ANOMALI

White Paper

AGENTIC SOC PLATFORM

Threat Hunting and Hypothesis-Led Identity Hunting: Improving Early Threat Detection Through Identity-Centered Investigation

EXECUTIVE SUMMARY

Traditional security operations rely heavily on alerts to surface malicious activity. While effective for known threats, alert-driven detection often fails to identify stealthy, low-signal attacker behavior early in the attack lifecycle. As adversaries increasingly exploit legitimate identities, native tools, and SaaS platforms, organizations require a more proactive detection model.

This whitepaper presents a hypothesis-led identity hunting approach powered by the Agentic SOC Platform. By integrating threat intelligence, identity behavior, endpoint telemetry, and AI-assisted analytics, the platform enables analysts to test informed hypotheses against behavioral baselines. The result is earlier detection of advanced threats, reduced attacker dwell time, and a more effective proactive security posture.

THE LIMITS OF ALERT-DRIVEN DETECTION

Alert-centric detection models introduce structural limitations:

- Heavy reliance on predefined detection logic
- Reactive investigations triggered only after suspicious events surface
- Limited visibility into identity misuse and subtle behavioral deviations

As attackers adopt living-off-the-land techniques and abuse legitimate credentials, malicious activity increasingly blends into normal operations. By the time alerts fire, attackers may already have established persistence or lateral access.

BUSINESS AND OPERATIONAL IMPACT

SECURITY OPERATIONS IMPACT

Without hypothesis-driven hunting, security teams experience:

- Inefficient or unfocused threat hunts
- Missed early indicators of compromise
- Delayed detection of lateral movement and SaaS abuse

BUSINESS IMPACT

These detection gaps lead to:

- Larger attack blast radius
- Increased remediation and recovery costs
- Reduced confidence in proactive security capabilities

Organizations remain dependent on alerts for visibility, limiting their ability to disrupt attacks early.

HYPOTHESIS-LED IDENTITY HUNTING: DEFINING THE USE CASE

Hypothesis-led identity hunting shifts detection from passive alert consumption to active investigation. Analysts begin with informed questions grounded in threat intelligence and identity behavior, then test those hypotheses across correlated datasets.

Rather than waiting for alerts to surface anomalies, this approach proactively evaluates whether observed activity meaningfully deviates from expected behavior or aligns with known adversary techniques.

The Agentic SOC Platform enables this model by unifying identity, endpoint, and intelligence data while supporting structured, repeatable hunt workflows.

HOW THE AGENTIC SOC PLATFORM POWERS THIS USE CASE

UNIFIED SECURITY DATA LAKE: BEHAVIORAL CONTEXT AT SCALE

Effective hunting requires broad visibility and historical context. The Agentic SOC Platform aggregates identity, endpoint, network, and SaaS telemetry into a unified data lake, preserving behavioral baselines over time.

Analysts can compare current activity against historical norms, enabling detection of subtle deviations that may indicate early compromise.

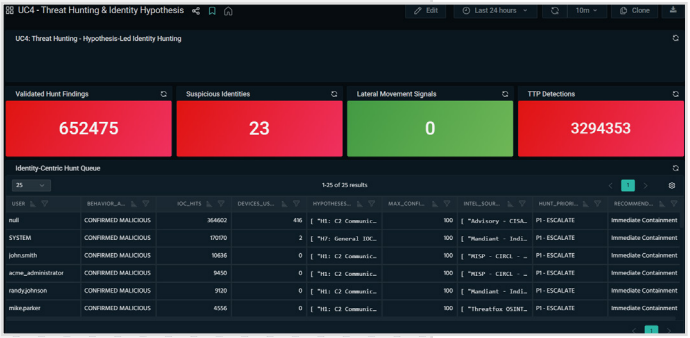

THREATSTREAM NEXT-GEN: INTELLIGENCE-DRIVEN HYPOTHESES

ThreatStream Next-Gen provides adversary techniques, indicators, and campaign context that inform hunt hypotheses. Intelligence guides analysts toward behaviors that matter, focusing investigations on techniques actively used in the wild rather than theoretical risks.

This ensures that hunting efforts are targeted, relevant, and aligned with real-world threats.

AGENTIC AI: ACCELERATED HUNT EXECUTION

Executing complex hunts across multiple datasets is time-consuming. Agentic AI assists analysts by correlating identity behavior, endpoint activity, and threat intelligence, helping validate or disprove hypotheses more efficiently.

AI-assisted analytics reduce manual effort while maintaining analyst control and interpretability.

IDENTITY-CENTRIC ATTACK SCENARIOS

Hypothesis-led identity hunting supports investigation of early-stage and stealthy techniques, including:

- Initial access through suspicious execution paths
- PowerShell-based download and execution
- Persistence via registry modification
- Privilege escalation and defense evasion
- Credential access and discovery activity
- Lateral movement and command-and-control behavior
- Data exfiltration and tunneling activity

By correlating identity context with endpoint behavior, these techniques can be detected earlier in the attack lifecycle.

BUSINESS AND SECURITY OUTCOMES

Organizations adopting hypothesis-led identity hunting achieve:

- Earlier detection of advanced and stealthy threats
- Reduced attacker dwell time
- Improved operational use of threat intelligence
- Stronger proactive security posture

Threat hunting becomes a repeatable, evidence-driven practice rather than an ad hoc exercise.

OPERATIONALIZING THE USE CASE

Implementation typically includes:

- Leveraging curated detection content such as Sigma rules
- Correlating identity and endpoint telemetry with expected log sources
- Executing hunts to evaluate detection logic
- Converting validated hunts into alert rules and monitoring workflows

This approach integrates into existing SOC processes and enhances, rather than replaces, alert-based detection.

CONCLUSION

Threat hunting is most effective when driven by hypotheses grounded in identity behavior and real-world threat intelligence. Alert-only models leave organizations blind to subtle, early-stage attacker activity.

By enabling identity-centric, hypothesis-led investigation, the Agentic SOC Platform empowers security teams to uncover stealthy threats earlier, reduce risk, and move from reactive detection to proactive defense.