Baringa

Asking the right cyber questions

The Australian directors' guide

In this guide

Asking the right questions

The directors we work with are becoming aware of the damage that poorly managed cyber risk can inflict. That it can disrupt critical services, erode customer confidence, trigger regulatory scrutiny and create long-tail strategic harm that only becomes visible months or years later.

Cyber is no longer only about cybercrime. State-linked espionage and foreign interference are at extreme levels and expected to worsen, accelerated by advances like AI. Espionage is a blended risk spanning cyber, people, access pathways, third parties, and physical security. This guide outlines the cyber security areas directors oversee and the responses they should expect from management.

The Australian Institute of Criminology estimates espionage cost Australia at least A$12.5 billion in 2023-24. This is likely an underestimate, as many serious and cascading impacts are difficult to detect or value. The good news is the same research estimates that tens of billions of dollars in additional costs may have been prevented through mitigation and counter-espionage activity.

This is a useful board lens: the value of security investment is often measured in losses avoided, not losses recorded.

For boards and directors, the challenge is not to become technical. It is to ensure cyber risk is governed with the same discipline as financial and operational risk – clear accountability, scenario-based understanding of material exposures, credible assurance, and evidence that investment is reducing risk where it matters most.

This guide is designed for directors and senior executives. It offers focused questions to help leaders engage more effectively on cyber matters and understand the characteristics of a well-supported response.

What to ask, and what to expect in response.

1 AI intellectual property and data exposure

As these tools become core intellectual property (IP), they also become attractive targets for cyber threat groups seeking to access client data, steal proprietary models, or compromise the integrity of AI-driven decision-making.

Director question:

What assurance do we have that our AI intellectual property and the customer data it relies on, are adequately protected?

Expected responses:

- We have a defined strategy for protecting AI development and deployment, including source code and model integrity, insider risk controls, and strong access governance.
- We have implemented controls to protect customer data and PII throughout AI data ingestion, processing, and output, including monitoring for unintended data exposure.
- We employ threat intelligence monitoring that focuses on known attack techniques targeting IP and customer data workflows.

2 The quantum transition

The Australian Signals Directorate (ASD) recommends organisations should start planning now to achieve quantum readiness by 2030. This involves replacing the quantum-vulnerable (RSA/ECC) public-key cryptography your organisation might use for critical security purposes like identity, certificates and secure connections.

Migration is a multi-year and dependency-heavy task because cryptography is embedded across platforms, vendors and third parties. Directors also need to be aware of the near-term threat that attackers can "harvest now, decrypt later" by collecting encrypted data today to crack down the track when capabilities mature.

Director questions:

Do we know where we rely on quantum-vulnerable cryptography?

What is our time-bound plan to move our critical services and data to post-quantum standards?

Expected responses:

- We're building a cryptography inventory, including a cryptographic bill of materials (CBOM), so we know exactly where our quantum-vulnerable cryptography is embedded.
- We're following ASD/ACSC's transition stages (Locate, Assess, Triage, Implement, Communicate/Educate).
- We've assigned decision-makers to unblock vendor or architecture constraints.
- We'll have a refined transition plan by end-2026. Implementation will begin by end-2028, starting with our critical systems and highest-value/long-life data, and be complete by end-2030.
- Our planning is anchored on mature post-quantum cryptography standards, including NIST's, and vendor-supported implementations.

3 Assets and dependencies

Core networks, internal identity and access, payment engines, ERP and finance platforms, and customer identity systems carry disproportionate risk. Without a clear, current understanding of how these critical assets and their dependencies interact, containment or recovery actions can have unintended consequences, such as halting customer channels or disrupting payments.

Director question:

Has our organisation identified its crown-jewel information and capabilities?

What measures are we using to prevent, detect and respond to threats like espionage and insider activity?

Expected responses:

- We have a defined list of crown jewels (data, models, IP, strategic plans, critical service designs), mapped to where they live and who can access them.
- We have an insider risk program prioritised around these crown jewels, including privileged access controls, monitoring and tight joiner-mover-leaver processes.
- Third-party and partner access is controlled, monitored and reviewed for crown-jewel areas.
- We can demonstrate detection and response coverage for high-risk exfiltration and misuse paths – not just policy compliance.
- We can show examples of when controls prevented or detected common espionage patterns, including insider-enabled collection.

4

Key technology assets

Organisations often overlook the physical or virtual assets that critical processes run on. Unauthorised access to client portals, and physical and virtual hardware are all well-known targets. But threat actors are just as likely to exploit vulnerabilities in the underlying operating system.

?

Director question:

How do we know what technology assets our critical business processes rely on?

Expected responses:

- We regularly update the Configuration Management Databases of the assets supporting business-critical functions, including physical servers, cloud-hosted resources and virtual environments.
- We map out asset dependencies to identify risks at operational levels.
- We implement asset tagging and monitor solutions to ensure complete visibility.

5 Critical business processes

After a cyber breach, getting core business functions back online as quickly as possible requires knowing which processes to recover first; and having the appropriate physical or virtual recovery infrastructure available.

Director questions:

How do we know we've prioritised the right business processes in our recovery plans?

When did we last test our ability to restore them?

Expected responses:

- We have clearly identified key business processes linked to our primary objectives and missions.
- We prioritise recovery efforts in line with our Cyber Incident Response Plan, supported by appropriate infrastructure and restoration capacity planning.
- Our last backup restoration testing was on [DATE].
- We implement asset tagging and monitor solutions to ensure complete visibility.

6 Third-party visibility

Supply chains are highly interconnected. Payments processors, market data providers, software vendors and managed service partners all sit within the attack surface. The compromise of a common supplier can become systemic risk. Yet supplier detection maturity often remains opaque.

Director question:

How do we assure the end-to-end cyber resilience of our most critical third parties, including their ability to prevent, detect, respond to, and recover from incidents that could impact us?

Expected responses:

- We regularly evaluate supplier cybersecurity posture through third-party audits, questionnaires and breach simulations.
- We participate in collaborative threat-sharing networks to enhance supply chain resilience.
- We use dedicated supplier risk monitoring tools that align with sector-specific regulatory expectations, such as Upguard, SecurityScorecard or BitSight.

7

CISO and CRO alignment

A key governance test is alignment between cyber risk (led by the CRO) and cyber operations (led by the CISO). These teams often use identical terms that mean different things. For example, “Control testing” in:

Governance/risk – means framework alignment, policy conformance, and whether documented controls exist and are followed.

Cyber operations – means validating controls under adversarial conditions, through attack emulation or red teaming, to prove they work in practice.

Just because cyber risk is ‘green’ on a board report, it doesn’t mean cyber controls actually work. Unless risk and cyber teams are working in lock-step, directors can miss the risks that matter most. This is when assurance programs drift – creating boardroom confidence that isn’t matched by real-world resilience.

Director questions:

Do the CISO and CRO share a single view of our top cyber risks?

Where is the evidence that our control testing and remediation are reducing those risks in line with enterprise risk appetite?

Expected responses:

- The CISO and CRO run a formal interlock with a shared cyber risk register aligned to enterprise risk, including agreed severity and decision rights.
- Cyber risk acceptance and exceptions are governed through risk appetite thresholds, with defined decision rights. We work through what the CISO can accept versus what must go to the CRO/ExCo/Board.
- We prioritise control testing based on the top risk scenarios affecting our critical services and material business impacts – not just generic framework coverage.
- We have examples where hands-on testing results directly changed our risk ratings, funding decisions or remediation timelines.
- Board reporting includes risk movement over time, top scenario exposure, operational cyber control effectiveness and remediation progress.
- Remediation actions have clear owners and expiry dates – not just technical metrics.

8

Cyber risk quantification

For boards, the challenge is achieving decision-grade clarity: understanding which cyber risks could materially impact critical services and making defensible trade-offs between investment and residual risk. Many organisations still rely on colour-coded heat maps, which can obscure scale, comparability and uncertainty, and make it difficult to explain why one investment matters more than another.

Some mature or highly regulated environments address this through quantitative analysis, but for many organisations that level of precision is not yet realistic. The underlying problem remains the same: without a consistent way to describe scenarios, assumptions, and impact in business terms, boards can struggle to judge whether spend is reducing the risks that matter most or simply increasing activity.

Directors should be alert to this gap and press for a clearer line of sight between material scenarios, control effectiveness, and risk acceptance decisions.

?

Director questions:

Can we quantify our top cyber risks in business terms?

What scenarios and assumptions are we using to quantify cyber risk?

Where is the evidence that demonstrates how our current controls and investments are reducing exposure in line with enterprise risk appetite?

Expected responses:

- We quantify our top cyber risks based on scenarios tied to critical services and material business impacts.
- We estimate impact as ranges with confidence levels and explicitly state assumptions and data gaps.
- Quantification uses a consistent enterprise risk methodology and is reviewed by both the CISO and CRO.
- We use quantification to inform investment priorities, risk acceptance decisions and resilience planning.
- We track exposure movement over time to account for changes in the threat environment or operational dependencies.
- We refresh quantification every few months – and after major incidents/changes.

9 Scenario planning

Institutions need to be aware of their most likely attack scenarios. These are not the necessarily the worst-case scenarios, which often have a low likelihood.

Director question:

What gives us confidence that we understand the most plausible attack scenarios against us and our supply chain?

Expected responses:

- We regularly validate realistic attack scenarios using frameworks such as MITRE ATT&CK or D3FEND.
- We proactively track emerging threat scenarios, including unknown threats within transformation projects, leveraging cyber threat intelligence.
- We use the Value-at-Risk methodology to quantify cyber risks to both business-critical processes and supply chains.

10 Cyber threat detection

Security platforms in large organisations ingest billions of events each month. Without careful tuning, analysts drown in false positives while real incidents slip through. Equally, without alignment to business priorities, detection rules may miss scenarios that matter most, like data exfiltration, insider fraud or payments systems disruption.

?

Director question:

Can we detect and contain the threats we care most about, quickly?

Expected responses:

- Our detection capabilities are tuned in to reduce alert fatigue while focusing on material business risks.
- We map detection use cases to critical business processes and regulatory obligations to ensure we are prioritising high-impact incidents.
- We regularly review our detection performance through red teaming, simulated attacks and post-incident analysis to validate effectiveness.

11 Cyber talent

High turnover, skill shortages and reliance on outsourced providers all pose risks to continuity and context retention.

Director question:

What steps are we taking to build, retain and test the skills of the teams underpinning our cyber detection capability?

Expected responses:

- We have a clearly defined skill development program, with regular training to ensure capability retention.
- We are monitoring team turnover and use proactive measures to mitigate retention risks.
- We document evidence of processes and capabilities where external partners are involved in cyber detection efforts.

12 Coverage gaps

Many organisations operate sprawling technology estates, spanning cloud, on-premise, retail infrastructure, payment systems, trading platforms and customer apps. Over time, detection coverage can drift as new assets emerge, or old ones are decommissioned. This can create unmonitored blind spots that leave high-value areas exposed.

?

Director questions:

Are we monitoring and protecting all of our technology estate?

Do we have blind spots?

Expected responses:

- We currently map detection coverage across all critical assets and broader asset classes, including core systems and customer-facing platforms.
- We regularly reconcile coverage against the full asset base, including identifying coverage gaps and assessing risks.
- Our assurance mechanisms include vulnerability scanning, security testing, such as Purple Teaming, and internal or external audits.
- We use automated asset discovery tools and continuous vulnerability assessments to prevent gaps emerging during technology transitions.

13 Control efficacy

Few cybersecurity frameworks and standards support structured reasoning or provide quantitative performance metrics that help assess the extent of protection they offer.

Director questions:

How do we check our control efficacy?

What are we learning from other organisations?

Expected responses:

- We use control standards, like the Cyber Assessment Framework, NIST CSF and ISO27001, for security benchmarking.
- We regularly engage in Purple Teaming exercises, leveraging threat intelligence in our sector.
- We study peer organisations to learn how we need to change based on their cyber incidents.

14

Cyber incident response and testing

As cyber incidents can escalate quickly into business crises, testing should evidence not only technical response and recovery, but also timely cross-functional decision-making under pressure (legal, operational and communications), including effective executive coordination and escalation to the board where the incident could materially impact the organisation.

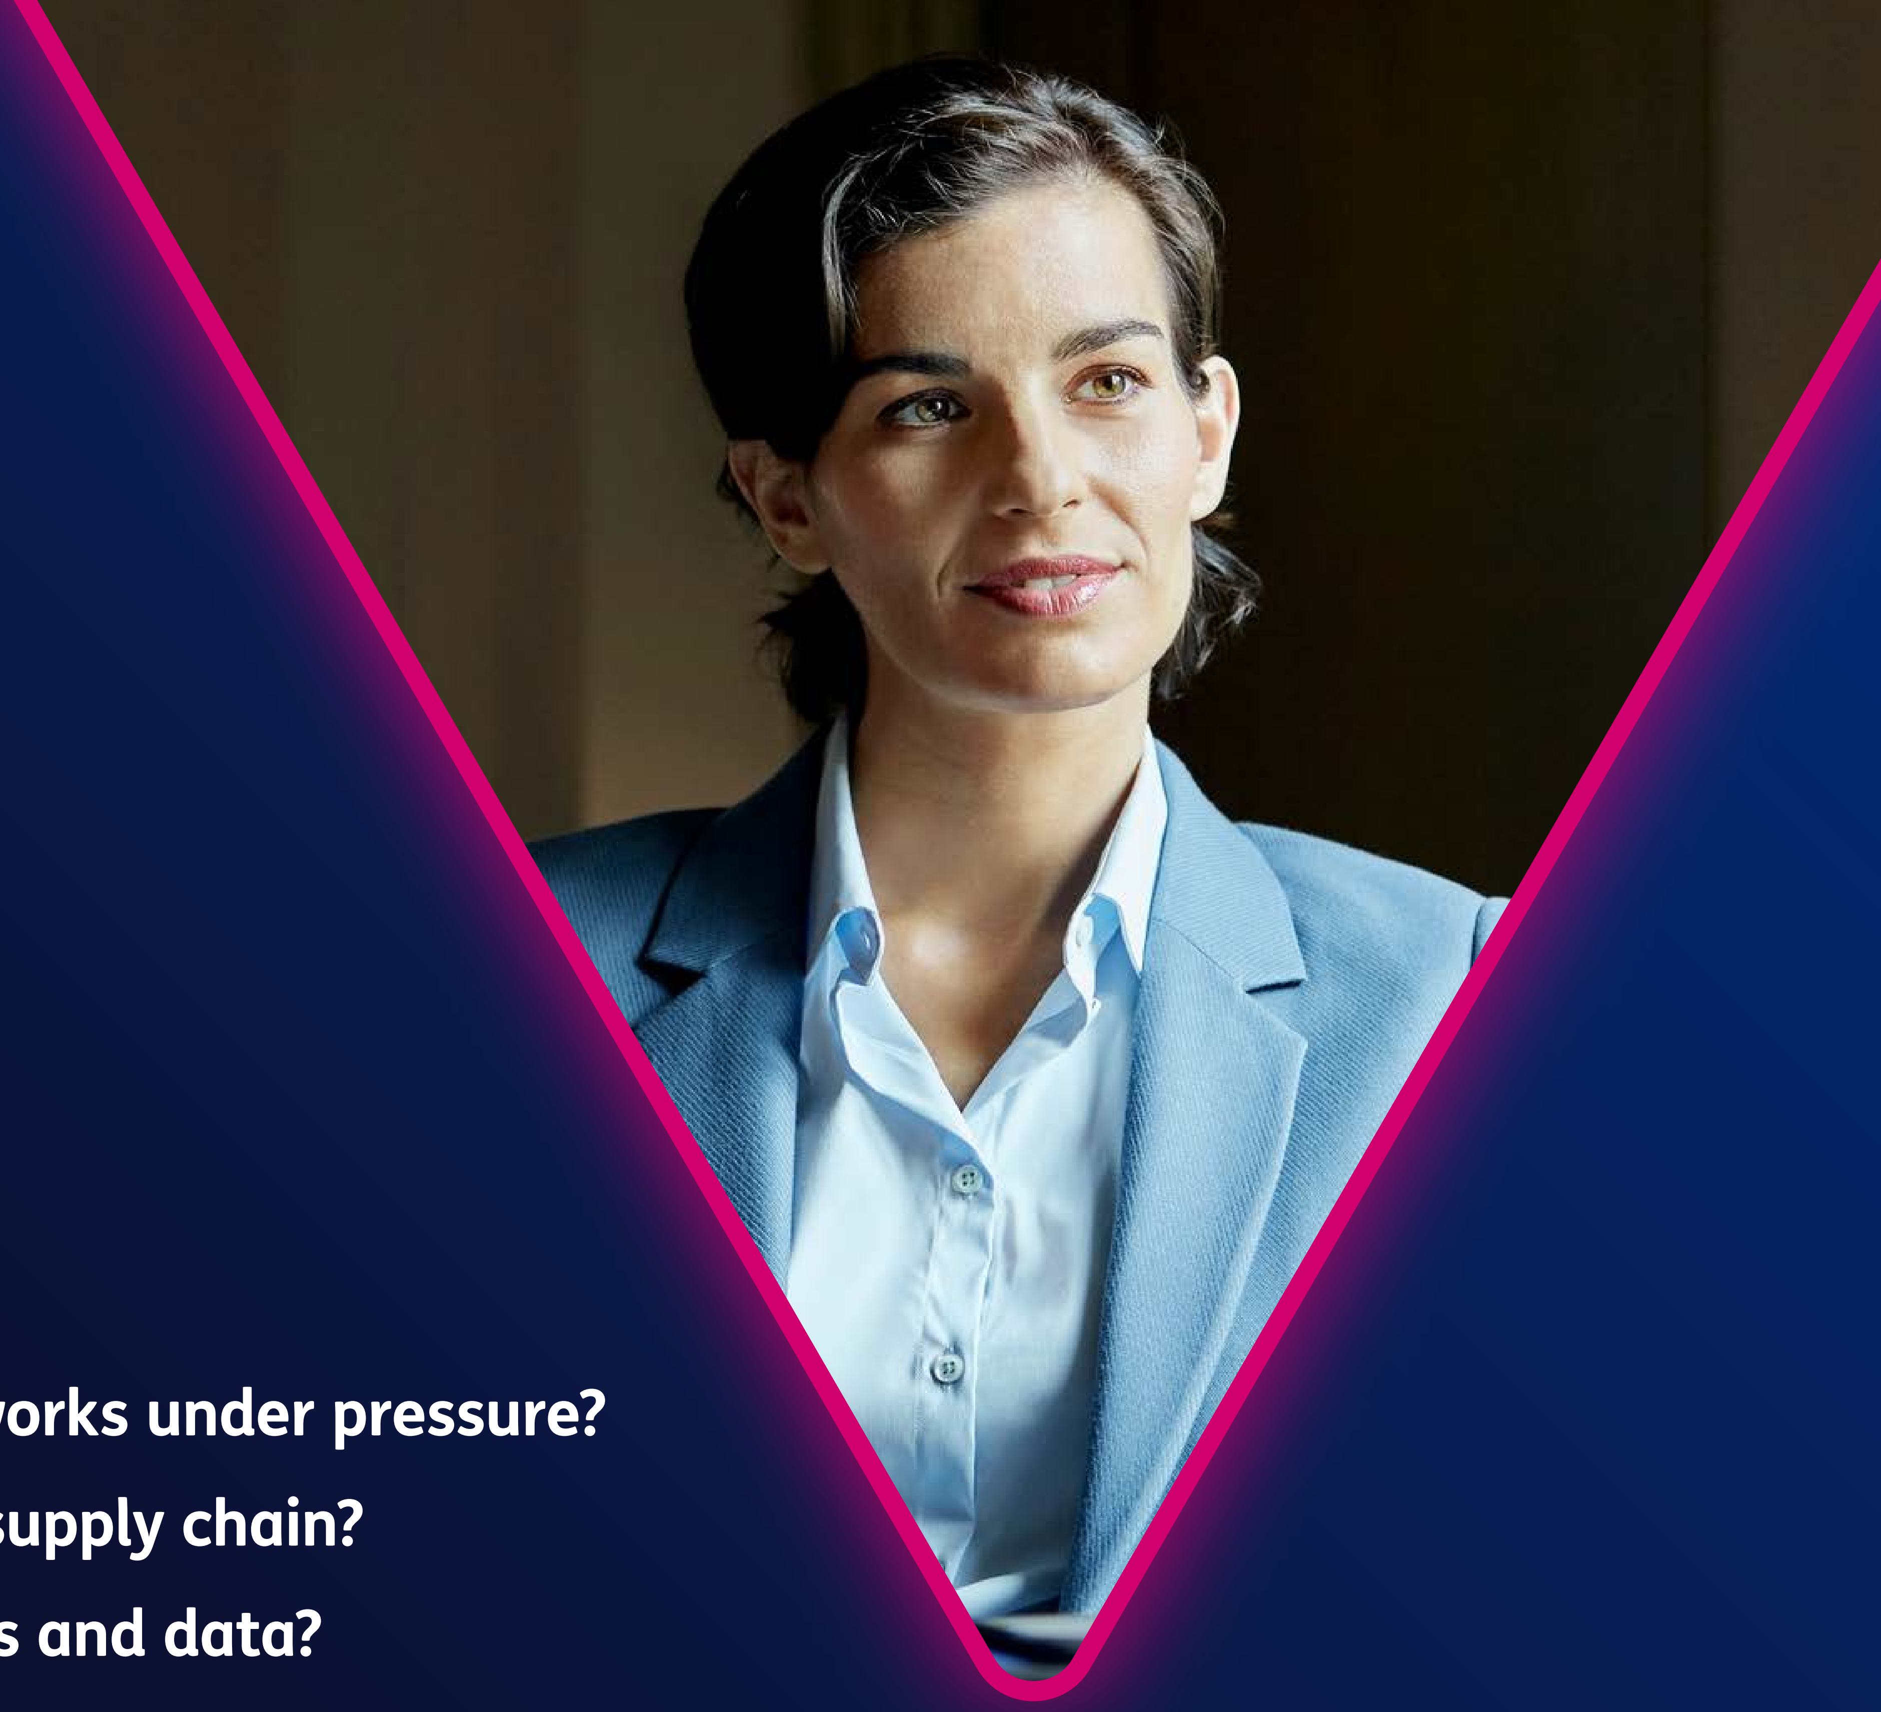

?

Director questions:

Have we proven incident response and recovery works under pressure?

Do we have an entity-wide plan covering our ICT supply chain?

When did we last prove restore for critical services and data?

Are decision rights and escalation thresholds rehearsed and clear?

When did we last test it end-to-end against plausible scenarios, and what evidence shows we improved?

Expected responses:

- We have a single, entity-wide incident response plan with clear roles, decision rights, and escalation, including critical ICT suppliers and recovery contingencies.
- We regularly test response and recovery end-to-end using realistic scenarios (for example ransomware, data theft, DDoS, supplier disruption), including restore testing for critical systems.
- Each exercise produces documented lessons learned, tracked actions with owners and dates, and evidence that improvements were implemented and re-tested.
- We benchmark selectively using external partners and exercises where it increases realism and confidence.

15 Cyber audits and reporting

Cyber reporting is only as strong as the independence and quality of assurance. If teams are assessing their own work, assurance can be difficult to defend. A common challenge is an assurance backlog, with audit findings and remediation actions piling up faster than the team can address them. This can result in high-severity issues staying open for long periods. Or issues being "closed" based on paperwork rather than validated fixes.

Directors should look for standard reviews, consistent quality gates, clear severity definitions, disciplined tracking – and a burn-down plan that gets and keeps the backlog under control.

Director questions:

What evidence can we provide that the validation of our cyber reports, control tests and audit outcomes is independent and effective?

What does our assurance backlog look like (volume, severity, age)?

Do we have a burn-down plan and a steady-state model to prevent recurrence?

Expected responses:

- We have a clear independent review model (e.g., second/third line and/or external) based on a defined scope for what must be independently validated.
- We use consistent review types: design review, operating effectiveness review, program assurance, control validation.
- Reviews follow a documented method, evidence standards and sign-off criteria.
- We have a review calendar and capacity plan. Throughput is measured and managed.
- Findings are consistently rated in terms of severity and risk. Closure is only accepted if a standard set of evidence requirements are met.
- Actions have owners, due dates, dependencies and escalation paths for overdue high-severity items.
- We have clear visibility of the backlog, including its size, severity and aging. Our time-bound burn-down plan shows when we will return to steady-state based on resourcing and prioritisation.

16 Alignment with regulatory frameworks and guidelines

Organisations can be compliant on paper while still being fragile in practice – especially when APRA, SOCI and Essential Eight requirements are treated as separate checklists. The risk isn't just gaps. It's duplication, inconsistent interpretations and weak evidence, which increases exposure during audits, incidents and regulatory scrutiny.

Directors should look for an integrated coherent program: a single control baseline tied to critical services and material risks. And one evidence library that maps to each framework. Otherwise, assurance drifts into a documentation exercise – rather than proof that controls work.

Director questions:

Do we have a single, traceable view of our alignment to applicable obligations and frameworks?

Where are the biggest gaps between documented alignment and operational effectiveness for critical services?

For our critical services, which control requirements are not yet met, what evidence supports that assessment, and what is the funded remediation plan?

Expected responses:

- We maintain a single enterprise control baseline mapped to all applicable obligations and frameworks.
- We can show what meets requirements, what does not yet meet them, and the evidence supporting that view.
- For critical services, we evidence operating effectiveness through testing, independent assurance, and dated remediation closure.
- We have a funded remediation roadmap prioritised by risk and criticality, with owners, dates, and measurable progress.

Conclusion: Resilience starts with asking the right questions

Boards are operating in an environment where cyber expectations are rising and tolerance for disruption is falling. Cyber frameworks, standards and guidelines can be valuable because they set a baseline for governance, control design and assurance – but they only protect the organisation if they translate into operational capability: clear accountability, exercised response and recovery, and evidence that controls work under realistic conditions.

At the same time, organisations are modernising technology estates to improve cost efficiency, productivity and speed. That increases reliance on interconnected platforms, identity services and third parties – expanding the attack surface and the pathways through which disruption can spread. Attackers are also using greater automation, including AI-enabled techniques, which can compress the time between exposure and impact.

Directors should therefore insist on clarity in business terms: which critical services matter most, what disruption is tolerable, what independent assurance exists, and what testing proves plans work in practice. The questions in this guide are designed to move the conversation from activity and spend to measurable risk reduction and demonstrable resilience.

How Baringa can help

Baringa helps boards and executives turn the right cyber questions into measurable risk reduction. We focus on quality governance outcomes: clearer accountability, credible assurance and evidence that investment is reducing exposure in critical services, sensitive data and key third parties.

Some of the areas where we support clients

- Board-ready governance and reporting that clarifies executive accountability and board oversight
- Postquantum and AI-era protection for cryptographic assets and high value IP
- Independent assurance and control validation grounded in disciplined review standards
- Scenario-led cyber risk quantification that informs genuine business decisions
- Operational resilience and crisis readiness aligned to critical services and real events
- SOC and detection performance uplift aligned to material business risks
- Third-party and supply chain assurance that exposes concentration risks and dependency gaps

This document: (a) is proprietary to Baringa Partners LLP ("Baringa") and should not be re-used for commercial purposes without Baringa's consent; (b) shall not form part of any contract nor constitute acceptance or an offer capable of acceptance; (c) excludes all conditions and warranties whether express or implied by statute, law or otherwise; (d) places no responsibility or liability on Baringa or its group companies for any inaccuracy, incompleteness or error herein; and (e) the reliance upon its' content shall be at user's own risk and responsibility. If any of these terms is invalid or unenforceable, the continuation in full force and effect of the remainder will not be prejudiced. Copyright © Baringa Partners LLP 2026. All rights reserved.

Reach out to our team

Douglas Foster
douglas.foster@baringa.com

Melanie Skyes
melanie.sykes@baringa.com

Chris Nott
christopher.nott@baringa.com