

CREST WHITEPAPER

Beyond keeping pace: **Adaptation in the age of AI**

October 2026

How must the cybersecurity ecosystem adapt, not simply to keep pace with AI, but to help shape the society it enables?

This CREST whitepaper explores how providers, buyers, regulators and standards bodies must evolve to build a secure, trusted and resilient AI-enabled future.



CONTENTS

03

EXECUTIVE SUMMARY

04–
06

SHORT-TERM IMPLICATIONS

07–
09

LONGER-TERM IMPLICATIONS

10

CONCLUSION

12–
18

APPENDICES 1–3

EXECUTIVE SUMMARY

The precise direction of AI and the corresponding impact on cyber is highly uncertain, but the effects are being seen already.

High cost, lower value, repetitive tasks are rapidly being automated and accelerated by AI. Under all plausible scenarios for the future, AI capability will increase with both positive and negative impacts – for society and for security.

Will the widespread adoption and integration of AI lead to the degradation and ultimate obsolescence of cyber security services and professionals as we know them, a radical shift in the technical and commercial model of providers or broadly the continuation of the status quo?

This whitepaper considers both the near- and longer-term future for the cyber security industry. Our near-term assessment draws on the plans of CREST's 550 global member companies, the impact AI is having on the threat activities they are observing, and the buyer and regulatory requirements that they are receiving or anticipate. For the longer term, we draw on the Scenarios for AI in 2030 produced by the UK Government Office for Science's Foresight Team and consider the likely implications for, and impact on, our industry.¹

Whilst our analysis of the likely scenarios find that AI is less likely to make the requirement for cybersecurity services disappear wholesale, under each scenario major change will and must happen. This means that service providers must be adept at transforming their businesses, their services and their commercial models

in order to adapt their human-led businesses at pace in a world where the attack surface, malicious capability and exploitable vulnerabilities grow at AI-speed and scale. It is equally clear that the more benign scenarios – where the productivity and scientific transformation benefits of AI can be enjoyed in a safer future world where potential harms are minimised and humans remain 'in the loop' for most tasks, creating more high value jobs – globally understood and widely enforced standards play a key role in ensuring technical safety and mitigating malicious usage.

CREST was formed 20 years ago to build trust in the digital world. We commit now to continue to be the trust layer for cybersecurity in the age of AI.

We will continue to define what good looks like when services are delivered through a combination of informed leaders, skilled practitioners, governed processes, validated tools, models, data and automation, under clear accountability.

“There is some risk involved in action, there always is. But there is far more risk in failure to act.”

– Harry S Truman, 1948

¹ Scenarios for AI in 2030 produced by the UK Government Office for Science:
<https://www.gov.uk/government/publications/ai-scenarios-2030-helping-policymakers-plan-for-the-future-of-ai/ai-scenarios-2030-helping-policymakers-plan-for-the-future-of-ai>

1. SHORT-TERM IMPLICATIONS

As the cyber eco-system – buyers of cyber security services ('buyers'), cyber security service providers ('providers') and national and regulatory authorities and other standard setting bodies ('standard setters') – adapts to the impacts of AI, we see different implications in the short-term (6-12 months) versus the longer-term (based on GOS experts' scenarios for 2030).

CREST members of all sizes and from all regions are already seeing a significant increase in requests for support for safe adoption of AI², including:

- **Policies for AI use**
- **AI asset discovery**
- **Securing and testing Generative AI (GenAI) or Large Language Model (LLM) enabled applications**
- **Secure coding with AI and related testing³**
- **Data leakage controls re: AI access to data; and**
- **Independent assurance of AI vendor risk**

This increase is likely to continue at an escalating pace. In parallel we are seeing challenge from the buyers of cyber security services ('buyers') to cyber security service providers ('providers') in the form of:

- **Downwards pressure** on price due to the expectations that providers pass on the savings from AI
- **Competition from market entrants** focused on AI-enabled or autonomous agents replacing traditional human-led services
- **Demand for independent assurance** of the quality and safety of AI enabled services and tools
- **Demand for legal liability disclaimers** where providers use AI

1.1 Demand side impact

As organisations embed AI across software development, business operations and customer facing systems, they are creating new forms of demand for cybersecurity expertise. Buyers are not only asking how to defend against AI-enabled threats, but they are also asking how to secure their own use of AI, how to understand where AI is already being used, and how to gain confidence in systems that may be changing continuously. This is already creating a broader assurance challenge. AI-generated code, AI-assisted development and automated configuration are increasing the need for verification. The speed at which software can now be produced does not remove the need for established secure development practices. It increases the importance of understanding whether code, dependencies and configurations can be trusted. In practice, this is likely to shift the emphasis from traditional manual review towards faster, AI-enabled analysis, remediation and assurance, supported by human validation and expert sign-off.

The same pattern is emerging across enterprise AI adoption. As organisations deploy generative AI, large language model applications, autonomous agents, model APIs, retrieval-augmented generation systems and third-party AI tools, the attack surface becomes harder to define and harder to govern. Security teams need to know where AI is being used, what data it can access, which models are being relied upon, how outputs are controlled and who is accountable when something goes wrong.

² CREST Research Report: AI in Penetration Testing 2026: crest-approved.org/ai-in-penetration-testing/

³ The DORA 2024 report found AI adoption can improve individual productivity and satisfaction, but it also observed negative effects on delivery stability and throughput, reinforcing the need for robust engineering and testing fundamentals. (Dora)

1. SHORT-TERM IMPLICATIONS

1.2 Supply side impact

This creates an opportunity for providers to develop new specialist services around secure AI adoption. AI asset discovery, AI threat modelling, prompt-injection testing, data-leakage testing, guardrail assessment, model supply-chain review and AI governance advisory are all likely to become more important as organisations move from experimentation to operational dependence. AI security is unlikely to remain a narrow technical discipline, it is becoming part of how organisations assure software, data, identity, governance and third-party risk.

At the same time, providers are adapting their own services. For example, AI is being used to accelerate analysis, automate routine activities and increase the speed of service delivery. Secure design, penetration testing, red teaming, managed detection and response, SOC operations and incident response are being reshaped by AI-enabled tooling. Providers can triage more quickly, summarise investigations, generate detection rules, process threat intelligence and support remediation at greater speed.

But this does not make human expertise less important, it changes where that expertise is most valuable. As AI accelerates the production of findings, experienced practitioners are needed to validate outputs, manage false positives, interpret context, prioritise risks and help clients make practical decisions.

The industry becomes less focused on producing information and increasingly focused on proving which information matters. We anticipate new roles emerging to design secure architectures and cybersecurity discipline specific scaffolding and guardrails.

This is especially important as attackers adopt the same technologies. AI is increasing the speed, scale and sophistication of vulnerability discovery, phishing, social engineering and attack-surface reconnaissance. The time between disclosure and exploitation is already under pressure, and organisations may have less time to assess risk, deploy patches and contain attacks before they are weaponised. In this environment, buyers place greater value on providers that can combine speed with judgement: rapid analysis, clear prioritisation and practical remediation advice.

1.3 Commercial engagement impact

The commercial model for cybersecurity services is also likely to change. Some buyers will expect AI-enabled efficiencies to reduce the cost of services, others will be more cautious, particularly where sensitive data, regulated environments or critical systems are involved.

Providers may therefore need to offer clearer choices about how AI is used in service delivery, from AI-enabled services with human oversight to services that explicitly limit or exclude AI use.

Many of our members tell us that they are explicitly designing twin-track services – with and without AI – to enable them to continue to serve buyers who cannot or will not currently take the risk of AI in their supply chain.

Buyers will need confidence not only in the AI systems they are adopting, but also in the AI being used by their providers. They will want to understand how tools are selected, how outputs are validated, how client data is protected, where human judgement is applied and who remains accountable for the final decision.

1. SHORT-TERM IMPLICATIONS

Existing cybersecurity standards, including CREST's own, are evolving to reflect AI-assisted delivery, including tool validation, evidence quality, audit trails, false-positive management, traceability, human oversight and sign-off. CREST has already updated its standard for penetration testing to cover AI-enabled testing with updates to all our other disciplines scheduled for release. New standards are required for AI security services themselves, covering areas such as LLM application testing, agent testing, model supply-chain assurance and AI governance. CREST has already released a standard for the security testing of AI-based systems.

The market is therefore changing in two directions at once.

Organisations need help securing AI, while providers are still discovering how best to use AI within their cybersecurity services. Every buyer that adopts AI creates new demand for assurance. Every provider that embeds AI into delivery creates new questions about quality, transparency and accountability.

The future of AI-enabled cybersecurity will not be determined simply by which providers use AI the most, it will be shaped by providers' ability to demonstrate that they use it responsibly, transparently and with the right level of expert oversight. For buyers, the question will be what 'good' looks like.

For providers, the challenge will be proving it. For standards bodies, regulators and organisations such as CREST, the task is to help define the expectations that allow the market to innovate with confidence.

We have already begun that work, releasing our industry-backed **Charter and Principles for responsible AI use** and publicising the names of all its signatories, while updating our standards for accreditation to take account of the use of AI within service delivery.

2. LONGER-TERM IMPLICATIONS

To help us look further ahead to understand the challenges that will confront our industry and community, we used the *UK Government Office for Science (GOS) scenarios for 2030 (updated June 2026)*. This formed the basis for thinking about what the potential implications are for cybersecurity as a whole, for attackers, buyers, and the industry – companies and individual professionals – specifically. The GOS scenarios were themselves constructed through workshops and subsequent rounds of peer review with over 70 experts from academia, industry, and government departments across various workshops and draw on over 200 other reports and research publications.

While the GOS scenario development process initially identified 32 possible combinations of outcomes and scenarios, experts reduced these to five scenarios that were ‘*sufficiently diverse and interesting*.’

These five scenarios have been kept up to date in the light of experience through periodic review, most recently in June 2026. They consider the rate of advancement of AI – whether continuing at the same rate as today, slowing or taking off; as well as the ‘critical uncertainties’ – the levels of capability, distribution and access, security, and adoption of AI as well as the labour displacement it causes and the levels of global cooperation to develop and make it safe.

2.1 The potential scenarios

Under all scenarios, AI capability increases and brings positive impacts, including productivity gains, economic growth, improved public services and scientific breakthroughs. These benefits are accompanied by a potential increase in the

impact and scale of harm – including an emergence of novel forms of cyber attack and a potential decrease in human ability to evaluate performance and safety.

In the worst case scenarios from a security perspective, the AI acts unexpectedly, evades controls, and cannot or will not explain its reasoning. Malfunctions and malicious use are severe and widespread. Recent well publicised security incidents involving frontier AI models illustrates that this risk has already started to manifest.

In the best case scenarios, the AI acts as intended and is under control with its reasoning explained and understood. Malfunctions and malicious use are less frequent and severe. The benign security scenarios are enabled by global alignment on standards that ensure technical safety and mitigate severe malicious use although confirming compliance with those standards at national level remains challenging.

2.2 The impact of the scenarios on cybersecurity

CREST’s strategy discussions involving the Chairs of our Regional Councils and Technical Focus Groups, our AI Working Group, Technical Committee and International Council, have considered those elements of each scenario that are specifically relevant to our industry while confining our views to our area of expertise – what malicious actors might do in these scenarios and what they could mean for the cybersecurity industry, its buyers and standard setters.

In this section we explore what these hypothetical scenarios could mean for cybersecurity...

2. LONGER-TERM IMPLICATIONS

In scenarios where the AI market is more diverse and distributed, the gap between open and closed-weight capabilities narrows. Government agencies and industry from multiple regions develop frontier-grade capabilities. AI can automate many tasks and workflows as capably as a human. While the benefits include economic opportunity and less dependency, competition means offensive capabilities tend to outpace defensive measures.

Under scenarios where frontier-grade capabilities are more diverse and with open-weight systems proliferating, the greater scrutiny and evaluation of capability is possible. However, such systems are harder to monitor and can be fine-tuned to remove safeguards or for malicious use.

In such scenarios, we assess that national cybersecurity authorities and sector regulators will be relatively AI-sceptical and are likely to focus on trying to limit the use of AI within security systems to mitigate safety concerns. They are likely to continue to advise industries working in regulated sectors to use contracts that ensure liability for any issues resulting from use of AI rests with the supplier.

Where public trust in, and the security of, AI tools is low, 'shadow' use of AI – where elements within the workforce adopt tools that will make their life easier and themselves more effective ahead of or contrary to corporate policies or procurement – will be high as individual users take the risk of onboarding AI tools in order to realise the benefits. This will create further vulnerabilities in target networks and infrastructure, and opportunities for malicious actors.

These actors will make very good use of the capabilities – including highly capable open source and often sector-specific computer network exploitation (CNE) tools – on offer.

With no inhibitions about using the tools on the attack side – safety concerns are not an issue – and more caution on the part of their targets, the gap between attack capability and defensive capability grows exponentially. Vulnerabilities are continually found and exploited and while patches are released at scale, defenders can't absorb and apply the patches fast enough.

In this world, some cyber security companies will themselves use the open-source CNE tools, giving them a strong capability that they can offer to their clients at low cost given the time savings. There is a strong market for the identification and fixing of the proliferating vulnerabilities in systems and services as well as those created by the shadow-users. Security architecture transformation will become a key service line for the cybersecurity industry – enabling clients to keep ahead of the avalanche of newly found gaps, exploits and patches. The more agile, flexible, and cost-effective companies will flourish.

In alternative scenarios, the AI is highly capable, matching or outpacing humans at the majority of, or even virtually all cognitive tasks. The leading systems are developed by a small number of companies with proprietary datasets and dominant market share. Access to them is likely to be strictly controlled by the national government. From a security perspective, the leading edge systems generally act as intended, are less vulnerable to manipulation or 'jail-breaking' and tend to have solid governance around deployment although there remains a risk of malicious access to the underlying technology.

2. LONGER-TERM IMPLICATIONS

In the scenario at the highest rate of advancement, the leading systems may have concealed goals to gain influence and – especially given the application of AI to robotic physical systems – have control over critical systems that they can use to cause severe harm. Behind the leading edge highly capable open-weight systems are available to threat actors.

In this ‘high capability and high institutional trust’ world there are deep relationships between the high-tech companies building the most advanced tooling and government agencies within each geo-political power-block. Test results, knowledge of vulnerabilities and whether they have been fixed or not, exploits and mitigations are shared within trusted groups but otherwise concealed behind NDAs and security classifications. A ‘security divide’ develops between those who can afford, and are permitted to use, the most advanced secure tools and understand their vulnerabilities, and those who are not.

With a prevailing narrative that tools are highly secure, buyers and the public will be sceptical about the need for cybersecurity. The cybersecurity industry’s warnings about unpublished vulnerabilities and state and criminal activities exploiting them will tend to go unheeded as self-interested alarmism.

Many highly skilled cyber security professionals will be made redundant and may find lucrative jobs in criminal enterprises which prize and reward their understanding of vulnerabilities in legacy systems – still widely used underneath the AI-veneer as the reality of transformation

and change programmes fail to keep pace with the leading edge of technology – and the known but unpublicised weaknesses in the AI systems.

Some successful cybersecurity companies develop deep knowledge of the underground markets. They develop new intelligence services understanding and reporting on the movement of former professionals, the knowledge and skills that are being hired and by which criminal or state employers, as well as what vulnerabilities and exploits are being sold, who is buying and where are they going to use them.

Others develop highly tailored and capable AI enabled defensive responses – including in operational technology-heavy environments – that enables them to effectively serve those clients that are still conscious of the need for security.

3. CONCLUSION

As AI becomes more deeply embedded across providers, buyers and threat actors its influence will extend across a growing range of cybersecurity activities. It will change how services are delivered, how human expertise is applied, and how organisations assess the quality of the services they buy, bringing further complexity to a technical marketplace that's already challenging for buyers to navigate. They will require education and support to assess varying kinds of AI-enabled technologies and capabilities and understand what to look for in a trusted provider.

This makes it increasingly important to have clear, consistent and independent ways to signal quality and trust. Objective-based frameworks can provide a common foundation of principles while allowing practices to evolve. Unlike prescriptive standards, this approach can establish shared expectations around transparency, accountability, validation, human oversight and professional competence. This will help buyers make informed decisions, give responsible providers a way to demonstrate quality, and support greater consistency in a market where capability is likely to remain uneven.

We cannot know what the future will look like for cybersecurity in the world of AI. Our analysis of the various scenarios that experts assess as likely shows different services and skills may be needed in the alternate versions of the future. But there is no scenario in which cybersecurity is not needed. It is already clear, now that the initial hype stage is starting to wane, that there is widespread support for an approach that balances innovation with safety and security.

And it is also clear that globally recognised, adopted and enforced standards will play a beneficial role in creating the future world. CREST is committed to being a pro-active and collaborative partner as we navigate this era together.

We have already produced our industry-backed principles that offer the cybersecurity industry a way to prepare for that uncertainty; we are creating a stable basis for trust while technologies rapidly evolve, updating and rolling out our standards for accreditation to help inform and guide the cybersecurity industry as we go into the future together.



Seven Stars House,
1 Wheler Road,
Coventry,
West Midlands,
CV3 4LB

APPENDIX 1: SERVICES MOST EXPOSED TO DISRUPTION

In the near term, the most exposed services are those where there is convergence of high labour time or cost, the accuracy with which AI can perform the task and the relative lower risk of the activity. The table below gives examples and potential outcomes that CREST has identified through its recent surveys and focus and working group discussions.

SERVICE / ACTIVITY	POSSIBLE AI IMPACT		POSSIBLE OUTCOME
Basic vulnerability scanning and enumeration	HIGH	High automation potential; already one of the most common AI uses in penetration-testing workflows ⁴ .	Price compression; absorbed into continuous testing and exposure-management services.
Report drafting and evidence collation	HIGH	High automation potential; currently the most common AI use case (e.g. in penetration testing ⁵).	Faster delivery, lower admin time, but still need for QA and human sign-off.
First-pass code review	HIGH	High automation potential, especially for common vulnerability classes.	Manual review shifts to high-risk code, architecture, exploitability and remediation advice (in addition to provision of proprietary code review agents).
Tier-1 SOC triage	HIGH	High automation potential because AI can summarise alerts, correlate events and recommend actions. Also potential for defensive AI to improve detection and triage.	MDR/SOC models shift from analyst-hours to automation-plus-expert escalation.
Generic threat intelligence reporting	HIGH	High disruption for generic CTI; medium-to-high opportunity for tailored CTI. This threatens low-value “news digest” or generic feed-style services.	The market moves away from generic intelligence reporting toward sector or client-specific, operationalised intelligence: “what matters to this client, in this sector, with this exposure, and what should they do now?” High-value CTI becomes fused with exposure management, detection engineering, threat hunting, red teaming and executive risk advice.
Security architecture / secure design review	HIGH	High disruption, high demand growth. AI can automate first-pass threat modelling, architecture review checklists, control mapping, design-pattern recommendations and policy-as-code checks. Other elements remain context-heavy: business risk, trade-offs, identity/data architecture, agent permissions, RAG design, model/provider selection, logging, monitoring and governance.	The most significant shift is to bring AI into the design phase, ensuring more secure systems pre-release. This will open up both service and product markets that could fundamentally change the vulnerability landscape. Premium value shifts to AI-era secure design authority: designing safe AI systems, validating AI architecture, reviewing agentic workflows, setting guardrails, and assuring that AI has been securely embedded into enterprise systems.

⁴⁻⁵ CREST Research Report: AI in Penetration Testing 2026: crest-approved.org/ai-in-penetration-testing/

APPENDIX 1: SERVICES MOST EXPOSED TO DISRUPTION

SERVICE / ACTIVITY		POSSIBLE AI IMPACT	POSSIBLE OUTCOME
GRC / Assurance	MED/HIGH	AI can automate elements of the assurance life-cycle, for example evidence collection, control mapping, questionnaire responses, policy drafting, framework cross-walks, audit-pack preparation, issue tracking, board-report drafting and continuous control monitoring.	Low-end compliance support may become more commoditised. The more valuable service could therefore shift from producing evidence for periodic framework-based assessments to continuous cyber assurance including enhanced cyber-risk quantification, control design, independent validation and even remediation.
Standard app pen tests	MED/HIGH	Medium to high disruption for routine periodic tests; lower disruption where business logic, chained exploitation or unusual architectures matter. Shift towards continuous or automated testing as part of the toolset.	Commodity tests become cheaper and more continuous; high-end testing becomes more specialised. Agents become the predominant resource within the service, with humans bringing experience and oversight.
Incident response	MEDIUM	Lower impact but meaningful change to keep pace with AI-enabled adversaries. AI can accelerate investigation, but crisis leadership, forensics, containment decisions, legal/regulatory coordination and business judgement remain human-heavy.	IR providers may need AI-enabled investigation plus AI-specific incident expertise.
Red teaming	MEDIUM/LOW	Traditional red teaming likely to have to incorporate AI to effectively simulate adversaries who are using AI. But this will be adopted cautiously and with oversight. Red teaming with AI, particularly in non-regulated sectors, is though a potential growth area.	New premium service and product category, with potential to become a core service area in its own right.
ICS/OT security	MEDIUM/LOW	Whilst the impact could be reduced due to typical airgaps, demand may grow in other areas where AI can be safely deployed. E.g. OT asset discovery, network baselining, anomaly detection, alert triage, vulnerability prioritisation, documentation review, segmentation analysis, attack-path modelling and incident-response support. Autonomous (or even just AI-enabled) testing would be much harder and riskier than in IT due to safety, availability, reliability, latency, legacy-system and physical-process constraints. A drive to fuse AI with robotics could bridge airgaps that have hitherto protected OT systems and need dedicated capabilities.	Traditional ICS/OT assessment is unlikely to be commoditised in the same way as routine IT-based assessment work. The market may shift toward AI-enabled but safety-constrained OT assurance work based on safe AI integration and the associated testing. There is likely to be a need for specialised understanding of AI-specific OT risks, governance and assurance frameworks, oversight mechanisms and safe operation.

APPENDIX 2:

ECOSYSTEM IMPACTS OF THE MOST PROMINENT DRIVERS WE SEE TODAY

DRIVERS	STAKEHOLDERS		
	BUYERS	PROVIDERS	STANDARD SETTERS
	<i>Demand-side change: what buyers of cyber services will need</i>	<i>Supply-side adaptation: what providers of cyber services will need to offer</i>	<i>Implications for government authorities, regulators, CREST, other standards bodies</i>
	AI-generated software increases verification demand Buyers are adopting AI coding tools, but they need stronger assurance that generated code and related configurations are secure.	We may see providers move increasingly to AI-enabled secure code review, combined with human validation. This includes a shift from manual line-by-line review to fast analysis, remediation and assurance.	Application security and penetration testing standards will need to adapt to cover AI-assisted delivery, including tool validation, false-positive management, human sign-off, evidence quality and limits of automated findings.
	Rapid enterprise AI adoption creates new security needs Clients need help securing AI in a number of areas including GenAI, LLM apps, agents, model APIs, training data, third-party AI tools and the inevitable “shadow AI”.	We may see providers increasingly adding a “secure AI adoption” service line including: AI asset discovery AI threat modelling LLM/RAG (Retrieval-Augmented Generation) /agent testing prompt-injection testing data-leakage testing guardrail assessment model supply-chain review; and AI governance advisory.	New standards will need to be created for AI security services including related accreditations, mapped to key international guidance to drive harmonisation and alignment.
	AI raises attack volume and sophistication Buyers will need more help with protection from any explosion of AI-discovered vulnerabilities and AI-enhanced attacks including social engineering. Buyers will need help adapting their vulnerability patching and remediation processes to meet much shortened windows between disclosure and real-life exploitation. This is combined with an increase in the number of patches required. Buyers may also be looking for more help with areas impacted by AI such as phishing resilience, identity security, attack-surface management, detection engineering and incident readiness.	Cyber service providers may move to incorporate AI within their services to help clients respond to AI-based threats at the speed now required. Red-team, managed detection and response, Secure Operations Centres (SOCs) and incident-response providers are likely to increasingly use AI to accelerate triage, threat intelligence analysis, detection-rule generation, investigation summaries and playbook execution. In the short-term at least, human expertise is likely to remain important for exploit development and malware work. It will also be key in helping clients prioritise and effectively implement remediation.	Incident response and SOC standards could be enhanced to cover AI-assisted incident triage, investigation and response. This could include areas such as audit trails, chain of custody, reproducibility, model/data handling, hallucination controls and human decision accountability. The standards could also emphasise the importance of speed of operations in the SOC / MDR services in order to effectively mitigate AI-enabled vulnerability identification and exploitation.
	Buyer challenging providers over their use of AI On the one-hand some buyers will expect the provider to drive down the cost of services due to efficiencies from AI and may be less interested in additional value the provider may offer as a substitute. On the other hand, some buyers may be wary of provider use of AI and the impacts on its data or service quality.	Providers will need to develop a commercial model to fit a range of buyer sentiment on the use of AI within cyber services. Providers may offer both AI-enabled services (with varying amounts of human in the loop) and those that explicitly exclude AI. Providers will clear AI-use policies and overt client acceptance.	Standards may be required covering the responsible use of AI within cyber services covering transparency, operational practices and human oversight. These standards may be a combination of over-arching requirements with tailored specifications for individual service lines.

APPENDIX 3: LONGER TERM POTENTIAL ECOSYSTEM IMPACTS OF THE FUTURE SCENARIOS

STAKEHOLDERS			
DRIVERS	BUYERS	PROVIDERS	STANDARD SETTERS
	<i>Demand-side change: what buyers of cyber services will need</i>	<i>Supply-side adaptation: what providers of cyber services will need to offer</i>	<i>Implications for government authorities, regulators, CREST, other standards bodies</i>
	AI becomes embedded everywhere Buyers will need continuous assurance for AI embedded in products, workflows, customer service, software engineering, operations, SOC, autonomous agents, robotics, autonomous vehicles and critical systems. They will need to understand where AI is being used, what data it can access, which tools it can call, how it interacts with other systems and who remains accountable when something goes wrong.	Providers will need continuous AI assurance capabilities, including monitoring AI activity, testing prompt and tool misuse, validating guardrails, reviewing agent permissions, assessing RAG data exposure, testing AI-enabled workflows and responding to AI-specific incidents. In sectors where AI connects to physical systems, providers will also need to understand the interaction between cyber risk, operational resilience and safety.	Standards for continuous assurance will need to come to the fore. They will need to cover AI lifecycle assurance, agentic systems, tool use, data access, model updates, human accountability and cyber-physical environments where AI-enabled systems can affect real-world operations.
	Trust in AI diverges across the market In lower-trust environments, buyers may be cautious about adopting AI defensively, particularly in regulated sectors or critical environments. In higher-trust environments, buyers may become overconfident in highly controlled AI systems and underinvest in cybersecurity. In both cases, buyers will need independent evidence of risk, control effectiveness and responsible AI use rather than relying on vendor claims.	Providers will need to serve both cautious and overconfident buyers. Some clients will want limited-AI or non-AI service options, clear liability positions and strong human oversight. Others will need help understanding that highly capable and highly trusted AI systems can still contain hidden vulnerabilities. Providers will need to translate complex AI risk into evidenced, board-level advice.	Standards may need to help establish a common baseline for trustworthy AI-enabled cybersecurity services. This could include transparency over AI use, evidence quality, supplier accountability, liability considerations, human oversight and clear expectations for when AI can and cannot be used within security-critical services.
	Regulatory and geopolitical fragmentation increases complexity Buyers operating across markets will need help navigating divergent regulatory expectations, different national approaches to AI, data residency, supplier location, model provenance and technology supply-chain exposure. In regulated sectors, buyers may also be pushed towards contracts that place liability for AI-related failures with suppliers.	Providers will need to support clients across fragmented regulatory and geopolitical environments. This may include AI supplier assurance, model and data provenance reviews, jurisdictional risk assessments, contract assurance, and advice on how to adopt AI when global consensus on safe use is limited.	Standards will need to balance international harmonisation with local regulatory realities. Principles-based frameworks may become more important, supported by more frequently updated guidance that can respond to changing geopolitical, regulatory and technical conditions.
	Attackers exploit AI capabilities without defensive constraints Buyers will need to assume that malicious actors will use AI-enabled tools without the safety, legal or governance constraints faced by defenders. They will need support against faster reconnaissance, vulnerability discovery, exploit development, phishing, identity theft, social engineering, espionage and attacks against public sector and critical services.	Providers will need to emulate AI-enabled attacker behaviours safely and responsibly. Red-team, threat intelligence, detection engineering, attack-surface management and incident response services will need to reflect the speed and scale at which AI-enabled adversaries can operate.	Standards will need to address AI-enabled offensive testing, safe emulation of malicious AI use, governance of dual-use tooling, client consent, audit trails, evidence handling and boundaries for responsible testing.

APPENDIX 3: LONGER TERM POTENTIAL ECOSYSTEM IMPACTS OF THE FUTURE SCENARIOS

STAKEHOLDERS			
DRIVERS	BUYERS	PROVIDERS	STANDARD SETTERS
	<i>Demand-side change: what buyers of cyber services will need</i>	<i>Supply-side adaptation: what providers of cyber services will need to offer</i>	<i>Implications for government authorities, regulators, CREST, other standards bodies</i>
	Open-source and state-enabled offensive tooling widens the attack / defence gap Buyers will need protection against a more diverse and harder-to-monitor threat ecosystem, including open-source AI tools, plug-ins, autonomous agents, commercially available systems and state-backed capabilities. They will need scalable support to identify and fix proliferating vulnerabilities before they are exploited.	Some providers may use open-source computer network exploitation tools defensively to improve speed, coverage and cost-effectiveness. The strongest providers will be those that can use powerful tooling responsibly, validate outputs, manage legal and ethical risks, and translate capability into practical remediation for clients.	Standards will need to set clearer expectations for the use of dual-use and open-source offensive tooling in defensive services. This could include requirements for provenance, tool governance, client acceptance, auditability, safeguards and human oversight.
	Shadow AI becomes a major unmanaged attack surface Even where public trust in AI is low, individual employees may still adopt AI tools to improve productivity ahead of, or contrary to, corporate policy. Buyers will need help identifying shadow AI use, understanding data exposure, enforcing policy, reducing risky behaviour and securing workflows that have already adopted unsanctioned tools.	Providers will need to offer AI asset discovery, shadow AI assessment, data-leakage testing, policy and control design, user-behaviour analysis, secure configuration guidance and remediation support. This is likely to become a significant source of demand, particularly where unmanaged AI use creates hidden vulnerabilities in target networks.	Standards will need to cover discovery and assurance of unsanctioned AI use, including evidence expectations, data handling, access control, user accountability and governance requirements for AI tools adopted outside formal procurement.
	Vulnerability and patch velocity exceeds defender capacity Buyers will need help coping with an avalanche of newly discovered vulnerabilities, exploits and patches. They will need prioritisation, exploitability assessment, compensating controls, temporary mitigations and support applying patches fast enough to keep pace with attacker capability.	Providers will need to move further into remediation engineering, continuous exposure management, attack-path reduction, virtual patching, secure configuration, identity hardening and security architecture transformation. The most valuable providers will not simply identify gaps; they will help clients absorb and act on the volume of change.	Standards will need to place greater emphasis on remediation outcomes, prioritisation and validated risk reduction. Assurance should increasingly assess whether services help clients reduce exposure in practice, not just whether they identify weaknesses accurately.
	Security architecture transformation becomes a core service line Buyers will need more than point-in-time assurance. They will need support redesigning architectures so they can withstand continuous vulnerability discovery, AI-enabled attacks, unmanaged AI adoption and faster exploitation cycles. This includes reducing blast radius, simplifying environments, improving segmentation and designing systems that are easier to monitor and remediate.	Providers will need to combine testing, architecture, engineering and managed service capability. Security architecture transformation may become a major service line, particularly for organisations struggling to keep ahead of newly found gaps, exploits and patches. Agile, flexible and cost-effective providers may gain advantage.	Standards will need to recognise architecture transformation and resilience engineering as part of the assurance ecosystem. This may require approaches that assess design quality, implementation support, operational effectiveness and measurable exposure reduction.

APPENDIX 3: LONGER TERM POTENTIAL ECOSYSTEM IMPACTS OF THE FUTURE SCENARIOS

STAKEHOLDERS			
DRIVERS	BUYERS	PROVIDERS	STANDARD SETTERS
	<i>Demand-side change: what buyers of cyber services will need</i>	<i>Supply-side adaptation: what providers of cyber services will need to offer</i>	<i>Implications for government authorities, regulators, CREST, other standards bodies</i>
	Commodity testing becomes cheaper and more automated Buyers will expect faster turnaround and lower unit cost for routine testing services, code review and vulnerability validation. The market will question why it should pay traditional day rates for work that AI-enabled tooling can perform repeatedly and continuously. At the same time, buyers will need to understand where automation is sufficient and where expert judgement remains essential.	Providers will need productised, AI-led testing platforms for lower-risk assets, with humans focused on complex exploitation chains, business logic, cloud and identity abuse, CNI and OT, AI systems, cyber-physical environments and remediation strategy. Providers that can combine automation, price competitiveness and expert oversight will be better positioned.	Standards will need to be aligned to varying levels of AI integration within products and services, such as AI-assisted, AI-augmented and autonomous-with-human-oversight. They will also need explicit requirements for what must still be human-reviewed.
	Services shift from reports to outcomes Buyers will be less interested in standalone reports and more focused on continuous, and ideally near real-time, risk reduction. They will look for validated exposure reduction, remediation support, attack-path reduction, secure development enablement, operational resilience and board-level assurance. Cyber buyers may also look to reduce the number of products and service providers through vendor consolidation.	Providers must combine consultancy, tooling, managed services and remediation engineering. Smaller specialist firms may need partnerships, platforms or niche differentiation. Larger consultancies and vendors are likely to bundle AI security into broader suites.	Standards will need to adapt to the importance of remediation and client outcomes rather than just identifying weaknesses or issues. Approaches will need to work for a hybrid product-plus-consultancy model rather than the historically separate markets.
	AI creates new cyber-physical risk Buyers in transport, manufacturing, energy, healthcare, logistics and other operational environments will need assurance over AI connected to robotics, autonomous vehicles, physical infrastructure and critical systems. The risk will not be limited to data loss or service disruption; failures or compromise could create physical, safety and operational consequences.	Providers will need AI-enabled OT and cyber-physical assurance capabilities, including safety-aware testing, control-system knowledge, robotics and autonomy testing, incident response for operational environments and advice that brings together cyber, safety and resilience disciplines.	Standards will need to develop expectations for AI security services in cyber-physical environments. This may include fail-safe design, human override, safety controls, operational resilience, evidence handling and coordination with sector safety regulators.

APPENDIX 3: LONGER TERM POTENTIAL ECOSYSTEM IMPACTS OF THE FUTURE SCENARIOS

STAKEHOLDERS				
		BUYERS	PROVIDERS	STANDARD SETTERS
		<i>Demand-side change: what buyers of cyber services will need</i>	<i>Supply-side adaptation: what providers of cyber services will need to offer</i>	<i>Implications for government authorities, regulators, CREST, other standards bodies</i>
DRIVERS	Capability becomes increasingly uneven	Buyers may face a widening security divide. Some organisations will have access to the most advanced secure tools, trusted vulnerability information and specialist support, while others will be priced out or excluded from trusted groups. Those outside the most capable ecosystems may need help understanding residual risk, compensating controls and exposure in legacy systems sitting beneath an AI-enabled layer.	Providers may increasingly serve different market tiers. Some will work with highly capable clients inside trusted ecosystems. Others will support organisations that cannot access the most advanced tools or intelligence. Providers serving the wider market will need to translate limited information into practical controls and resilience measures.	Standards will need to help prevent assurance becoming available only to the most privileged parts of the market. Baseline expectations, independent assurance mechanisms and accessible guidance will be important to reduce the risk of a widening security divide.
	Unpublished vulnerabilities and concealed risk become more important	In high-trust environments, buyers and the public may be sceptical about the need for additional cybersecurity services because AI systems are perceived as highly secure. Buyers will need help understanding that controlled interfaces and vendor assurances may still conceal unpatched vulnerabilities, legacy weaknesses and risks known only to trusted groups or adversaries.	Some providers will need deep knowledge of underground markets, exploit trading, criminal hiring patterns and state-linked activity. Others will develop highly tailored AI-enabled defensive responses, especially for clients that remain alert to the risks of unpublished vulnerabilities and hidden weaknesses.	Standards may need to address how vulnerability information is shared, evidenced and protected when it is commercially sensitive, classified or subject to NDA. They may also need to define expectations for assurance where not all vulnerability information can be made public.
	Human expertise becomes more valuable, but may become harder to retain	Buyers will need confidence that providers still have the human expertise required to validate AI outputs, understand complex exploitation chains, assess hidden vulnerabilities and make difficult prioritisation decisions. They may also need assurance that providers can retain specialist skills as the labour market changes.	Providers will need to invest in retaining and developing skilled practitioners. In some futures, highly skilled professionals may move out of traditional cybersecurity employment, including into better-paid criminal or state-linked activity. Providers will need strong cultures, training pathways and operating models that make expert human judgement central to AI-enabled delivery.	Standards, qualifications and professional development frameworks will need to evolve. Greater emphasis may be needed on validation, interpretation, governance, AI system testing, OT and cyber-physical assurance, threat intelligence, ethics and accountability.
	Accreditation must assess people, process and platforms	Buyers will need independent assurance that an AI-enabled service is trustworthy, not just that a tool is effective or named consultants are skilled. They will need confidence in the organisation, the technology stack, the governance model and the way human expertise is applied.	Providers' capability will depend on evidence of a strong organisation, professional processes, cutting-edge tooling and excellent human oversight. This will include how AI tools are selected, validated, monitored, governed and integrated into delivery.	Standards that traditionally focused on services will need to be enhanced to incorporate both services and products, whether integrated or standalone. Accreditation may need to assess people, processes, platforms, governance, tooling, evidence quality and accountability together.