

PALO Governance for ChatGPT

Development brief for press and media

27 October 2026 | Independent project by Fabrizio Degni

PALO Governance is a plugin candidate designed to help people prepare and review AI governance documentation in ChatGPT. Six skills organize lifecycle assessments, agent action designs, evidence reviews, maturity assessments, KPI/KRI and AI FinOps registers, and data governance dossiers. Each skill uses a packaged snapshot of PALO reference material.

The connected 0.2.0 draft adds a read-only PALO Knowledge Reader through the Model Context Protocol (MCP). The project has reached package preparation and portal validation milestones. Authenticated integration remains incomplete. The candidate has not been submitted for public review, approved or published.

What the plugin produces

- Lifecycle Assessment: a case brief, a control/evidence matrix across six phases and actions for the next review.
- Agentic Action Design: draft Action Claims and Effect Contracts, delegation limits, approval requirements and tests for proposed actions.
- Evidence Review: claim-to-evidence comparisons, contradictions, incident timelines and unresolved findings.
- Maturity Assessment: a deployment and governance profile that preserves missing evidence and required assessment items.
- KPI, KRI and AI FinOps: indicator definitions, cost attribution and reconciliation plans, including undefined values and uncertain billing.
- Data Governance: data-flow inventories, provenance and fitness requirements, disclosure constraints and evidence gaps.

These outputs support accountable human review. The plugin does not grant production authority, execute target business actions or certify compliance. PALO's separate runtime components have their own implementation and qualification boundaries.

Why the architecture matters

Each of the six skill directories carries its own copy of 26 canonical references. The development process verifies resource paths and hashes before freezing the archive. This reduces dependence on host access to sibling directories and makes the reference snapshot traceable. The connected draft contains 185 files.

The action-design workflow separates authorization, execution, observed effect and business outcome. An execution receipt alone cannot establish all four. A proposed action specifies evidence to collect later. An unresolved past attempt stays unknown or inconclusive until observations resolve it.

The Reader integration uses delegated OAuth permissions for guidance and knowledge retrieval. Microsoft Entra shows consent for only `palo:guide` and `palo:knowledge:read`. Service admission remains a separate check. The current single-tenant configuration also limits who can authenticate, regardless of the listing's intended geographic availability.

Evidence available as of 27 October

The OpenAI Platform draft shows Checks passed for all six skills and Domain verified for the Reader endpoint. It also shows Configuration incomplete, Needs attention and Not published. These states are reproduced in real process screenshots in the companion presentation.

A corrected local Codex harness tested the static 0.1.0 package across 18 cases: 16 passed within their stated scope, one failed and one was blocked by infrastructure. The failed case echoed an explicitly synthetic API placeholder. No real credential was supplied. The blocked case timed out during workspace routing before output. These results do not establish authenticated ChatGPT or Reader behavior for version 0.2.0.

GLM-5.3 High reviewed an extracted manifest, six skills and a shared contract. Gemini 4 with high reasoning reviewed three critical skills. Neither review covered the complete corpus or tested the installed runtime. Corrections addressed local reference isolation, routing and evidence-state handling. A reported invisible-character defect was rejected after a byte-level check found no such character. Neither reviewer approved the final connected archive.

Current blocker and remaining work

The Microsoft sign-in attempt returned AADSTS500113, indicating that the application had no registered reply address. The exact callback for this OpenAI connection remains required. OpenAI Help Center acknowledged escalation to a support specialist. That acknowledgement establishes a support handoff, not a technical resolution or publication approval.

The remaining release work includes the exact callback, admission of the intended ChatGPT client at the Reader service, authenticated tool discovery and calls, five positive and three negative connected tests, reviewer access, a recorded demonstration, public privacy and terms pages, and required publisher attestations. The prepared connected test cases have not run. A submission receipt must precede any claim that the candidate has entered public review.

Editorial use and contact

Suggested description: “PALO Governance is an independently developed ChatGPT plugin candidate for preparing AI governance documentation with traceable references and explicit evidence gaps.”

The accompanying slides and PNGs document development, including unresolved findings. They are suitable for a development story. They do not announce availability, imply OpenAI endorsement or demonstrate a completed authenticated integration. Screenshots are actual captures with crops that exclude private contact and infrastructure details. Code excerpts are abbreviated and labelled; they are not a deployment recipe.

Project: <https://paloframework.org/>

Author and media contact: Fabrizio Degni, <https://www.linkedin.com/in/fdegni/>

Official process references: <https://developers.openai.com/plugins/deploy/submission> and <https://developers.openai.com/plugins/build/auth>

Source records: package manifests and hashes, local test results, reviewer transcripts, Microsoft Entra permission capture, OpenAI Platform captures and support escalation acknowledgement. The companion evidence index maps these records to the slides. No direct quotation attributed to Fabrizio Degni or OpenAI is proposed in this brief.