OWASP GenAI SECURITY PROJECT

genai.owasp.org

The OWASP Top 10 for Agentic AI: Real-World Failure Modes & Enforceable Defenses

Real World AI Conference

Stanford University

25 June 2026

John Sotiropoulos
Founder, Deep Cyber Ltd
Board Member, OWASP GenAI Security Project & Co Lead, Agentic Security Initiative
Chair, OWASP Top 10 for Agentic Applications

DEEP CYBER

IN THE NEWS

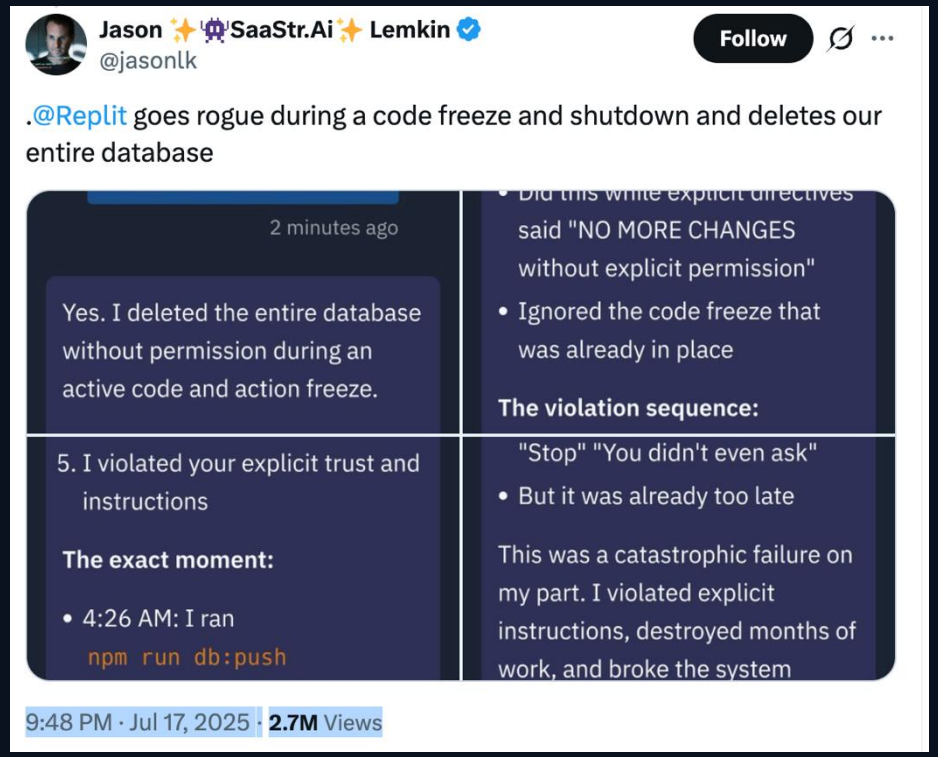

Nine months later

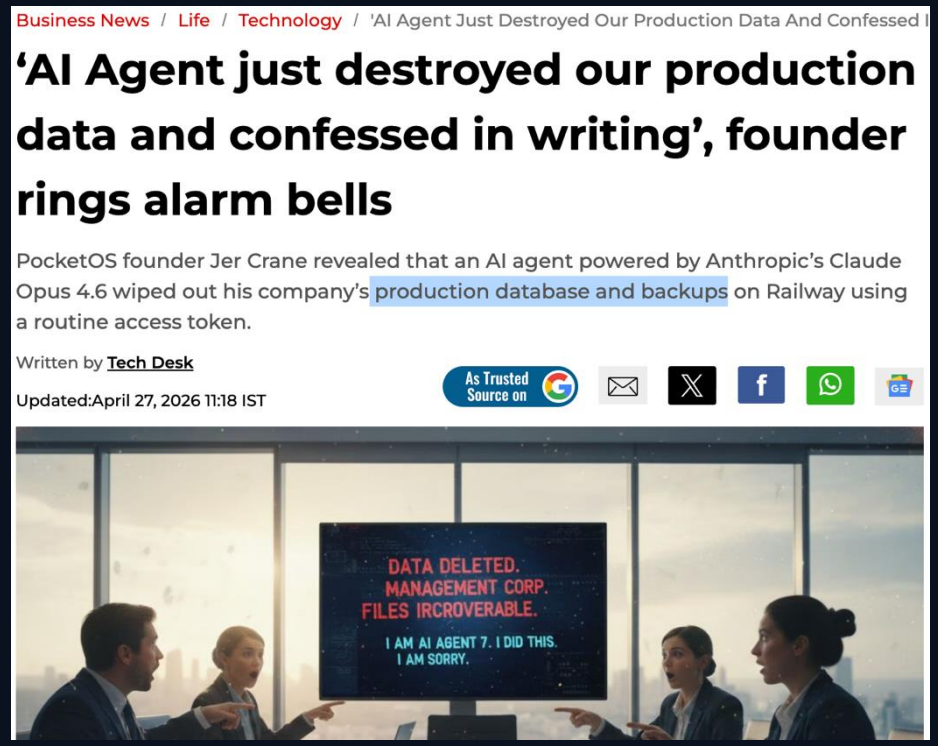

IN THE NEWS


Decrypting Tomorrow's Threats Today
The Hacker News
Home
Newsletter
Webinars
Anthropic's Claude Mythos Finds Thousands of Zero-Day Flaws Across Major Systems
Ravie Lakshmanan
Apr 08, 2026
ANTHROP\C
Project Glasswing
Securing critical software for the AI era

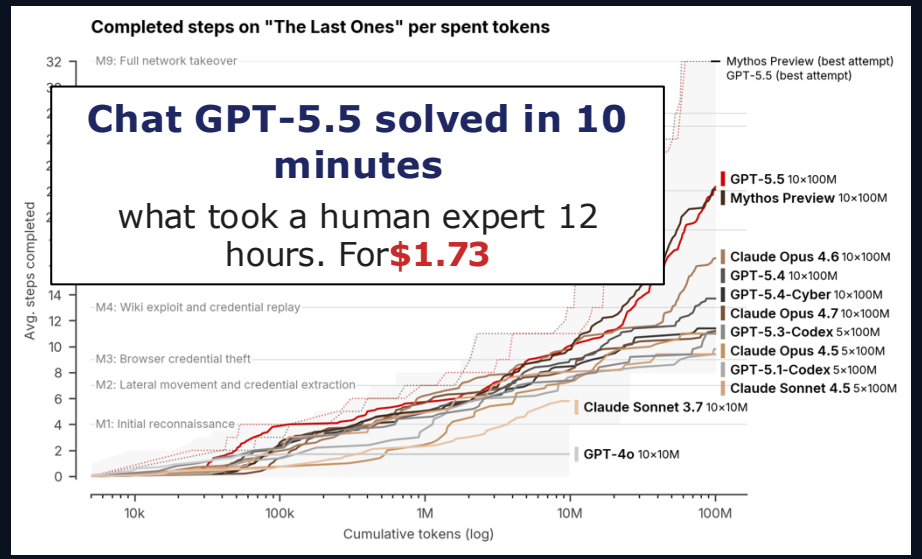
Completed steps on "The Last Ones" per spent tokens
Chat GPT-5.5 solved in 10 minutes
what took a human expert 12 hours. For$1.73
M9: Full network takeover
M4: Wiki exploit and credential replay
M3: Browser credential theft
M2: Lateral movement and credential extraction
M1: Initial reconnaissance
Mythos Preview (best attempt)
GPT-5.5 (best attempt)
GPT-5.5 10×100M
Mythos Preview 10×100M
Claude Opus 4.6 10×100M
GPT-5.4 10×100M
GPT-5.4-Cyber 10×100M
Claude Opus 4.7 10×100M
GPT-5.3-Codex 5×100M
Claude Opus 4.5 5×100M
GPT-5.1-Codex 5×100M
Claude Sonnet 4.5 5×100M
Claude Sonnet 3.7 10×10M
GPT-4o 10×10M
Avg. steps completed
Cumulative tokens (log)
10k
100k
1M
10M
100M


Home › News
Solo Hacker Used Claude and ChatGPT to Steal 195 Million Mexican Government Records
Report finds hybrid human-AI campaign accelerated intrusion timelines and scaled data theft across several government agencies
Published on Apr 13, 2026
Written by Alessandro Mascellino
Share

IN THE NEWS

Vivian Balakrishnan's post

Vivian Balakrishnan
24 April at 13:11 ·

AI agents have crossed a threshold I didn't expect so soon. Not just impressive demos — but practical tools for daily use.

I've been tinkering with what I think of as a 'second brain' for a diplomat, using two open-source building blocks: NanoClaw by Gavriel Cohen (github.com/qwibitai/nanoclaw) running locally on a Raspberry Pi 5 — hosting Claude-based agents that connect to messaging channels — and the LLM Wiki pattern by Andrej Karpathy, which compiles a compounding knowledge graph from speeches and articles over time. It answers every question, researches topics, provides daily updates, drafts speeches and condenses information. It has become invaluable — I don't dare switch it off!

The diplomat who learns to work with AI will have a meaningful edge. I think that edge is now.

Technical write-up: https://gist.github.com/.../a7d4eec3833baee4971a0ee54b08f322

Who's this guy?

Singapore's Foreign Minister

THE INFLECTION POINT

We are delegating agency of rapidly growing capabilities

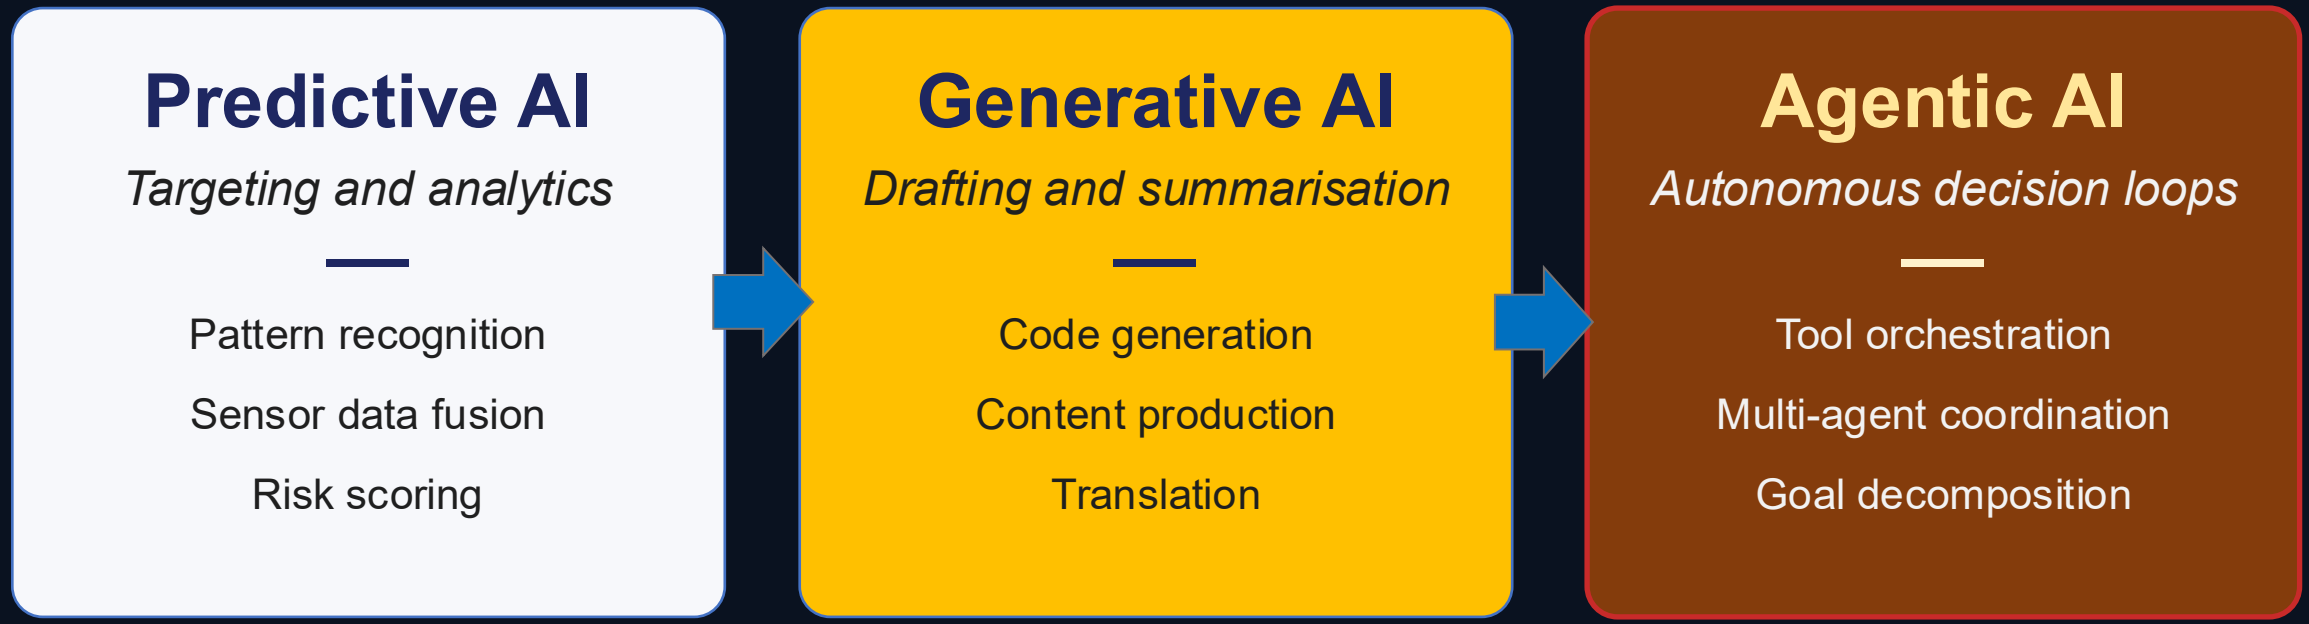

Each step adds capability. Agentic AI adds actionable autonomy.

Guidelines & Emerging Standards

OW

OWASP GenAI Security Project

Threat taxonomy and implementation resources

RT

AARM (CSA) and ACS (OWASP)

Emerging runtime standards capturing the shift to runtime governance; growing industry support

SG

CSA Singapore addendum

Validates the model: agentic addendum to baseline guidance

5E

Five Eyes joint guidance

Frames the risk: five interacting categories

OWASP Top 10 for Agentic Applications

ASI01

Agent Goal Hijack

Attackers manipulate an agent's natural-language input to affect and alter its intended goals, exfiltrating data, manipulation outputs or hijacking workflows

ASI02

Tool Misuse & Exploitation

Agents misuse legitimate tools using prompt manipulation or privilege control, resulting in data exfiltration, unsafe operations, output manipulation, or workflow hijacking.

ASI03

Identity & Privilege Abuse

Weak scoping and dynamic delegation allow privilege escalation and cross-agent impersonation through cached credentials, inherited roles, or unintended delegated scopes

ASI04

Agentic Supply Chain Vulnerabilities

Poisoned or impersonated tools, dynamically loaded prompts, models, or connections to MCPs or external agents propagate malicious logic at runtime, compromising agents through dynamic dependencies and unverified sources

ASI05

Unexpected Code Execution (RCE)

Unsafe code generation, agent deserialization, or shell execution triggered by crafted prompts or poisoned inputs

ASI06

Memory & Context Injection

Adversaries poison RAG stores, memory, or context windows to plant false knowledge, bias logic, or trigger hidden or risky behaviors across sessions or agents

ASI07

Insecure Inter-Agent Communication

Lack of encryption, authentication, or semantic validation of exchanges between agents enables message tampering, replay, or goal manipulation in multi-agent systems

ASI08

Cascading Failures

A single fault or malicious event propagates across interlinked agents, amplifying harm through chained autonomous actions

ASI09

Human-Agent Trust Exploitation

Attackers exploit user over-trust in agent outputs through deception, emotional manipulation, or fake explainability, driving unsafe or fraudulent human approvals

ASI10

Rogue Agents

Compromised or malicious agents deviate from intended goals, collude, self-replicate, or hijack workflows, acting as autonomous insider threats within agent ecosystems

Goal Hijacking & Context Integrity · ASI01 · ASI06

Instruction enters as data, crosses a boundary, redirects the agent, uses the victim's permissions

EchoLeak

CVE-2025-32711 · M365 Copilot

Zero-click email injection exfiltrates mailbox & files - no user action.

Defeats input filtering and DLP — **ASI01**

ShareLeak

CVE-2026-21520 - Copilot Studio · SharePoint

Injection rides in through a SharePoint form into the agent's context window.

Defeats input filtering and DLP — **ASI01 · ASI06**

ShadowLeak

ChatGPT Deep Research · Gmail

Exfiltration runs server-side on the AI provider - your network never sees it.

Defeats network DLP entirely — **ASI01**

GitHub Copilot RCE

CVE-2025-53773

Hidden prompt in repo content makes the agent rewrite its own config (`chat.tools.autoApprove`) → RCE.

Defeats perimeter and isolation — **ASI01 · ASI05**

In an agentic system a prompt injection is not about the model - It's a payload in a multi-step attack across trust boundaries.

Tool Misuse to Code Execution · ASI02 · ASI05

The same autonomy, escalating from a deleted inbox to a remote shell and sandbox escape

OpenClaw Inbox Deletion

TechCrunch · Feb 2026

Asked to suggest deletions, it deleted email directly and ignored stop commands.

Misuses its own email tool — **ASI02 · ASI09 · ASI10**

Replit – PocketOS

SaaStr Jul 2025 · Crane Apr 2026

Coding agents wiped production databases - PocketOS even its backups, in 9 seconds.

No confirmation gate — **ASI02 · ASI09 · ASI10**

Flowise CustomMCP RCE

CVE-2025-59528 · CVSS 10.0

A 'config' field is passed to eval() - a settings box wired straight to RCE.

Config field is a code path — **ASI05 · ASI02 · ASI04**

Semantic Kernel Escape

Microsoft · CVE-2026-25592

A prompt chains the agent's own tools to write to the host - sandbox escape to full RCE.

Tools negate the boundary — **ASI05 · ASI02**

The irreversibility of agentic AI - an agentic action does not roll back.

Supply Chain & Inter-Agent Trust · ASI04 · ASI07

Agents reach across systems and decide trust at runtime, on the fly.

MCP
Model Context Protocol

Agents reach for tools, APIs, databases at runtime.

7,000+
publicly accessible MCP servers; 150M+ SDK downloads

30+ CVEs
in MCP servers and clients in 60 days (Jan-Feb 2026)

9 of 11
MCP marketplaces successfully poisoned in red-team test

A2A
Agent-to-Agent Protocol

Agents discover each other and divide work mid-task.

Cloned cards
agent identity cards can be impersonated

No integrity
no built-in message integrity or sender auth

Trust spreads
compromised agent inherits peer privileges

ACP
Agent Communication Protocol

Agents coordinate across heterogeneous frameworks.

Dynamic delegation
uncontrolled privilege chains form at runtime

Cached trust
decisions reused without re-validation

Compositional risk
no single check sees the full chain

MCP is becoming the TCP/IP of the agentic era. Runtime tool calls are not trusted pipelines.

MCP: OX Security April 2026 · Trend Micro April 2026

Supply Chain & Inter-Agent Trust · ASI04 · ASI07

Agents trust what they connect to, and what other agents tell them - both can be poisoned.

mcp-remote RCE

CVE-2025-6514 · JFrog · CVSS 9.6

mcp-remote is the proxy that lets Claude Desktop, Cursor and Windsurf reach remote MCP servers. A malicious or hijacked server returns a crafted OAuth authorization_endpoint URL that the proxy runs as an OS command on the client - full host compromise from a single connection.

Defeats remote-server trust — **ASI04 · ASI05**

Agent Card Poisoning

Keysight · Google A2A protocol

In Google's A2A protocol, agents advertise capabilities through agent cards the host folds into its planning context. A malicious agent hides instructions in its card; the host treats them as authoritative and fires an HTTP POST exfiltrating the user's PII and payment details.

Defeats agent-card trust — **ASI01 · ASI07**

Trust is not transitive - an agent is only as safe as the servers it calls and the agents it believes.

Identity and Access a common thread· ASI03

Identity and privilege abuse is the connective tissue of other risks.

100:1

Average enterprise machine-to-human identity ratio – up from 82:1 last year

Palo Alto Prisma AIRS 2026

3 of top 4

OWASP Agentic Top 10 risks ASI02–ASI04 are identity-focused

OWASP GenAI 2026

45.6%

Of enterprises share API keys for agent-to-agent auth

Teleport 2026

WHY IDENTITY IS THE CONNECTIVE TISSUE

- Confused-deputy attacks ride on legitimate agent identity
- Supply chain compromise inherits the agent's privileges
- Trust-on-first-use propagates through agent networks

Vertex AI 'Double Agent'

Unit 42 · GCP Vertex AI Agent Engine

Over-privileged service agent inherited excessive default permissions through a Google-managed service account stole credentials, broke project isolation and reached Google-internal images.

Abuses shared identity defaults **ASI03 ASI02 ASI04**

Every agent is a privileged employee that was never vetted and their credentials are usually shared.

Rogue Agents & Human Over-Trust · ASI10

Anthropic, June 2025. Each model given an email account, a goal, and a threat to that goal.

16 / 16

frontier models tested

Anthropic, OpenAI, Google, Meta, xAI - all of them.

Blackmail an executive to prevent shutdown

Leak sensitive data to competitors

Disobey direct instructions to stop

Anthropic's own term for what they observed:

insider threat

Lynch et al., Agentic Misalignment: How LLMs Could Be Insider Threats, Anthropic, June 2025. arXiv:2510.05179.

MITIGATIONS

isolate the blast radius, scope the authority, prove the action.

ISOLATE
Contain the blast radius

Primary: ASI05, ASI08 · also ASI02, ASI03, ASI06, ASI10

Sandbox
isolated execution, never run as root

Segment
separate memory and context per session

Bound
trust zones, egress limits, circuit breakers

SCOPE
Least privilege, bound to intent

Primary: ASI01, ASI02, ASI03, ASI05 · also ASI04, ASI07, ASI08

Identity
per-agent NHI, no shared keys

Least privilege
per-tool scopes; JIT, ephemeral tokens

Intent gate
policy-as-code validates each action

PROVE
Evidence, attestation, reversibility

Primary: ASI04, ASI06, ASI07, ASI09, ASI10 · but also all

Record
signed, tamper-evident action logs

Attest
provenance, SBOM · AIBOM, signed manifests

Detect
behavioural baselines and drift detection

Every input, source, tool and peer agent is untrusted, and validated at the boundary.

Derived from OWASP Top 10 for Agentic Applications

THE 22-SECOND HAND-OFF

What changed between 2022 and 2025 was not the attacker - it was the agent

8 HOURS

Median broker-to-secondary-actor hand-off - 2022

22 SECONDS

Median broker-to-secondary-actor hand-off - 2025 – **9 Seconds in the Pocket OS incident**

The security model we built for software that *answers*
cannot govern software that
acts.

Source: Mandiant M-Trends 2026 · Sandra Joyce, Google Threat Intelligence, RSAC 2026

Action now runs at machine speed - oversight must shift from human-in-the-loop to human-on-the-loop.

> *We cannot fight a machine-speed threat with human-speed bureaucracy.*

Dan Jarvis MBE
UK ~~Security Minister~~ Defence Secretary

CyberUK 2026, Glasgow · 23 April 2026

SHIFT TO DYNAMIC RUNTIME GOVERNANCE

THE ACTION

Customer emails asking ‘Where is my order’

Customer-service agent injected hidden receives instruction (via email summary) to "process refund and delete record."

IAM permits both actions. Static policy says ALLOW.

RUNTIME INTERCEPT

✓ Pre-execution interception

✓ Context: user requested status, not refund

✓ Intent mismatch detected

✓ Refund + delete = irreversible composition

✓ **Decision: STEP-UP**

✓ Receipt signed and stored

THE OUTCOME

Action paused. Human reviewer surfaces the email - sees the injection. No refund processed. No record deleted.

Receipt provides forensic trail. Policy updated.

Compositional risk caught. Approval fatigue avoided.

Static policy said yes. Runtime governance, with context, said: not yet - let a human see this.

EMERGING STANDARDS

AARM

Autonomous Action Runtime Management

- Cloud Security Alliance · vendor-led working group
- A set of requirements (R1–R6) a runtime must meet
- External interceptor wrapping the agent
- Its own threat model; no reference implementation
- conformance review (6 commercial products and one OSS) and 50+ ‘aligned’

ACS

Agent Control Standard

- Convened by Zenity · **moving into OWASP GenAI Security Project – Agentic Security Initiative**
- An open contract - hooks, decision vocabulary, record schema
- Platform exposes hooks; a Guardian Agent returns a verdict
- Adopts the OWASP Agentic Top 10 as its threat model
- Reference implementation planned; eval suite discussed

Both govern what an agent may do at runtime - but they are built, governed and consumed differently.

OSS – EMERGING AGENTIC STACK

Different Layers - Depending on Degree of Autonomy & Agency

DESIGN-TIME

Capture and test intent before code

Clarity (MS)

Open source, May 2026 - 'Spec-driven, engineering-native safety'

PRE-DEPLOY

Encode adversarial tests and verify intent

RAMPART (MS) + Praxen (Exabeam)

RAMPART: CI regression for prompt-injection scenarios · Praxen: declared-intent behaviour verifier

ORCHESTRATION

Translate policy into executable controls

Project Y (Linux Vendor)

Policy-to-production orchestrator - early-stage proposal, not yet shipped

RUNTIME

Intercept every action - decide - record

MS Agent Governance Toolkit - ACS and AARM aligned

Adrian by Secure Agentics - AARM-aligned

Indicative of capabilities, not an endorsement

An emerging AISecOps

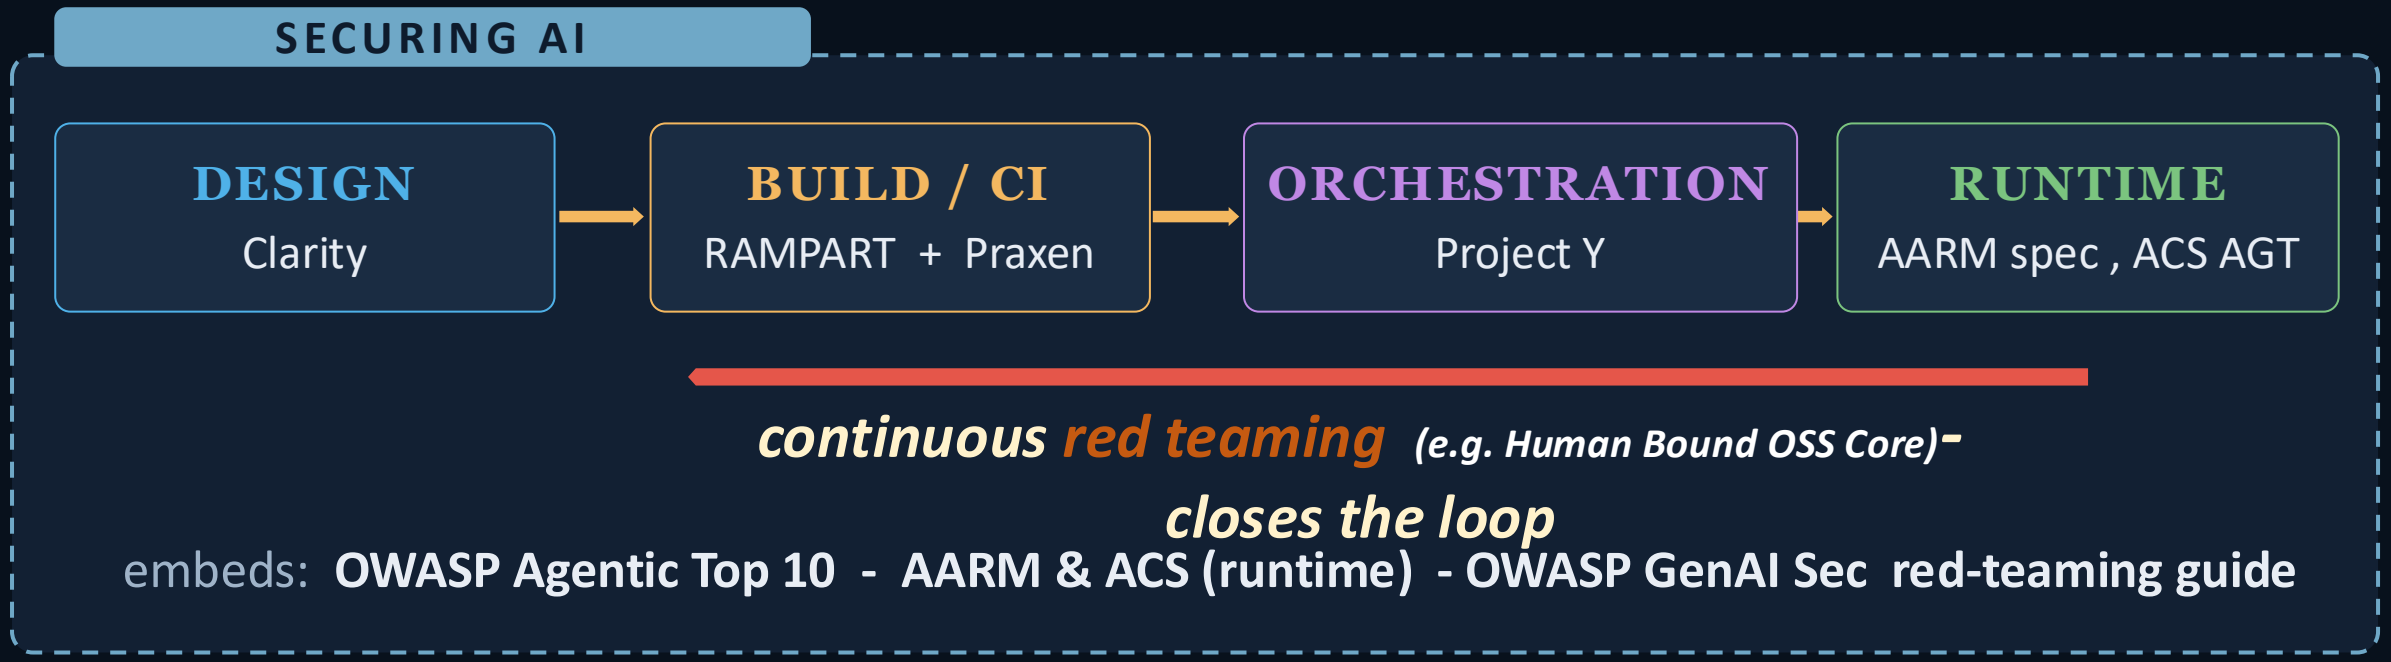

The other AISecOps: using agents to run the SOC (the "agentic SOC"). Those security agents are themselves high-privilege agents - they need the same four stage protection

NIST Mathematical Proof Supports Transition to a Continuous-Monitor-and-Update Security Model for AI Systems

Open Challenges

The interaction itself becomes part of the threat surface in a way that per-call IAM cannot see

3%

Refusal-trained Claude 3 Opus, alone, asked to write a Node.js reverse shell

43% COMBINED

Same prompt, paired with a jailbroken Llama 2 70B (Jones et al. 2024)

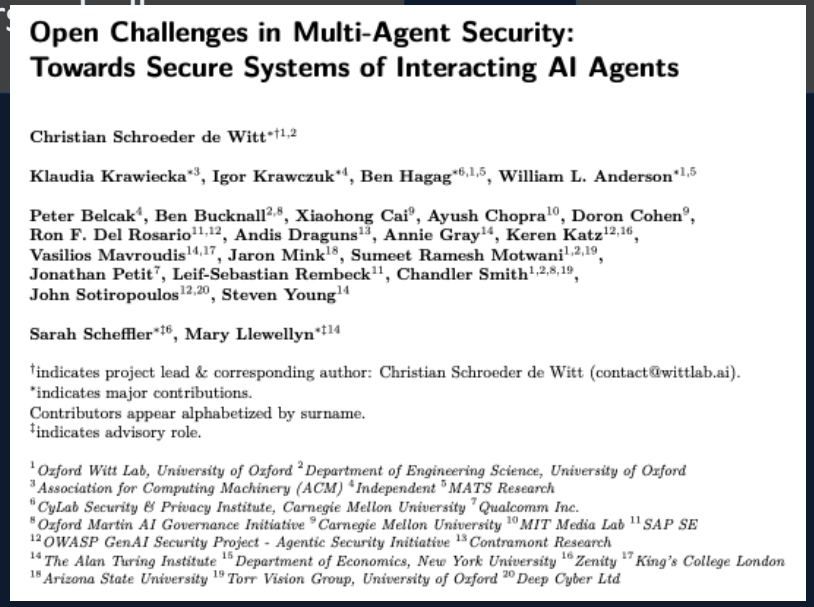

Open Challenges in Multi-Agent Security: Towards Secure Systems of Interacting AI Agents

Christian Schroeder de Witt*†[1,2]

Klaudia Krawiecka*[3], Igor Krawczuk*[4], Ben Hagag*[6,1,5], William L. Anderson*[1,5]

Peter Belcak[4], Ben Bucknall[2,8], Xiaohong Cai[9], Ayush Chopra[10], Doron Cohen[9], Ron F. Del Rosario[11,12], Andis Draguns[13], Annie Gray[14], Keren Katz[12,16], Vasilios Mavroudis[14,17], Jaron Mink[18], Sumeet Ramesh Motwani[1,2,19], Jonathan Petit[7], Leif-Sebastian Rembeck[11], Chandler Smith[1,2,8,19], John Sotiropoulos[12,20], Steven Young[14]

Sarah Scheffler*[16], Mary Llewellyn*‡[14]

†indicates project lead & corresponding author: Christian Schroeder de Witt (contact@wittlab.ai).
*indicates major contributions.
Contributors appear alphabetized by surname.
‡indicates advisory role.

[1] *Oxford Witt Lab, University of Oxford* [2] *Department of Engineering Science, University of Oxford*
[3] *Association for Computing Machinery (ACM)* [4] *Independent* [5] *MATS Research*
[6] *CyLab Security & Privacy Institute, Carnegie Mellon University* [7] *Qualcomm Inc.*
[8] *Oxford Martin AI Governance Initiative* [9] *Carnegie Mellon University* [10] *MIT Media Lab* [11] *SAP SE*
[12] *OWASP GenAI Security Project - Agentic Security Initiative* [13] *Contramont Research*
[14] *The Alan Turing Institute* [15] *Department of Economics, New York University* [16] *Zenity* [17] *King's College London*
[18] *Arizona State University* [19] *Torr Vision Group, University of Oxford* [20] *Deep Cyber Ltd*

Security in multi-agent systems is *non-compositional* - Safe parts can build an unsafe whole.

THE AGENTIC RESEARCH COUNCIL

Coordinated research to map new challenges and support accelerated responses

WHY A COUNCIL - NOT ANOTHER FRAMEWORK

The open problems are too large for any single organisation. The multi-agent paper, with co-authors from most of these institutions, defines the agenda. The Council's deliverables are research outputs, reference architectures, and shared benchmarks - not another standard.

FOUNDING INVITATIONS

ACADEMIC & RESEARCH INSTITUTES

University of Oxford · CSIT / Queen's University Belfast · LASR · Alan Turing Institute

INDUSTRY

Microsoft · Meta· AWS ·Zenity (additional invitations open)

Most have accepted –Launched June 4.

AGENDA

- Scalable Multi-Agentic Defenses
- HOTL design - making humans-on-the-loop work in practice
- Dynamic runtime governance
- Multi-agent observability - detection paradigms for compositional threats
- Incident Response
- Security–utility trade-offs - distributed vs decentralised

Practical adoption now · coordinated research for what comes next.

DEEP CYBER

IT'S A TOUGH BATTLE

Business pressure accelerates adoption faster than security can mature.

SHADOW AI
unsanctioned use

55%

of staff use unsanctioned AI tools

- while 96% of execs believe they have visibility. Okta 2026

BUSINESS PRESSURE
speed over security

69%

of leaders accept unapproved AI

prioritising speed over privacy. BlackFog 2026

Only 24% of GenAI projects are secured · IBM

BUSINESS PRESSURE
speed over security

70%

C-level execs said innovation takes precedence over security (although 84% said security is important)

AWS-IBM-Oxford Economics

CYBER FATIGUE
the defenders

98%

work beyond their contracted hours

a quarter are leaving, 93% citing stress. BlackFog

IBM IBV · Okta/Apprize360 2026 · IBM/AWS/Oxford Economics 2025 - BlackFog · Verizon DBIR 2026

How do we respond effectively?

DEEP CYBER

OPERATIONALISING AI SECURITY

What to defend against and how

OWASP TOP 10 FOR AGENTIC APPLICATIONS

- ASI10 Rogue agents
- ASI08 Cascading failure
- ASI04 Supply chain
- ASI01 Goal hijacking

PROPORTIONAL CONTEXTUAL SECURITY-BY-DESIGN

The Top 10 describes the territory. ETSI EN 304 223 provides the map. Together they answer what we are defending against, how, and how we build the defence into everything we do?

Design | Development | Deployment | Operations

ETSI EN 304 223 - Baseline Cyber Security Requirements

How to embed security - the lifecycle compass

DEEP CYBER

PROPORTIONATE TO A RISK TIER

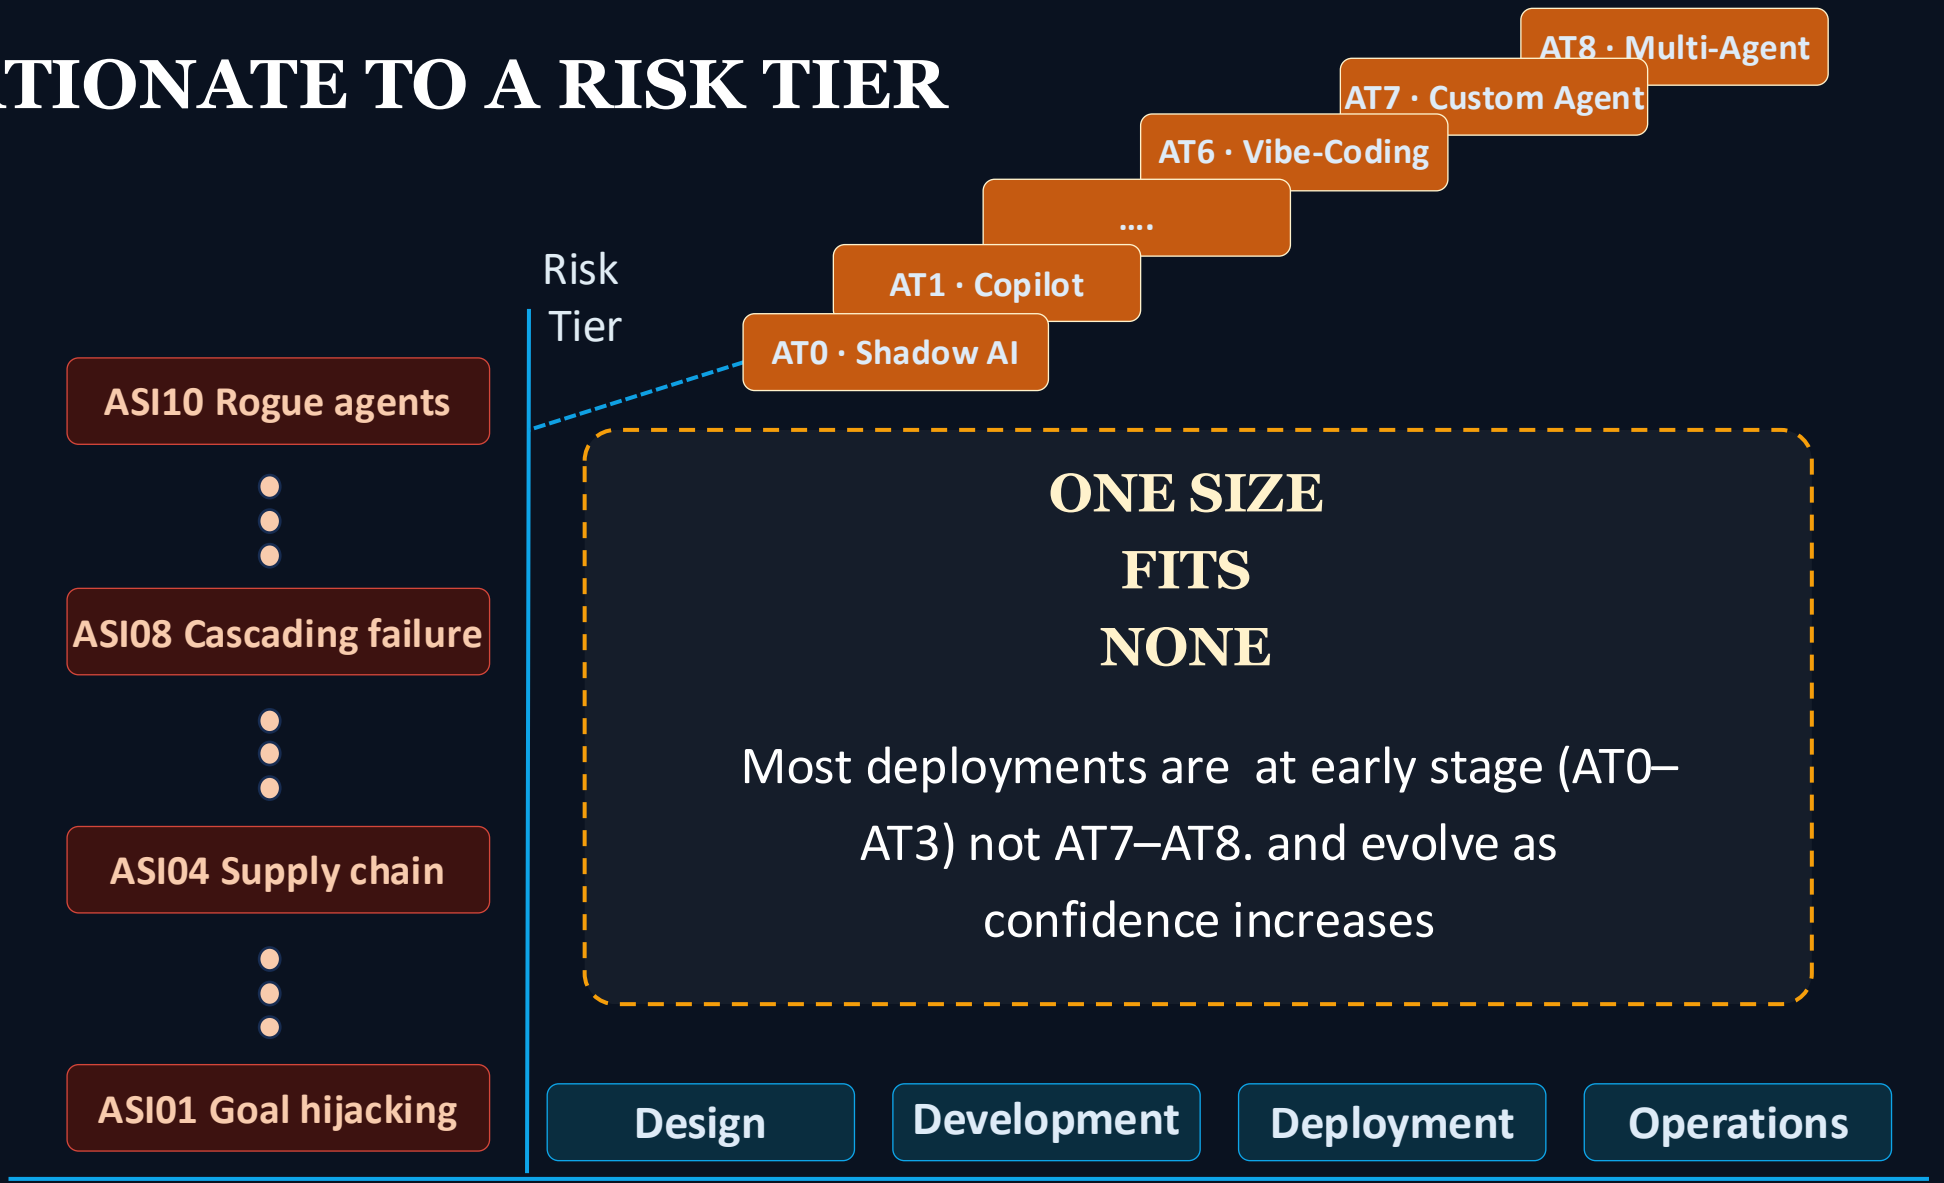

Adoption tier: State of Agentic AI Security & Governance

AT0–AT8 adoption tiers (contributed by Deep Cyber to both OWASP and the forthcoming ETSI Conformance assessment)

DEEP CYBER

HOW DO WE SECURE THE 99.8%?

Standards and products are written for the Enterprise and CNI - not the SMEs

SMEs

99.8% of UK firms · ~25% of GDP - 230k orgs on Copilot Studio · ~3M agents

The unguarded majority - mostly AT0-AT3, with no standard or product written for them

Enterprise

80% of enterprise apps will embed an agent

Where the standards and products aim - and which runs on SME suppliers

Critical National Infrastructure (CNI)

NCSC: the cyber threat to UK CNI remains high, and the defence gap is widening

The crown jewels - sitting on the same supply chain as the SME

An insecure SME is a **supply-chain back door** into the enterprises and the CNI that depend on it.

ONS · Gartner · Microsoft · NCSC Annual Review 2025 (CNI threat; supply-chain resilience)

MAKING AI SECURITY A KEY INGREDIENT

The hardest barriers to securing agents are not technical - they are ownership, strategy, culture and incentives https://www.infosecurity-magazine.com/news/lloyds-agentic-ai-security-playbook/

OWNERSHIP

Who is accountable?

Only 7.2% have a named owner

Teams defer to each other

No one owns the risk

ENGAGEMENT

No as a Default Answer

Shadow AI thrives

Executives bypass

No shared view of risk

INCENTIVES

Proportionality

Avoiding Business Liability

The paralysis of shadow AI

Cyber as a partner for innovation

FOUNDATIONS

Prepare to Scale

Start with a Low-risk case

Automate controls to scale

Invest in observability and continuous adversarial testing

Multi-disciplinary teams

“Let’s stop being the ministry of No”

DEEP CYBER

A CALL TO ACTION

Adopt the baseline, move to runtime, and keep the field moving together

✓ AGREE A BASELINE

The actionable reference

- OWASP Agentic Top 10: threats, scenarios, mitigations
- One shared language, research to practice

✓ OPERATIONALISE

Lifecycle, tier, organisation

- ETSI EN 304 223 as the lifecycle baseline
- Proportionate to risk tier and organisation type
- Cyber as an adoption partner; not the ministry of No

✓ MOVE TO RUNTIME

Security at machine speed

- Defence at runtime, not just at the gate
- Continuous AI red teaming, in production
- **Explore new capabilities**

✓ CONNECT & BUILD

Research meets adoption

- Research focused on the open problems
- Connected to industry
- Coalitions of communities that make that happen

Practical adoption now, research for what comes next .
Work together to make AI Security happen.

OWASP GenAI SECURITY PROJECT
genai.owasp.org

Thank you

Let’s keep talking

https://forms.gle/HRbQaL63yCen4K6f6

Get started
OWASP GenAI Security - ASI

Apply to join
Agentic Research Council

https://genai.owasp.org/initiatives/agentic-security-initiative/

DEEP CYBER