JIEGOU

JIEGOU OPERATOR BRIEF

10-Layer AI Governance Self-Assessment

Operator-Grade Diagnostic for AI Governance Maturity

VERSION	v1 · 2026-05-20
FROM	JieGou — AI Operations Partner for engineering-led mid-market IT teams ($50M–$1B revenue)
FOR	CIOs · CTOs · CISOs · Heads of IT · Heads of AI / Engineering Leadership
STATUS	Public artifact — free to use, cite, and circulate

jiegou.ai

NO SALES-CALL CONDITION

A structured 30-question self-assessment of your organization's AI governance posture across 10 layers — Identity & Access, Audit Trail, Data Governance, Human Oversight, Model Governance, Tool Governance, Compliance, Cost Controls, Observability, and Incident Response.

You'll get a per-layer score (0-10), an overall score (0-100), a grade (A-F), per-layer recommendations for weak spots, and an industry benchmark comparison. Time to complete: **20-30 minutes**.

This is the same framework JieGou uses internally to architect our own platform and to assess every customer engagement. We share it externally for two reasons. First, the AI governance gap is real and most organizations underweight it until an audit, incident, or board question forces the conversation. Second, you should be able to evaluate JieGou the same way we evaluate ourselves — the framework is the same on both sides of the table.

Why this exists — the gap it measures

A Fortune 500 governance survey published in late 2025 found:

> ***80% of Fortune 500 companies had deployed AI agents in production. Only 14% had received security approval for those deployments.***

That gap is not a vendor problem. It's a governance-maturity problem. The technology shipped faster than the organizational practices, and the practices most companies need don't yet have a canonical form. This assessment is one attempt at canonical form — drawn from regulatory frameworks (SOC 2, EU AI Act, NIST AI RMF), from production deployments JieGou operates today, and from gaps we've watched customers hit when their AI estate scaled past one or two workflows.

Three uses the result supports:

1. **Internal baseline** — know where you actually are, layer by layer
2. **Board / audit defense** — concrete evidence of where you are mature, where you have gaps, and what you're doing about them
3. **Peer benchmark** — see where your maturity sits relative to organizations of comparable size

The assessment does **not** require any JieGou product to use. The framework is the framework whether you build, buy, or operate AI workflows.

The 10 layers — what each one is and why it matters

#	Layer	What it covers	Why it matters
1	**Identity & Access**	RBAC for AI tools, per-agent identity (not just per-user), SSO/SAML for AI consoles	Without agent-level identity, you can't tell who did what — humans and agents collapse into one audit subject
2	**Audit Trail**	Action logging with full attribution, decision traceability back to source data and prompts, evidence export for compliance reviews	"Why did the AI do that?" must have an answer; "show me the evidence" must be exportable
3	**Data Governance**	PII / PHI detection, data residency controls, encryption at rest for keys and credentials	Regulators ask data-residency questions before they ask AI-quality questions
4	**Human Oversight**	Approval gates for consequential AI actions, graduated autonomy (not binary on/off), low-confidence escalation policies	Shadow Mode discipline scales; "trust the AI" doesn't
5	**Model Governance**	Multi-provider evaluation, certified / approved model registry, provider-portability (workflows don't rewrite when you switch LLMs)	Provider lock-in becomes provider risk when prices, policies, or capabilities shift
6	**Tool Governance**	Access controls on which tools AI agents can invoke, human-approval gates on high-impact tool calls, certified / tested integrations	AI without tool boundaries is AI with arbitrary code execution
7	**Compliance**	Regulatory framework mapping (SOC 2 / EU AI Act / HIPAA / GDPR), tracking dashboards, documented governance policies	The EU AI Act enforcement period starts in phases through 2026-27; mapping is no longer optional
8	**Cost Controls**	Per-agent / per-workflow token budgets, rate limiting on AI requests, departmental cost allocation	The first unbounded-LLM-loop bill is the one that makes Finance care about governance
9	**Observability**	Real-time AI activity monitoring, workflow-health dashboards, output-quality tracking over time	"Is the AI still working as well as last quarter?" needs a metric, not a hunch

#	Layer	What it covers	Why it matters
10	**Incident Response**	Dead-letter queue / retry mechanism for failed operations, incident tracking for AI failures, vendor risk register	When (not if) an AI vendor has an outage, your incident response posture decides the blast radius

How scoring works

Per-question scoring (4 levels):

Answer	Value	What it means
Not implemented	0	Capability doesn't exist in your environment
Basic / ad hoc	1	Exists informally, inconsistently, or for some workflows but not others
Moderate / documented	2	Implemented + documented; covers most production workflows
Comprehensive / enforced	3	Implemented + enforced + audited; covers all production AI workflows

Per-layer scoring (3 questions per layer, weighted):

Each question has a weight (2 or 3) reflecting its importance within the layer. Layer score = (weighted sum of normalized answers) × 10 → produces a 0-10 score.

Layer score	Status
8.0 – 10.0	**Strong** — mature; maintain + extend
5.0 – 7.9	**Adequate** — operational gaps in specific areas
2.0 – 4.9	**Weak** — significant gaps; prioritize remediation
0.0 – 1.9	**Missing** — layer is essentially absent

Overall scoring:

Overall score (0-100)	Grade	Interpretation
90 – 100	**A**	Industry-leading governance maturity
75 – 89	**B**	Solid foundation, gaps in specific layers
55 – 74	**C**	Material governance gaps; remediation priority
35 – 54	**D**	Governance significantly underweighted relative to AI deployment
0 – 34	**F**	Critical gap — board / audit / regulator exposure

Industry context (median scores by company size, from our internal benchmarks):

Company size	Median overall score
Enterprise (10,000+ employees)	~42
Mid-market (500-10,000)	~35
SMB (50-500)	~25
Startup (<50)	~20

A "B" grade (75+) puts you in the top decile across all sizes. This isn't because most organizations are bad — it's because AI deployment outran AI governance for everyone.

The 30 questions

For each question, pick the answer that best matches your current state. Be honest. The point of self-assessment is operational truth, not an aspirational story. If the answer is "ad hoc," score it as ad hoc.

Layer 1 — Identity & Access

#	Question	Weight
1.1	Do you have role-based access control for AI agent management?	3
1.2	Do your AI agents have their own identity (separate from the user who created them)?	3
1.3	Is SSO/SAML configured for your AI tooling?	2

Layer 2 — Audit Trail

#	Question	Weight
2.1	Are AI agent actions logged with full attribution?	3
2.2	Can you trace AI decisions back to source data and prompts?	3
2.3	Can you export audit evidence for compliance reviews?	2

Layer 3 — Data Governance

#	Question	Weight
3.1	Do you have PII/PHI detection in your AI workflows?	3
3.2	Are data residency controls configured for AI processing?	3
3.3	Are API keys and credentials encrypted at rest?	2

Layer 4 — Human Oversight

#	Question	Weight
4.1	Do you have approval gates before AI agents take consequential actions?	3
4.2	Do you use graduated autonomy levels (not just on/off)?	3
4.3	Do you have escalation policies for when AI confidence is low?	2

Layer 5 — Model Governance

#	Question	Weight
5.1	Do you evaluate multiple LLM providers for each use case?	3
5.2	Do you have a certified/approved model registry?	2
5.3	Can you switch LLM providers without rewriting workflows?	3

Layer 6 — Tool Governance

#	Question	Weight
6.1	Do you have access controls on which tools AI agents can use?	3
6.2	Do specific tools require human approval before agent invocation?	3
6.3	Are your MCP / tool integrations certified or quality-tested?	2

Layer 7 — Compliance

#	Question	Weight
7.1	Do you have a regulatory framework mapping for your AI systems (SOC 2, EU AI Act, etc.)?	3
7.2	Do you have a compliance dashboard tracking AI governance controls?	2
7.3	Do you have documented AI governance policies?	3

Layer 8 — Cost Controls

#	Question	Weight
8.1	Do you have per-agent or per-workflow token budgets?	3
8.2	Do you have rate limiting on AI agent requests?	2
8.3	Do you track and allocate AI costs by department or team?	3

Layer 9 — Observability

#	Question	Weight
9.1	Do you monitor AI agent activity and performance in real time?	3
9.2	Do you have metrics and dashboards for AI workflow health?	2
9.3	Do you track AI output quality over time?	3

Layer 10 — Incident Response

#	Question	Weight
10.1	Do you have a dead letter queue or retry mechanism for failed AI operations?	2
10.2	Do you have incident tracking for AI failures or misbehavior?	3
10.3	Do you maintain a vendor risk register for your AI providers?	3

Per-layer remediation playbook

For any layer scoring **Weak** or **Missing**, the first three remediation moves are below. These are ordered by leverage: do #1 before #2 before #3 in each layer.

Identity & Access

1. Implement RBAC with at least 3 distinct roles (admin / operator / viewer minimum)
2. Give each AI agent its own scoped identity and per-tool permissions
3. Configure SSO/SAML for centralized authentication across your AI tooling

Audit Trail

1. Log every AI agent action with full attribution (who initiated, what tool was called, what data was read, when, with what result)
2. Ensure AI decisions can be traced back to specific input data + prompt versions
3. Set up automated evidence export for compliance reviews (JSON, S3, or forward-to-SIEM)

Data Governance

1. Deploy PII/PHI detection in all AI workflows processing personal data
2. Configure data residency controls for regulated data categories (HIPAA / GDPR / industry-specific)
3. Encrypt all API keys and credentials at rest with envelope encryption (KMS / Secrets Manager)

Human Oversight

1. Add approval gates before AI agents take consequential actions (outbound communication, ERP writes, financial transactions)

2. Implement graduated autonomy levels instead of binary on/off controls (Shadow Mode → Supervised → Trusted, with per-workflow gating)
3. Create escalation policies for low-confidence AI outputs (confidence threshold → human review)

Model Governance

1. Evaluate multiple LLM providers with structured bakeoffs for each new use case
2. Maintain a registry of approved models with version-pinning for production use
3. Design workflows to be model-agnostic — workflows should not require rewrites when you switch primary providers

Tool Governance

1. Implement tool-level access controls in your MCP / integration layer (allowlist, not denylist)
2. Require human approval for high-impact tool invocations (financial systems, external communications, data modification)
3. Certify and quality-test tool integrations before production use; treat each integration as a release-gated component

Compliance

1. Map your AI systems to relevant regulatory frameworks (SOC 2 controls, EU AI Act risk tiers, sector-specific frameworks)
2. Build a compliance dashboard tracking governance controls against framework requirements
3. Document AI governance policies and make them auditable — written policies, version-controlled, with review cadence

Cost Controls

1. Set per-agent or per-workflow token budgets to prevent runaway costs from infinite loops or prompt-injection abuse
2. Implement rate limiting on AI agent requests (per-user, per-account, per-workflow)
3. Track and allocate AI costs by department or team for accountability and budget planning

Observability

1. Monitor AI agent activity and performance in real time (latency, error rate, throughput)
2. Build dashboards for AI workflow health and success rates (success rate, exception rate, time-to-completion)

3. Track AI output quality over time to detect model drift, prompt regressions, and data-distribution shifts

Incident Response

1. Implement a dead-letter queue for failed AI operations with retry policies
2. Set up incident tracking for AI failures and misbehavior with named accountable owners
3. Maintain a vendor risk register for all AI providers (uptime SLAs, data handling, sub-processor list, exit strategy)

What to do with your result

Three paths — all free as part of the JieGou Operations Partner relationship; the framework belongs to you whether or not you ever buy anything from us.

Path A — Run it yourself, use the report internally. Pick your answers, do the math (the formulas above), share with your team and board. The framework is yours to use. No need to send us anything.

Path B — Walk through it with us (45-min call). We'll facilitate the assessment, ask follow-up questions where "ad hoc" or "moderate" needs disambiguation, and produce a written report with your scores + tailored recommendations for your top 3 weak layers. Calendar link or async scheduling — your preference.

Path C — Send us your raw answers; we send you a written analysis (within 5 business days). Email your 30 answers (questionId → 0/1/2/3) to `assessments@jiegou.ai`. We produce a written report with scoring, status interpretation, prioritized recommendations, and (if relevant) peer-benchmark context for your size segment.

There's no Path D where we sell you something to score higher. Whichever path you pick, the report is the deliverable — not a sales call disguised as an assessment.

What we DON'T do with your answers

To make the path-B and path-C options unambiguous:

- **We do NOT store your answers without explicit permission.** If you pick Path C, we ask permission before retaining anything beyond the analysis window.
- **We do NOT use your answers for sales targeting.** The point of the assessment is operational truth, not lead qualification.

- **We do NOT share answers with other customers, partners, or aggregators.** Per-organization data stays per-organization. Aggregated, anonymized benchmarks (e.g., "median enterprise score is 42") are produced only with explicit opt-in.
- **We do NOT condition the report on a sales call.** Path C produces a written report regardless of whether you ever talk to a JieGou rep.
- **We do NOT upsell from weak-spot findings.** If your weak layer maps to something JieGou doesn't operate (e.g., HR-side identity controls), we'll say so explicitly and recommend the right kind of partner.
- **You can request deletion** of your answers + report at any time. Email `privacy@jiegou.ai`. Confirmed within 5 business days.

This list models our internal "we don't do that" exclusion-list discipline — explicit boundaries that protect both sides of the relationship.

Sourcing — where the 10 layers come from

We didn't invent these layers from scratch. They're the synthesis of:

- **Regulatory frameworks** — SOC 2 Common Criteria (especially CC6 Logical & Physical Access, CC7 System Operations), EU AI Act risk tiers and conformity assessments, NIST AI Risk Management Framework, ISO/IEC 42001 AI Management System
- **Production-deployment experience** — JieGou's own platform runs on these 10 layers in production; the layer definitions reflect what we've had to build to operate AI for paying customers
- **Customer pattern recognition** — what customer audits and security reviews actually ask for; what board questions actually surface; what enforcement actions on competitors have established as precedent

The framework will evolve. Likely additions in v2: explicit prompt-governance and prompt-injection-response (currently distributed across layers 4 and 6), explicit model-supply-chain layer (currently in layer 5), explicit AI-DLP layer (currently in layer 3). When v2 ships, your v1 scores remain usable with documented translation.

A word on the broader posture

This assessment frames AI governance as **architecture**, not policy. Policy is the easier half — write the doc, get it reviewed, file it. Architecture is the harder half — build the systems that make the policy automatically enforceable rather than relying on human discipline at scale.

The organizations that come through audits well are the ones that built audit evidence into the operating layer, not the ones with the longest policy documents. This assessment scores architecture more heavily than policy on purpose. A layer can have a 100-page policy doc and still score "missing" if the architecture doesn't enforce it; conversely, a layer can have a 2-sentence policy and score "strong" if the architecture makes the policy automatic.

The 10 layers are independent enough to remediate individually but interlocking enough that strong-layer + missing-layer organizations look different in audit than weak-across-the-board organizations. Treat the layer pattern as a map, not just an average.

Closing

This assessment is the framework JieGou uses to think about its own platform. It's also the framework we'd use to evaluate yours. We give it away free because if your AI governance posture is mature enough to not need us, we'd rather you know that and skip the sales process; and if it's not, then the conversation we'd want to have is informed by an honest baseline, not a marketing-style "AI maturity benchmark" designed to manufacture deficits.

If a 45-min walk-through is useful — or if you'd like to send your answers for a written analysis — we're at `assessments@jiegou.ai`.

If you'd just like to run it internally and never tell us your result, that's the right answer too. The framework belongs to you either way.

—

JieGou — AI Operations Partner for Engineering-Led Mid-Market IT

10-Layer AI Governance Self-Assessment · v1 · 2026-05-20 · Tier 1 deliverable · framework drawn from SOC 2 + EU AI Act + NIST AI RMF + production-deployment experience · ~20-30 min to complete · three engagement paths offered, no sales-call condition