Written evidence submitted by Richard Holland to the Cyber Security and Resilience Public Bill Committee (CSRB07)

Submitted by:

A Field Chief Information Security Officer (Field CISO) advising UK public-sector, housing, education and regulated organisations on cyber risk, resilience and operational cyber security.

MP Summary (for Members of Parliament)

This submission provides practitioner-led evidence on how the Cyber Security and Resilience (Network and Information Systems) Bill can be strengthened to deliver measurable, real-world cyber resilience rather than compliance alone.

The Bill rightly seeks to improve cyber resilience across the UK, but in its current form it risks reinforcing a compliance-driven approach that has repeatedly failed to prevent serious disruption. Evidence from public-sector and regulated environments shows that organisations can be compliant with standards, certifications and audits while remaining unable to detect, respond to or recover effectively from cyber attacks.

A central weakness of the Bill is that it does not distinguish clearly between resilience that is documented and resilience that is practised. Cyber resilience is an operational capability that must be exercised regularly under realistic conditions. Where this does not occur, organisations experience delayed detection, unclear decision-making and prolonged recovery during real incidents, even where policies and plans exist.

The Bill also places insufficient emphasis on operational outcomes. It does not require evidence that controls work under attack conditions, nor does it prioritise measurable indicators such as detection speed, response effectiveness or recovery capability. This risks incentivising organisations to optimise for audit success rather than for real-world resilience.

Cyber Essentials and Cyber Essentials Plus are valuable baseline hygiene schemes, but evidence shows they are increasingly misinterpreted as indicators of resilience for medium and large organisations. The Bill does not currently address the false sense of assurance this creates or require proportionate, risk-based assurance at scale.

Framework complexity and recognising sector differences is another material risk. The UK already operates multiple overlapping cyber security frameworks without a coherent national approach without sector-informed profiling. Further, framework proliferation risks increasing administrative burden while diverting resources away from practical resilience improvements. The Scottish national security operation centre (SOC) experience suggests that resilience is better served when frameworks first recognise sector difference and operational reality, with common outcomes aligned where appropriate, rather than

enforcing uniformity from the outset. The Act would therefore benefit from a clearer expectation that cyber security frameworks are applied proportionately and with explicit recognition of sector context as well as reducing the number of them and reduce complexity.

The Bill's scope focuses on Critical National Infrastructure and Managed Service Providers, but does not sufficiently reflect systemic economic and societal impact. Cyber disruption affecting large retailers, logistics providers and service platforms can have national consequences despite these organisations falling outside traditional CNI classifications.

A significant omission is software supply-chain risk. Modern software relies heavily on third-party and open-source components developed internationally. The Log4j vulnerability demonstrated how difficult it is for organisations to identify affected systems, trace dependencies and ensure vulnerabilities are remediated at source. The Bill does not adequately address this systemic risk.

The submission also highlights the need for greater alignment with the EU Cyber Resilience Act, given the EU's role as a major trading partner and defence collaborator, and for the Bill to reflect geopolitical realities where cyber attacks are rapid, state-linked and require swift defensive action.

The Bill underestimates the regulatory capacity required to oversee cyber resilience effectively, particularly as AI-enabled attacks increase the speed and volume of incidents. Expanding regulatory responsibility without equivalent investment in specialist capability risks limiting the Act's effectiveness.

Finally, the Bill also does not explicitly address the cyber security risks introduced by artificial intelligence. AI is already accelerating the speed, scale and sophistication of cyber attacks, while organisations are increasingly dependent on AI-driven security and operational systems. The absence of AI considerations creates a gap in the Bill's ability to reflect the modern threat environment. Rather than introducing AI-specific compliance frameworks, the submission recommends that the Bill address AI risk through practised resilience. This includes requiring organisations to exercise AI-enabled attack scenarios, monitor AI-driven systems for abnormal behaviour, detect data poisoning or ethical bias, and demonstrate the ability to remediate, retrain or recover affected models safely.

The submission concludes with clear recommendations to improve the Bill, including making practised resilience a statutory expectation, focusing regulation on measurable outcomes, reducing framework fragmentation, addressing software supply-chain risk, enabling rapid defensive action during incidents, and ensuring regulatory capability keeps pace with modern cyber threats.

1. Introduction and context

I submit this evidence in my professional capacity as a Field Chief Information Security Officer, with over 25 years' experience, providing independent, practitioner-led advice to Boards and executive teams across UK public services, housing associations, higher education institutions and other regulated environments. My work focuses on cyber resilience, supply-chain risk, managed security services and the operational realities of defending organisations against modern cyber threats.

This submission identifies areas where, in its current form, the Cyber Security and Resilience (Network and Information Systems) Bill risks falling short of its stated objective of materially improving national cyber resilience. The intention is to support the Committee in strengthening the Bill so that it delivers practical, timely and proportionate improvements to the UK's cyber posture.

2. Cyber resilience must be practised, not merely defined

While the Act places appropriate emphasis on cyber resilience, it does not sufficiently distinguish between resilience that is documented and resilience that is exercised. In operational terms, cyber resilience is not a static attribute achieved through compliance or certification, but a capability that must be regularly tested and refined under realistic conditions.

Evidence from real-world cyber incidents consistently shows that organisations with comprehensive policies, certifications and governance structures still experience prolonged disruption when incidents occur. In many cases, incident response and business continuity plans exist but have never been tested end to end, particularly under conditions that involve system unavailability, loss of supplier support or senior decision-making under time pressure.

For example, across public-sector and regulated environments, it is common for organisations to conduct periodic tabletop exercises limited to technical teams, while senior leadership and Boards are not involved. When serious cyber incidents occur, this results in delayed escalation, uncertainty over decision authority, inconsistent communication and hesitation around recovery actions. These delays materially increase operational disruption, financial loss and reputational damage.

Lessons arising from the Scottish National Security Operations Centre, reinforces the point that resilience must be practiced. Centralised detection and monitoring capability can provide valuable visibility, but it does not, in itself, ensure effective response or recovery. Where individual organisations had not practised incident response, escalation and decision-making internally, alerts and intelligence did not consistently translate into timely or decisive action during incidents. Evidence from post-incident reviews also shows that organisations often assumed that participation in a central SOC reduced the need to rehearse their own

response and recovery arrangements. In practice, this led to delays where roles, decision authority and recovery priorities had not been exercised in advance. The result was uncertainty at critical moments, despite the presence of monitoring capability

Similarly, recovery assumptions are frequently untested. Organisations may formally define recovery time objectives, yet have never validated whether critical systems can actually be restored within those timeframes following a destructive cyber attack. In practice, dependencies on suppliers, legacy systems or shared infrastructure are often only discovered during live incidents, when remediation options are limited.

The Act does not currently require organisations to demonstrate that resilience plans have been exercised, lessons identified and improvements implemented. As a result, organisations may appear compliant while remaining operationally fragile. This gap undermines the Act's objective of improving real-world resilience rather than theoretical preparedness.

The Act would therefore benefit from explicitly recognising that cyber resilience must be demonstrated through regular exercising of detection, response and recovery capabilities, including involvement of senior leadership, rather than inferred from the existence of plans, policies or certifications alone.

3. Risk of reinforcing compliance, frameworks and address size and complexity rather than operational outcomes

a. Risk of reinforcing compliance rather than operational outcomes

The Act, as drafted, risks reinforcing an assurance model that prioritises compliance with controls, policies and frameworks over demonstrable operational cyber resilience. In practice, many organisations that are considered "compliant" continue to experience significant cyber incidents because compliance assessments largely focus on whether controls exist, not whether they function effectively during real attacks.

Evidence from post-incident reviews across public-sector and regulated organisations shows a recurring pattern. Organisations are able to demonstrate alignment to recognised standards, completion of annual audits and the existence of formal risk registers, yet still experience prolonged periods of compromise or service disruption. In many cases, security controls are technically present but are not effectively monitored, tuned or integrated into operational processes.

For example, it is common for organisations to meet compliance requirements for logging and monitoring, yet lack the capability to detect malicious activity in a timely manner because logs are not actively reviewed, correlated or acted upon. During incidents, this results in attacks remaining undetected for extended periods despite formal compliance with control requirements.

Similarly, organisations may be able to evidence the existence of incident response plans, escalation procedures and governance structures as part of compliance assessments, yet discover during live incidents that these arrangements are unclear in practice. Decision-making authority may be ambiguous, responsibilities between internal teams and suppliers may be poorly defined, and response actions may be delayed while compliance obligations are interpreted rather than decisive action being taken.

The Act does not currently require organisations to demonstrate that controls operate effectively under attack conditions, nor does it sufficiently emphasise outcomes such as detection speed, response effectiveness or recovery capability. As a result, there is a risk that organisations focus on satisfying regulatory requirements and audit cycles rather than investing in the operational capabilities that materially reduce cyber risk.

Without a clearer distinction between compliance and resilience, the Act may inadvertently incentivise organisations to optimise for audit success rather than for their ability to withstand, respond to and recover from cyber attacks. This risks undermining the Act's stated objective of improving real-world cyber resilience.

b. Risk-based frameworks and sector suitability

The Act anticipates the use of risk-based cyber security frameworks but does not sufficiently address how such frameworks should be selected, interpreted or applied in a manner that reflects sector-specific risk and operational reality. This lack of clarity creates a material risk that organisations will prioritise framework alignment over effective risk reduction.

Evidence from public-sector and regulated environments shows that when frameworks are applied without appropriate sector or organisational context, security effort is often misdirected. Organisations may demonstrate formal alignment to a recognised framework while remaining exposed to their most significant operational risks.

For example, organisations have been required to adopt comprehensive, control-heavy frameworks originally designed for highly regulated or safety-critical environments, despite operating in sectors with limited budgets, legacy technology constraints and different threat profiles. In such cases, significant effort is expended on documenting control compliance and producing assurance artefacts, while critical gaps in detection, response capability or supplier dependency remain unaddressed.

Conversely, other organisations operating in high-risk environments have adopted lighter-touch frameworks that provide insufficient depth for their threat exposure. In post-incident reviews, it is common to find that frameworks were followed as written, yet did not require meaningful assessment of lateral movement detection, incident escalation across suppliers, or recovery from destructive attacks. The framework was technically applied, but operational risk was not reduced.

A further practical issue arises where organisations are required to align simultaneously to multiple frameworks with different structures, terminology and assurance models. In these cases, security teams often focus on mapping controls between frameworks to satisfy assurance requirements rather than investing time in improving real-world resilience. This mapping activity frequently becomes the dominant security workload, particularly in public-sector organisations with constrained resources.

The Act does not currently provide sufficient guidance on how frameworks should be adapted to sector risk, organisational size or systemic impact, nor does it require that framework adoption results in demonstrable improvement in operational outcomes. Without such clarity, there is a risk that framework selection becomes a compliance exercise driven by regulatory interpretation rather than an informed assessment of cyber risk.

The Act would therefore benefit from explicitly recognising that framework selection and application is a technical and sector-specific exercise requiring expert input. Clear expectations that frameworks must be applied proportionately and assessed against operational outcomes would reduce misalignment, avoid unnecessary burden and accelerate the delivery of meaningful cyber resilience.

c. Framework proliferation and complexity without sufficient recognition of sector differences

The Act does not sufficiently address the challenge of applying cyber security frameworks across sectors with materially different risk profiles, operational models and public impact. While framework consolidation may appear attractive in principle, recent public-sector experience suggests that effectiveness depends first on recognising sector difference.

Lessons from the Scottish National SOC highlight that organisations operating within the same central security ecosystem nonetheless faced very different operational realities. Local authorities, health bodies and other public-sector organisations differed significantly in their tolerance for disruption, decision-making structures, legacy technology and dependency on third parties. Applying uniform assurance expectations did not always reflect these differences and, in some cases, obscured the risks that mattered most during live incidents.

More broadly, the UK already operates multiple cyber security frameworks across government and the public sector, including those used by the Ministry of Defence, the NHS and other departments. These frameworks exist because sector context matters. Attempts to overlay or consolidate frameworks without first acknowledging sector-specific risk can result in compliance activity that does not translate into operational resilience.

Evidence from public-sector assurance activity shows that organisations frequently spend significant effort mapping controls between frameworks rather than addressing the most pressing operational risks. This is particularly acute where frameworks are applied prescriptively rather than adapted to sector context.

The Act does not currently provide sufficient clarity on whether framework coherence should be achieved through consolidation or through sector-informed profiling. The Scottish experience suggests that resilience is better served when frameworks first recognise sector difference and operational reality, with common outcomes aligned where appropriate, rather than enforcing uniformity from the outset.

The Act would therefore benefit from a clearer expectation that cyber security frameworks are applied proportionately and with explicit recognition of sector context, before pursuing broader consolidation. This would reduce misalignment, improve adoption and ensure that framework use supports, rather than distracts from, real-world resilience.

d. Limitations of Cyber Essentials at scale

Cyber Essentials and Cyber Essentials Plus provide an important baseline for basic cyber hygiene and are well suited to small organisations and entry-level supply-chain assurance. However, the Act does not sufficiently recognise the limitations of these schemes when applied to medium and large organisations or those delivering high-impact services.

In practice, Cyber Essentials assessments are point-in-time validations of configuration and control presence. They are typically conducted annually and focus on preventative measures such as patching, access controls and malware protection. While these controls are necessary, they do not assess whether an organisation can detect malicious activity, respond effectively to an incident or recover critical services following a cyber attack. Further, larger organisation with complex environment and 1000s of end points will find it incredible difficult to maintain and remove all vulnerabilities and so seek to ringfence and segregate any end point which requires CE, therefore allow any machine outside the ringfence to not adhere to the same stringent requirement which doesn't remove the risk especially from unilateral movement.

Evidence from public-sector and regulated environments shows that organisations holding Cyber Essentials or Cyber Essentials Plus certification have still experienced ransomware incidents, data compromise and prolonged service disruption. In post-incident reviews, it is common to find that baseline controls were technically in place and compliant at the time of assessment, yet attackers were able to operate undetected for extended periods due to the absence of effective monitoring, alerting and response capability.

For example, organisations may meet Cyber Essentials requirements for patch management and endpoint protection, yet lack the capability to identify lateral movement, privilege escalation or data exfiltration once an attacker gains an initial foothold. Similarly, Cyber Essentials does not assess whether incident response plans have been exercised, whether recovery objectives are achievable in practice, or whether suppliers and managed service providers are prepared to support recovery during a real incident.

At Board and procurement level, Cyber Essentials certification is frequently interpreted as evidence that an organisation is “secure” or “resilient”. This misinterpretation can result in over-confidence, reduced scrutiny and under-investment in operational resilience capabilities. The Act does not currently address this risk or provide guidance on the appropriate use of Cyber Essentials as a baseline rather than a measure of resilience.

Without clearer differentiation, there is a risk that the Act legitimises the use of baseline certification as a substitute for proportionate, risk-based cyber resilience in larger or higher-impact organisations. This would undermine the Act’s objective of improving real-world resilience rather than compliance alone.

The Act would therefore benefit from explicitly positioning Cyber Essentials and Cyber Essentials Plus as foundational hygiene measures for smaller organisations, while setting clearer expectations that medium and large organisations demonstrate resilience through risk-based approaches, operational capability and exercised response and recovery.

4. Scope and systemic economic impact

While the Act’s focus on Critical National Infrastructure and Managed Service Providers is appropriate, it does not sufficiently account for organisations whose disruption would cause significant economic or societal impact despite not meeting traditional definitions of CNI. This represents a material gap in addressing modern, digitally interconnected risk.

Evidence from recent cyber incidents across the UK and internationally demonstrates that disruption to large consumer-facing organisations can have effects comparable to, and in some cases exceeding, those associated with traditional infrastructure outages. Large retailers, logistics providers, food supply chains and payment platforms are deeply embedded in everyday economic activity. Cyber disruption affecting such organisations can rapidly cascade into shortages, delayed services, loss of consumer confidence and significant financial harm.

For example, cyber incidents affecting major retail or logistics operations have resulted in widespread service disruption, manual workarounds across supply chains and knock-on impacts for smaller dependent businesses. While such organisations are not formally designated as CNI, their digital platforms function as critical enablers of national economic activity. The absence of proportionate resilience obligations for these organisations creates a disparity between impact and regulatory expectation.

Similarly, large service platforms supporting housing, benefits administration or consumer finance may not fall within traditional CNI classifications, yet disruption can directly affect vulnerable populations and place immediate pressure on public services. In these cases, cyber incidents create both economic and societal consequences, even though the organisations involved sit outside the current scope of the Act.

The Act's current approach risks underestimating systemic risk by equating "criticality" solely with sector classification rather than with real-world impact. This may lead to situations where organisations with high systemic importance are subject to minimal resilience expectations, while others with lower economic impact are more tightly regulated.

An impact-based approach, considering factors such as scale, dependency, substitutability and potential for cascading effects, would better reflect the realities of a digital economy. Without such an approach, the Act risks leaving significant sources of national risk insufficiently addressed, contrary to its stated objective of strengthening overall cyber resilience.

5. Managed service providers and systemic dependency

Managed Service Providers, including managed detection and response, cloud service operators and outsourced IT providers, are now critical enablers of cyber resilience across the UK economy. However, the Act does not sufficiently address the systemic risk created by widespread dependency on a relatively small number of providers with highly variable operational maturity.

Evidence from public-sector and regulated environments shows that organisations increasingly rely on managed services for core security and operational functions, often assuming that the presence of a contract equates to resilience. In practice, many organisations have limited visibility into how their providers would perform during a real cyber incident, particularly under conditions of scale, concurrency or supplier stress.

For example, post-incident reviews frequently reveal that while monitoring services were in place, alerts were delayed, poorly prioritised or not escalated effectively during active attacks. In some cases, organisations believed incidents were being actively managed by their provider, only to discover that responsibility boundaries were unclear or that response actions required explicit customer approval, resulting in critical delays.

A further recurring issue arises during large-scale or sector-wide incidents. When multiple customers of the same provider are affected simultaneously, provider capacity becomes a limiting factor. Organisations often assume guaranteed response, yet discover during incidents that provider resources are shared and prioritised, with no clear transparency on how competing demands are managed.

Dependency risk is also amplified by subcontracting and offshore delivery models. In practice, organisations may contract with a UK-based provider but rely on security operations, development or support functions delivered through complex supply chains. During incidents, this can introduce additional latency, communication challenges and uncertainty over accountability, particularly where data sensitivity or national security considerations are involved.

The Act does not currently require organisations or providers to demonstrate exercised operational resilience, transparency of dependencies or readiness to operate under stress. Without such expectations, reliance on managed services may concentrate risk rather than reduce it, particularly in sectors where multiple organisations depend on the same small group of providers.

The Act would therefore benefit from clearer expectations that managed service providers demonstrate operational resilience in practice, including their ability to respond effectively during concurrent incidents, manage dependencies transparently and support customers during high-impact cyber events. Without addressing this, the Act risks underestimating a key source of systemic cyber risk.

6. Software supply-chain and open-source risk

The Bill does not adequately address the practical difficulty organisations face in managing cyber risk arising from complex software supply chains, particularly where software products incorporate third-party and open-source components developed outside the direct control of the consuming organisation.

Modern software used across the UK economy is frequently developed internationally, including in the United States, and commonly relies on open-source libraries maintained by distributed communities rather than single accountable vendors. While open-source software provides significant innovation and economic benefit, it also introduces systemic cyber risk when secure development practices, maintenance obligations and accountability are unclear.

The Log4j vulnerability provides a clear illustration of this challenge. The affected code library was developed and maintained as an open-source component, embedded deep within a wide range of commercial and bespoke software products. When the zero-day vulnerability was disclosed, many organisations were unable to quickly determine whether they were affected because they did not have visibility into the full software composition of the systems they relied upon.

Even where organisations were able to identify affected software, remediation proved difficult. Responsibility for fixing the vulnerability often sat several layers removed from the organisation, requiring action by software vendors, who in turn depended on upstream maintainers. In many cases, organisations had limited leverage or assurance that vulnerabilities were being addressed at source, despite the significant operational and national risk posed.

The Act does not currently provide sufficient mechanisms to address this class of risk. In particular, it does not place adequate emphasis on software transparency, secure development lifecycle assurance, or accountability for third-party and open-source

components embedded within critical software. As a result, organisations may remain exposed to systemic vulnerabilities that are outside their immediate control, despite otherwise strong cyber security practices.

Without addressing software supply-chain risk explicitly, the Act risks focusing on organisational controls while leaving a significant source of modern cyber risk insufficiently mitigated. This gap is increasingly material given the scale of software reuse, the speed at which vulnerabilities can be exploited, and the difficulty of tracing and remediating issues across complex supplier ecosystems.

The Act would therefore benefit from stronger recognition of software supply-chain risk, including expectations around transparency of software components, accountability for remediation at source, and proportionate assurance of secure development practices across suppliers and their dependencies.

7. EU alignment and international interoperability

The Act does not give sufficient consideration to alignment with the European Union Cyber Resilience Act, despite the EU remaining one of the United Kingdom's largest trading partners and a close collaborator in defence, security and technology supply chains. This lack of alignment risks creating unnecessary complexity for organisations operating across UK–EU markets and may delay the delivery of practical cyber resilience improvements.

Evidence from organisations operating in regulated and public-sector supply chains shows that many UK-based organisations already align their cyber security practices to EU requirements because of commercial necessity. Technology manufacturers, software vendors, defence suppliers and digital service providers frequently supply both UK and EU markets. In practice, these organisations are required to meet EU cyber security expectations relating to secure-by-design principles, product lifecycle assurance and vulnerability management, even where UK regulation does not explicitly require the same measures.

Where UK and EU regimes diverge, organisations are often forced to maintain parallel assurance processes. For example, suppliers may be required to demonstrate secure development lifecycle practices and vulnerability disclosure mechanisms to satisfy EU obligations, while simultaneously evidencing organisational controls and governance structures to satisfy UK regulatory expectations. This duplication consumes security and engineering capacity without proportionate benefit, delaying the implementation of controls that would otherwise reduce risk.

In defence-adjacent and national security supply chains, misalignment creates additional friction. UK organisations frequently participate in multinational programmes where cyber security requirements are harmonised across allied nations. Divergent regulatory

approaches increase contractual complexity and create uncertainty over which standards take precedence, particularly where software and hardware components are developed, maintained or updated across borders.

From an operational perspective, the lack of alignment also affects incident response and vulnerability management. The EU Cyber Resilience Act places strong emphasis on coordinated vulnerability handling, transparency and remediation at source. Where UK legislation does not mirror these expectations, organisations may face inconsistent reporting obligations and uncertainty about how to prioritise remediation actions across jurisdictions during active cyber incidents.

The Act's current approach risks positioning UK cyber regulation as an additional compliance layer rather than a complementary one. This may discourage investment, slow adoption and undermine the Act's objective of improving cyber resilience in practice rather than in theory.

Greater alignment with the principles and outcomes of the EU Cyber Resilience Act would reduce friction for UK organisations, support international interoperability and strengthen shared resilience across economically and strategically linked sectors. Without such alignment, the Act risks increasing regulatory burden while delivering limited additional resilience benefit.

8. Geopolitics and the need for rapid response

The Act does not adequately reflect the impact of geopolitical events on cyber threat activity, nor does it sufficiently recognise the need for organisations to act at speed when responding to cyber incidents driven by state-sponsored or geopolitically motivated actors. This represents a material gap in the Bill's ability to support national cyber resilience in periods of heightened international tension.

Evidence from recent years demonstrates a clear correlation between geopolitical events and increased cyber activity targeting government, public-sector and economically significant organisations. During periods of conflict, diplomatic tension or sanctions, cyber operations are frequently used as a tool to disrupt services, gather intelligence or undermine public confidence. Such activity is often rapid, coordinated and designed to exploit hesitation, uncertainty or procedural delay within target organisations.

In practice, organisations facing geopolitically driven cyber attacks often need to take decisive defensive actions within hours rather than days. These actions may include isolating systems, suspending services, blocking access from specific regions, altering operational configurations or engaging specialist incident response support at short notice. Delays caused by uncertainty over regulatory expectations, reporting obligations or contractual authority can materially increase the scale and impact of an attack.

Post-incident reviews commonly show that organisations hesitated to take decisive action during the early stages of an incident due to concerns about regulatory compliance, data protection implications or the perceived requirement to seek approval from multiple internal and external stakeholders. In several cases, this hesitation allowed attackers to deepen persistence, expand lateral movement or exfiltrate data, resulting in greater disruption and recovery time.

The Act does not currently provide sufficient clarity on how organisations can take rapid, good-faith defensive action during live cyber incidents without fear of regulatory penalty. Nor does it clearly align with the principles underpinning the United Kingdom's cyber force, which recognise cyber as an operational domain requiring speed, adaptability and proportionate response.

Without clearer recognition of the need for agility during geopolitically driven cyber incidents, there is a risk that the Act inadvertently constrains the very behaviours required to defend systems effectively. This risks undermining national resilience at precisely the moments when cyber threats are most acute.

The Act would therefore benefit from explicitly supporting rapid, proportionate defensive action during cyber incidents, aligned with national cyber defence principles, while maintaining appropriate accountability and oversight once incidents are stabilised.

9. Regulatory capacity and pace of technical change

The Act underestimates the scale of regulatory capability required to oversee cyber resilience effectively, particularly given the pace of technological change and the increasing sophistication of cyber threats. While the expansion of regulatory responsibility, including to the Information Commissioner's Office, is well intentioned, it represents a material increase in scope, technical complexity and operational demand.

Evidence from existing regulatory and incident-response activity indicates that cyber incidents now occur with greater frequency, complexity and speed than traditional regulatory models were designed to handle. In practice, regulators are increasingly required to assess not only whether organisations had appropriate policies and governance in place, but whether technical controls operated effectively in real time, whether response actions were appropriate under pressure, and whether systemic risks were adequately managed.

For example, during significant cyber incidents, regulators may receive large volumes of incident notifications within a short period. Assessing these incidents meaningfully requires specialist technical expertise to distinguish between minor events and those with genuine systemic or national significance. Without sufficient cyber-specialist capability, there is a risk that regulatory effort is consumed by processing notifications rather than prioritising the incidents that pose the greatest risk.

A further challenge arises in post-incident assessment. Modern cyber incidents frequently involve complex attack chains, cloud services, managed service providers and third-party software dependencies. Evaluating whether an organisation acted reasonably and proportionately during such incidents requires deep understanding of threat actor behaviour, detection limitations, response trade-offs and recovery constraints. Traditional regulatory skill sets focused on governance, compliance and legal interpretation are not sufficient on their own to make these assessments consistently or at speed.

The pace of technical change compounds this challenge. The increasing use of artificial intelligence by threat actors is accelerating attack automation, phishing effectiveness and exploitation speed. At the same time, defensive technologies and operating models are evolving rapidly. Regulators without continuous access to current threat intelligence, tooling and specialist cyber expertise risk lagging behind the environments they are tasked with overseeing.

There is also evidence that organisations facing active cyber incidents are sometimes uncertain how regulators will interpret their actions, particularly where rapid defensive measures have operational or data protection implications. This uncertainty can lead to overly cautious behaviour, delayed response or excessive internal escalation, increasing the impact of incidents. Clear, informed and timely regulatory guidance is therefore an essential component of national cyber resilience.

The Act does not currently provide sufficient assurance that regulatory capacity, skills, tooling and operating models will scale in line with the expanded responsibilities it introduces. Without explicit consideration of these factors, there is a risk that regulatory effectiveness is constrained by bandwidth and capability rather than policy intent, undermining confidence in the Act and slowing its practical impact.

The Act would therefore benefit from explicit recognition that effective cyber resilience regulation requires sustained investment in specialist technical capability, access to current threat intelligence and operating models that reflect the speed and complexity of modern cyber threats.

11. The Act does not adequately address the cyber security risks introduced by artificial intelligence

The Act does not explicitly address the cyber security risks introduced by the increasing use of artificial intelligence by both threat actors and defenders. This represents a significant omission, given the rapid adoption of AI technologies across the UK economy and the accelerating use of AI-enabled techniques in cyber attacks.

From an operational perspective, artificial intelligence is already materially changing the cyber threat landscape. Threat actors are using AI to automate reconnaissance, generate highly targeted phishing campaigns, identify vulnerabilities at scale and adapt attack

techniques more rapidly than traditional manual methods. This has lowered the barrier to entry for sophisticated attacks and increased both the volume and speed of cyber incidents affecting organisations of all sizes.

Evidence from recent incidents shows that AI-enabled phishing and social engineering attacks are significantly harder for users and traditional security controls to detect. Messages are increasingly personalised, context-aware and linguistically accurate, reducing the effectiveness of awareness training and basic filtering controls. The Act does not currently consider how organisations are expected to adapt resilience measures to account for this shift.

Artificial intelligence also introduces new risks on the defensive side. Many organisations now rely on AI-driven security tooling for threat detection, prioritisation and automated response. While these tools provide clear benefits, they also introduce dependencies on opaque models, training data quality and vendor-managed algorithms. In practice, organisations often lack visibility into how AI-driven decisions are made, how models are updated, or how bias, drift or model failure might affect security outcomes during an incident.

There is also growing evidence of emerging attack techniques that specifically target AI-enabled systems, including data poisoning, model manipulation and abuse of automated decision-making. These risks are not theoretical. They are already being explored in the context of security tooling, fraud detection systems and automated access controls. The Act does not currently address how such risks should be governed, assessed or mitigated.

A further challenge arises from the widespread integration of AI capabilities into third-party products and managed services. Organisations may be consuming AI-driven functionality without clear understanding of where AI is used, how it is secured, or how failures would be handled during a cyber incident. This compounds existing supply-chain risk and complicates incident response, particularly where automated actions have operational or data protection implications.

The absence of explicit consideration of AI risk creates a gap between the Act and the reality of the modern cyber threat environment. Without addressing AI-enabled attack acceleration, AI-driven defensive dependency and emerging risks such as model poisoning, the Act risks being backward-looking at a time when cyber threats are evolving most rapidly.

The Act would therefore benefit from explicit recognition that artificial intelligence materially alters cyber risk, and from clearer expectations around governance, transparency and resilience where AI is used within security controls, operational systems and critical services.

11. Recommendations to improve the Act

Based on the evidence and examples set out in this submission, the following changes are recommended to ensure the Act delivers measurable, real-world cyber resilience rather than compliance alone.

Make exercised cyber resilience a statutory expectation

The Act should explicitly require organisations within scope to **demonstrate that cyber resilience arrangements are practised**, not merely documented.

This should include a clear expectation that organisations regularly exercise cyber incident detection, response and recovery under realistic conditions, and can evidence that lessons identified have been acted upon. Without this requirement, resilience remains theoretical and untested, as demonstrated by repeated post-incident failures across public-sector and regulated organisations.

Require outcome-based evidence of resilience, not just control compliance

The Act should require organisations to evidence **operational outcomes**, not simply the presence of controls, policies or certifications.

This should include measurable indicators such as the ability to detect incidents in a timely manner, respond effectively under pressure, and restore priority services within achievable recovery timeframes. This change is necessary to avoid reinforcing a compliance-led assurance model that has repeatedly failed to prevent prolonged disruption despite formal adherence to standards.

Clearly position Cyber Essentials as baseline hygiene only

The Act should explicitly position Cyber Essentials and Cyber Essentials Plus as **baseline cyber hygiene schemes suitable for small organisations and entry-level assurance only**.

For medium and large organisations, or those with high systemic or societal impact, the Act should require proportionate, risk-based resilience assurance beyond Cyber Essentials. This would address the false sense of security currently created when baseline certification is misinterpreted as evidence of resilience.

Establish expert-informed selection and application of risk-based frameworks

The Act should require that the selection and application of cyber security frameworks be **informed by structured input from cyber security practitioners and sector experts**.

Framework designation should be proportionate to sector risk, organisational scale and systemic impact. This would prevent misalignment where organisations either over-engineer compliance or under-protect critical operational risks, as evidenced across public-sector and regulated environments.

Reduce framework proliferation through a coherent national model

The Act should commit to **reducing framework fragmentation**, particularly across government and the public sector.

This should be achieved through a **single, outcome-focused UK cyber resilience framework**, aligned to existing international standards, with sector-specific profiling rather than bespoke divergence. This would reduce duplication seen across MOD, NHS and other public-sector frameworks, freeing resources to improve operational resilience, whilst recognising the experience from the Scottish National SOC.

Extend scope based on systemic economic and societal impact

The Act should move beyond a purely CNI-based scope and incorporate **impact-based criteria**, enabling proportionate resilience obligations for organisations whose disruption would cause significant economic or societal harm.

This change is necessary to reflect the realities of a digitally interconnected economy, where large retailers, logistics providers and service platforms function as de-facto critical infrastructure despite falling outside traditional classifications.

Strengthen requirements for managed service provider resilience

The Act should require both organisations and managed service providers to **demonstrate exercised operational resilience**, including their ability to respond effectively during concurrent or large-scale incidents.

This should include transparency of dependencies, response capacity under stress and clarity of accountability during incidents. Without this, reliance on managed services risks concentrating systemic cyber risk rather than reducing it.

Explicitly address software supply-chain and open-source risk

The Act should include explicit provisions addressing **software supply-chain risk**, including expectations for transparency, traceability and accountability for third-party and open-source components embedded in critical software.

This should include requirements to ensure vulnerabilities are identified and remediated at source wherever possible. The Log4j vulnerability demonstrated that organisational controls alone are insufficient to manage this class of systemic risk.

Align the Act more closely with the EU Cyber Resilience Act

The Act should pursue alignment, where appropriate, with the principles and outcomes of the EU Cyber Resilience Act to reduce duplication, improve interoperability and strengthen shared resilience across UK–EU supply chains, particularly in defence and technology sectors.

Misalignment risks increasing compliance burden without delivering additional resilience benefit.

Enable rapid, good-faith defensive action during cyber incidents

The Act should explicitly support organisations in taking **rapid, proportionate defensive action** during live cyber incidents without fear of regulatory penalty, provided actions are taken in good faith to reduce harm.

This should be aligned with the principles underpinning the UK cyber defence force, recognising cyber as an operational domain where speed and adaptability are essential to effective defence.

Ensure regulatory capability scales with responsibility

Any expansion of regulatory responsibility introduced by the Act must be matched by **explicit investment in specialist cyber capability, threat intelligence and modern operating models**. Without this, regulatory effectiveness risks being constrained by capacity and skills gaps, particularly as AI-enabled cyber threats accelerate incident volume and complexity.

Recommendation on addressing artificial intelligence risk through practised resilience

The Act should explicitly recognise that artificial intelligence materially accelerates cyber risk and introduces additional failure modes, including data poisoning, model manipulation and ethical bias. Rather than creating AI-specific compliance frameworks, the Act should address these risks through practical resilience expectations. Organisations should be required to consider AI-enabled attack scenarios within cyber resilience exercising, including automated phishing, rapid exploitation and the degradation of AI-driven decision-making. Where AI or algorithmic models are used within security controls, operational systems or critical services, organisations should be expected to demonstrate the ability to monitor for abnormal model behaviour, detect potential poisoning or bias, and take timely corrective action. This should include the capability to suspend, retrain, roll back or replace affected models and to recover services safely where automated decisions have undermined trust, accuracy or security. By focusing on monitoring, remediation and recovery of AI-enabled systems as part of routine resilience practice, the Act can remain adaptable while ensuring resilience expectations reflect the realities of an AI-accelerated threat environment.

12. Conclusion

The Cyber Security and Resilience Bill is necessary and timely. However, without addressing practised resilience, systemic software supply-chain risk, geopolitical realities and regulatory capability, there is a risk that it prioritises compliance over capability.

Addressing these issues would materially strengthen the Bill's ability to deliver durable improvements in UK cyber resilience and better protect the UK economy, public services and national security.

January 2026