

SAMPLE

AI Architecture Review Checklist

The opening of AI-Architecture-Review-Checklist.docx, one of 4 files in the kit.

<https://solasai.net/products/ai-architecture-checklist/>

AI Architecture Security Review Checklist

47 pass/fail checks across eight categories, each with what to look for, how to test it, the evidence to keep, why it carries its severity, and the OWASP Top 10 entries it addresses

Product: AI Architecture Security Review Checklist, version 3.0, 19 October 2026. **Classification:** Purchaser use only, customize before use. Solas AI, solasai.net.

How to run the review

Run the checklist against any system that sends data to or takes action from a language model: a chatbot, a copilot feature in an existing product, a retrieval pipeline, an agent with tools. Run it twice: at design review, when most checks are answered from the architecture document and the vendor contract, and before launch, when the implementation checks need a config export, a code path, or a test run. One reviewer with read access to the repository, the cloud console, the model provider dashboard, and the vendor terms can complete it in a working session for a system of ordinary size.

Each check is pass, fail, or not applicable. Not applicable is reserved for checks whose subject does not exist in the system under review (no file uploads, no tool calling, one tenant); it is not a way to defer a hard check. Record the evidence location for every pass, because the next reviewer, the enterprise buyer's security team, and the penetration tester will ask for it. The Review Tracker workbook shipped with this kit holds the 47 checks with result drop-downs, counts the failures by severity, and applies the release gate.

Severity model

Critical: a failure that, on its own, leaks another customer's data, lets an attacker control the model's instructions or actions, or exposes a credential. The system does not ship with a Critical fail. High: a failure that an attacker can turn into one of the above with a second step, or that a security questionnaire or audit will treat as a finding. Medium: a weakness that degrades safety, quality, or operability and should be fixed within a release or two. Low: good practice.

Release gate

Every Critical check passes. No more than two High checks are deferred, each with a named owner, a date, and a compensating control written down. Medium fails go to the backlog with an owner. The reviewer records one of three outcomes: approved for production, approved with conditions (the deferred High items and their dates), or blocked until the named checks pass.

How the mapping works

Each check names the entries of the OWASP Top 10 for LLM Applications (2025 edition, LLM01 to LLM10) and, where the system has agentic features, the OWASP Top 10 for Agentic Applications (version 2026, published December 2025, ASI01 to ASI10) that it addresses. The mapping table at the end runs the other way, from each OWASP entry to the checks that cover it, so a reviewer can answer "what do we do about prompt injection" in one place. The OWASP lists are the vocabulary most security questionnaires and penetration testers now use; speaking it saves a meeting.

1. Data flow and privacy (7 checks)

1.1 A data flow diagram shows every path customer data takes through the AI pipeline

The sample ends here. The full document runs about 6,955 words; the kit holds 4 files. Version 2.2, 19 October 2026. Buy it at <https://solasai.net/products/ai-architecture-checklist/>, 14-day refund for any reason.