whilst production and recovery systems are unavailable, minimising intolerable harm while recovery is underway.
Recovery Sequencing	10	**Invocation and Orchestration**	Clear conditions, authority, and processes exist for invoking cyber recovery, enabling timely transition from incident response to recovery when predefined thresholds are met.
	11	**Tiered Recovery**	Recovery is executed in a prioritised sequence based on criticality, dependencies, and feasibility, balancing customer harm, market impact, and system stability as services are progressively restored.
	12	**Disconnection and Reconnection**	Suppliers and counterparties that are critical to service recovery and/or invoking contingent processes are managed through pre-agreed disconnection, assurance, and reconnection arrangements that support safe recovery.
	13	**Data Repair and Reconciliation**	Data and system states are reconciled and validated during and after recovery to restore accuracy, completeness, and confidence before resumption and reconnection.
	14	**Service Resumption**	Clear protocols are in place to support the safe reintroduction and operation of recovered services, standing down any contingent processes invoked during recovery and the approach to manage customer reconnection.
Assurance and Continuous Improvement	15	**Pre-incident Testing**	Pre-incident testing provides confidence that cyber recovery requirements can be met, and validate that recovery capabilities, processes, and data safeguards are ready to operate under severe stress.
	16	**During Incident**	Assurance during an incident provides timely and credible confidence that recovery actions are mitigating the incident effectively and that services can be safely resumed, balancing speed, safety, and stakeholder needs without overstating certainty.
	17	**Continuous Improvement**	A continuous feedback cycle supports the evolution of cyber recovery strategy, requirements, and execution over time, strengthening firm-level recovery capabilities and broader sector preparedness.

1 Governance and strategy

This section sets out the strategic foundations for effective cyber recovery, establishing how boards and senior management should govern, prioritise, and direct recovery preparedness under scenarios that could threaten firm viability. It outlines the role of leadership in aligning recovery ambition with risk appetite, defining a clear and credible cyber recovery strategy, and making explicit choices about the service levels, data and capabilities needed to deliver them.

1.1 Board oversight and risk appetite

Principle 1: Boards oversee and senior management own cyber recovery preparedness, understand its limits, and align it with risk appetite.

Boards and senior management play a central role in overseeing cyber recovery preparedness. They set the tone, provide direction, and seek assurance that the organisation's recovery capabilities and limitations are **well understood and broadly consistent** with its stated risk appetite and wider obligations. To support this outcome, firms may find it helpful to establish a comprehensive, clearly documented cyber recovery strategy (**Principle 2**) that facilitates effective collaboration across the organisation. All relevant business functions (including technology, operations, risk, compliance, and communications) should be involved from the outset in defining requirements and designing recovery approaches that reflect genuine business priorities.

Boards and senior management should seek to **understand, and where appropriate challenge, any assumptions and limitations** within recovery plans to comprehend the potential scope and scale of disruption, including which recovery steps are deterministic and which rely on more uncertain factors (non-deterministic) such as identification and eradication of a threat actor in the firm's environment. Periodic reporting should provide a clear status of the progress to deliver the strategy including results from assurance and testing.

1.2 Cyber recovery strategy

Principle 2: A cyber recovery strategy sets a clear, documented direction for how a firm will withstand and recover from cyber incidents, defining the scenarios it prepares for, the recovery objectives it intends to achieve, and requirements across services, data and supporting capabilities to meet these objectives.

A cyber recovery strategy provides the **unifying intent** that shapes subsequent design, prioritisation, and operating decisions. This strategy will be different for each individual organisation, based on risk appetite, business requirements, legal and regulatory framework, technology environment, etc.

Firms should aim to minimise the scope of their cyber recovery strategy. As shown in **Figure 1**, complexity increases as scope grows which in turn increases the difficulty of delivering a viable cyber recovery capability.

Figure 1. Minimising scope increases the likelihood of success

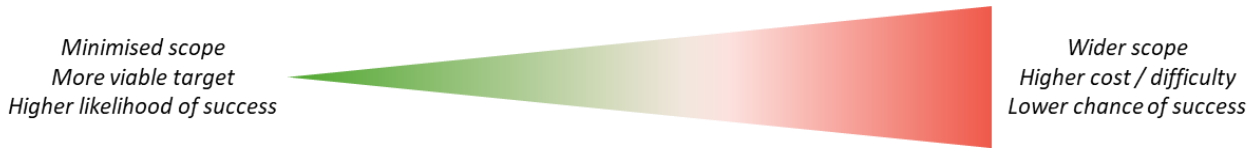

Recovery efforts are more likely to falter when firms have not explicitly agreed on '*what they are trying to achieve*' in the event of a cyber incident that threatens the firm's viability. A credible strategy therefore also

articulates the overarching purpose of recovery, the outcomes that matter most for the firm and its stakeholders, and the degree of disruption the firm is prepared to tolerate.

The strategy also **establishes, at a high level, the scenarios**, supported by risk assessments, that the organisation chooses to prepare against. Without explicit and realistic scenarios, recovery planning may default to optimistic assumptions or narrow threat models that do not reflect actual systemic or firm-specific risk. Defining relevant scenarios (for example, by drawing on real-world events as outlined in **Appendix A**) helps firms test their assumptions, understand key vulnerabilities, and ensure that recovery plans, capabilities, and investments address genuinely material exposures rather than hypothetical technical failures.

Documented cyber recovery requirements should be included in the strategy to set out **clear, actionable expectations for the capabilities the firm must have in place to support recovery**, shaped by the firm's overall recovery strategy, the outcomes of recovery scoping (**Principle 3**), and application of relevant scenarios.

Such requirements typically span business, operational, and technical domains. They describe, among other things:

- the services and data that should be recoverable under the selected scenarios;
- the conditions and constraints under which recovery is intended to be achievable; and
- the safeguards and controls required to support execution of recovery activities.

To remain flexible and proportionate, cyber recovery requirements are usually framed in terms of '*what must be possible*' (e.g. recover prioritised services to a defined minimum level, operate manually for a period, restore from a verified clean state), rather than mandating specific technology solutions. This allows firms to apply the requirements across diverse environments while adapting implementation choices over time. This requires careful alignment to the evolving business architecture to remain accurate. To achieve this, recovery should be considered through the system development lifecycle.

1.3 Prioritisation and scoping

Principle 3: A scope for recovery is established that focuses on firm survivability and prioritises key services (or parts of services) for restoration following any disruption.

In scenarios that threaten firm viability, restoring all Important Business Services (IBS)[2] within their time-based IToLs is unlikely to be achievable. Firms can plan for survivability by defining a **minimum viable** scope to deliver critical outcomes while minimising intolerable harm, with explicit assumptions and limitations documented (including what will not be recovered immediately).

When determining a minimum viable scope, firms may find it useful to consider three aspects:

(i) Prioritisation should recognise that **not all IBSs are equal** in such circumstances: some IBSs may be more critical to the firm's survivability, and each IBS may only need to be delivered at a '*reduced*' capacity or functionality to meet its **minimum level of service**. Where time-based IToLs are likely to be breached, firms could place **greater emphasis on other IToL dimensions**, such as value, volume, subcategories of services (i.e. types of payments), and who is affected, to minimise intolerable harm while stabilising the firm and the wider system. In wholesale markets, for example, firms may elevate market-critical settlement flows ahead of broader client payment activity to contain contagion, while

[2] For Important Business Service (IBS) definition, see CMORG Guidance for Firm Operational Resilience v3.0 (Apr 2025)

deferring non-essential services. Minimum level of service is also **not a static concept**: firms should anticipate step changes over time (t_1, t_2, …, t_n) as recovery progresses, and re-baseline the minimum viable scope accordingly.

(ii) In addition to IBSs, stricken organisations are likely to depend on **internal enabling services** and elements of core infrastructure to support response and recovery. These may include identity and access management, communications tooling, key risk and finance platforms, and critical third-party services. Considering these dependencies explicitly helps avoid designing a recovery scope that cannot be implemented in practice.

(iii) A **minimum viable data set** could be identified and maintained containing the essential records required to perform the objectives defined in their recovery strategy. This data should be stored securely, isolated from primary systems, in a readily accessible, human-readable format where appropriate. The aim is that, even if core IT systems are unavailable, staff still have access to the information needed to continue vital business processes manually or via alternate means. Examples might include workaround data (CSV files, spreadsheets), reference data, key business records (such as database extracts or transaction logs), identity and access information (users/groups), and technical documents or artefacts such as architectural diagrams required to restore critical systems.

2 Cyber recovery capabilities

The cyber recovery capabilities defined in this section provide a guide to options for technical building blocks that can be used and combined to support a firm's recovery objectives and prioritised scoping. These capabilities articulate current industry thinking on the '*how to*' of recovery recognising that defining the '*what to*' must be performed prior.

This is a novel and developing area and it is not a simple or direct process to verifiably achieve the outcome of recovery from an incident that threatens the viability of the firm. The guidance therefore should not be interpreted as a mandated architecture; firms can choose to select and implement these or other capabilities.

2.1 Capabilities lifecycle

Principle 4: **Cyber recovery capabilities are designed, implemented, tested, operated, and maintained as an end-to-end lifecycle, ensuring they remain effective.**

Cyber recovery capabilities should be managed within a **coherent, end-to-end lifecycle**, ensuring that documented requirements are translated into executable, and sustainable arrangements in practice. A pivot to a recovery architecture is disruptive and, in some cases, may not be feasible to implement within the existing environment, e.g. prohibitive complexity or costs. The scale and the scope of change must be realistic and align with a firm's recovery strategy, dependencies and risk appetite (**Principle 2**).

Effective lifecycle management helps ensure recovery capabilities are not treated as static or one-off solutions, but are designed, operated, and maintained in a way that supports disciplined recovery execution over time. This includes ensuring that recovery arrangements remain aligned to the firm's operating model, dependencies, and governance structures, and continue to reflect changes in risk, technology, and operational complexity. The lifecycle of cyber recovery capabilities typically encompasses the following interconnected stages:

- **Design and solution architecture**: Recovery requirements are translated into end-to-end capability designs that account for system dependencies, operational constraints, and the conditions under which recovery would need to be executed.
- **Build and implementation**: Designed recovery capabilities are implemented in a controlled manner, with appropriate security, segregation, and integration into existing technology, processes, and controls.
- **Testing and validation**: Recovery capabilities are validated through a range of structured testing (**Principle 15**) to assess whether they meet documented requirements and can be executed as intended, with findings used to remediate gaps.
- **Operational readiness and execution**: The firm ensures that recovery capabilities can be invoked and managed effectively during an incident, supported by clear roles, decision-making authorities, recovery documentation, and escalation arrangements.
- **Maintenance, learning, and evolution**: Recovery capabilities are maintained and adapted over time to reflect changes in the firm's environment, threat landscape, the advent of new capabilities and lessons learned from testing, exercises, or actual incidents (**Principle 17**).

2.2 Operating model and processes

Principle 5: A formalised operating model and operating processes enable cyber recovery capabilities to be sustained in steady state during normal operations and executed effectively under disrupted conditions.

Where a firm is taking an internalised approach, cyber recovery is designed to be a '*last resort*' measure requiring a **specialised operating model** to continue to operate and perform recoveries in the likely absence of the general operating environment. The specialised operating model will likely contain mandatory deviations relative to the general operating environment given the need to continue operating *after* the disruption of standard services. A robust technical operating model should be designed and implemented to allow secure and predictable operational running of the environment in the event where standard processes and functions are not available. Any necessary deviations and compromises need to be reviewed carefully in the design phase.

> **Example**: Critical control areas like log file management, entitlements and authentication, privileged user management, encryption key management, end user device protection will all have enterprise capable service models within the standard operating environment. However, some or all of these capabilities may be compromised in the event of a successful cyber attack and therefore the operating model needs to be designed to allow these to operate within the recovery environment. This may be achieved via shadow services, which are isolated and provide the necessary capabilities to recover.

Cyber recovery arrangements are also expected to be supported by documented processes that provide clarity on responsibilities, sequencing, decision-making, and dependencies (**Section 3**). Operating processes are a critical enabler of cyber recovery because capabilities that are technically sound **may still fail if they cannot be operated, governed, and coordinated under stress**. Without such processes, recovery activities risk becoming fragmented, delayed, or inconsistent at precisely the point when speed and discipline matter most.

The importance of operating processes becomes most apparent during prolonged cyber incidents, where normal working assumptions may not hold. Staff availability may be reduced, escalation paths compressed, normal communication mechanisms compromised or inoperable and recovery timelines extended, requiring sustained execution over days or weeks. In these conditions, reliance on tacit knowledge or informal coordination is unlikely to be sufficient. Documented processes provide a **stable reference** that supports continuity of recovery activity even as conditions deteriorate or personnel change.

This principle also recognises the need for operating processes to support both **ongoing maintenance in steady state** and **execution during live recovery**. In steady state, processes underpin the upkeep, readiness, and assurance of recovery capabilities. During an incident, they support controlled activation, prioritisation, and coordination across technology, business, and governance functions. Together, these ensure that cyber recovery capabilities remain usable in practice, not merely well-designed in theory.

2.3 Data protection

Principle 6: Recovery data is protected to the level necessary to prevent interference, unauthorised access, and loss of integrity.

This principle encompasses preventive properties that keep recovery data safe from accidental or deliberate actions such as **immutability** (protection from interference or deletion) and **secrecy** (protection from unauthorised access). This also includes **preservation of integrity** for protection from unauthorised alteration that would undermine trust (**Principles 8, 13 and 16**). Taken together, these properties define what "protected" means for recovery data: it can be read when appropriately authorised, but it cannot be changed, exposed, or

corrupted in ways that would compromise restoration outcomes. The emphasis is on baseline protection objectives, not on mechanisms or processes.

The importance of this principle is twofold. First, it reduces the likelihood that a threat actor can simultaneously compromise production and recovery assets, thereby preserving the firm's ability to restore services when under stress. Second, it sustains stakeholder confidence, internal and external, that the data underpinning recovery retains confidentiality and structural trustworthiness, enabling accountable, timely decisions.

In circumstances where normal safeguards may be bypassed, this principle establishes a preventive stance whereby recovery assets are **resilient to change, resistant to exposure, and reliable in structure**. It sets the expectation that protection is not situational or ad hoc, but embedded and durable, so that confidence in recovery decisions does not depend on the status of production controls during an incident. When combined with the joint effect of enabling encryption, immutability, isolation, and data cleanliness provides a firm with confidence that data maintains referential integrity and remains trustworthy by design.

2.3.1 Immutability

Immutability can be used to ensure that data is protected from accidental or deliberate change for use in recovery activities, through use of one or more storage mechanisms that enforces 'write-one, read-many' (WORM) behaviour to create a form of **immutable storage**. This capability ensures that data necessary for recovery is stored and locked in such a way that it can be read as required (based on a robust entitlement model) but cannot be changed or deleted until its retention period has expired.

Immutability can be achieved through a variety of techniques, both individually and in combination (**Deep Dive: Cyber Recovery Vault** or CMORG Data Vaulting Reference Architecture[3] for more details). Determinations on which to use should take account of recovery throughput requirements:

- Enterprise disk or object storage platforms that enforce WORM semantics **at the firmware or platform level**, including SAN and NAS systems and enterprise tape libraries. Many vendors offer different levels of immutability, often summarised as '*governance mode*', where a privileged delete is possible, and '*compliance mode*' where a delete is not possible, even with privileges. Compliance mode is considered important to keep data fully protected even in the event of a directory services breach.
- Backup platforms that implement immutability **within the backup system**, preventing deletion or modification of backup data for a defined period.
- Object storage systems (on-prem or cloud-based) that enforce **immutability per object**, typically using retention and legal-hold constructs.
- **Storage snapshots or Point-in-Time (PIT) copies** that are locked against deletion or rollback for a defined retention period.
- Physical media (tape, optical) that is inherently immutable once written and **physically isolated** from online environments.
- Data stores where immutability is enforced through **cryptographic chaining**, hashing, or append-only logs.
- **Data escrow** where recovery-critical data is stored with independent third parties under contractual immutability and access controls.

[3] CMORG (2022), Data Vaulting Reference Architecture v1.0, April

2.3.2 Secrecy

Encryption can be used to ensure that data is protected from unauthorised viewing **even after access management processes have been bypassed**. In principle, where encryption would materially reduce overall risk based on the data being transmitted and processed it should be encrypted, and key management processes made highly robust.

Encryption can be enabled at multiple levels, including '*at rest*' in the storage layer (protecting against theft of media), '*in flight*' (protecting against reading of packets of data being moved) and critically within data objects themselves (protecting against unauthorised reading of content once written). Key management should be enterprise-level and highly protected to avoid keys being made available to unauthorised actors, and to ensure that legitimate users are not locked out of legitimate access. Key vaults should be considered as critical infrastructure and protected through use of immutable storage, with frequent data synchronisation points to avoid having encrypted backups available for recovery that are invalidated by a lack of keys. Key management may be a candidate for critical accelerated data reporting since having fully up to date keys minimises avoidable data loss.

Example: Encryption of data at rest can be provided through storage-level capabilities, including controller-firmware-based solutions and self-encrypting drives. Encryption in flight can be enabled through secure tunnelling mechanisms and protocol-based encryption features. There are also various methods for encrypting the data itself, such as database-level transparent encryption, tokenisation, and other cryptographic approaches.

Deep Dive: Cyber Recovery Vault (CRV)

Cyber recovery requires trusted, uncompromised data spanning business, infrastructure, and system artefacts so that recovery, from enabling manual workarounds to full rebuilds, can be executed with confidence. A **Cyber Recovery Vault** is one method of enabling this. It should ingest predetermined data types, retain multiple historical copies to mitigate long-term stealth tampering, and use strong integrity controls to verify that data has not been altered. The vault should also detect attempts to ingest damaged or suspicious data using mechanisms such as entropy checks and reconciliation. Because it concentrates sensitive material, robust encryption and hardened cyber defences are essential.

A CRV should be scalable to ingest large datasets (such as multi-terabyte database files) with sufficient throughput to meet business needs, such as daily or end-of-day loads, or the ability to support peak demand. It should also remain reachable when needed, e.g. cloud-hosted vaults may be accessible when corporate networks are compromised, though the reverse may also be true.

A firm's appetite for data loss may be defined using Recovery Point Objectives (RPO) in a mechanical way, typically from 0 (no data loss) to hours. In an operating CRV, the amount of data loss should clearly relate to: (i) the impact on the service(s) of missing data (e.g. missing transactions or payments); and (ii) the ability to recover the system back to normal operation (e.g. a complete database backup file, complete end-of-day batch files). Depending on the focus of the use of the CRV, one or both goals should be considered.

Vault implementations vary: some use storage-level logical air-gapping (e.g., disk snapshots), segregated backup solutions, tape-based architectures, or data-diode-protected ingestion paths. Typical content includes immutable identity directory data, database files and redo logs, binaries, system images, and libraries. Ideally, all are retained over extended periods so firms can revert to a clean pre-incident state aligned with the vault design.

2.4 Isolation

Principle 7: Exposure to disruptive events is minimised through approaches that distance or separate critical recovery assets from resultant impacts.

Isolation refers to the **deliberate separation** of critical recovery assets from production and other shared environments to reduce the risk of simultaneous compromise during a cyber incident. Isolation can be achieved through a range of approaches, each designed to minimise risk while supporting a safe and effective recovery, as shown in the examples below:

- segregation of **critical payment/market connectivity stacks** (e.g., SWIFT) from general IT, with strict internet restrictions and tiered admin paths, to enable staged resumption and safer reconnection;
- **one-way transfer mechanisms** (data diodes) to export telemetry from sensitive domains to SOC tooling without creating inbound paths, maintaining visibility during recovery;
- **digital segregation patterns** for cloud-hosted vaults using ephemeral ingress/egress zones and independent management planes, minimising standing connectivity; or
- backup/replication over **isolated networks with severable links** and WORM/object-lock, ensuring last-resort copies can't be reached from compromised enterprise segments.

Deep Dive: Isolated Environment

An isolated environment can be used individually or jointly for testing, forensics, inoculation, or recovery purposes. It provides a **critical layer of protection** for restoring and verifying systems or business processes either before or independently of disruptive events. By being logically segregated from production, it reduces the risk that live operations and recovery assets are compromised together. This separation enables controlled restoration and verification and offers a trusted setting for recovery activities without further exposure.

When appropriately scoped, it can host **carefully selected critical business and technical services** needed to support recovery. This allows firms to validate data, processes, and system integrity ahead of an incident or during live recovery, strengthening confidence in outcomes. In recovery-vault models, the environment may hold pre-positioned data, (recovery-)media, and supporting infrastructure, enabling rapid recovery while maintaining strong isolation. Firms may also run a production service from this environment, potentially making it the new production environment, requiring a transition from isolated to connected.

To remain effective, the environment should avoid **inherent dependencies on production or external shared services**, such as identity, entitlement, time, or other control services, that could undermine isolation. Isolation typically includes network segregation, separate compute and storage, and discrete operational, vaulting, and recovery services. Limited connectivity for data transport should be tightly secured and controlled from inside the environment, minimising duration and pathways of access. Physical isolation, such as removing offline backup media, can further protect against overwrite or compromise.

Isolated environments should be **rapidly deployable, usable, and rebuildable**, allowing applications and data to be safely reconstructed without endangering golden backups or production. They can include recovery tooling aligned with the firm's technical stack, mandatory services (AD, LDAP), anomaly detection, and compute and storage to build, test and validate recovered environments. Capacity planning should reflect the firm's recovery approach including the use of multiple isolated environments to move applications through testing toward production readiness, with use varying between crisis and non-crisis needs.

Isolation may be achieved not only within the firm, but through deliberate reliance on independent third parties, sector utilities, or external service providers, where separation of governance, identity, infrastructure, and assurance materially reduces correlated cyber risk and supports trusted recovery. Examples include:

- legally and operationally independent **third-party escrow or vaulting** of recovery-critical data;
- **sector utilities and FMIs** as stable external reference points;
- temporary **stand-in or surrogate** service provision by third parties; or
- **segregated cloud tenants or providers** for recovery-critical assets.

2.5 Trusted recoverability

Principle 8: Data used for recovery is subject to controls that detect tampering and confirm its correctness, accuracy, and readiness for use.

Trusted recoverability is concerned with **establishing confidence that recovery data can be relied upon in practice**, rather than assuming that protective measures alone are sufficient. Firms should be able to determine whether data intended for recovery remains trustworthy, complete, and usable, particularly where adversarial activity may be subtle, delayed, or designed to evade preventive controls. This principle therefore focuses on the ability to detect loss of trust and to confirm fitness for use before recovery decisions are taken.

Trusted recoverability relies on the **timely detection of data tampering or corruption** across the recovery lifecycle, ensuring that any unauthorised modification, contamination, or manipulation is identified before it can contaminate restoration processes. By maintaining strong visibility of unexpected changes and patterns affecting recovery assets, firms reduce the risk of restoring compromised data and undermining subsequent recovery efforts. This includes detection at every stage (pre-ingestion, post-ingestion, and during recovery) supported by intelligent scanning and learning processes capable of identifying unusual content or deviations in change volumes. Importantly, the focus is on detecting signals that something is wrong, rather than the specific tools used to produce those signals. These controls may operate within backup tools, the network data plane, or the vault itself, with recursive scanning helping pinpoint when and where compromised data entered the environment.

> **Example**: Many backup and storage platforms now include native anomaly-detection capabilities as part of their baseline functionality designed to identify unusual patterns or potentially malicious changes within data. These capabilities may operate during data ingestion, within the storage environment itself, or as part of recovery-time validation.

Recovery depends not only on data being protected, but on it being **fit for purpose under stress**. Without confidence in data quality and usability, recovery execution may be delayed, partial, or fail entirely. Controls that establish data cleanliness and readiness are therefore critical for timely, accountable recovery outcomes.

Proactive data cleanliness checking provides confidence that the data held in the vault has the necessary integrity to facilitate recovery. If the data in the vault is corrupted, recovery may be impossible. While multiple levels of programmatic anomaly detection may be employed, looking for suspect binaries, change rates, etc, attacks which corrupt data by flipping values, adding or removing key records and other subtle methods will not be detected. Data centric data validation and testing for unexpected change may be employed to ensure that data retains necessary integrity.

> **Examples**: Synthetic transactions, data checksums, entropy detection, canary data, reconciliation.

2.6 Interim contingency measures

Principle 9: **Pre-planned contingency arrangements are established and maintained that allow critical functions and obligations to be fulfilled whilst production and recovery systems are unavailable, minimising intolerable harm while recovery is underway.**

Standard recovery processes, especially those involving recovery of foundational infrastructure services, are lengthy and non-deterministic. Whilst recovering to a pre-defined recovery scope is an urgent priority, the functions listed below are typically those that may need to operate before any initial technical recovery succeeds, often without enterprise identity services (e.g. AD), core infrastructure, standard collaboration tooling, and potentially without confidence in internal data integrity.

Requirements

- **Executive command and governance**, to establish authority and decision-making immediately to operate command and control functions and invoke recovery processes.
- **Situational awareness and investigative support**, to scope impacts, determine incident blast radius, inform containment decisions, and undertake forensic activities.

Capabilities

- **Secure out-of-band communications**, to enable coordination when normal channels are unavailable or untrusted.
- **Interim business contingencies**, to support manual or alternate processing, surrogate capabilities, and business decision making (e.g. exposure tracking, risk limits, financial position estimation).
- **Recovery orchestration and assurance**, to allow recovery to begin safely before systems are restored.

Processes

- **Workforce access and welfare**, to allow the firm to operate humanly, even if digitally constrained, in order to maintain staff continuity (e.g. payroll functions), provision emergency credentials, or issue safety communications.
- **Regulatory and market obligations**, to fulfil time-critical obligations that cannot wait for recovery such as liquidity, capital, prudential or market reports, incident notifications, and stakeholder engagement including customers.
- **Evidence and decision recording**, to facilitate accountability and traceability through methods of record-keeping, and learning through (early) lessons capture.

In order to meet time-critical demands, an interim capability should be pre-planned, tested, and capable of operating out-of-band or isolated from the production environment. Such a capability is frequently fed with necessary data, to a pre-determined level of detail and scope, stored in a highly secure logical location which will be quickly accessible during the disruptive period and recovered without access to institutional standard services.

> **Example**: Data may be stored in alternative public cloud providers or secure escrow parties, in easily accessible and widely supported formats (such as CSV files or spreadsheets) to allow instant access. Time-series or event-streaming mechanisms may be employed to transfer data securely. This capability leverages other core properties such as immutability (**Principle 6**) and isolation (**Principle 7**).

For collaboration tooling, most firms rely on integrated IP / application-based communication and chat tools which may be susceptible to disruption, especially if the communication processes are hosted on premise. An alternative method of communication will need to be established very quickly, capable of running on unaffected devices, with an alternative directory service so that phone numbers and other contact information are available. The high levels of reliance placed on peer-to-peer and peer-to-group collaboration tools may also be affected and will need to be re-established near immediately to support coordination of response. Pre-planning and testing will be necessary for even simplistic communications to be established, with agreed levels of compliance with standard message retention and recording requirements, as well distribution of key contacts to locations which are isolated from general production (e.g. cannot use primary enterprise document repository to hold this content).

Example: A backup collaboration and productivity setup on a separate domain provides email, messaging, calls, video conferencing, and office tools for essential staff. Leadership can quickly coordinate and share updates. Additionally, an out-of-band notification system ensures crisis teams can send secure, trackable instructions to personnel via SMS, calls, or email.

3 Recovery sequencing

Recovery sequencing considers the ordering and dependency management required to safely re-establish the recovery scope defined in the recovery strategy (**Principle 3**). This requires a clear understanding of the capabilities required to achieve that objective, including consideration of both response (e.g. interim contingency measures) and recovery activities.

The approach should consider the non-deterministic nature of cyber disruptions and be sufficiently adaptable for prevailing conditions at point of recovery that may influence prioritisation decisions i.e. weekday vs. weekend, end of month, market conditions, etc.

All aspects of recovering sequencing scope described in this section should be rehearsed and/or exercised to gain confidence in their effectiveness (**Principle 15**).

3.1 Invocation and orchestration

Principle 10: Clear conditions, authority, and processes exist for invoking cyber recovery, enabling timely transition from incident response to recovery when predefined thresholds are met.

BAU incident management processes may not be effective in managing recovery from cyber incidents that threaten a firm's viability. Consideration should be given to the invocation and orchestration needed to manage response and recovery activities, including:

- Any specific preconditions required to trigger recovery invocation.

  Example: Pre agreed factors to demonstrate sufficient cyber containment has been achieved to enable recovery to proceed safely.

- Pre agreed decision making and escalation protocols to aid the timeliness and effectiveness of recovery activities. This should include clarity on the role of senior management/board during response and recovery.

  Examples: Who is authorised to make recovery prioritisation decisions, how will senior management be engaged and what is expected of board members during recovery.

- Processes required to support the invocation and operation of response and recovery activities.

  Examples: Communication protocols, interim contingency measures, managing third party dependencies, and achieving forensic requirements to demonstrate recovery assurance.

- Agreed exit criteria to resume post recovery operation in line with agreed cyber recovery strategy. This should include defining how those criteria will be demonstrated.

3.2 Tiered recovery

Principle 11: Recovery is executed in a prioritised sequence based on criticality, dependencies, and feasibility, balancing customer harm, market impact, and system stability as services are progressively restored.

Recovery should be prioritised in tiers based on business impact, e.g. IBS ITOLs, and aligned with the scope defined in the recovery strategy.

Tiered recovery is a structured approach that restores business services in a deliberate, prioritised sequence. It focuses first on the services most critical to firm viability, taking into account inter-service dependencies, customer impact, market stability, and the practical feasibility of restoration.

By progressively bringing services back online in controlled tiers, the organisation ensures that recovery is both safe and sustainable. A disciplined sequencing enables leadership to balance speed with assurance, reducing harm and restoring confidence as the firm transitions from crisis to stability.

As shown in **Figure 2**, beyond business prioritisation, sequencing also needs to consider technology dependencies for both invocation of any interim contingency measures and service recovery:

1. Adequate **protection** capabilities need to be deployed to reduce the risk/blast radius of cyber disruption and support recovery capabilities (**Section 2**).
2-4. The restoration of management control planes (e.g. authentication, security), **communication / collaboration capabilities and core infrastructure** will precede restoration of **business applications and data**.
5. Providing **security assurance** sufficient to give **third parties confidence to reconnect**.
6. The effectiveness of business **response plans and interim contingency measures** to provide a level of mitigation whilst recovery is underway.
7. Consideration of **data repair and reconciliation**.
8. The ability to **safely reintroduce recovered services** and exit contingencies / workarounds.

3.2.1 Desktops, laptops and collaboration services

Essential workers need to be able to communicate and collaborate effectively to orchestrate response and recovery. BAU capabilities may be impacted by the event, therefore consideration should be given to:

- The sequencing of any contingency communication and collaboration capabilities required to support response and recovery activities.
- The capabilities that colleagues involved in response and recovery activities will require, the potential for those capabilities to be impacted by the event, and what contingency capabilities are required to mitigate (**Principle 9**).

> **Examples**: End user devices, access credentials, documentation e.g. playbook access, and tools required for response and recovery.

Figure 2. Recovery sequencing workflow

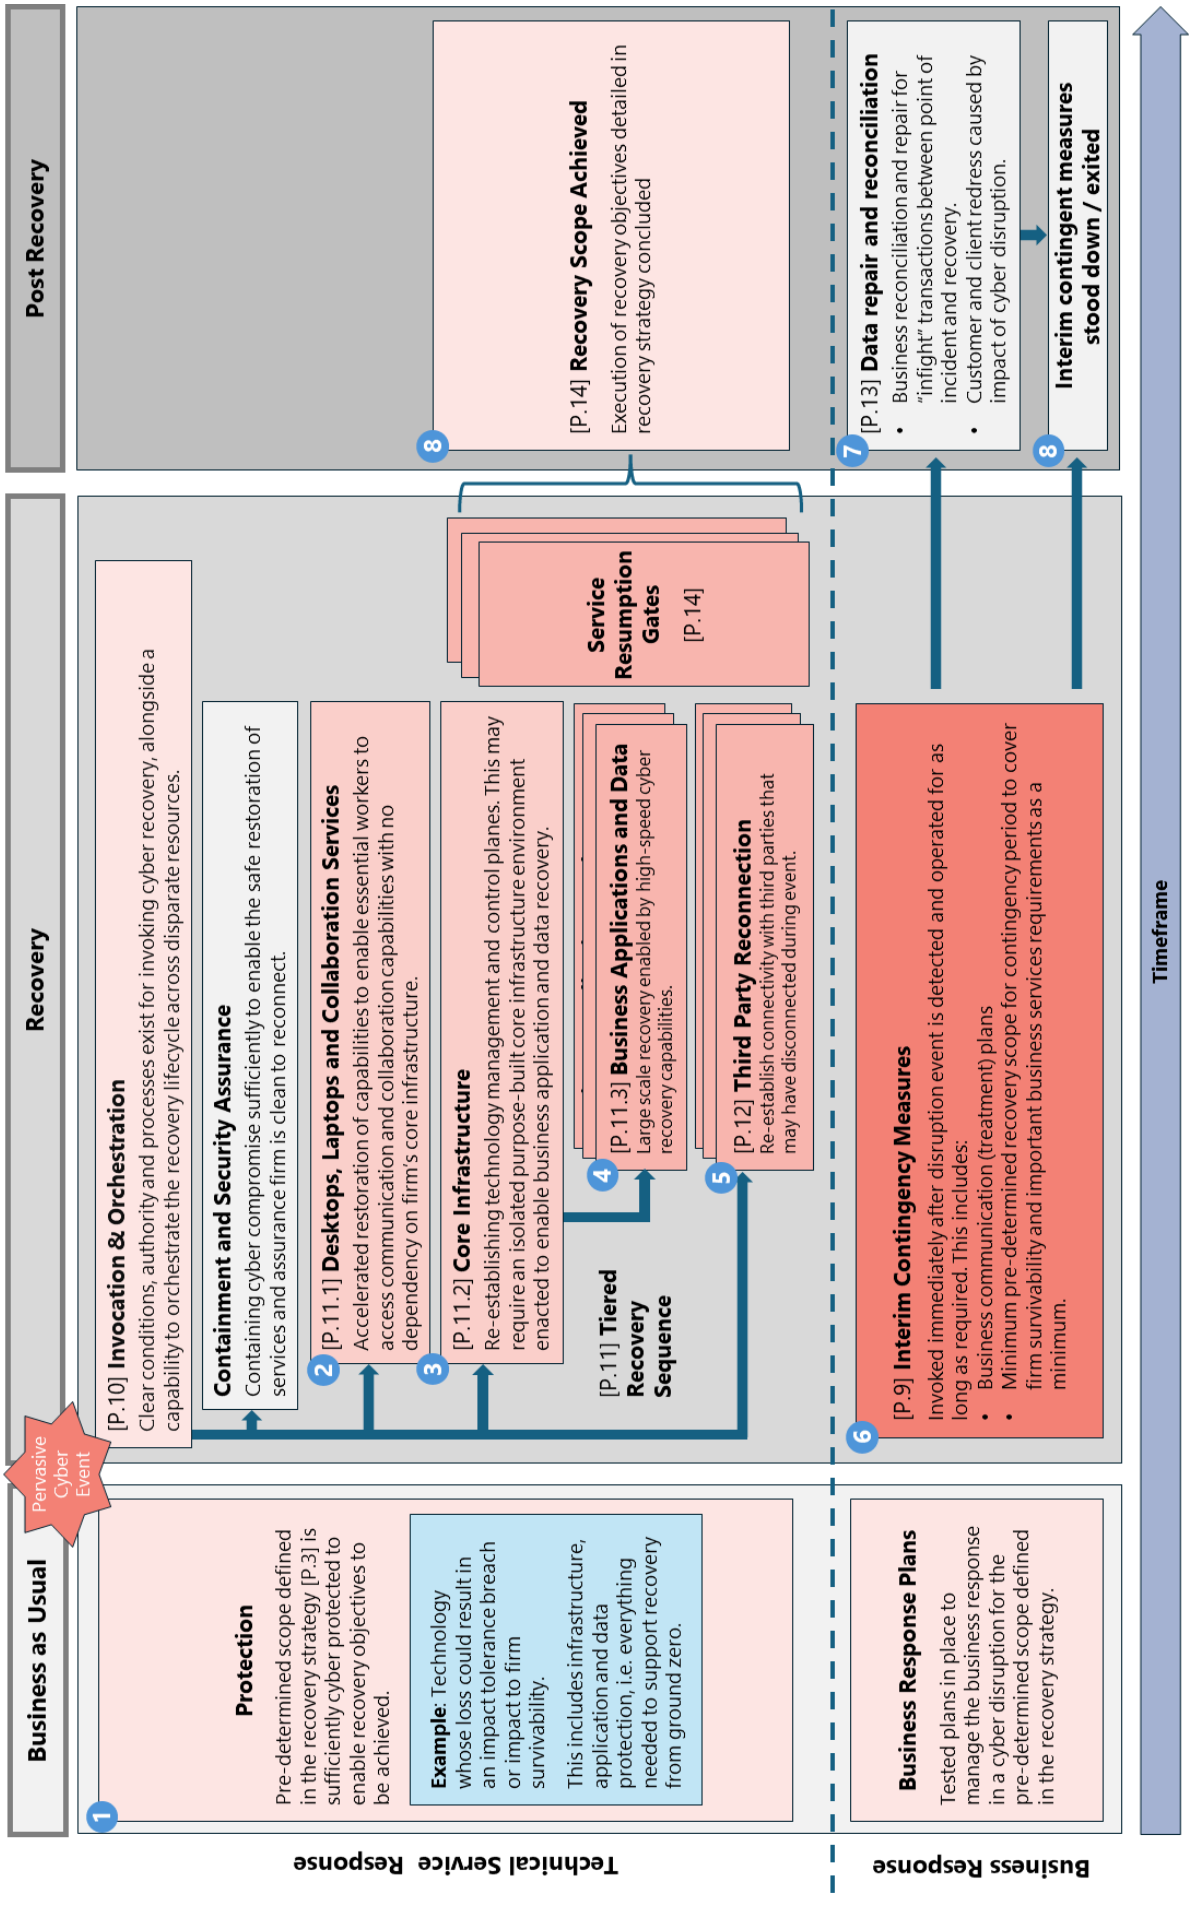

3.2.2 Core infrastructure

Core infrastructure is a foundational enabler for the operation of business services. These infrastructure services could be compromised in the event, and so need to be considered within the recovery sequencing:

- Firms need to consider core infrastructure services required to support both the recovery and operation of services in scope, and the recovery sequencing and interdependencies required to re-instantiate those services.

  > **Examples**: Network services, directory services such as AD, privileged access management solutions, security tooling, virtual machine/container hosting platforms, software build and release capabilities.

- Sequencing should consider potentially circular dependencies at the core infrastructure layer e.g. recovery of DNS is dependent on directory services, but those directory services depend on DNS to operate. Where firms are unable to gain confidence in mitigating those dependencies, they may need to consider alternative recovery approaches e.g. tertiary isolated environments (**Principle 7**).

3.2.3 Business applications and data

Technology services may need to be recovered at scale from bare metal/tertiary environments including infrastructure, application, middleware and data layers:

- Recovery sequencing should consider the scope and interdependence of business applications and data required to achieve the recovery strategy.

  > **Examples**: Supporting technology services that may not be directly consumed by business users/services e.g. middleware/messaging utilities. Management/control planes beyond core infrastructure needed to support the ongoing operation of recovered services such as batch processing/housekeeping requirements.

- Consideration should also be given to understanding any constraints in achieving the recovery strategy, both technology constraints (e.g. throughput) and nontechnology constraints (e.g. resources and processes required to execute recovery).
- Data volumes and reconciliation between systems is a key consideration, larger more complex systems will take longer to recover and reintegrate.
- Recovery sequencing should consider any post recovery technical and business repair/validation steps required to provide confidence to re-enable services e.g. data integrity and reconciliation across the service chain.

3.3 Disconnection and reconnection

Principle 12: Suppliers and counterparties that are critical to service recovery and/or invoking contingent processes are managed through pre-agreed disconnection, assurance, and reconnection arrangements that support safe recovery.

Recovery sequencing should consider third party dependencies to support recovery, and those third parties' response to the firm experiencing a cyber incident:

- The nature of the third-party relationship should be considered. The approach may differ between third party counterparties and suppliers, and with the context of the third-party arrangement.

> **Example**: The approach may be different between a counterparty in the context of hedge fund positions versus reconciliation with a payment FMI.

- Planning should also consider partial disconnection of high-risk connections, degradation of service of operation versus total loss, and the potential for a coordinated sector response where there is industry concentration or a reliance on capabilities from other firms.
- Seek to have pre agreed disconnection/reconnection principles with critical third parties. The CMORG Reconnection Framework[4] may be helpful in providing further guidance for firm's consideration.

> **Examples**: Considerations could include threat to the third party based on connectivity/data exchange, the security assurance required to enable reconnection in the event of disconnection or service closure, that reconnection may be iterative through the recovery lifecycle and agreeing incident management and communication protocols.

- Consider third party dependencies that may not be critical in normal operation but could be in a resilience stress scenario e.g. SaaS hosted code repos.
- Specific consideration should be given to any counterparty dependencies required to support interim contingency measures given limited real-world experience in establishing trust in an alternative processing environment.

3.4 Data repair and reconciliation

Principle 13: Data and system states are reconciled and validated during and after recovery to restore accuracy, completeness, and confidence before resumption and reconnection.

Recovery sequencing should consider managing backlog, controlled replay, and reconciliation activities necessary to safely resume operations. The approach should consider:

- Repair and reconciliation actions may need to be iterative, with early value lying in decision support and extending to enabling 'limited service' resumption prior to achieving the full recovery over time.
- The scope of data that needs to be repaired and reconciled, recognising that some data will be less relevant based on duration of outage and market movement in the intervening period.
- The sources of data required to perform reconciliation e.g. leveraging FMI, custodian and counterparty positions to reconcile internal positions.
- Recovered systems are likely to return at different recovery points, effective reconciliation requires handling of misaligned recovery points rather than assuming uniform restoration across the estate.
- The processes and capabilities needed to perform that repair and reconciliation, and the availability of those capabilities in a cyber disruption, including personnel, access to external data sources and tooling capabilities.
- The capacity required to manage reconciliation in a cyber recovery scenario, noting it is likely to be more material than in BAU operation.
- Activity required to reflect transactions processed through interim contingencies back into recovered technology services e.g. balance and transaction reporting.

[4] CMORG (2025), Reconnection Framework v3.0, July

3.5 Service resumption

Principle 14: Clear protocols are in place to support the safe reintroduction and operation of recovered services, standing down any contingent processes invoked during recovery and the approach to manage customer reconnection.

Real world experiences indicate there is often a lag from services being recovered to consumers gaining confidence in resuming their use. Recovery sequencing should anticipate the post recovery activity required to support the safe reintroduction of recovered services, including:

- Approach to re-establish trust with customers, clients, counterparties, third parties and regulators to support timely transition to resumption of service post recovery.
- The levels of assurance and practices that may be employed to give senior management confidence to agree service resumption. Consider framing resumption decisions around bounded risk, stated assumptions, and known unknowns, rather than implying certainty or completeness.
- Resumption may need to be phased progressively across a series of defined states, rather than a binary go/no go decision, to balance residual risk with the levels of compensating controls available.

  Example: Early service resumption may rely less on typical preventative controls and more on intensified monitoring and rapid rollback capability. Heightened detection and conditional operating thresholds may provide safeguards while residual uncertainty remains.

- Good practices described in reconnection frameworks including emphasis on reconciliation, ring-fencing inflight items, end-to-end testing, and heightened monitoring as part of safe resumption.
- How any contingencies deployed would be stood down post-recovery, including any reconciliation or repair activity.

4 Assurance and continuous improvement

This section sets out a high-level approach for supporting financial institution's in assuring their cyber recovery capabilities, both in terms of testing and validating those capabilities before an incident crystallises (pre-incident) and in recovering effectively while an incident is ongoing or being closed out (during/post-incident). It also sets out the complementary objective of embedding continuous improvement, ensuring that insights from testing, incidents, the evolving threat landscape, business-led changes, and technological advancements are systematically used to strengthen recovery arrangements over time. It acknowledges that testing cannot replicate a viability-threatening cyber scenario in full. Full end-to-end testing of all systems, data, and dependencies at enterprise scale is generally not feasible. Instead, effective assurance relies on progressive, layered validation, evidence building, and controlled testing that reflects real constraints while still supporting decision makers with credible and validated insight.

The following considerations are relevant throughout to inform testing programmes and firms' approaches to build and maintain credible, evidence-based assurance in their ability to respond and recover from cyber incidents that could threaten firm viability.

- **Align assurance with recovery sequencing:** Assurance activities should map to the firm's end-to-end recovery model whether involving invocation, containment, interim workarounds, rebuild, or full restoration. Testing and validation should assess each stage individually and the transitions between them, providing confidence in both discrete steps and overall recovery flow
- **Measuring success:** All assurance should be grounded in predefined and transparent success measures. These could include time-based objectives and functional outcomes, enabling the organisation to evaluate the effectiveness of each test and build an aggregated, evidence-based view of recovery readiness.
- **Consider customer or member impacts throughout:** Across all phases, assurance should assess how and when customers or members may need to act e.g. disconnecting, authenticating, or reconnecting. Proactive communication planning and user-centric testing help maintain trust and ensure that customer actions support, rather than hinder, recovery.
- **Maintain defensible evidence and documentation:** All assurance activities including assumptions, readiness checks, execution steps, and outcomes should be documented in a manner that meets regulatory, board, and internal audit expectations. Evidence should support both immediate decision-making and retrospective review.

4.1 Pre-incident testing

Principle 15: **Pre-incident testing provides confidence that cyber recovery requirements can be met, and validate that recovery capabilities, processes, and data safeguards are ready to operate under severe stress.**

This phase should be grounded in rigorous, realistic preparation that acknowledges the increasing severity of the threat environment. Testing should be calibrated using relevant scenarios that assume disruption has occurred, ensuring recovery plans are validated under conditions that mirror the pressure and uncertainty of a real incident. While full end-to-end, enterprise-scale replication is generally not feasible, the range of testing techniques set out in this section can be used to build progressive, layered validation of a firm's preparedness.

Testing should involve broad organisational participation spanning business, technology, security, operations, and governance teams to reflect the multidisciplinary nature of genuine crisis response. By challenging

assumptions, exposing hidden dependencies, and rehearsing decision-making across the organisation, testing helps build the confidence and muscle memory required to execute recovery effectively when it matters most.

Figure 3 sets out the spectrum of testing approaches that can be employed depending on the firm's cyber maturity and sophistication of its existing testing capabilities.

Figure 3. Types of testing

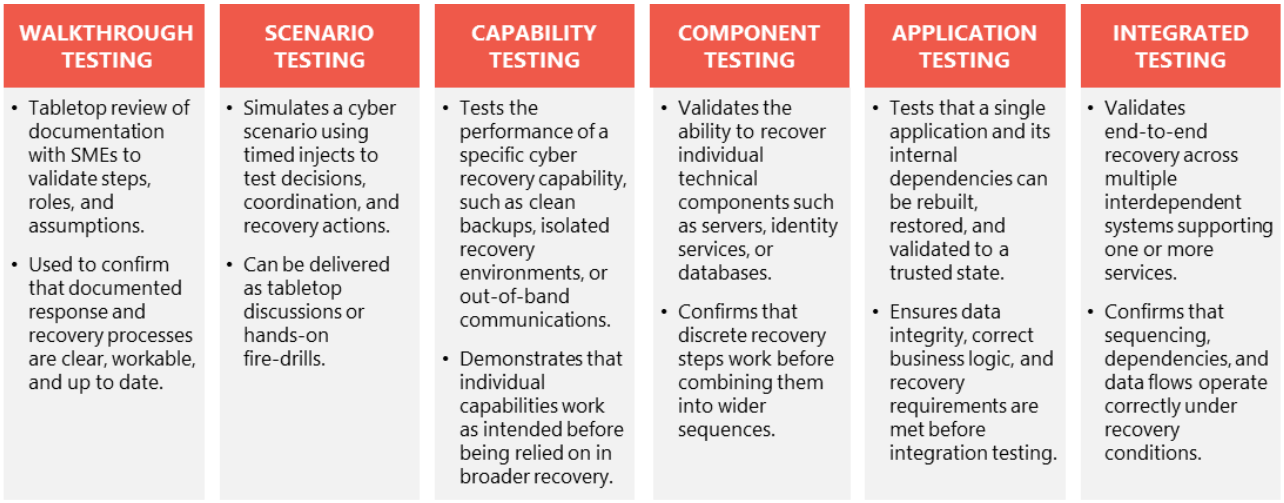

4.1.1 Walkthrough testing

Walkthrough testing centres around **tabletop validation of cyber recovery documentation** and examining the unique aspects of responding to a cyber incident. This testing has an overall goal of increasing confidence that plans will work as intended in a real cyber incident and that the right SMEs are clear on their roles. Walkthrough testing will typically centre on a set of recovery documentation aligned to the firm's cyber recovery capabilities and recovery sequencing.

4.1.2 Scenario testing

Scenario testing **simulates a cyber scenario through a series of timed injects** to validate recovery processes, while also exercising key elements of the response processes where relevant. Cyber scenario tests may be delivered as **tabletop exercises**, which use structured discussion and predefined questions to explore decision-making and actions, or as **fire drills**, which rehearse documented response and recovery plans in a more operational, hands-on format. In addition to identifying vulnerabilities and validating roles, responsibilities, and escalation paths, cyber scenario testing can be used to assess specific cyber recovery capabilities, the provision of services within recovery scope, the execution of communication protocols, plans, and broader opportunities for enhancing contingency arrangements.

While participation from both business and technical teams is expected, cyber scenario testing can also serve as a controlled environment to involve customers, members, or third parties, particularly where disconnection and reconnection processes are in scope of the scenario test.

Example: A Board-level tabletop examines unique aspects of responding to a targeted cyber attack including strategic decisions when dealing with a high degree of disruption. The facilitators inject multiple scenario changes and developments, focused on recovery elements, that present Board members with a ransomware scenario resulting in operational disruption to the firm's IBSs, e.g. decision-making process on invoking alternative isolated environments.

4.1.3 Capability testing

Capability testing validates the **performance of individual cyber recovery capabilities** that underpin the firm's approach to cyber recovery (as detailed in **Section 2**), and that these capabilities work as intended before being relied upon as part of the firm's overall recovery. Effective capability testing should consider: (i) documented process for the capability being tested; (ii) availability of a suitable test or production environment; (iii) relevant teams to undertake the testing; and (iv) pre-requisites in place.

> **Example**: Testing undertaken will vary depending on the type of capability, but examples might include penetration testing of a Cyber Recovery Vault or Isolated Environment, Red and Purple Team exercises, response-team testing focused on communicating solely through out-of-band channels, and validation of interim contingency measures by technical and business teams, including external parties where relevant.

4.1.4 Component testing

Component testing demonstrates the **ability to recover the separate IT components of an application through discrete testing**. Component testing involves validation of the discrete recovery processes in test or production environments. This should include testing of the individual IT components and data from the full sequence of recovery (e.g. server, identity services, database, etc.). For the effective testing of an individual component, its role and position in the full sequence of recovery, including any pre-requisites, should be understood.

Effective component testing should consider: (i) documented recovery process for the IT component being tested; (ii) availability of a test environment that sufficiently mirrors the production environment; (iii) relevant technical teams to undertake the testing; and (iv) pre-requisites in place.

> **Example**: Testing undertaken by a database support team to demonstrate that they can successfully restore a database and data from an immutable backup to a test environment. The success criteria of this test are the assurance on the integrity of the data recovered.

4.1.5 Application testing

Application testing validates that **individual applications** including their data, configuration, and internal dependencies **can be recovered to a trusted state** following a cyber incident. It demonstrates that the application can be rebuilt, restored, and functionally verified in isolation before being incorporated into broader integrated or end-to-end recovery testing.

This type of testing focuses on recovering the application and its immediate prerequisites (e.g., configuration baselines, application-level data, service accounts, certificates, and internal API interactions) and confirming that it operates as expected within a representative test environment. It bridges the gap between capability testing, which assesses discrete technical components, and integrated testing, which validates recovery sequencing across multiple applications.

A typical approach includes preparing recovery artefacts and documentation, deploying the application into a representative non-production environment, executing the recovery steps, verifying functionality and data integrity, and capturing outcomes to support remediation. Effective application testing should build on earlier testing techniques by using recovery documentation, trusted deployment artefacts, representative environments, clean validated data, clear success criteria, and engagement from relevant technical teams. Evidence should be captured in an audit-ready format to support internal and regulatory assurance.

4.1.6 Integrated testing

Integrated testing validates the **coordinated recovery of multiple interdependent systems and applications**, confirming that individual recovery capabilities can be executed in anger, **under severe stress, and** in the correct **recovery sequencing**. While full end-to-end replication of all systems, data, and dependencies at enterprise scale is generally not feasible, integrated testing together with the other techniques in this section enables firms to build progressive, layered validation of their preparedness and their ability to restore to pre-determined service levels and progress toward full recovery.

Rather than assessing components in isolation, this form of testing verifies how capability and application and component-level testing by verifying whether these work together in the right order, respecting dependencies and data flows to support the firm's cyber recovery strategy. Where full end-to-end execution cannot be achieved in a single exercise, firms may test combinations of tests in stages and consolidate results to build a holistic view of recovery performance and timelines.

Effective integrated testing should build on earlier testing techniques, drawing on validated component-level capabilities, confirmed dependencies, and representative environments to exercise end-to-end recovery in a realistic and controlled manner. It should also bring together the relevant business and technical teams and ensure that clear, auditable evidence is captured to support internal and regulatory assurance.

4.2 During incident

Principle 16: **Assurance during an incident provides timely and credible confidence that recovery actions are mitigating the incident effectively and that services can be safely resumed, balancing speed, safety, and stakeholder needs without overstating certainty.**

During a cyber incident, assurance activities focus on providing rapid, evidence-based confidence that recovery actions are working as intended. Unlike pre-incident preparation, assurance during an incident should be promptly validate transitions between recovery sequencing phases, confirming containment, integrity, and trustworthiness and incorporating where appropriate the use of third parties for validation and surge capacity.

4.2.1 Data integrity

Data integrity assurance validates that **data has not been corrupted, tampered with, or lost** as a result of a cyber incident, providing confidence that recovered systems can be trusted before they are brought back online. During a cyber incident, particularly one involving ransomware, destructive malware, or data poisoning (where threat actors subtly corrupt data over time), there is a risk that data has been modified, encrypted, or deleted. Before restoring services, firms need to demonstrate that the data underpinning those services is accurate, complete, and has not been compromised.

In circumstances that are expected to exceed the time-based aspects of a firm's IToLs, the focus shifts to validating data integrity of services within recovery scope, followed by full recovery in line with defined targets. This may involve one or more IBSs and could include validation following a full bare metal restore.

Data integrity validation typically involves comparing recovered data against known-good baselines such as pre-incident snapshots, reconciliation markers, or external records; verifying files and databases through hash and checksum checks to detect unauthorised modification; confirming transactional completeness, reconciled balances, and intact referential integrity; and working with incident response and forensic teams to understand the scope of compromise and determine which data sets were affected or remain trustworthy. The depth of integrity validation will depend on the nature of the incident and the criticality of the data. For high-value or regulated data, independent verification or third-party attestation may be appropriate

Example: Following an incident affecting a firm's payments processing environment, the recovery team restores the core payments database from immutable backup. The immediate focus is on validating data integrity to support delivery of a minimum level of service for the affected payments services. Before reconnecting to upstream and downstream systems, the team validates data integrity by comparing account balances against the previous end-of-day reconciliation, verifying transaction logs against a reference copy held by a third-party processor, and running checksum validation on critical database tables. Findings are documented and approved by the appropriate authority before services are resumed. Full data integrity validation across all affected systems follows as part of the recovery to defined targets.

4.2.2 External assurance

In a significant cyber incident, firms often need **independent, third-party validation** to support internal actions, reconnect decisions, regulatory notifications, and communications with customers, counterparties, and markets. External assurance is not about providing absolute certainty, which is unattainable, but about delivering credible, evidence-based negative assurance from a trusted external firm. This gives stakeholders confidence that the actions taken are sufficient, the environment is understood, and appropriate monitoring remains in place.

External assurance should be structured, bounded, and defensible. Providers should help articulate what has been done, what has been observed, and what safeguards remain active, without implying unattainable guarantees such as the complete removal of all threats. Effective assurance begins with clear scoping by defining the systems, applications, identities, data sets, and network boundaries under review, establishing the time period the assessment applies to, and confirming that the chosen scope is appropriate for the reconnection decision.

Assurance should also follow a layer-plus-one approach, extending assessment beyond the compromised layer to adjacent layers such as hosting, identity, and neighbouring network segments, reducing the risk of hidden persistence or lateral movement. Findings must be evidence-based, drawing on reviewed logs, completed forensic analysis, integrity checks, reconciliations against known-good sources, adjusted detections, updated monitoring, and confirmation that compensating controls are effective. Providers should make clear which evidence sets were reviewed, to what depth, and with what limitations.

All statements should be time-bounded and explicitly tied to the assessed environment to avoid any implication of future certainty. The assurance should also outline remaining monitoring, whether re-baselining has been completed, available rollback or containment options if new indicators emerge, and any assumptions, residual gaps, or conditions underpinning the assessment. This transparency helps boards and regulators understand the remaining risks and how they are being managed.

Example: *"As of 18:00 GMT on 10 December 2025, within the scope of the payments processing domain, we have independently validated that the firm has: (i) restored from backups previously verified as clean, (ii) completed integrity reconciliations against scheme and custodian datasets within agreed tolerances, (iii) re-imaged affected servers from clean baselines, (iv) rotated relevant credentials and renewed certificates, and (v) implemented heightened monitoring with tuned detections. We have not observed evidence of malicious persistence, lateral movement, or data tampering in the scoped environment since 12:10 GMT on 9 December 2025. Adjacent network segments and associated identity groups have been scanned and spot-tested with no indicators of compromise detected. Continuous monitoring remains in place, and should new indicators emerge, the firm has defined rollback procedures and agreed stakeholder notification pathways."*

4.2.3 Reconnection

Within a modern IT ecosystem, it is typical for technology supporting IBSs to interact with third-party capabilities. Following a cyber incident, it is critical for an affected firm to **provide assurance to third parties that it is demonstrably safe for reconnection to take place** (as part of the overall reconnection process described within the CMORG Reconnection Framework and the considerations set out in **Principle 12**).

From an assurance perspective, the affected firm should focus on demonstrating that reconnection risks are effectively managed by showing it has a clear understanding of the incident, that the threat actor has been removed with independent attestation where appropriate, and that systems have been rebuilt to a trusted state.

4.3 Continuous improvement

Principle 17: **A continuous feedback cycle supports the evolution of cyber recovery strategy, requirements, and execution over time, strengthening firm-level recovery capabilities and broader sector preparedness.**

Continuous improvement applies across the entire recovery lifecycle, updating strategy, capabilities, sequencing, and assurance processes in response to testing outcomes, lessons identified from incidents, the evolving threat landscape, business-led changes, and technological advancements.

A **formalised continuous improvement** process should be established to ensure that lessons from recovery activities are consistently captured, assessed, prioritised, and translated into tangible enhancements. This process should extend beyond technical controls to include strategy, governance, operating models, dependencies, and decision-making arrangements. Where appropriate, improvement activities should be overseen through established operational resilience or cyber governance forums, with visibility to senior management and, where material, the Board.

Post-incident reviews should capture evidence-based learnings that can be used to strengthen recovery capabilities. Importantly, firms should also share appropriate incident insights with the broader financial services sector and relevant authorities to help uplift collective resilience and reduce the likelihood and impact of similar incidents across the ecosystem.

Continuous improvement should also account for **change outside of discrete incidents**, including material technology transformations, changes to third-party dependencies, organisational restructuring, or emerging attack techniques. Cyber recovery strategies, requirements, and documentation should be reviewed periodically (at least annually) to confirm they remain aligned to the firm's current business model, service offering, and risk appetite. Where gaps are identified, firms should make explicit, risk-based decisions on whether to remediate, accept, or defer them.

Leadership engagement is an important enabler of effective continuous improvement. Boards and senior management should receive clear, outcome-focused reporting on the maturity and limitations of cyber recovery capabilities, progress against agreed improvement actions, and any material residual risks. This supports informed oversight and helps ensure that cyber recovery preparedness continues to develop in line with the firm's broader resilience objectives.

Taken together, a disciplined continuous improvement cycle helps firms avoid complacency, ensures recovery arrangements remain credible and executable, and incrementally strengthens the firm's ability to recover from significant cyber incidents over time.

Next steps

This guidance has set out a shared cyber recovery framing where an incident threatens a firm's viability and conventional resilience measures may no longer be sufficient. It is principles-based and non-prescriptive by design, recognising that cyber recovery remains an evolving discipline and that no single architecture, capability set, or operating model will be appropriate for all firms. Throughout this guidance, emphasis has been placed on prioritisation and realism, including an acknowledgement of uncertainty, non-determinism, and the limits of what can be fully tested or assured in advance. Firms are therefore encouraged to interpret and apply the guidance proportionately, in line with their own risk appetite, business model, technical complexity, and operating constraints.

To support this, the guidance concludes with a selection of **reflective questions** tailored to its intended audiences. These questions are not intended to be exhaustive, nor to represent required actions or benchmarks. Instead, they are designed to act as **practical prompts** helping senior management and those with operational accountability to test assumptions, surface hidden dependencies, and identify areas where recovery arrangements (including external dependencies) may be fragile in practice. Used thoughtfully, they can support internal discussion, scenario exercises, investment decisions, and continuous improvement by encouraging firms to consider how their cyber recovery approach would operate under severe stress, and where its known limits lie. Firms may use these questions as a starting point for their own analysis and dialogue, rather than as a checklist to be completed or a standard to be met.

Reflective questions for senior management

- **Q.** Does my firm have a clear understanding of its cyber recovery strategy, explicitly scoped to survivability, with known limits?
- **Q.** Does my firm have a clearly defined minimum viable recovery scope for viability-threatening scenarios, with known dependencies, constraints, and failure modes?
- **Q.** If my firm's production environment and disaster recovery arrangements are no longer trusted (e.g. through assumed compromise), how do we reestablish trust and have recovery assets ready for use when needed?
- **Q.** What testing has my firm performed to demonstrate our ability to recover under severe stress, and what are we unable to prove through testing?

Reflective questions for those with operational accountability

- **Q.** Does my firm understand the sequence of services it will attempt to recover under this scenario?
- **Q.** Does my firm understand the dependencies that teams will require throughout recovery (e.g. skills, access, authority, or coordination), as well as the tools, platforms, or automation that may be unavailable, degraded, or untrusted, and the practical workarounds that teams would use instead?
- **Q**. Have we sufficiently anticipated the levels of assurance associated with safe resumption of our critical dependencies (including third parties)?
- **Q.** During recovery, what circumstances might cause teams to pause, slow down, or decline to proceed due to lack of confidence, and what evidence or signals would be needed to regain enough trust to continue or adjust approach with agreed decision makers?

Appendix A: List of stressed impacts

The table below provides a non-exhaustive list of 'stressed impact' prompts that, if extrapolated, could persistently threaten availability and/or require protracted restoration times and complements the CMORG Dynamic Scenario Library (DSL)[5] with a selection of broader, but shallower examples.

Whilst not all of these examples have led to the invocation of cyber recovery processes, there is potential that if these occurred at scale or in combination, they could result in material impacts and may therefore be worth consideration within broader firm scenario planning.

Impact domain	Impact scenario	Real-world examples
Loss of Integrity • Core infrastructure • Business apps & data	**Ransomware/Malware** • Large databases • Hypervisor / HCI layers	**VMWare ESXi Attack (2024)** In 2024, a threat actor exploited a hypervisor-level flaw to escalate privileges on VMware ESXi hosts, then deploy Black Basta ransomware. This resulted in mass encryption of hypervisors, with broad impact across virtualisation, database, and integrated infrastructure layers.
Loss of Integrity • End user devices • Collaboration services	**Ransomware/Malware** • Mass endpoints	**CNA Financial Ransomware Attack (2021)** In 2021, CNA Financial, one of the largest U.S. insurance companies, was hit by the Phoenix CryptoLocker ransomware, which spread across thousands of endpoints and internal systems, bypassing standard defences and forcing a company-wide shutdown.
Loss of Integrity • Containment • Eradication assurance • End user devices • Core infrastructure • Business apps & data	**Insider threats**	**DPRK IT Worker Campaign** A large-scale operation in which the threat actor, posing as freelance developers and remote IT staff, infiltrates firms by using fake identities, proxy networks, and front companies to secure legitimate jobs. Once inside, they leverage company-issued endpoints to steal source code, sensitive data, and cryptocurrency, with earnings funnelled to support sanctioned weapons programs, highlighting the insider-like risk of hiring unvetted remote workers.
Software Component Vulnerabilities • Eradication assurance • Core infrastructure • Business apps & data	**Zero-day within key open-source component**	**Log4Shell (ApacheLog4j)** A critical zero-day in the open-source Apache Log4j library enabled remote code execution in countless Java-based applications used by the financial sector. Because Log4j was embedded across enterprise systems, the flaw exposed widespread vulnerabilities and prompted urgent global patching directives from security authorities.
Software Component Vulnerabilities • Eradication assurance • Core infrastructure • Business apps & data	**Malicious poisoning of widely used library**	**NPM Package poisoning (2018) & XZ Utils (2024)** In 2018, the widely used event-stream NPM package was hijacked when a new maintainer inserted a malicious dependency that secretly targeted the Copay Bitcoin wallet, stealing cryptocurrency keys from downstream users. More recently, in 2024, a contributor slipped a backdoor into XZ Utils, a core Linux compression library, enabling remote code execution through compromised authentication.
Trust Infrastructure Compromise • Core infrastructure • Business apps & data	**Full forest-level AD loss**	**NotPetya Incident (2017)** A destructive wiper campaign masquerading as ransomware that spread rapidly through compromised Ukrainian tax software before escalating worldwide. Using exploits like EternalBlue and stolen credentials, it crippled firms by encrypting or wiping Windows systems and making recovery nearly impossible.
Trust Infrastructure Compromise • End user devices • Core infrastructure • Business apps & data	**Kerberos (TGT) ticket compromise**	**Ryuk Ransomware Campaign (2019)** In 2019, Ryuk ransomware operators targeted corporate networks by gaining access through phishing or stolen credentials, then using tools like Mimikatz to extract Kerberos TGTs from memory. With these stolen tickets, they performed pass-the-TGT attacks, impersonating high-privilege accounts without knowing their passwords, enabling lateral movement across the network and deployment of ransomware at scale.
Trust Infrastructure Compromise • Core infrastructure • Business apps & data	**PAM compromise (intentional / unintentional)**	**SolarWinds SUNBURST Attack (2020)** In 2020, the SolarWinds SUNBURST attack saw the threat actor compromise highly privileged accounts managed by Privileged Access Management (PAM) systems, enabling them to move laterally, access sensitive networks, and deploy malicious payloads undetected for months.
Trust Infrastructure Compromise • End user devices • Core infrastructure • Business apps & data	**Internal CA compromise / External Certificate provider**	**Comodo Group Hack (2011)** In 2011, a threat actor compromised the Comodo Group, a trusted certificate authority, and issued themselves fraudulent SSL certificates for major domains, including mail.google.com, login.live.com, and addons.mozilla.org. These certificates could have been used for man-in-the-middle attacks, allowing the attackers to intercept and decrypt communications intended for these services.
Software Component Vulnerabilities • Core infrastructure	**Network edge device zero-day (vendor-wide)**	**Ivanti Connect Secure & Ivanti Policy Secure (2024)** In 2024, Ivanti Connect Secure and Ivanti Policy Secure appliances, commonly used for secure remote access, were found to have critical zero-day vulnerabilities. Exploitation of these flaws allowed attackers to implant web shells and harvest credentials, compromising client networks.

[5] CMORG (2026), Dynamic Scenario Library v1.2, March

Appendix B: Key reference material

Cross Market Operational Resilience Group (CMORG)

- Data Vaulting Reference Architecture v1.0 (Apr 2022)
- Cloud-Hosted Data Vaulting - Good Practice Guidance v1.0 (Jan 2025)
- Guidance for Firm Operational Resilience v3.0 (Apr 2025)
- Reconnection Framework v3.0 (Jul 2025)
- Dynamic Scenario Library v1.2 (Mar 2026)

UK National Cyber Security Centre (NCSC)

- Incident Management Guidance (Sep 2019)
- Small Business Guide: Response & Recovery (Oct 2020)
- Mitigating Malware and Ransomware Attacks (Sep 2021)
- Putting staff welfare at the heart of incident response (May 2022)
- Cyber Incident Exercising (Dec 2023)
- Guidance on effective communications in a cyber incident (Oct 2024)
- Ransomware-resistant backups (Nov 2024)
- Cyber Governance Code of Practice (Apr 2025)
- Decommissioning assets (May 2025)
- Cyber Assessment Framework: Principle D1. Response and recovery planning (Aug 2025)
- Supply chain security guidance (Oct 2025)
- How to prepare for and plan your organisation's response to severe cyber threat: a guide for CNI (Jan 2026)
- Improving your response to vulnerability management (Feb 2026)

International Organization for Standardization (ISO)

- ISO/IEC 27035-1:2023 – Information technology – Information security incident management - Part 1: Principles and process (Feb 2023)
- ISO/IEC 27040:2024 - Information technology - Security techniques - Storage security (Jan 2024)
- ISO/IEC 27031:2025 - Cybersecurity - Information and communication technology readiness for business continuity (May 2025)

National Institute for Science and Technology (NIST)

- NIST SP 800-184 – Guide for Cybersecurity Event Recovery (Dec 2016)
- NIST Cybersecurity Framework (CSF) 2.0 – Recover Function (Feb 2024)
- NIST SP 800-61 Rev. 3 – Computer Security Incident Handling Guide (Mar 2025)

Cyber Risk Institute (CRI)

- CRI Profile v2.0 (Feb 2024)