upwind

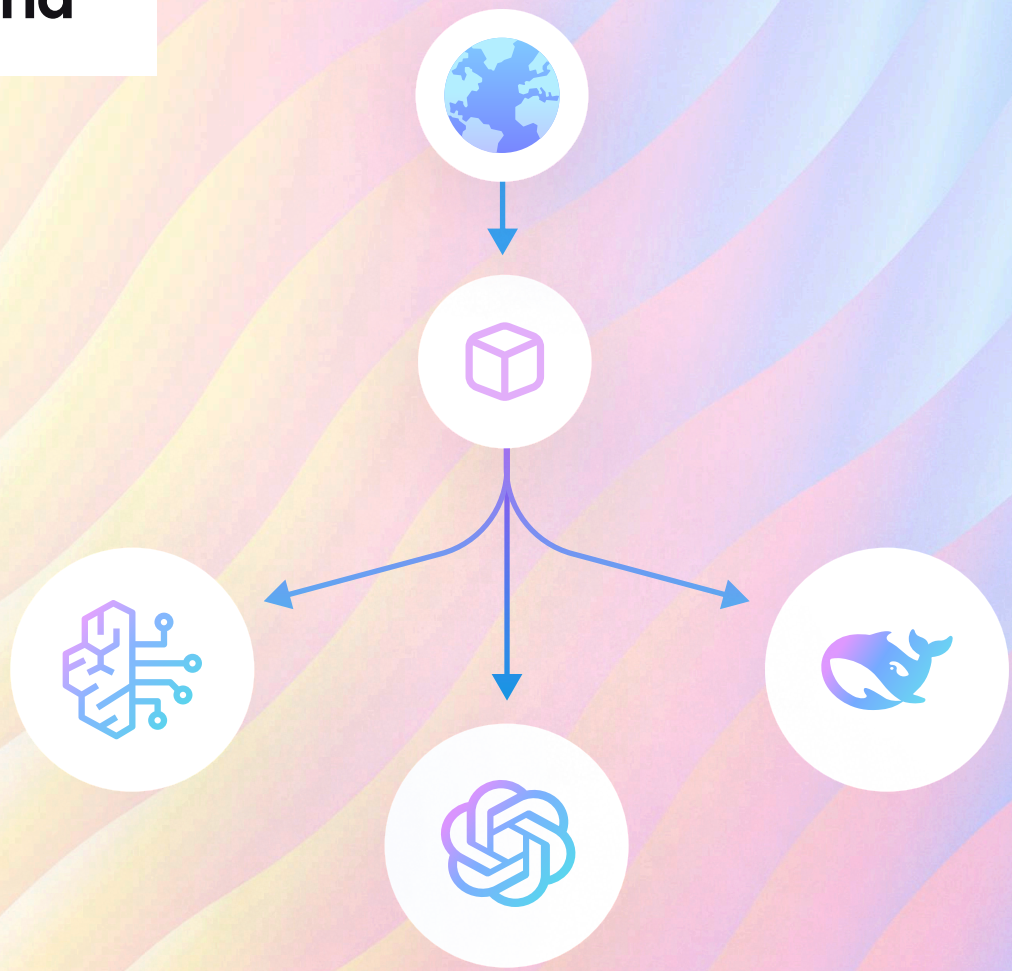

WHITE PAPER

Why Securing AI Workloads Requires Runtime

As enterprises race to adopt GenAI and AI-powered applications, security strategies must evolve. The continued Cloud attack surface expansion and emergence of GenAI are reshaping security from the ground up, and security leaders must adopt an entirely new approach to security cloud infrastructure and applications - one that will detect and respond to the dynamic risks posed by AI workloads.

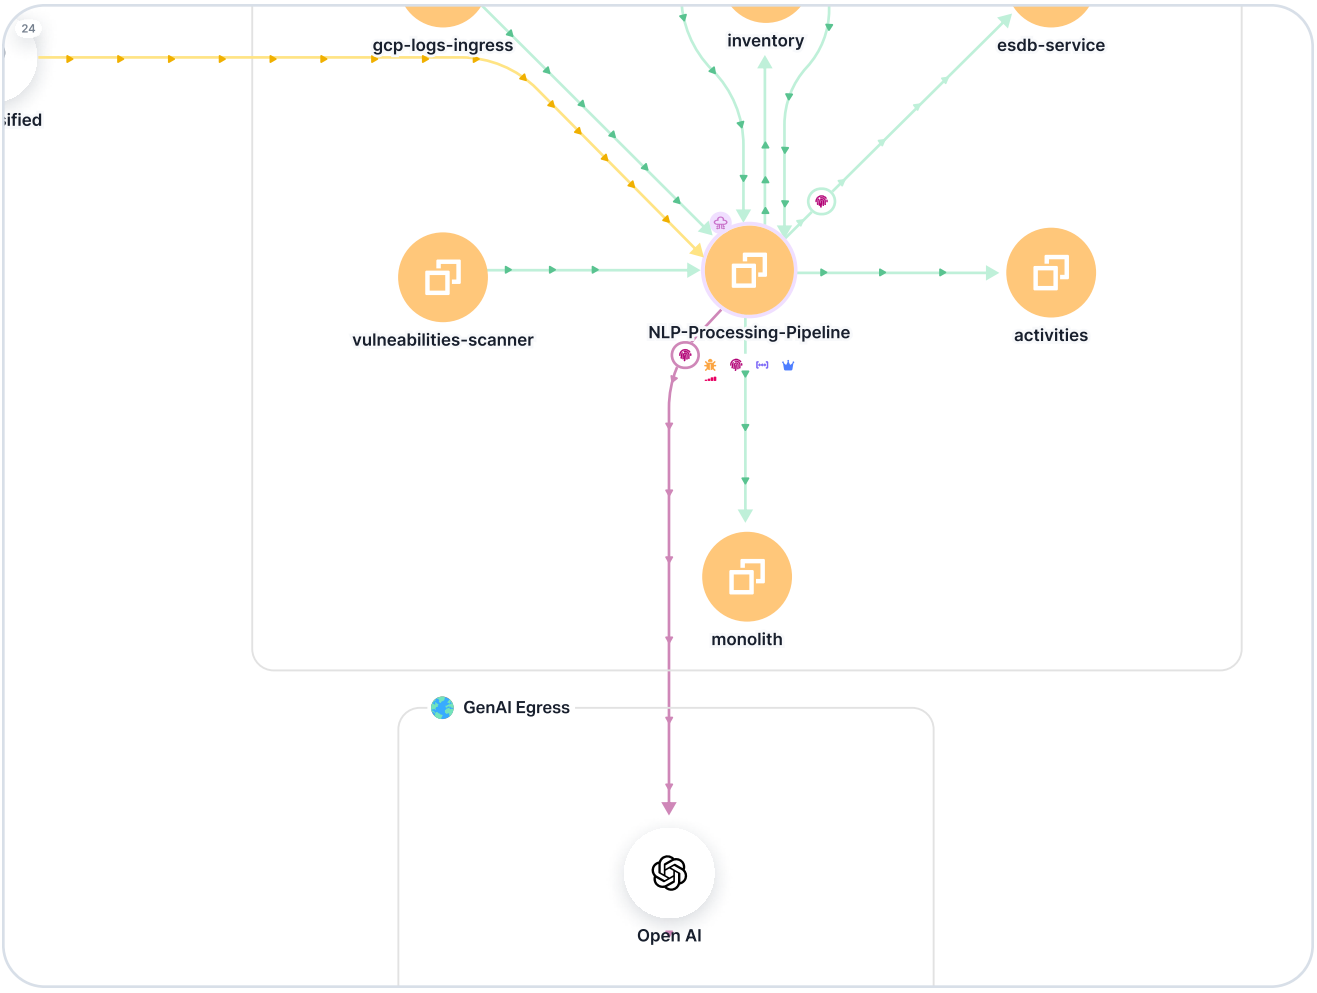

The white paper outlines the emerging challenges in securing AI workloads, and dives into why an entirely new, dynamic approach is required in order to stay ahead of increasingly sophisticated attackers and the complex AI attack surface.

The Evolving Threat Landscape of AI Workloads

The adoption of AI introduces new security challenges that extend beyond traditional infrastructure concerns. Threat actors are already targeting GenAI APIs, manipulating models, and exploiting runtime behavior.

Emerging AI Threats Include:

Data Leakage

AI models process and return sensitive data, increasing the likelihood of exposure without real-time enforcement.

Prompt Injection and Model Manipulation

Crafted inputs can override model logic, extract hidden data, or generate harmful outputs.

Unauthorized API Usage

AI models process and return sensitive data, increasing the likelihood of exposure without real-time enforcement.

Abuse of Open-Source Models

Integrating unvetted AI packages can introduce hidden vulnerabilities or malicious behavior.

AI workloads are dynamic and ephemeral by nature. These characteristics demand continuous monitoring, behavioral analysis, and active control during execution.

AI Security Requires Context from Layers 3, 4, and 7

True AI workload protection requires insight into behavior across all critical network and application layers. Security cannot depend solely on what code is written—it must understand how that code runs, where it communicates, and what it does during execution.

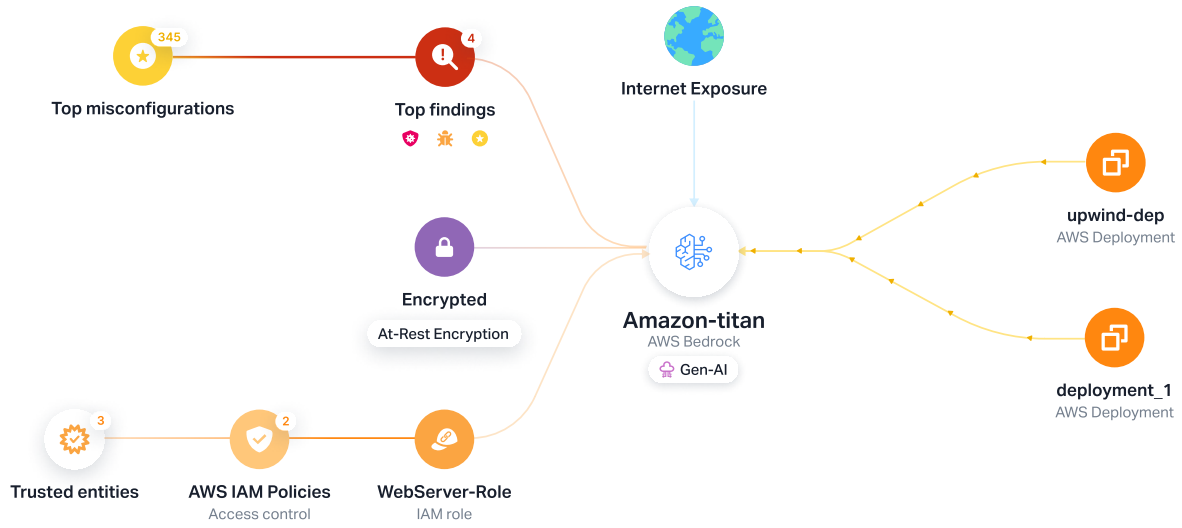

Layered Runtime Telemetry

Layer	What It Reveals	Importance to AI Security
3 Network	IP-level traffic, routing, and connectivity	Identifies unauthorized external communication and data exfiltration attempts
4 Transport	Protocols, ports, and session patterns	Flags anomalies like beaconing, lateral movement, or reverse shells
7 Application	API calls, payloads, request behavior	Detects misuse of AI APIs, prompt injection, and access violations

Runtime context from these layers enables correlation between network behavior and application intent. This is essential for identifying misuse of AI services, understanding the risk of each interaction, and blocking active threats before they escalate.

Purpose-Built Runtime Security for AI Workloads

Upwind provides a runtime-native security platform tailored to cloud-native and AI environments. It continuously monitors workloads across the entire execution lifecycle and correlates runtime telemetry with CI/CD and infrastructure context.

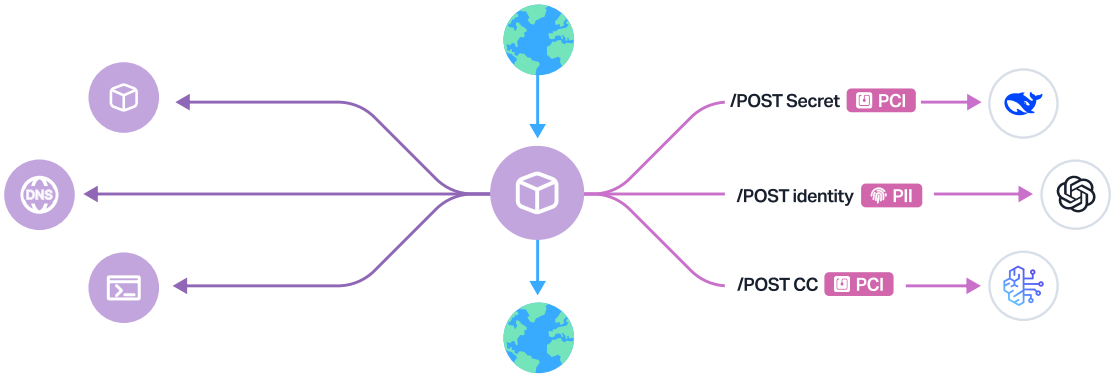

Upwind's Dynamic AI Security Includes:

Live Runtime Threat Detection

Catch attacks as they happen — identify privilege escalations, unauthorized API activity, and data exfiltration in real time.

Deep Context-Aware Security Insights

Correlate threats with workload identity, image versions, network posture, and developer activity for high-fidelity alerts.

Full-Stack Traffic and Behavior Monitoring

Continuously observe Layers 3, 4, and 7 to uncover anomalous system behavior, suspicious traffic, and risky app logic.

Code-to-Cloud CI/CD Risk Mapping

Connect runtime vulnerabilities back to specific code changes for faster root cause analysis and remediation.

AI Usage and Policy Enforcement

Reveal which services use GenAI, what data is shared, and ensure compliance with enterprise AI governance policies.

Upwind helps security teams focus on what's real—actual exposure, actual behavior, and actual threat impact.

Case Study: Preventing AI API Misuse in Production

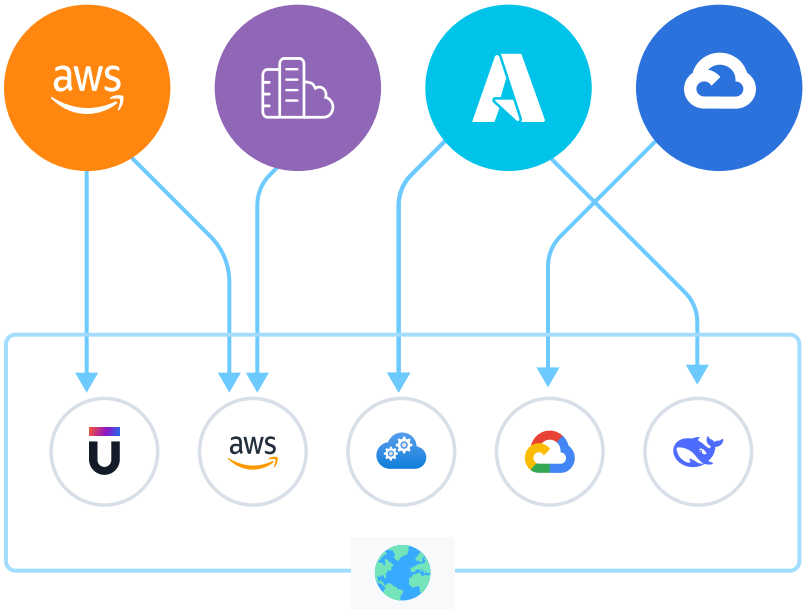

A global software company using Upwind detected abnormal outbound traffic from a containerized microservice. Upwind flagged unsanctioned connections to a third-party LLM API originating from a newly deployed model-assist feature.

With Layer 7 analysis, Upwind was able to identify the unintended traffic and note that sensitive data was being sent to the LLM API. The team blocked the request in real time - preventing data leakage and unauthorized use of AI resources.

This was caught by employing the exact strategies outlined in this paper - correlating real-time traffic and and monitoring for abnormal behaviors at runtime.

Conclusion

Upwind defends against advanced threats with AI-specific threat policies, real-time AI communication path mapping, and sensitive data discovery - offering deep visibility into data egress patterns to minimize exposure risks.

Secure Your AI Workloads with Upwind

Upwind delivers full-stack AI security by operationalizing these best practices:

- Real-time visualization of AI communication paths across layers 3, 4, and 7
- Sensitive data discovery using AI and regex-based payload inspection
- Threat detection with dynamic baselines and AI-focused policies
- Customizable CSPM rules for AI-centric cloud posture hardening

AI Inventory

AI Runtime activity

AI Threat detection

AI Security Posture

Want to know more about Upwind's AI security solution? Visit www.upwind.io or send us a note at hello@upwind.io to schedule a brief demo and see real-time security in action.