CROWDSTRIKE INPUT ON THE UK CALL FOR EVIDENCE ON SECURE AI

28 February 2026

I. INTRODUCTION

In response to the introduction of the UK Government's Secure AI Infrastructure call for information on 29 January 2026, CrowdStrike offers the following views and recommendations.

We approach these questions from the standpoint of a leading international, AI-native, cloud-native cybersecurity provider that defends globally distributed enterprises from globally distributed threats. CrowdStrike offers insights informed by multiple practice areas: cyber threat intelligence; proactive hunting, incident response and managed security services; and an AI-powered software-as-a-service cybersecurity platform and marketplace. Accordingly, this perspective is informed by CrowdStrike's role in protecting organizations from data breaches and a variety of other cyber threats.

II. COMMENTS

CrowdStrike appreciates the UK government's efforts to ensure the secure development and deployment of advanced AI systems, and the opportunity to provide comments on this important initiative. As an organization that has deployed AI at scale across tens of millions of endpoints for prevention for more than a decade, we have accumulated significant expertise in securing AI systems and infrastructure, while also observing how threat actors target these environments.

III. THREAT LANDSCAPE, CHALLENGES, AND APPROACHES

As Crowdstrike has identified in the 2026 *CrowdStrike Global Threat Report*, the current threat landscape demonstrates that advanced AI infrastructure will operate in a machine-speed adversary environment.[1] Attacks by AI-enabled actors increased by 89% year-on-year, while average breakout time fell to 29 minutes, with the fastest observed at 27 seconds. In parallel, cloud-conscious intrusions rose 37%, with a 266% increase among state-nexus actors. These trends indicate that AI models, training data, agents, and supply chains should be treated as high-value assets that require robust

[1] CrowdStrike. (2026, February 24). *CrowdStrike 2026 global threat report: The definitive threat intelligence report for the AI era* [Report]. https://www.crowdstrike.com/en-us/global-threat-report/

protection. Security architectures must therefore assume rapid lateral movement, legitimate credential abuse, and cross-domain evasion as baseline conditions rather than edge cases.

Identity and trust relationships now sit at the centre of intrusion tradecraft. Valid account abuse accounted for 35% of cloud incidents, and adversaries increasingly subvert hybrid identity, SaaS tokens, and federated trust to obtain persistent access. At the same time, state-nexus actors exploited edge devices in 40% of cases involving vulnerability exploitation. AI systems compound these risks: model development depends on complex software supply chains and open-source dependencies, while agentic and prompt-driven systems introduce novel manipulation vectors. Protecting model weights alone will be insufficient without securing identity control planes, developer ecosystems, edge infrastructure, and runtime AI interaction layers.

Secure AI infrastructure must therefore be designed as an integrated, cross-domain security architecture that preserves confidentiality and integrity of model weights and sensitive data without materially degrading performance. This requires strong identity governance (including non-human identities), hardened software supply chains, and unified telemetry across endpoints, cloud, SaaS, and orchestration environments. Government-supported research and pilots should explicitly evaluate how high-assurance protections — such as isolation, encryption-in-use, and runtime monitoring — can be deployed at scale without undermining availability, scalability, or economic competitiveness.

The Security Challenge and Current Approaches

The protection of AI infrastructure presents unique cybersecurity challenges. The risks to model weights, sensitive data, and system configurations are substantial and increasing as AI capabilities advance. From our perspective, the most significant threats and attack vectors include:

1. **Manipulation of AI Systems and Model Behaviour:**
   Adversaries may seek to manipulate AI systems directly, including through techniques such as prompt injection, indirect instruction manipulation, or abuse of agentic workflows. Unlike traditional exploits, these techniques can compromise the reliability, integrity, or safety of AI outputs without requiring theft of model weights. Where AI systems are integrated with enterprise tooling or granted delegated authority, such manipulation may result in unauthorised actions, policy bypass, or downstream system compromise. This introduces a

new risk that requires robust security controls.

2. **Exposure of Sensitive Data Processed by AI Systems:**
   AI environments process and store significant volumes of sensitive information, including proprietary datasets, operational data, intellectual property, and potentially personal or regulated data. Compromise of identity credentials, API tokens, or cloud control planes may allow adversaries to access or exfiltrate this data while blending into legitimate workflows. As AI systems scale across distributed cloud and hybrid environments, maintaining confidentiality and enforcing identity controls becomes increasingly important.

3. **Supply Chain and Development Pipeline Compromise:**
   AI development can include open-source dependencies, pretrained components, containerised environments, CI/CD pipelines, and third-party orchestration tools. Compromise of upstream software providers, repositories, or development environments can introduce persistent, low-visibility access into AI training or deployment systems. As AI components are integrated across global supply chains, ensuring the security of these systems becomes essential to maintaining integrity.

III. RECOMMENDATIONS - Capabilities to Strengthen Protection

AI systems depend on a complex tech stack that includes software, hardware, GPUs, cloud workloads, training data, etc. Each layer of this stack, as well as the vendors that produce it, must ensure security from development through deployment and use. Further, security may be a layer of the stack in its own right. As the UK works to develop and deploy the most advanced AI systems, ensuring comprehensive security will be essential to maintaining trust in AI technologies.

AI has transformed cybersecurity, and has the power to further strengthen cybersecurity outcomes. Leveraging best-in-class cybersecurity technologies deploying AI is essential to meeting constantly-evolving threats. In addition to leveraging AI for cybersecurity, we must simultaneously address the security of AI systems themselves and the data leveraged by AI.

The following are high-level principles CrowdStrike uses to frame our approach in protecting AI systems.[2]

- *Data Operations*: Ensuring the integrity of AI models through carefully curated training data. This includes rigorous processes for protecting our corpus against adversarial machine learning attacks.

- *Continuous Improvement*: Constant refinement of models to adapt to new threats. Our adversarial pipeline, for instance, allows us to generate new adversarial samples to train our machine learning models, increasing their effectiveness against evolving threats.

- *Privacy-by-Design*: Developing AI systems with Privacy-by-Design principles in mind. This helps to leverage AI in a manner designed to respect user privacy while delivering robust security.

- *Transparency and Accountability*: Clear documentation of AI systems' capabilities and limitations. This transparency is crucial for building trust with users and complying with emerging AI regulations.

More broadly, AI systems must be defended through the 'last-mile;' that is, enterprises, applications, and users. Security teams increasingly require specialized capabilities to achieve this. An emerging security category called AI Detection and Response (AIDR) enables security teams to: monitor AI behavior and interactions; secure AI agent identities and access; discover shadow AI; prevent sensitive data leakage to AI systems; protect cloud-based AI applications; and implement guardrails on AI system actions.

Additionally, standard cybersecurity practices also apply broadly to AI security. We recommend several approaches to strengthen both cyber and AI security:

- **Zero Trust Architecture:** Zero Trust is an incredibly impactful concept for increasing cybersecurity in AI environments. By removing implicit trust and continuously validating every stage of digital interaction, organizations can significantly reduce the risk of unauthorized access to model weights and training data. Zero Trust principles can be applied across the AI infrastructure

[2] *The Evolving Role of AI in Data Protection*, January 29, 2025. https://www.crowdstrike.com/en-us/blog/the-evolving-role-of-ai-in-data-protection/

and extend beyond traditional networks.

- **Identity Threat Detection and Response (ITDR):** Identity has become a primary attack vector for sophisticated adversaries. In AI environments, this extends beyond human identities to include AI Agents. Human, and non-human, identities can have access to sensitive resources and must be governed effectively. We recommend implementing Identity Threat Detection and Response solutions that continuously monitor user and system activity, detect unusual behavior, and alert security teams to potential compromise.

- **Advanced Telemetry and Observability:** When responding to a security incident or event involving AI systems, every second counts. Organizations should implement advanced telemetry and observability solutions specifically designed for AI environments, providing granular visibility into activities across the AI infrastructure ecosystem - **often a Next-Generation Security Information and Event Management (SIEM) solution.** This should include monitoring of data access patterns, compute resource usage, model behavior, and network communications to detect anomalies that may indicate compromise. High-fidelity telemetry is essential for threat hunting across AI environments and enabling rapid response to emerging threats.

It is also worth noting, the most significant threat to AI infrastructure comes from sophisticated threat actors operating unlawfully across jurisdictions. Effective defense requires cross-border data flows and threat intelligence sharing to identify and counter emerging threats. Data localization requirements can be detrimental to AI security, as they may prevent the use of cutting-edge cybersecurity solutions that rely on global visibility. We encourage policies that support appropriate cross-border data flows for security purposes while maintaining privacy and regulatory compliance.

IV. CONCLUSIONS

The secure development and deployment of advanced AI systems will depend on recognising that AI infrastructure constitutes a high-value, high-complexity attack surface requiring purpose-built security approaches. Protecting application deployments and sensitive data is necessary but insufficient without addressing identity control planes, software supply chains, and the resilience of distributed compute environments. As adversaries operate at increasing speed and sophistication, secure AI infrastructure must be designed with integrated, cross-domain visibility,

strong governance of human and non-human identities, and safeguards that preserve integrity without undermining performance or innovation. A coordinated programme led by DSIT, AISI and NCSC presents a timely opportunity to embed these principles into the UK's AI ecosystem, ensuring that the UK remains both competitive and trusted as a location and partner for AI development.

V. ABOUT CROWDSTRIKE

CrowdStrike (NASDAQ: CRWD), a global cybersecurity leader, has redefined modern security with the world's most advanced cloud-native platform for protecting critical areas of enterprise risk – endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft, and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting, and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity, and immediate time-to-value.

CrowdStrike: We stop breaches.

Learn more: https://www.crowdstrike.com/.

VI. CONTACT

We would welcome the opportunity to discuss these matters in more detail. Public policy inquiries should be made to:

Drew Bagley
VP & Counsel, Privacy and Cyber Policy

Līga Rozentāle, CISM
Director, Public Policy EU/International

Email: policy@crowdstrike.com

###