


Getting Started with Charlotte AI:

From Insight to Impact



Learn how to accelerate, automate, and scale AI-driven workflows with Charlotte AI. This guide shares best practices, real-world workflows, and tactical strategies to help your security team move fast, work efficiently, and reclaim analyst time.

Table of Contents

Accelerate Your SOC with Charlotte AI	4
Section 1: Getting Started with Charlotte AI	5
Section 2: Best Practices When Using Charlotte AI	7
Section 3: Maximizing Time Savings with AI-Powered Automation	9
Phase 1: Accelerate Workflow Creation with Prompts	10
Phase 2: Embed AI Reasoning in Falcon Fusion SOAR	11
Phase 3: Deploy Charlotte AI Agents	20
Section 4: Measuring Impact	22
Section 5: Scaling Your Agentic SOC	23

Accelerate Your SOC with Charlotte AI

Reclaim time. Redirect focus. Reimagine what's possible.

AI-accelerated adversaries have compressed attack timelines from days to seconds. In the modern threat landscape, speed is your most critical asset. CrowdStrike® Charlotte AI™ bridges the gap between human expertise and machine speed. Trained on millions of elite analyst decisions and powered by a unified, AI-native platform, Charlotte AI turns hours of manual investigation into fast, decisive action, helping security teams stay ahead.

Charlotte AI acts as a force multiplier, not an analyst replacement. By accelerating time to insight and handling repetitive tasks, Charlotte AI frees your analysts to focus on high-priority threats.

Getting Started: Qualifying CrowdStrike Falcon® platform customers now receive **free Charlotte AI credits that renew monthly**. This guide provides real-world workflows and best practices to immediately deploy those credits so that you can maximize impact, reclaim analyst time, and elevate your team to move at machine speed.

Section 1

Getting Started with Charlotte AI

What Is Charlotte AI?

Charlotte AI is the brain of your agentic SOC. It powers all of the agentic and generative AI capabilities across the Falcon platform, from natural language interactivity to mission-ready agents that offload time-intensive work.

Key Benefits of Using Charlotte AI

- **Master the Falcon Platform:** Charlotte AI spans your Falcon platform modules so you can operate CrowdStrike's unified platform with confidence from Day One. Ask direct questions and quickly surface cross-domain context, threat intelligence, and answers from documentation — no complex queries required.
- **Level Up Every Analyst:** Standardize how your team works. Charlotte AI guides investigations, recommends next steps, and explains what's happening in plain language, turning ad hoc analysis into consistent, documented work.
- **Automate Tasks with Specialized Agents:** Offload repetitive, time-intensive work — such as triage, data normalization, enrichment, and exposure prioritization — to out-of-the-box agents. This allows analysts to focus on high-impact decisions and actions, shrinking time to detect and respond.
- **Governance by Design:** Maintain complete analyst control with role-based permissions, transparent audit logs, and bounded autonomy — you decide what gets automated and when.

What Is a Charlotte AI Credit?

Charlotte AI credits act as the currency for AI operations. Credits are aggregated into a universal pool, allowing your entire team to draw from a shared balance when using Charlotte AI. Falcon administrators assign and manage Charlotte AI roles, deciding which users have permission to view or interact with Charlotte AI to ensure credits are used strategically.

Key Features and Capabilities

Use your credits to explore these features.

Charlotte AI Feature	Description
Charlotte AI Conversations	Query Falcon platform data in plain language (investigate detections, analyze vulnerabilities, summarize threat intelligence, build advanced search scripts, etc.).
Charlotte AI Promptbooks	Create, save, and share reusable prompt libraries across teams.
Agentic Workflows	Activate AI reasoning in security orchestration, automation, and response (SOAR) workflows to unlock adaptable automation.
Command-Line Explanations	Generate step-by-step explanations of complex command lines and scripts.
Investigate with Charlotte AI	Generate an in-depth incident summary and next-step suggestions to understand the scope of an attack.
AI-Generated Parsers	Build and edit data parsers based on sample logs provided.
Live Asset Query Generation	Generate operating system (OS) queries for live asset query search.
Cloud Asset Search	Search cloud assets in plain language.
CrowdStrike Agents	
Detection Triage Agent	Automatically classify and prioritize new detections to filter noise.
Agentic Response Agent	Automate analysis by generating guiding questions and answers.
Malware Analysis Agent	Analyze malware by correlating files, hashes, code similarities, sandbox results, and threat intelligence. This agent also generates YARA detection rules and next-step recommendations.
Hunt Agent	Auto-summarize results from threat hunting queries, surfacing threats by adversary, malware family, or MITRE ATT&CK® technique.
Workflow Generation Agent	Turn user prompts into CrowdStrike Falcon® Fusion SOAR workflows.
Exposure Prioritization Agent	Explain vulnerabilities in plain language, validate their real-world exploitability, and analyze potential business impact.

Section 2

Best Practices When Using Charlotte AI

Charlotte AI gives you a fast, measurable way to operationalize AI across your security operations. Start with these best practices to maximize the value of your credits and scale your operations.

1. Automate Reporting Cycles

Automate the manual analysis and reporting your team already delivers — recurring summaries, executive updates, or in-depth retrospectives. Use Charlotte AI to collect data, analyze trends, and tailor outputs for specific audiences. Measure the hours saved per week to demonstrate immediate ROI.

2. Delegate Routine Tasks to Agents

Identify the high-volume, low-complexity tasks clogging your queue. Deploy specialized agents to handle triage, enrichment, and exposure analysis automatically, or ask Charlotte AI to generate Falcon Fusion SOAR workflows to implement complex automation across the platform.

Pro Tip:

Automate one recurring task this week. Measure the time saved, share the results, and use that success to build momentum toward broader AI adoption.

3. Standardize Analysis with Promptbooks

Don't let effective prompts live in one analyst's head. Build shareable promptbooks for common scenarios, such as exposure analysis, vulnerability analysis, or threat intelligence summarization. Share promptbooks across your team to give analysts ready-to-run starting points, reducing guesswork and cycles of trial and error.

4. Manage Governance and Access

Ensure responsible AI usage and manage credit reserves by defining who can see and use Charlotte AI across your team. For more details on Charlotte AI roles, see the [Charlotte AI documentation](#).

5. Track Usage and Prove ROI

Track adoption across your team with the Charlotte AI credit usage dashboard. Use these insights to quantify productivity gains, analyze analyst engagement, and build the business case for purchasing additional AI credits.

Section 3

Maximizing Time Savings with AI-Powered Automation

As organizations explore how to bring AI into their security operations, the most pressing question they ask is often “What can we implement today?”

Charlotte AI provides a practical, accessible path to operationalizing AI. Though they are not exhaustive, the following use cases represent high-value ways Charlotte AI helps organizations scale their impact around the clock. Each step enhances — not replaces — analyst judgment, giving your SOC a controlled, progressive path to machine-speed operations.

No hype. No coding. No black box. Just practical, AI-powered automation that compounds value over time.

Incremental Steps	How You Do It	Outcomes
Phase 1: Accelerate Workflow Creation with Prompts	Use natural language prompts to generate Falcon Fusion SOAR workflows	• Build or modify workflows in minutes or seconds • Lower the technical barrier to building automation • Customize existing playbooks or build new ones without complexity
Phase 2: Embed AI Reasoning in Falcon Fusion SOAR	Add the “Charlotte AI-LLM Completion” action to Falcon Fusion SOAR workflows	• Transform rule-based scripts into adaptive workflows • Generate structured analysis and reporting on demand • Produce persona-tailored summaries, insights, and content
Phase 3: Deploy Charlotte AI Agents	Orchestrate Charlotte AI’s expert agents to offload time-intensive tasks	• Accelerate mean time to respond (MTTR) with automated triage and analysis • Eliminate alert fatigue by filtering noise • Deliver accurate, consistent Tier-1 response at scale

Phase 2:

Embed AI Reasoning in Falcon Fusion SOAR

Reporting is one of the most time-consuming tasks in the SOC — and can be a critical drain on team resources. Security teams often lose hours each week manually collecting data, analyzing it, and summarizing findings.

Charlotte AI enables teams to turn static playbooks into intelligent, adaptive automation. By adding the “**Charlotte AI-LLM Completion**” action to Falcon Fusion SOAR workflows, security teams can embed large language models (LLMs) directly into their playbooks.

This capability moves automation beyond “if/then” logic, enabling Falcon Fusion SOAR workflows to:

- **Analyze Context:** Reason over structured or unstructured data
- **Adapt on the Fly:** Handle unexpected edge cases so automation doesn't break
- **Generate Custom Outputs:** Create tailored outputs for specific audiences (e.g., language translation, executive summaries, formatted notes)

How to Build This: For step-by-step guidance on configuring this action, please refer to [Charlotte AI documentation](#).

Here are four real examples of high-impact workflows to try today.




Image shows a Charlotte AI agentic workflow in Falcon Fusion SOAR. Using the “Charlotte AI-LLM Completion” Action, security teams can perform inference within structured workflows, specifying instructions for what models should do, what data to leverage, and how to structure output.

Workflow Example 1: Automating Monthly False Positive Reviews

Pain Points:

- Reviewing false positives consumes hours each month, diverting analysts from higher-value investigations.
- Inconsistent reviews lead to missed patterns or redundant reclassification.

Agentic Workflow Solution:

Build a workflow that reviews false positive alerts, identifies reduction opportunities, validates that no true threats were missed, and delivers a monthly summary with actionable recommendations.

Expected Outcomes:

- Automates monthly reviews to tune detection rules and surface real threats
- Enables fast, data-driven retrospectives that help improve precision and reduce false positives
- Strengthens SOC consistency and auditability over time

Workflow Guidance:

Required Modules	CrowdStrike Falcon® Insight XDR, CrowdStrike Falcon® Next-Gen SIEM
Data Inputs	Pulls security detection data from false positives, including: • When and where the alert happened • Which program triggered the alert • Which commands were run • Which user and computer were involved • Network connections that were made

Prompt Instructions
(Overview)

Instruct the LLM to act as a highly analytical security analyst and follow these instructions:

- Review false positives to spot true threats
- Identify recurring patterns and rule gaps
- Recommend tuning or exclusions to reduce noise
- Summarize results and automatically email a monthly HTML report

Workflow Steps

STEP 1

Kick Off Workflow

Schedule to initiate on the first day of every month at 7:00 a.m. ET

STEP 2

Gather Data

Pull prior-month false positive detections

STEP 3

Conduct Readiness Check

Confirm all inputs downloaded

STEP 4

Perform Analysis

Generate a detailed report explaining patterns identified and recommended exclusions or rule refinements

STEP 5

Send the Report

Structure output in HTML and email the security team

Workflow Example 2: Automating Exposure Reporting

Pain Points:

- Translating complex vulnerability data into clear, business-relevant insights can be time-consuming and inconsistent.
- Prioritization is often driven by CVSS scores alone, without considering exploitability, asset criticality, or exposure.

Agentic Workflow Solution:

Build an agentic workflow to automate executive vulnerability briefings, translating technical findings into prioritized remediation recommendations for executives and highlighting risk, exposure, and business impact.

Expected Outcomes:

- Delivers a standardized report that translates risk data into business impact
- Reduces manual reporting while improving clarity, alignment, and stakeholder buy-in for remediation efforts

Workflow Guidance:


Required Module	CrowdStrike Falcon® Exposure Management
Data Inputs	Pull vulnerability data from Falcon Exposure Management, including: Vulnerability details: • CVE IDs • CrowdStrike's ExPRT.AI rating (proprietary severity score) • Exploitation status Environment context: • How many systems have each vulnerability • Host information (criticality, internet exposure, location) Key filters applied: • Only "open" vulnerabilities (not yet fixed) • Only "High" or "Critical" ExPRT.AI ratings • Only "Actively used (Critical)" exploit status

Prompt Instructions
(Overview)

Instruct the LLM to:

- **Translate technical data:** Turn CVEs and CVSS scores into business impact language
- **Categorize risks:** Group vulnerabilities by type (e.g., “Remote Code Execution Risks” or “Authentication Bypass Issues”)
- **Assess significance:** Provide a risk assessment, avoiding alarmism
- **Recommend next steps:** Suggest clear next steps, prioritized by impact
- **Format:** Generate a polished executive report

Workflow Steps



Workflow Example 3:

Automating Data Egress Pattern Analysis

Pain Points:

- High event volumes make it hard to spot suspicious or policy-violating transfers.
- Investigations often start only after exfiltration or misuse occurs.

Agentic Workflow Solution:

Build a workflow that ingests CrowdStrike Falcon® Data Protection logs to analyze data egress across users, destinations, and content types. Generate a weekly report that surfaces risky behaviors, highlights compliance issues, and provides actionable recommendations for security leadership.

Expected Outcomes:

- Detects potential data exfiltration and misuse earlier
- Automates policy tuning and risk reduction through continuous analysis
- Delivers weekly summaries of risk patterns, violations, and prevention actions

Workflow Guidance:

Required Module	Falcon Data Protection
Data Inputs	Pull event data from Falcon Data Protection, including: • Files transferred (names, types, sizes) • Where data is going (cloud services, USB drives, email, web) • Who is transferring data (usernames, departments) • When transfers happen (dates, times, frequency) • Types of sensitive data (source code, financial records, customer data, credentials) • Policy violations or compliance flags • Risk indicators

Prompt Instructions
(Overview)

Instruct the LLM to:

- **Analyze patterns:** Identify trends and anomalies in data transfers
- **Assess risk:** Determine which transfers pose the highest security or compliance risk
- **Identify threats:** Flag behaviors suggesting data theft, insider threats, or policy violations
- **Categorize findings:** Group similar events (e.g., “multiple users uploading source code to GitHub”)
- **Recommend actions:** Suggest mitigations and policy adjustments
- **Report results:** Generate and email a summary report with risk ratings and key insights

Workflow Steps

STEP 1

Kick Off Workflow

Initiate every
Monday at 6:00
a.m. ET

STEP 2

Collect Falcon Data
Protection Events

Pull the last
seven days of
data egress
events

STEP 3

Analyze the Data

Analyze the data
and generate a
report with:

- Executive
summary and key
findings
- Notable events
and key patterns
- Overall risk
assessment and
recommendations

STEP 4

Send the Report

Email the plain-
text report to the
security team

Workflow Example 4: Assisting Phishing Detection and Response

Pain Points:

- Email remains a top attack vector, but fragmented logs and disconnected playbooks slow down investigation and response.
- Manual triage consumes hours and increases the risk of missed or delayed detections.
- Analysts struggle to correlate behavior, attachments, and URLs across disparate tools.

Agentic Workflow Solution:

Build a workflow to analyze user-specific email activity. Leverage third-party email telemetry ingested into Falcon Next-Gen SIEM to identify suspicious behavior, automate response actions, and recommend next steps aligned to organizational policy.

Expected Outcomes:

- Cuts phishing triage time, reduces alert fatigue, and improves prioritization
- Delivers consistent, policy-aligned verdicts and recommendations

Workflow Guidance:

Required Module	Falcon Next-Gen SIEM
Data Inputs	Query email security event logs from Falcon Next-Gen SIEM, including: • Vendor source (e.g., Mimecast, Abnormal Security) • Sender and recipient addresses • Email subject lines • Embedded URLs • Email attachments • Timestamps

Prompt Instructions
(Overview)

Instruct the LLM to:

- **Analyze patterns:** Identify trends and anomalies across sender behavior, URL use, and attachments
- **Assess risk:** Determine which exchanges represent the highest security or compliance risk
- **Identify threats:** Flag potential spear-phishing or credential harvesting activity
- **Categorize findings:** Assign verdicts (malicious, suspicious, clean) with confidence levels
- **Recommend actions:** Suggest mitigations and policy improvements
- **Take action:** Automatically contain the user or escalate cases based on model verdict and confidence

Workflow Steps

STEP 1

Kick Off Workflow

Run on demand with an individual user's email

STEP 2

Collect Data

Retrieve email telemetry and metadata filtered for the target user

STEP 3

Analyze the Data

Generate a structured report with search results, verdict, confidence level, and explanation of reasoning

STEP 4

Automated Actions

If "Malicious" and "High Confidence," add sender to the organization's blocklist; if "Suspicious," send a contextual alert to the on-call security analyst

Phase 3:

Deploy Charlotte AI Agents

Effective triage and investigation are the backbone of any SOC. Yet even in mature SOCs, analysts vary in how they assess severity and record findings. Charlotte AI offloads this effort with specialized agents for detection triage and analysis, bringing expert-level consistency, speed, and precision to every detection.


Trained on years of the CrowdStrike Falcon® Complete Next-Gen MDR team's real-world decisions, these agents allow you to scale high-quality investigations across the SOC without adding headcount.

Detection Triage Agent

The Charlotte AI Detection Triage Agent automatically evaluates first-party Falcon platform detections (endpoint, identity, or cloud detections). It reasons like a Falcon Complete Next-Gen MDR analyst, reviewing event context to classify detections. It also produces a recommendation for every detection, rapidly filtering noise and prioritizing the queue.

What the Detection Triage Agent provides:

- **Detection verdict:** A judgment on whether the detection is a true or false positive
- **Confidence level:** An assessment of how strongly the agent supports its verdict
- **Priority score:** A ranked assessment that helps analysts prioritize high-severity detections first
- **Recommended next steps:** Suggestion to close out or escalate the detection



Response Agent


The Charlotte AI Response Agent takes investigations further by generating and executing structured rounds of follow-up questions and answers.

What the Response Agent provides:

- **Investigative questions and answers:** The agent drives investigations by generating targeted questions and then answers them by leveraging a user's Falcon platform modules
- **Explanation of reasoning and analysis:** It generates a detailed explanation of its reasoning and explains how it obtained answers to guiding questions
- **Flexible depth:** Users control the depth of analysis (spending 1, 3, or 6 credits per round), allowing them to perform deeper analysis when severity or context demands it

Did you know?

Analysts can collaborate with Charlotte AI's Response Agent in real time. Within the Response Agent's investigation canvas, analysts can add questions, provide answers, and prioritize which questions get answered first. This enables analysts to provide Charlotte AI with critical context unique to their organization.



Pro Tip: Use Falcon Fusion SOAR to automate response actions triggered by the outputs of the Charlotte AI Detection Triage Agent and Response Agent. Pairing agentic investigation with automated response can reduce MTTR, minimize manual handoffs, and bring expert reasoning and consistency to your investigations.

Measurable Performance:

> 98% agent decision accuracy, trained on the decisions of Falcon Complete Next-Gen MDR experts¹

5 min. — Estimated time savings per triaged detection²

10 min. — Estimated time savings per investigation with Charlotte AI's Agentic Response agent³

¹ Accuracy rating is a measure of Charlotte AI triage decisions that match the expert decisions from the CrowdStrike Falcon Complete Next-Gen MDR team.

² Time savings represents the amount of time an analyst would have spent triaging detections but can now use that time for other skilled work while Charlotte triages the detections. Individual results may vary based on factors such as total alert volume.


³ The time savings of more than 10 minutes per investigation is an estimate based on Agentic Response's ability to automate tasks that would otherwise require more than 10 minutes of manual effort by a human analyst. This should not be interpreted as a guarantee that this will lead to a 10-minute reduction in the total investigation time or mean time to respond (MTTR).

Section 4

Measuring Impact

The Charlotte AI Credit Usage dashboard provides a centralized view of how Charlotte AI is transforming your security operations, helping you measure, manage, and maximize your ROI.

With the dashboard, you can see credits spent and remaining, uncover adoption patterns, and identify top users, giving you a snapshot of how much analyst capacity Charlotte AI has returned to your team.



Charlotte AI credit usage dashboard

Section 5

Scaling Your Agentic SOC

You've taken the first step toward operationalizing AI by utilizing your complimentary Charlotte AI credits. Now it's time to scale, turning early wins into enterprise-wide automation.

Next Steps to Scale Adoption:

- **Share High-impact Wins:** Promote consistency and collaboration by sharing effective workflows, promptbooks, and reports across teams. Build on proven successes to accelerate adoption.
- **Uncover New Use Cases:** Activate Charlotte AI to streamline more workflows — connecting detection triage and analysis to automated response actions, implementing new reporting initiatives, or identifying additional modules to unlock more mission-ready agents.
- **Expand Your Capacity:** Contact your CrowdStrike account team to increase your monthly Charlotte AI credits and explore advanced features from Charlotte AI.

Learn more about Charlotte AI:

<https://www.crowdstrike.com/en-us/platform/charlotte-ai/>

About CrowdStrike

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

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

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

CrowdStrike: We stop breaches.

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

Follow us: [Blog](#) | [X](#) | [LinkedIn](#) | [Facebook](#) | [Instagram](#)

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