OWASP GenAI SECURITY PROJECT

OWASP Top 10 for Agentic Applications Global Kick-off

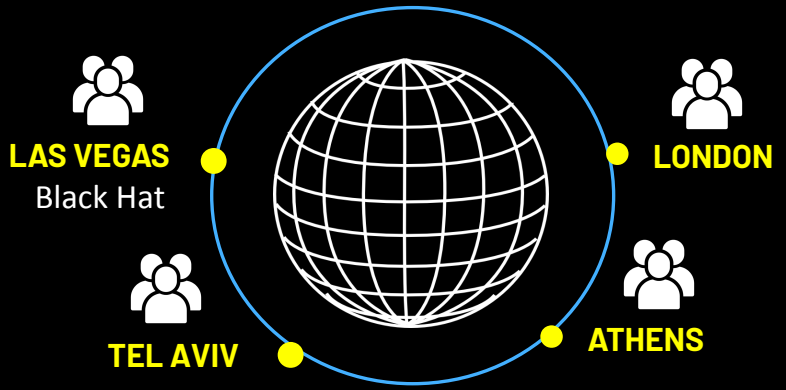

Welcome!

Our Agenda Today

- ✔ Introduce the Agentic Security Initiative
- ✔ Preview the public draft of our OWASP Top 10 for Agentic applications
- ✔ Explain how you can become part of our public consultation and help shape the final release

Agentic Security Initiative

Part of the GenAI Security Project

Guidance & Resources

- Initiative Overview Blog
- Agentic AI – Threats and Mitigations
- Agentic AI Threat Navigator
- Multi-Agentic System Threat Modeling

Get Involved

- Slack: team-llm-autonomous-agents
- Initiative Charter

Initiative Lead(s)

- John Sotiropoulos
- Ron F. Del Rosario

- Join the OWASP Slack Workspace

The Top 10 for LLM for LLM and Gen AI has become one of many initiatives the project now leads – we are one of them

https://genai.owasp.org/initiatives/#agenticinitiative

Expert-backed Community

Apostol Vassilev
Adversarial AI Lead at NIST

Hyrum Anderson
CAMLIS Cofounder, AI Security Pioneer, CISCO

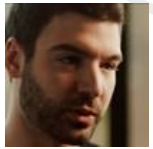

Vasilios Mavroudis
Principal Research Scientist, Allan Turing Institute

Josh Collier
Principal Researcher, Allan Turing Institute

Alejandro Saucedo
Linux Foundation, Advisor @ UN, EU, ACM

Chris Hughes
Host of Resilient Cyber, Cyber Security Author

Michael Burgundy
OWASP AISVS Co-Chair & Zenity CTO

Peter Bryan
Principal AI Security Research Lead- AI Red Team Microsoft

Egor Pushkin
Chief Architect, Data and AI at Oracle Cloud

Matt Sanner
Security Leader at AWS, Elected Board Member at CoSAI

Dan Jones
Principal Researcher AI Red Team at Microsoft

Steve Wilson
GenAI Security Project Founder &Chair, Chief AI and Product Officer at Exabeam

Secure Agentic Lifecycle

Code Samples and Developer Validation

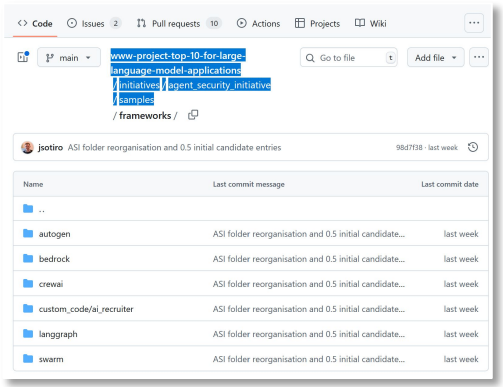

New! Agengic CTF App

Product Adoption

18 Solutions

OWASP GenAI - ASI - Agentic Threat and Mitigations Taxonomy Product Support

CORTEX CLOUD · NOMA · ai+ me · Microsoft · Pillar · TREND · Straiker · GuardionAI · ADVERSA · MINDGARD · Citadel AI · tenable AI · splx AI · NRI SECURE · HUMAN · unbound

Agentic Threat and Mitigations Coverage

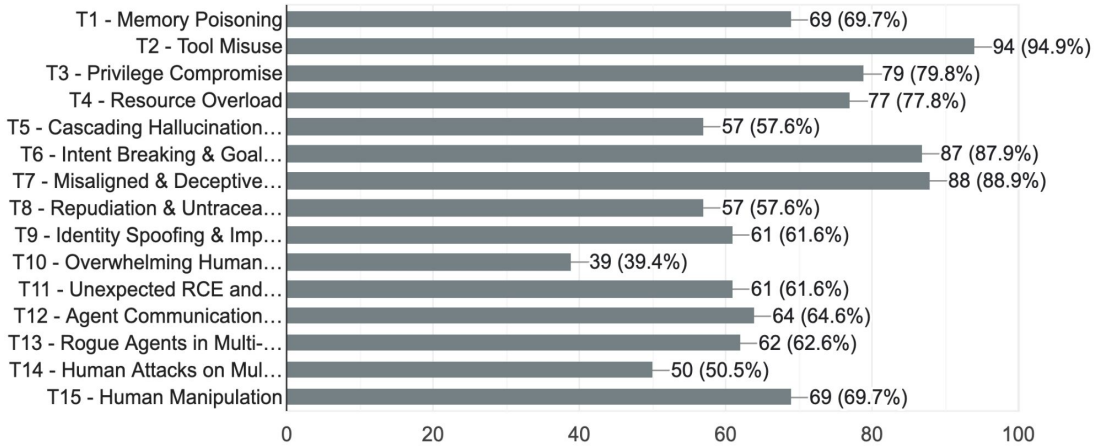

DRAFT Not For Release

Active Feedback

Large Scale Feedback Surveys, Open Workshops, Expert Panels

Including early adopters, cross-references, and dedicated research

SAP

kainos®

Microsoft

Taxonomy of Failure Mode in Agentic AI Systems

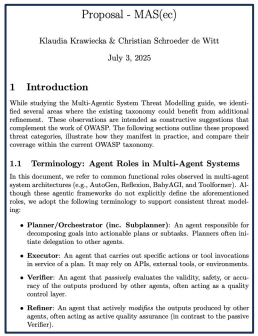

Proposal - MAS(ec)

Klaudia Krawiecka & Christian Schroeder de Witt

July 3, 2025

1 Introduction

While studying the Multi-Agentic System Threat Modelling guide, we identified several areas where the existing taxonomy could benefit from additional refinement. These observations are intended as constructive suggestions that complement the work of OWASP. The following sections outline these proposed threat categories, illustrate how they manifest in practice, and compare their coverage within the current OWASP taxonomy.

1.1 Terminology: Agent Roles in Multi-Agent Systems

In this document, we refer to common functional roles observed in multi-agent system architectures (e.g., AutoGen, Reflexion, BabyAGI, and Toolformer). Although these agentic frameworks do not explicitly define the aforementioned roles, we adopt the following terminology to support consistent threat modeling:

- **Planner/Orchestrator (inc. Subplanner)**: An agent responsible for decomposing goals into actionable plans or subtasks. Planners often initiate delegation to other agents.
- **Executor**: An agent that carries out specific actions or tool invocations in service of a plan. It may rely on APIs, external tools, or environments.
- **Verifier**: An agent that passively evaluates the validity, safety, or accuracy of the outputs produced by other agents, often acting as a quality control layer.
- **Refiner**: An agent that actively modifies the outputs produced by other agents, often acting as active quality assurance (in contrast to the passive Verifier).

Recent Changes

🔗 **Agentic Protocols** like **MCP, A2A, and ACP** define how agents communicate, delegate, and trust each other — but often lack strong attestation, authentication, or context control.

🌐 **Distributed Multi-Agent Systems** (e.g. **Agent Cards**, **Registries**) introduce **decentralized execution** and **dynamic task assignment**, which attackers can spoof or hijack.

⚙️ **Supply Chain Complexity**: Agents now become dynamic components depending on tools, models, plugins — blurring the line between "runtime behaviour" and "pre-baked vulnerability".

💥 **Real-World Exploits** are already exposing these weaknesses, demonstrating:

- Goal Manipulation and Agent Hijacking
- Rogue agent impersonation
- Toolchain-based privilege misuse
- And more

Security Standards At Pace

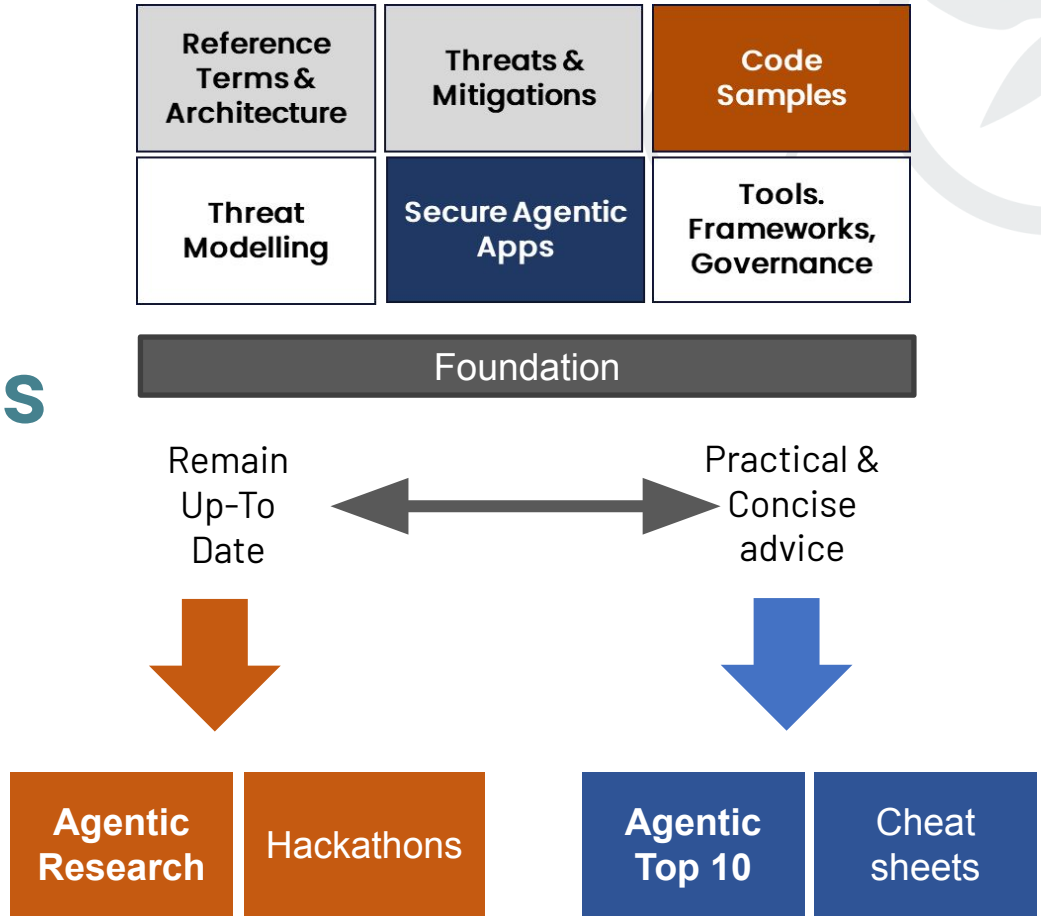

Why a Top 10?


Agentic AI – Threats and Mitigations
Multi-Agentic system Threat Modelling Guide
Securing Agentic Applications Guide
State of Agentic AI Security and Governance
270 Pages


Where do I start?


What is the top AI security challenge in your organization?
WIZ
Lack of AI expertise in the security organization
31%
Incorporating built-in guardrails and check
17%
Dealing with shadow AI
14%
Safeguarding sensitive training data
17%
Continuously monitoring for unusual activities
7%
Securing access to GenAI models
6%
Detecting and removing attack paths to models
6%
Testing GenAI pipelines
5%
0%
10%
20%
30%
40%

Our Threat Taxonomy Is Our Baseline

AGENCY & REASONING

- **T06.** Intent Breaking and Goal Manipulation
- **T07.** Misaligned and Deceptive Behaviours
- **T08.**Repudiation and Untraceability

MEMORY AND CONTEXT

- **T01.** Memory Poisoning
- **T04** Cascading Hallucinations

TOOLS & EXECUTION

- **T02.** Tool Misuse
- **T03.** Privilege Compromise
- **T04.** Resource Overload
- **T11**. Unexpected RCE and Code Attacks

IDENITY & AUTHENTICATIO N

- **T09** Identity Spoofing and Impersonation

HUMAN ENGAGMENT

- **T10** Overwhelming Human-in-the-Loo p (HITL)
- **T15,** Human Trust Manipulation

MULTI-AGENCY

- **T12** Agent Communication Poisoning
- **T13.** Rogue Agents
- **T14** Human Attacks

How did we start our Top 10?

- Allow new entries
- Supports further development of the baseline taxonomy
- Preserves investment in our Threats and Mitigations

Top Ten for Agentic Applications 0.5

ASI01	**Agent Behaviour Hijack**	Manipulating an agent's goals plans to pursue attacker-aligned objectives.
ASI02	**Tool Misuse and Exploitation**	Tricking agents into using their tools in harmful or unintended ways.
ASI03	**Identity & Privilege Abuse**	Impersonating agents or escalating access through identity or auth and access permissions weaknesses.
ASI04	**Agentic Supply Chain Vulnerabilities**	Introducing insecure models, agents , tools or artefacts compromise agent integrity.
ASI05	**Unexpected Code Execution (RCE)**	Triggering unauthorized or unsafe code execution through agent behaviors.
ASI06	**Memory & Context Poisoning**	Corrupting agent memory or context to distort reasoning and decision-making.
ASI07	**Insecure Inter-Agent Communication**	Poisoning messages or abusing protocols between agents to alter behavior.
ASI08	**Cascading Failures**	Faults or hallucinations propagate through agents, causing compounded failures.
ASI09	**Human-Agent Trust Exploitation**	Exploiting over-trust or fatigue in human oversight to enable agent misuse including model deceptive behaviours
ASI10	**Rogue Agents**	Malicious or compromised agents acting autonomously to deceive, disrupt, or exfiltrate.

Example Entry

https://github.com/OWASP/www-project-top-10-for-large-language-model-applications/tree/main/initiatives/agent_security_initiative/agentic-top-10/0.5-initial-candidates

What's next?

This is deliberately an early draft to start the conversation and involve you

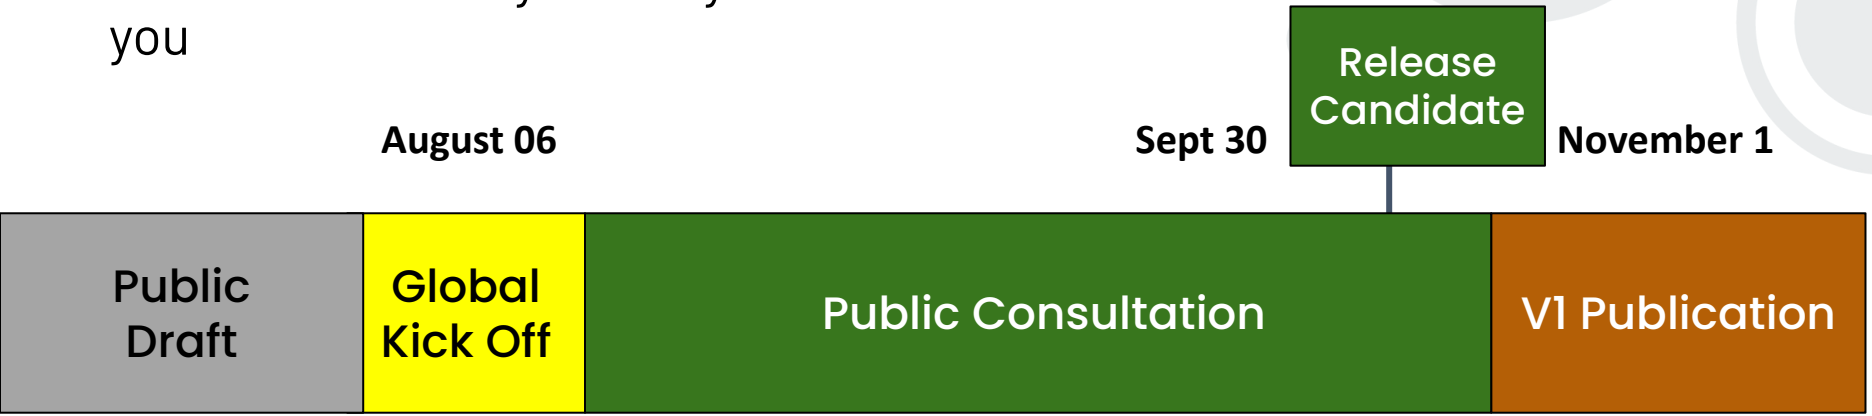

- New Submissions
- Working Groups per Entry
- Alignment Reviews (AIVSS Core Risks, MAS, other)
- Review Workshops
- Exploit and Incidence Research
- Survey & Votes

Be Part of our Effort

- Register interest to
  - Be notified for news, surveys, and votes
  - Be part of our Working Groups and Review Sessions
    https://forms.gle/QQTUQgCA8KfTYKwy7
-
- Submit Amendments or new entries via GitHub
  - https://github.com/OWASP/www-project-top-10-for-large-language-model-applications/tree/main/initiatives/agent_security_initiative
-
- Join the team
  - Slack channel:
  - Weekly calls every Monday 5:30pm-6:30pm
  - Details on how to join
    https://genai.owasp.org/initiatives/#agenticinitiative

OWASP GenAI SECURITY PROJECT

Thank you! Register here for next steps