Five decisions for SME leaders to reduce risk and protect revenue

Executive White Paper

Date: 18/12/2025

Version: 1.0.1

Restriction: PUBLIC

by: Denis B. - Senior IT Consultant

Brimbor Consulting

Introduction

UK data shows **43%** of businesses identified a cyber breach/attack in the last 12 months (≈612k firms), with phishing dominant and an average cost of **£1,600** for the most disruptive incident (all businesses). Cyber security sits with the board, per the NCSC's 10 Steps and its Toolkit for Boards. *Source: gov.uk*

Busy Reader Summary

Map essential revenue functions and blind spots. Output: 1-page map + gap register with owners/dates. Anchor to NCSC 10 Steps and CAF B5.

- https://www.iso.org/standard/75106.html
- https://www.ncsc.gov.uk/collection/cyber-assessment-framework/caf-objective-b/principle-b5-resilient-networks-and-systems
- https://www.ncsc.gov.uk/pdfs/guidance/principles-for-ransomware-resistant-cloud-backups.pdf

Guarantee connectivity. Check ISP SLA/fix-time and right-to-exit.

- https://www.ofcom.org.uk/internet-based-services/network-security/resilience-guidance

Clear and separate **policies for access** and visitors/externals. Protect admin workstations/interfaces; log entries; escort visitors; no unauthorised devices.

- https://www.ncsc.gov.uk/collection/cyber-assessment-framework/caf-objective-b/principle-b5-resilient-networks-and-systems

Governance **baseline**. Executive owner (non-IT), incident plan, and to assure clients/insurers.

- https://ico.org.uk/action-weve-taken/complaints-and-concerns-data-sets/data-security-incident-trends/
- https://www.ncsc.gov.uk/cyberessentials/overview

Table of content

Why it matters

Cyber risk is now a routine business risk

Cyber incidents are no longer exceptional events affecting only large or highly digital organisations. In the UK, nearly half of businesses experienced a cyber breach or attack in the last 12 months, and this figure rises sharply with company size. For medium and large enterprises, cyber disruption is now a *probable* scenario, not a hypothetical one.

While the average cost of the most disruptive incident may appear modest when viewed in isolation (£3,550 for businesses and £8,690 for charities), this figure masks the real business impact: operational downtime, lost revenue, delayed customer delivery, reputational damage and management distraction. Phishing-led incidents, which dominate current attack patterns, exploit human and process weaknesses rather than technical sophistication, making them both frequent and difficult to eliminate entirely.

Scale increases exposure, not immunity

Larger organisations are targeted more often precisely because they are more complex. Multiple systems, legacy architectures, distributed teams and third-party dependencies increase the attack surface and the likelihood that a single failure propagates into wider disruption.

As organisations grow, informal controls and ad-hoc technical decisions no longer scale safely. What worked when systems were smaller and teams more centralised becomes a source of fragility. At this point, cyber risk transitions from an operational concern to a structural one, embedded in architecture, configuration, and decision-making models.

Cyber security is a governance issue, not an IT one

The UK's National Cyber Security Centre is explicit: cyber security sits within the Board's remit. This is not about expecting directors to become technical experts, but about recognising that architecture choices, risk acceptance, monitoring capability and incident readiness are business decisions with strategic consequences.

When cyber security is treated as "an IT problem", organisations tend to focus on tools rather than outcomes. Conversely, when it is governed at board level, priorities shift toward

resilience, clarity of accountability and informed trade-offs between cost, risk and operational continuity.

Resilience, not prevention, is the business objective

No organisation can guarantee it will never suffer a cyber incident. The meaningful differentiator is how often disruption occurs and how quickly the business can recover when it does.

The outcome to aim for is twofold:

- **Reducing the probability of material disruption** by addressing systemic weaknesses in architecture, configuration, and user exposure.
- **Minimising time to revenue recovery** through effective detection, decision-making and incident management.

Organisations that perform well in these areas experience cyber incidents as contained events rather than existential crises. For leadership teams, this directly translates into protected revenue, preserved customer trust and reduced strategic distraction.

Cyber resilience does not improve through awareness alone. It improves when leadership makes a small number of explicit, reviewable decisions about how the organisation creates revenue, where it is fragile, and which risks are consciously accepted or reduced. The following five decisions define that leadership stance, and each is an area where independent, external support can help accelerate clarity, challenge assumptions and turn intent into evidence.

Decision 1: Map essential functions and blind spots

(technology, process, physical environment, suppliers)

Why it matters

Disruptions rarely begin as “IT problems”. They start as interruptions to revenue-generating workflows (taking orders, processing payments, fulfilling deliveries, paying staff, supporting customers) and then cascade through overlooked weaknesses. A misconfigured admin interface, an unescorted visitor, or a single supplier outage can rapidly amplify a local failure into a business-wide disruption.

This is why the NCSC’s Cyber Assessment Framework (CAF B5) explicitly requires **resilience** by design. It recognises that administrative devices and interfaces are frequently targeted, that continuity must be planned rather than assumed and that these are **governance** concerns. The NCSC’s 10 Steps reinforce this by linking resilience directly to architecture, configuration, monitoring and incident readiness — not just to security tooling.

What leaders need to decide

The first leadership decision is to make revenue dependencies explicit and bounded. This means identifying a small set of **Essential Functions (EFs)** (typically three to seven) that directly create, protect or enable revenue. Examples include e-commerce checkout, invoicing runs, point-of-sale systems, customer support lines, warehouse or logistics platforms, or payroll.

For each essential function, leadership must require a structured catalogue of potential **break points**, across four domains:

- **Technical**: legacy operating systems or applications, privileged access paths (RDP, VPNs, cloud consoles), and single points of failure such as authentication, DNS, or payment gateways.
- **Process**: single approvers, batch-dependent activities, or gaps in out-of-hours coverage.

- **Physical**: server rooms, communications racks, keys and badges and any unlogged or weakly controlled access.
- **Suppliers and externals**: telecommunications, cloud and SaaS providers, managed service partners, payment processors, logistics providers and visitor workflows (registration, escorting, and device rules).

Each blockage should then be classified by acceptable downtime (critical, important, or supporting) to create a shared, provisional tolerance model that will be refined in later decisions.

How the decision is taken at board level

This decision is not about technical detail; it is about visibility and accountability. In practice, it is best captured on a single page that maps essential functions across revenue, operations, IT, and physical lanes. For each essential function, three to five interruptors are listed alongside a clearly named control owner.

Crucially, the board (or delegated executive committee) must explicitly record whether each gap is accepted or remediated, by when, and what evidence will demonstrate completion. Evidence is not a policy statement; it is a configuration screenshot, a log extract, or a contractual clause. External advisors are often valuable here to ensure the mapping is realistic, challenge optimistic assumptions, and translate technical controls into board-level assurance.

Illustrative example

In a retail SME, card-present point-of-sale transactions may be identified as an essential function. Interrupters could include a broadband outage, router misconfiguration, reuse of admin credentials, or a contractor connecting an unmanaged laptop in a back-office area. Controls might include mobile connectivity failover, reserved network capacity for POS traffic, a segregated admin workstation with no email or web access, and formal visitor escort procedures supported by signage and access logs.

What “good” looks like as a deliverable

A standard, repeatable output is a one-page **Essential Function Map** paired with a gap register showing owners, remediation actions, required evidence, and dates. When aligned

with CAF B5 and the NCSC 10 Steps, this format is immediately recognisable to auditors, insurers and regulators, and provides a defensible basis for risk discussions.

Ongoing assurance

At a minimum, leadership should expect quarterly updates showing changes to the map, evidence that the highest-impact gaps have been closed and an executive note highlighting any new administrative interfaces or suppliers introduced.

A common pitfall to avoid is conflating visitors with general physical security. Visitor access is a distinct risk domain and requires its own explicit policies, controls, and evidence — not informal practices.

Decision 2: Set pragmatic recovery objectives and prove they work

(RTO / RPO, with enforced restore testing and offline or immutable copies)

Definitions

- **RTO (Recovery Time Objective)**: the maximum acceptable time to restore a service after disruption. It answers the question: *how long can the business tolerate this function being unavailable?*
- **RPO (Recovery Point Objective)**: the maximum acceptable data loss measured in time. It answers the question: *how much data can we afford to lose if we have to restore?*

These are business decisions, not technical preferences. They express revenue tolerance and operational reality.

Why it matters

Most serious outages are not caused by missing backups; they are caused by **failed or unproven recovery**. Data exists, but cannot be restored fast enough, cleanly enough, or at all.

This is why recognised standards focus on objectives and proof. ISO 22301 requires organisations to exercise and test their continuity arrangements to validate that recovery strategies actually work (Clause 8.5), with RTO and RPO derived from the Business Impact Analysis (BIA). Similarly, the NCSC is explicit: organisations should maintain recent offline backups and design them to resist ransomware and other destructive actions. This includes the ability to restore earlier versions, protect encryption keys and prevent unauthorised deletion.

In governance terms, backups without *tested* restores are assertions, not controls.

What leaders need to decide

Leadership must set **pragmatic, tiered recovery objectives** that reflect how the business operates, not aspirational “zero-loss” targets that are neither affordable nor enforceable. Typical tiers include:

RTO tiers

- **Tier 1**: ≤ 4 hours (payments, checkout, customer support lines, ERP posting)
- **Tier 2**: ≤ 24 hours (management reporting, data warehouses)
- **Tier 3**: ≤ 3–5 days (archives, historical data)

RPO tiers

- **Tier 1**: 0–15 minutes (transactional systems)
- **Tier 2**: ≤ 4 hours (line-of-business applications)
- **Tier 3**: ≤ 24 hours (bulk or reference data)

Retention expectations should also be explicit: typically, 30–90 days online, at least 7–30 days in immutable or offline form, and 6–12 months archived depending on audit and regulatory needs.

Finally, leaders must require **evidence cadence**: at minimum, quarterly full restore tests for Tier 1 systems and semi-annual tests for Tier 2 systems.

How the decision is taken at board level

This decision is enacted by approving an **RTO/RPO matrix mapped to each Essential Function** identified in Decision 1. The board (or delegated authority) should explicitly mandate that there is, at all times, **at least one backup copy that is offline or immutable**.

Equally important is the approval of restore runbooks. These define who performs the restore, where it is performed, how long it should take, and what evidence must be collected. External advisors often add value here by validating that objectives are realistic, ensuring restore tests are meaningful, and confirming that evidence would stand up to audit, insurer or incident scrutiny.

Illustrative example

For a services SME, customer billing may be classified as an essential function with an RTO of four hours and an RPO of fifteen minutes. Backups might include hourly snapshots retained for seven days, daily full backups retained for thirty days, immutable object-locked copies for twenty-one days, and a weekly offline tape.

Each quarter, the organisation restores a recent snapshot into a quarantined network, recording timings, logs, and screenshots. The test passes if the database opens correctly, integrity checks succeed, and restoration completes within the defined RTO.

What "good" looks like as a reference architecture

A commonly accepted pattern is **3-2-1 plus immutability**: three copies of data, on two different media, with one offsite and one immutable or offline. Immutability can be implemented through object-lock or WORM (Write Once, Read Many) storage, or through backup vaults with delayed deletion and dual control. Separation of duties is essential: backup administrators should not hold domain-level delete rights, and destructive actions should require out-of-band authentication. These patterns directly align with NCSC guidance.

Ongoing assurance

Leadership should routinely request evidence from restore exercises: screenshots and logs, actual restore durations versus RTO, data currency versus RPO, the version restored, and confirmation of immutable policy enforcement. An issue log should track gaps, remediation actions, and the date of the next test.

Common pitfalls include backup accounts with unrestricted delete rights, lack of versioning, and treating backups as complete without ever rehearsing a full restore.

Decision 3: Guarantee availability

(connectivity SLAs, exit rights, and documented 4G/5G failover)

Why it matters

For most organisations, connectivity is the single point through which revenue flows. When broadband fails, payment terminals stop, online checkouts stall, VoIP degrades, and cloud platforms become unreachable. Availability failures therefore translate immediately into lost revenue and operational paralysis.

The regulatory context matters here. In the UK, Ofcom codes of practice give organisations explicit rights when connectivity underperforms, including the ability to exit contracts without penalty if minimum guaranteed speeds are not met. At the same time, national infrastructure coverage has reached a point where cost-effective resilience designs are feasible: widespread optical fibre network (FTTP) and near global 4G, with rapidly expanding 5G, change what "reasonable" availability looks like. Leadership decisions should be informed by this reality, not by historical assumptions.

Definitions

- **Primary link**: the main connectivity service used under normal conditions (typically FTTP or a leased line).
- **Failover**: an automatic switch to an alternative connection when the primary link fails.
- **Exit rights**: contractual rights allowing termination or migration when agreed performance thresholds are not met.
- **SLA (Service Level Agreement)**: documented performance commitments, including fix times and compensation.

What leaders need to decide

Leadership must decide how availability is guaranteed in practical terms, starting with a clear choice of primary connectivity and a tested fallback. In most cases, FTTP is sufficient; leased lines are justified where symmetrical bandwidth or stricter SLAs are required.

Failover should be explicit and automatic. A 4G or 5G router with seamless switchover is now a standard expectation, with a realistic target of restoring connectivity within 60 to 120 seconds. During failover, bandwidth must be actively managed: essential services such as POS, SaaS applications, and VoIP should be prioritised, while non-essential or high-volume traffic is restricted or blocked.

Contractually, leaders must require documentation of provider fix-time commitments, compensation mechanisms and the practical steps to exercise exit rights — including who within the organisation is authorised to trigger them.

How the decision is taken at board level

This decision is operationalised by reviewing existing ISP contracts against Ofcom codes of practice and producing a concise “service card” for each site or provider. This card captures guaranteed speeds, escalation contacts, fix-time expectations, and the exit path if service levels are not met.

The board or executive committee should then approve a connectivity runbook. This runbook mandates regular (quarterly) failover testing during core operating hours, using a controlled window, and requires evidence such as page-load times, transaction completion times and system logs. External advisors can support by validating contract terms, designing failover architectures, and ensuring tests reflect real business usage rather than synthetic checks.

Illustrative example

In a multi-site hospitality group, each location may use FTTP as the primary link with a 5G failover on a dual-WAN router. The POS network is placed on a dedicated VLAN with a guaranteed 500 kbps per terminal during failover, while VoIP is constrained to a low-bandwidth codec.

During quarterly tests, the primary circuit is deliberately taken down. Card payments are verified end-to-end, primary connectivity is restored, and router logs are exported as evidence.

What “good” looks like as a reference architecture

A common edge pattern uses a dual-WAN router combining FTTP and 4G/5G, policy-based routing, per-VLAN quality of service, and throttled guest Wi-Fi during failover. For higher-criticality sites, organisations may add provider diversity (for example, two fixed-line providers on separate networks, or a leased line paired with FTTP) with mobile connectivity acting as a tertiary option.

Ongoing assurance

Leadership should periodically review failover test evidence, confirm that contractual exit rights remain enforceable, and validate that segmentation rules still reflect current business priorities.

Typical pitfalls include assuming failover will “just work”, lacking a documented runbook, holding old contracts without clear exit rights, or allowing all traffic to flow during failover; resulting in universal degradation instead of controlled continuity.

Decision 4: Physical security ≠ visitor and external access

(two distinct policies, one leadership responsibility)

Why it matters

Not all disruptions are “cyber” in the technical sense. Low-tech intrusions, such as an unlocked server room, an open electrical panel or an unescorted visitor, can stop operations just as effectively as malware. These incidents are often accidental, well-intentioned, and entirely human. That is precisely why they must be governed, not assumed away.

Physical security and visitor or external access are frequently conflated into a single informal process. In practice, they represent different risk domains and must be addressed separately, while remaining under clear executive ownership. The NCSC’s Cyber Assessment Framework (CAF B5) explicitly requires protection of administrative devices and interfaces, and planning for continuity in the event of disruption. The Government Security Handbook reinforces this by linking governance, segmentation, and reduction of blast radius, including for physical and human-initiated events.

A reminder from experience

In one organisation, a server room was accessible without escort. A senior delegated director (privileged in the business, but not an IT role) had the habit of entering the server room to investigate a service issue and unplugged power cables from a server “to see if it would work again”. Yes, a hard power-off can occasionally clear a blocked system, but it also introduces bigger risks: un-flushed RAID/controller write-caches and broken write-ordering, dirty shutdowns of file systems and databases, potential uncontrolled failover, and the loss of diagnostic logs needed to fix the root cause. The minutes “saved” at the front end were repeatedly offset by hours of rebuilds, integrity checks and degraded performance across shared storage.

The incident was not malicious. The failure was structural. Access should not have been possible without escort, and the racks themselves should have been locked. Following recommendations, the server room was permanently locked, rack doors secured, and key custody formalised. Access became a managed process: keys held by executive support, clear

rules on who may enter alone or accompanied, phone confirmation when needed, and mandatory escort for all external interveners. This type of control is simple, inexpensive, and highly effective.

Definitions (shared language for leadership)

- **Physical security**: controls protecting spaces, infrastructure, and equipment (rooms, racks, power, comms).
- **Visitor / external access**: processes governing non-employees or non-routine access (contractors, suppliers, auditors).
- **Privileged paths**: systems and devices used to administer, configure, or recover core services.

What leaders need to decide

Leadership must require a clear, enforceable baseline of physical and human controls, expressed in concrete terms.

Physical controls should include controlled access to server rooms and racks (badges or keys with logs), cameras where appropriate, and an inventoried key system. Administrative workstations should be dedicated, locked down, and used only for privileged tasks (no email, no web browsing, and secured storage when not in use).

Visitor and external access controls must be explicit and visible. This includes registration on arrival, clearly marked visitor badges, mandatory escorting, and well-signposted restricted zones. No unauthorised devices should be connected to the network: switch ports should be dedicated or locked down, USB storage disabled where not explicitly required, and guest Wi-Fi fully isolated from internal systems.

Privileged path hardening completes the picture. Multi-factor authentication is mandatory, administrative systems should not have direct internet access, logging must be comprehensive, and elevation of privilege should occur through a controlled bastion or jump host.

How the decision is taken at board level

This decision is formalised by approving **two distinct policy documents**, with **two named owners**, and a single point of executive arbitration. Typically, facilities or HR own the

visitor and reception policy in an “easy-read” format, while IT or security own the administrative and privileged-access policy. A joint forum ensures neither domain undermines the other.

Quarterly assurance should include sampled visitor logs, photographic evidence of locked racks, and reviews of privileged access. External advisors can help translate lived incidents into proportionate controls and ensure policies are usable, not theoretical.

Illustrative example

In a logistics environment, visitors arriving at a loading bay receive a visitor badge, are escorted, and are restricted to designated “amber” zones. Badges are returned and logged on departure.

Administrative workstations are in a secured room, USB ports are disabled, and all privileged actions are performed via a bastion system that records sessions and requires strong authentication.

Ongoing assurance

Leadership should periodically verify that physical access rules are followed in practice, that visitor Wi-Fi remains segmented and rate-limited, and that administrative systems have not drifted into everyday use.

Common anti-patterns include a single, forgotten “visitor procedure” buried in a PDF, unlocked racks “because it’s inconvenient”, and admin workstations treated like standard office computer. These are not technical failures; they are governance gaps.

Decision 5: Governance and compliance

(executive ownership, incident readiness, and Cyber Essentials as a baseline)

Why it matters

Cyber resilience breaks down most often not because controls are missing, but because decisions are slow, fragmented, or improvised under pressure. Appointing an executive owner outside of IT aligns cost-risk trade-offs, accelerates cross-team decisions, and removes ambiguity when time matters.

A pre-approved incident plan replaces improvisation with intent. It defines who decides, who communicates, and what takes priority (before reputational, legal, or financial pressure distorts judgement).

Cyber Essentials provides a government-backed baseline recommended by the NCSC. It reassures customers and insurers, supports access to certain public-sector markets, and creates a common language for assurance. Monitoring ICO enforcement trends further informs how and when to notify, shaping communication strategies before an incident occurs.

Definitions

- **Executive owner**: a senior leader accountable for cyber risk decisions and trade-offs, not day-to-day technical delivery.
- **Incident plan**: a pre-approved framework for detection, escalation, decision-making, communication, and recovery.
- **Cyber Essentials (CE / CE+)**: a UK government-backed scheme covering five core control areas, with CE+ adding independent technical verification.

What leaders need to decide

Leadership must formally appoint an **executive owner**, typically a CFO, COO, or Legal Director, with a written charter defining decision rights, budget authority, and arbitration responsibilities for SLAs and risk acceptance.

An **incident plan** must be approved at executive level. This includes activation criteria, a decision matrix spanning operations, payments, HR, legal, and customer impact, and pre-

drafted communication templates for email, web, and key stakeholders. Notification thresholds for the ICO and coordination with the NCSC should be explicit.

Cyber Essentials should be adopted as the organisational baseline. For a disciplined SME, preparation typically takes six to eight weeks. The scope should cover the entire organisation, with controls enforced across firewalls, secure configuration, user access control, malware protection, and patch management. Where feasible, CE+ provides stronger assurance through independent testing.

How the decision is taken at board level

This decision is enacted through a short governance pack: the executive owner's charter, the incident plan and communication templates, and a Cyber Essentials readiness or gap report.

Progress should be reviewed quarterly, with clear visibility of outstanding gaps, remediation timelines, and any changes to regulatory or insurer expectations. External advisors can support by stress-testing incident plans, validating CE readiness, and ensuring evidence meets external scrutiny.

Illustrative example

In a B2B services firm, the CFO acts as executive owner. Playbooks exist for ransomware, phishing, and connectivity loss, defining who contacts whom and what messages are issued at one hour, four hours, and twenty-four hours.

Cyber Essentials Plus readiness includes hardened endpoints, Multi Factor Authentication across all services, patching within fourteen to thirty days, privilege reviews, and independent scans and tests.

Sector note (regulated environments)

In financial services, regulators such as the FCA require firms to identify important business services, define impact tolerances, and demonstrate the ability to operate within those tolerances through mapping and scenario testing. While not all organisations are regulated to this level, the direction of travel is clear and increasingly influential beyond the sector.

Ongoing assurance

Leadership should expect regular confirmation that ownership remains clear, incident plans are current, and Cyber Essentials controls have not drifted.

A common failure mode is governance “on paper only”: an owner without authority, an incident plan that has never been exercised, or Cyber Essentials treated as a one-off compliance task rather than an operating baseline.

Choosing the right parameters without overspending

Cyber resilience does not require perfection. It requires coherence. The most effective approach is sequential and revenue-led:

- start from how the organisation makes money (**Decision 1**)
- assign realistic recovery objectives (**Decision 2**)
- calibrate connectivity accordingly (**Decision 3**)
- lock down human and physical access (**Decision 4**)
- and anchor everything in executive ownership and a recognised baseline (**Decision 5**)

Evidence matters. If it cannot be shown (screenshots, logs, tickets, reports) it cannot be relied upon. Controls only exist once they produce evidence that is reviewed, at least quarterly, at executive level.

Standard patterns should be used first. The 3-2-1(+immutable) backup model, dual-WAN connectivity with mobile failover, privileged administrative workstations, and Cyber Essentials or CE+ are widely recognised by insurers, auditors, and regulators because they work and are repeatable.

A minimal operating model (who does what)

- **Executive owner (non-IT)**: arbitrates risk and cost, sets priorities, leads external communication.
- **Risk / IT / Security**: maps functions and gaps, sets RTO/RPO, implements immutable backups, maintains runbooks and tests, delivers CE/CE+.
- **Operations / Facilities / HR**: enforces physical and visitor policies, escorts, and access logs.
- **Finance / Procurement**: manages SLAs, Ofcom exit rights, penalties, and supplier contingency plans.
- **Comms / Legal**: owns templates, client and press messaging, ICO/NCSC notifications, and evidence retention.

What to ask for in the next 90 days

- **By Day 14**: one-page Essential Function Map and gap register aligned to CAF B5 and the NCSC 10 Steps.
- **By Day 30**: approved RTO/RPO grid, 3-2-1(+immutable) backup policy, first restore drill evidence pack.
- **By Day 45**: ISP service card (SLA and exit rights), mobile failover tested with QoS, test report.
- **By Day 60**: two approved policies (physical security and visitors), privileged admin workstation built, audit sample completed.
- **By Day 90**: incident plan and communication templates approved, Cyber Essentials readiness report delivered.

Brimbor Consulting *can support or accompany organisations across each of these decisions (from executive facilitation and mapping, through technical validation and evidence packs, to Cyber Essentials readiness and incident preparedness) with a focus on proportionate controls, clarity of ownership, and outcomes that stand up to scrutiny.*

Ready to take control of your cyber resilience?

→ Start a conversation

This document provides general guidance for executives on cyber and operational resilience. It is not legal, regulatory, tax or insurance advice and does not create any advisory, fiduciary or client relationship. The content is provided "as is", without warranties of completeness or fitness for a particular purpose. Implementation decisions remain the reader's responsibility and must be validated against applicable laws, contracts and sector rules; where necessary, seek advice from qualified counsel or regulated professionals. Neither the author nor Brimbor Consulting accepts any liability for anything arising from reliance on this document.