

Your AI is live. Do you know what's now exposed?

The AI Exposure Checklist: 10 Questions Every Security Leader Should Be Able to Answer Confidently

AI moves fast, and security visibility doesn't always keep up. Most organizations can't fully see what their AI systems expose, access, or connect to once they're live. Here's what needs to be on your radar:

High blind spot

1 Do you know every AI endpoint that's currently exposed to the internet?

Not what you launched six months ago, what's live right now. AI features get added quietly by engineering teams, third-party integrations, and SaaS tools. By the time security catches up, the surface has already grown.

High blind spot

2 Have you mapped what data your LLMs and agents can access, via RAG, APIs, and tools, and what they will return or execute if prompted?

LLMs can be manipulated into surfacing data they weren't designed to share. If your model has access to customer records, internal documents, or regulated data, that access needs to be scoped and tested, not assumed to be safe.

Often overlooked

3 Do you know how your AI workflows interact with each other after they're live, including agent-to-agent and tool-chaining behavior at runtime?

AI components rarely operate alone. They call APIs, pass data between services, and trigger downstream actions. The risk often lives in the connections, not the models themselves, and those interactions are rarely documented end to end.

High blind spot

4 Are your AI-facing APIs enforcing authentication and rate limits consistently?

One misconfigured API that feeds an LLM can undo everything else. Inconsistent auth, missing rate limits, and unenforced scopes are common, and they expand your attack surface without any visible warning sign.

Often overlooked

5 Could a user manipulate your AI into taking actions it shouldn't?

Prompt injection isn't theoretical. If your AI can take actions — send messages, query databases, trigger workflows — a crafted input can redirect those actions. This is now a primary attack vector against AI-enabled products.

Often overlooked

6 Do you have clear data ownership for everything your AI systems touch?

Customer data, partner data, regulated data... once it moves into an AI pipeline, ownership blurs fast. If you can't trace where data goes and who's accountable for it, you can't make defensible compliance claims.

High blind spot

7 Does your security tooling actually see your AI layer, or only what sits below it?

Most traditional security tools were built for static infrastructure and known attack patterns. They test components in isolation. AI applications behave dynamically and are shaped by real-world usage, meaning exposure can grow in places your current tooling won't flag.

Often overlooked

8 Are you testing AI behavior under adversarial conditions, not just expected scenarios?

Standard QA and penetration testing don't cover how an AI system behaves when a user is actively trying to break it. Adversarial inputs, edge case prompts, and chained exploits require a different kind of evaluation.

Often overlooked

9 When your AI stack changes, does your security posture update automatically?

New model versions, updated integrations, and changing configurations all affect your risk profile. Point-in-time assessments go stale quickly. If visibility isn't continuous, you're always reacting to yesterday's picture.

High blind spot

10 If a regulator asked you to prove your AI systems handle sensitive data appropriately, could you?

In finance, healthcare, education, and SaaS, that question is no longer theoretical. It's already being asked. "We believe our systems are secure" isn't an answer. You need documented visibility, not confidence based on partial information.

Most teams who work through this list find at least three questions they can't answer with confidence. Not because they've been negligent, but because most security programs weren't designed with AI in mind.

Chronoloq helps teams close that gap by continuously surfacing how AI systems actually behave: what data they can reach, how boundaries shift, and where risk accumulates quietly over time.

Book a meeting to find out what's exposed in your AI environment

