March 2026

ATARC

Beyond the SOC:
AI-Powered SecOps and Federal Strategies for Modern Cyber Defense

A Federal Practitioner Roundtable White Paper
(Chatham House Rules)

IN PARTNERSHIP WITH

Google Public Sector | accenture

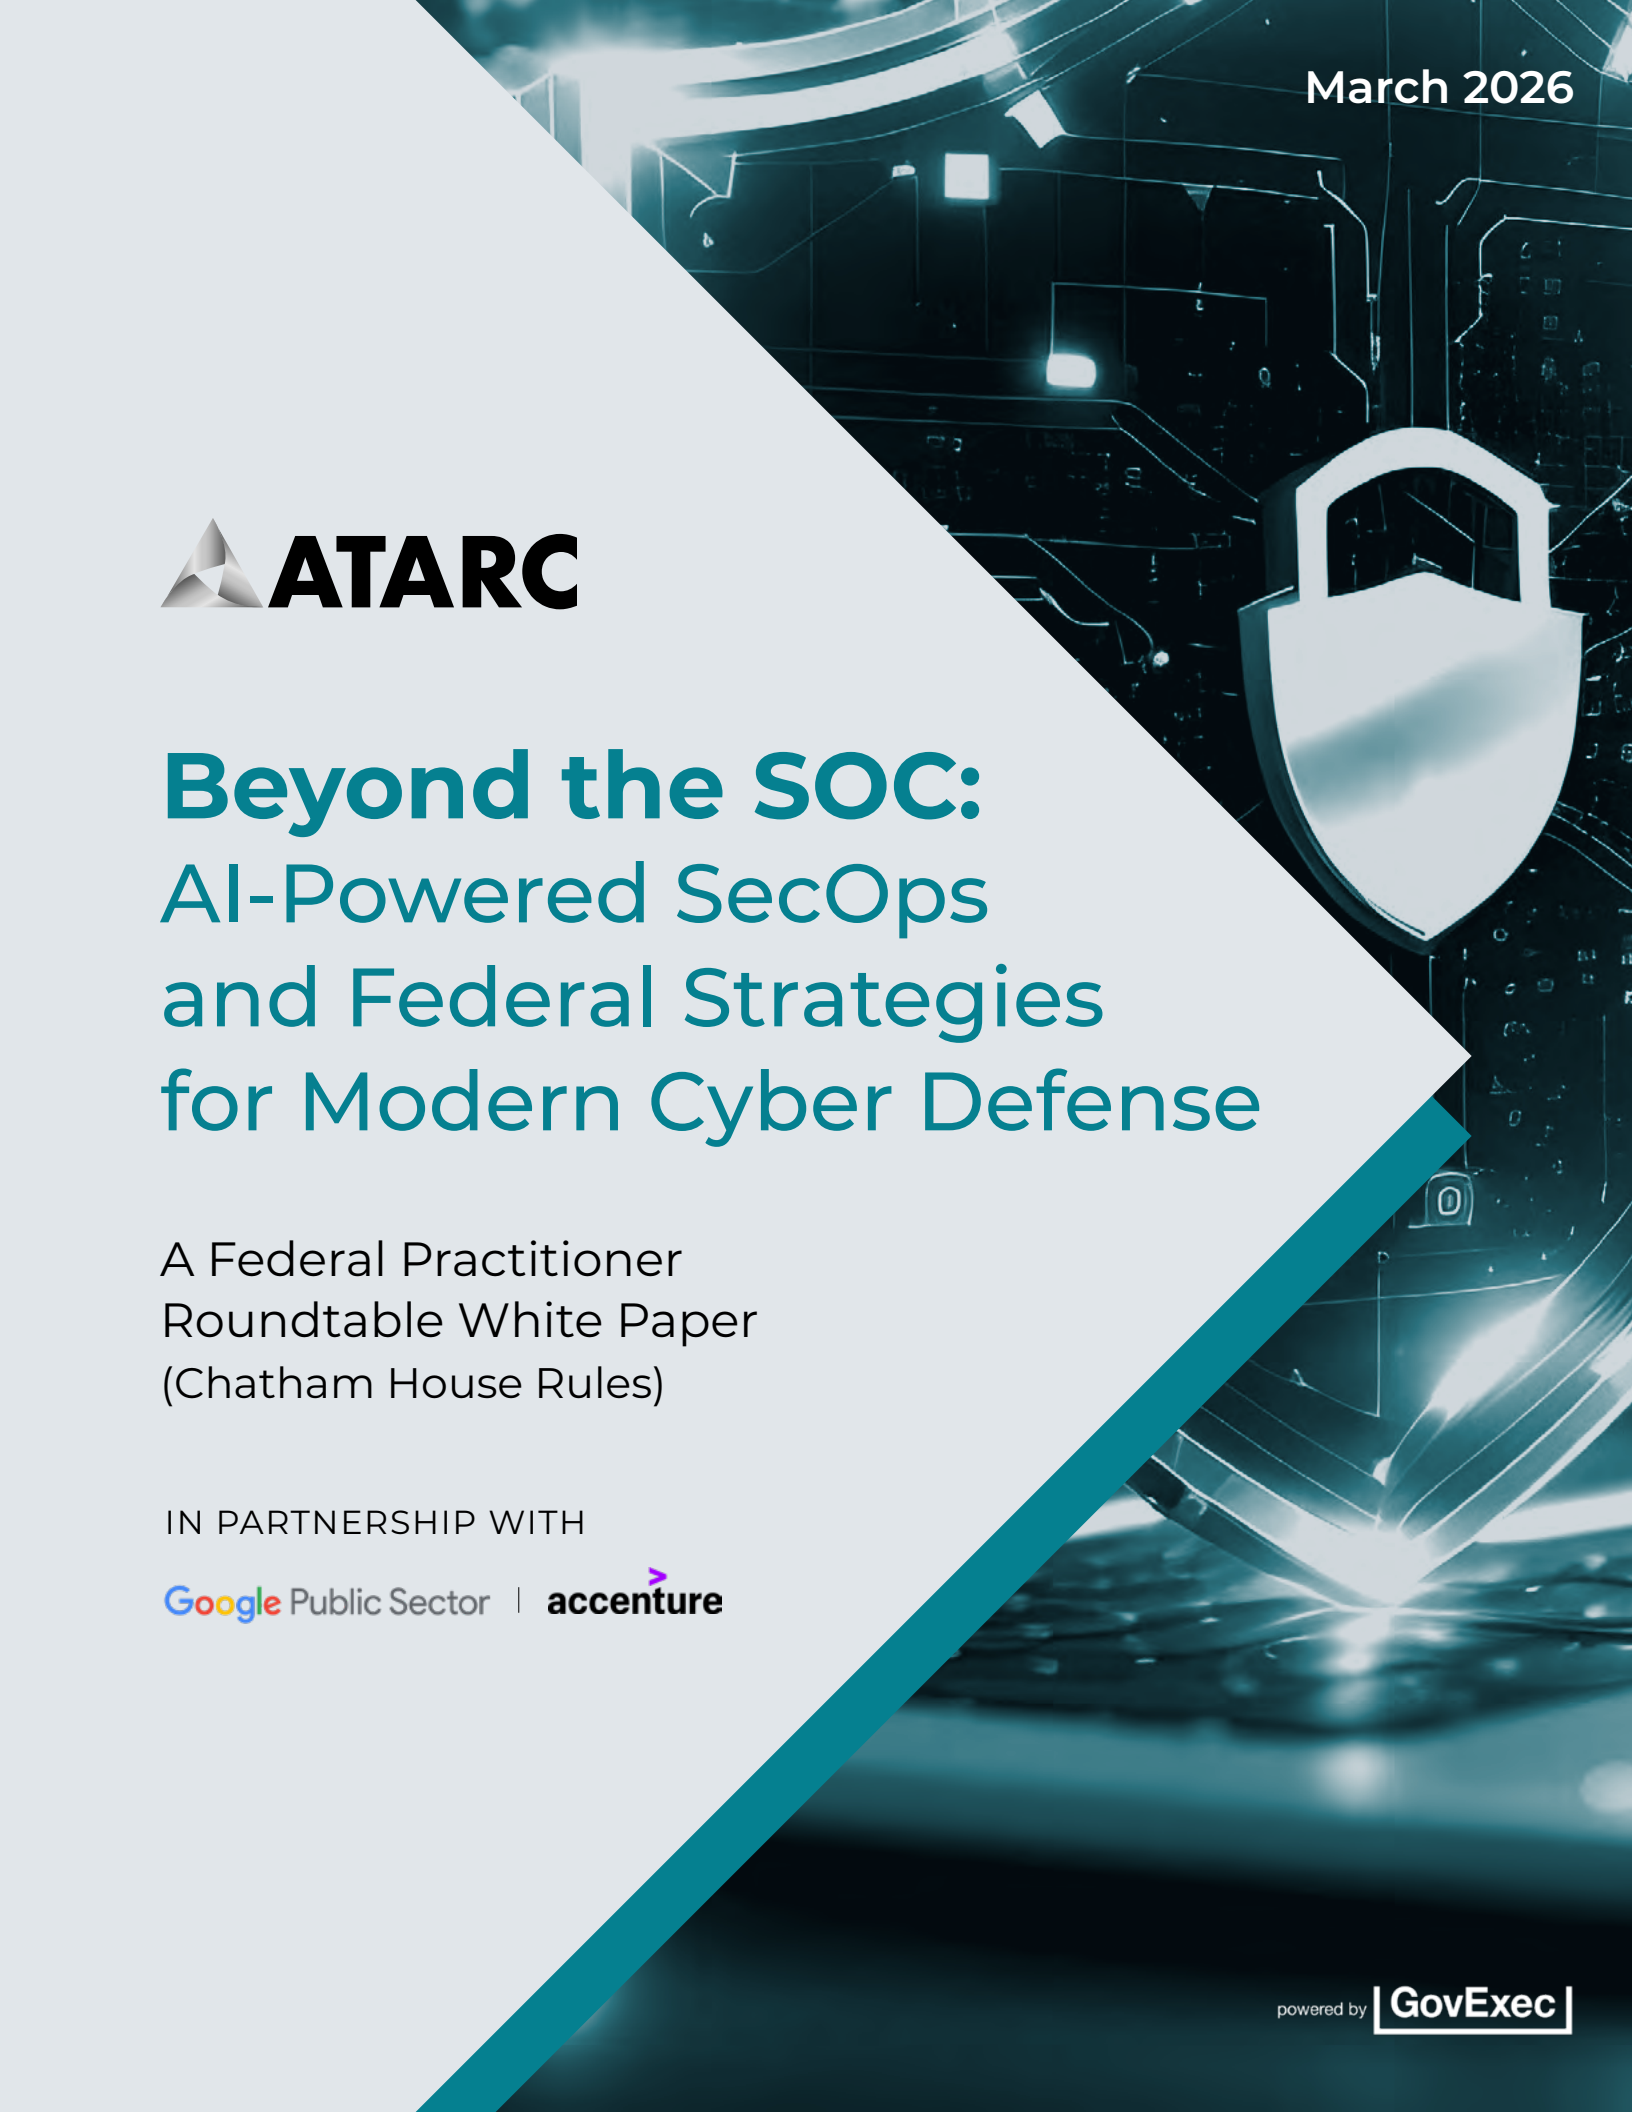

powered by GovExec

Executive Summary

Federal cybersecurity operations are undergoing a fundamental transformation. For more than a decade, the Security Operations Center (SOC) has served as the nerve center for detecting and responding to threats. Today, however, the scale, speed, and sophistication of adversaries—combined with the explosion of data—are pushing traditional models to their limits.

This roundtable revealed a clear shift underway. **Agencies are moving beyond the concept of a centralized SOC toward a more distributed, data-driven, and increasingly autonomous model of cyber defense.** Artificial intelligence is central to this transition, but not in the way it is often described. Rather than replacing analysts, AI is reshaping how work is performed, how decisions are made, and how quickly agencies can respond to threats.

At the same time, practitioners emphasized that the challenges are as significant as the opportunities. Data growth is accelerating faster than budgets. Storage, retention, and access costs are becoming dominant concerns. Acquisition models struggle to keep pace with technological change. And perhaps most importantly, agencies are grappling with how to trust and govern AI systems that are inherently probabilistic.

The future of SecOps in government will not be defined by a single platform or capability. It will be defined by how well agencies integrate data, automation, human expertise, and risk-based decision-making into a cohesive operational model.

The End of the Traditional SOC Model

The traditional SOC was designed for a different era—one in which data volumes were manageable, threats evolved at a slower pace, and human analysts could reasonably be expected to investigate and respond to incidents in near real time.

That model is breaking down.

Across agencies, practitioners described environments in which telemetry is growing exponentially. Every device, application, and interaction generates data. Investigations that once involved a handful of systems now span hundreds of sources. The result is a level of complexity that exceeds human capacity to process.

In this environment, the SOC is no longer a place. It is becoming a function—distributed across systems, platforms, and increasingly, automated processes.

This shift is being driven by necessity. By the time a human analyst identifies a threat, analyzes it, and determines a course of action, the adversary may have already achieved their objective. Data exfiltration, lateral movement, and persistence can occur in minutes or seconds.

The implication is clear: detection and response must happen at machine speed.

The Data Explosion and Its Consequences

At the center of this transformation is data. Every participant in the discussion pointed to data—its volume, cost, and complexity—as the defining challenge in modern SecOps.

Data is no longer just an input to security operations. It is the environment in which those operations take place.

Agencies are required to retain large volumes of data for compliance, investigation, and historical analysis. These requirements are not optional, and they often extend for years. At the same time, new tools and capabilities—particularly those powered by AI—are generating even more data.

This creates a compounding effect. The more advanced the tools, the more data they produce. The more data that exists, the more expensive it becomes to store, process, and analyze.

Practitioners described a growing tension between capability and cost. Cloud platforms offer scalability, resilience, and advanced analytics, but they also introduce consumption-based pricing models that are difficult to predict and control. Storage is only one part of the equation. Data movement, access, and processing all contribute to the total cost.

As a result, agencies are beginning to rethink their architectures. Hybrid approaches are gaining traction, with data distributed across cloud and on-premises environments based on cost, sensitivity, and access requirements. At the same time, there is increasing emphasis on data optimization—deduplication, normalization, and filtering—to ensure that only relevant data is retained and analyzed.

AI as an Operational Necessity, Not a Luxury

Artificial intelligence is often discussed as an emerging capability. In SecOps, it is rapidly becoming a necessity.

The sheer volume of data makes it impossible for human analysts to keep pace. AI is being used to triage alerts, correlate events, and identify patterns that would otherwise go unnoticed. More advanced implementations are beginning to automate aspects of investigation, running iterative analyses across large datasets and surfacing potential attack paths.

This does not eliminate the need for human expertise. On the contrary, it increases its importance.

AI systems excel at processing large volumes of data and identifying correlations. They are far less reliable when it comes to interpretation, context, and judgment. Practitioners highlighted the risk of false conclusions, particularly in areas such as attribution, where subtle distinctions can have significant implications.

As a result, the emerging model is one of collaboration. AI handles scale and speed. Humans provide validation, context, and decision-making.

This partnership is essential, but it also introduces new complexities. Agencies must determine how much autonomy to grant AI systems, how to measure their performance, and how to ensure that their outputs are trustworthy.

The Rise of the AI-Driven SOC

A new model is beginning to take shape: the AI-driven SOC.

In this model, AI is not simply an add-on to existing tools. It is embedded throughout the operational workflow. Data is ingested, processed, and analyzed in near real time. Investigations are initiated and advanced by automated processes. Human analysts focus on higher-order tasks, such as validating findings, refining models, and responding to complex incidents.

Some agencies are already experimenting with this approach. Early implementations have demonstrated the ability to significantly reduce investigation time, enabling analysts to process more cases and identify threats more quickly.

However, these systems are still evolving. Current AI models have limitations, particularly in areas such as context management and long-term reasoning. They can struggle to connect disparate events into a coherent narrative, especially when those events span large datasets and extended timeframes.

To address these challenges, agencies are augmenting AI with additional capabilities. Graph-based data models, for example, are being used to represent relationships between entities, improving the ability to detect complex attack patterns. Memory architectures are being developed to retain and reference historical context. Adversarial validation techniques are being introduced to challenge and refine AI-generated conclusions.

These approaches reflect a broader trend: the recognition that AI is not a single solution, but a component within a larger system.

Governance, Risk, and the Limits of Trust

As AI becomes more integrated into SecOps, questions of governance and trust become more pressing.

Unlike traditional systems, AI does not operate on deterministic rules. Its outputs are influenced by training data, model architecture, and context. This makes it inherently less predictable and more difficult to validate.

For federal agencies, this presents a significant challenge. Security decisions often have high consequences, and there is limited tolerance for error. At the same time, the speed of modern threats requires rapid action.

Balancing these factors requires a new approach to risk.

Practitioners emphasized the importance of aligning security capabilities with mission requirements. Rather than attempting to eliminate all risk—a goal that is neither realistic nor achievable—agencies must define acceptable levels of risk and design their systems accordingly.

This requires engagement at the highest levels of leadership. Decisions about risk tolerance cannot be made solely within IT or security teams. They must be informed by mission priorities, operational requirements, and resource constraints.

In practice, this means shifting the conversation from technology to outcomes. Instead of asking what tools to deploy, agencies must ask what level of protection is required, what level of risk is acceptable, and what trade-offs are necessary to achieve those objectives.

The Acquisition and Budgeting Challenge

Even as technology evolves rapidly, federal acquisition processes remain relatively static.

Practitioners described significant challenges in aligning long-term budgeting cycles with the pace of innovation. Planning for capabilities three years in advance is increasingly difficult in an environment where new technologies emerge on a monthly basis.

This mismatch creates risk. Agencies may invest in solutions that are outdated by the time they are deployed, or they may miss opportunities to adopt more effective approaches.

At the same time, the cost structure of modern technologies adds complexity. Consumption-based pricing, fluctuating licensing costs, and the rapid evolution of capabilities make it difficult to forecast expenses accurately.

Addressing these challenges will require changes in both policy and practice. More flexible funding models, greater emphasis on outcomes rather than specific technologies, and closer collaboration between technical and acquisition teams will all be necessary.

The Human Element

Despite the focus on technology, the discussion repeatedly returned to the importance of people.

AI can augment human capabilities, but it cannot replace the experience, intuition, and judgment that analysts bring to their work. In many cases, the effectiveness of AI systems depends on the expertise of the individuals who design, train, and use them.

At the same time, the workforce is evolving. New tools are lowering the barrier to entry, enabling individuals with limited experience to perform tasks that previously required specialized expertise. This creates both opportunities and risks.

Training and education will be critical. Agencies must ensure that their workforce understands not only how to use AI tools, but also their limitations. They must be able to interpret outputs, identify errors, and make informed decisions.

Equally important is the need to foster a culture of adaptability. As technologies continue to evolve, the ability to learn and adjust will be as important as technical proficiency.

Conclusion

Federal cybersecurity is at a turning point. The traditional SOC model, while still relevant, is no longer sufficient to address the scale and complexity of modern threats. A new model is emerging—one that is more distributed, more automated, and more dependent on data and AI.

This transition will not be simple. It requires changes in technology, architecture, governance, and culture. It also requires a willingness to rethink long-standing assumptions about how security operations should function.

The agencies that succeed will be those that embrace this complexity and take a holistic approach. They will invest in data as a strategic asset, integrate AI into their workflows, and align their capabilities with mission needs.

Most importantly, they will recognize that the goal is not to eliminate risk, but to manage it effectively in an environment where change is constant and uncertainty is inevitable.

The future of SecOps is not beyond the SOC—it is beyond the boundaries that have traditionally defined it.

About This Report

This white paper reflects insights from a federal practitioner roundtable conducted under Chatham House Rules. The perspectives presented capture shared experiences, challenges, and emerging strategies across government, without attribution to specific individuals or organizations.