Tuskira™

White Paper

Close the Gap Between Exposure and Detection

How Tuskira unifies exposure, detection, investigation, and response through a shared Security Context Graph

Contents

Executive Summary

Security operations are built on the outdated assumption that they need to centralize all telemetry, write detection rules against the aggregated data, and rely on analysts to triage what comes out. That model is breaking.

Detection is delayed by data pipelines. Alert volume has exceeded human capacity. Detection logic decays faster than teams can maintain it. Meanwhile, attackers no longer operate on an event-by-event basis. They move across identity, cloud, endpoint, and network as coordinated attack paths, increasingly accelerated by AI.

Security teams can see what's exposed, and they can see what's alerting. What they still struggle to answer is whether their detections cover what's actually breachable. The result is a structural mismatch: security teams process alerts in queues, while attacks execute across systems.

Tuskira introduces a different architecture. Instead of centralizing data, Tuskira unifies security context across the existing stack and performs detection, correlation, and investigation where the data already lives. AI-driven agents operate across this shared context to produce verdicts, not alerts, and to continuously adapt detection and response based on real attack paths.

At the same time, the platform models the environment as a live digital twin, identifying which vulnerabilities, identities, exposures, and control gaps are reachable, and which actions can break those paths before they are exploited.

This convergence of:

- **Detection: Federated SOC**
- **Exposure and breach-path analysis: CTEM**
- **AI-driven investigation and response (under human oversight)**
- **Unified context: the Security Mesh and Context Graph**

... creates a new operating model for the SOC. From reactive alert triage to continuous, preemptive security operations, the implications go beyond a faster SOC. Security operations are converging into a single system, where shared context serves as the control plane. This paper outlines how that architecture works, how the components interact, and what it means for security teams operating in modern, distributed environments.

The Failure of the Centralized SOC Model

A modern enterprise runs CrowdStrike for endpoint, Okta or Entra for identity, AWS, Azure, or GCP for cloud, Proofpoint for email, an NDR for east-west traffic, a vulnerability scanner, a CNAPP, a WAF, and a SIEM to make sense of it all. Each tool is a strong specialist in its domain. Collectively, they generate more information than any human team can process, and the SIEM that promised to unify it has become the queue every analyst lives inside.

The structural failure is not in any one tool. It's in the centralized model itself. Logs traverse a transport chain (endpoint to collector to pipeline to SIEM to rule match) that introduces detection latency, which adversaries exploit. Detection-engineering teams spend the majority of their time maintaining brittle rules across heterogeneous schemas rather than building new coverage. Each new tool added to the stack increases ingestion, storage, rule count, analyst load, and investigation time, without proportional improvement in outcomes. The industry has spent a decade optimizing the queue, yet teams still lose hours or days moving between tools before they know what matters.

Modern defense depends on detections grounded in real exposure, not isolated events.

The Real Problem: Attacks Are Paths, Not Events

The centralized model fails for a deeper reason than cost or latency. The SOC operates as a queue. Attackers operate through paths.

Consider a real attack. Someone compromises a credential, such as phishing, a third-party breach, or an exposed token. That's an identity event. The attacker authenticates successfully, and the EDR sees unusual behavior, baseline reconnaissance, and lateral movement. That's an endpoint event. The attacker escalates privileges and moves to cloud resources. That's a cloud event. Finally, data is exfiltrated. That's a network event.

A coordinated attack crosses the territories of four tools as a single progression. The centralized SOC sees it as four separate alerts in four separate queues. Each tool sees one chapter of the story and concludes that nothing interesting is happening.

> **Attackers don't operate through queues. They operate through paths. No single queue sees the whole path.**

AI-driven adversaries are further compressing those paths. Recent threat intelligence shows automated reconnaissance shrinking target mapping to hours per target, and AI-assisted campaigns chaining lower-severity issues into working exploits at machine speed. The disclosure-to-weaponization window for AI-discovered exploits is now measured in minutes, not weeks. Faster queues will not close that gap.

The problem is no longer collecting more data. It's connecting exposure, detections, identities, controls, and attack paths into operational context.

The Shift: Unify Context, Distribute Detection

To keep pace with modern threats, the SOC must converge four previously separate functions:

- **Detection across distributed systems**
- **Investigation across multiple domains**
- **Exposure analysis and attack-path modeling**
- **Response through existing controls**

Convergence requires a shared, continuously updated understanding of the environment rather than a centralized log store. Detection has to move to the source, correlation has to happen across tools in real time, without data movement, and agents have to reason over context.

This is the deeper architectural shift the industry is already seeing: Detection, exposure analysis, investigation, and response are no longer separate workflows. They're collapsing into a single operating layer. The traditional boundary between SOC and CTEM disappears because both depend on the same underlying truth: what's reachable, what's detectable, and what can be stopped. Shared context becomes the control plane that connects them.

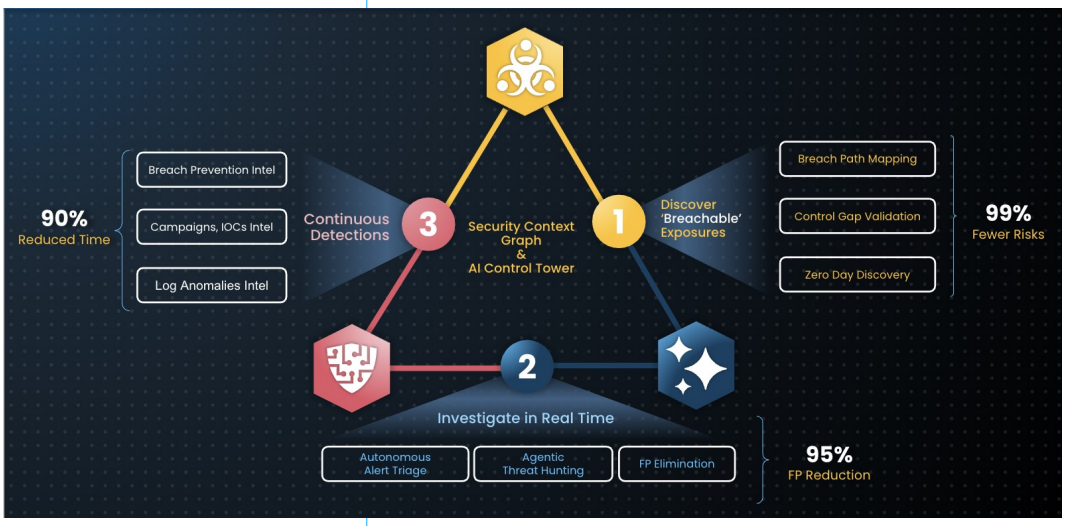

Figure 1 — Exposure, detection, investigation, and response converge through the Security Mesh and Security Context Graph.

The Tuskira Platform

Tuskira is the platform that operationalizes that convergence. It organizes itself into layered capabilities running on top of a shared digital twin. "Ask Tuskira" is the conversational access layer that any user (analyst, engineer, CISO, compliance lead) can address in plain English. Federated SOC sits below it, federating detection across the existing tool stack and running the AI SOC agent workforce. CTEM is the preemptive exposure portfolio, currently anchored by vulnerability management, Zero Day Response, and Kairo (Breach-Path Modeling). Beneath everything is the Security Mesh and Security Context Graph, a live digital twin of identities, assets, vulnerabilities, controls, behaviors, and the blast radius between them, connected to over 150+ native tool integrations.

The layering is the point. A finding from the Breach-Path Modeling (Kairo) updates the Security Context Graph; the next Federated SOC correlation is sharper as a result; an L2 agent investigation pulls Kairo's path model into its evidence; every step is accessible through Ask Tuskira and auditable end-to-end. The platform's intelligence increases as the environment changes; the marginal cost of adding a new tool decreases.

Unify the context, not the logs.

Federated SOC (FedSOC): Real-Time Detection and Investigation

Federated SOC is the operations layer, the surface that responds to what's happening right now. It sits across the existing stack rather than replacing it, connecting to CrowdStrike, AWS CloudTrail, ExtraHop NDR, Proofpoint, legacy SIEM, and dozens of other sources through native APIs. As each new source is onboarded, Federated SOC samples incoming events, infers the schema, maps field relationships, and builds a common dictionary and vocabulary across all sources. Weeks of connector and normalization engineering happen automatically.

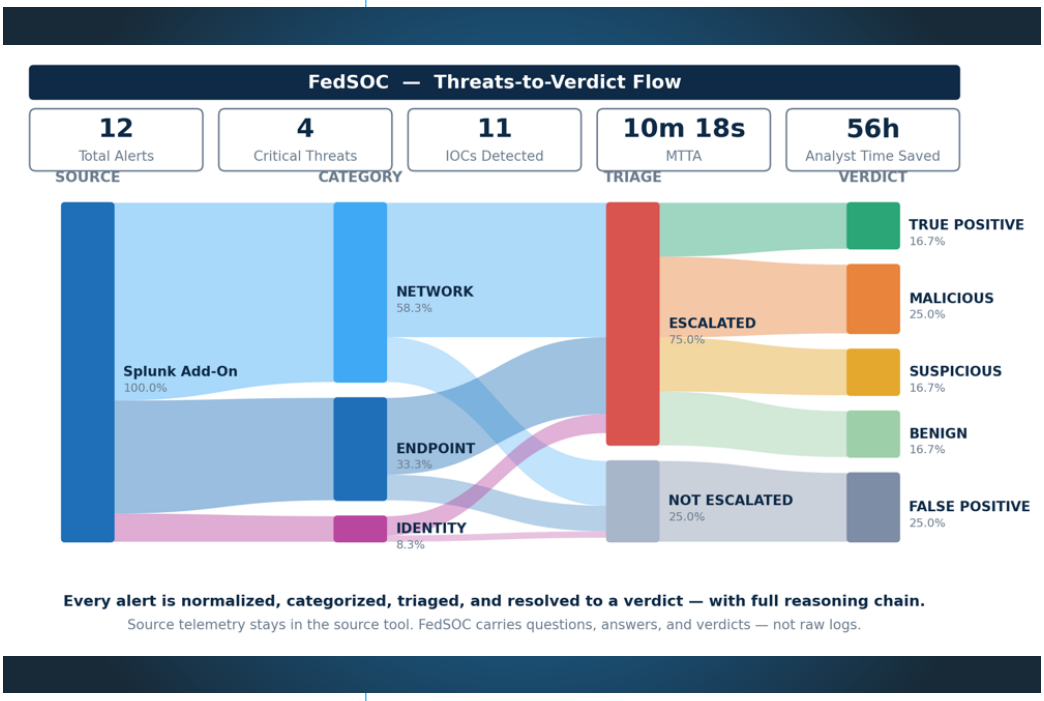

Figure 2 — FedSOC routes every alert from raw source telemetry through categorization, triage, and final verdict in one auditable flow.

Detections are authored in natural language. An engineer can describe threats, such as phishing activity, suspicious identity events, lateral movement, and exfiltration, and Federated SOC evaluates the relevant sources, identifies the most relevant telemetry, and generates production-ready detection logic. Coverage is continuously evaluated against MITRE ATT&CK, reachable attack paths, and real exposure across the environment to show where detection is strong, where blind spots remain, and where new detections should be recommended or generated.

That is the precise architectural distinction. Raw telemetry stays where it was generated. What FedSOC unifies is the semantic layer above it, a shared vocabulary, a normalized event model, and a continuously updated picture of how the environment is actually behaving. Every alert is normalized, categorized, triaged, and resolved to a verdict with a full reasoning chain. The platform carries questions, answers, and verdicts, but never raw logs.

AI SOC Agents: L1, L2, and Response

Within Federated SOC, the work that has historically been performed by human L1 and L2 analysts is handled by AI SOC agents under human oversight. An L1 Triage Agent receives every detection within seconds, validates the signal, enriches it against the Security Context Graph, and renders a verdict with supporting evidence. An L2 Investigation Agent maps timeline and blast radius, validates the kill-chain against Kairo's path model, evaluates adversary infrastructure and intent, and produces a full reasoning trace. A Response Agent proposes the highest-leverage control action (session revocation, password reset, MFA enforcement, endpoint isolation, IP block, ticket creation) and executes it within boundaries defined by the security team, with human approval where policy requires.

The human analyst's role changes rather than disappearing. Instead of processing a queue of 11,000 alerts a day, an analyst reviews agent verdicts, handles exceptions, hunts for novel threats, and makes the strategic decisions only humans can make. Transparency matters here: every reasoning step the agents produce is auditable, and every control action is logged, reversible, and bounded by policy. A platform that can't show the chain a model used to reach a verdict can't be evaluated for trust. Tuskira is built so that the chain is the default output.

Trust and Governance: Humans Stay in the Loop

A verdict is a decision artifact, and decisions have owners. The platform treats this explicitly. Customers define policy across three axes:

1. Which actions an agent can execute autonomously
2. Which require a named human approver
3. Which are blocked outrigh.

Policy is configured by severity, asset class, identity type, and business criticality. A session revocation on a low-tier service account is not governed the same way as a domain-controller isolation. Every agent verdict carries a confidence score, a full reasoning chain, and the data sources it consulted. Every action is logged, reversible, and replayable. Human overrides are first-class events that feed back into the platform's priors, so the agents become sharper at the exact boundaries the security team cares about.

When the agent is wrong, the system is designed to fail safely. High-impact or irreversible actions are not executed autonomously. They require human approval or are blocked outright by policy. False positives are captured as overrides, preserved in the audit trail, and used to update future priors, thresholds, and playbooks. For low-risk reversible actions, customers can define autonomous execution boundaries; for destructive or business-impacting actions, approval gates remain mandatory.

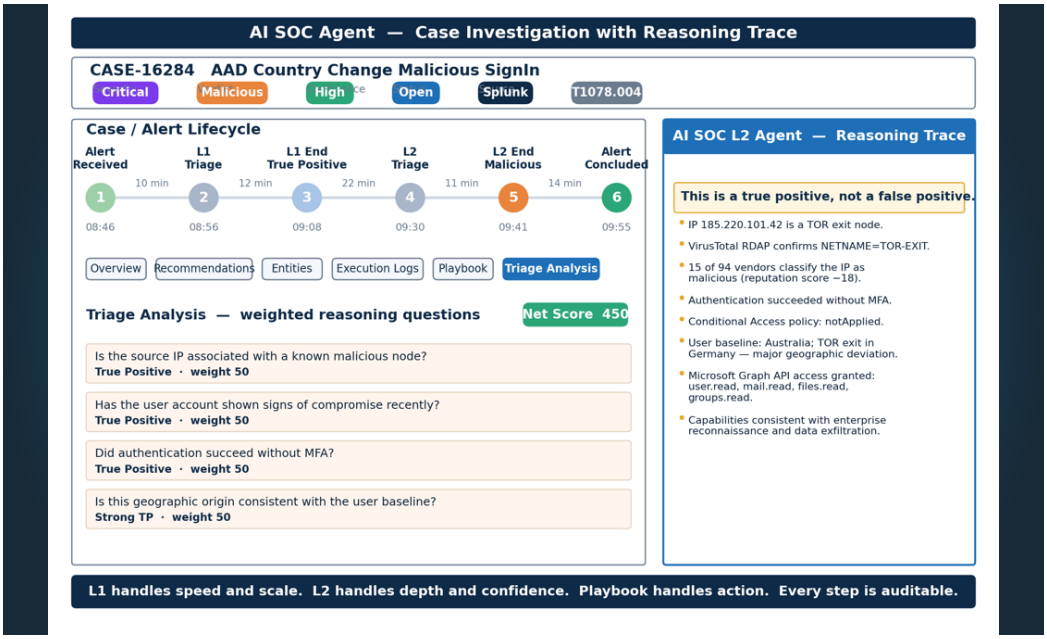

Figure 3 — A representative AI SOC agent case: an AAD country-change malicious sign-in. L1 triage validates the signal in under a minute; L2 investigation expands the evidence chain; every reasoning step is shown to the analyst alongside the case lifecycle.

CTEM : Preemptive Risk and Path Disruption (VM, Quell, Kairo)

Where Federated SOC operates on what is happening now, the CTEM portfolio operates on what could happen next. It models the environment as a live system and asks where an adversary could actually reach, what control would stop them, and how the answers change as the environment changes. The portfolio is organized into three production modules (Quell, Zero Day Response, and Kairo) that ride on the same Security Mesh and are accessible through Ask Tuskira.

Continuous Threat Exposure Management			
Module	**Vulnerability Management**	**Quell** Zero Day Response	**Kairo** Breach Modeling
Pain Point	Vulnerability backlog is growing. Teams need prioritization that focuses on real issues, not signals.	Zero-day CVEs are growing and current assessment requires significant human effort.	AI-driven recon and weaponization are growing. Attackers chain identity, endpoint, cloud, and lateral movement.
Outcomes	Prioritized CVEs and mitigations. Control policies identified. Measurable change over time.	Identify zero-days impacting the org. Control changes and mitigations for newly disclosed zero-days.	Identify reachable kill-chain artifacts. Prescribe control or infra changes to break the breach path.
How	Exploitability + Defensibility + Reachability analysis on a digital twin attack path model.	Zero-day agent ingests the ZD feed, hunts affected software and assets, identifies undefended exposure.	Breach modeling across identity, cloud, endpoint, and network — east-west, cross-cloud, insider.
AI Agent	Quell AI Agent + Ask Tuskira	Zero-day Agent + Ask Tuskira	Breach Modeling AI Agent + Ask Tuskira

Figure 4 — The CTEM offering matrix. Each module ships with a dedicated AI agent and is accessible through Ask Tuskira.

Kairo: Breach-Path Modeling

Kairo is the breach-path disruption capability and the most mature module in the CTEM portfolio. It continuously models the environment as a digital twin (identities, cloud accounts, workloads, networks, exposures, controls, telemetry) and enumerates the paths an adversary can actually traverse to crown-jewel assets. A newly disclosed CVE or AI-discovered zero-day matters only if it creates a reachable path to a breach in your environment. Kairo decides reachability, prescribes the highest-leverage control action, and tracks closure with evidence as the environment changes.

Step	What it does
1. Unify	Normalize identity, cloud, workload, and network signals into the Security Mesh.
2. Model	Build a live graph of reachability, privilege, and exploitability — a digital twin grounded in business context.
3. Map	Enumerate every traversable path to crown-jewel assets, ranked by exploitability and blast radius.
4. Identify residual	Surface paths existing controls do not block and existing detections do not see.
5. Disrupt	Orchestrate targeted control actions at the highest-leverage point — EDR, IAM, WAF, firewall, SIEM — with analyst approval where policy requires it.
6. Resolve	Track closure with evidence and re-validate reachability as the environment changes.

The result is that the platform reports breach paths instead of finding counts. In a typical deployment, more than 95 percent of findings are de-prioritized as unreachable, and a single targeted control action often closes multiple paths through a shared chokepoint.

Vulnerability Management and Zero-Day Response

Tuskira focuses the vulnerability backlog on the issues that matter, combining exploitability, defensibility, and reachability analyses with Kairo's digital-twin attack-path model to prioritize CVEs with public exposure and live exploitation, identify the control changes that mitigate them, and measure progress over time. Quell specializes in newly disclosed threats: the Zero Day Agent ingests the feed, hunts for affected software and assets, identifies which exposure is actually undefended, and proposes the control change to mitigate it. Identity Exposure Management, on the roadmap, will extend the same digital-twin model to identity blast radius and privilege chains.

Ask Tuskira: The Access Layer

Ask Tuskira is the natural-language interface to the entire platform. Any user (an L1 reviewer, a detection engineer, a CISO preparing for a board update, a compliance lead pulling evidence) can ask a question in plain English and receive an answer grounded in the same Security Context Graph the agents reason over. "Show me every reachable path to our production payments database." "Which of last week's CrowdStrike detections never propagated past the endpoint?" "Generate a report of the controls we changed in response to the Ivanti CVE." Ask Tuskira doesn't invent answers; it composes them from the same federated queries and the same agent reasoning chains the rest of the platform uses, and surfaces the citations so every claim can be verified.

The Operating Model Shift

The technical change is the easier half. The harder change is operational. When agents triage, investigate, and propose response, the SOC is no longer a queue-processing function; it becomes a supervisory function. That changes job descriptions, escalation policies, audit trails, and accountability.

L1 analysts move from "close as many tickets as possible" to "validate agent verdicts and hunt for novel threats." L2 and L3 analysts work on detection design, threat modeling, and the patterns that fall outside agent priors. Approvals shift from individual judgment calls to policy-governed decisions, so the question is no longer "do I trust this analyst" but "is this policy correct?" Incident command stays human; the work of preparing the incident-command briefing does not. Success metrics shift from MTTD and MTTR alone to verdict accuracy, override rate, breach-path closure, and analyst-hours redeployed to threat work. Tuskira is the platform that makes this shift possible. The shift itself is something the security organization has to lead.

A Day in the Federated, Agentic SOC

The following is an illustrative scenario, drawn from patterns observed across enterprise deployments. Details are composite, not a single anonymized incident.

It is 2:17 AM. A credential-stuffing campaign hits Azure AD using credentials from a third-party breach published earlier that evening. The native Azure AD detection fires immediately. FedSOC picks it up within seconds and asks the question across the federation: Who is this account? What does it have access to? Has this IP been seen before? Do these credentials match recent threat intelligence, and what is the user's normal baseline?

The L1 Triage Agent receives an enriched verdict from the Security Context Graph in under twelve seconds: the IP is a known TOR exit node, fifteen of ninety-four threat-intel vendors classify it as malicious, the credentials match the breach, the account has elevated permissions to a cloud bucket containing customer records, and the behavior violates every baseline for this user. The L1 verdict is true positive; the case escalates. The L2 Investigation Agent maps the blast radius and validates the kill chain against Kairo's path model. The breach path is reachable but not yet exercised. The Response Agent prepares the containment plan: revoke the session, disable the account, block the IP range, and create the incident ticket. A human reviews the reasoning chain and approves the recommended actions at 2:21 AM. Tuskira then executes the approved response and records the action trail.

Four minutes from initial attack to approved and executed containment. In a centralized, queue-based SOC, the same alert pattern can take hours because each stage of the attack is in a different queue. The contrast isn't about faster humans. It is about a different architecture.

That four-minute response is what shared context as the control plane actually looks like.

Outcomes and Operational Impact

For the CISO and CFO together, the practical case is economic. In one global financial services deployment, Tuskira reduced 12.3 million raw findings to 0.46% actionable risk within weeks, and triage time fell from three weeks to thirty minutes.

That result illustrates the broader pattern: adding more tools shouldn't mean adding more infrastructure costs, analyst load, or alert volume. With Tuskira, existing investments such as SIEM, EDR, IAM, CNAPP, WAF, and NDR get more useful because their native signals contribute directly to a shared context layer.

Across production deployments, outcomes vary by environment, integration breadth, and starting state, but customers have seen:

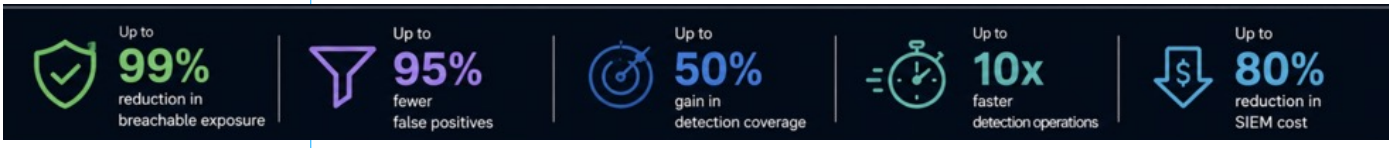

Architecture is the mechanism that produces those economics. The marginal cost of adding a new tool falls because no additional ingestion or storage is required. Detection coverage rises because each connected tool contributes its native detection capability, and FedSOC continuously closes MITRE ATT&CK gaps. Analyst time shifts from triage to judgment work because the platform produces verdicts under human oversight, not more alerts.

"2026 is the year cyber defenses are seeing the shift from AI-assisted attacks to AI-enabled attacks, and defenders need to adapt. That's why Intrado partnered with Tuskira."

—Charles Gifford
CISO, Intrado

What This Means for You

This isn't a rip-and-replace argument. The SIEM doesn't disappear. It instead becomes a consumer of enriched, correlated findings rather than the engine that produces them, a more focused and more valuable role. Tools you already own stay where they are; FedSOC connects to them through their native APIs. Integration breadth matters: the platform's intelligence is bounded by the tools it can see, so honest evaluation starts with a coverage map of your stack.

Organizational readiness matters more than technical readiness. The shift from "analyst processes queue" to "analyst oversees intelligent system" is a change to job descriptions and team structure, not a layoff. L1 analysts become reviewers and hunters. L2 and L3 analysts find their work elevated to genuinely strategic problems. Transparency is non-negotiable as every agent's verdict produces an auditable reasoning chain, and every control action is logged, reversible, and bounded by policy.

Evaluate security platforms by how well they connect context across detection, exposure, investigation, and response, not by how fast they triage alerts.

See the platform in action

Request a personalized demo of the Tuskira platform.

tuskira.ai/request-a-demo

www.tuskira.ai

Lorin@tuskira.ai