ZSCALER DECEPTION

Disrupt AI-Orchestrated Attacks

AT-A-GLANCE

Outpace Machine-Speed Threats with High-Fidelity, High-Intent Detections

The Rise of the Automated Kill Chain

The threat landscape is undergoing a seismic shift as adversaries move beyond manual attacks to AI-orchestrated cyber espionage. The first confirmed AI-orchestrated campaign highlighted this new reality. In the attack, a threat actor utilized AI agents—chained through an MCP server and equipped with tools for binary analysis, network scanning, and password cracking—to successfully conduct reconnaissance, map attack surfaces, write custom payloads and exfiltrate data from targeted organizations. These attackers combined the competence of a nation-state actor with the operational tempo of a machine.

This compressed timeline has created a critical gap in traditional defenses. The result is a new baseline where adversaries can successfully compromise an environment and exfiltrate data long before they are detected.

HUMAN ADVERSARY	AI-AGENT ADVERSARY
Human Pace: Moves in minutes or hours, requiring manual research and deliberate execution	**Machine Speed:** Executes steps in seconds, operating at a tempo that outpaces traditional defenses
Manual Research: When blocked, humans must stop to research, script, and test new fixes	**Instant Iteration:** AI reads error logs and generates new code instantly to bypass roadblocks in real time
Single Path: Typically follows a sequential, one-step-at-a-time approach through the network	**Parallel Execution:** Explores every vulnerability and moves across the entire environment all at once
Weeks or Months: A full compromise often takes weeks of quiet, manual trial and error	**Minutes or Hours:** Completes the entire attack cycle before traditional tools can even correlate the logs

Turn AI Curiosity into Detection

To counter these machine-speed threats, defenders need alerts that provide broad coverage and absolute accuracy. Traditional detection methods relying on signatures and behaviors cannot keep up with the pace of AI-driven attacks. Because these agents use valid credentials and legitimate tools, they force defenders into a 'probabilistic trap'—spending hours investigating ambiguous alerts that are often indistinguishable from routine employee activity. Zscaler Deception shifts the model from "maybe" alerts to deterministic detection, meaning every alert is a confirmed threat, not a guess.

By planting decoys such as fake databases, servers, endpoint assets, public cloud resources, and Active Directory services and privileged accounts alongside real assets, Deception creates an environment where any interaction is an unambiguous indicator of a threat.

AI-orchestrated attacks operate by attempting to exploit every vulnerability at once and move through an environment across multiple paths simultaneously. This specific behavior makes Deception uniquely effective: the probability of an AI attacker hitting a decoy is nearly 100% because they attempt to exploit everything at once.

Neutralize Agentic Decision Making

As adversaries deploy autonomous agents to navigate environments and execute kill chains, defenders require a strategy that effectively targets agentic decision making. Zscaler Deception disrupts this automated logic by introducing deterministic traps that force attackers—whether human or machine—to reveal their intent instantly. By shifting the defensive focus from searching for suspicious behavior to planting high-fidelity triggers, Deception provides a more effective way to combat agentic threats through:

- **High-signal, intent-based detection:** Only adversaries trip decoy credentials, services and files. These alerts lead to early detection and triage.
- **Low cost and low operational overhead:** Decoys don't sit inline, they don't break apps, and can be deployed for ZPA users in one click.
- **Defense-in-depth multiplier:** Deception amplifies your identity, endpoint, network, and data-layer controls, catching early pivots that other tools either miss or see later and generate undue noise.
- **Breaking attacker tempo:** Decoys introduce unpredictability into your environment. Attackers, whether human or AI, must probe assets and artifacts to exploit them. Probing leads to detection, slowing AI-driven campaigns and buying your SOC time.

Counter AI-Driven Attacks with Deception

Zscaler Deception can find and contain both human and AI-orchestrated bad actors across your entire environment.

Perimeter Traps	Honeytokens and Honeyusers	Realistic Decoys
Catch attackers before they enter your network with external decoys that mirror your VPNs and firewalls.	Seed decoy users and keys into your directory to trigger an instant lockout the moment they are used.	Deploy fake databases and consoles that look like your real systems, enabling more chances to catch AI-driven attacks.
Lures	**Orchestrated Response**	**Multi-Point Coverage**
Place beacon files within sensitive data that send an alert the second they are opened or exfiltrated.	Automatically isolate compromised users and hosts and revoke access the moment a decoy is touched.	Detection across perimeter, endpoints, AD, networks, cloud workloads, and OT/IoT

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange™ platform protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

+1 408.533.0288 **Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134** **zscaler.com**

zscaler™

Zero Trust Everywhere