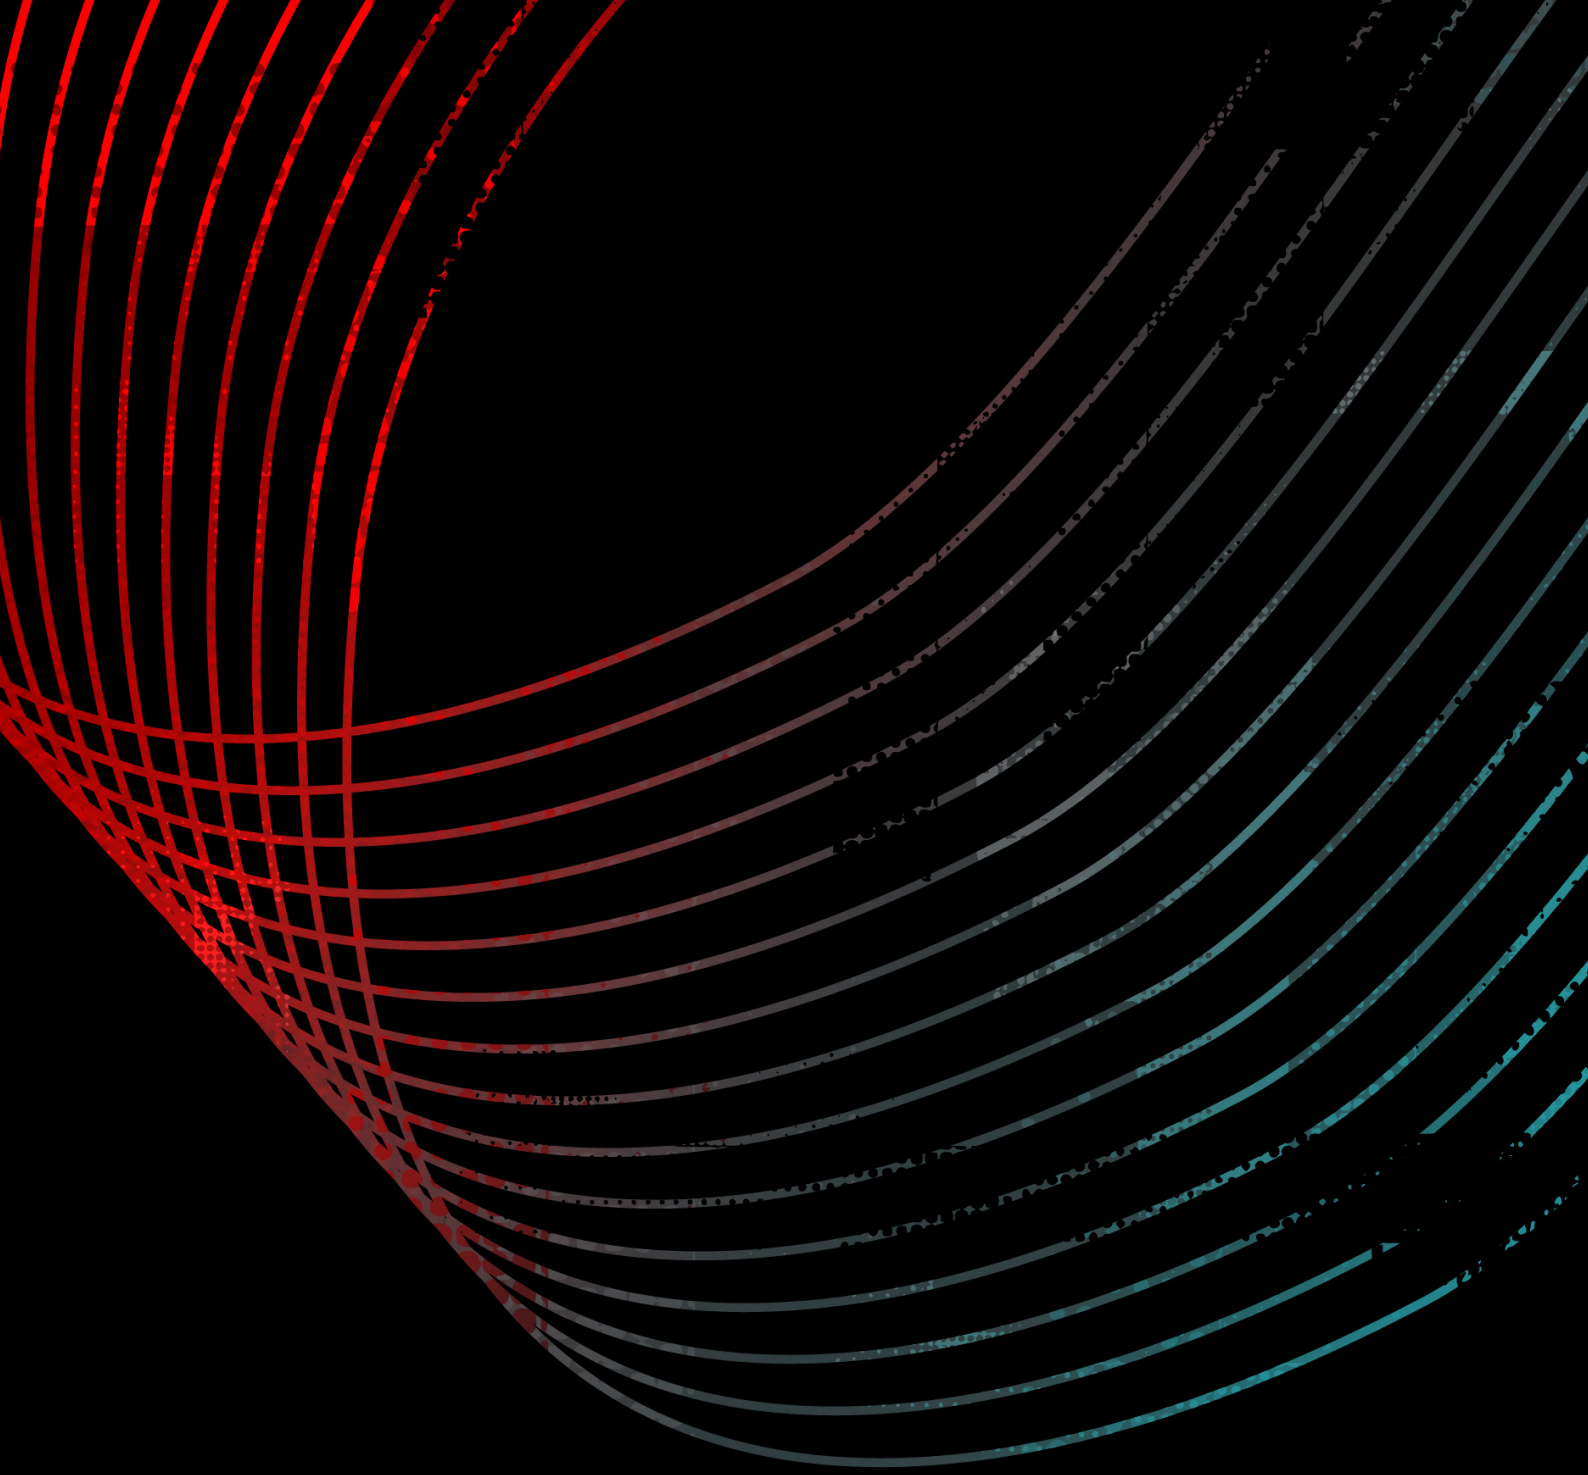

CROWDSTRIKE

The Agentic SOC Guide

A Four-Step Journey to AI-Powered Security Operations

Table of Contents

AI Is Reshaping Cybersecurity

Adversaries are using AI to attack at machine speed

51s

Fastest breakout time in 2024[1]

AI adoption expands the enterprise attack surface

33%

of surveyed organizations report addressing GenAI security risks[2]

Defenders are mobilizing agents across the SOC

30%

of cybersecurity teams have integrated AI tools[3]

AI is changing the rules of engagement in cybersecurity — and they're changing faster than most security teams can react.

Attacks now unfold in seconds, not hours. Adversaries are weaponizing AI, deploying intelligent agents across every phase of the kill chain to amplify speed, scale, and sophistication.

But the same technology fueling the threat is also redefining defense. AI is evolving the SOC from reactive, manual processes to proactive, agentic operations. But many organizations have yet to adopt AI, leaving their SOCs struggling to keep pace with AI-powered adversaries.

This marks a defining inflection point for cybersecurity. The time for transformation isn't in the future — it's now. Security teams that embrace AI-driven operations will gain a decisive advantage, evolving from reactive firefighting to proactive, predictive defense.

[1] CrowdStrike 2025 Global Threat Report

[2] State of AI: McKinsey Report 2024

[3] ISC 2025 AI Pulse Survey

Three Growing Gaps in Security Operations

The modern SOC is at a breaking point. Adversaries move at machine speed, while defenders are constrained by manual processes, fragmented tools, and limited human capacity. Incremental automation — the scripts and playbooks of the past decade — can no longer close the gap.

Three critical threats to the SOC are:

The Labor Gap: There are more than 4 million unfilled cybersecurity positions worldwide,[4] and the shortage is growing each year. Even elite analysts are burning out under relentless alert volume and cognitive overload.

The Skills Gap: 90% of organizations report that their teams lack the advanced expertise required to counter today's threats[5] — from cloud exploitation to AI-powered attacks.

The Response Gap: The fastest recorded eCrime breakout time was just 51 seconds in 2024, according to the CrowdStrike 2025 Global Threat Report — far faster than human analysts can triage, investigate, and contain manually.

Together, these gaps have created an unsustainable imbalance — a widening divide between attacker speed and defender capability.

> **"The legacy SOC is trying to fight a 21st-century war with 20th-century weapons.**
>
> **Security teams have become modern-day firefighters — constantly responding to the next alert, overwhelmed by data volume, and struggling to keep pace as attack timelines compress."**
>
> George Kurtz
>
> **CEO, CROWDSTRIKE**

[4] ISC2 Cybersecurity Workforce Study 2024: Employers Must Act as Cybersecurity Workforce Growth Stalls and Skills Gaps Widen

[5] KPMG Security Operations Center Survey 2024: The time to transform is now, p.17

The Agentic SOC:

A New Operating Model for the AI Era

The answer to today's speed and scale gaps isn't another dashboard or incremental playbook — it's a new operational model.

The agentic SOC represents the next evolution in security operations — one where intelligent AI agents reason, decide, act, and learn across domains, operating alongside human expertise to deliver continuous, adaptive defense.

How AI Solves the Three Critical Gaps in the SOC

The Labor Gap: AI augments analyst capacity with autonomous triage, investigation, and response — scaling human expertise without adding headcount.

The Skills Gap: AI codifies expert tradecraft and delivers guided reasoning, empowering every analyst to perform at an elite level.

The Response Gap: AI accelerates detection-to-containment, operating at machine speed to stop breaches in seconds — not minutes or hours.

Where Agents Add Value — and Where Humans Lead

Some tasks are best left to machines; others demand human intuition, experience, and judgment. The key is balance. This matrix maps core SOC functions to their automation potential, showing how AI and humans work together under defined guardrails — and the outcomes that transform security operations.

SOC FUNCTION	AUTOMATION POTENTIAL	AI ROLE	HUMAN ROLE	OUTCOME
Alert Triage and Prioritization	HIGH	Analyze context, enrich alerts, and suppress false positives	Validate escalations and refine rules	Reduction in noise — analysts focus on true positives
Investigation and Correlation	HIGH	Stitch together telemetry, build attack narratives, and assess scope	Review complex incidents and apply judgment on intent	Machine-speed investigation with expert-level consistency
Reporting and Knowledge Capture	HIGH	Auto-generate incident summaries, metrics, and playbook updates	Ensure accuracy, tone, and institutional learning	Accelerated reporting and institutionalized expertise
Response and Containment	MEDIUM	Execute playbooks — isolate hosts, disable accounts, block processes	Approve or override actions in high-risk cases	Faster containment, governed by policy and oversight
Threat Hunting and Intelligence Integration	MEDIUM	Surface anomalies and correlate intel to drive proactive hunts	Define hypotheses, validate insights, and adapt models	Continuous learning and proactive defense
Governance and Strategy	LOW	Get AI assistance with analytics and recommendations	Define policy, assess risk, and set operational priorities	AI informs and humans decide

Table 1. *Core SOC functions and their automation potential*

Rule of thumb: Automate where context and consistency matter most, and keep humans in command where judgment, escalation, or accountability are critical.

How AI Transforms Roles Across the SOC

AI will change how your SOC operates. Understanding how each role will change will help you prepare for the future of work.

	Existing Reality	Future Reality
 **CISO**	Works hard to manage risk amid fragmented visibility and manual reporting Insights often arrive after the fact, making it difficult to stay ahead of emerging threats.	Gains unified, real-time visibility across the enterprise AI-driven analytics surface emerging risks instantly. Governance shifts from reactive compliance to proactive risk optimization.
 **SOC Manager**	Balances growing alert volumes, talent gaps, and operational demands Maintaining consistency across shifts and geographies is a constant challenge, and scaling often means adding more people, not outcomes.	Manages a coordinated ecosystem of human and AI agents Operations become faster, more accurate, and inherently scalable. Standardization and performance improve automatically as AI learns.
 **Analyst**	Spends much of the day sorting alerts and performing repetitive investigations Valuable time and expertise are consumed by noise, leaving little room for creative problem-solving or skill growth.	Focuses on high-value analysis and proactive threat hunting AI agents autonomously triage, investigate, and summarize incidents. Work becomes strategic, creative, and impactful — not reactive.
 **Security Engineer**	Devotes significant effort to integrating tools, maintaining scripts, and managing complexity Innovation is slowed by technical debt and the limits of manual automation.	Designs and defines guardrails for custom agents, then deploys and orchestrates them through a unified platform Easily build new workflows, integrate intelligence, and orchestrate secure, autonomous operations across the SOC.

Why the Agentic SOC Matters

The agentic SOC represents more than an evolution in tooling — it's a fundamental shift in how security operations think, act, and scale. The legacy SOC, limited by human time and manual process, can no longer keep pace with machine-speed adversaries who exploit automation, AI, and global scale to outmaneuver defenders. Every second now matters.

The agentic SOC bridges the gap between today's overwhelmed operations and tomorrow's autonomous, adaptive defense. It delivers scale without compromise — automating the tasks that slow teams down while preserving and amplifying human expertise. In this model, analysts aren't replaced; they're empowered. AI agents handle the repetitive, time-consuming work of triage, correlation, and containment, freeing humans to focus on strategy, judgment, and innovation.

The result:

A SOC that moves at the speed of the adversary — or faster — guided by AI precision and human command.

The Path Forward:
A Framework for Agentic SOC Transformation

Moving from overwhelmed, human-centered operations to a human-in-command, agentic SOC doesn't happen overnight. It's a structured journey that builds capability in layers: Get the data right, prove value with governed automation, customize agents to your unique workflows, and then orchestrate agents in a coordinated system. Each step is measurable and designed to raise confidence while reducing risk.

The four phases at a glance are:

1

Build an AI-Ready Foundation

Make your data ready for AI. Humans and agents alike need a unified, clean, and contextual view of your environment to reason and act effectively in real time.

2

Identify Quick Wins and Ready-to-Use AI Opportunities

Start by auditing where AI and agentic technology can deliver value immediately. Focus on use cases that require minimal customization — areas where out-of-the-box AI can automate repetitive, high-impact tasks or augment existing workflows.

3

Customize and Evolve Agents to Scale Your Expertise

Once you've captured early wins with ready-to-use AI, extend that value by customizing agents to reflect how your team thinks, decides, and operates. Infuse domain expertise, playbooks, and decision logic so your AI agents evolve alongside your organization.

4

Orchestrate a Multi-Agent Defense

As your ecosystem of agents matures, connect and coordinate specialized agents across workflows, keeping humans in command while enabling adaptive, cross-domain defense at machine speed.

PHASE 1: BUILD AN AI-READY FOUNDATION

Data is the foundation for AI. Without a unified, high-quality data layer, AI lacks the context it needs to see, decide, and act at machine speed. Fragmented data tells only part of the story, forcing teams to manually stitch together events across tools. To unlock the full potential of the agentic SOC, organizations must first build an AI-ready data foundation that delivers clean, complete, real-time visibility across the attack surface. This starts with unifying and operationalizing data across every domain — endpoint, cloud, identity, network, and more. Instead of pushing AI into fragmented tools, the foundation must bring data into a unified data layer, in one format and in real time.

Modern SOCs increasingly treat telemetry as a shared service — Collected once, enriched once, and delivered everywhere. By consolidating telemetry within a single security data layer, teams eliminate silos and give AI agents the full context needed to reason, correlate, and act. Streaming ingestion, schema normalization, and contextual enrichment ensure that every signal carries the who, what, where, and when — the critical ingredients for machine-speed decisions.

An AI-ready data layer doesn't just improve detection — it simplifies operations, cuts data management costs, and ensures every agent operates with complete, current intelligence. The result is a SOC that's faster, smarter, and ready for autonomy — one where data fuels every decision, and AI never operates in the dark.

Getting Data In — and Making It Work

Building an AI-ready foundation requires more than collecting data — it demands getting it in fast and making it usable the moment it arrives.

Endpoint agents, cloud APIs, identity logs, and network sensors all produce telemetry with their own formats, timestamps, and levels of fidelity. Traditional data ingestion models were not built for such high volumes of data or this level of complexity. The SOC needs real-time, reliable ingestion that also allows analysts to filter, route, redact, and transform data in motion as well as effectively manage costs.

Getting data in is only half the battle. With legacy security information and event management (SIEM) tools, teams ingested all data to log everything and search later. But to keep up with threats in the AI era, teams need to detect, investigate, and respond to threats in real time. To make data actionable, they need an environment that unifies data and context. With a unified data model, teams can correlate across domains to connect users, assets, and events into a complete picture of an attack and apply real-time analytics to drive faster, more precise decisions.

Migration Considerations

Teams moving from a legacy SIEM need to migrate data without compromising their operations. Developing a migration strategy is key to ensuring a seamless transition. To reduce risk and verify fidelity, many organizations adopt a parallel-run approach, routing the same telemetry to legacy and modern platforms during validation.

LEGACY SIEM TO AGENTIC SOC

Migration Timeline and Phases

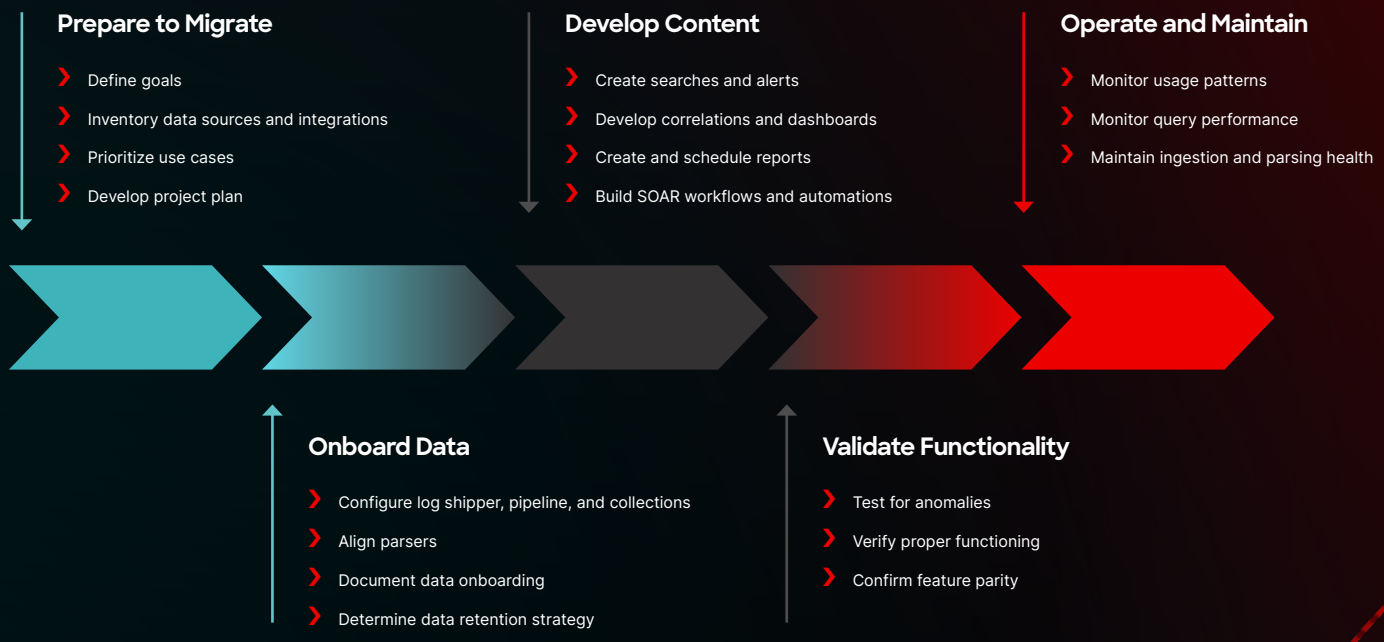

PHASE 2: IDENTIFY QUICK WINS AND READY-TO-USE AI OPPORTUNITIES

With an AI-ready data layer established, security teams can now operationalize intelligence through agentic automation. Phase 2 focuses on deploying pre-built, out-of-the-box agents designed to handle discrete, high-volume SOC tasks. These agents act as specialized digital teammates, executing well-defined tasks such as alert triage, CVE and exploitability analysis, correlation rule generation, and incident documentation — faster and more consistently than human analysts alone. Delivered by trusted platforms, these agents apply elite, proprietary reasoning at machine speed.

Unlike traditional playbook automation, these agents reason with rich, contextualized data and adapt their actions to evolving threat signals. They don't just execute workflows, they understand them, using built-in reasoning and context-awareness to learn over time.

But not all tasks will be suited to AI agents. Just as important as identifying where to apply agents is knowing when not to use them. The key to success is identifying tasks that are repeatable, bounded, and context-dependent — areas where precision and speed are paramount. For example, alert triage or correlation rule building benefit from autonomous reasoning and complex pattern recognition, while tasks requiring nuanced judgment or adversarial interpretation remain best handled by analysts. Conversely, tasks that are highly deterministic or require simple, repeatable execution can more likely be handled through traditional rule-based automation or security orchestration, automation, and response (SOAR) playbooks. In such cases, deploying AI agents can introduce unnecessary complexity, cost, unpredictability, and maintenance overhead without delivering proportional value. The goal is not to replace existing automation, but to extend it — reserving agents for problems best suited to AI.

By transferring repetitive and procedural workloads to agents, SOCs can shift human focus to higher-order reasoning and strategic long-term initiatives — tasks that demand extensive critical thinking, contextual interpretation, and coordinated response across teams. This shift elevates the role of human analysts from reactive firefighters to strategic overseers and agent orchestrators, scaling operations and driving long-term operational resilience.

Prioritize Where to Deploy AI Agents

To help security teams identify where to start, consider mapping common security tasks against organizational impact and implementation difficulty.

- **High-impact, low-difficulty** tasks should be prioritized as quick wins to demonstrate immediate value.
- **High-impact, high-difficulty** tasks can follow once data maturity and trust models are established.
- **Low-impact, low-difficulty** tasks may remain better suited to traditional rule-based automation.
- **Low-impact, high-difficulty** areas should generally be avoided to prevent wasted effort or risk.

AI Agent: Task Prioritization Framework

	LOW DIFFICULTY / LOW OVERHEAD	HIGH DIFFICULTY / HIGH OVERHEAD
HIGH IMPACT	**Quick Wins (deploy first)** Tasks that deliver immediate ROI and measurable analyst relief with minimal integration complexity *Examples:* Alert triage, enrichment correlation, CVE analysis, malware classification	**Strategic Investments** High-value use cases that require advanced reasoning, orchestration, or data maturity before deployment *Examples:* Multi-agent workflows, adaptive response coordination
LOW IMPACT	**Maintain with Rule-Based Automation** Tasks that are repetitive, well-structured, or deterministic — better handled through SOAR or scripts *Examples:* Notification triggers, routine data normalization	**Avoid or Defer** Tasks with low impact but high complexity or ambiguity — automating these may increase operational risk or overhead without meaningful benefit *Examples:* Incident escalation decisions, legal/compliance notifications, attribution analysis

Change Management and Best Practices for Deploying Agents

Deploying out-of-the-box agents isn't just a technical exercise — it's a shift in your organization's operating model. To realize the full value of agentic automation, SOCs should approach deployment as a phased change program, emphasizing transparency, control, and continuous improvement.

1. **Start Small, Scale Fast:** Begin with one or two low-risk, high-volume use cases — such as alert triage or enrichment — to validate accuracy and performance before expanding across workflows. This "pilot-to-scale" model allows teams to measure return on investment (ROI) early while refining guardrails.

2. **Maintain Human Oversight:** Every agent action should remain observable, explainable, and reversible. Establish clear escalation and approval workflows so analysts can supervise agent decisions with transparency and explainability. Make agent reasoning explainable — every action should be traceable to a data source, rule, or contextual insight.

3. **Align Automation with Business Value:** Prioritize use cases with measurable impact such as reducing analyst burnout, improving mean time to respond, or reducing exposure. Automation that drives measurable outcomes reinforces strategic alignment and will enable teams to drive buy-in and executive support.

Example Quick Wins with Agent-based Automation

Below are a few examples of high-impact, low-effort tasks well-suited to agent-based automation with out-of-the-box agents.

AUTOMATION OPPORTUNITY	SOC FUNCTION	DESCRIPTION	EXAMPLE OUTCOMES
Triage Analysis	Alert Triage and Prioritization	Classify alerts as true/false positives, prioritize by confidence and context, and recommend next steps	• **Faster mean time to respond (MTTR)** • **Reduced alert fatigue** • **Improved prioritization accuracy**
Investigation Analysis	Investigation and Correlation	Autonomously drive investigations, correlate evidence, summarize findings, and draft analyst reports	• **Less burnout** • **Improved accuracy and consistency**
Hunting Query Analysis	Threat Hunting and Intelligence Integration	Generate and refine hunting queries across endpoint, identity, and network telemetry based on recent threat intel and indicators of compromise (IOCs)	• **Faster hunts** • **Broader coverage** • **Higher-fidelity detection hypotheses**
Malware Analysis	Threat Hunting and Intelligence Integration	Detonate and analyze malware, correlate behavioral patterns against threat intel; classify malware families, generate YARA rules	• **Rapid malware classification** • **Faster containment** • **Enriched threat intel**
CVE Analysis	Alert Triage and Prioritization	Map new CVEs to the environment, assess exploit availability and impact, and recommend prioritized remediation actions	• **Reduced patch backlog** • **Faster risk mitigation**
Risk Exploitability Analysis	Alert Triage and Prioritization	Evaluate exploit likelihood using threat intel, attack path simulation, and business asset context	• **Reduced exposure window** • **Improved risk visibility**
Correlation Rule Generation	Investigation and Correlation	Generate and optimize detection rules	• **Faster mean time to detect (MTTD)/MTTR**
Data Normalization and Transformation	Investigation and Correlation	Transform diverse log formats into normalized schemas; enrich with context from assets, users, and threat intel	• **Consolidation savings** • **Improved accuracy** • **Enhanced correlation and visibility**
SOAR Workflow Generation	Response and Containment	Build and optimize automation playbooks	• **Faster MTTR** • **Time savings**

PHASE 3: CUSTOMIZE AND EVOLVE AGENTS TO SCALE YOUR EXPERTISE

Once foundational automation is in place, the next evolution of an agentic SOC is customization — building agents that reflect not just what your SOC does but tailored to your team's needs. As SOCs mature in their use of out-of-the-box agents, many discover workflows that reflect the unique way their organization operates — the policies or workflows that define their unique day-to-day operations. These are ideal candidates for custom agents designed around your specific data, tools, and decision models.

Creating custom agents doesn't mean starting from scratch. It means designing reasoning and workflow logic that aligns to your environment — mapping each agent's knowledge, instructions, context, authorized actions, and built-in guardrails. Custom agents capture that logic once and apply it consistently, preserving institutional knowledge while accelerating repetitive cognitive tasks.

Teams typically identify a need for custom agents when:

- **Out-of-the-box agents lack relevance to your environment:** Reasoning is required over internal datasets or third-party telemetry not covered by out-of-the-box agents.
- **Processes are unique to your SOC:** Examples include specialized escalation or proprietary data correlation.
- **Human reasoning creates a bottleneck:** Analysts repeatedly apply judgment patterns that could be codified and reused.

To ensure scalability and trust, design custom agents with four principles in mind:

Composable, not monolithic: Build small, task-specific agents that can interoperate across workflows.

Governed by design: Every agent should operate within defined scope, visibility, and controls.

Human-in-command: Analysts remain the orchestrators, providing oversight, feedback, and escalation authority.

Outcome-driven design: Target time-intensive, reasoning-heavy tasks where automation can reclaim analyst capacity and improve consistency.

Before You Build: Assess Readiness for Custom Agents

Before developing custom agents, evaluate whether your SOC is ready across four dimensions:

- [ ] **People**
  - Analysts understand how to supervise, validate, and refine agent reasoning.
  - Roles are defined for feedback, escalation, and oversight.
- [ ] **Process**
  - Key workflows are documented, consistent, and repeatable.
  - Decision logic and escalation paths are clearly defined.
- [ ] **Data**
  - Telemetry is unified, trustworthy, and context-rich.
  - Data sources provide the depth that agents need to reason effectively.
- [ ] **Tooling and Environment**
  - Infrastructure supports safe, iterative testing and version control.
  - Governance frameworks exist for auditing and refining agent decisions.

If your SOC meets these criteria, you're ready to start encoding institutional knowledge into custom AI agents that amplify human reasoning and accelerate security outcomes.

SOCs that invest in custom agents move beyond efficiency gains — they build a living system of agents that mirrors their best analysts' reasoning, scales it across the organization, and evolves continuously. Each agent becomes a codified form of expertise — a force multiplier that strengthens resilience and response maturity over time.

PHASE 4: ORCHESTRATE A MULTI-AGENT DEFENSE

As agents accumulate, the SOC evolves into an agentic ecosystem — a coordinated network of interoperable, context-aware agents that reason, act, and adapt together across workflows and tools. The next evolution of SOC transformation is orchestration — where agents collaborate to reason, act, and adapt as a coordinated system.

Where earlier phases focused on single-agent efficiency, agentic automation introduces collaboration at scale — linking multiple agents to think, act, and adapt together across tools, data, and workflows.

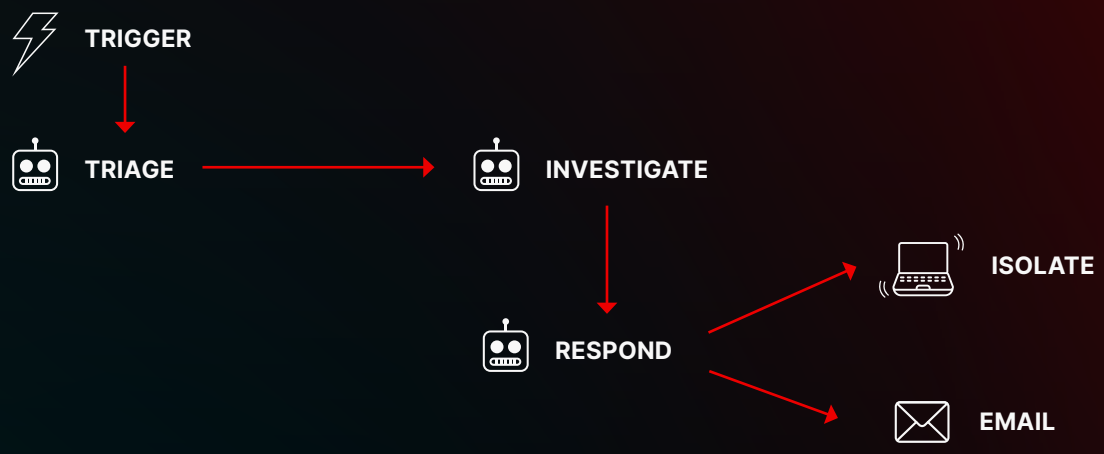

In this model, multi-agent workflows become force multipliers. A triage agent can automatically pass findings to an investigation agent for in-depth analysis correlation, which then triggers a response agent to recommend containment actions in the endpoint detection and response (EDR) or identity and access management (IAM) system. A compliance agent can then summarize incident artifacts for audit, while a correlation rule agent authors new detection rules for similar alerts. Each agent contributes its expertise to a broader ecosystem of reasoning — sharing context, validating outcomes, and acting in concert to achieve faster, more complete responses.

Through these interlinked workflows, the SOC moves from manual coordination between tools to adaptive, cross-domain defense driven by intelligent collaboration. This shift doesn't just accelerate response — it transforms the SOC into a continuously learning, adaptive defense system that scales expertise and strengthens resilience as it grows.

Principles for Building AI That's Safe, Secure, and Accountable

Even as agents gain autonomy, humans must remain firmly in command. Successful agentic SOCs will balance automation with control, ensuring every decision is transparent, traceable, and reversible. To sustain trust and safety, organizations should implement structured oversight and continuous governance around agentic workflows:

1. **Define Agent Boundaries and Roles:** Set explicit limits for what agents can and cannot do. Enforce principle of least privilege across all agent actions. Distinguish between recommendation and execution roles.
2. **Prioritize Agent Explainability and Auditability:** Every agent decision should be interpretable and traceable, from data source to outcome. Visibility into reasoning builds analyst confidence and ensures compliance with operational and regulatory standards.
3. **Maintain Human-in-the-Loop Review:** Keep analysts embedded in the feedback cycle. Encourage analysts to flag, correct, or augment agent outputs. Require human validation for high-impact actions (such as containment or escalation), use review queues for decision verification, and regularly monitor agent performance for signs of drift. These interactions form a feedback loop that strengthens accuracy and ensures automation evolves in alignment with human judgment.

Key Challenges — and How to Mitigate Them

Prepare for challenges. As you introduce autonomy, the same forces that unlock speed and scale can also amplify risk. Success hinges on four foundations: data quality, governance, skills, and securing the AI itself — with humans firmly in command.

1. **Data Readiness and Visibility Gaps**

   AI agents are only as strong as the data they see. Fragmented telemetry and siloed tools limit their ability to reason holistically across the attack surface.

   → SOLUTION: UNIFY TELEMETRY, NORMALIZE SCHEMAS, AND ENRICH HUMAN EXPERTISE AND THREAT INTELLIGENCE.

2. **Trust, Control, and Governance**

   Autonomous action requires strong governance frameworks. Defining clear boundaries, auditability, and escalation paths ensures AI-driven actions remain explainable and compliant.

   → SOLUTION: DEFINE GUARDRAILS, ENFORCE AUDITABILITY, AND REQUIRE EXPLAINABLE ACTIONS.

3. **Human-AI Collaboration and Skill Evolution**

   SOC teams must evolve from task execution to AI supervision and orchestration. Success depends on retraining analysts to interpret, guide, and improve AI behavior — not compete with it.

   → SOLUTION: ESTABLISH AN ANALYST UPSKILLING PROGRAM, AND KEEP HUMANS IN THE LOOP SO ANALYSTS CAN CONTINUE TO DEVELOP CYBERSECURITY EXPERTISE WHILE GOVERNING AI.

4. **Security of the AI Itself**

   As SOCs adopt agentic systems, adversaries will target the AI models, data pipelines, and integrations. Protecting the AI operating environment — from model poisoning to prompt injection — becomes critical.

   → SOLUTION: APPLY THREAT MODELING, SEGMENT AI ENVIRONMENTS, AND MONITOR FOR DRIFT AND POISONING.

What Success Looks Like — and How to Prove It

Broader Coverage and Proactive Defense

Before: Threat coverage was limited by human scale; intelligence integration lagged behind adversary innovation.

After: AI agents deliver continuous coverage across endpoints, cloud, identity, and data. The result: Threat intelligence and detections update in real time, operationalizing emerging intelligence at pipeline-native speed to block attacks before they spread.

→ MEASURE IT: INCREASED TELEMETRY COVERAGE, REDUCED DWELL TIME, AND FEWER MISSED DETECTIONS

Improve Quality and Consistency of Investigations

Before: Analysts faced alert fatigue, inconsistent investigations, and knowledge silos between shifts or regions.

After: AI agents and agentic workflows encode elite analyst reasoning. Every investigation follows best practices automatically, with human experts focusing on strategic decisions and edge cases. The result: consistent, expert-grade investigations, every time.

→ MEASURE IT: LOWER FALSE POSITIVES, FASTER TRIAGE-TO-RESOLUTION, AND CONSISTENT PLAYBOOK ADHERENCE

Accelerate and Scale Response

Before: Manual triage and incident handling slowed MTTR, especially under surge conditions.

After: Autonomous triage, correlation, and remediation allow response to scale infinitely, across workloads, environments, and geographies. The result: The SOC shifts from responding to incidents to orchestrating an intelligent defense network.

→ MEASURE IT: IMPROVED MTTR, THROUGHPUT PER ANALYST, AND TIME SAVED PER INVESTIGATION

From Vision to Reality:
Powering the Agentic SOC with CrowdStrike

CrowdStrike provides the unified platform and AI architecture that enable secure, scalable, and agentic security operations. This roadmap is designed to help security leaders build AI maturity step by step, starting with a unified, AI-ready data layer, advancing through agentic automation, and culminating in a collaborative ecosystem of interoperable AI agents. Each phase builds upon the last, turning intelligence into action and automation into advantage, all powered by the CrowdStrike Falcon® platform.

PHASE 1: CONSOLIDATE ON CROWDSTRIKE'S AI-READY DATA LAYER

CrowdStrike's unique data advantage is the foundation of its AI leadership — a continuously expanding universe of high-fidelity telemetry that fuels intelligence, automation, and protection at scale. The Falcon platform ingests trillions of security events every day, normalizing and enriching them with elite human intelligence and a decade of adversary tradecraft — creating the richest, most battle-tested data set in cybersecurity.

At the core of this ecosystem is **CrowdStrike Enterprise Graph®** — the architectural backbone of the Falcon platform. It continuously connects users, devices, workloads, intelligence, and adversary behaviors into a single, contextual understanding of the enterprise. This living, dynamic model gives both humans and AI the complete picture needed to uncover attack paths, identify emerging risks, and take action with confidence.

Consolidate on an AI-ready Data Layer

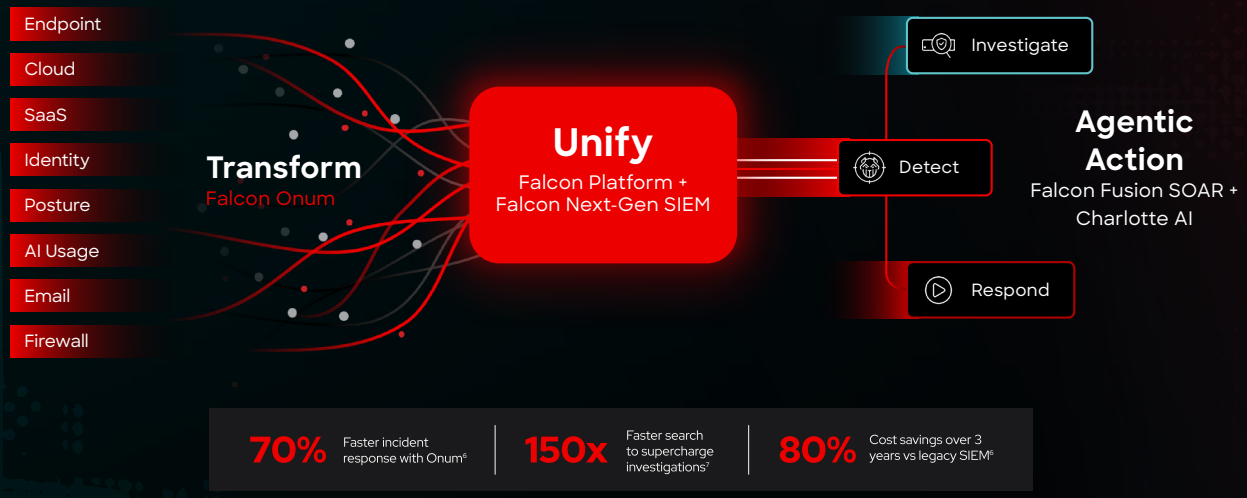

70% Faster incident response with Onum[6]

150x Faster search to supercharge investigations[7]

80% Cost savings over 3 years vs legacy SIEM[6]

Together, these innovations form the data and intelligence layer that makes CrowdStrike uniquely capable of delivering **AI-powered security in real time — across every endpoint, identity, workload, and cloud.**

[6] Numbers are projected estimates of average benefit based on company's own internal analysis and recorded metrics provided by customers during pre-sale motions that compare the value of CrowdStrike with the customer's incumbent solution. Actual realized value will depend on the customer's module deployment and environment.

[7] Results are from a case study. Individual results may vary.

Building on the industry's richest data and intelligence, the Falcon platform was purpose-built to turn that data advantage into action — uniting intelligence, automation, and AI into one continuous defense system. This transformation starts with **real-time, AI-ready telemetry** flowing through Falcon Onum, our high-speed pipeline that refines and enriches data in motion. It collects, filters, enriches, and routes telemetry in real time, optimizing how data flows across the ecosystem. Falcon Onum ensures that every element of the Falcon platform operates on consistent, high-fidelity data, enabling autonomous and AI-driven security operations at scale.

The data from Onum and the Falcon platform is unified and operationalized in Falcon Next-Gen SIEM — the engine of the agentic SOC. Falcon Next-Gen SIEM powers lightning-fast search at petabyte scale, precise detections, faster response, and continuous insight. By combining cross-domain telemetry, adversary intelligence, and real-time analytics, Falcon Next-Gen SIEM empowers teams to anticipate, investigate, and stop threats at machine speed.

Next-Gen SIEM in Action: Automated Parser Creation

AI-generated parsers in Falcon Next-Gen SIEM solve one of the hardest problems in traditional SIEMs: getting data in. Instead of days spent hand-coding and tuning custom log parsers, AI automatically analyzes sample logs and builds accurate, normalized parsers in minutes, all compliant with the CrowdStrike Parsing Standard (CPS).

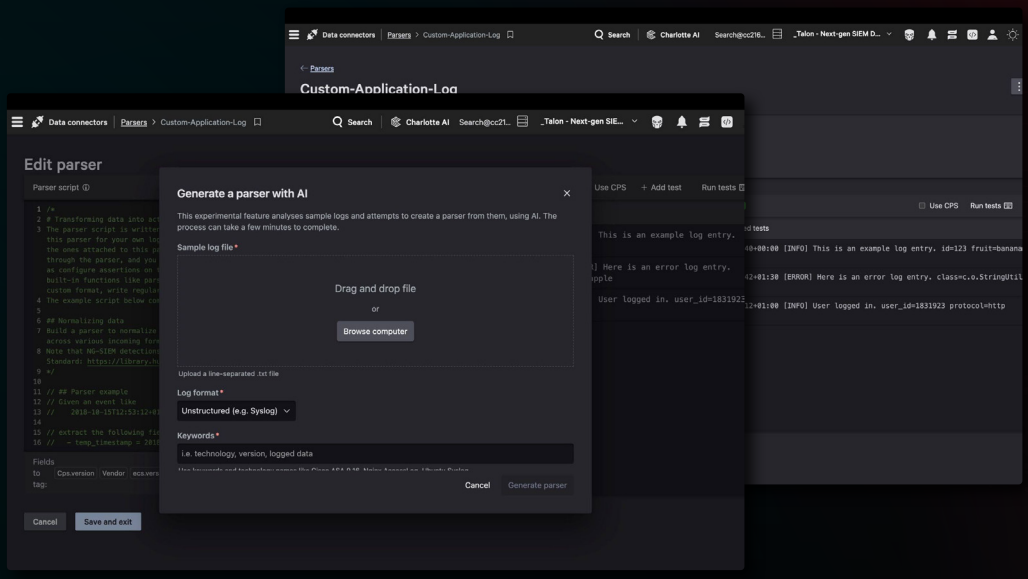

How It Works

1. Upload or select a sample log
2. AI models inspect log structure and identify key fields, delimiters, and event types
3. A CPS-compliant parser is automatically generated
4. Validate the parser against test data
5. Deploy the parser immediately or edit further

PHASE 2: DEPLOY CROWDSTRIKE'S WORKFORCE OF TURNKEY AGENTS FOR COMMON TASKS

Charlotte AI

Charlotte AI is CrowdStrike's AI security analyst, powering all agentic and generative AI capabilities across the Falcon platform. From natural language interactions to command-line explanations and Tier 1 triage analysis, Charlotte AI is purpose-built to help teams accelerate security operations, automate repetitive work, and reclaim time.

Your SOC is a team — analysts, tools, and systems all working to stop threats — and Charlotte AI is the brain that makes the team smarter, faster, and more effective. And the best part: Every customer can now try Charlotte AI in their environment, at no additional cost.

Get Faster Answers

Ask questions and get actionable insight in seconds

Level Up Investigations

Unlock AI-powered insights across Falcon modules

Move at Machine Speed

Offload time-intensive work with mission-ready agents

CrowdStrike's Agentic Workforce, Powered by Charlotte AI

To further offload repetitive tasks, CrowdStrike delivers a workforce of turnkey AI agents. These agents extend Charlotte AI's intelligence and capabilities across the SOC, going beyond agentic triage, detection, and response to take on specialized tasks that traditionally require human time and attention. From cutting through CVE overload to streamlining data transformations, these agents tackle the SOC's highest-friction challenges with consistent, expert-informed automation.

PROVEN AGENTIC OUTCOMES

>98%
Decision accuracy for detection triage, benchmarked to Falcon Complete

+15
minutes saved per investigation

3x
faster mean time to detect (MTTD)

TYPE	AGENT
Threat Detection and Response	**DETECTION TRIAGE AGENT** — Analyzes detections, identifies true threats, and filters out noise
Threat Detection and Response	**RESPONSE AGENT** — Drives investigations, recommends next steps, takes authorized action
Exposure Management	**EXPOSURE PRIORITIZATION AGENT** — Summarizes vulnerabilities, validates exploitability, and maps impact to assets
Threat Hunting and Intelligence	**MALWARE ANALYSIS AGENT** — Defends against entire malware families, not just single samples
Threat Hunting and Intelligence	**HUNT AGENT** — Executes hunting queries and persistently scans the environment for emerging threats
Threat Hunting and Intelligence	**SEARCH ANALYSIS AGENT** — Interprets natural-language questions, explores relationships, and delivers clear insights
SIEM and SOAR	**CORRELATION RULE GENERATION AGENT** — Generates, validates, and optimizes rules to improve detection and classification
SIEM and SOAR	**DATA TRANSFORMATION AGENT** — Converts plain-language transformations into executable queries in CrowdStrike Falcon® Fusion SOAR
SIEM and SOAR	**WORKFLOW GENERATION AGENT** — Converts plain-language prompts into executable Falcon Fusion SOAR workflows
SIEM and SOAR	**DATA ONBOARDING AGENT** — Streamlines the entire data pipeline creation process through natural-language automation
SIEM and SOAR	**FOUNDRY APP CREATION AGENT** — Quickly creates custom security applications without the need for coding

Agent Spotlight: Malware Analysis

The Malware Analysis Agent enables security teams to analyze malware samples and defend against malware families at scale. It researches hashes, extracts configuration files, compares code similarities, and recommends responses in seconds. By coordinating sandboxes, YARA engines, and malware search, it delivers instant attribution, generates YARA detection rules automatically, and enables analysts to defend against entire malware families, not just single samples.

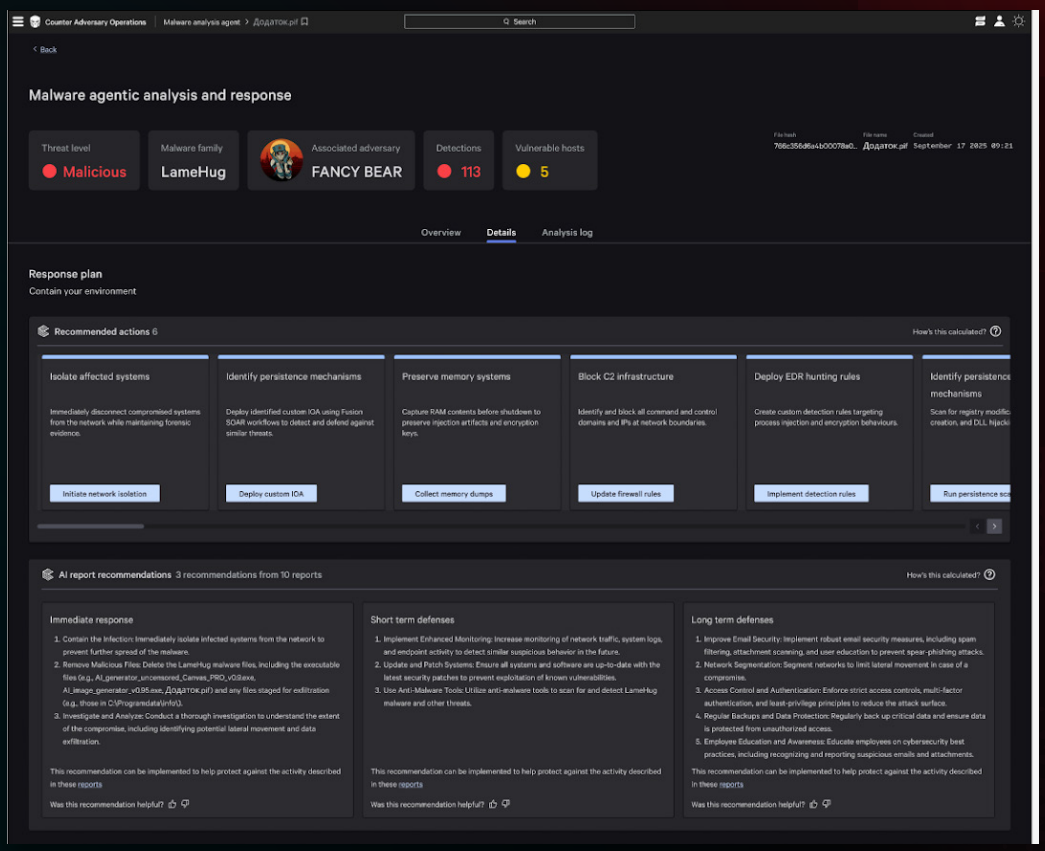

How It Works

1. A suspicious file is detected
2. The Malware Analysis Agent performs static file analysis and classifies the malware
3. The agent executes the file in a sandbox to perform dynamic file analysis
4. It generates YARA detection rules
5. It identifies related files across CrowdStrike's global telemetry
6. It generates recommendations to rapidly neutralize related threats

PHASE 3: BUILD CUSTOM AGENTS WITH CHARLOTTE AI AGENTWORKS

If CrowdStrike's out-of-the-box agents don't address a key workflow for your specific requirements, Charlotte AI AgentWorks, the industry's first no-code security agent development platform, empowers security teams to quickly build, test, and deploy custom AI agents — no code, no AI expertise required. Charlotte AI AgentWorks turns every security team into an agent builder: Simply describe your agent's mission, data sources, and authorized actions in natural language, and Charlotte AI handles the rest, building an agent that operates with elite reasoning at machine speed, with Falcon-grade governance — and always under your control.

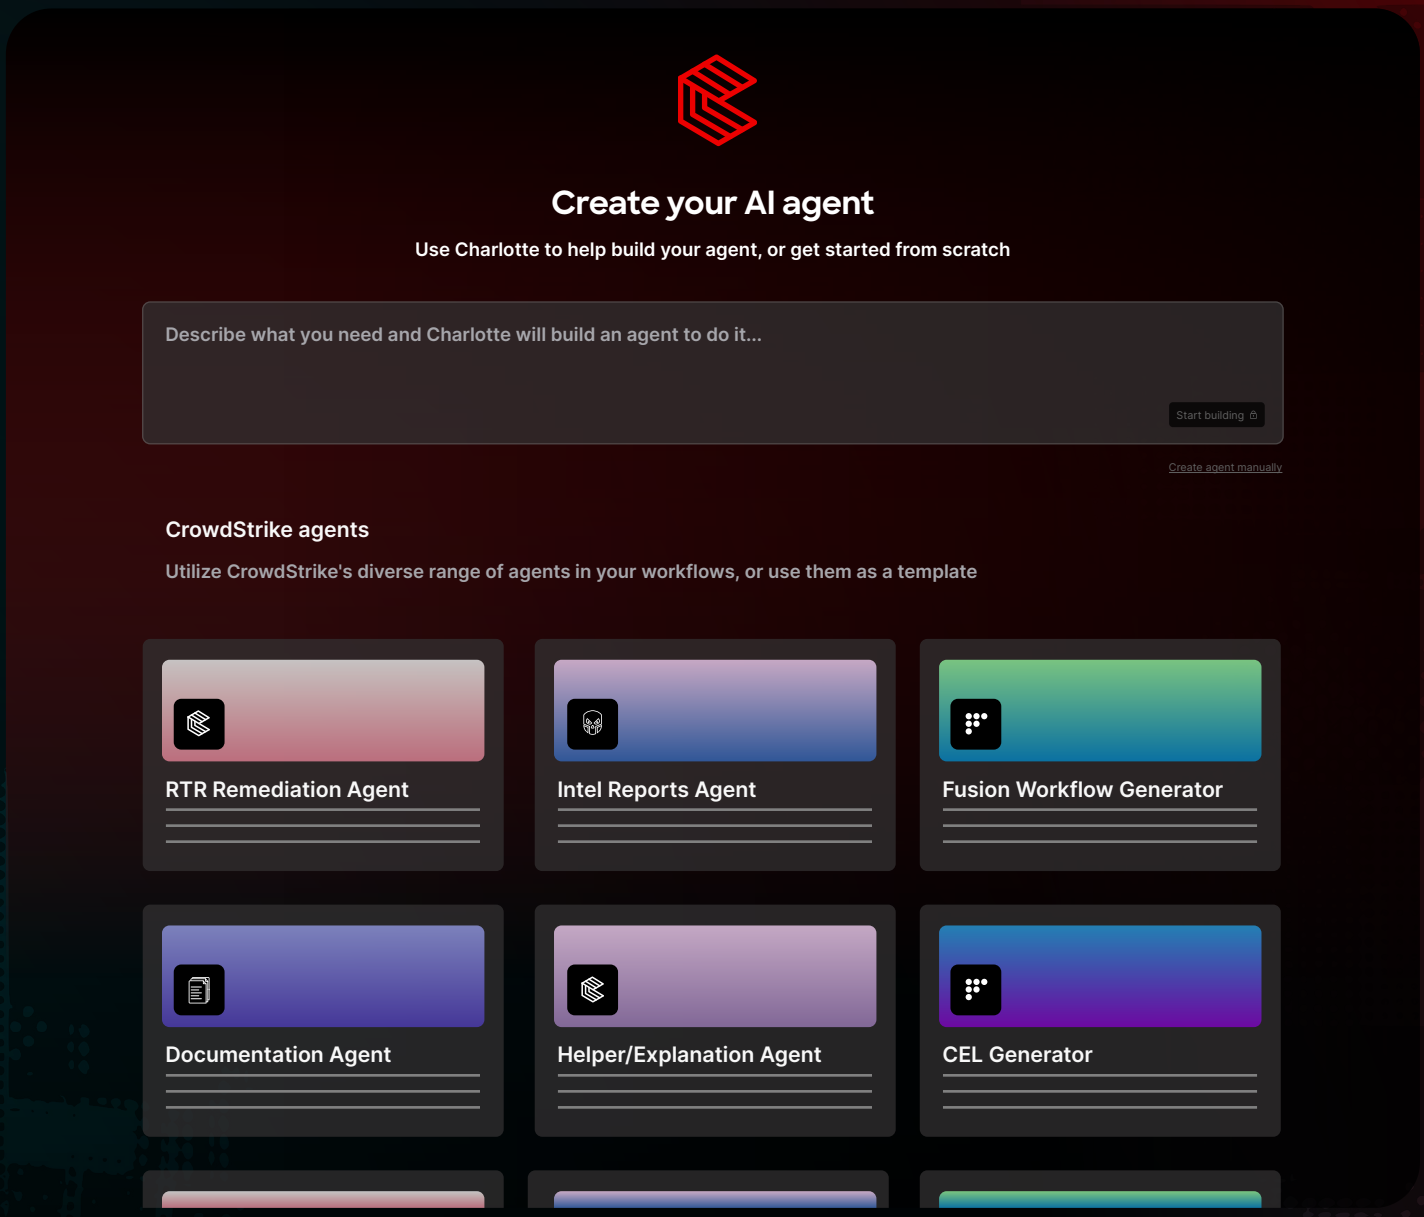

PHASE 4: ORCHESTRATE AN AGENTIC ECOSYSTEM WITH CHARLOTTE AGENTIC SOAR

Specialized agents are powerful on their own, accelerating investigations, generating workflows, and scaling expertise at machine speed. But their full potential is realized when they work together. Charlotte Agentic SOAR bridges the precision of deterministic automation with the adaptability of intelligent agents, enabling workflows that reason, adapt, and act with speed and context.

Where traditional SOAR stops at predefined playbooks, Charlotte Agentic SOAR introduces intelligence into orchestration itself. It blends the structure of rule-based automation — ideal for predictable, repeatable tasks — with the adaptive reasoning of AI agents that interpret context, weigh options, and decide dynamically. Together, they create a unified system that operates with the reliability of automation and the flexibility of human judgment.

CrowdStrike delivers this next evolution of SOAR through the combined power of Falcon Fusion SOAR, Charlotte AI, and Charlotte AI AgentWorks.

- **Falcon Fusion SOAR** provides the orchestration backbone — the connective tissue that unites agents, integrations, and analyst oversight into dynamic yet coordinated workflows. It gives analysts an intuitive interface to design and direct agentic workflows, defining the flow of data, reasoning, and action across the SOC with transparency, precision, and intent.
- **Charlotte AI** powers CrowdStrike's mission-ready agents, applying elite judgment and contextual reasoning to automate high-impact tasks, from malware analysis and exposure prioritization to detection triage and investigations.
- **Charlotte AI AgentWorks** empowers teams to build custom agents through a no-code interface, extending automation to specialized, domain-specific tasks.

Together, these capabilities form the foundation of Charlotte Agentic SOAR — a unified, intelligent ecosystem where human expertise, automation, and AI-driven decision-making converge. By connecting the reliability and control of deterministic workflows with the flexibility and reasoning of agentic workflows, Charlotte Agentic SOAR turns static automation into adaptive, collaborative defense, empowering the SOC to think, act, and adapt at machine speed — while keeping humans firmly in command.

Orchestrate Agents Across Your SOC: Charlotte Agentic SOAR

- Reimagine productivity with agent-to-agent orchestration
- Evolve from reactive response to continuous intelligence
- Maintain oversight with guardrails and governance

Agentic SOAR Brand Impersonation Workflow

This agentic workflow automates the detection and response to brand impersonation campaigns such as phishing and smishing. By combining AI-driven agents with workflow automation, it enables security teams to move from manual triage to intelligent, coordinated response — all within minutes.

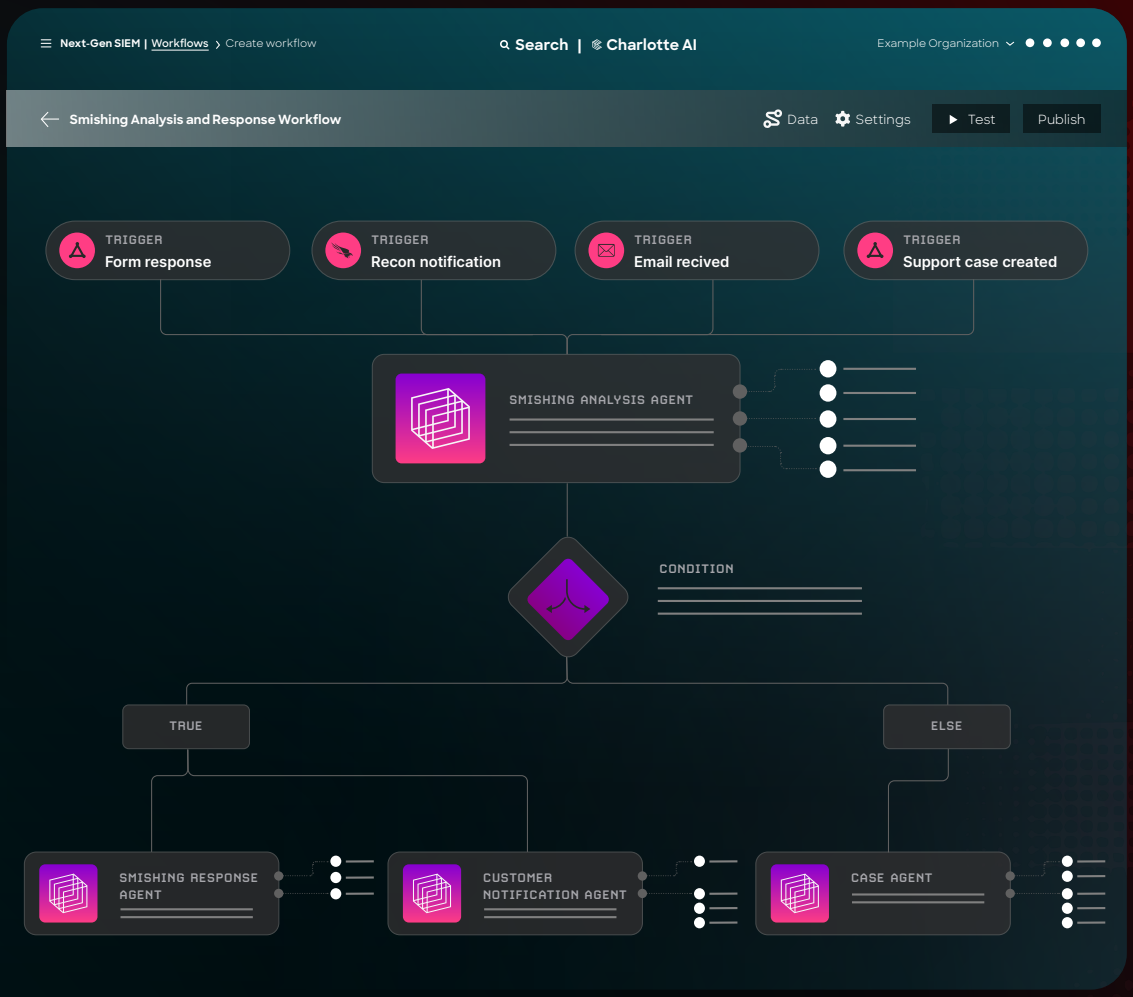

How It Works

1 2 3 4 5 6 7

Detection

Reports of fake texts and phishing emails are received across multiple channels

Workflow Creation

The Workflow Generation Agent builds an initial Falcon Fusion SOAR workflow from a plain-language prompt

Triage

The Smishing Analysis Agent reviews messages, correlates data, and determines if the threat is real

Case Management

If uncertain, a Case Agent automatically compiles findings and routes them to review channels

Automated Response

When confirmed, the Smishing Response Agent executes takedowns and removes malicious infrastructure

Customer Notification

The Customer Notification Agent alerts affected users

Continuous Improvement

All actions are logged, and the workflow refines itself over time to be faster and more accurate

But automation alone isn't enough — it must also be secure. When organizations attempt to link AI agents through unsupported or open-source connectors, they expose themselves to prompt injection, credential theft, and supply chain compromise — the very tactics adversaries now weaponize in the AI era. CrowdStrike's Model Context Protocol (MCP) eliminates that risk by enabling secure, governed, agent-to-agent communication across Charlotte AI, CrowdStrike-delivered agents, customer-built agents, and trusted third-party agents. Every interaction is encrypted, auditable, and governed at enterprise scale through MCP, creating a coordinated, AI-speed defense that teams can trust, with the visibility and control SOC leaders demand.

How CrowdStrike Ensures Responsible AI with Charlotte AI and CrowdStrike Agents

Charlotte AI and CrowdStrike agents operate with built-in guardrails that enable security teams to deploy AI with confidence and control:

Trusted Data: Every response is grounded in Falcon platform data or user-defined data, not generic or public sources.

Traceable Answers: Each decision includes inspectable response details so teams can understand Charlotte AI's reasoning and recreate its work at every step.

Role-Based Access: Charlotte AI and CrowdStrike agents are governed by the Falcon platform's access controls, ensuring data and actions are aligned to user permissions and organizational policy.

Bounded Autonomy: Users define what gets automated, when, and under what conditions. No surprises, no loss of control.

Expert MDR Validation: Charlotte AI and CrowdStrike agents get smarter with every new threat, trained and validated through a powerful feedback loop incorporating decisions by the world's top SOC analysts and incident responders.

Together, these safeguards enable responsible, accountable AI that teams can trust to act at machine speed without compromising oversight or integrity.

>> DID YOU KNOW?

Charlotte AI and CrowdStrike agents are trained on millions of real-world triage decisions from Falcon Complete Next-Gen MDR analysts.

This creates a continuous feedback loop where human expertise refines each agent's capabilities, and the agents amplify analyst speed, accuracy, and decision-making across the SOC.

CrowdStrike Falcon® Complete Next-Gen MDR is the industry's leading managed detection and response (MDR) service and the gold standard in detection and response, fusing elite human intelligence with advanced AI automation to deliver 24/7 protection that strengthens customer defenses while continuously enhancing the Falcon platform.

CrowdStrike's expert analysts deliver full-enterprise defense across endpoints, identities, and cloud workloads, leveraging Falcon Next-Gen SIEM for unified visibility and cross-domain threat correlation.

By unifying AI-powered detection, investigation, and expert-led response, CrowdStrike delivers the speed, scale, and expertise required to stop breaches in the AI era.

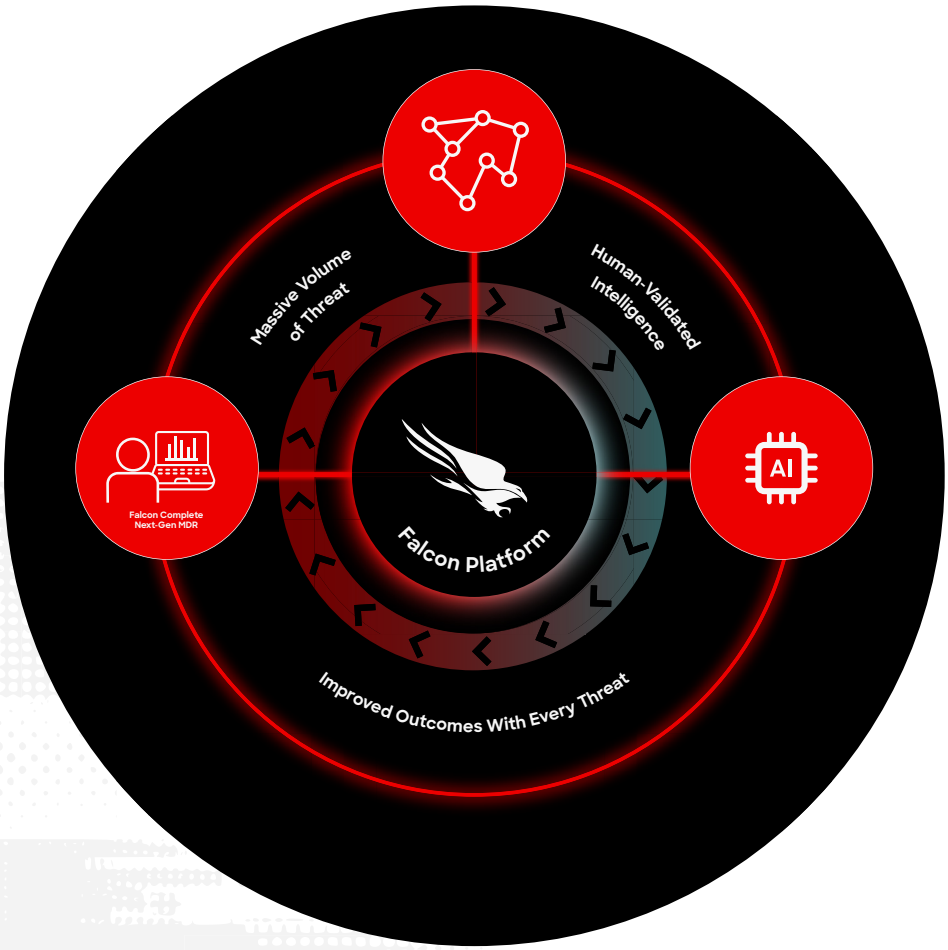

Need more guidance to operationalize AI?

Technology alone does not solve every security challenge, and teams often need guidance and expertise to fully operationalize AI in the SOC to increase speed, precision, and scale in the response lifecycle. CrowdStrike offers professional services to help you on your journey to the agentic SOC. These experts partner with your team to assess current capabilities, identify modernization opportunities, and design a roadmap for AI-driven efficiency. From finding opportunities to infuse AI into security operations, to optimizing processes and workflows, to integrating Charlotte AI and agentic automation, we help you securely accelerate detection, investigation, and response. Backed by the Falcon platform's unified data and intelligence, CrowdStrike enables your SOC to operate with precision, speed, and confidence — ready to stop breaches in the AI era.

Learn more

Next Steps

Try Falcon Next-Gen SIEM with FREE 10GB/day of third-party data ingestion

See how easy it is to bring in data from tools in your security ecosystem with Falcon Next-Gen SIEM. CrowdStrike Falcon® Insight XDR customers can ingest up to 10GB/day of third-party data to extend powerful detection, investigation, and response capabilities beyond the Falcon platform.

Try now

Get started with Charlotte AI today

Activate elite judgment at machine speed with Charlotte AI, the brain of the agentic SOC. Unlock natural language interactivity, AI-powered insights, and CrowdStrike's mission-ready agents — all designed to help you accelerate investigations, automate time-intensive tasks, and reclaim time.

Start here

Request information about CrowdStrike Services for SecOps readiness in the AI era

Contact us

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: www.crowdstrike.com

Follow us: Blog | X | LinkedIn | Facebook | Instagram | YouTube

Start a free trial today: www.crowdstrike.com/free-trial-guide