

AGENT-AWARE VULNERABILITY MANAGEMENT

Rank each flaw by what the **agent** can do with it

Six inputs, the OWASP AIVSS score, and where the new ISO/IEC 27090 fits.

A vulnerability score describes the flaw. It says nothing about the agent behind it.

The same prompt injection leaks an answer in a read-only assistant. In an agent that can pay suppliers, act without review and remember what it was told, it can move money and keep doing it. CVSS gives both the same number.

Six inputs to an agent-aware priority

- 1** **Exploitability**
CVSS v4.0, EPSS, KEV, red team results
- 2** **Autonomy**
Can it act without a person approving?
- 3** **Permissions**
Whose authority, and how far it reaches
- 4** **Tools**
What it can call, and with what privilege
- 5** **Memory persistence**
Does what it is told survive the session?
- 6** **Business impact**
What is at the end of the agent's reach

One CVSS 4.3 finding, scored with the agent factors

Read-only assistant

5.1

Factor sum 1.5 of 10. Stays
Medium.

Finance operations agent

7.9

Factor sum 6.5 of 10.
Autonomy, write tools and
long-term memory push it to
High.

Illustrative case, OWASP AIVSS v0.8 formula, threat multiplier 0.97, no mitigation credit.

10

agent factors OWASP AIVSS v0.8 adds on top of a CVSS v4.0 base score

Autonomy, tools, natural language control, context, non-determinism, opacity, persistent memory, identity, multi-agent, self-modification. Each scored 0, 0.5 or 1. Strong, validated mitigations can take off at most a third.

56

pages in ISO/IEC 27090:2026, published in October 2026

Guidance on understanding, detecting and mitigating security threats to AI systems across their life cycle, from JTC 1/SC 27, the committee behind ISO/IEC 27001 and 27002. It describes threats and consequences. It leaves the ranking to you.

How ISO/IEC 27090 maps to the six inputs

Exploitability	How each AI threat is carried out and detected
Autonomy	Consequences of compromise during operation
Permissions	Identifying what needs protecting
Tools	Prompt and output injection
Memory	Operational inputs can become training data
Business impact	Connecting threats to their potential impact

Many agent flaws have no patch. Cut the agent's reach instead.

Approval for new bank accounts, a tool that drafts but cannot release, memory kept per session. The finance agent drops from 7.9 to 6.8 without touching the flaw.

Five questions for the board

- 1 Do we record each agent's autonomy, permissions, tools, memory and impact tier?
- 2 When a flaw is found, how do we know which agents it reaches?
- 3 How many open findings would change band if we scored the agent too?
- 4 Which agent findings have no patch, and what reach have we cut?
- 5 Which ISO/IEC 27090 threats have we never tested for?

Score the **agent** as well as the flaw.

Five fields in the AI system register, kept by the agent's owner, and ISO/IEC 27090 as the test plan.



Terence Kok · AI & Innovation Strategist

SAVE · COMMENT · FOLLOW