



KubeCon



CloudNativeCon

Europe 2026

#KubeCon #CloudNativeCon

From Logs to Decisions: Autonomous AI Agents for Real-Time Kubernetes Threat Response

Willem BERROUBACHE, Orange



kubectl describe human Willem

Willem Berroubach *Telco Cloud & Security Architect @ Orange*



 **Securing the 5G Core:** Designing monitoring and security for mission-critical Telco Cloud & K8s environment.

 **AI-Driven Security:** Building AI solutions to protect networks and customers.

 **Production & Edge:** Built K8s prod clusters for Orange France and led 5G Edge deployments as PM.

 **Tech Leadership:** Kubernetes trainer for the Orange Group (100+ engineers upskilled).





KubeCon



CloudNativeCon

Europe 2026

Disclaimer:

The following slides are based
on a true story
(and it's not a fairy tale)



Cloud-Native Security: Expectation vs. Reality

Based on true (and terrifying) audit events



KubeCon



CloudNativeCon

Europe 2026

[CRITICAL]:

Cloud-Native logo found ...

... on top of a smoking legacy stack.

[VULN]:

Hardcoded pw: admin

200+ CVEs hiding in "certified" partner apps.

[MISCONFIG]:

YAML: « Please, don't look »

(Trust us, you really don't want to).

[FATAL]:

Leaving infrastructure wide open to breaches.

A direct threat to user data and customer trust.



Beyond Audits: Why Runtime Monitoring is Non-Negotiable

📄 The Auditor checks your SecurityContext & RBAC

🚩 The Hacker exploits your running Container

🗺️ Static Manifests are just the map

🏔️ Runtime Events are the real-world terrain

👁️ Full-Stack Observability is seeing the Pod breach

🚒 Automated Detection is the immediate Remediation



🚀 The threat is live. Your defense should be too.



The Supervision Crisis: Why Security Breaks at Scale



KubeCon



CloudNativeCon

Europe 2026

Data Tsunami 🌊



More logs
More events

Alert Fatigue 😞



Everything fires
Nothing matters

Context Blindness 🧩



Event visible
Threat invisible

Discovery Deadlock 🔍



Missing Context
Blocked Decisions

Remediation Lag 🛠️



Manual fixes
Slow response
Rising risk

Silence the noise.
Understand the signal.
Automate the response.

AI isn't just a tool it's the brain our infrastructure was missing



KubeCon



CloudNativeCon

Europe 2026



AGENTS.md




LangChain



LLM




Kubeflow



Agentic AI
Foundation


Agentic AI

The Intelligent Core of Supervision



Closed-Loop Intelligence: From Real-Time Inference to Human Feedback

REAL-TIME INFERENCE




TRAINING PIPELINE

Kubeflow Pipelines v2



FEEDBACK LOOP



Alarm example!

ALARMS29

REMEDIATIONS16

ML MODEL

ALL 29

CRIT 9

HIGH 14

MED 3

LOW 3

< CORR 11

namespace...

CRITICAL

< CORRELATED

ML -0.06

PROPOSED

Shell spawned in container (no TTY)

default/kubewsg-store-5786c67984-fbzdr

15:36:56

HIGH

CHAIN x3

REMIEDIATED

Shell spawned in container (no TTY)

default/kubewsg-store-5786c67984-fbzdr

15:07:46

CRITICAL

CHAIN x2

PROPOSED

Shell spawned in container (no TTY)

default/kubewsg-store-85f85c46b8-dqbx5

15:07:46

MEDIUM

PROPOSED

Launch Sensitive Mount Container

<NA> / <NA>

15:07:46

HIGH

ML -0.06

PROPOSED

Shell spawned in container (no TTY)

default/kubewsg-store-5786c67984-fbzdr

15:07:05

CRITICAL

CHAIN x2

Shell spawned in container (no TTY)

default/kubewsg-store-85f85c46b8-dqbx5

15:06:07

HIGH

CHAIN x2

ML -0.10

Sensitive file read in container

default/kubewsg-store-5786c67984-fbzdr

15:06:17

LOW

PROPOSED

Launch Sensitive Mount Container

<NA> / <NA>

15:06:07

CRITICAL

CHAIN x2

Shell spawned in container (no TTY)

default/kubewsg-store-85f85c46b8-dqbx5

15:04:25

CRITICAL

Shell spawned in container (no TTY)

DISMISS

OVERVIEW

ATTACK

RESPONSE

Namespace

Pod

Container

ML Score

Detection

False Positive

Time

default

kubewsg-store-5786c67984-fbzdr

kubewsg-store

-0.062

RULE+CORRELATED

LOW

2026-03-04T14:36:56.683508513Z

THREAT ANALYSIS

What happened:

Justification:

Impact:

FP reason:

MITRE ATTACK

T1190 Exploit Public-Facing Application

Initial Access

T1190 - Exploit Public-Facing Application

T1059.004 - Command and Scripting Interpreter: Unix Shell

T1105 - Ingress Tool Transfer

T1496 - Resource Hijacking

10.192.104.212:30180

https://orang...LeapToudev7 Orange OpenSearch Dashboards Google Claude-opus...soft Foundry

KubeSwag - Cloud Native Merch

Security Dashboard — KubeCon EU 2026

 **KubeSwag**

 Shop

 My Orders

 Account

 Support



Cloud Native Swag

Premium merchandise for Kubernetes enthusiasts, SREs & DevOps engineers

Search products... 

Best Sellers

8 products



Kubernetes Logo Tee NEW

Organic cotton, unisex

\$29.99 



Cloud Native Hoodie HOT

Fleece-lined, zip-up

\$59.99 



kubecti apply -f coffee.yaml

Ceramic, 350ml

\$14.99 



CNCF Sticker Pack

25 die-cut vinyl stickers

\$9.99 

Practical outcomes

Operationalizing Security Intelligence at Scale



Contextual Fidelity

- Signals validated against cluster state and **toolsets**.
- High-fidelity triggers **eliminate** operational noise and off-hours **overload**.



Assisted Defense

- **Fast containment** through Runtime blocking and security mechanisms.
- **Human-in-the-loop** validation ensures precise and safe mitigation.



Forensic insight

- **Event correlation** to reconstruct attack paths and history.
- Supplies the **actionable forensic history** required for deep investigation.

Finalizing robustness, ready to deploy soon!

AI Risks

Hallucinations

Cascading Errors

Model Drift

Overconfidence Effect

KEEP
HUMAN
IN
THE
LOOP !!

Human Safeguards

**Ethical
Judgment**

Context Awareness

Human Oversight

Validation

Agentic AI: A New Operating Paradigm



KubeCon



CloudNativeCon

Europe 2026



Scale x Speed

Humans are great at judgment.

Terrible at processing.

Agents do both: at any scale,



From Tools to Reasoning

Classic automation executes rules.

Agentic AI understands goals, chooses tools, adapts to context.



A New Stack Primitive

Agents aren't plugins.

They're a new layer

Containers changed deployment

Agents change operations.

Agentic AI is a new paradigm. The real challenge is how we adopt it.

Thank you ! / For updates:



From Logs to Decisions: Autonomous AI Agents for Real-Time Kubernetes Threat Response

Willem Berroubache & Networks Security and Supervision team!!!

Lead Architect 5G Security Monitoring & AI For Security, Orange

Code & resources coming soon:
Connect on LinkedIn to get notified



#KubeCon #CloudNativeCon

linkedin.com/in/willem-b
willem.berroubache@orange.com