The GRASP Framework

Five pillars to evaluate the enterprise-readiness of an AI project

AppSecAI

Pillar Definitions

Pillar	The Question	Without It
G Governance	Who controls it, what are the rules, and what does it cost?	*People adopt AI tools on their own. No one tracks what's running, who approved it, what the rules are, or what it's costing.*
R Reliability	Does the system stay up, and are results consistent?	*It works until it doesn't. A service goes down, a model gets updated, and output quality drops without anyone noticing.*
A Assurance	Are results accurate, relevant, and solving the right problem?	*The system produces thousands of outputs, but no one has checked whether any of them are actually correct or useful.*
S Scalability	Can it handle the full workload, not just a pilot?	*One team runs a successful pilot. Meanwhile, the other 500 problems across the org are still handled by hand.*
P Protection	Is your data safe — and do you trust the providers?	*Your data flows to third-party models with no agreements in place. It could end up training someone else's product.*

appsecai.io/owasp-april-2026

Slides · GRASP Templates · Book · Project OASIS

↻ FLIP FOR EVAL MATRIX

OWASP Bay Area · April 16, 2026

AppSecAI
appsecai.io

GRASP Evaluation Matrix

Rate any AI system 1–5 on each pillar. The goal isn't a perfect 25 — it's matching the score to the risk and purpose of the project.

AppSecAI

	1 None	2 Basic	3 Working	4 Strong	5 Leading
G Governance	**Blind** No visibility. Every session starts from scratch.	**Aware** Someone knows AI is being used; ad hoc oversight.	**Visible** Dashboards, approval gates, cost visible per team.	**Policy-Driven** Automated rules; budget controls and cost alerts.	**Autonomous** Full policy engine, configurable autonomy, cost governance.
R Reliability	**Fragile** No error handling; results vary wildly.	**Reactive** Basic error handling; manual recovery.	**Robust** Fallbacks, retries, consistency checks.	**Consistent** Self-heals; drift tracked; regressions caught.	**Resilient** Full redundancy; drift detection + automated rollback.
A Assurance	**Unproven** No validation of outputs.	**Spot-Checked** Manual review of some outputs.	**Fit for Purpose** Systematic validation; solving the right problem.	**Outcome-Driven** Evidence tied to business outcomes; drift corrected.	**Accountable** Continuous measurement; efficacy vs objectives; SLAs.
S Scalability	**Manual** One use case, by hand.	**Team** Works for one team.	**Multi-Use** Multiple teams; some automation.	**Enterprise** Org-wide; cost-effective at scale.	**Industrial** Full workload; cost-per-outcome optimized.
P Protection	**Exposed** Unvetted providers; data may train models.	**Basic** Trusted models; opt-outs; basic agreements.	**Policy-Based** Providers vetted; data policy enforced.	**Defended** Data scanned before leaving; prevention rules.	**Self-Hosted** All AI internal; no data leaves perimeter.

Total Score: 5–10 Critical · 11–14 At Risk · 15–17 Developing · 18–21 Strong · 22–25 Leading

↻ FLIP FOR DEFINITIONS

OWASP Bay Area · April 16, 2026

AppSecAI
appsecai.io