

Source dossier

Inspection is not a shield — Tenable and OpenAI vet the agent marketplace, and Article 25(4) exempts the suppliers

Edition 2026-09-05 · narrative: Agent Incidents and Liability · compiled 2026-09-05

Claims index

Claim	Statement	Status	Sources
CLAIM-2026-09-05-01	Tenable and OpenAI announced the CyberAgents Exchange AI Inspector on 3 September 2026 (wire timestamp 17:57 ET, dateline Columbia, MD) — a security review for community-contributed AI agents, skills, MCP servers and multi-agent playbooks, combining frontier assessment using OpenAI GPT cyber models, skills inspection powered by Tenable One AI Exposure, and expert review by Tenable researchers. It is expected to be available in September 2026.	verified	tenable2026tenable , globenewswi2026 wire
CLAIM-2026-09-05-02	As of the 3 September 2026 release, the CyberAgents Exchange listed more than 100 community-submitted AI components; Tenable states more than 40,000 customers globally.	verified	tenable2026tenable , globenewswi2026 wire
CLAIM-2026-09-05-03	The CyberAgents Exchange launched on 4 August 2026 at Black Hat USA with more than 50 AI components across agents, skills, MCP servers and playbooks, as a free, open-source community with no fees for listing or using agents.	verified	tenable2026tenable b
CLAIM-2026-09-05-04	AI Act Article 25(1)(a): a distributor, importer, deployer or other third party is considered a provider of a high-risk AI system, subject to the Article 16 provider obligations, where 'they put their name or trademark on a high-risk AI system already placed on the market or put into service, without prejudice to contractual arrangements stipulating that the obligations are otherwise allocated'.	verified	europeancom2024a rticle, euartificia20 24responsib
CLAIM-2026-09-05-05	AI Act Article 25(4), first sentence: 'The provider of a high-risk AI system and the third party that supplies an AI system, tools, services, components, or processes that are used or integrated in a high-risk AI system shall, by written agreement, specify the necessary information, capabilities, technical access and other assistance based on the generally acknowledged state of the art, in order to enable the provider of the high-risk AI system to fully comply with the obligations set out in this Regulation.'	verified	europeancom2024a rticle, euartificia20 24responsib

Claim	Statement	Status	Sources
CLAIM-2026-09-05-06	AI Act Article 25(4), second sentence: 'This paragraph shall not apply to third parties making accessible to the public tools, services, processes, or components, other than general-purpose AI models, under a free and open-source licence.' The written-agreement duty therefore does not reach a supplier distributing an agent, skill or MCP server under a free and open-source licence.	verified	euartificia2024responsib, europeancom2024article
CLAIM-2026-09-05-07	AI Act Article 25(2): where a third party becomes the provider under Article 25(1), the initial provider that placed the system on the market is no longer considered the provider of that specific system, and must closely cooperate with the new provider — making available technical documentation for conformity assessment, disclosing known limitations and failure modes, and giving targeted technical access for testing.	verified	europeancom2024article, euartificia2024responsib

Only claims marked **verified** may carry a post. **single-source** may appear as hedged context; **disputed** must not be published at all.

Tenable Uses OpenAI GPT Cyber Models to Help Defenders Inspect Community-Built AI Components

PRIMARY

Publisher	Tenable
Published	2026-09-03
Retrieved	2026-09-05
Cite key	tenable2026tenable
URL	https://www.tenable.com/press-releases/tenable-uses-openai-gpt-cyber-models-to-help-defenders-inspect-community-built-ai-components

Assessment. The announcing party's own release. Dateline 3 September 2026, Columbia, MD, issued around the OpenAI Cyber Summit. Every figure and component-type list in this edition is taken from here or from the identical wire copy, not from coverage. Note the discrepancy corrected in this edition: aggregators date the announcement 4 September; the release and the wire timestamp are 3 September.

Claims drawn from this source. CLAIM-2026-09-05-01, CLAIM-2026-09-05-02

Dateline. September 3, 2026, Columbia, MD.

Tenable announced a collaboration with OpenAI to create the CyberAgents Exchange AI Inspector, a new security review process for select AI agents, skills, MCP servers and multi-agent playbooks on the CyberAgents Exchange, powered by Tenable.

The AI Inspector combines frontier assessment using OpenAI GPT cyber models with skills inspection powered by Tenable One AI Exposure and expert review from Tenable researchers, to help security teams safely accelerate enterprise adoption of agentic AI.

Figures. Over 100 community-submitted AI components in the CyberAgents Exchange. Over 40,000 Tenable customers globally.

Availability. Exchange Inspector is expected to be available in September.

Quote. Eric Doerr, Chief Product Officer at Tenable: "Agentic AI will only reach its potential in the enterprise if security teams can trust the components being introduced into their environments." He describes the goal as building a more rigorous way to inspect community-built AI components.

Limitations. The release makes no warranty about the outcome of the review. It states that the process helps teams safely accelerate adoption and lets security teams trust components before deployment. It does not address allocation of regulatory responsibility, and does not mention licensing.

Wire copy: Tenable Uses OpenAI GPT Cyber Models to Help Defenders Inspect Community-Built AI Components

PRIMARY

Publisher GlobeNewswire

Published 2026-09-03

Retrieved 2026-09-05

Cite key globenewswi2026wire

URL <https://www.globenewswire.com/news-release/2026/09/03/3356323/0/en/tenable-uses-open-ai-gpt-cyber-models-to-help-defenders-inspect-community-built-ai-components.html>

Assessment. Wire distribution of the identical release, retrieved independently of the Tenable page as a cross-check on the dateline and the figures. Timestamp 3 September 2026, 17:57 ET. The two retrievals agree on component types, the September availability, and the Doerr quote. It is the timestamp that settles the coverage-date error.

Claims drawn from this source. CLAIM-2026-09-05-01, CLAIM-2026-09-05-02

Dateline. September 03, 2026, 17:57 ET.

Security review combining frontier assessment, skills inspection powered by Tenable One AI Exposure, and expert researcher evaluation.

Component types reviewed. AI agents; skills; MCP servers; multi-agent playbooks.

Model naming. "OpenAI GPT cyber models" — no specific model name or version is given.

Availability. Expected September 2026.

Quote. "Agentic AI will only reach its potential in the enterprise if security teams can trust the components being introduced into their environments" — Eric Doerr, Chief Product Officer, Tenable.

Caveat. The release does not state what the review guarantees. "Safely accelerate enterprise adoption" is promotional framing, not a warranty that risk or vulnerabilities are eliminated.

Tenable Launches Industry's First Open-Source AI Agent Exchange

PRIMARY

Publisher	Tenable
Published	2026-08-04
Retrieved	2026-09-05
Cite key	tenable2026tenableb
URL	https://www.tenable.com/press-releases/tenable-launches-industrys-first-open-source-ai-agent-exchange

Assessment. The exchange's own launch release, one month earlier. It supplies the comparator for the 100+ figure and, decisively for this edition, establishes that the components are distributed free and under open-source licences — which is what puts them inside the Article 25(4) carve-out.

Claims drawn from this source. CLAIM-2026-09-05-03

Dateline. August 4, 2026, Las Vegas — Black Hat USA, Booth #2639.

Scale at launch. More than 50 AI components.

Component types. Agents (e.g. Navi); MCP Servers (e.g. SentinelOne Purple AI); Skills (e.g. SOC-Hunter, Recorded Future MITRE APT Attack Path Analysis); Playbooks (e.g. The Hounds Pack).

Licence and cost. Open-source licences. Described as a free, open-source community with no fees for listing or using agents.

Quotes. Vlad Korsunsky, Tenable CTO: "With the CyberAgents Exchange, we're creating a collaborative 'town square' where cybersecurity practitioners can build, test, improve and share the best in agentic defense." Seth Fogie, Director of Security, Baptist Memorial Health Care: "Its core philosophy of working together is what made the threat intelligence community so invaluable." Braden Preston, VP of Product Management, SentinelOne. Jamie Zajac, CPO, Recorded Future: "Intelligence doesn't just make agents smarter. It makes them traceable, auditable and repeatable."

Note. The launch release does not name a specific SPDX licence identifier.

tenable/cyberagents-exchange — community-contributed agents, skills, MCP servers, and playbooks

PRIMARY

Publisher	GitHub
Published	2026-08-04
Retrieved	2026-09-05
Cite key	github2026tenablecy
URL	https://github.com/tenable/cyberagents-exchange

Assessment. The repository behind the exchange. Cited for the practical instruction in the edition — that a commit is pinnable and a marketplace listing is not. Used for discovery and reference only; no figure in this edition is drawn from it.

Claims drawn from this source. CLAIM-2026-09-05-03

Full text not archived — discovery/headline reference only.

Article 25: Responsibilities along the AI value chain

PRIMARY

Publisher European Commission AI Act Service Desk

Published 2024-07-12

Retrieved 2026-09-05

Cite key europeancom2024article

URL <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-25>

Assessment. The Commission's own reproduction of the article. Used for the verbatim text of 25(1)(a) and the first sentence of 25(4). The second sentence of 25(4) — the free-and-open-source carve-out that the whole edition turns on — was taken verbatim from the artificialintelligenceact.eu reproduction and is consistent with the Commission text.

Claims drawn from this source. CLAIM-2026-09-05-04, CLAIM-2026-09-05-05, CLAIM-2026-09-05-06, CLAIM-2026-09-05-07

Article 25(1)(a), verbatim. "they put their name or trademark on a high-risk AI system already placed on the market or put into service, without prejudice to contractual arrangements stipulating that the obligations are otherwise allocated"

Article 25(4), first sentence, verbatim. "The provider of a high-risk AI system and the third party that supplies an AI system, tools, services, components, or processes that are used or integrated in a high-risk AI system shall, by written agreement, specify the necessary information, capabilities, technical access and other assistance based on the generally acknowledged state of the art, in order to enable the provider of the high-risk AI system to fully comply with the obligations set out in this Regulation."

Article 25(2), substance. Where Article 25(1) applies, the initial provider that placed the system on the market is no longer considered the provider of that specific AI system. The initial provider must closely cooperate with the new provider and make available the information and technical access necessary to fulfil the obligations, in particular conformity assessment: technical documentation, disclosure of known limitations and failure modes, and targeted technical access for testing. It does not apply where the initial provider explicitly excluded the system from becoming high-risk.

Article 25(5). Without prejudice to intellectual property rights, confidential business information and trade secrets.

Responsibilities along the AI value chain - Article 25

SECONDARY

Publisher EU Artificial Intelligence Act

Published 2024-07-12

Retrieved 2026-09-05

Cite key euartificia2024responsib

URL <https://artificialintelligenceact.eu/article/25/>

Assessment. Reused cite key from an earlier edition. Retrieved again for the second sentence of Article 25(4), the open-source carve-out, which is the load-bearing quotation in this edition. Secondary grade: it is a faithful reproduction, not the Official Journal, so it is cross-checked against the Commission service desk text above.

Claims drawn from this source. CLAIM-2026-09-05-04, CLAIM-2026-09-05-05, CLAIM-2026-09-05-06

Article 25(4), second sentence, verbatim. "This paragraph shall not apply to third parties making accessible to the public tools, services, processes, or components, other than general-purpose AI models, under a free and open-source licence."

The carve-out is drafted as an exemption from the paragraph, not as a reallocation of the obligation. The provider's own duties under Article 16 are untouched by it. What the provider loses is the counterparty the paragraph would otherwise have given them.