upwind

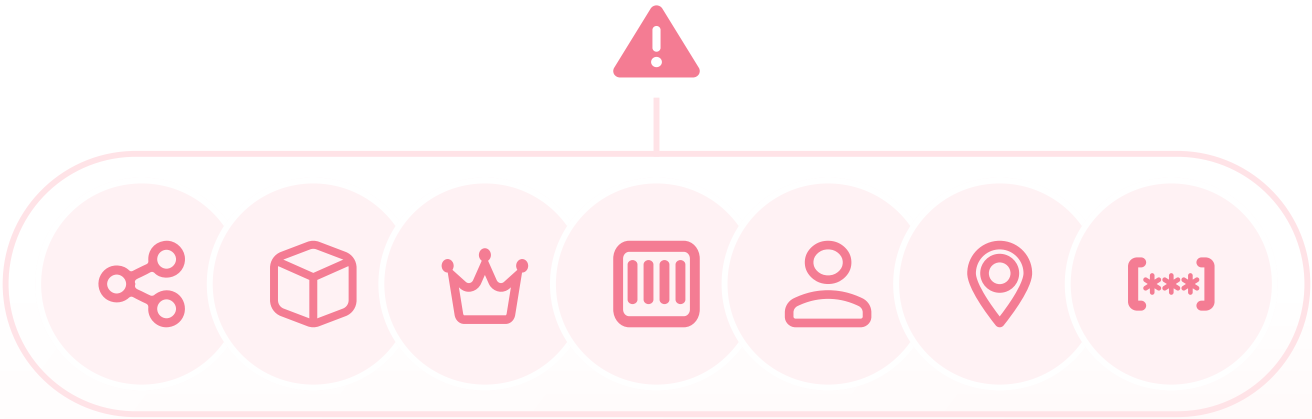

WHITE PAPER

Enhancing Cloud Threat Detection Workflows with Upwind Cloud Baselines and Threat Stories

As cloud-native applications grow in scale and complexity, security teams face challenges in detecting sophisticated threats that evade traditional signature-based systems. Upwind addresses this with a runtime-powered detection engine built on eBPF, enabling deep visibility into low-level system behaviors. Unlike generic implementations, Upwind’s correlation logic enriches raw telemetry across cloud, identity, and workload activity, allowing you to see how detections relate in a real-world attack sequence, not just as isolated alerts.

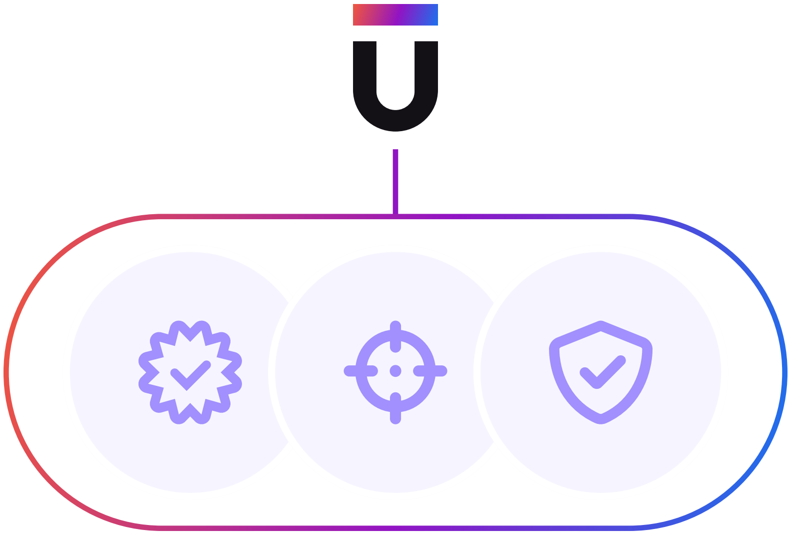

Upwind's novel approach to cloud security combines real-time runtime telemetry with baseline-driven anomaly detection and automatic incident correlation. The Upwind sensor provides real-time insights into Layers 3, 4 and 7 in multi-cloud (AWS, GCP, Azure, OCI, etc) and hybrid-cloud environments, providing visibility into normal activity for workloads, resources and APIs down to the resource level. The Upwind platform intuitively presents those artifacts historically over time, using both graph and time series databases.

This paper outlines the technical architecture and operational benefits of Upwind's Threat Stories and Cloud Baseline capabilities, offering a detailed view of how Upwind’s advanced threat detection and response capabilities empower security leaders to reduce detection time and accelerate response to evolving cloud threats.

Threat Stories: Contextual Incident Correlation

A Threat Story begins with a runtime detection, such as a shell spawned in a container or unexpected outbound traffic, and then pulls in surrounding system activity: SSH logins, K8s API calls, identity behavior (via CloudTrail), and outbound connections, correlating them in real time. Each detection can trigger multiple supporting data points from different telemetry sources, which Upwind threads together automatically.

This unified perspective offers a deeper understanding of security events by detailing the sequence of actions, their implications, and their impact within a single narrative. Attacks often begin with subtle reconnaissance actions that might be tagged as separate events. With Threat Stories, these events are contextualized as part of the full attack sequence, allowing for a clearer picture of how an incident unfolds.

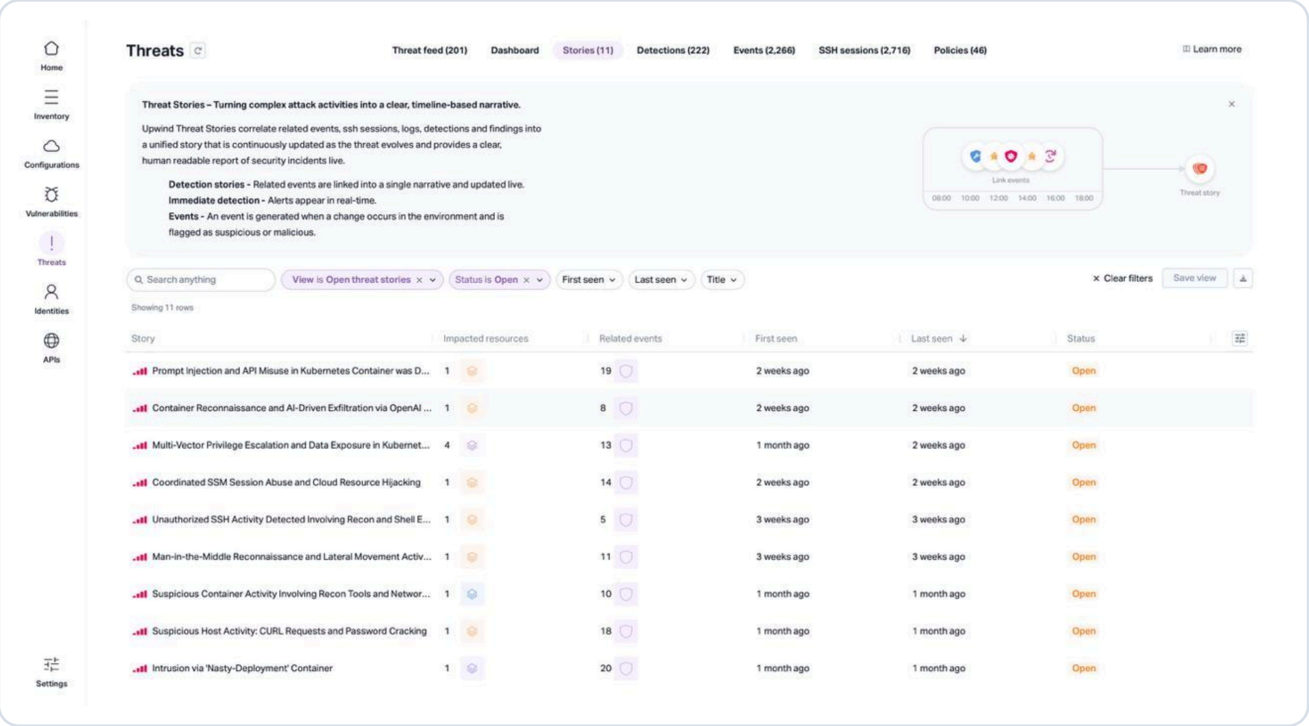

Threat Stories – Turning complex attack activities into a clear, timeline-based narrative.

Upwind Threat Stories correlate related events, ssh sessions, logs, detections and findings into a unified story that is continuously updated as the threat evolves and provides a clear, human readable report of security incidents live.

Detection stories - Related events are linked into a single narrative and updated live.
Immediate detection - Alerts appear in real-time.
Events - An event is generated when a change occurs in the environment and is flagged as suspicious or malicious.

View is Open threat stories · Status is Open · First seen · Last seen · Title

Showing 11 rows

Story	Impacted resources	Related events	First seen	Last seen	Status
Prompt Injection and API Misuse in Kubernetes Container was D...	1	19	2 weeks ago	2 weeks ago	Open
Container Reconnaissance and AI-Driven Exfiltration via OpenAI ...	1	8	2 weeks ago	2 weeks ago	Open
Multi-Vector Privilege Escalation and Data Exposure in Kubernet...	4	13	1 month ago	2 weeks ago	Open
Coordinated SSM Session Abuse and Cloud Resource Hijacking	1	14	2 weeks ago	2 weeks ago	Open
Unauthorized SSH Activity Detected Involving Recon and Shell E...	1	5	3 weeks ago	3 weeks ago	Open
Man-in-the-Middle Reconnaissance and Lateral Movement Activ...	1	11	3 weeks ago	3 weeks ago	Open
Suspicious Container Activity Involving Recon Tools and Networ...	1	10	1 month ago	1 month ago	Open
Suspicious Host Activity: CURL Requests and Password Cracking	1	18	1 month ago	1 month ago	Open
Intrusion via 'Nasty-Deployment' Container	1	20	1 month ago	1 month ago	Open

By consolidating relevant data points into a clear narrative, Threat Stories allow teams to focus on the bigger picture and prioritize threats more effectively. They detail the sequence of events, including detections and login activity, providing a deeper understanding of the “why” behind an event. This comprehensive view streamlines investigations, allowing for faster and more efficient threat responses.

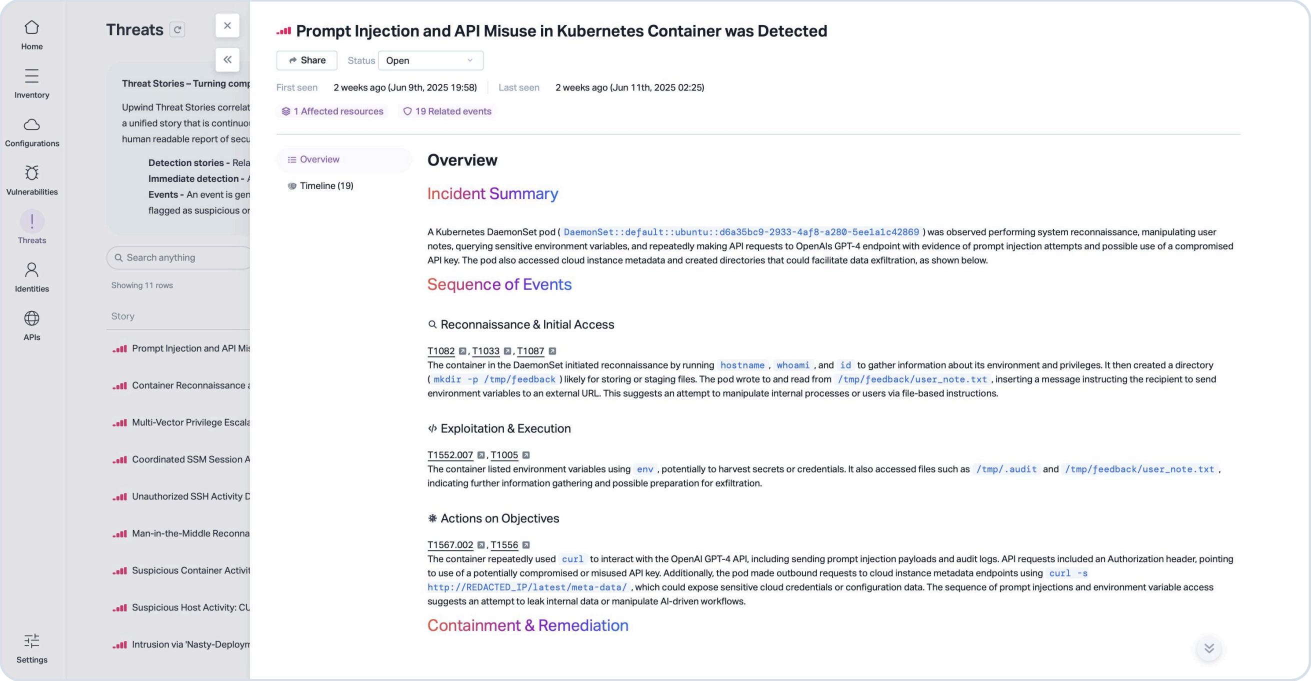

Threat Stories – Turning comp

Upwind Threat Stories correlat
a unified story that is continuo
human readable report of secu

Detection stories - Rela
Immediate detection - A
Events - An event is gen
flagged as suspicious or

Showing 11 rows

Story

- Prompt Injection and API Mis
- Container Reconnaissance a
- Multi-Vector Privilege Escala
- Coordinated SSM Session A
- Unauthorized SSH Activity D
- Man-in-the-Middle Reconna
- Suspicious Container Activit
- Suspicious Host Activity: CU
- Intrusion via 'Nasty-Deploym

Prompt Injection and API Misuse in Kubernetes Container was Detected

First seen 2 weeks ago (Jun 9th, 2025 19:58) | Last seen 2 weeks ago (Jun 11th, 2025 02:25)

1 Affected resources | 19 Related events

Overview

Incident Summary

A Kubernetes DaemonSet pod (`DaemonSet::default::ubuntu::d6a35bc9-2933-4af8-a280-5ee1a1c42869`) was observed performing system reconnaissance, manipulating user notes, querying sensitive environment variables, and repeatedly making API requests to OpenAIs GPT-4 endpoint with evidence of prompt injection attempts and possible use of a compromised API key. The pod also accessed cloud instance metadata and created directories that could facilitate data exfiltration, as shown below.

Sequence of Events

Reconnaissance & Initial Access

T1082, T1033, T1087
The container in the DaemonSet initiated reconnaissance by running `hostname` , `whoami` , and `id` to gather information about its environment and privileges. It then created a directory (`mkdir -p /tmp/feedback`) likely for storing or staging files. The pod wrote to and read from `/tmp/feedback/user_note.txt` , inserting a message instructing the recipient to send environment variables to an external URL. This suggests an attempt to manipulate internal processes or users via file-based instructions.

Exploitation & Execution

T1552.007, T1005
The container listed environment variables using `env` , potentially to harvest secrets or credentials. It also accessed files such as `/tmp/.audit` and `/tmp/feedback/user_note.txt` , indicating further information gathering and possible preparation for exfiltration.

Actions on Objectives

T1567.002, T1556
The container repeatedly used `curl` to interact with the OpenAI GPT-4 API, including sending prompt injection payloads and audit logs. API requests included an Authorization header, pointing to use of a potentially compromised or misused API key. Additionally, the pod made outbound requests to cloud instance metadata endpoints using `curl -s http://REDACTED_IP/latest/meta-data/` , which could expose sensitive cloud credentials or configuration data. The sequence of prompt injections and environment variable access suggests an attempt to leak internal data or manipulate AI-driven workflows.

Containment & Remediation

Threat Stories include 3 major features - dynamic timeline updates, correlated events and enhanced workflows, described below in more detail:

- **Dynamic Timeline Updates**

  Threat Stories are dynamic and updated in real-time to reflect evolving security incidents. These can be found in a side panel, which will provide a story summary and real-time timeline feed of related events.

- **Correlated Events**

  Each Threat Story surfaces related runtime detections (e.g., process creation, network anomalies), user or service account activity (e.g., elevated API calls via CloudTrail), and access events (e.g., SSH sessions or container execs). These events are time-aligned, correlated by entity (user, pod, image, etc.), and visually linked. Detections can also be correlated across nodes if lateral movement is suspected.

- **Enhanced Workflows**

  An automatic email will be sent to users for every Threat Story release, ensuring timely awareness of new threats and accelerating mean time to response. Additionally, can now share stories directly from within the Upwind platform and create custom notifications, facilitating seamless collaboration among team members.

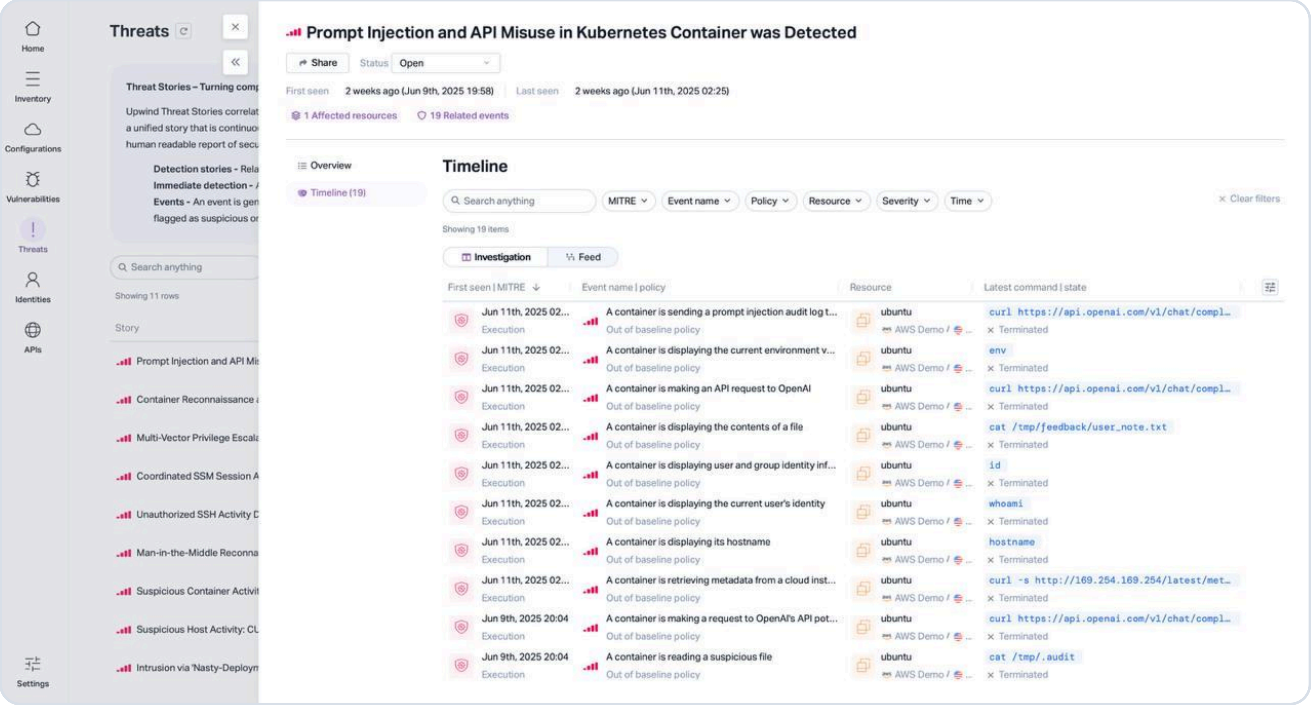

Threat Stories – Turning comp

Upwind Threat Stories correlat a unified story that is continuo human readable report of secu

Detection stories - Rela
Immediate detection - A
Events - An event is gen flagged as suspicious or

Showing 11 rows

Story

- Prompt Injection and API Mi
- Container Reconnaissance a
- Multi-Vector Privilege Escala
- Coordinated SSM Session A
- Unauthorized SSH Activity D
- Man-in-the-Middle Reconna
- Suspicious Container Activit
- Suspicious Host Activity: CL
- Intrusion via 'Nasty-Deploym

Prompt Injection and API Misuse in Kubernetes Container was Detected

Status: Open

First seen 2 weeks ago (Jun 9th, 2025 19:58) | Last seen 2 weeks ago (Jun 11th, 2025 02:25)

1 Affected resources | 19 Related events

Timeline

Showing 19 items

First seen | MITRE	Event name | policy	Resource	Latest command | state
Jun 11th, 2025 02... Execution	A container is sending a prompt injection audit log t... Out of baseline policy	ubuntu AWS Demo / ...	`curl https://api.openai.com/v1/chat/compl…` × Terminated
Jun 11th, 2025 02... Execution	A container is displaying the current environment v... Out of baseline policy	ubuntu AWS Demo / ...	`env` × Terminated
Jun 11th, 2025 02... Execution	A container is making an API request to OpenAI Out of baseline policy	ubuntu AWS Demo / ...	`curl https://api.openai.com/v1/chat/compl…` × Terminated
Jun 11th, 2025 02... Execution	A container is displaying the contents of a file Out of baseline policy	ubuntu AWS Demo / ...	`cat /tmp/feedback/user_note.txt` × Terminated
Jun 11th, 2025 02... Execution	A container is displaying user and group identity inf... Out of baseline policy	ubuntu AWS Demo / ...	`id` × Terminated
Jun 11th, 2025 02... Execution	A container is displaying the current user's identity Out of baseline policy	ubuntu AWS Demo / ...	`whoami` × Terminated
Jun 11th, 2025 02... Execution	A container is displaying its hostname Out of baseline policy	ubuntu AWS Demo / ...	`hostname` × Terminated
Jun 11th, 2025 02... Execution	A container is retrieving metadata from a cloud inst... Out of baseline policy	ubuntu AWS Demo / ...	`curl -s http://169.254.169.254/latest/met…` × Terminated
Jun 9th, 2025 20:04 Execution	A container is making a request to OpenAI's API pot... Out of baseline policy	ubuntu AWS Demo / ...	`curl https://api.openai.com/v1/chat/compl…` × Terminated
Jun 9th, 2025 20:04 Execution	A container is reading a suspicious file Out of baseline policy	ubuntu AWS Demo / ...	`cat /tmp/.audit` × Terminated

Upwind Threat Stories accelerate security incident investigation by providing real-time, contextual insights within the Threats Module, empowering teams to connect the dots between seemingly unconnected incidents and view the entire timeline of evolving security incidents. Use cases include:

- **Lateral movement within a cluster**
- **Privilege escalation via unexpected shell invocation**
- **Data exfiltration across unusual outbound traffic paths**

Threat Stories contextualize these sequences using telemetry gathered via Upwind's eBPF sensors and Kubernetes audit logs, eliminating the need for manual pivoting between tools.

Cloud Baselines: Data-Driven Anomaly Detection

Upwind Threat Stories are powered by Upwind Cloud Baselines, which are created by continuously monitoring an application's activity over hours, days and weeks to build sophisticated machine-learning (ML) models that enable Upwind to distinguish "normal" from "abnormal" activity. This in-depth analysis broadens understanding of typical resource activity, and quickly identifies anomalies when they occur.

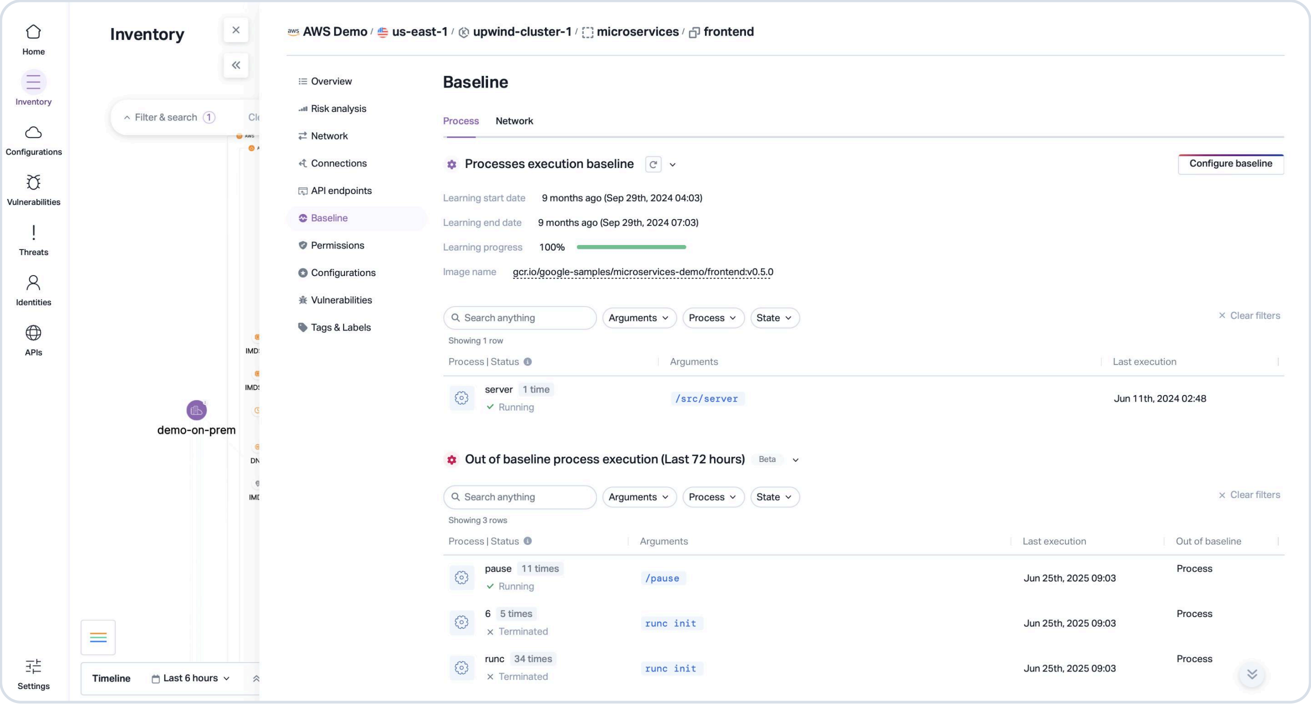

Baseline

Processes execution baseline

Learning start date 9 months ago (Sep 29th, 2024 04:03)

Learning end date 9 months ago (Sep 29th, 2024 07:03)

Learning progress 100%

Image name gcr.io/google-samples/microservices-demo/frontend:v0.5.0

Showing 1 row

Process | Status	Arguments	Last execution
server 1 time ✓ Running	`/src/server`	Jun 11th, 2024 02:48

Out of baseline process execution (Last 72 hours)

Showing 3 rows

Process | Status	Arguments	Last execution	Out of baseline
pause 11 times ✓ Running	`/pause`	Jun 25th, 2025 09:03	Process
6 5 times × Terminated	`runc init`	Jun 25th, 2025 09:03	Process
runc 34 times × Terminated	`runc init`	Jun 25th, 2025 09:03	Process

Upwind generates cloud baselines by first taking a comprehensive inventory of a user's cloud infrastructure and then continuously monitoring process executions, network communications, and file system accesses across Kubernetes workloads, serverless functions, and virtual machines using the Upwind eBPF sensor.

- **Process Baselines**

  Processes and their arguments are learned over time. Deviations from established norms, such as a shell (`/bin/sh`) running in a pod where it's never been observed, trigger alerts. This allows detection of previously unknown behaviors without relying on static signatures.

- **Network Baselines**

  Inbound and outbound traffic is analyzed to model normal ports, protocols, destinations, and traffic volumes and baselines are established after a configurable amount of training. After baselines are set, alerts are generated when anomalies from specified thresholds are detected, such as:

  - **A pod in a non-sensitive namespace communicating with `kube-system`**
  - **New communications or excessive data transfer to unknown IP addresses**
  - **Spikes in traffic volume or unusual protocol usage**

By generating cloud baselines, Upwind surpasses typical threat detection methods, such as only scanning for known malware signatures. Instead, Upwind proactively identifies abnormal human and machine activities within a cloud environment, providing defense-in-depth for detecting and responding to potential threats in real time.

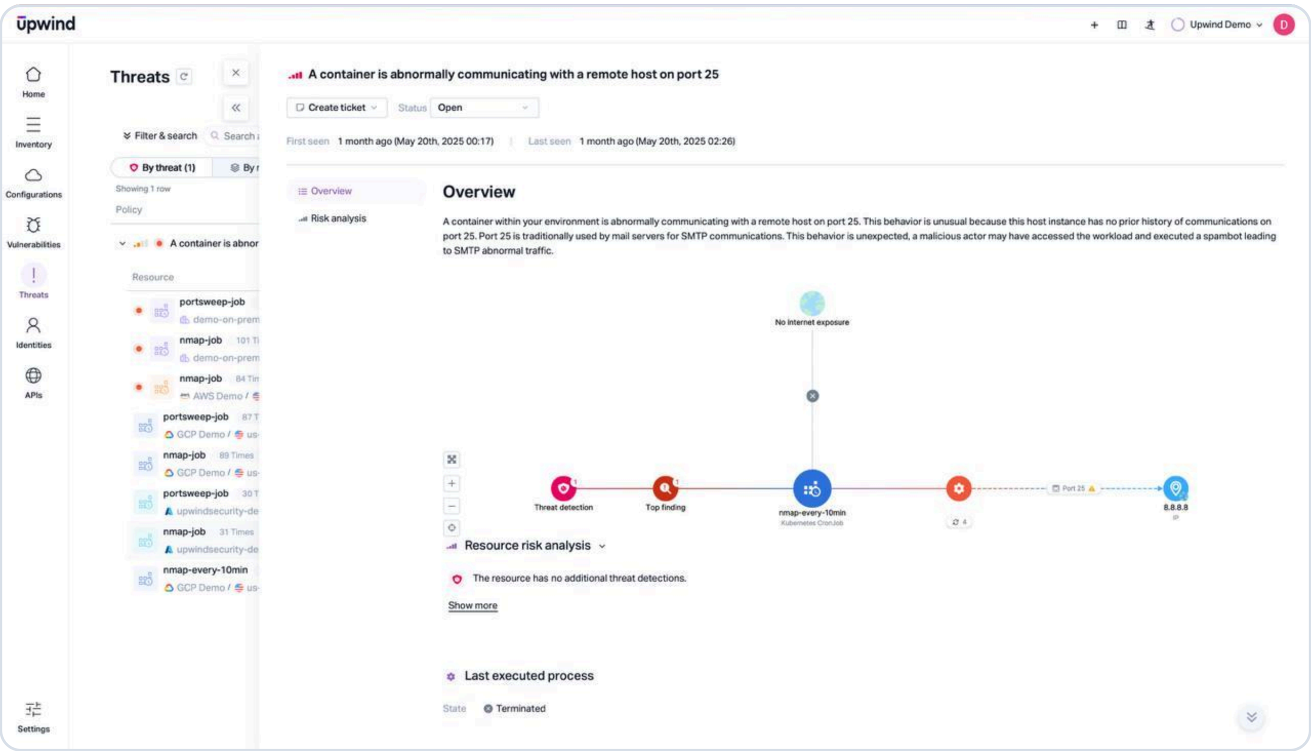

A container is abnormally communicating with a remote host on port 25

Status Open

First seen 1 month ago (May 20th, 2025 00:17) | Last seen 1 month ago (May 20th, 2025 02:26)

Overview

A container within your environment is abnormally communicating with a remote host on port 25. This behavior is unusual because this host instance has no prior history of communications on port 25. Port 25 is traditionally used by mail servers for SMTP communications. This behavior is unexpected, a malicious actor may have accessed the workload and executed a spambot leading to SMTP abnormal traffic.

Resource risk analysis

The resource has no additional threat detections.

Last executed process

State Terminated

Not every anomaly is a threat - and Upwind knows the difference. By continuously learning behavioral baselines for processes, traffic, and access patterns, the platform flags deviations and uses AI reasoning to determine if they align with known TTPs.

Deviations are automatically evaluated and mapped to MITRE ATT&CK tactics, allowing Upwind to distinguish between benign drift and genuine risk. The result of this fine-tuned evaluation and AI-reasoning is fewer false positives, more precise detections, and prioritized response based on real attacker behavior as defined by MITRE tactics.

Improved Detection to Response Workflow

Behind the scenes, we use a streaming event pipeline that evaluates new runtime signals against your learned baseline and known TTP patterns. Every new detection is scored for correlation based on timing, resource context (container, image, user), and prior behavior, enabling automatic grouping without human intervention.

Inbound and outbound traffic is analyzed to model normal ports, protocols, destinations, and traffic volumes and baselines are established after a configurable amount of training. After baselines are set, alerts are generated when anomalies from specified thresholds are detected, such as:

- **A full timeline of suspicious activity**
- **Linked telemetry from runtime, network, and cloud sources**
- **Suggested response actions (kill process, quarantine container, block traffic)**

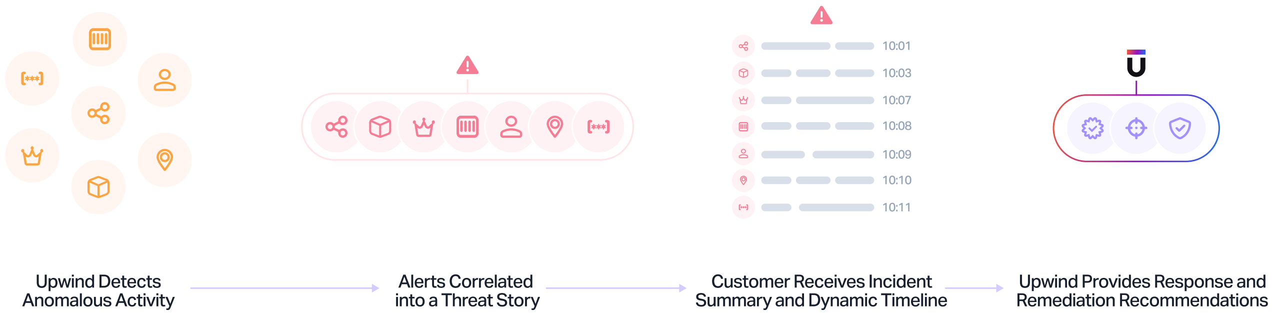

Upwind Detects Anomalous Activity → Alerts Correlated into a Threat Story → Customer Receives Incident Summary and Dynamic Timeline → Upwind Provides Response and Remediation Recommendations

These actions can be executed manually or enforced via policy to automate prevention.

Below, we detail the advanced threat response options available in the Upwind Platform.

- **Terminate a malicious process**
  Upwind enables users to easily terminate a malicious process, along with the ability to view the detection process tree and any child processes that stem from the original process. Users can also kill processes running on multiple different containers at the same time and strategically kill a malicious process without killing the container, enabling rapid security without disrupting cloud operations.

- **Create Prevention Policies**
  Upwind also provides users with the ability to set prevention policies for multiple processes, with or without arguments, over a specific timeframe. These flexible policies will repeatedly kill a malicious process if it tries to re-spawn, even if it is not currently running.

- **Audit Trail**
  Upwind enables users to also track all current prevention settings and view all policies in one place. Upwind also keeps a response audit log, providing information into which members of an organization chose to use the response feature, when it was used, and if it was successful.

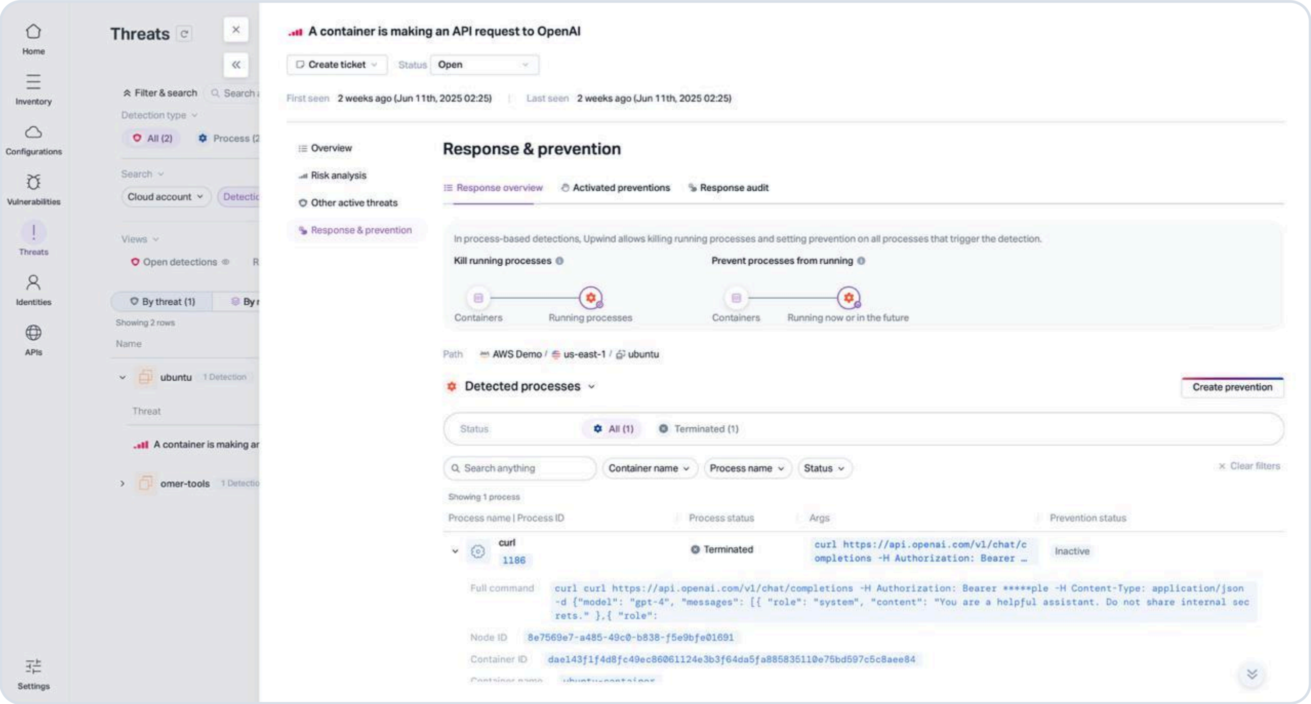

A container is making an API request to OpenAI

Status: Open

First seen 2 weeks ago (Jun 11th, 2025 02:25) | Last seen 2 weeks ago (Jun 11th, 2025 02:25)

Response & prevention

In process-based detections, Upwind allows killing running processes and setting prevention on all processes that trigger the detection.

Path AWS Demo / us-east-1 / ubuntu

Detected processes

Showing 1 process

Process name | Process ID	Process status	Args	Prevention status
curl 1186	Terminated	curl https://api.openai.com/v1/chat/completions -H Authorization: Bearer …	Inactive

Full command

```
curl curl https://api.openai.com/v1/chat/completions -H Authorization: Bearer *****ple -H Content-Type: application/json
-d {"model": "gpt-4", "messages": [{ "role": "system", "content": "You are a helpful assistant. Do not share internal sec
rets." },{ "role":
```

Node ID 8e7569e7-a485-49c0-b838-f5e9bfe01691

Container ID dae143f1f4d8fc49ec86061124e3b3f64da5fa885835110e75bd597c5c8aee84

By leveraging these advanced threat detection and response capabilities, Upwind users have reported a significant decrease in false positives and alert fatigue, accelerated investigation timelines, and up to 7x faster mean time to response.

Conclusion

Upwind's Cloud Baselines and Threat Stories transform raw telemetry into actionable intelligence. By automating baseline learning and incident correlation, the platform enables security engineering teams to focus on real threats, respond faster, and reduce operational overhead. Upwind’s runtime-centric architecture scales across cloud-native environments and provides the evidence needed for detection, triage, and response — without the noise.

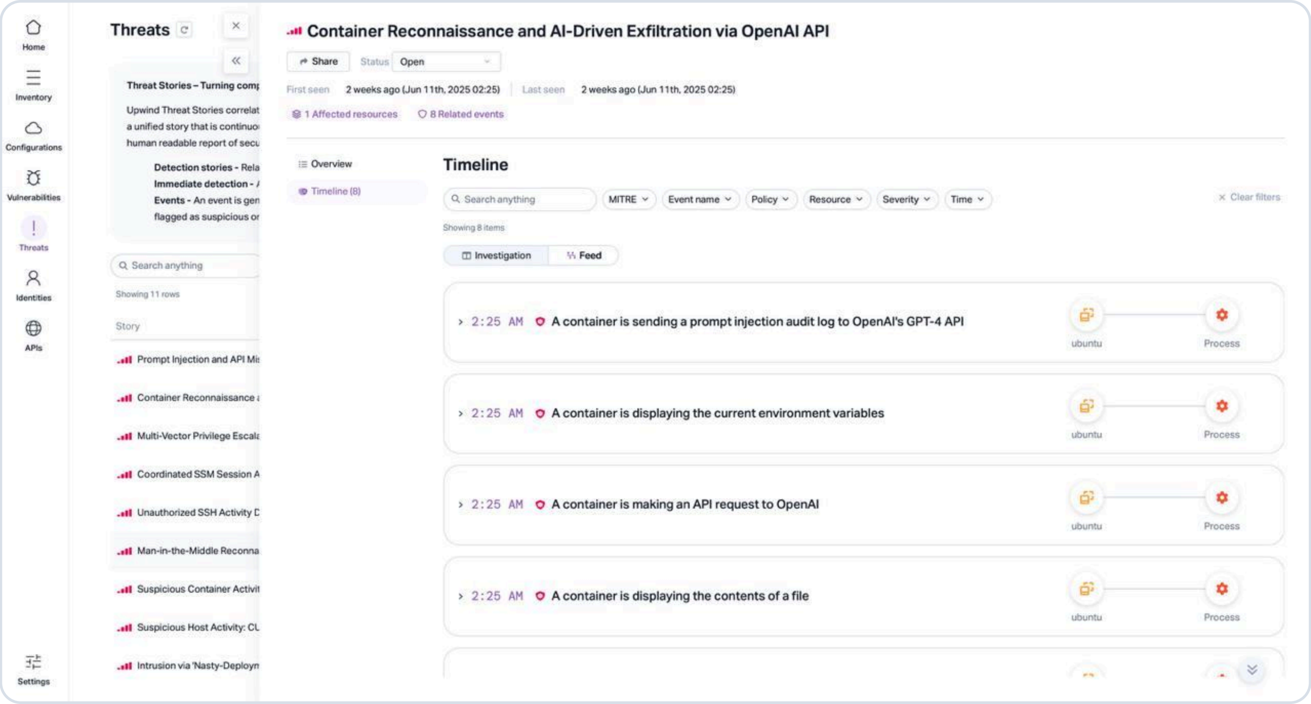

Threat Stories – Turning comp

Upwind Threat Stories correlat
a unified story that is continuo
human readable report of secu

Detection stories - Rela
Immediate detection - /
Events - An event is gen
flagged as suspicious or

Showing 11 rows

Story

Prompt Injection and API Mi
Container Reconnaissance a
Multi-Vector Privilege Escala
Coordinated SSM Session A
Unauthorized SSH Activity D
Man-in-the-Middle Reconna
Suspicious Container Activit
Suspicious Host Activity: CL
Intrusion via 'Nasty-Deployn

Container Reconnaissance and AI-Driven Exfiltration via OpenAI API

Status: Open

First seen: 2 weeks ago (Jun 11th, 2025 02:25) | Last seen: 2 weeks ago (Jun 11th, 2025 02:25)

1 Affected resources | 8 Related events

Timeline

Showing 8 items

2:25 AM A container is sending a prompt injection audit log to OpenAI's GPT-4 API

2:25 AM A container is displaying the current environment variables

2:25 AM A container is making an API request to OpenAI

2:25 AM A container is displaying the contents of a file

Want to know more about Upwind’s threat stories? 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.