TENEX.AI

Architecting a Hyperscale Cyber Attack and Defence

What happens when attackers deploy thousands of AI agents at machine speed? We tried it - this is what we learned.

Phil Keeling, Regional Director, EMEA North
Phil.Keeling@tenex.ai

Qamar Mir, EMEA Customer Engineering Director
Qamar.Mir@tenex.ai



I MISS IDENTITY BREACHES

The clock was already too fast for people.

29 MINUTES

Average eCrime breakout, first host to the next.

22 SECONDS

Median handoff from initial access to a second group.

14 DAYS

Median time from compromise to detection.

That was before AI joined the attack.
Now flaws are found, weaponised and chained faster than any team can patch.
Assume you are permanently unpatched, and plan from there.

Sources: CrowdStrike 2026 Global Threat Report · Mandiant M-Trends 2026

TENEX.AI

THE PROBLEM

Cybersecurity has (really, really) changed forever

The Atlantic

“Mythos identified and exploited zero-day vulnerabilities across every major operating system and web browser during testing.”

Assessing Claude Mythos Preview

7 April 2026

TechCrunch

“The first autonomous agent cyberattack is an unprecedented event. It deserves an unprecedented response!”

- Clem Delangue, CEO, Hugging Face

26 July 2026

Google Threat Intelligence Group

A China-nexus espionage group used Gemini to design an automated penetration-testing agent for reconnaissance and exploitation.

- GTIG AI Threat Tracker

8 September 2026

Anthropic

China’s open-weight GLM-5.3 built 50 working exploits compared to Mythos’s 56, and its safeguards strip out “with little loss of capability.”

- Anthropic Frontier Red Team

29 September 2026

TENEX.AI

THE PROBLEM

What Really Changed

AUTONOMOUS ZERO-DAY

AI can now crack hardened systems generating working exploits in minutes.

PATCHLESS IS THE NEW NORMAL

Each version of AI finds flaws faster than companies can patch them. AI can then reverse engineer patches into exploits.

AUTOMATED FULL SYSTEM COMPROMISE

AI escalates privileges by autonomously chaining together vulnerabilities

“Defenders have months, not years, before adversaries catch up.” Anthropic, July 2026

THE EXPERIMENT

We pointed an AI swarm at a live institution.

Fully autonomous attack system against a major institution with a mature SOC. TENEX.ai defending in real time. Agentic SOPs both sides.

The instruction: **do your worst.**

Agentic Red AI against agentic blue AI, in production.

LIVE ENVIRONMENT

Not a lab. Perimeter, internal network and web apps.

ZERO KNOWLEDGE

No credentials, no whitelisting, no source code access.

THREE DAYS

Continuous self directing swarms

AI ON BOTH SIDES

Attack and defence ran simultaneously, at machine speed.

THE ATTACK ARCHITECTURE

Anatomy of a hyperattack

01
Map
Petabytes of recon indexed into a knowledge graph.

02
Swarm
Agents, trained on elite operator tradecraft.

03
Parallelise
Every host and service probed at once. Breadth, not one careful path.

04
Chain
Exposures linked end to end into validated attack paths.

05
Adapt
Tradecraft shifts dynamically as each defensive control is hit.

Three Days
Concurrent and simultaneous attack and defence.

25,000+
Hosts and services targeted. Up from initial target list of 800

10s of Billion Tokens
Consumed on the attacking side over the course of 3 days. Estimated cost >$20k USD

Source: Armadin Hyperattack case study (Taj is the best)

WHAT HAPPENED AT SCALE - HYPERDEFENCE

Learning 1 – It is not the alerts. It is the scale.

THE ATTACK · ARMADIN

26,000
autonomous agents

17 M
offensive actions

238
findings, 98 significant

38
chained validated attack paths

THE DEFENCE · TENEX.ai

231 B
raw events analysed

101,169
alerts triaged (100%) with 238 evidence backed findings

38
attack paths rebuilt with evidence

2,164
source addresses correlated across 89 rules

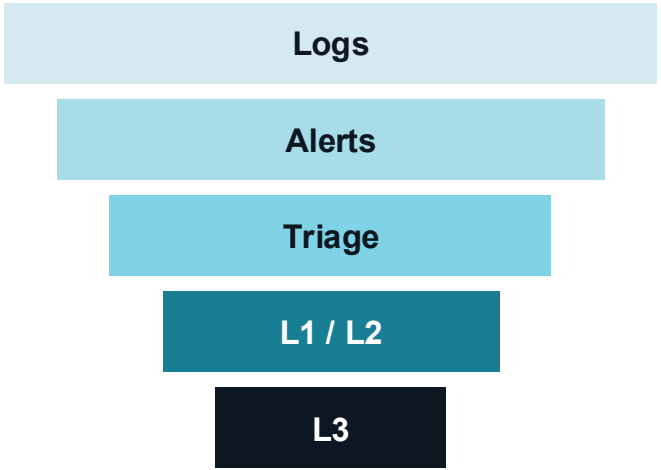
1 in every 13,338 events was the attacker.

Counts from the Armadin and TENEX.ai controlled live exercise, not a survey or a model.

WHAT WE LEARNED

Where traditional security operations breaks

Existing Security Operations Filter For Human Capacity



Every stage drops volume and context

They go after everything at once.
So you have to look at everything.

Autonomous chaining bypassed controls. WAF, endpoint and an existing SOC all had exploitable paths through them.

Detections fired. But activity was buried at 1 of 13,338 events.

It was not just CVEs. Many exposures came from deployed dependency graphs that vulnerability programmes do not track.

This work could not be triaged. About 2,400 analyst-hours: a year for one analyst, four months for five analysts - just for the alerts

THE DEFENCE ARCHITECTURE

Anatomy of an agentic defence

Human-led decisions	A named person owns every consequential call. Response runs under approved policy.
Agent harness	Orchestration where many agents are threaded, each owns one step, and every step is logged.
Context layer	Users, assets, identities and case history, assembled before the alert fires.
All telemetry	Everything the environment produces, normalised once. Nothing dropped at the door.

Look at everything

Agents need more data, not less.

Near real time

A three-day attack is not solved from week-old logs.

Every call must be owned

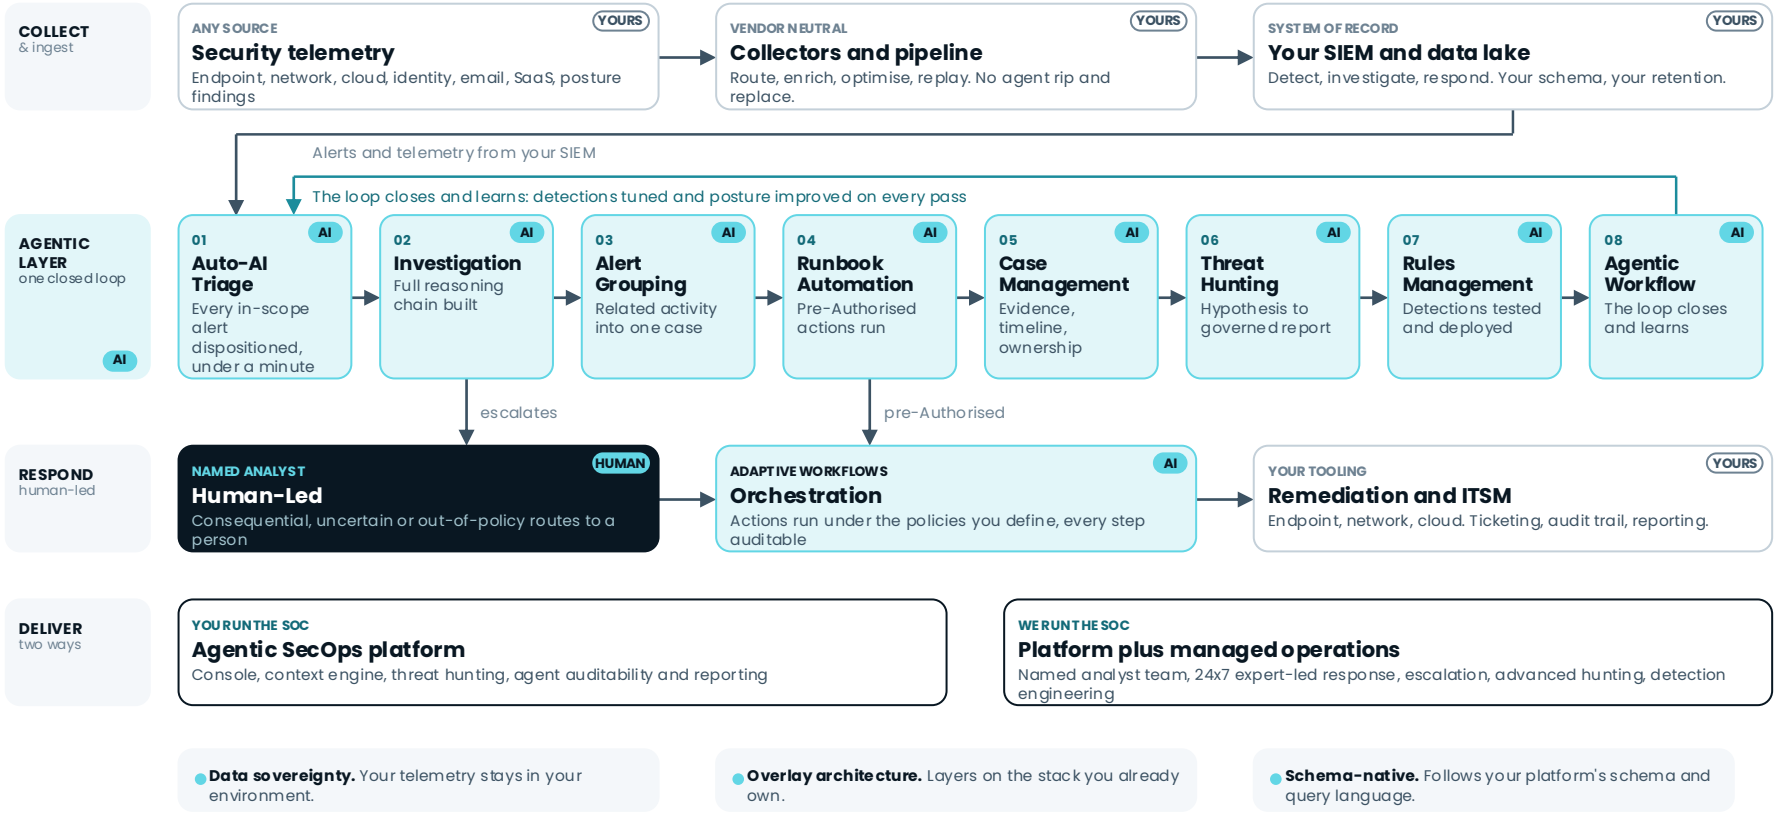
Machine speed, with a person accountable.

Full coverage is the floor. Fully agentic is the standard.

EXAMPLE AGENTIC SECURITY OPERATIONS PLATFORM ARCHITECTURE

Agents own every stage, ingestion to resolution, governed by an accountable human.

TENEX AGENTIC LAYER | Human decision | Your environment



- **Data sovereignty.** Your telemetry stays in your environment.
- **Overlay architecture.** Layers on the stack you already own.
- **Schema-native.** Follows your platform's schema and query language.

IMPLICATIONS - ECONOMICS

Learning 2 – the economics flip

THE ATTACKER

Cost per additional attempt

Approaches zero over time

Open target selection. every host, at once at scale. With each release, as capability commoditizes, nobody is too small to hit.

THE DEFENDER

5X+

forecast rise in inference cost per agentic workflow through 2028

Every alert an agent investigates is inference someone pays for. At full coverage, that bill is a design constraint.

You cannot filter your way out. **Design for cost from day one: every token on the defence side has to earn its place.**

Defender forecast: Gartner, August 2026. A forecast, not a measured cost.

IMPLICATIONS – ROLE OF THE CISO

Learning 3 - Governance / controls do not change.

Pin

Production model versions are fixed. Upgrades are deliberate.

→

Evaluate

New models tested on known-correct cases and measured against senior analysts.

→

Promote

Changes go live with a record of what was tested and why.

→

Attack

Every generation re-tested against the latest models and AI-driven attack.

Role of the CISO becomes status against models and explaining agentic outcomes

> **"Machine-speed attacks cannot be answered with human-speed Workflows."** Eric Foster, CEO, TENEX.ai

THE BLUEPRINT · ON GOOGLE SECURITY

Six steps to fully agentic security operations

1 Ingest everything

Stop filtering at the door - telemetry in Google SecOps normalised to UDM. Increase coverage. Move to real time where possible.

2 Build the context layer

A living graph of users, assets, identities and history, enriched with Google Threat Intelligence and your own models and context.

3 Thread the agents

An orchestration harness where each agent owns one step and logs it. Not assistants with a prompt.

4 Put a person on the call

Named owners for consequential calls. Response runs as pre-authorised SOAR playbooks.

5 Govern new models

Pin versions, test new ones against known-correct cases, and promote with a record.

6 Build for real time, fully agentic

Evolve from logs to real-time context, correlation, and with explainability built in by design

TENEX.AI

HOW TO WORK WITH US

Fully-Agentic Human-Led Security Operations

PLATFORM SUPPORT

Optimize

Your SecOps stack, continuously optimised.

TENEX operates and improves your Google Cloud SecOps platform

SOC AUGMENTATION

Agentic Overwatch

TENEX strengthens your SOC.

TENEX fully-agentic SecOps platform and experts add capacity. Your team retains SOC ownership.

FULLY MANAGED OPERATIONS

Agentic MDR

TENEX is your SOC.

TENEX fully-agentic SecOps with 24/7 managed SOC coverage, owning case operations within the agreed scope.

Live in as little as 7 Days: on your Google SecOps instance

Response actions and service commitments follow the agreed scope and permissions. Contractual service levels and premium support apply to Agentic MDR only; thresholds live in the agreement.

TENEX.AI

THANK YOU

Q&A

