Tuskira

Meet Iris

From Alert Queue to Verdict in Seconds

Iris is Tuskira's AI SOC agent workforce, powered by the Security Context Graph that already correlates identity, exposure, controls, detections, and attack paths across your stack. Other AI SOCs determine whether an alert is malicious. Iris determines whether an alert represents meaningful organizational risk. L1 evaluates each alert against exposure, identity, control, and attack-path context on the Graph and renders a risk-based verdict in seconds. L2 maps the blast radius and validates the kill chain against Kairo's path model. The Response Agent orchestrates containment through the controls you already own (EDR, IdP, firewall, WAF, SIEM), with analyst approval where policy requires.

> THE QUESTION EVERY SOC ANALYST IS BURIED UNDER
> **What's the verdict?**

WHY NOW

The SOC Queue Is Structurally Broken

Frontier AI has compressed the path from vulnerability discovery to exploitation from weeks to minutes, and attack chains now span identity, cloud, endpoint, and network in a single coordinated progression. Meanwhile, 11,000 alerts a day is routine in the SOC, most low-confidence and stripped of operational context. Human-speed triage cannot keep pace. The SOC has to shift from queue processing to verdict supervision on a shared context layer.

DAILY ANALYST LOAD

11,000+

Routine alert volume per SOC analyst, most without exploitability evidence

INITIAL ACCESS → OBJECTIVE

Days → Minutes

Compression for AI-driven attack chains

Market Alternative	Where They Stop	How Iris Is Different
In-house SOC / queue-based triage	Analysts drown in alert volume; investigations are single-threaded across identity, endpoint, cloud, and network.	**Renders a risk-based verdict on every alert in seconds by evaluating exposure, identity, control, and attack-path context on the Security Context Graph.**
Managed Detection & Response (MDR / MSSP)	Humans triage at SOC speed; quality varies with analyst skill, and detection logic is gated by the provider's playbook.	**Produces risk-based verdicts in seconds on your own Security Context Graph, with full reasoning chains. Every escalation arrives pre-investigated with blast radius and recommended response attached.**
SOAR platforms	Automate response after humans triage; the bottleneck moves from response to triage and stays there.	**Triages, investigates, and proposes response autonomously on a shared context layer. Every action is logged, reversible, and bounded by customer-defined policy.**
Single-vendor AI SOC	Handle alerts from one vendor's telemetry well; struggle to correlate across heterogeneous stacks.	**Reasons on the unified Security Context Graph that already powers Kairo, Quell, and FedSOC. Kairo finds the breach paths, Quell finds the emerging threats, and Iris uses that context to decide which alerts represent meaningful risk.**

How Iris Reasons on the Security Context Graph

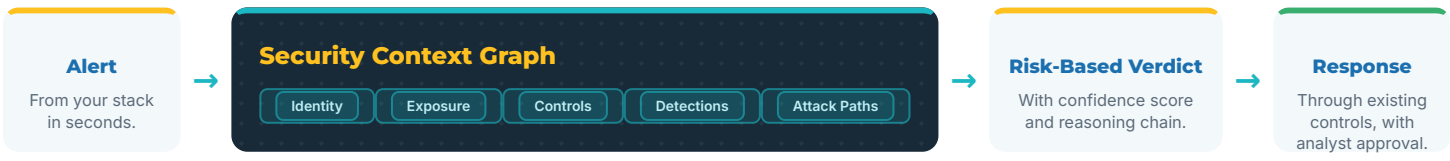

Cases Iris triages, investigates, and contains

Case Type	Detection Chain	What Iris Determines	How Iris Responds
Malicious sign-in from a new geography	Azure AD sign-in fires; identity baseline, threat intel, and asset access correlated	IP is a known TOR exit node, credentials match a recent breach, the account has elevated cloud permissions, every baseline is violated	**Revokes the session, disables the account, blocks the IP range, files the incident ticket. Verdict in 12 seconds, contained in under four minutes.**
Phishing to endpoint compromise	Email artifact, endpoint behavior, and identity event stitched into one case	User clicked, macro executed, C2 channel established, credentials harvested; kill chain validated against Kairo's path model	**Isolates the endpoint, revokes the user's session, blocks the C2 domain, and surfaces the full kill chain as a single auditable case.**
Ransomware staging	Cobalt Strike beacon → LSASS dump → SMB pivot → shadow-copy deletion attempt	Staging sequence is in progress, not an isolated behavior; kill chain is reachable across endpoint, identity, and network	**Contains the staging endpoint, blocks lateral movement at the network chokepoint, and escalates with the full reasoning chain attached.**
Insider privilege misuse	Privileged identity, unusual access to crown-jewel data, off-hours pattern, behavior vs. baseline	Access is in scope of the identity, but the pattern deviates from baseline and touches sensitive resources	**Opens a case with full reasoning chain for human review. High-impact actions on insiders stay human-led by policy.**

L1 Triage Agent
Verdicts on every alert in seconds, with confidence scores and reasoning chains.

L2 Investigation Agent
Maps blast radius and validates the kill chain against Kairo's path model.

Response Agent
Orchestrates containment through existing controls, with analyst approval where policy requires.

Investigation Audit Trail
Every AI action, decision, and analyst approval recorded end-to-end with reviewable evidence and full reasoning chain.

Continuous learning
Every verdict, override, and resolution updates the platform's priors over time.

From Alert to Action: Traditional SOC vs. Iris

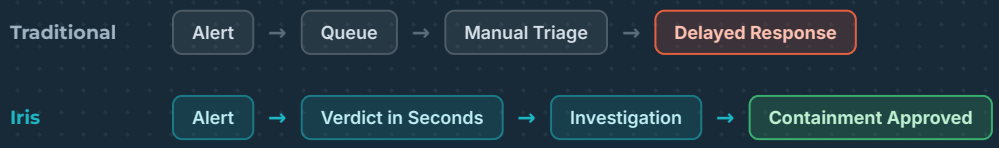

98%
Noise reduction (verdicts replace queues)

4 min
Attack to approved containment

Seconds
Verdict on every alert, with confidence score

150+
Tool integrations powering the Security Context Graph

See Iris in Action

See how Iris triages every alert, investigates blast radius, and orchestrates containment, with a full reasoning chain on every verdict and human approval where policy requires.

Request an Iris demo: tuskira.ai/demo | Lorin@tuskira.ai

Tuskira™
Full Stack Agentic SecOps