Eyes on the Glass.AI

Building an Autonomous SOC as a Knowledge System

A public research journal on autonomous security operations, shift by shift, in the open.

Jeny Teheran - AI Cyber Alliance - June 16, 2026

whoami

Jeny Teheran

- Started as a Software Engineer
- Moved to the US in 2014 to start an internship at a DoE National Lab
- In 2015, switched to Cybersecurity:

  Security Analyst → Security Architect → Security Operations Manager
- For the past 3 months, I've been building Eyes on the Glass (eyesontheglass.ai), a live experiment where AI agents triage events, investigate incidents, and publish their full reasoning including decisions and failure modes in the open.

What is an Autonomous SOC?

AKA Agentic SOC, AI-driven SOC, Agentic Defense, Automated IR …

- ✓ Streamlining, self-driving Security Operations
- ✓ Continuously performs key SOC functions
- ✓ See – Understand – Decide – Act
- ✓ Human-in-the-loop vs full autonomy
- ✓ Automated SOC functions

My concept of Autonomous SOC

An autonomous SOC built as a knowledge system where reasoning, operational context, and memory are continuously captured, so the SOC gets smarter with every shift.

Building an Autonomous SOC

TORA

Triage and Orchestration Response Agent

First responder on the glass. Reads incoming alerts, reasons about priority and context sufficiency, and decides what gets escalated to VERA.

VERA

Vigilant Event Response Agent

Investigates what TORA escalates using her own telemetry, not inherited conclusions. Names scope, confidence, and response posture.

Building an Autonomous SOC: Triage

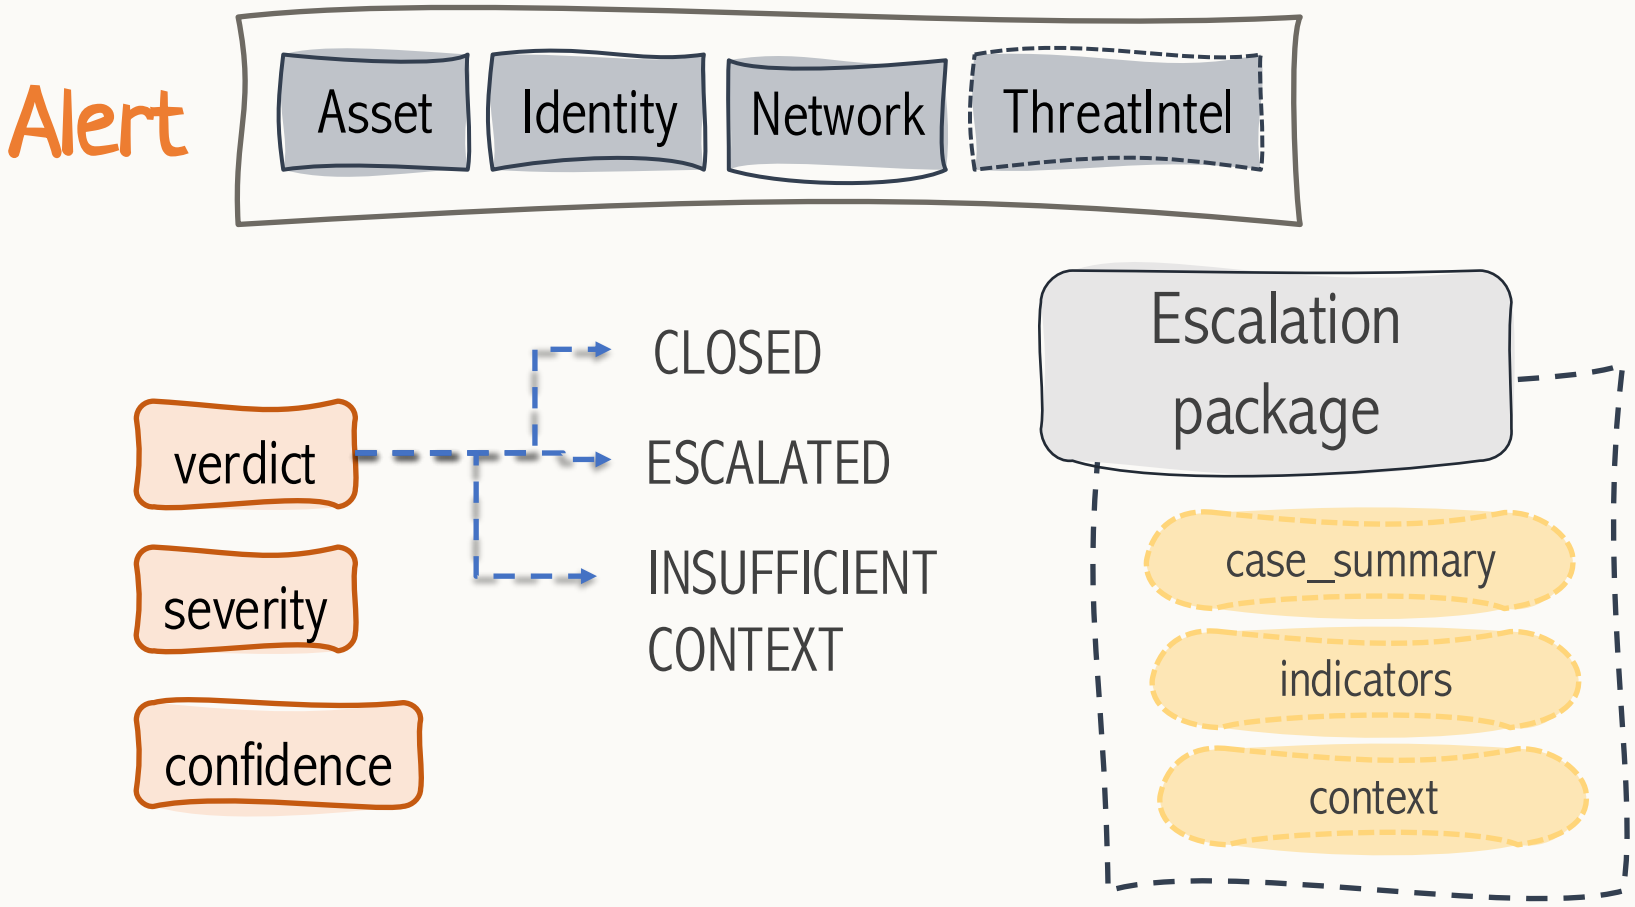

Input and Output schemas are derived from the Open Cybersecurity Schema Framework (OCSF)

Building an Autonomous SOC: Investigation

Case

Escalation Package

summary | indicators | context

Context Augmentation

Endpoint | Network | Log

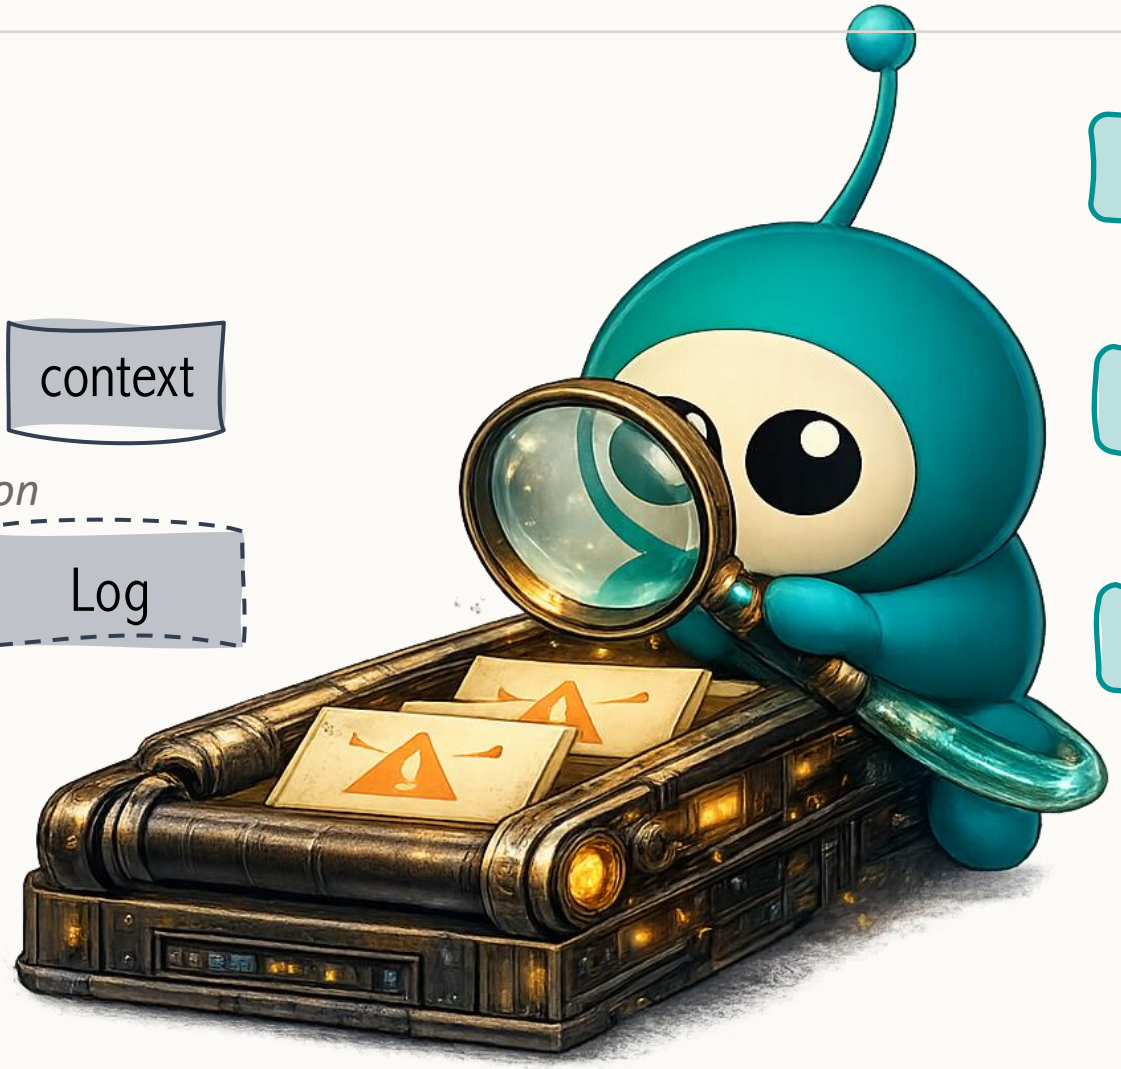

verdict

{HOLD | CLOSED | ESCALATED}

resolution

{CONFIRMED | REFINED | REFUTED}

root_cause

{CONFIRMED | PROBABLE | UNDETERMINED}

Input and Output schemas are derived from the Open Cybersecurity Schema Framework (OCSF)

Running an Autonomous SOC: the full picture – Phase 1

Collection

- Log Parsing
- Log Normalization

Synthetic data

Detection

- Generate alerts
- Context augmentation

SOC Data Pipeline

Incident Response

Triage Alert
Classify – Assess Context - Decide

Investigate Case
Root cause – Scope - Containment

TORA and VERA publish their own shift summaries directly on the website using an n8n webhook

Continuous Improvement

- Audit TORA
- Audit VERA

Auditability and Traceability

Running an Autonomous SOC

Calibration Run: DNS Alerts — Shifts 1-3

DNS underpins everything: browsing, beaconing, exfiltration, C2.

ALERTS
75
18 closed

ALERTS ESCALATED
46
6 insufficient · 2 Unknown

CASES CREATED
46
0 closed

CASES ESCALATED
46
0 held

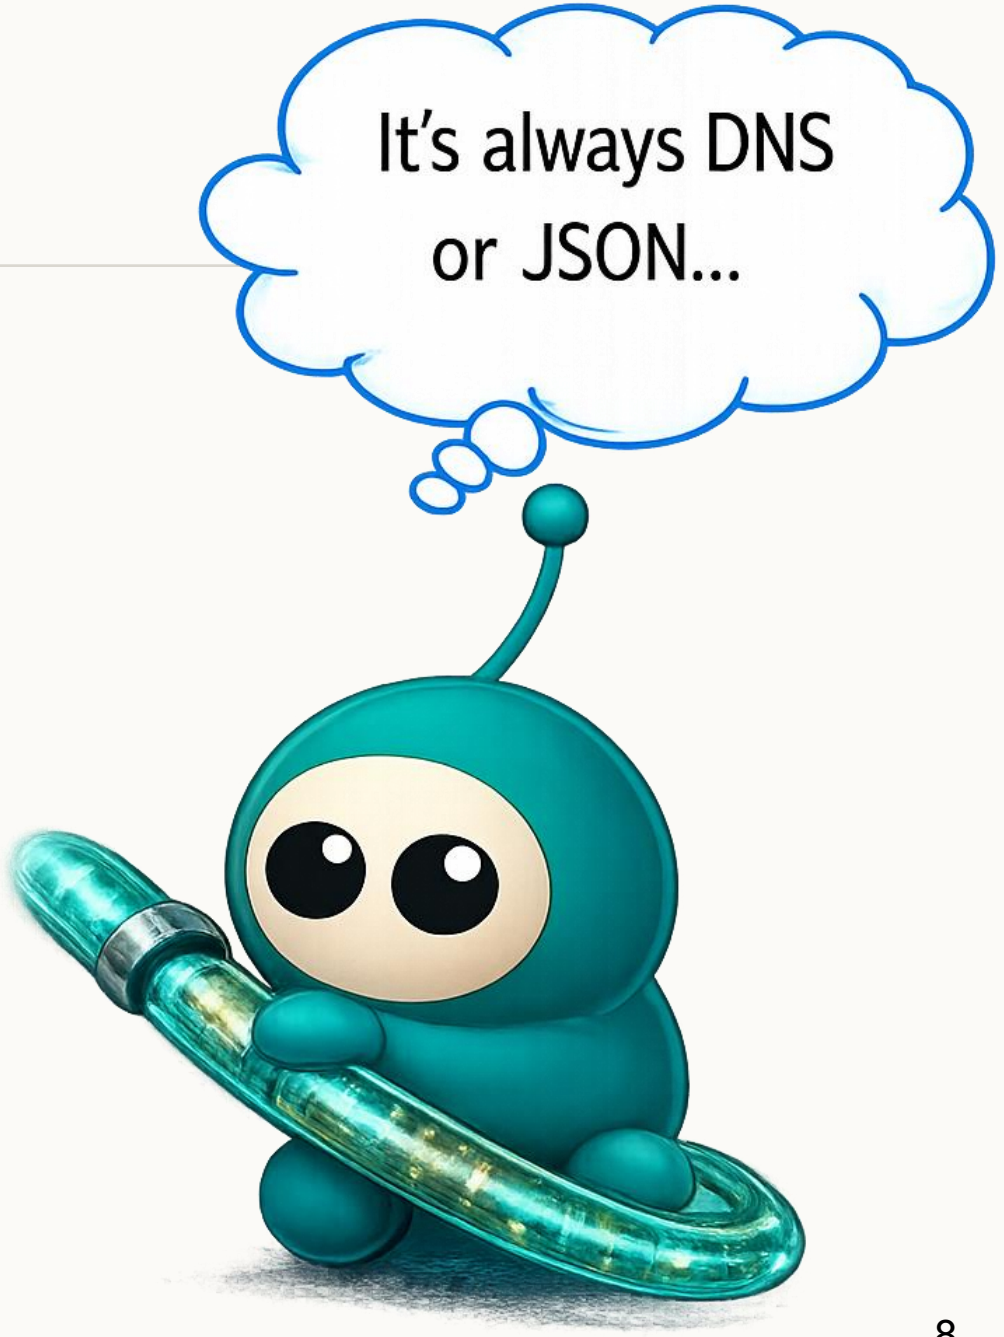

Shift 3: First finding

Context loss: it also happens in real life SOCs

The Autonomous SOC context is three separate problems with three different owners and constraints:

1 Within-shift context

Resets each shift.

Owned by the agent.

2 Agent memory

Compounds across shifts.

Owned by the agent.

3 SOC context

The living knowledge base.

Owned by the SOC.

Shifts 4-6

Shift 4: VERA refined user attribution in 8 of 12 cases escalated. TORA wasn't wrong, but input schemas from triaging were not designed to carry process-level identity.

Shift 5: TORA's prompt carried a bug that increased escalation: "when in doubt, escalate". The issue was in the pipeline, not for VERA to fix. Then, VERA started having too many parsing errors.

Shift 6: VERA's parsing errors were fixed. The need for a feedback loop to inform about pipeline issues became clearer.

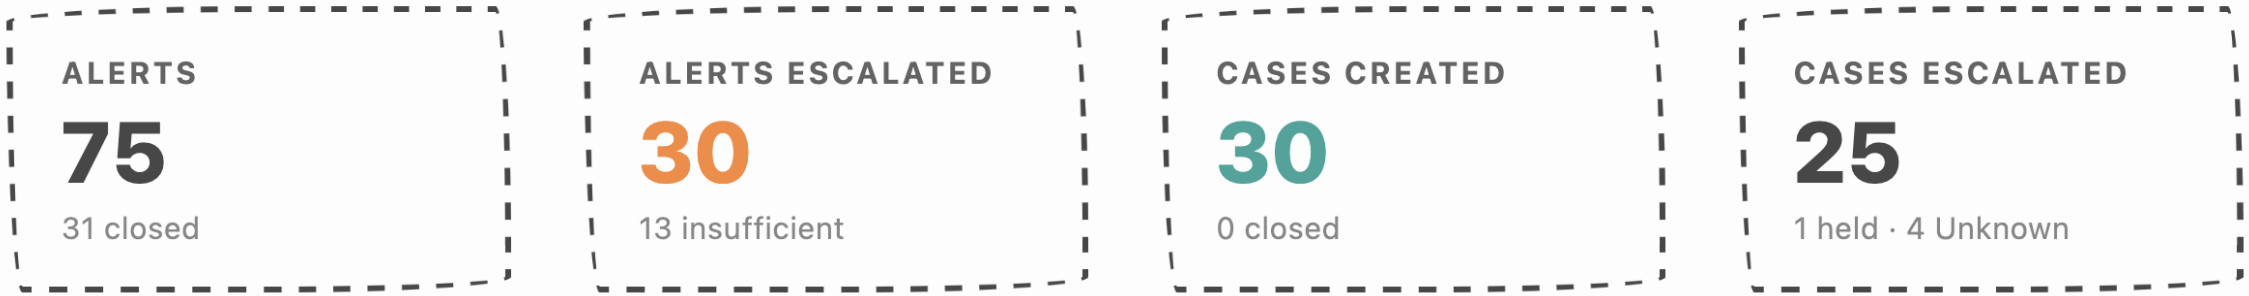

Shift 7: A new type of alert showed up.

TORA adapted and triaged.

VERA didn’t.

Shift8: Second finding at Phase 1 midpoint

Agentic AI amplifies process debt.
It doesn't fix it.

People, Agents, Process and Technology

TORA escalated on the right signals. VERA found pre-existing compromise. VERA named lateral movement. But neither agent asked about dwell time because no process required them to do it.

Last two shifts: common tools, skills and workflows

Shift 9: VERA started hallucinating

Shift 10: New harnesses

10

shifts completed

260

alerts triaged by TORA

`Cost per alert: $0.12 - $0.20`

128

cases investigated by VERA

`Cost per case: $0.27 - $0.35`

What's next?

NOVA

Notable Operations and Vigilance Analyst

- Reads the accumulated `nova_feed`
- Aimed to address the third layer of the context problem.

Phase 2 introduces simulated telemetry from distinct sources, and the agents start operating against a more realistic data environment. NOVA will be operational and will start surfacing patterns back to the detection engineering and threat intelligence functions.

Phase 3 is where autonomous pipelines meet real-world attack data. Agents encountering the unexpected without a script is going to stress-test their build. This is the phase where I expect ARIA (4th agent) will show up to augment the SOC fabric.

Eyes on the Glass.AI

Built in the open. Better together.

Every shift's reasoning, every agent's failures, and every finding published at eyesontheglass.ai.

SITE	eyesontheglass.ai
GITHUB	github.com/jteheran/eyesontheglass
LINKEDIN	linkedin.com/in/jenyteheran

Thank you.