One attestation, and the discipline behind it

JieGou — the diligence trust page · attestation + published honesty, together

Security diligence usually gets you two kinds of evidence: a certificate, or a promise. Here is the certificate — and the operating discipline that makes it more than paper.

The attestation. JieGou's **SOC 2 Type II** examination is complete: an independent audit of the JieGou Platform over **March 23 – June 23, 2026**, against the Security trust services criteria, returning an **unqualified opinion with no exceptions noted**. Type II attests that controls *operated over a period* — access control, change management, audit logging, encryption, incident response — tested as they actually ran, not as they were described. The full report is available under NDA. Continuous compliance monitoring runs between audit periods; the Type II cycle renews annually.

The discipline. A point-in-time attestation tells you controls ran for ninety audited days. What it can't tell you is what a company does when something fails *outside* the audit window — which is where most vendors' evidence ends and self-attestation begins. Ours doesn't. JieGou publishes findings **against itself**, on a running basis, because that record is expensive to fake and impossible to backfill:

- **We retitled a published essay because our own thesis failed verification** — and opened the corrected essay by saying so. *(Public: shyanming.substack.com/p/the-selection-problem.)*
- **Our research program pre-registers hypotheses and publishes its own nulls** — four refuted hypotheses in one recent series, recorded verbatim in the permanent record, kept, not buried.
- **When our system fails silently, the incident becomes a named control.** Every silent-failure path we catch in our own operations is converted into a logged, bounded, monitored one, with the incident named in the fix commit. This is not occasional: our own instrumentation has caught and converted multiple such incidents in each of the last several weeks — most recently a delivery-gate error class that reported success on failure, an ephemeral-key rotation flaw, and a build-blocking configuration fault. No customer reported any of them; our operation of the system did.
- **Our ungoverned baseline caught an error in our governed layer — and we wrote the failure mode up as a hazard class** that names our own product's risk surface.

The full dated record, with receipts, is the companion page **"How We're Wrong on Purpose"** *(jiegou.ai/downloads/wrong-on-purpose-v1.pdf)* — each item public or inspectable in a diligence session.

Why the two belong together. The attestation is a third party swearing the controls ran. The discipline is the reason they'll still be running when no auditor is looking: a company that pays the cost of publishing its own errors has no cheaper option left than fixing them. In a market where nearly every AI-governance claim is self-attested, this is the only pairing we know of that arrives with its own evidence — a signature from outside, and a paper trail from inside.

Attestation issued 4 August 2026. Honesty record current as of 5 August 2026 — dated items age out; ask for the current page. Full SOC 2 Type II report: under NDA via jiegou.ai/security.