Akamai

Cyber Resilience in the AI Era: Securing with Microsegmentation

Peeravas Sansupakorn
Solution Engineer
Akamai Thailand

Introduction

58%

Increase in ransomware attacks in 2025

24 days

Average downtime following a ransomware incident

US$276 billion

Estimated cost of ransomware damages annually by 2031

US$5.5 million

Recovery costs in 2025

Akamai

Ransomware extortion tactics have evolved significantly over recent years

Akamai

Single extortion

Infiltrating businesses with ransomware (encrypting data and demanding ransom for decryption)

Double extortion

Adds the threat of exposing exfiltrated customer information if not paid

Triple extortion

Adds using DDoS attacks to disrupt business operations as extra pressure to force the victim to pay the ransom

Quadruple extortion

Adds the sending of messages to harass business partners, employees, customers, high-level executives, and media to inform them of the breach and pressure the primary victim

Akamai

Ransomware Risk Reduction

What's the Impact?

Financial
Ransoms, fines, and downtime

Data Security and Integrity Risks
Double extortion tactics, compliance violations

Reputation and Trust Erosion
Your own business, but also your third-parties

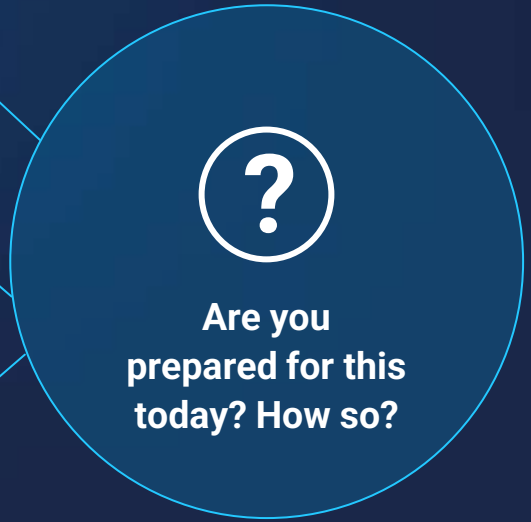

Attacks can happen

Most Frequently Exploited Vulnerabilities

Among the Mandiant incident response investigations performed in 2024, the most frequently exploited vulnerabilities affected security devices, which are, due to their function, typically placed at the edge of the network. Three of the four vulnerabilities were first exploited as zero-days. While a broad selection of threat actors have recently targeted edge devices, Mandiant also specifically noted an increase[3] in targeting from Russian[4] and Chinese[5]cyber espionage actors.

Most Frequently Exploited Vulnerabilities

Initial Infection Vector, 2024

Ransomware-Related

[* M-Trends – 2025 Mandiant Report]

Ransomware Killchain

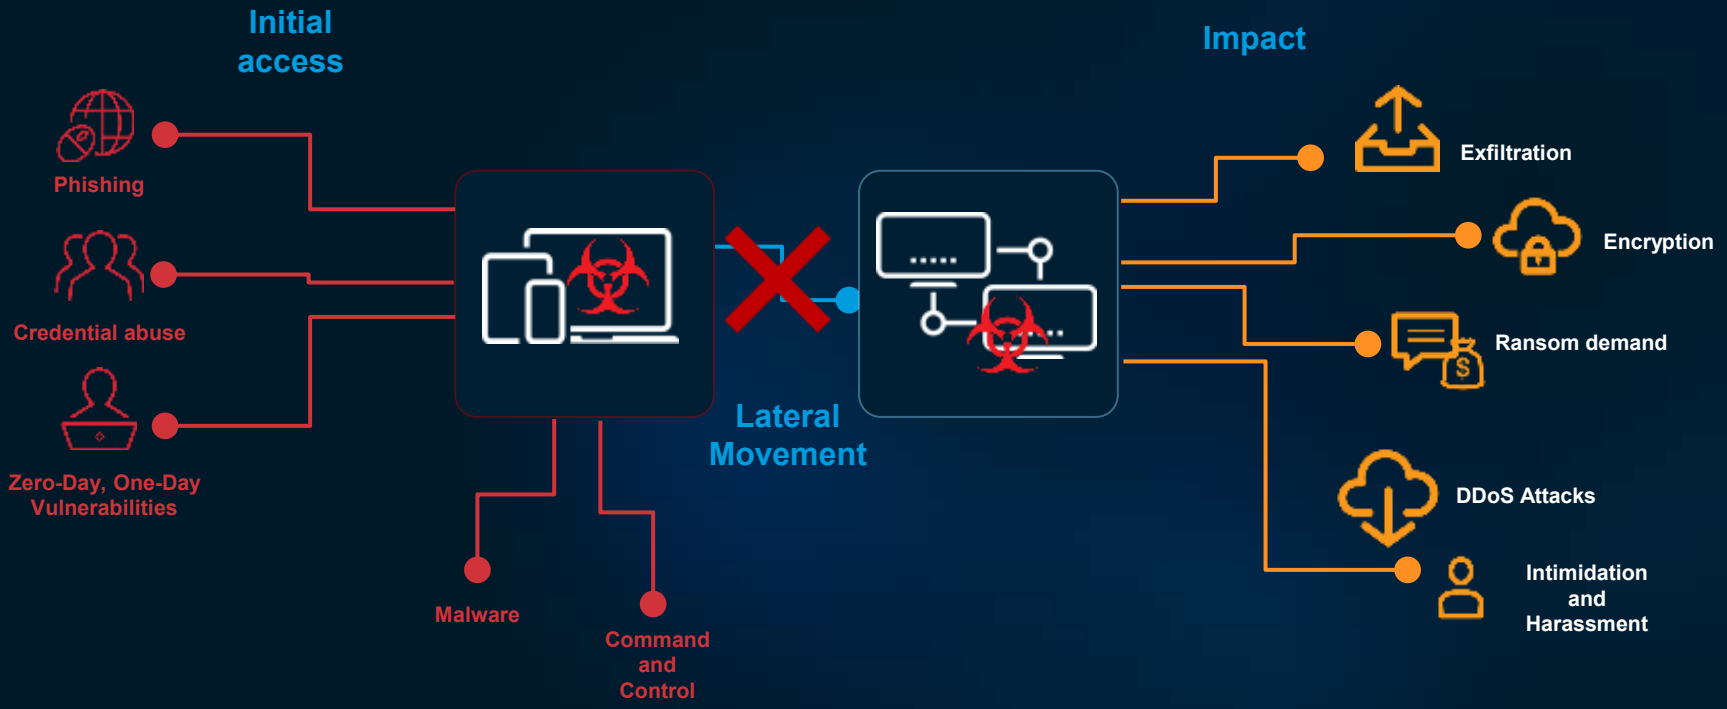

Securing Workloads
in the Mythos Era

May 2026

The half-life of an unpatched flaw just collapsed.

Frontier AI models are now competitive with elite human researchers at finding and exploiting software flaws. The strategic priority shifts from preventing every breach to ensuring the inevitable breach cannot move.

1000s

Zero-day vulnerabilities surfaced by Mythos Preview in its first weeks of testing

75

Bugs Palo Alto Networks found in a single AI-assisted cycle, vs. 5–10 per month previously

2×

Estimated near-term cybersecurity spend increase needed (Bain & Company)

The Mythos moment: a new threat baseline

On 7 April 2026, Anthropic released Claude Mythos Preview under Project Glasswing — a frontier model gated to vetted defensive partners.

Mythos is a signal, not the threat itself.
Comparable cyber capabilities already exist in other frontier models — and more are coming.

What this means in practice:

- Faster vulnerability discovery, including in decades-old code
- Lower attacker skill floor for routine, unhardened systems
- AI chains low-severity flaws into workable exploit paths

THE CAPABILITY SHIFT

"AI capabilities have crossed a threshold that fundamentally changes the urgency required to protect critical infrastructure. The old ways of hardening systems are no longer sufficient."

Project Glasswing partner statement, April 2026

Before frontier AI vs. now

Capability	Pre-Frontier-AI Era	Mythos Era (2026+)
Time to find zero-day in mature code	Weeks to years; expert humans only	**Hours to days; partially autonomous**
Cost per novel exploit	High; bounded by skilled labor	**Dramatically lower**
Volume of usable vulnerabilities surfaced	Trickle, researcher-prioritized	**Flood — thousands per scan run**
Attacker barrier to entry	Significant expertise required	**Lowered for routine systems**
Defender posture required	Patch, monitor, respond	**Assume breach; contain by default**

Source: Anthropic, Bain & Company, World Economic Forum, Axios, May 2026.

Why the perimeter cannot carry this alone

Three structural problems make patch-and-perimeter defense increasingly fragile:

Patch velocity has a ceiling

Even mature teams cannot patch faster than AI-assisted attackers can find new flaws — especially across legacy and OT systems that are difficult or impossible to update.

Detection without containment is half a control

Alerts on lateral movement are useful only if the pathways themselves are constrained. Without segmentation, the alert often arrives after the damage is done.

Hybrid estates multiply trust boundaries

On-prem, cloud, Kubernetes, SaaS, and OT each create lateral pathways. Legacy firewalls cannot see, let alone govern, east-west traffic between them.

Akamai Guardicore: containment-first by design

What it is

A software-based microsegmentation platform that enforces Zero Trust on east-west traffic across hybrid environments. It operates at the process and workload level — not at IPs or VLANs — so policy survives infrastructure change.

Core capabilities

- Granular process-to-process policy enforcement
- Continuous application dependency mapping
- Policy simulation before enforcement
- Multi-Factor Segmentation for sensitive services

Rethink How We Segment

Before: using Traditional Hardware Firewalls (Old Ship)

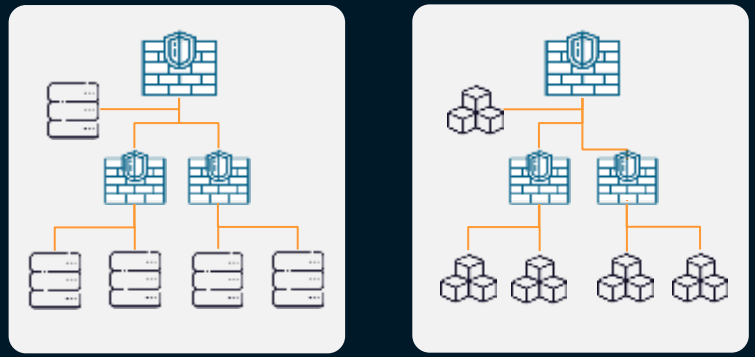

Physical & Virtual (cloud) firewall appliances created network choke points

After: using Akamai Guardicore Segmentation (New Ship)

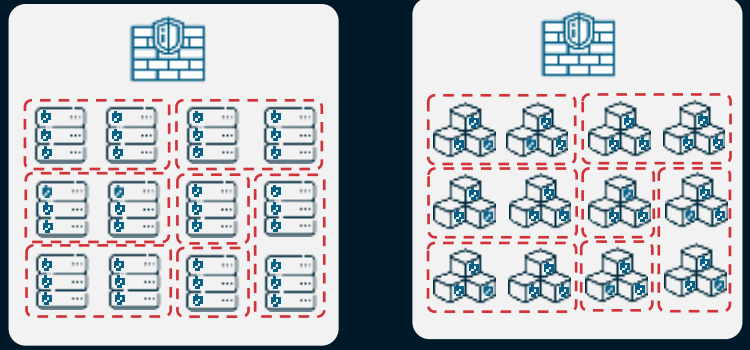

Software-based Agents provide granular control (e.g., Process, Service, User, FQDN)

The AI-powered enhancements that matter

Built from analysis of 500+ real-world segmentation projects.

01

Continuous Discovery

Real-time mapping of workloads, dependencies, and process behavior.

02

AI That Understands Apps

Models infer intent and generate explainable, enforcement-ready policies.

03

Proof-Driven Enforcement

Simulate against live traffic before policies go live — removes the fear of breakage.

04

Continuous Risk Containment

Runtime assurance, exposure analysis, and response unified with segmentation.

05

Delegated App-Owner Workflows

App Owner Portal routes approvals to people with real context.

How We See It

Stop Threat by Controlling Lateral Movement

Gain visibility
to expose blind spots and the context to control

Segment
between apps, workloads, endpoints, and devices

Control access
to apps based on identity

Operationalize at scale

Detect and respond
to threats and changes in security posture over time (breach, zero days, bad hygiene)

Micro-segmentation: Controlling Lateral Movement

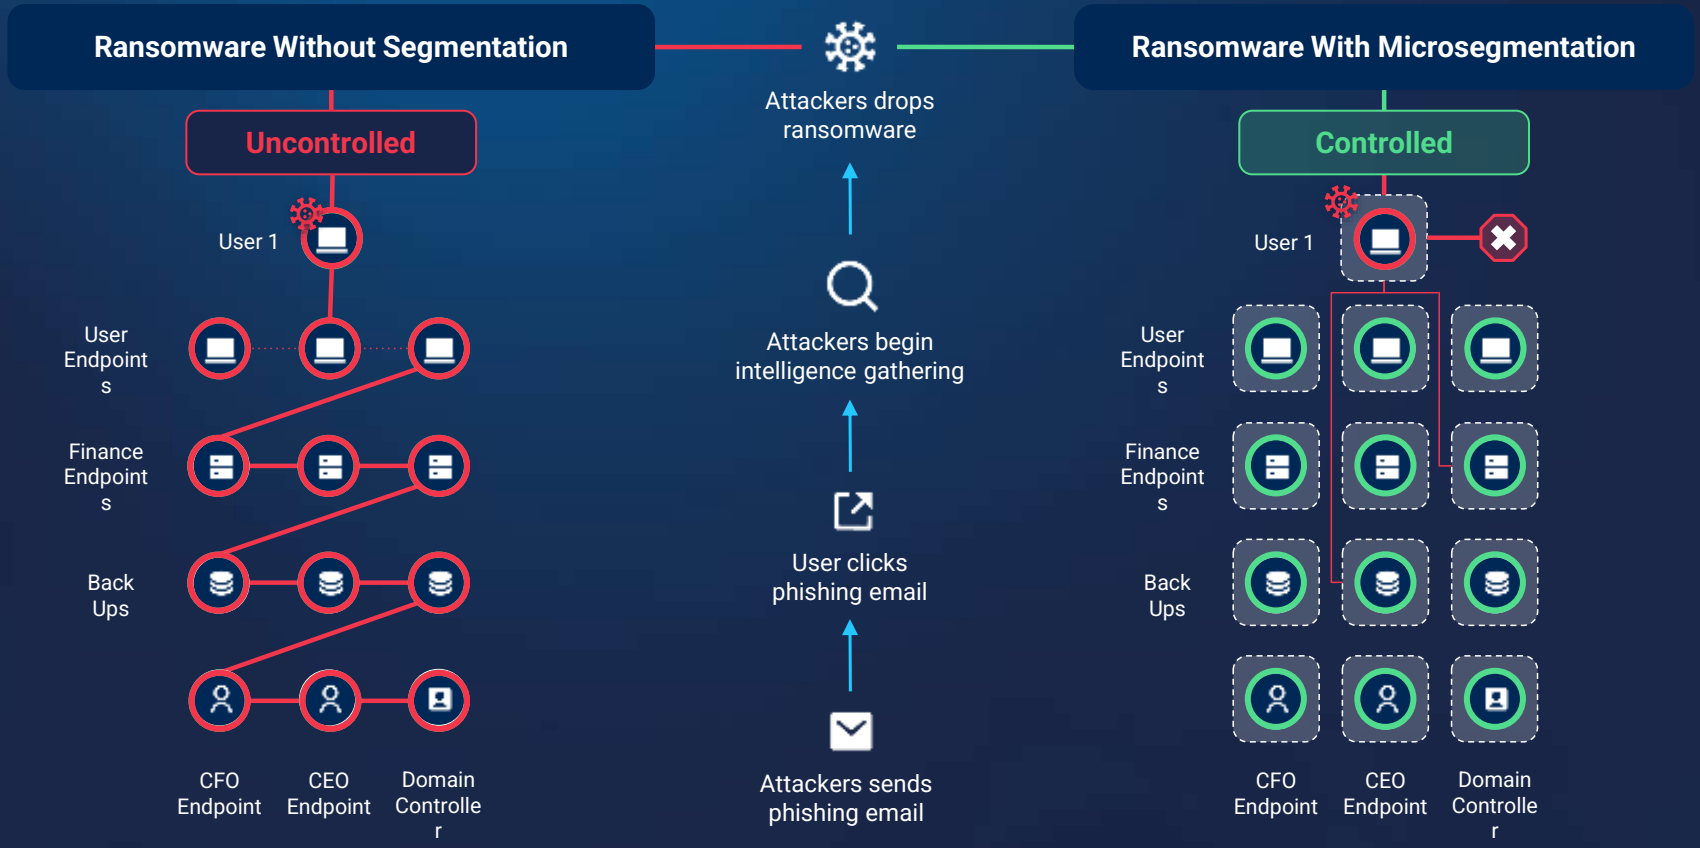

How We Do It: Controlling Lateral Movement

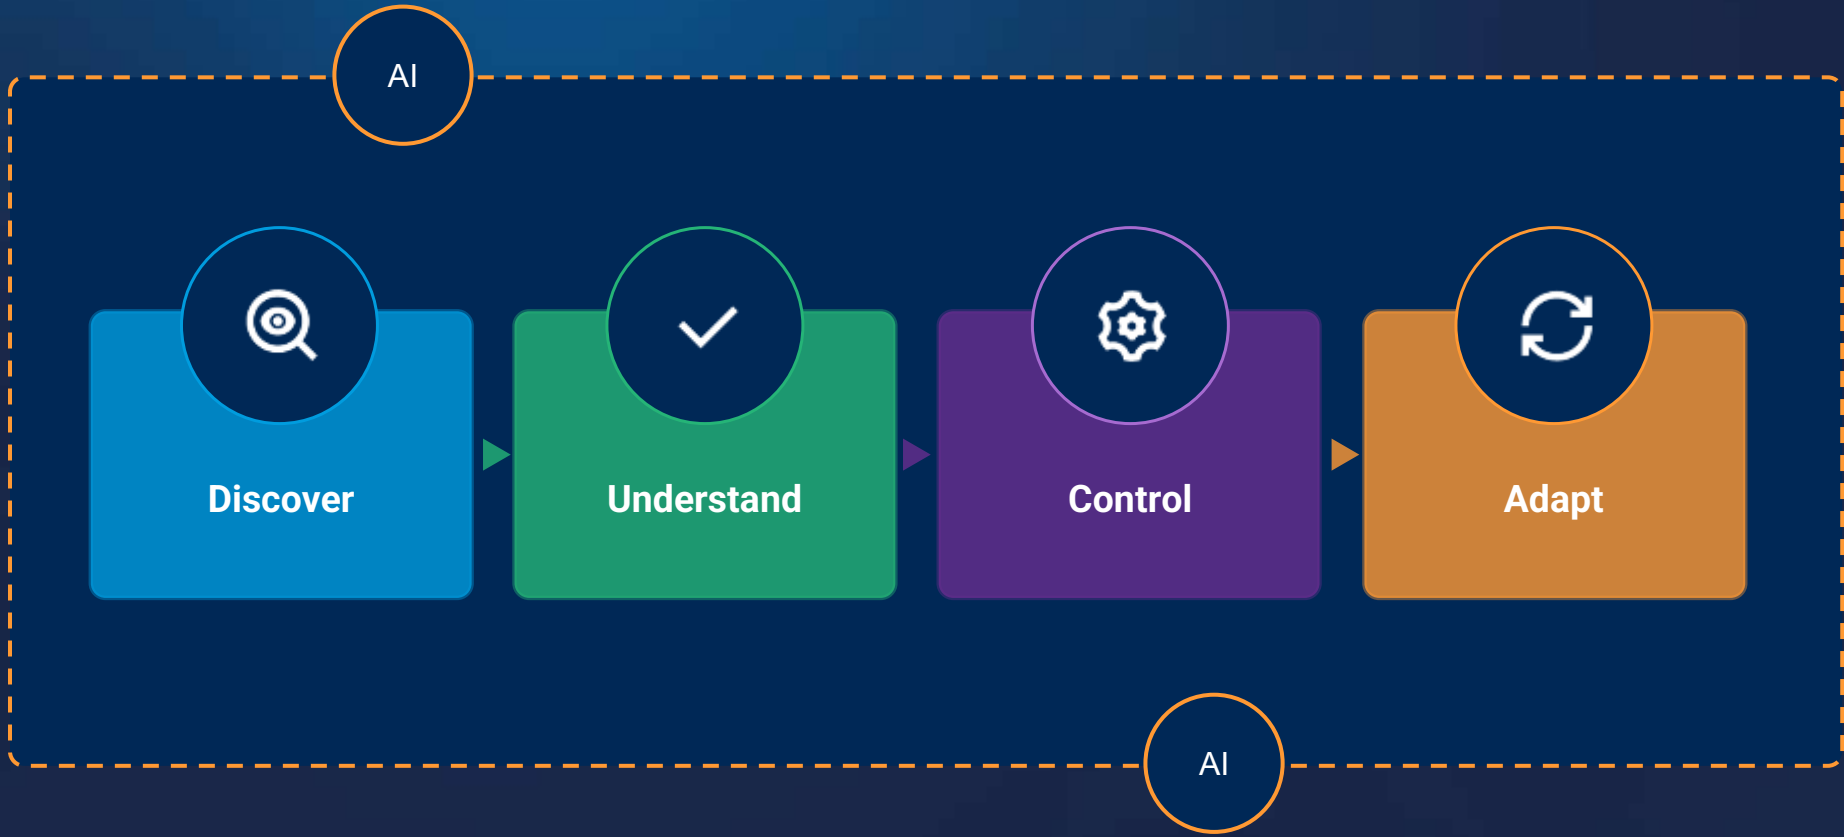

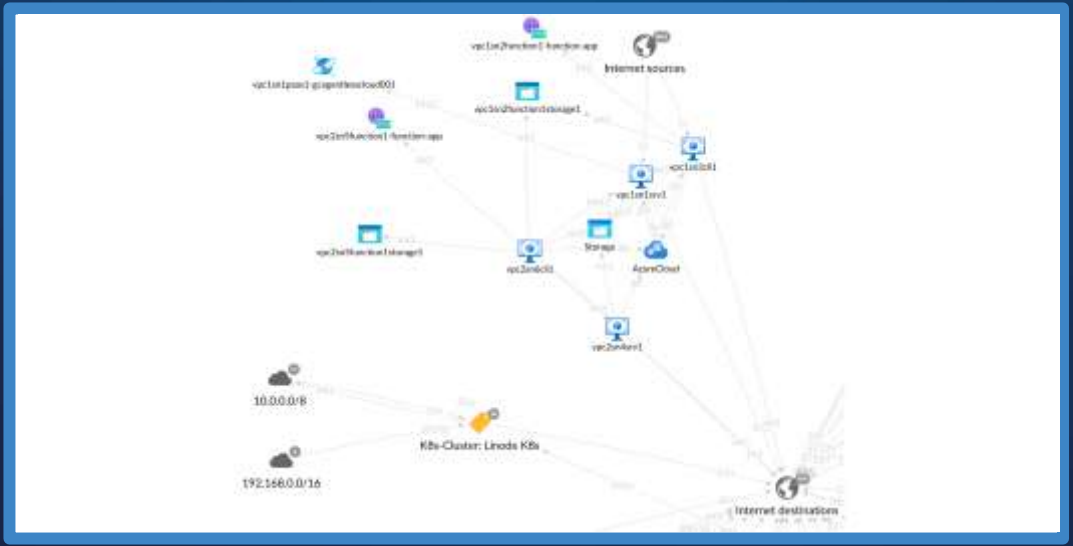

Discover communicating assets across environments

Asset Discovery and Protection includes on-prem data centers, cloud instances, legacy OSes, IoT/OT devices, Kubernetes clusters, and more – without ever having to change consoles.

Full visualization

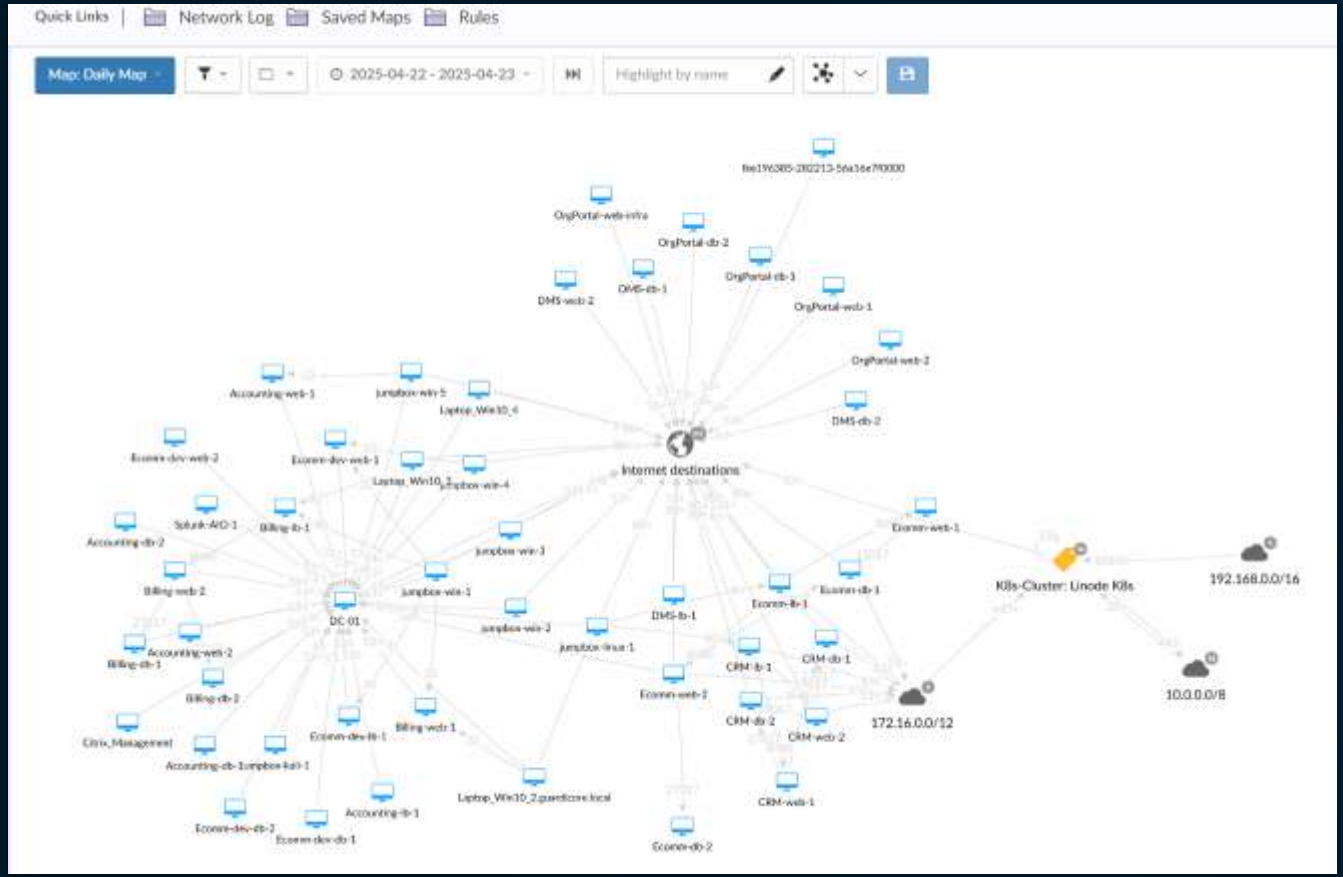

Understand what these connections are and if they're necessary…or suspicious

Map Application Dependencies and Relationships

Process-level visibility allows users to identify servers with similar roles (same tier)

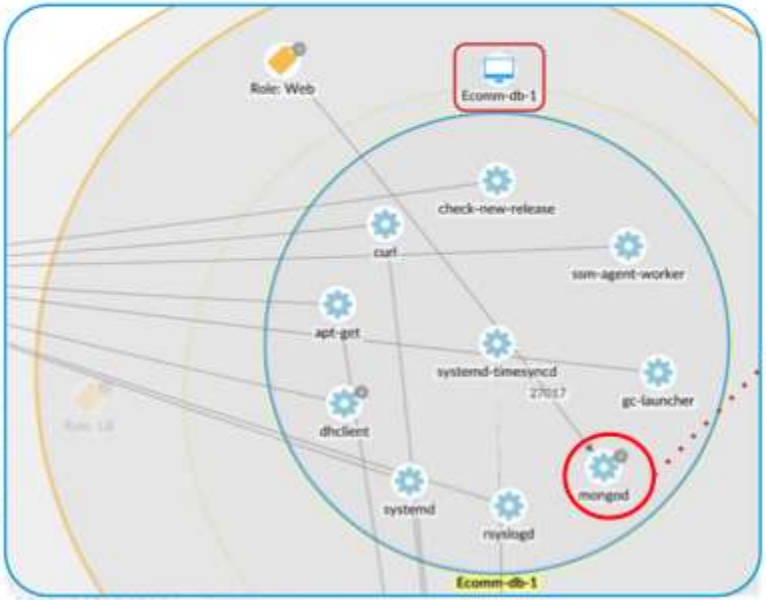

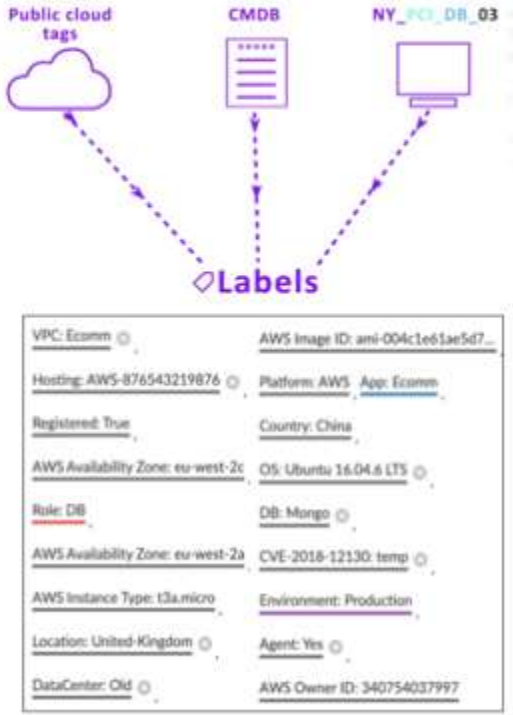

Akamai

AI Labeling for better automation

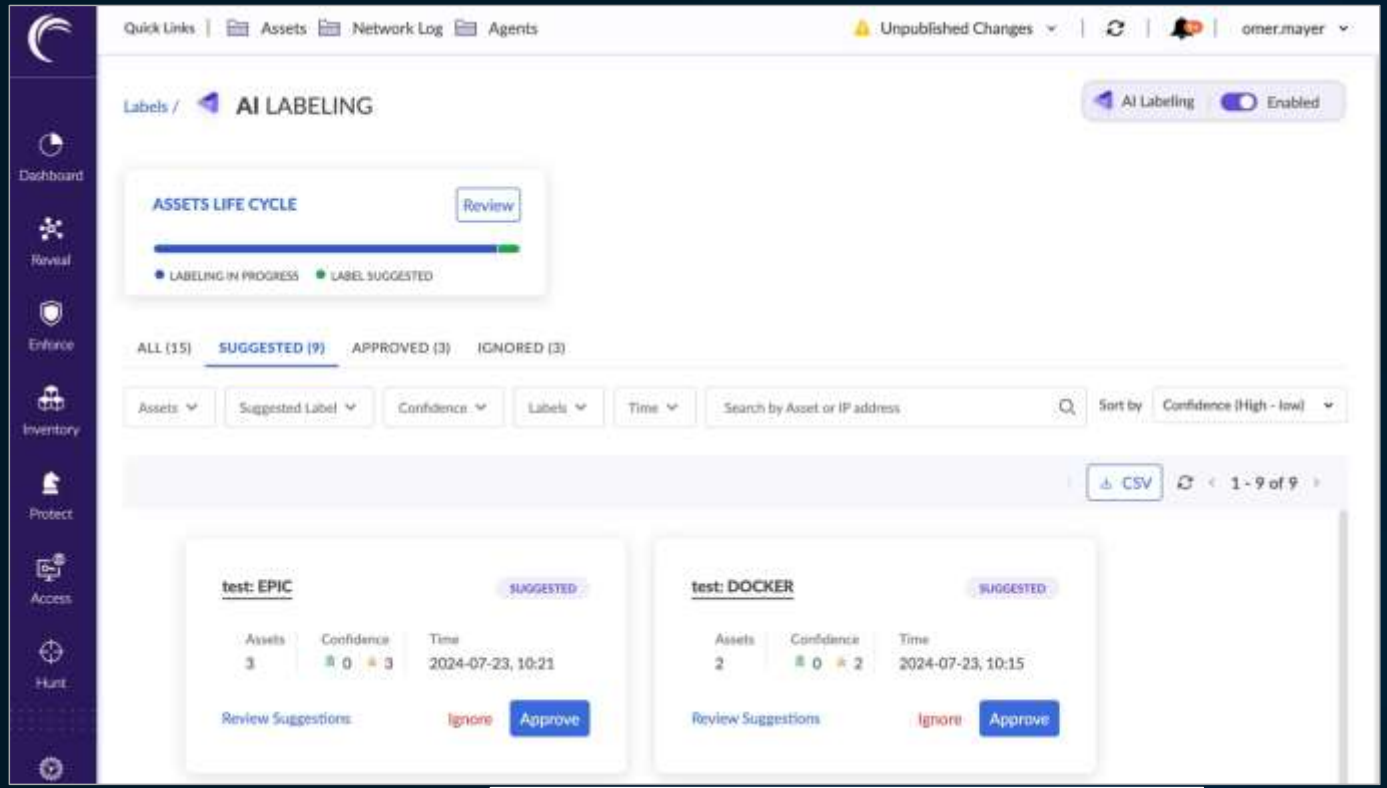

AI Labels Lifecycle Management

Control asset communications with tailored policies

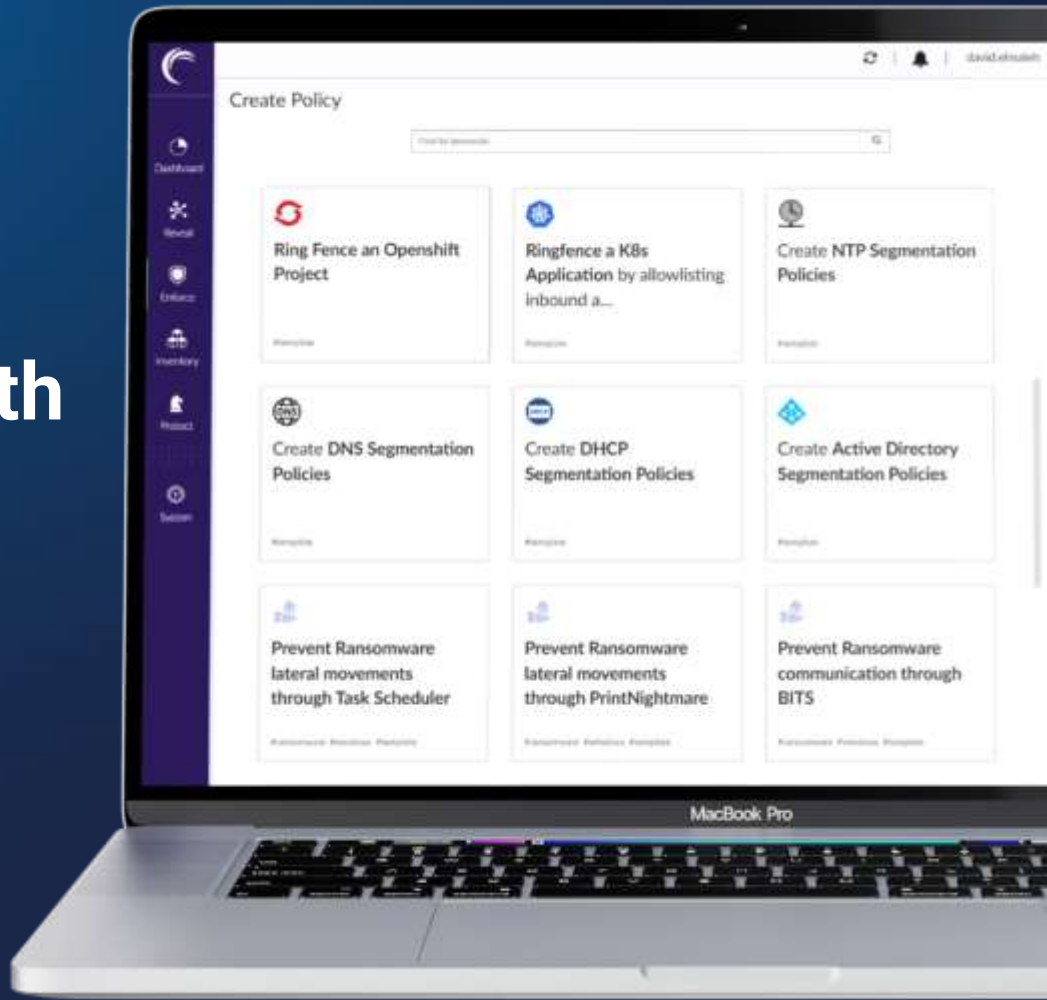

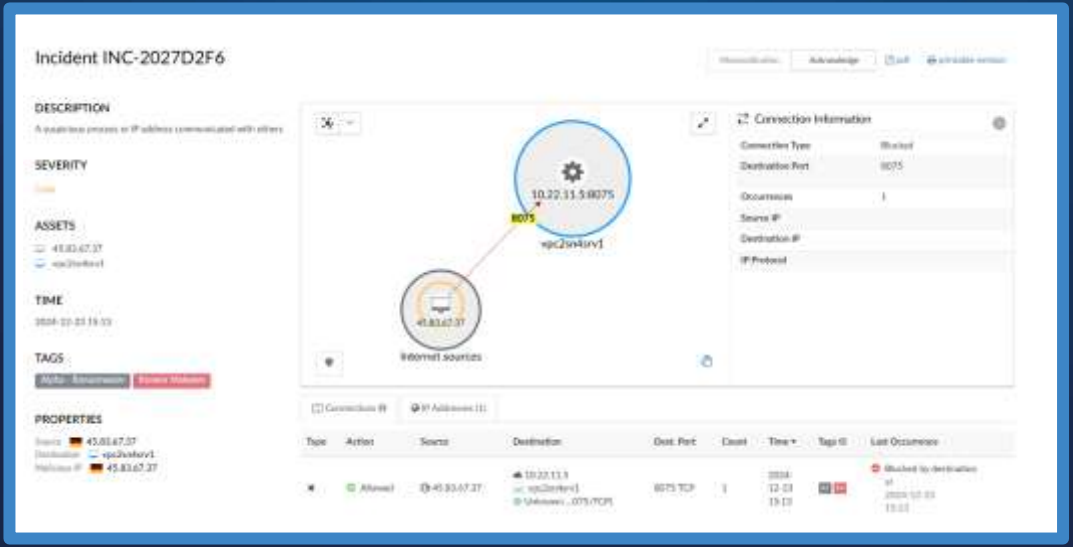

And control potential threats with integrated incident response

Policies: Agnostic to the Infra

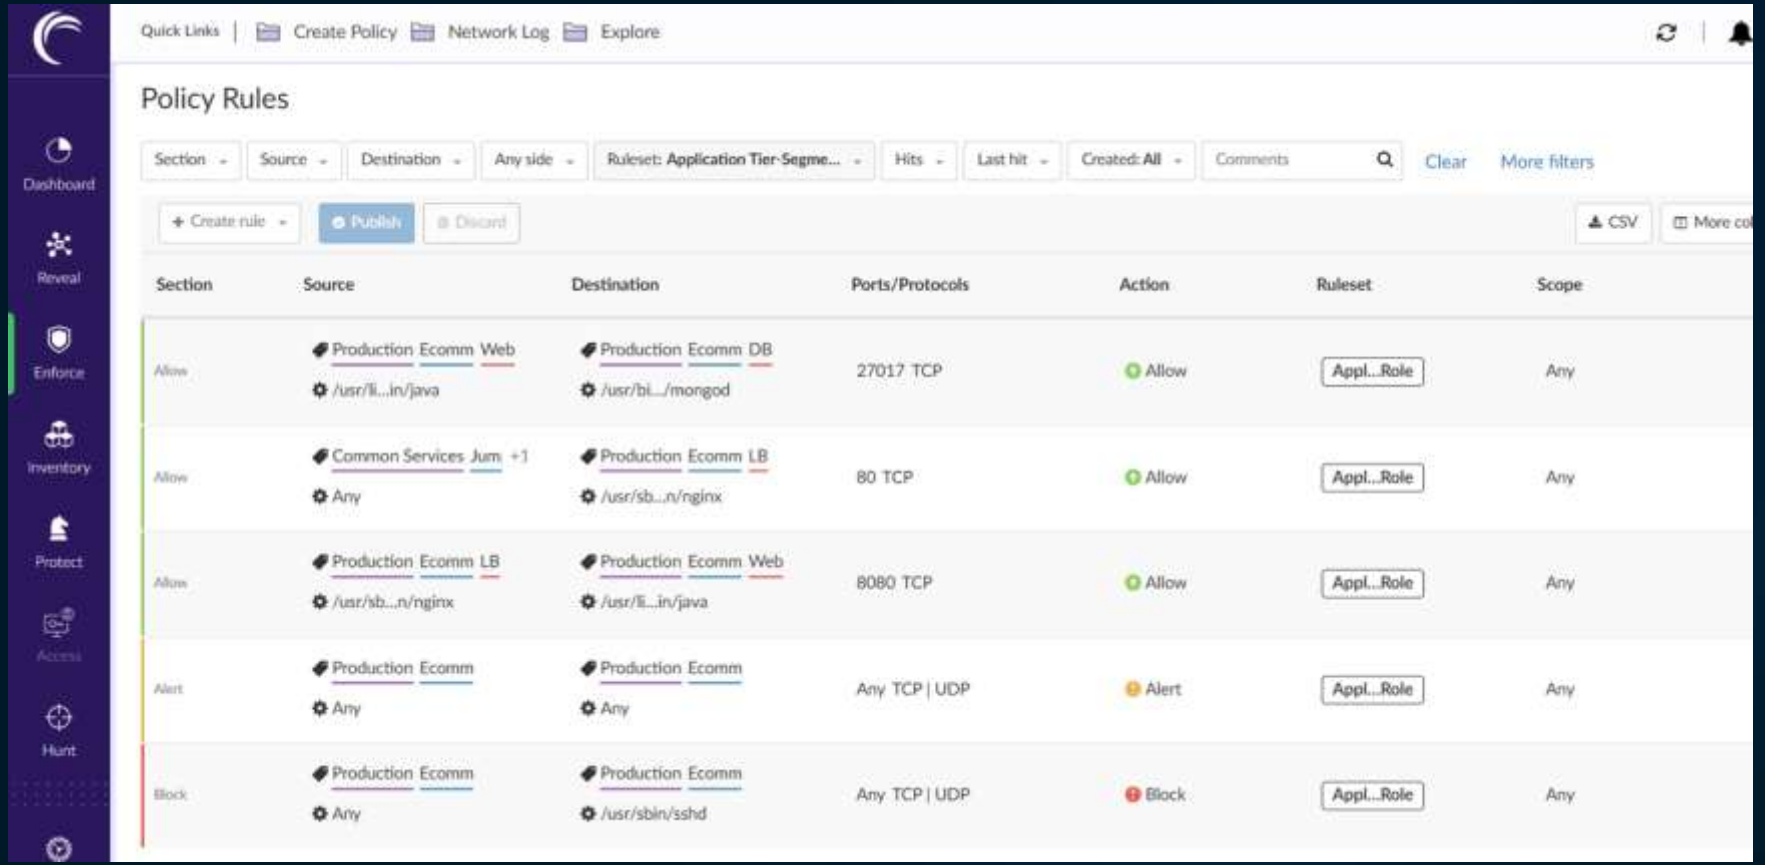

Adapt to the latest threats with guidance from our security experts

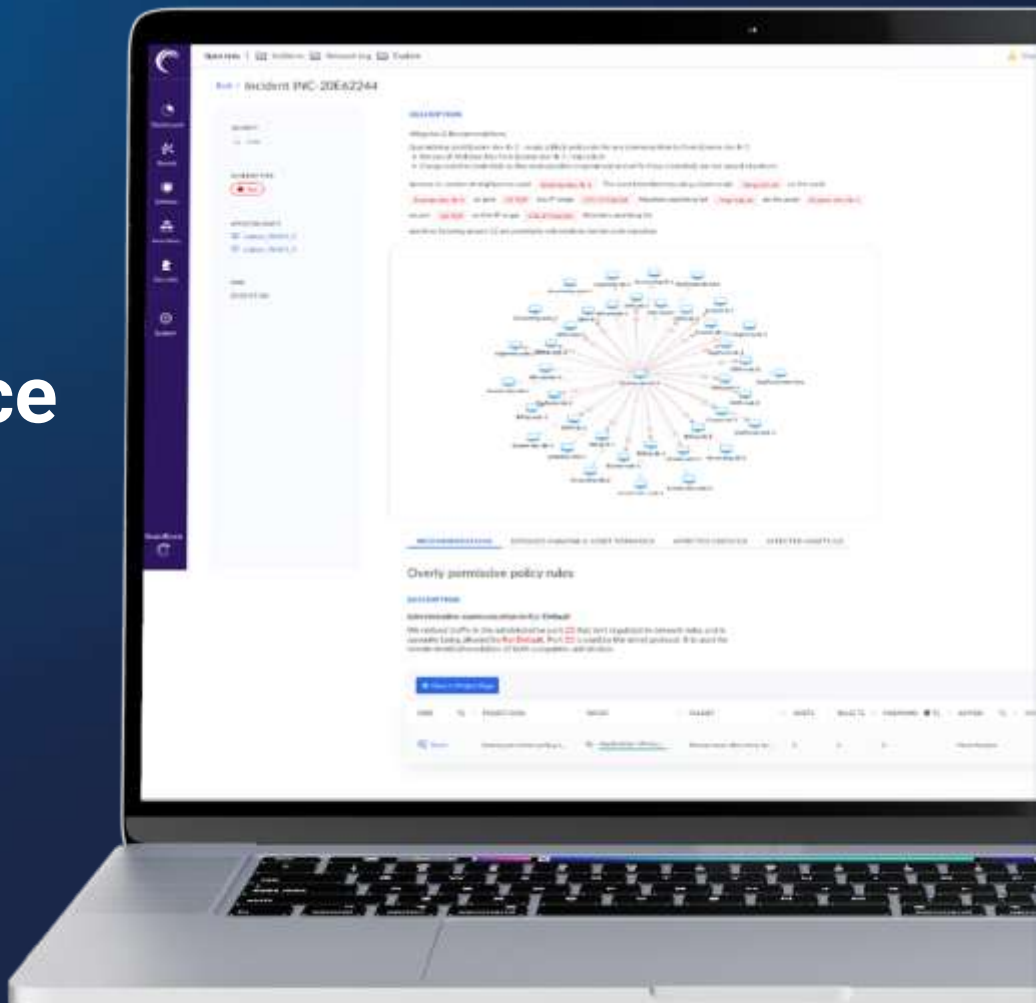

Four defensive properties for the Mythos era

01

Lateral movement is bounded by default

Compromise of one workload cannot reach others without explicit policy permission. Ransomware blast radius shrinks to its smallest viable footprint.

02

Policy written against application intent

Modeling process behavior — not IPs — means policies survive autoscaling, container churn, and cloud migration without leaving gaps.

03

Proof-driven enforcement ends hesitation

The most common failure mode of segmentation is policies that are designed but never enforced. Simulation removes that excuse.

04

One plane across the hybrid estate

Energy, utilities, manufacturing, and other OT-heavy verticals get consistent enforcement where it is most needed.

Strategic adoption: six moves that pay off quickly

1. **Ring-fence the crown jewels first**

   Domain controllers, backups, financial systems, regulated data stores — the assets AI-driven attack chains will target.

2. **Deploy Essential Policies early**

   Pre-built templates close common attack vectors (RDP, SMB, lateral admin) without full app mapping.

3. **Layer Multi-Factor Segmentation**

   On sensitive ports and services, require authenticated identity even when the network path is valid.

4. **Operationalize the App Owner Portal**

   Delegated approvals turn a serial bottleneck into a parallel workflow — the single biggest unlock for enforcement.

5. **Measure containment, not coverage**

   Blast-radius reduction and time-to-contain matter more than agent install count. Make this the board metric.

6. **Bring AI workloads under policy from day 1**

   Inference services, vector stores, and agent frameworks all generate sensitive east-west flows. Don't retrofit later.

Customer Story: Global Systems Integrator with 300,000 protected assets

Eliminate Lateral Movement w/ Speed and at Scale

Challenge

Ransomware impact to business operations

Unable to respond quickly to ransomware attack and slow recovery time after attack

Large Scale deployment requiring 300,000 endpoints and workloads

Board-level initiative with high visibility

Project must be completed within 12 months

Outcome

Deployed 280,000 agents in two weeks and enforced segmentation policies on all within only 8 days

On top, protecting 20,000 workloads across the globe

Solution limited lateral movement at scale

Using automation to apply quarantine policy to assets with suspicious behavior from any asset.

100% of endpoints and workloads in enforcement mode

THE BOTTOM LINE

Assume the breach.
Make sure it cannot move.

The Mythos moment makes one thing concrete: the cost of finding and exploiting unknown software flaws is collapsing. Defenders cannot win on patch velocity alone. The highest-leverage architectural choice now is to ensure the breach that will happen cannot become the breach that makes the news.

Containment > Prevention

Application intent, not IPs

One plane across the estate