ARCTIC WOLF

THE ARCTIC WOLF

Essential Guide to the Agentic SOC

A Vendor-Neutral Blueprint for Designing, Building, and Operating AI-Driven Security Operations

TABLE OF CONTENTS

Executive Summary 3

Chapter 1: What Is an Agentic SOC? 4

Chapter 2: Why Traditional SOC Models Fail 5

Chapter 3: The Agentic SOC Operating Model 6

Chapter 4: Core Architectural Requirements 8

Chapter 5: Operational Data and Intelligence Requirements 10

Chapter 6: Trust, Safety, and Reliability Architecture 11

Chapter 7: Expected Outcomes of an Agentic SOC 12

Chapter 8: Why Most AI-Driven SOC Approaches Fall Short 13

Chapter 9: Implementation Realities — Build vs Assemble 15

Chapter 10: Implementation Challenges 16

Chapter 11: Agentic SOC Maturity Model 17

Chapter 12: How to Build an Agentic SOC 18

Chapter 13: How to Evaluate Agentic SOC Solutions 19

Chapter 14: Common Pitfalls in Pursuing the Agentic SOC 21

Chapter 15: The Future of the Agentic SOC 22

Chapter 16: Conclusion 24

Appendix A: Key Concepts and Terminology 24

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Executive Summary

Modern security operations are breaking down not because organizations lack tools, but because the operating model behind those tools no longer matches the scale, speed, and complexity of current threats.

Traditional security operations centers still depend on humans to interpret alerts, drive investigations, and decide what to do next. That dependence has become the primary constraint.

The Agentic SOC addresses that constraint by shifting core execution to AI agents. These systems do more than assist analysts. They investigate, correlate evidence, test hypotheses, and, within defined boundaries, take response actions. Humans remain essential, but their role shifts from primary executor to overseer, validator, and strategic decision-maker.

This guide defines the Agentic SOC, explains why traditional models fail, outlines the architecture and data requirements needed to make the model real, and provides frameworks for implementation, maturity assessment, and solution evaluation.

Readers will leave with a precise definition of what the model is and what distinguishes it from approaches that use similar language, a sequenced framework for building toward it, and an evaluation methodology rigorous enough to separate genuine capability from well-presented marketing.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 1: What Is an Agentic SOC?

Security operations has always centered on a simple question: who does the work? In the traditional SOC, humans do. Technology generates alerts and surfaces evidence, but human analysts still do the investigation. The Agentic SOC changes that by making AI the primary execution layer.

1.1 Defining the Agentic SOC

The Agentic SOC is a security operations model in which AI agents execute the primary workflows of detection, investigation, and response.

This is different from AI-assisted security, where the analyst still owns the work.

In an Agentic SOC, agents initiate investigations, gather and correlate evidence, evaluate likely explanations, and reach conclusions. When policy allows, they can also take response actions. Humans do not disappear; they move into oversight, approval, exception handling, and strategic control.

The defining shift is that execution is no longer limited by human bandwidth. Work starts when data arrives, not when an analyst becomes available.

1.2 Core Characteristics of the Agentic SOC

An Agentic SOC is defined by AI-led execution, continuous operation, parallel investigation, evidence-based reasoning, environment-aware decision-making, and governed autonomy.

These elements work together rather than as separate features.

Instead of waiting in queues, investigations begin immediately. Instead of following one analyst through one path, multiple agents can examine endpoint, identity, network, and historical context at the same time. Decisions are based on evidence and environment-specific context, then constrained by governance that determines what can happen automatically and what requires human approval.

1.3 What the Agentic SOC Is Not

The Agentic SOC is not simple automation because automation follows predefined logic and cannot reason through changing conditions.

It is not SOAR because orchestration coordinates actions but does not independently investigate and decide. It is not a chatbot interface on top of security tools, because conversation alone does not change the operating model.

It is also not unconstrained autonomy. The model depends on bounded autonomy, where agents act within clear limits and humans retain accountability.

Redefining the Operating Model

The Agentic SOC is not a feature added to the traditional SOC. It is a different operating model in which AI becomes the primary execution layer and humans provide direction, control, and oversight. To understand why that shift matters, it is necessary to examine why the traditional model can no longer keep up.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 2: Why Traditional SOC Models Fail

Traditional SOCs struggle not because organizations stopped investing, but because the design of the model itself no longer scales. More tools, more alerts, and even better analytics do not solve a workflow built around human throughput.

2.1 The Evolution of Security Operations

Security operations evolved in stages. SIEM-centric SOCs improved visibility. Detection-centric SOCs improved signal quality.

Automation and orchestration reduced manual work. AI-assisted SOCs improved prioritization and analysis. Each stage added capability, but none changed the core assumption that humans execute investigations. The Agentic SOC is the first model that changes that assumption.

2.2 Structural Limitations of Traditional SOCs

Traditional SOCs are constrained by tiered teams, alert overload, fragmented tools, and linear workflows.

Alerts move through handoffs, queues, and escalations, creating latency and losing context. Analysts must stitch together data across endpoint, network, cloud, and identity systems, often under time pressure. At the same time, modern attacks span multiple domains and span multiple domains and move so fast that step-by-step human investigation can't keep up. These are structural limits, not process defects.

2.3 Why AI Assistants Are Not Enough

AI assistants help, but they do not solve the core problem.

They can prioritize alerts, summarize findings, and suggest next steps, yet the analyst still has to review the case, decide, and act. Assistants only operate when prompted by the analyst, whereas agents run constantly and without prompting. That leaves an assist-to-execution gap. The system may become more efficient at the margin, but it remains bounded by human capacity. AI that supports the old model does not become a new model.

Why Incremental Improvement Is Not Enough

Traditional SOC improvements have all been layered onto the same basic design. Visibility improved. Signal quality improved. Efficiency improved. Analysis improved. But humans still sit at the center of execution.

The Agentic SOC breaks from that pattern by moving execution to AI. That is why it can address the limits incremental improvement cannot.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 3: The Agentic SOC Operating Model

The Agentic SOC operates differently at a system level. Its value is not only speed. It is the redesign of how work is started, distributed, completed, and supervised.

3.1 From Human-Led to AI-Led Operations

In the traditional SOC, humans drive detection, investigation, and response.

In the Agentic SOC, AI systems become the primary operators. They act on signals immediately, gather context, connect evidence, and determine likely outcomes. Humans shift into supervision, validation, and exception handling. This allows scale without surrendering control, provided the system is governed and explainable.

3.2 The Modern Investigation Lifecycle

The Agentic investigation lifecycle is continuous and integrated.

Detection happens in real time. Triage starts immediately. Parallel investigation pulls from multiple data sources at once. The system then evaluates hypotheses, builds a structured conclusion, and either escalates, continues monitoring, or executes a controlled response. The key difference is continuity. The workflow does not pause between stages waiting for human attention.

3.3 Parallel vs Sequential Execution

Traditional investigations are usually sequential: one analyst works one path at a time.

The Agentic SOC decomposes an event into multiple investigative threads that specialized agents can pursue simultaneously. This increases speed, but more importantly it increases depth. Endpoint behavior, identity activity, network signals, and prior history can all be examined at once and combined into one case.

3.4 Continuous vs Queue-Based Processing

Queue-based security assumes delay is acceptable.

The Agentic SOC rejects that assumption. Work starts as soon as a signal appears and continues without waiting in a backlog. This reduces dwell time, improves coverage, and creates a steadier operating rhythm. The system becomes always-on rather than analyst-available.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 3: The Agentic SOC Operating Model, Cont'd...

3.5 Collapse of the Tiered SOC Model

Tiered SOC structures were built to distribute scarce human expertise.

They also create handoffs, latency, and inconsistency. In an Agentic SOC, the primary execution layer is unified and AI-led, so the traditional Tier 1 to Tier 3 structure is no longer the core operating mechanism. Humans still matter, but they sit above the workflow as reviewers, approvers, and strategic decision-makers rather than inside every routine investigation.

The Architecture That Makes It Real

An operating model alone is not enough. The Agentic SOC requires architecture that can support continuous execution, context-aware reasoning, coordinated agents, and governed action.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 4: Core Architectural Requirements

The Agentic SOC depends on a set of interlocking architectural foundations. Weakness in any one of them does not create a partial Agentic SOC. It creates a system that cannot reliably deliver the model.

4.1 Unified Security Data Layer

The first requirement is a unified data layer that ingests, normalizes, enriches, and correlates telemetry from endpoints, networks, cloud systems, identity platforms, and other relevant sources.

Without that layer, agents work on incomplete or inconsistent inputs, which produces fragmented investigations and unreliable decisions.

4.2 Security Knowledge Layer (Context Engine)

The data layer provides visibility.

The knowledge layer provides understanding. It models the relationships among users, devices, identities, applications, behaviors, and history so the system can reason in context rather than on isolated events. This layer must be environment-specific. Normal behavior is organization-dependent, and the system needs that context to decide accurately.

4.3 Multi-Agent Architecture

The Agentic SOC requires multiple specialized agents rather than one generalized model.

Different agents handle detection, enrichment, investigation, response, and validation. They must coordinate, share outputs, and operate under central orchestration so investigations move toward resolution rather than producing disconnected insights.

4.4 Autonomous Investigation Engine

The autonomous investigation engine owns the investigation lifecycle.

It gathers evidence, generates and tests hypotheses, updates its conclusions as new information arrives, and builds a structured case record. This is what turns analysis into end-to-end investigative ownership.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 4: Core Architectural Requirements Cont'd...

4.5 Governance and Bounded Autonomy Layer

Autonomy must be constrained.

Governance defines what actions agents may take, under what conditions, at what confidence levels, and when human approval is required. This layer also needs policy enforcement, monitoring, validation, rollback, and consistent controls across all agents. Without it, autonomy is either unsafe or too restricted to matter.

4.6 Human-in-the-Loop Framework

Human involvement must be designed, not improvised.

The human-in-the-loop framework defines escalation paths, approval checkpoints, override mechanisms, and accountability boundaries

4.7 Continuous Learning System

The system must continuously improve by learning from investigation outcomes, analyst feedback, changing environments.

Performance has to be monitored and tuned over time. Otherwise, the system degrades as threats and environments evolve.

4.8 Integration and Execution Layer

The Agentic SOC must be able to act.

That requires API-driven integration with security and infrastructure systems so agents can isolate devices, disable accounts, change configurations, and execute other governed actions in real time. Observation alone is not enough.

Architecture as Prerequisite

These architectural components are interdependent. Data without context is insufficient. Agents without governance are unsafe. Investigation without execution is incomplete. The Agentic SOC requires the full system.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 5: Operational Data and Intelligence Requirements

Architecture determines whether the model can exist. Data determines whether it works. The accuracy, relevance, and structure of the data feeding the system are what separate theoretical capability from reliable performance.

5.1 Data Quality as the Primary Differentiator

Raw telemetry is not enough.

The Agentic SOC depends on curated, high-signal data that has been cleaned, validated, enriched, and structured for reasoning.

Data quality and data relevance are distinct but equally important. Clean data that does not reflect the threats and behaviors that matter in a given environment is still insufficient.

5.2 Real-World Investigation Data

The system must learn from real investigations, validated outcomes, and actual attack behavior, not only synthetic datasets or generalized threat feeds.

Real-world investigation data gives the system pattern recognition grounded in operational reality and increases confidence that its conclusions reflect how threats actually appear in practice.

5.3 Embedded Operational Expertise

Data alone does not capture expert judgment.

The system also needs embedded operational expertise: investigative methods, prioritization logic, and reasoning patterns drawn from experienced analysts.

This allows agents to do more than process telemetry. It allows them to approach security problems with discipline that resembles real practitioners.

5.4 Continuous Data Reinforcement

Every investigation should improve the system.

New outcomes, behaviors, and observed patterns must feed back into the data and decision layers so the system adapts over time.

This includes both local learning inside one environment and broader intelligence drawn across environments, balanced carefully with environment-specific context.

Data as the Operational Foundation

In the Agentic SOC, data is not a passive record. It is the substrate for reasoning and execution. That makes data engineering, data quality, and real-world intelligence central to system performance.

The Trust Problem

Good data enables decisions, but it does not create trust on its own. Organizations still need confidence that the system can act safely, explain itself, and stay aligned with policy.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 6: Trust, Safety, and Reliability Architecture

Once AI moves into operational execution, trust becomes the gating requirement. Organizations will not allow autonomous investigation and response unless the system behaves safely, predictably, and transparently.

6.1 Why Trust Is the Primary Barrier

The technical path to the Agentic SOC is increasingly feasible.

The harder issue is organizational willingness to let the system act.

That reluctance comes from two concerns: the cost of being wrong and the opacity of many AI systems. If a system cannot explain its conclusions and cannot be trusted to stay within acceptable risk boundaries, it will not be allowed to operate at meaningful levels of autonomy.

6.2 Required Safety Mechanisms

Trust begins with safety controls.

Inputs and outputs must be validated. The system must be tested against adversarial conditions. Runtime monitoring must detect abnormal behavior. Rollback must exist when actions need to be reversed.

These controls do not remove all risk, but they make risk manageable and recoverable.

6.3 Explainability Requirements

A trustworthy Agentic SOC must show its work.

Every material decision should be tied to evidence, supported by traceable reasoning, and captured in full audit logs.

Explainability is necessary for analyst review, governance, compliance, and continuous improvement. Without it, the system remains a black box and trust erodes.

6.4 Governance Lifecycle

Governance is not a launch activity.

It spans pre-deployment validation, controlled rollout, continuous monitoring, periodic model evaluation, and ongoing threshold tuning.

Trust is built over time as the system proves that it can act safely, escalate appropriately, and remain aligned with organizational policy.

Capability Without Confidence Is Not Enough

High capability does not create adoption by itself. Confidence does. Safety, explainability, and governance are what convert technical potential into operational permission.

What a Governed System Actually Produces

Once the system is trusted and governed, organizations can evaluate it on what matters most: the outcomes it produces in operations, analyst experience, and risk reduction.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 7: Expected Outcomes of an Agentic SOC

The value of the Agentic SOC is measured by outcomes, not by whether a platform claims to use AI. The model matters because it changes operational performance in ways earlier approaches could not.

7.1 Operational Outcomes

The first outcomes are operational.

Investigations begin immediately, which reduces investigation time and compresses dwell time. Backlog is no longer the accepted norm because signals are processed continuously rather than deferred.

This also improves coverage. Every signal can be evaluated, not only the subset a human team has time to inspect.

7.2 Efficiency Outcomes

The Agentic SOC removes repetitive tasks from human analysts and reallocates human effort toward oversight, edge cases, and higher-value judgment.

That increases productivity and allows organizations to scale outcomes without scaling headcount at the same rate.

7.3 Quality Outcomes

The model also improves quality.

Investigations follow a more consistent method, use broader evidence, and apply context-aware reasoning across domains.

That reduces variability, lowers the chance of missed threats, and improves the quality of decisions.

7.4 Experience Outcomes

Analysts receive structured conclusions instead of disconnected raw signals.

Context is already assembled, which reduces tool switching, back-and-forth, and cognitive load. The result is a cleaner operating experience built around clarity rather than manual assembly.

7.5 Business Outcomes

From a business perspective, the model can shorten time to value, reduce operational burden, improve return on security investment, and lower cyber risk by identifying and addressing threats earlier and more consistently.

It changes security operations from a labor-scaling problem into a system-scaling one.

Compounding Returns

These outcomes reinforce one another, creating a compounding effect over time. Organizations that build the model well will continue to pull ahead of those that do not.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 8: Why Most AI-Driven SOC Approaches Fall Short

Most AI-driven SOC offerings improve pieces of the workflow but do not change the workflow itself. They can look advanced while still delivering optimization rather than transformation.

8.1 AI as Assistant vs Operator

The most common failure is treating AI as an assistant rather than an operator.

If analysts still own execution, the system remains bounded by human throughput. That distinction matters more than how sophisticated the assistant appears.

8.2 Lack of Multi-Agent Orchestration

A second failure is relying on a single model or loosely connected components rather than a coordinated multi-agent system.

Security operations spans multiple functions and domains, and without specialization plus orchestration, execution becomes fragmented.

8.3 Absence of Context and Knowledge Layer

Many systems analyze telemetry without a true context engine.

Without knowledge of relationships, norms, and history inside the environment, they produce excess noise, miss subtle threats, or both. Context is not an enhancement. It is a requirement for accurate reasoning.

8.4 Weak or Missing Governance

Some systems improve analysis but do not build governance into the architecture.

That forces organizations into one of two poor choices: unsafe autonomy or overly restricted operation that restores human bottlenecks. Effective autonomy depends on meaningful control.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 8: Why Most AI-Driven SOC Approaches Fall Short, Cont'd...

8.5 Poor Data Foundations

AI systems operating on noisy, fragmented, low-relevance data will produce noisy, fragmented, low-reliability outcomes.

Weak data foundations limit model performance no matter how advanced the interface looks. The Agentic SOC requires curated, correlated, real-world operational data.

8.6 Fragmented Architectures

Many organizations assemble AI-driven security from separate tools connected by integration layers.

That often adds latency, inconsistency, and complexity while preventing any one system from owning the full workflow from detection through response. Integration can connect tools. It cannot by itself create a unified operating system.

Separating Claims from Capability

The practical test is not whether a vendor uses AI, but whether the system meets the architectural, data, governance, and execution requirements of the Agentic SOC. Anything less produces partial alignment at best.

The Build-or-Assemble Decision

Once organizations understand the gap between claims and capability, they face a practical choice: build the system internally, assemble it from components, or adopt a more integrated approach.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 9: Implementation Realities — Build vs Assemble

The Agentic SOC is not a normal integration project. It is a system-level build. That makes implementation strategy a question of feasibility and long-term sustainability, not just preference.

9.1 The Hidden Complexity of Building an Agentic SOC

Building internally often looks easier than it is.

The real difficulty lies in engineering a unified data layer, training and validating models on environment-specific and real-world investigation data, coordinating specialized agents, and enforcing governance across the entire system. Because these elements are tightly coupled, weakness in one area undermines the others.

9.2 Why DIY Approaches Fail

DIY efforts often show early promise but struggle in production.

Integration burden grows. Operational overhead rises. Specialized expertise proves hard to sustain. Outcomes become inconsistent, which weakens trust. That creates the common DIY paradox: visible capability without reliable, scalable performance.

9.3 Practical Considerations

Organizations deciding how to implement should weigh time to value, scalability, operational sustainability, internal expertise, and alignment with long-term strategy.

The question is not whether the organization can build something. It is whether it can build, operate, govern, and improve the full system over time.

Making the Decision Deliberately

Three questions provide useful orientation. First, does the organization have sustained engineering capacity — not just to build the initial system, but to maintain, retrain, and evolve it over time? Second, does the solution under consideration provide the core architectural requirements as an integrated system, or as separately sourced components the organization will be responsible for connecting? Third, regardless of path, which capabilities will remain under direct organizational control — the policies governing autonomous action, the confidence thresholds determining when humans are engaged, and the criteria by which outcomes are evaluated? Those must not be delegated.

Where Implementation Actually Breaks Down

Even when the build-versus-assemble choice is made thoughtfully, many efforts still falter in execution. The next challenge is not deciding to pursue the model, but managing the obstacles that derail implementation.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 10: Implementation Challenges

The hardest implementation problems are rarely isolated technical defects. They are points where traditional assumptions collide with the requirements of the new model.

10.1 Data Fragmentation

Security data is often scattered across disconnected systems, each with its own formats, schemas, and assumptions.

That makes it difficult to establish the unified, real-time visibility the Agentic SOC requires. Until fragmentation is addressed, context and reasoning remain limited.

10.2 Integration Complexity

Even when the right systems exist, connecting them into a real-time operational fabric is difficult.

APIs differ, workflows conflict, latency accumulates, and changes in one component ripple across the rest of the system. Integration complexity is a system challenge, not a connector problem.

10.3 Model Reliability

Models may perform well in controlled settings but struggle in live environments.

Threat variation, incomplete data, drift, and ambiguous cases can all weaken reliability.

For the Agentic SOC, reliability matters more than isolated model accuracy because the system must perform consistently under operational conditions.

10.4 Organizational Resistance

Implementation also changes roles, workflows, and accountability, which can trigger resistance from teams accustomed to the traditional model.

Analysts may worry about loss of control or job relevance. Leaders may hesitate to shift accountability toward AI-led execution. Adoption requires organizational change as well as technical change.

10.5 Governance Gaps

Many organizations reach technical capability before they establish the policies, thresholds, approval paths, and controls needed to govern it.

That creates either unsafe freedom or constrained deployment that limits value. Governance has to mature alongside capability.

Challenges as Diagnostic Signals

These challenges should be treated as signals about where the model is not yet aligned. They are useful because they reveal which foundations still need work.

Mapping the Progression

Organizations need a way to translate those signals into a view of current state and next steps. That is the purpose of a maturity model.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 11: Agentic SOC Maturity Model

The Agentic SOC is not adopted in a single leap. Organizations progress through stages that reflect different levels of execution, integration, and governance.

Level 1: AI-Assisted SOC

At Level 1, AI improves prioritization, visibility, and summarization, but humans still execute investigations. Outcomes may improve modestly, yet backlogs, variability, and workload remain.

This stage is useful, but it does not change the operating model.

Level 2: Automated Workflows

At Level 2, organizations automate repeatable tasks through predefined playbooks and orchestration. Efficiency improves for known scenarios, but the system cannot adapt to new conditions or reason through ambiguity.

This stage reduces manual effort without creating true agentic execution.

Level 3: Partial Agentic Operations

At Level 3, some parts of the workflow become agentic. The system may autonomously triage, enrich, and investigate selected use cases, but coverage is uneven and architecture, context, or governance is still incomplete.

This is a transitional stage that often gets mistaken for the destination. The most recognizable hallmark is that investigations run autonomously for well-understood, high-volume threat types but revert to manual workflows for anything novel or cross-domain. That pattern indicates partial transformation, not completion.

Level 4: Fully Agentic SOC

At Level 4, AI agents own the full investigation lifecycle within a unified, governed architecture. Investigations run continuously, context is shared across domains, and human involvement is focused on oversight, approval, and exceptions.

This is the point at which the operating model has truly changed.

The Full Model

The maturity model matters because it distinguishes incremental progress from actual transformation. It helps organizations see whether they have improved the old system or built the new one.

Assessment Into Movement

Assessment should drive action. The point is not to label maturity for its own sake, but to identify what must change next in data, architecture, governance, and operations.

The Build Sequence

Once maturity is understood, organizations need a practical sequence for moving forward. The transition works best when each layer is established in the order the next layer depends on.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 12: How to Build an Agentic SOC

Building an Agentic SOC is a coordinated transformation, not a single deployment. Sequence matters because each phase depends on the integrity of the one before it.

Phase 1: Establishing the Data Foundation

Begin with the data layer. Ingest, normalize, enrich, and correlate telemetry across core security domains so the system has a shared operational foundation.

Without this phase, everything above it is unstable.

Phase 2: Building the Knowledge Layer

Next, build the context engine that models relationships, historical behavior, and environment-specific norms.

This gives the system the context required for meaningful reasoning.

Phase 3: Deploying the Agent Framework

Then deploy the agent framework: specialized agents, task decomposition, communication paths, and orchestration logic.

The goal is coordinated execution, not isolated agent demos.

Phase 4: Enabling Autonomous Investigation

Once agents and context exist, enable autonomous investigation. The system must be able to collect evidence, test hypotheses, update conclusions, and build structured cases without step-by-step human direction.

This phase is where automation gives way to reasoning.

Phase 5: Implementing Governance

As execution expands, implement governance: policies, confidence thresholds, approval rules, auditability, monitoring, and rollback.

Governance must shape autonomy rather than trail it.

Phase 6: Integrating Response Capabilities

After that, integrate governed response actions with the relevant systems so the platform can act as well as analyze.

Response must remain tightly coupled to policy and control.

Phase 7: Continuous Optimization

Finally, optimize continuously. Tune models, policies, and workflows. Monitor drift. Learn from outcomes.

Expand autonomy only as the system proves reliable.

Why Sequence Determines Outcome

Skipping ahead usually creates fragile systems. Organizations that build in sequence create stronger alignment across data, context, agents, governance, and execution.

Evaluating What Has Been Built

Once a system exists, whether internally built or externally sourced, it must be evaluated as a whole operating model rather than as a collection of attractive features.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 13: How to Evaluate Agentic SOC Solutions

Evaluating Systems, Not Features

The right evaluation lens is system capability, not feature presence. The question is not whether the solution includes AI, automation, or orchestration. It is whether those pieces combine into continuous, coordinated, governed execution.

13.1 Architectural Evaluation

Start with architecture.

Assess whether the solution has a unified data layer, a context engine, coordinated agents, an investigation engine, and integrated execution capabilities. Also determine whether the design is truly unified or simply assembled from fragmented components.

13.2 Agent Capabilities

Then assess the agents themselves.

Are they specialized? Are they orchestrated? Can they own investigations end to end, or do they merely generate outputs for humans to finish? These questions reveal whether AI is operating or only assisting.

13.3 Governance and Control

Evaluate how the system constrains autonomy.

What actions can it take? Under what confidence levels? When does it escalate? How are approvals, audit trails, and overrides handled? A system without clear answers here is not ready for meaningful autonomy.

13.4 Data and Context

Examine the data foundation.

What data is used, how is it normalized and enriched, and how environment-specific is the context layer? A solution without strong data and context will struggle regardless of its interface or agent branding.

13.5 Outcome Validation

Finally, validate outcomes.

Can the system explain decisions, demonstrate measurable results, and show consistent performance over time and across cases? Outcome validation is where
claims either hold up or fail.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 13: How to Evaluate Agentic SOC Solutions, Cont'd...

What Rigorous Evaluation Reveals

Rigorous evaluation often shows that many solutions have strong components but weak system alignment. They may excel in one use case or one domain while falling short of continuous, end-to-end operation. This is also why demo-driven evaluation is one of the most common failure modes in procurement. A demonstration shows best-case behavior on prepared inputs — not how the system performs when data is incomplete, when threat patterns are novel, or when governance boundaries are tested.

Every dimension of this framework should be assessed against real operational data and realistic investigation scenarios, not curated demonstrations.

The Risk of Partial Alignment

Partial alignment is dangerous because it can look like completion. Organizations may overestimate capability, expand trust too quickly, and discover only later that the system does not hold together under real operating conditions.

Evaluation as a Control Point

Evaluation should therefore be treated as a control point. It forces precision about architecture, autonomy, context, governance, and measurable outcomes before the organization commits at scale.

The Cost of Getting This Wrong

A poor decision here affects more than tooling. It shapes the future direction, cost structure, and risk posture of the security program.

Pitfalls That Derail Even Informed Organizations

Even organizations that understand the model can go wrong in execution. The most common pitfalls follow a predictable pattern.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 14: Common Pitfalls in Pursuing the Agentic SOC

Most failures come from misalignment between ambition and foundation. Organizations reach for agentic outcomes while keeping too much of the old model intact.

Confusing Automation with Agentic Systems

One common mistake is treating automation as the same as agentic execution. Automation is useful, but it handles fixed logic, not adaptive investigation.

Confusing the two causes organizations to overstate how far they have progressed.

From Automation to Generalized AI

A second mistake is assuming generalized AI will naturally solve the problem once automation hits its limits. Generalized models can be impressive, but without specialization, context, orchestration, and governance, they do not become a reliable operating model.

From Intelligence to Uncontrolled Execution

As systems gain analytical power, some organizations push toward action before they build adequate controls. That creates either unsafe autonomy or a rapid retreat to conservative restrictions that restore human bottlenecks.

From Governance Gaps to Data Limitations

Even when governance improves, data weaknesses often surface next. Fragmented, low-signal, or poorly contextualized data quietly limits the quality of decisions and increases friction throughout the system.

From Data to Workflow Misalignment

Organizations may also keep workflow assumptions that belong to the old SOC, such as queues, handoffs, or tool-by-tool investigation patterns. That misalignment prevents the architecture from delivering its full value.

The Compounding Nature of Misalignment

These pitfalls reinforce one another. Small gaps in governance, data, context, or workflow design accumulate until the system feels more complex without becoming more effective.

The remedy is building intentionally so that data, agents, governance, workflows, and human roles align to the same operating logic. That requires realism about current state, discipline about sequence, and precision about what the Agentic SOC actually is.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 15: The Future of the Agentic SOC

The Agentic SOC is not a distant concept. It is the direction security operations is already moving because the pressures driving it are structural and persistent.

From Augmentation to Execution

The first visible shift is from AI augmentation toward AI execution. Assistance is proving valuable, but it does not solve the labor and latency limits of the traditional SOC. Execution does.

That is why the role of AI will continue moving from support toward operational ownership.

Increasing Autonomy, Bounded by Control

Autonomy will increase, but so will the need for precision in control. The future is not unrestricted AI. It is increasingly adaptive autonomy bounded by policy, confidence, reversibility, and risk.

That creates a spectrum of autonomy rather than a binary choice.

From Reactive to Continuous Operations

Security operations will continue moving away from alert-driven, queue-based workflows toward continuous analysis and investigation.

The system will increasingly evaluate the environment as an ongoing process rather than reacting only after alerts appear.

Deepening Context Through Expanded Knowledge Layers

Knowledge layers will grow richer, incorporating more relationships, more history, and broader operational context.

That will improve reasoning and reduce dependence on static rules, but it will also increase the importance of data quality.

Convergence of Security Functions

As context and execution expand, functions that were once separated across tooling and teams will begin to converge. Detection, investigation, response, and related workflows will operate within a more unified system.

This does not eliminate specialized expertise. It changes where that expertise is applied.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 15: The Future of the Agentic SOC, Cont'd...

Human Roles in an Agentic Future

Humans will remain essential, but their work will continue shifting toward oversight, governance, exception handling, and strategic direction. Expertise becomes more leveraged because it is applied to the hardest decisions rather than the highest volume of tasks.

Organizations will need to adapt skills, roles, and operating expectations accordingly.

The Emerging Standard

Over time, the Agentic SOC will move from emerging model to expected standard for advanced security operations.

Not every organization will move at the same pace, but the direction is clear.

The Defining Shift

The defining shift is simple: security execution moves from humans supported by tools to systems governed by humans.

That is the core operating change around which the future of the SOC will be organized.

Closing Perspective

The Agentic SOC exists because current conditions require it. The pace and complexity of modern threats have already made the old model increasingly unsustainable.

Organizations can move toward the new model deliberately or be pushed toward it by necessity.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Chapter 16: Conclusion

The argument of this guide is straightforward. The traditional SOC is constrained by a human-centered execution model that no longer scales. The Agentic SOC addresses that limit by shifting core execution to AI agents operating continuously, in parallel, and within governed boundaries.

That shift is only real when architecture, data, governance, and workflows all align to support it. Without those foundations, organizations get enhanced versions of the traditional SOC rather than a true operating model change.

The path forward is demanding, but the direction is increasingly clear. Security operations will be defined less by who can accumulate the most tools and more by who can build systems that can investigate, decide, and act at the speed the environment requires.

And it is no longer a question of whether security operations moves in this direction. The only question is whether organizations get there by design or by necessity.

Appendix A: Key Concepts and Terminology

Agentic SOC

A security operations model in which AI agents execute the core workflows of detection, investigation, and response, with humans focused on oversight, governance, and exceptions.

Agent (Security Context)

A specialized AI component responsible for a defined function such as detection, enrichment, investigation, response, or validation, operating independently and in coordination with other agents.

Autonomous Investigation

The ability of the system to initiate, conduct, and complete an investigation by collecting evidence, evaluating hypotheses, and reaching conclusions without step-by-step human direction.

Continuous Security Operations

A model in which analysis and investigation happen continuously across the environment instead of waiting in alert-driven queues.

Knowledge Layer

A contextual representation of the environment that models relationships among users, devices, identities, applications, behaviors, and history.

Context-Aware Reasoning

Decision-making that evaluates activity in relation to organizational behavior, entity relationships, and historical patterns rather than as isolated events.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

Appendix A: Key Concepts and Terminology, Cont'd...

Bounded Autonomy

A controlled form of autonomous action in which system behavior is limited by policy, confidence thresholds, approval rules, and risk constraints.

AI-Led Execution

An operating model in which AI systems, rather than humans, drive the primary workflows of security investigation and response.

Automation vs Agentic Execution

Automation executes predefined steps under predefined conditions. Agentic execution reasons through context and adapts its actions dynamically.

AI Assistance vs AI Operation

AI assistance helps humans work faster or better. AI operation means the system itself initiates, progresses, and completes the workflow.

Workflow vs Investigation Lifecycle

A workflow is a fixed sequence of steps for a known scenario. An investigation lifecycle is a dynamic process that adapts as evidence is gathered and interpreted.

Fragmented vs Unified Architecture

Fragmented architecture connects separate tools with separate data models and logic. Unified architecture provides a shared data, context, and execution framework across the system.

Detection vs Continuous Investigation

Detection identifies potentially suspicious events. Continuous investigation extends beyond alerting to ongoing, contextual analysis of behavior across the environment.

Data Normalization, Enrichment, and Correlation

Normalization standardizes data, enrichment adds context, and correlation links related events so raw telemetry becomes usable for reasoning and action.

Governance, Control, and Validation

Governance defines policy, control enforces it, and validation confirms that system decisions and actions are appropriate and effective.

Model Accuracy vs Operational Reliability

Model accuracy measures performance under defined conditions. Operational reliability measures whether the full system produces correct, contextually appropriate outcomes consistently over time.

Closing Note

These terms are interdependent. Used together, they define the operating model, design requirements, and evaluation standard of the Agentic SOC.

ARCTIC WOLF

©2026 Arctic Wolf Networks, Inc., All Rights Reserved. Arctic Wolf, Arctic Wolf Platform, Arctic Wolf Security Operations Cloud, Arctic Wolf Managed Detection and Response, Arctic Wolf Managed Risk, Arctic Wolf Managed Security Awareness, Arctic Wolf Incident Response, and Arctic Wolf Concierge Security Team are either trademarks or registered trademarks of Arctic Wolf Networks, Inc. or Arctic Wolf Networks Canada, Inc. and any subsidiaries in Canada, the United States, and/or other countries.