ZERO Networks

How to Build an AI-Ready Network Architecture

A Framework for Benchmarking AI Readiness

Table of Contents

03 Overview and Maturity Model at a Glance

04 Stage 1: Initial

05 Stage 2: Developing

06 Stage 3: Intermediate

07 Stage 4: Advanced

08 Stage 5: Holistic

09 Why Organizations Stall Before Holistic Maturity

How Zero Networks Accelerates the AI Security Journey

ZERO. Networks

Overview

AI has collapsed the time between vulnerability discovery and exploit. Tools like Anthropic's Mythos and OpenAI's Daybreak can find and weaponize vulnerabilities at expert-human level — autonomously. Attackers move laterally in as little as **27 seconds**[1], and organizations take an average of **247 days to detect and contain a breach.**[2] Shadow AI and ungoverned agents are already proliferating inside enterprise environments — invisible and unrestricted.

Detection and response can't close that gap. Attackers need to be contained and paralyzed, not chased.

This framework maps five stages of microsegmentation maturity against that reality — showing where the exposure lives at each stage, and what it takes to eliminate it.

Each maturity stage includes:

- The security controls in place and their limitations
- The AI-specific threat exposure at that stage
- How Zero's AI Segmentation capabilities address the gap

The AI Threat Context

Anthropic's Claude Mythos Preview can find and exploit software vulnerabilities at expert-human level — autonomously. OpenAI's Daybreak followed in May 2026, bringing the same capability with broader availability. Mozilla's agentic pipeline found 271 of 423 Firefox vulnerabilities in a single cycle. What previously took expert humans weeks now takes hours — and the window is shrinking to minutes.

Meanwhile, roughly 80% of enterprises have AI tools actively running inside their environment, with ~0% of AI access actually blocked.[3] The gap between visibility and control is where AI risk lives.

Maturity Model at a Glance

Stage	Controls	AI Exposure
1 - Initial	Perimeter firewalls	Flat network, all AI ungoverned
2 – Developing	VLANs / ACLs / Internal FW	Zone-based but no AI identity controls
3 – Intermediate	Application ringfencing	Intra-ring AI lateral movement
4 – Advanced	Per-asset host controls	Static policies lag AI speed
5 – Holistic	MFA microseg of all assets & IDs	AI agents governed; AILM blocked

ZERO. Networks

1 CrowdStrike, 2026 Global Threat Report
2 IBM, Cost of a Data Breach Report 2026
3 IBM, Agentic AI Security Guide

Stage 1: Initial

Your network boundary is protected. Everything inside is open.

Flat network, all AI ungoverned; every shadow tool and agent has free run of the environment

Typical Controls	Traditional perimeter firewalls blocking north-south traffic. Flat internal network with minimal east-west controls.
Security Outcome	Blocks external threats at the edge. Once an attacker is inside, lateral movement across the entire network is largely unrestricted.

CISO Reality

- Effective for blocking external threats at the perimeter
- No meaningful control over east-west traffic once an attacker is inside
- A single compromised endpoint becomes a launchpad across the environment

AI Security Exposure at This Stage

- Shadow AI tools (ChatGPT, Gemini, Copilot, and thousands of others) run unchecked — no network-layer enforcement to distinguish sanctioned from unsanctioned AI
- AI agents operating inside the network have no identity-based least-privilege boundary; they can reach any system they can address
- Mythos-class AI can discover and exploit vulnerabilities; with a flat internal network, a single exploit immediately translates into unrestricted lateral movement
- No visibility into what AI tools are communicating with or what data is leaving the environment

How Zero Networks Addresses the AI Gap

- AI SaaS Control: Block all unsanctioned AI destinations at the network layer with a single rule — no proxy or DLP infrastructure required
- AI Lateral Movement Control (AILM): Establish default-deny identity-based least privilege, eliminating the flat network that makes AI-discovered exploits so damaging
- AI Agent Control: Auto-discover agents running across managed assets and apply the same identity-bound controls you give users

Next step

Implementing VLANs, ACLs, and internal firewalls to begin separating network zones — or skip straight to Holistic with automated microsegmentation.

ZERO. Networks

Stage 2: Developing

Your network is zoned. But zones don't stop AI agents — or the attackers they enable.

Zones exist but AI has no identity constraints; manual policy management can't keep pace with shadow AI

Typical Controls	VLANs, ACLs, and internal firewalls separating user, server, OT, and cloud zones. Policies are hardware-dependent and manually maintained.
Security Outcome	Better containment between departments or environments. Static, manual rule management creates complexity and operational overhead over time.

CISO Reality

- Improved containment between departments and environments
- Common in data centers, OT, and hybrid cloud architectures
- Often hardware-dependent and manually maintained, increasing operational overhead
- Static rule sets become difficult to update as the environment changes

AI Security Exposure at This Stage

- VLANs segment user populations but don't govern AI tool access — employees in any zone can still reach unsanctioned AI services
- AI agents operating within a zone have no per-identity constraints; trust within a segment remains broad
- Manual rule management means AI-related policy updates (adding new sanctioned tools, blocking emerging threats) are slow to deploy and easy to misconfigure
- Mythos-class exploits can pivot between zones through misconfigurations or overly permissive inter-VLAN rules — a common artifact of manual policy management

How Zero Networks Addresses the AI Gap

- AI SaaS Control: Enforce a centrally managed allow-list of sanctioned AI destinations across all VLANs simultaneously — updates propagate instantly, no per-zone configuration
- AILM: Replace static inter-zone rules with identity-based least-privilege connections, reducing the attack surface that AI-discovered inter-VLAN paths exploit
- Compliance & Risk Engine: Natural-language queries across billions of connections reveal which AI tools are communicating across zone boundaries

Next step

Implementing application or workload ringfencing to isolate groups of assets — or skip straight to Holistic with automated microsegmentation.

ZERO. Networks

Stage 3: Intermediate

Your critical applications are ringfenced. AI agents inside the ring are still trusted implicitly.

Ringfencing helps but AI agents inside a ring still have broad intra-ring access

Typical Controls	Ring-fencing of application groups (PCI systems, production workloads, dev/test environments, sensitive server clusters). Traffic into the group is controlled; trust within the group remains open.
Security Outcome	Strong improvement over flat networks. Useful for risk reduction and often aligned with CMDB/ application ownership models. If one system inside the ring is compromised, lateral movement within the group remains possible.

CISO Reality

- Meaningful improvement over flat network architecture
- Useful transitional step for risk reduction and compliance alignment
- Often maps well to application ownership and CMDB models
- Broad trust within each ring leaves residual exposure to intra-group lateral movement

AI Security Exposure at This Stage

- AI agents operating within a ringfenced application group still have broad access to every other system in that ring — a compromised agent or prompt-injected coding tool can move laterally within its segment
- Sanctioned AI controls may exist at the perimeter of a ring but not within it — sensitive internal systems remain reachable once AI is inside the ring
- AI-discovered vulnerabilities in a single system within a ring enable intra-ring lateral movement, which may be sufficient to reach high-value targets
- No per-identity controls on AI agents — they are governed by network position, not verified identity

How Zero Networks Addresses the AI Gap

- AI Agent Control: Extend least-privilege boundaries inside each ring — every agent is governed by identity, not just network position
- LLM Protection: Segment model infrastructure at the network layer; only authorized systems within the ring can reach AI/ML workloads
- AILM: Eliminate intra-ring lateral movement by applying per-identity network controls to every connection, not just connections crossing ring boundaries

Next step

Implementing more granular segmentation to isolate individual assets — or skip straight to Holistic with automated microsegmentation.

ZERO. Networks

Stage 4: Advanced

Every asset has a policy. But static policies can't keep pace with AI's attack speed.

Per-asset controls exist but static policies can't adapt to AI agent behavior in real time; non-human identities lack governance

Typical Controls	Host-level policy controls between individual workloads. Every server, endpoint, and cloud workload receives its own access policy. Dramatically reduces blast radius.
Security Outcome	Prevents lateral movement between individual assets. Enables least-privilege networking. Without automation, policy creation and maintenance become operationally intensive.

CISO Reality

- Prevents lateral movement between individual workloads
- Dramatically reduces blast radius of any single compromise
- Enables true least-privilege networking
- Without automation, policy/labeling creation and maintenance are operationally heavy — and rule sprawl creates coverage gaps

AI Security Exposure at This Stage

- Static per-asset policies can't adapt in real time to the behavioral signals of AI agents — a compromised agent operating within its allowed policy may go undetected
- Policy management overhead is high; AI-related updates (new agent, new sanctioned SaaS, new internal AI workload) require manual policy changes that lag behind deployment
- Identity-based controls may be incomplete — policies based on IP or host identity don't govern user-to-AI or AI-to-asset connections with the same precision as MFA-verified access
- AI agents are non-human identities; without agent-specific identity binding, they inherit overly permissive host-level policies

How Zero Networks Addresses the AI Gap

- AI Agent Control: Apply identity-based least-privilege to every agent — auto-discover all agents, bind each interaction to a verified identity, and block any out-of-boundary connection with immediate SOC alert
- AILM + MFA: Extend just-in-time MFA to AI-to-asset connections, ensuring that every connection a privileged AI agent makes requires explicit authorization
- Compliance & Risk Engine: Continuously score AI agent behavior against NIS2, CIS, and PCI-DSS; surface live gaps between current AI access patterns and policy intent

Next step

Implementing comprehensive microsegmentation of all assets and identities with dynamic MFA.

ZERO. Networks

Stage 5: Holistic

Every connection is verified. Every identity is governed. AI is a controlled, visible participant — not an unmanaged risk.

AI agents are governed by granular identity-based controls; Mythos-class exploits lead to dead ends – even when attacks are AI-accelerated

Typical Controls	Comprehensive microsegmentation combining identity verification (MFA), per-session authorization, least-privilege network access, host-level controls, user-to-asset and workload-to-workload segmentation, and continuous adaptive policy enforcement.
Security Outcome	Every connection is explicitly verified. Every asset is protected. Every identity — human or AI — is governed. Trust is never assumed. This is where segmentation becomes a business enabler, not just a firewall project.

CISO Reality

- Comprehensive protection across users, devices, AI agents, and workloads
- Dynamic Zero Trust enforcement — policies evolve automatically to meet need
- Intelligent automation eliminates manual rule management overhead
- Every connection is verified; every identity is governed; trust is never assumed

AI Security Exposure at This Stage

- AI risk is governed by the same identity-based least-privilege framework as all other access — no special-casing required
- Shadow AI is blocked at the network layer; sanctioned tools are explicitly permitted
- AI agents are treated as non-human identities: auto-discovered, identity-bound, and subject to just-in-time MFA for privileged access
- Mythos-class AI may find vulnerabilities, but every exploit path leads to a dead end — no lateral movement is possible without explicit, verified authorization

How Zero Networks Addresses the AI Gap

- Full AI Segmentation platform: AI SaaS Control, AI Agent Control, LLM Protection, AILM, and Compliance & Risk Engine operating as a unified control plane
- Natural-language queries across billions of connections for continuous AI risk visibility and audit-ready evidence
- Automated policy lifecycle management — policies adapt to behavioral signals without manual intervention
- Live scoring against NIS2, CIS, DORA, and PCI-DSS; real-time gap identification as AI adoption evolves

Next step

You've achieved the leading edge of segmentation maturity. Maintain and extend AI governance as AI adoption grows — new agents, new models, and new AI surfaces require continuous visibility and control (deterministic, human-on-the-loop automation).

ZERO. Networks

Why Organizations Stall Before Holistic Maturity

Most segmentation programs — and the AI governance efforts layered on top of them — fail because they rely on:

Manual rule writing: Static firewall constructs can't keep pace with AI agent proliferation, new model deployments, or Mythos-speed vulnerability discovery.

Ticket-driven policy changes: AI adoption moves faster than change management cycles. Shadow AI tools appear and disappear faster than policies can be written.

Poor visibility: Without automatic discovery of AI agents and services, you can't govern what you can't see.

Lack of identity integration: Network-position-based controls treat AI agents as any other host — they don't apply identity-based least-privilege to non-human actors.

Phased implementation complexity: CISA's microsegmentation guidance involves multiple phases. Without automation, each phase introduces new operational overhead and creates temporary exposure windows.

How Zero Networks Accelerates the AI Security Journey

Unlike legacy segmentation vendors that depend on heavy professional services and manual policy engineering, Zero Networks uses automation-first microsegmentation — and extends that automation-first approach to AI:

Automated Policy Creation: Learns legitimate communication patterns — including AI agent behavior — and builds controls automatically. No manual rule writing.

AI SaaS Control: Sanction ChatGPT, Gemini, Copilot. Auto-block every other AI destination at the network layer. One rule, enforced everywhere.

AI Agent Control: Auto-discover every agent. Apply identity-based least-privilege to every interaction. Block and alert on any out-of-boundary connection.

LLM Protection: Segment model infrastructure so only authorized systems can reach it. Code that can't be reached can't be tampered with.

AI Lateral Movement Control (AILM): Default-deny, identity-bound least privilege removes every path an exploit needs. Tools like Mythos and Daybreak find the door; AILM makes sure it opens onto an empty room.

Compliance & Risk Engine: Natural-language queries across billions of connections. Live NIS2, CIS, DORA, and PCI-DSS scoring. Audit-ready evidence on demand.

ZERO. Networks



Benchmark your AI readiness, uncover your biggest risks, and get actionable next steps in a shareable report with Zero's AI Readiness Assessment

Take the Assessment



Stop lateral movement and contain any attack — human or AI — with automated, identity-driven microsegmentation

Book a Demo

ZERO. Networks