Sponsored content | White paper

Stop breaches fast: How threat hunting pinpoints bad actors early in their attacks

CIO

Sponsored by zscaler™

Attackers are changing how they operate — and they're doing it faster than most organizations can adapt. Instead of relying on obviously malicious files, infrastructure, or noisy activity on devices, adversaries now hide in legitimate web resources, leveraging trusted software-as-a-service (SaaS) tools, and blend in with operational network traffic. These tactics enable attackers to establish a foothold while evading detection, creating an early-stage visibility gap.

Zscaler Threat Hunting closes this gap. It helps organizations detect these nascent behaviors early in the attack life cycle, when risk and the remediation cost are lower. By analyzing cloud-scale internet telemetry across the Zscaler Zero Trust Exchange™, the service surfaces subtle signals in Zscaler Internet Access (ZIA) traffic. Zscaler combines these detections with human-led investigation and proprietary intelligence from its security research team ThreatLabz to give security teams a clearer view of developing threats.

The new visibility challenge

Finding bad actors early in their attack is challenging when their early-stage activity increasingly blends into normal traffic patterns. We're seeing the following types of hard-to-spot behaviors in attacks these days:

- Users are unknowingly redirected to "typo-squatted," compromised, or newly registered domains.
- Malware is delivered from legitimate cloud storage, such as Cloudflare R2.
- Credential harvesting pages mimic trusted brands and trending technologies such as generative AI (genAI).
- Covert command-and-control (C2) channels "check in" through rare domains or proxies through trusted cloud infrastructure.

These activities can unfold without triggering obvious device-level indicators, especially during reconnaissance, delivery, and exfiltration stages. Security teams need insight into these upstream behaviors to intervene earlier and reduce the window of opportunity for an attacker.

How Zscaler Threat Hunting works

Zscaler Threat Hunting analyzes global ZIA telemetry — billions of daily transactions across user, application, and web destinations — to identify anomalies within seemingly legitimate traffic. Using Zscaler's proprietary TRACER (telemetry, refine, analyze, context,

escalate, revise) methodology, the hunt team correlates behaviors that may indicate an emerging threat, such as unusual download patterns, connections to rare domains, or traffic consistent with early C2 activity.

Because Zscaler sees internet-bound requests before any code executes on an endpoint, it can detect malicious behavior in the earliest moments of the attack life cycle. For example:

- **Weaponization:** Spotting impersonation, compromised sites, and typo-squatted domains
- **Delivery:** Identifying compromised remote management tools that bypass endpoint protection
- **Command and control:** Detecting beaconing patterns that suggest that a foothold has been established

Each finding feeds into analyst review while simultaneously enriching signals with threat intelligence from hundreds of tracked adversary groups. The result is higher-fidelity detections delivered proactively to customer security teams.

Human expertise at cloud scale

Machine learning provides the speed needed to sift through massive volumes of telemetry, but human expertise is what distinguishes threat hunting from automated detection. Zscaler's team evaluates the context around each anomaly — whether a behavior represents a known malicious pattern, a variant of an existing campaign, or a previously unseen threat actor technique.

This continual refinement cycle helps the team:

- Improve signal-to-noise ratio, increasing efficacy
- Build new behavioral models based on live campaigns
- Improve the speed and accuracy of downstream investigations
- Identify both broader and targeted threats appearing across customers

Customers benefit from detections shaped by both global visibility and expert analysis, capabilities most organizations cannot replicate with local tools alone.

Detecting real-world threats earlier

Zscaler Threat Hunting routinely uncovers sophisticated campaigns that exploit user trust and internet reputation, including:

- **Fake CAPTCHA campaigns**, where users are tricked into pasting malicious PowerShell code into terminal windows. Zscaler identifies these attacks through identification of suspicious JavaScript files.
- **Malware hosted in cloud infrastructure**, where adversaries hide payloads in legitimate cloud infrastructure such as Cloudflare R2. Zscaler detects unusual file-retrieval patterns that precede endpoint execution.
- **GenAI impersonation scams**, where attackers exploit trending AI tools to lure users into sharing credentials and corporate secrets. Zscaler surfaces these attacks through approximate matching analysis.

Early detection of these behaviors gives organizations time to intervene — blocking malicious infrastructure, revoking tokens, or isolating suspicious activity before attackers can do much damage.

Complementing endpoint hunting and SOC workflows

Zscaler Threat Hunting is designed to work alongside existing endpoint detection and response programs, not compete with them, and provides visibility into different parts of the cyberattack kill chain.

Together, they create broader visibility across the attack surface, improve coverage at every life cycle stage, and accelerate investigations by providing broader context. Zscaler's validated findings help security operations center (SOC) teams prioritize what matters most, regardless of where the threat originates.

Discover how Zscaler Threat Hunting can help you find and stop bad actors early in their attacks.