PERMISO

PERMISO AI EXPOSURE COVERAGE

35+ AI security exposures detected and mapped to the OWASP Top 10 for Agentic Applications and OWASP Top 10 Agentic Risk frameworks

COVERAGE AT A GLANCE

Permiso detects 35+ AI security exposures across cloud and SaaS environments, covering identity, data, agent and tool configurations, model governance, prompt security, network boundaries, audit controls, and safety monitoring. Each exposure is severity-ranked and mapped to both OWASP AI security frameworks.

35+	11	15	11
TOTAL EXPOSURES	CRITICAL	HIGH	MEDIUM + LOW

Why OWASP mapping matters: As enterprises deploy AI agents into production, security and compliance teams need a common criteria for evaluating risk. The OWASP Top 10 for Agentic Applications 2026 and the OWASP Top 10 Agentic Risk provide that criteria. Mapping Permiso's detection capabilities to both frameworks shows exactly which classes of AI risk are covered and at what depth.

EXPOSE CATEGORIES

CATEGORY	EXPOSURES	CRITICAL	HIGH	MEDIUM/ LOW	EXAMPLES
Identity	4	2	2		Over-privileged AI execution roles, cross account AI role access, standing privileges without JIT
Data	6	2	4		Public or unencrypted AI data, sensitive data in training, data residency violations, PII access
Agent & Tools	7	3	3	1	Unrestricted tool execution, agent tool overreach, agent self- modification, A2A control gaps
Composite	4	4			Data + autonomy + egress, prompt to-tool-to-infra paths, autonomous admin agents
Model	4		1	3	Unapproved model usage, unpinned versions, missing ownership, model drift
Prompt	3	1	1	1	Prompt injection exposure, missing prompt logging, shared prompt stores
Audit	2	1	1		Agent can disable logs, incomplete AI audit trails
Network	3		1	2	Unrestricted internet egress, public AI endpoints, cross-VPC access
Governance	3			3	No AI risk register, missing ownership and accountability, changes outside change management
Safety	1			1	No AI output safety monitoring

OWASP TOP 10 FOR AGENTIC APPLICATIONS 2026

The OWASP Top 10 for Agentic Applications 2026 defines the most critical security risks specific to AI agent deployments. Permiso detects exposures across 9 of the 10 categories.

OWASP CATEGORY	EXPOSURES	SEVERITY BREAKDOWN	WHAT PERMISO DETECTS
ASI01 Agent Goal Hijack	4	2C 1H 1M	Prompt injection exposure, prompt-to-tool-to-infra attack paths without validation, missing prompt logging, absent output safety monitoring
ASI02 Tool Misuse & Exploitation	5	4C 1H	Unrestricted tool execution via wildcard IAM, agent tool overreach, public or unencrypted AI data stores, data combined with autonomy and egress, unrestricted internet egress
ASI03 Identity & Privilege Abuse	5	2C 3H	Over-privileged AI execution roles, cross-account AI role access, standing privileges without JIT, excessive data access by AI, data residency violations
ASI04 Agentic Supply Chain Vulnerabilities	6	1C 2H 3M	Sensitive data in training, training data poisoning risk, unapproved model usage, unpinned model versions, model drift without review, AI changes outside change management
ASI06 Memory & Context Poisoning	1	1M	Shared prompt stores accessible across agent boundaries
ASI07 Insecure Inter-Agent Communication	1	1H	Agent-to-agent control gaps with no policy enforcement on cross-agent interactions
ASI08 Cascading Failures	5	3H 1M 1L	Incomplete AI audit trails, missing output and decision logging, no human-in-the-loop controls, cross-VPC and cross-account network access, no AI risk register
ASI09 Human-Agent Trust Exploitation	2	1H 1M	Human access to AI runtime roles without segmentation, public AI endpoints
ASI10 Rogue Agents	5	2C 1H 2M	Autonomous admin agents with broad permissions, agents that can disable logs,agent self-modification, missing AI and model ownership

OWASP TOP 10 AGENTIC RISK

The OWASP Top 10 Agentic Risk framework defines the broader risk categories for agentic AI systems. Permiso detects exposures across 9 of the 10 categories.

OWASP CATEGORY	EXPOSURES	SEVERITY BREAKDOWN	WHAT PERMISO DETECTS
A1 Prompt Injection	3	3C	Prompt injection exposure with unfiltered user input, prompt-to-tool-to-infra paths without validation, autonomous admin agents with no guardrails
A2 Broken AuthZ	2	1H 1M	Human access to AI runtime roles without proper segmentation, public AI endpoints
A3 Training Data Poisoning	2	1C 1H	Sensitive data (PII, secrets) used in training without classification, training data poisoning risk
A4 Insecure Agent Lifecycle	7	1H 5M 1L	Agent self-modification, missing AI and model ownership, AI changes outside change management, shared prompt stores, model drift without review, no AI risk register
A5 Suplly Chain Abuse	4	1C 1H 2M	Cross-account AI role access, unapproved model usage, cross-VPC and cross-account network access, unpinned model versions
A6 Excessive Agency	8	5C 3H	Unrestricted tool execution, agent tool overreach, over-privileged execution roles, data + autonomy + egress, prompt-to-infra paths, A2A control gaps, no human-in-the loop, standing privileges
A7 Sensitive Data Exposure	4	1C 3H	Public or unencrypted AI data stores, data residency violations, unrestricted internet egress, excessive data access by AI
A8 Unsafe Outputs	1	1M	No AI output safety monitoring
A10 Insufficient Logging	4	1C 3H	Agents that can disable logs, incomplete AI audit trails, missing output and decision logging, missing prompt logging

HOW THIS CONNECTS TO ENTERPISE AI AGENT SECURITY

These exposures map directly to the five AI agent security use cases that enterprises are evaluating today. Each OWASP category connects to one or more of the core security controls organizations need in place.

ENTERPRISE USE CASE	OWASP AGENTIC APPLICATIONS	OWASP AGENTIC RISK	WHAT PERMISO DETECTS
Agent Discovery Inventory, and Lifecycle	ASI04, ASI08, ASI10	A4, A5	Unmanaged agents, missing ownership and accountability, no AI risk register, AI changes outside change management, unpinned model versions
Least-Privilege and Access Control	ASI02, ASI03	A2, A6	Over-privileged AI roles, cross-account access, standing privileges without JIT, unrestricted tool execution, agent tool overreach, excessive data access
Behavioral Monitoring and Anomaly Detection	ASI01, ASI010	A1, A6, A8	Prompt injection exposure, prompt-to-infra paths, autonomous admin agents, agent self-modification, no human-in-the-loop,no output safety monitoring
Agent-to-Agent Communication Security	ASI07, ASI09	A2, A5	Agent-to-agent control gaps, human access to AI runtime roles, cross-VPC and cross-account network access, public AI endpoints
Compliance and auditability	ASI08	A3, A7, A10	Agents that can disable logs, incomplete audit trails, missing prompt and decision logging, data residency violations, sensitive data in training, public or unencrypted data

See How Permiso Covers Your AI Agent Attack?

35+ exposures across 10 categories, mapped to both OWASP AI security frameworks, severity-ranked, and out-of-the-box ready.

To see these exposures in your environment, visit **permiso.io** or request a technical walkthrough.