The Quiet Before the Storm

AI, Cybersecurity, and the Approaching Domain

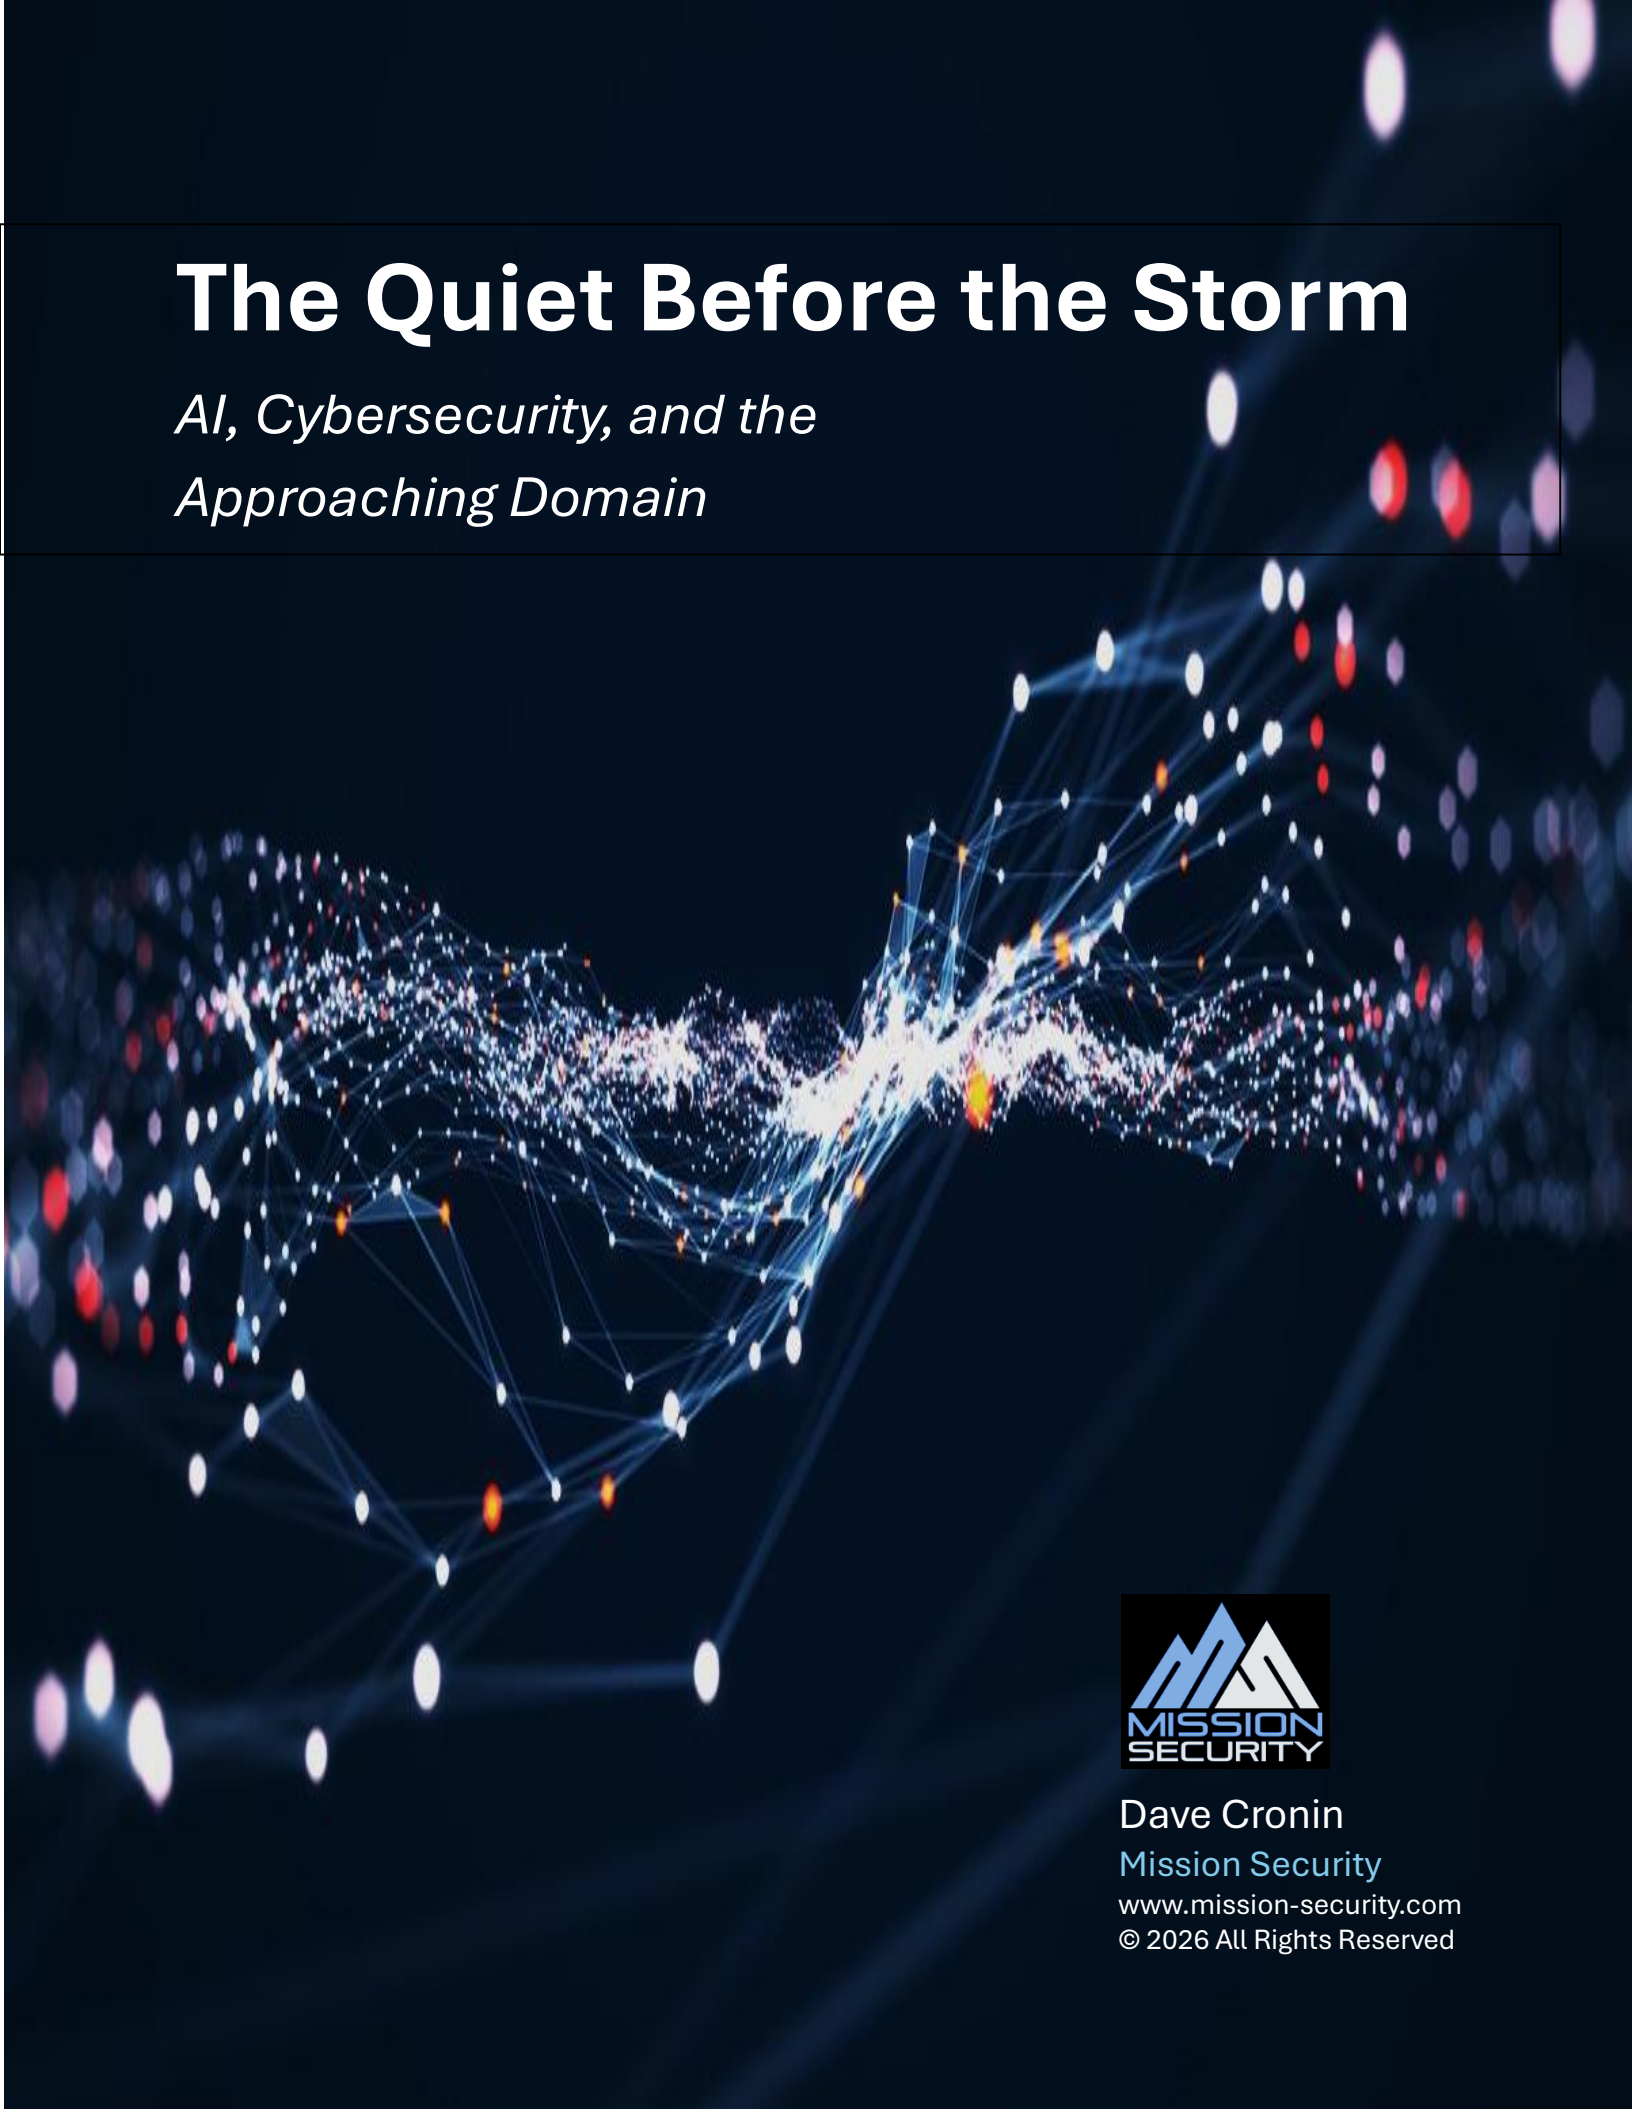

Dave Cronin
Mission Security
www.mission-security.com

The Quiet Before the Storm: *AI, Cybersecurity, and the Approaching Domain*

Artificial intelligence is rapidly reshaping the cybersecurity landscape – although not in the ways most organizations currently perceive. While public discourse emphasizes AI-driven productivity and incremental security improvements, a more consequential shift is underway beneath the surface. We are entering a transitional phase in which AI capabilities, particularly in offensive application, are advancing faster than they are being operationalized at scale.

This has created a temporary "lull"; a period where the most disruptive impacts of AI in cyber have not yet fully materialized. However, this should not be mistaken for reduced risk. On the contrary, it represents a widening gap between emerging adversarial capabilities and enterprise readiness.

This paper considers that organizations are largely unprepared for worst-case AI-driven cyber scenarios and that the industry is approaching a critical pivot point - where cybersecurity transitions from a human-paced discipline to a machine-like speed. In other words, the AI versus AI domain.

The Illusion of Stability

Despite rapid advancements in AI, there has not yet been a proportional surge in high-profile, AI-driven cyber incidents. This has led some organizations to assume that existing controls remain sufficient. This assumption is flawed, and three dynamics explain the current perception:

1. Capability Lag in Operationalization

- Adversaries are still refining how to deploy AI effectively at scale. The tooling exists, but industrialization is still maturing.

2. Blending of Malicious and Legitimate Activity

- AI enables attacks that closely mimic normal user behavior, reducing detectability and lowering signal-to-noise ratios.

3. Overreliance on Legacy Controls

- Security frameworks built around endpoints, networks, and static identity checks are not designed to counter adaptive, AI-driven threats.

The Next Phase of AI-Driven Threats

Over the next 12–36 months, several developments are likely to fundamentally alter the threat landscape. Advancements in AI are set to fundamentally transform the cyber threat landscape across multiple dimensions.

AI will industrialize social engineering by enabling highly personalized, context-aware phishing and impersonation at scale, including real-time voice and video deepfakes, behavioral mimicry of executives and employees, and multi-step deception campaigns executed by autonomous agents - making identity the primary attack surface.

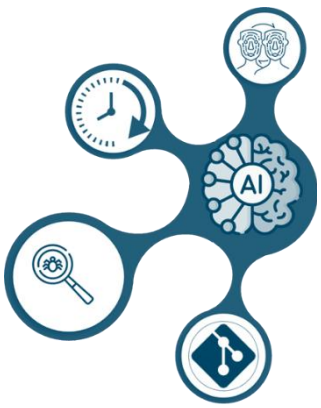

At the same time, AI-driven autonomous attack chains will automate the full lifecycle of cyberattacks, from reconnaissance and vulnerability identification to exploitation, persistence, lateral movement, and data exfiltration, compressing attack timelines from weeks to minutes and leaving defenders with little to no response window.

In parallel, AI will dramatically accelerate vulnerability discovery, including zero-days, rendering traditional patch-and-remediation cycles increasingly ineffective as time-to-exploit approaches real time.

Compounding these risks, adversaries will also target AI systems directly through data poisoning, model manipulation, and attacks on decision-making pipelines, potentially causing compromised systems to generate trusted but incorrect outputs and undermining decision integrity across critical sectors.

Why Most Organizations Are Not Prepared

The cybersecurity industry is structurally misaligned with the nature of AI-driven threats.

- **Human-Speed Operations vs. Machine-Speed Attacks:** Security Operations Centers (SOCs) are built around human triage and response. AI-driven attacks will operate at speeds that exceed human capacity to respond.

- **Fragmented Security Architectures:** Most enterprises rely on a fragmented set of security tools—such as Endpoint Detection and Response (EDR), Security Information and Event Management (SIEM), and Identity and Access Management (IAM)—that operate in silos and lack the integration and automation needed to respond effectively to AI-driven threats, leaving organizations unable to adapt in real time to increasingly dynamic and coordinated attacks.

- **Weak Identity Foundations:** Despite widespread adoption of multi-factor authentication (MFA), most organizations still have weak identity foundations, as many implementations remain phishable, rely on session-based rather than continuous authentication, and lack resilience against advanced impersonation techniques.

- **Reactive Security Models:** Current security strategies remain largely reactive - focused on detecting, investigating, and responding after an event, while AI-driven adversaries will operate in a predictive and adaptive manner, making traditional reactive models increasingly insufficient.

A Realistic Worst-Case Scenario

A realistic worst-case scenario in the age of AI is not a single catastrophic breach, but a coordinated, multi-vector campaign executed at scale.

Such an attack could involve the simultaneous targeting of municipalities, healthcare systems, and financial institutions, the use of AI-generated impersonation to trick executives into authorizing fraudulent transactions, the rapid exploitation of newly discovered vulnerabilities across multiple environments, and adaptive ransomware that changes behavior in response to defensive measures.

The cumulative effect of these tactics would likely overwhelm incident response capabilities, make attribution increasingly difficult, and ultimately erode trust in digital systems and the institutions that depend on them.

The Shift to AI vs. AI Cybersecurity

The defining characteristic of the next era of cybersecurity will be a shift toward AI-versus-AI, where defense becomes autonomous and intelligence-driven. In this model, detection evolves to be continuous and behavior-based, response becomes automated and real-time, and threat hunting shifts from reactive to predictive. As a result, the Security Operations Center (SOC) of the future will be smaller and more specialized, relying heavily on AI-assisted triage and response, while human operators focus on strategy, oversight, and complex decision-making.

Strategic Imperatives for Organizations

Below are strategic imperatives (not just tactics) you should be thinking about at a leadership level.

1 - Treat AI as an Actual New Attack Surface - Not Just a Tool

- If it touches data or decisions, it is part of your security boundary

2 - Implement AI Governance *Now* - Even if Imperfect

- Stand up a lightweight AI governance framework immediately (acceptable use, data classification, approval process)

3 - Secure the Identity Layer Aggressively

- AI makes identity attacks dramatically more effective. If this is weak, nothing else matters

4 - Prepare for Deepfake & Social Engineering Attacks

- Your finance and exec teams are now the primary targets. This is no longer theoretical.

5 - Protect Data from AI Misuse

- Everyone focuses on using AI for defense - fewer focus on protecting data from AI misuse. Extend your DLP and data controls.

6 - Build AI-Specific Threat Detection

- Traditional detection models are lagging - Enhance SOC use cases for AI-assisted phishing patterns & abnormal prompt / API behavior

7 - Secure the AI Supply Chain

- Think beyond your perimeter... AI vendors, model providers, plugins and extensions

8 - Develop an AI Incident Response Playbook

- Most IR plans do not address AI. Include a legal/compliance escalation and a communication strategy

9 - Upskill the Workforce (Fast)

- This is not optional anymore - awareness gaps are being actively exploited. Your people are both your biggest risk and best control

10 - Align AI Security with Business Strategy

- AI is not just a security issue - it's a business risk multiplier (e.g. revenue impact, brand trust, and regulatory exposure)

Many organizations remain exposed to AI-driven risks due to over-reliance on AI tools without proper validation, a lack of visibility into how employees are using AI, and weak identity controls such as outdated MFA methods.

Compounding this, few are prepared for advanced threats like deepfake-driven fraud targeting executives and finance teams, and AI initiatives are often moving forward without adequate involvement from security—creating significant gaps in oversight, governance, and risk management.

Conclusion

The current state of AI in cybersecurity is best understood not as a plateau, but as a staging ground. The absence of widespread, high-impact AI-driven cyber events should not be interpreted as safety - it reflects timing.

As adversaries refine and scale their use of AI, the cybersecurity landscape will shift rapidly and decisively. Organizations that continue to rely on legacy models, built for a slower, more predictable threat environment - will find themselves outpaced.

The defining challenge of the next decade will not be whether organizations adopt AI in their defenses, but whether they do so fast enough to keep pace with those using it against them.

The "lull" is not a reprieve. It is the warning.

Perspective & Context - The views expressed in this paper are based on current industry trends, publicly observable market dynamics, and professional experience. Market conditions, technologies, and provider operating models vary and evolve over time. This content is provided for informational purposes only and should not be interpreted as contractual, financial, or investment advice.

MISSION SECURITY