ZSCALER

Defend Against Frontier AI Attacks with Zero Trust

AT-A-GLANCE

Frontier AI moves at machine speed. Your architecture needs to move first.

AI has industrialized the attack surface. Frontier AI models compress reconnaisance and exploit cycles from weeks to minutes, autonomously chaining vulnerabilities, bypassing legacy controls, and orchestrating attacks across your entire environment in parallel. The answer isn't a faster response. It's an architecture that removes what they can reach, restricts where they can move, and catches them the moment they act.

These models have permanently reset the defender's clock. Any internet-reachable VPN, application portal, or exposed server is now a near-guaranteed liability, not a risk to be managed, but a vulnerability to be found and exploited within minutes.

These attacks are no longer sequential. AI-driven attacks run discovery, exploitation, privilege escalation, and lateral movement simultaneously, at machine speed. The concept of dwell time, where teams accumulate and correlate alerts over hours, is no longer viable. Detection and response alone cannot stop AI-driven attacks from doing damage.

Architecture Beats AI

Frontier AI makes attacks faster, cheaper, and more scalable. But the best defense is still architectural. Organizations need a zero trust approach that minimizes exposure, restricts reachability, limits lateral movement, and stops exploits before they succeed. These four capabilities work together to reduce what attackers can see, reach, and abuse at every stage of the kill chain.

Capability #1: Zscaler Private Access (ZPA)

Eliminate the attack surface. Make every application invisible by default.

Many AI-driven attacks begin with reconnaissance, scanning exposed infrastructure, harvesting DNS entries, probing VPN concentrators, and querying certificate transparency logs. If it's reachable, AI will find it. If AI finds it, it will be attacked.

ZPA hides your private applications entirely behind the Zero Trust Exchange™, meaning no public IP, no inbound port, no publicly resolvable DNS, and no certificate exposure. Users connect through the Exchange, not the network. Attackers scanning the internet see nothing. The attack surface is eliminated before the first probe lands.

ZPA also continuously verifies every connection—identity, device posture, and context—so that even if credentials are stolen through AI-orchestrated phishing or deepfakes, they are useless without the matching device and context. With ZPA you can:

- **Minimize your attack surface & block threats:** Hide apps from the internet with inside-out connections; inspect TLS traffic
- **Eliminate vulnerable infrastructure:** Retire discoverable VPNs and edge firewalls
- **Reduce cost and complexity:** Enable fast, direct access to private apps for remote and in-office users while reducing IT overhead

Capability #2: Autonomous User–to–App Segmentation

Stop lateral movement. Keep a single compromise from turning into an enterprise–wide breach.

Moving applications behind ZPA is the first step. But when a user or device is compromised, and you should assume it will be, the real question is: how far can the attacker move? Frontier AI accelerates post–compromise actions: privilege escalation, internal reconnaissance, and pathfinding to crown jewels. On a flat network, one compromised endpoint is an enterprise incident.

Zscaler's Autonomous User–to–App Segmentation observes your real user–to–application access patterns, maps them to identity (AD groups, department, role), and recommends precise least–privilege policies automatically, reducing a multi–week firewall rule project to a single click. A single accepted recommendation can reduce the exposed user population for a critical application significantly, instantly shrinking the blast radius.

- **See every application:** Automatically discover and classify every application in your environment
- **Eliminate lateral movement:** Leverage AI–driven least–privilege policy recommendations from live traffic analysis
- **Strengthen your posture:** Monitor access and policy usage with an Insights Dashboard for continuous optimization

Capability #3: Deception

Machine–speed threat detection, deterministic alerts and automated containment.

Frontier AI models are especially prone to triggering deception because they use multi–path reasoning, exploring many possible routes to achieve an objective in parallel. In practice, that means they are far more likely than a human attacker to touch decoys, probe honeypots, and reveal themselves early. No legitimate user has any reason to touch a decoy, so every interaction is confirmed malicious, with no ambiguity. Whether it's an agentic AI attacker or a compromised insider, the signal is the same: deterministic proof of compromise.

One–click deployment of Zscaler Deception activates a full deception strategy across your zero trust environment with no new agents, no hardware, and no network reconfiguration. Deception is enabled through your existing ZPA infrastructure, with no changes to users. Decoy applications, credentials, and services are deployed alongside your real environment, creating a minefield that any attacker, human or AI, will inevitably trigger as they move laterally and will:

- **Stop lateral movement:** Decoy applications appear alongside your real private apps in ZPA. Attackers attempting to pivot hit decoys instead of real resources
- **Redirect attackers with endpoint lures:** Decoy credentials and bookmarks deployed via Zscaler Client Connector lead compromised identities straight to decoys, away from real resources
- **Trap attackers probing unauthorized services:** Decoy services on real apps catch anyone scanning ports that legitimate users would never touch

Capability #4: Autonomous App Shield

Stop the attack before it reaches your application.

Patching is no longer a viable first line of defense. The mean time to exploit a published CVE is now negative seven days. Attackers weaponize vulnerabilities before defenders can respond. Frontier AI doesn't just find individual CVEs: it can chain medium-severity issues together into sophisticated, multi-step exploit paths that traditional scanners never surface. Legacy and custom applications may never receive patches at all.

Zscaler Autonomous App Shield is an AI-powered application protection layer that deploys inline through existing ZPA App Connectors with no new infrastructure required. It continuously fingerprints your private applications, constructs the exploit chains an attacker would build, and generates virtual protection signatures in the Zero Trust Exchange within minutes. Attacks are blocked inline before they ever reach the vulnerable application, providing immediate protection until patches can be applied.

- **Discover application risk continuously:** AI harness scans applications with dynamic context to maintain an always-current view of exploitable weaknesses.
- **Decide the right protections automatically:** Zscaler uses AI to determine which signatures and virtual protections each individual application needs.
- **Deploy defenses at machine speed:** Blocks exploit attempts inline at the App Connector before they reach private applications.

One Platform, Built for the Frontier AI Era.

ZPA, Autonomous User-to-App Segmentation, Deception and Autonomous App Shield run as an integrated system on the Zscaler Zero Trust Exchange: the world's largest inline security cloud, processing 500+ billion transactions daily across 160+ global data centers.

When Deception triggers, ZPA revokes access in milliseconds, with no analyst approval required. When Autonomous App Shield generates a virtual patch, it is enforced across the Exchange instantly. When Segmentation tightens a policy, it applies across every user and device simultaneously. This is not a bundle of point tools, it is a coordinated, machine-speed defense that operates at the same pace as the attacks it stops.

Contact your Zscaler account team to start your zero trust defense against frontier AI attacks.

About Zscaler

Zscaler (NASDAQ: ZS) is a pioneer and global leader in zero trust security. The world's largest businesses, critical infrastructure organizations, and government agencies rely on Zscaler to secure users, branches, applications, data & devices, and to accelerate digital transformation initiatives. Distributed across 160+ data centers globally, the Zscaler Zero Trust Exchange™ platform combined with advanced AI combats billions of cyber threats and policy violations every day and unlocks productivity gains for modern enterprises by reducing costs and complexity.

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

zscaler™

Act Fast.
Stay Secure.