

# OWASP GenAI Compliance for Humans and AI Agents

The OWASP GenAI Data Security framework defines 16 risk categories. This guide connects those risks to Hoop.dev controls that secure your data and improve engineering productivity.

- 11 of 16 OWASP risks addressed through a single gateway
- Covers humans, AI agents, and non-human identities
- Change the risk profile of data-in-transit with data masking and in-session guardrails

---

**16**

OWASP risk categories

**4**

Direct control matches

**4**

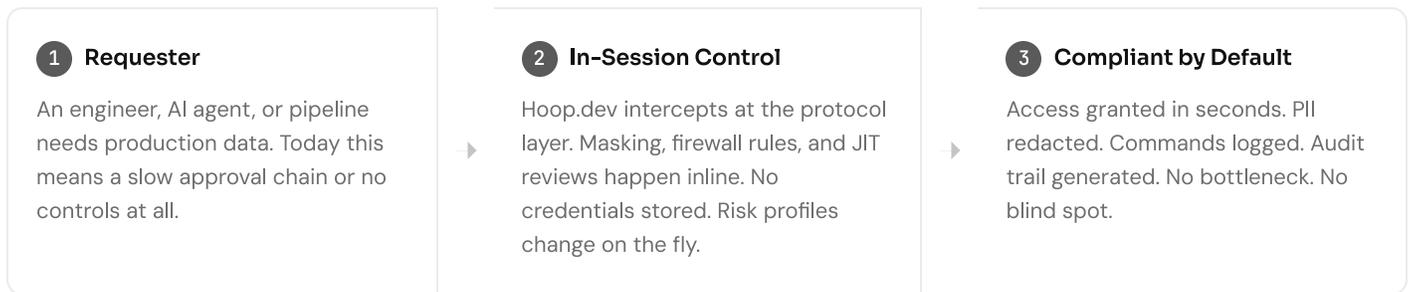
Contributing controls

**3**

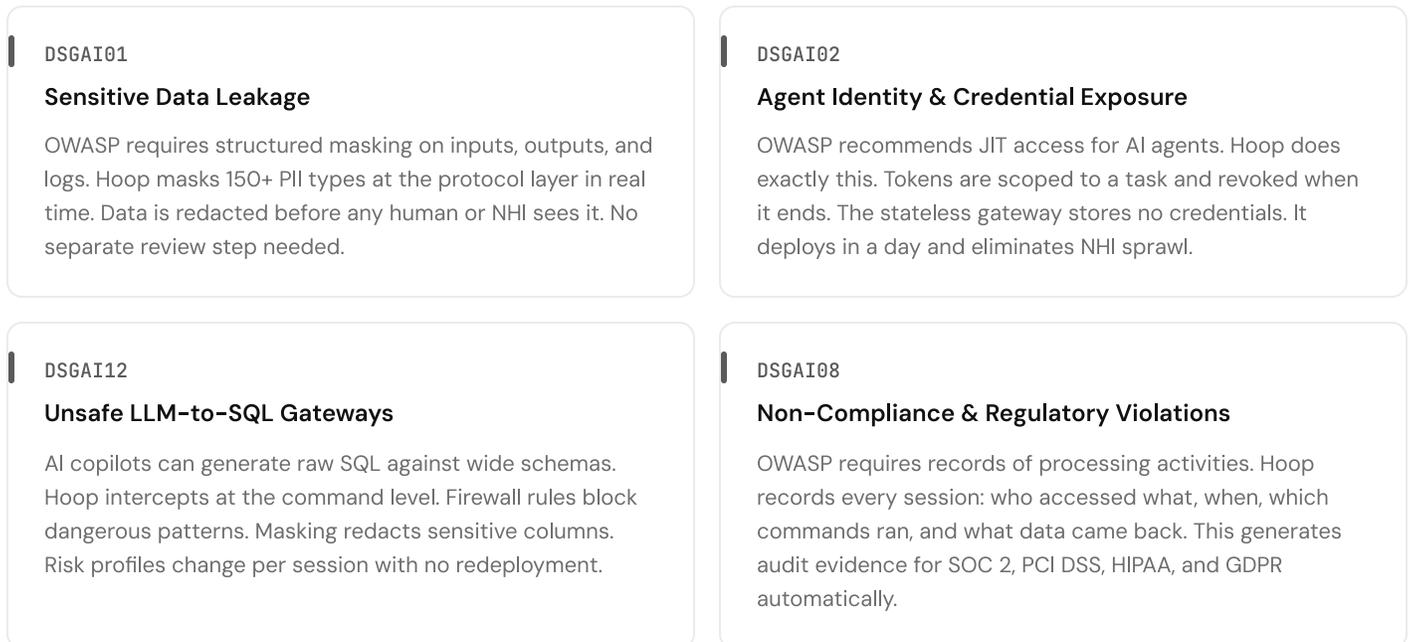
Adjacent references

# Restricting AI context doesn't solve the problem, it just delays the solution

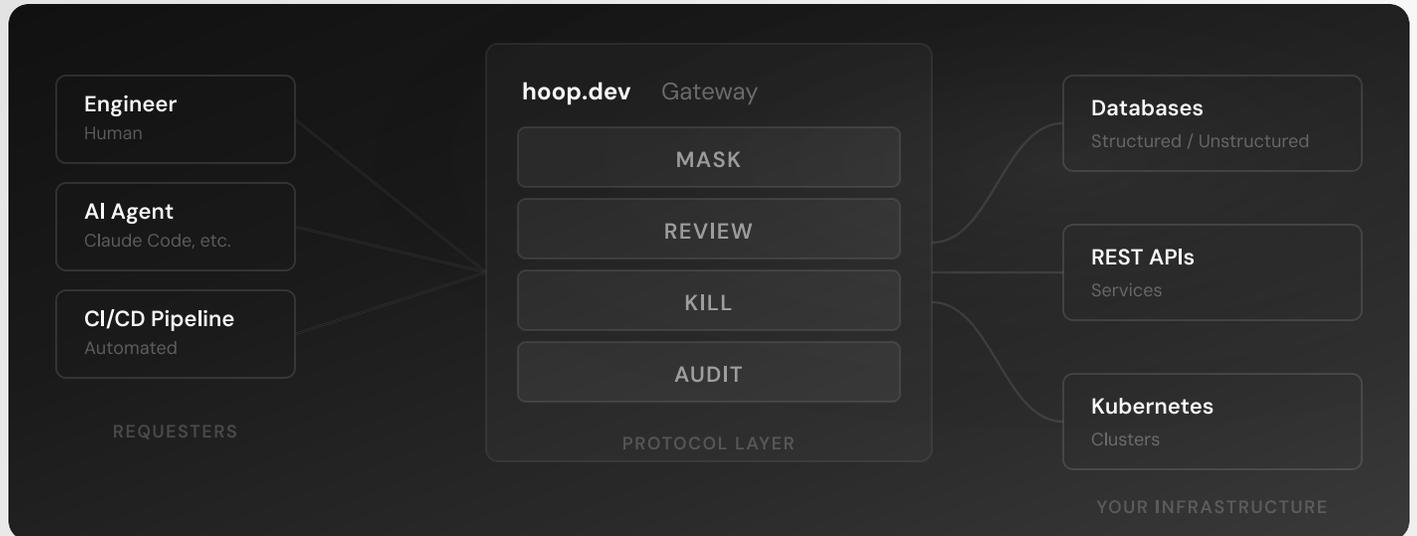
Data leakage, credential sprawl, and uncontrolled SQL generation all existed before generative AI. AI agents just make these failures happen faster. Most teams respond by locking down access. That slows engineers and frustrates security teams without fixing the root cause. The OWASP GenAI framework calls for stronger controls at the data layer. Hoop.dev satisfies those requirements with in-session governance at the infrastructure boundary.



## DIRECT CONTROL: WHERE HOOP ELIMINATES THE BOTTLENECK



# One gateway for humans and non-human identities, down to the command level



## CONTRIBUTING CONTROL: PART OF THE MITIGATION STACK

**DSGAI03**  
**Shadow AI & Unsanctioned Data Flows**  
Hoop doesn't control which AI tools people use. It controls what data those tools can reach. If access goes through Hoop, the data is already redacted before it hits any unsanctioned tool.

**DSGAI06**  
**Tool, Plugin & Agent Data Exchange**  
AI agents call tools that touch databases and clusters. Hoop sits at that boundary. It limits what data comes back through masking and logs every interaction at the command level.

**DSGAI07**  
**Data Governance & Lifecycle**  
Group-based RBAC, JIT reviews, and time-bounded sessions ensure the right identity gets the right data for the right duration. Risk profiles adjust on the fly with no redeployment.

**DSGAI15**  
**Over-Broad Context & Prompt Over-Sharing**  
AI systems pull production data to improve answers. Hoop's masking redacts PII before it leaves the database layer. Data minimization happens upstream, inline, before the model ever sees it.

# OWASP GenAI risk-to-control mapping

RISK ID	RISK NAME	HOOP CONTROL	TIER
DSGAI01	Sensitive Data Leakage	AI data masking (150+ PII types)	<b>DIRECT</b>
DSGAI02	Agent Identity & Credential Exposure	JIT access, stateless gateway	<b>DIRECT</b>
DSGAI12	Unsafe LLM-to-SQL Gateways	Command interception, firewall policies	<b>DIRECT</b>
DSGAI08	Non-Compliance & Regulatory Violations	Immutable session recording	<b>DIRECT</b>
DSGAI03	Shadow AI & Unsanctioned Data Flows	Data minimization at boundary	CONTRIBUTING
DSGAI06	Tool & Agent Data Exchange Risks	Context minimization, observability	CONTRIBUTING
DSGAI07	Data Governance & Lifecycle	RBAC, JIT reviews, time-bounded sessions	CONTRIBUTING
DSGAI15	Over-Broad Context Windows	Upstream masking at infra layer	CONTRIBUTING
DSGAI11	Cross-Context Conversation Bleed	Per-user session isolation	ADJACENT
DSGAI14	Excessive Telemetry Leakage	Masking applies to audit logs	ADJACENT
DSGAI16	Endpoint Assistant Overreach	Access gate for all requesters	ADJACENT

**The key insight:** Compliant access should be the fastest path, not a bottleneck. Every OWASP GenAI risk traces back to what humans and NHIs can reach. In-session controls at the protocol layer let you manage what every identity can see and do. You can change the risk profile on the fly.

## CONTROL GLOSSARY

How Hoop.dev gateway controls map to OWASP GenAI framework language

**MASK**



### AI Data Masking

Redacts 150+ PII types in real time at the protocol layer. Data is filtered before any human or NHI sees it. Maps to DSGAI01 and DSGAI15.

**REVIEW**



### JIT Access Reviews

Time-bounded sessions with approval gates. No standing credentials. Tokens scoped to a task and revoked when done. Maps to DSGAI02 and DSGAI07.

**KILL**



### Firewall Policies & Command Interception

Blocks dangerous query patterns. Controls outbound traffic. Risk profiles change per session with no redeployment. Maps to DSGAI12 and DSGAI06.

**AUDIT**



### Immutable Session Recording

Full request/response traceability for every human and NHI session. Generates audit evidence for SOC 2, PCI DSS, HIPAA, and GDPR. Maps to DSGAI08.