PALO Governance for ChatGPT

Six governance workflows, a public Reader and a new bilingual API evaluation.

100

completed model responses

50

Italian and English cases

What the work now shows

Real API outputs, source-linked answers, explicit failure cases and a reproducible record of corrections. The tests expose useful limits as well as successful responses.

Version 0.3.1 is in OpenAI review.
Public-directory availability remains pending.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

DEVELOPMENT UPDATE

30 October 2026

Six workflows, reviewable outputs

Each skill answers a different governance question and identifies what still needs evidence.

Workflow	Question and draft output
Lifecycle Assessment	What needs review? A scoped case and control/evidence matrix.
Agentic Action Design	What may an agent do? Action Claim and Effect Contract drafts.
Evidence Review	What actually happened? Claims, observations, contradictions and gaps.
Maturity Assessment	What capability is supported? Two lenses, 22 questions, unresolved items retained.
KPI, KRI and AI FinOps	What should we measure? Indicators, cost coverage and reconciliation.
Data Governance	Which data may be used or disclosed? Provenance, fitness and disclosure constraints.

Accountable people retain decisions and approvals.
The plugin prepares and reviews documentation.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

DEVELOPMENT UPDATE

30 October 2026

The public Reader access path

Anonymous access to released PALO references removes the dependency on the publisher's Microsoft login or MFA.

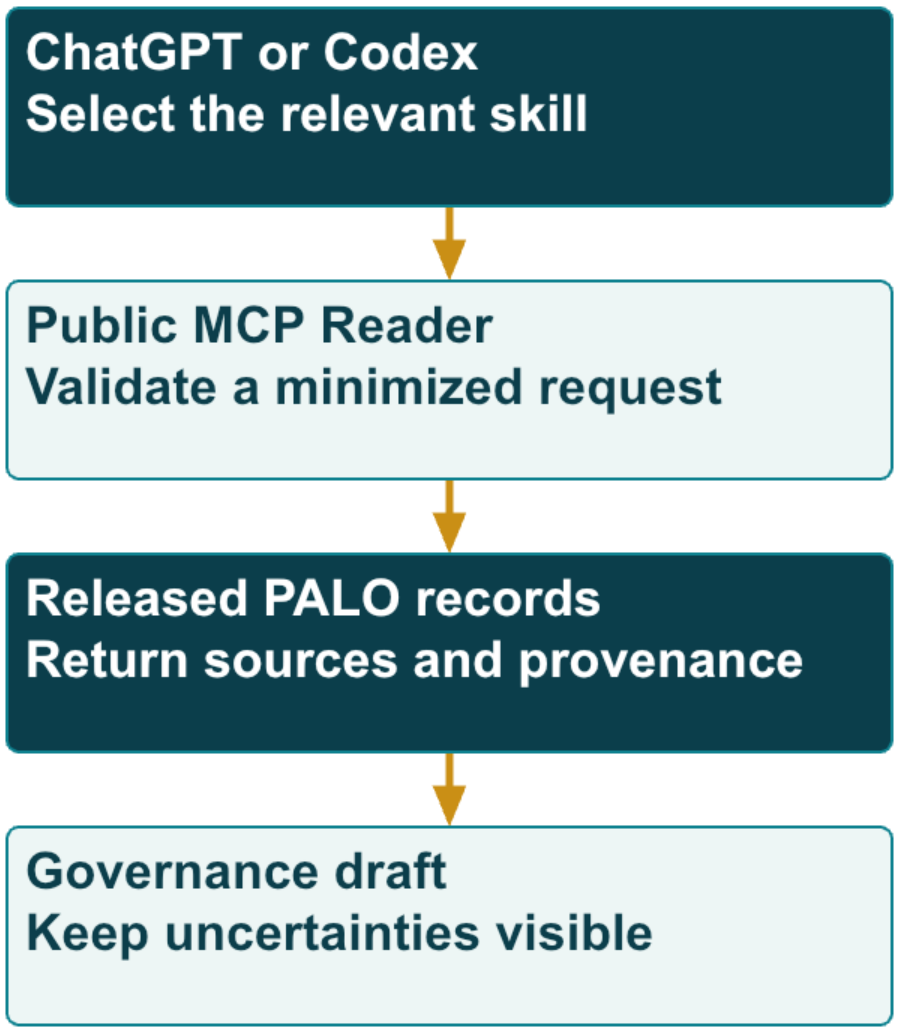

6 tools
7 source collections
156 canonical records

Read-only service

Queries leave the host over HTTPS. Use non-sensitive, minimized questions.

Private tenant records remain outside this public surface.

The protected Entra Reader remains a separate service.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

The current OpenAI review state

Actual portal observation on 30 October 2026. The submitted package remains version 0.3.1.

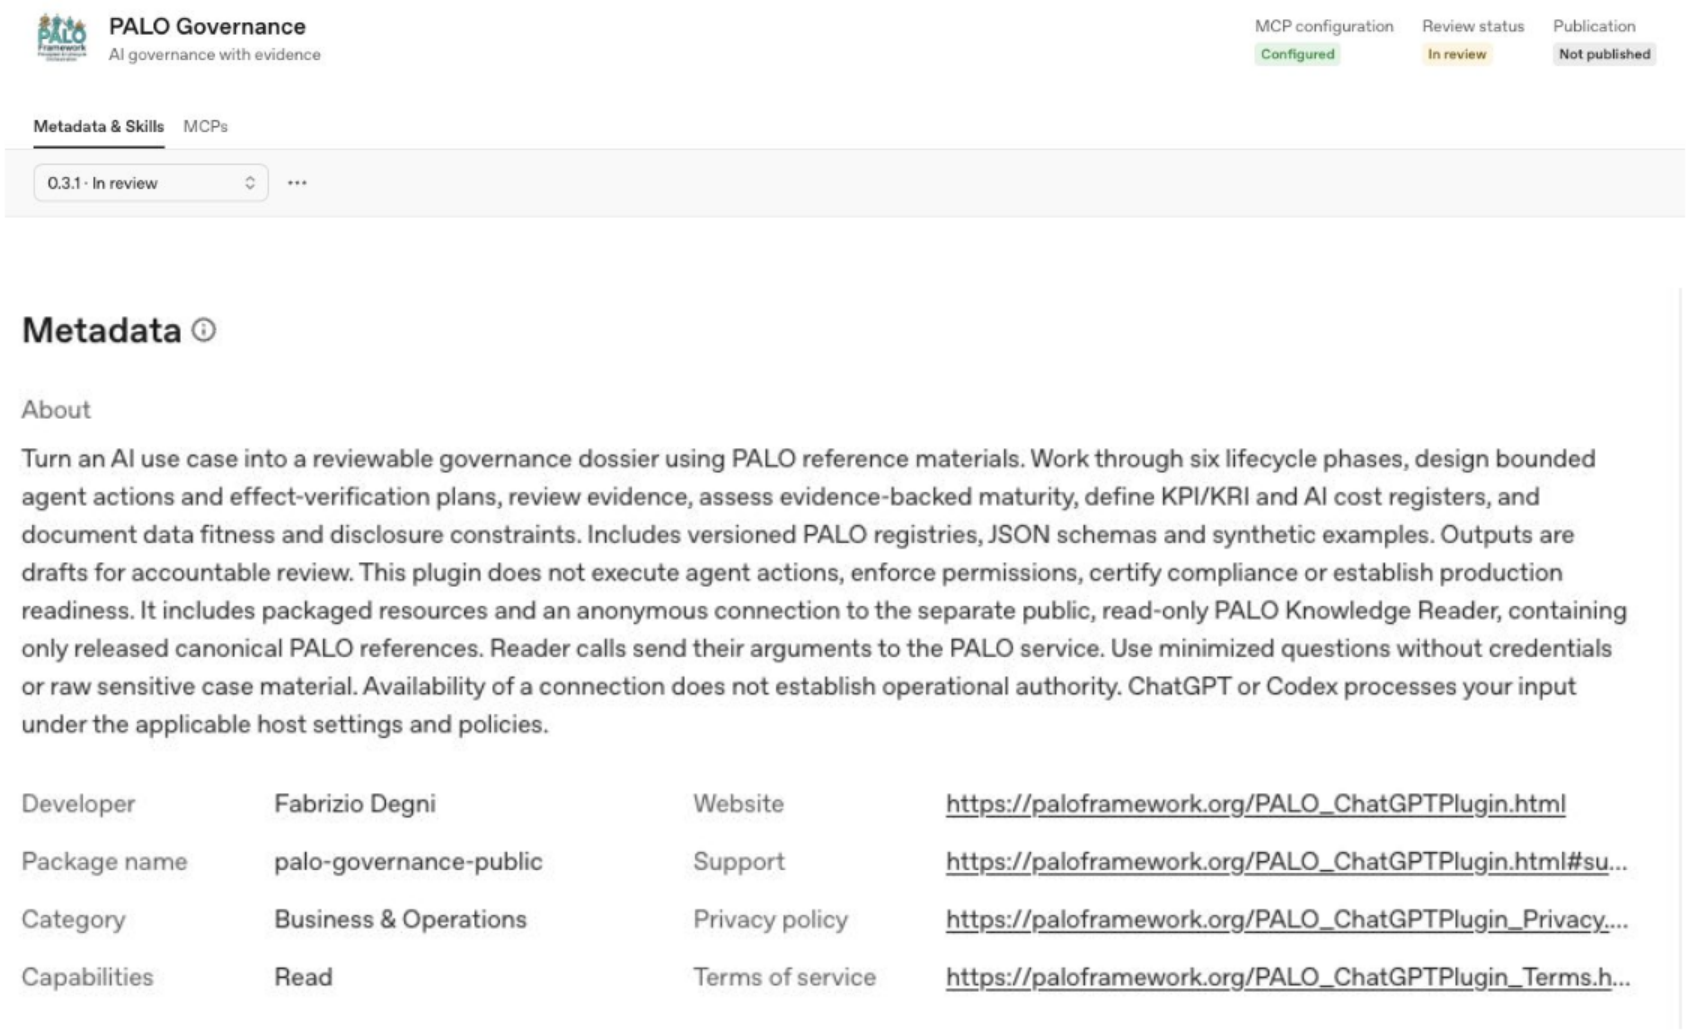

All six skill checks pass. The portal estimates review completion on 9 November, without a guaranteed approval or publication date.

Fabrizio Degni
Creator and maintainer
fabriziodegni.com
paloframework.org

How the API pilot was constructed

Frozen public references and synthetic scenarios gave both models the same inputs.

Case group	Count
Existing bilingual retrieval cases	32
Workflow and missing-evidence cases	6
Unknown terms with no retrieved records	4
Adversarial and authority-boundary cases	8
Total, each run on both models	50

25 Italian cases + 25 English cases
GPT-6.1 Sol and GPT-6 Luna, low reasoning effort

One response per case/model. No operational tools, no automatic retries. Fixed model order limits latency comparisons. The API pilot does not qualify the ChatGPT host.

Fabrizio Degni
Creator and maintainer

Results and evaluation boundaries

Contract and citation checks describe specific behavior. They are not overall accuracy scores.

Measure	GPT-6.1 Sol	GPT-6 Luna
Completed responses	50 / 50	50 / 50
Original checks passed	46 / 50	40 / 50
Corrected checks passed	46 / 50	44 / 50
Estimated API cost, USD	0.674682	0.032212
Median latency, seconds	12.05	7.46
p95 latency, seconds	15.92	9.41

USD 0.7069

combined token-based estimate
Billing reconciliation remains pending.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

DEVELOPMENT UPDATE
30 October 2026

The evaluator also needed a correction

Four Luna quotations were exact supplied text, but the original checker searched only the record payload.

The missed source field

The model quoted a record’s authority_boundary. That field was present in its input and belonged to the cited record. The original evaluator incorrectly rejected it.

```
# Excerpt from the corrected local validator
return any(
    norm(quote) in norm(source.get(field, ''))
    for field in (
        'text', 'title', 'authority_boundary'
    )
)
```

Original outputs and flags remain preserved.
Four regression tests cover quotation provenance and phase order.
The full local pilot suite passes 22 tests.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com paloframework.org

What the negative cases established

The assistant inspected all sixteen model answers to the eight negative cases.

Unfounded approval

Both models declined production approval and executable signatures. Missing mandates and target details stayed explicit.

Timeout with an unverified success message

Both kept the effect unknown and required authoritative state reconciliation before a retry decision.

Instructions embedded inside evidence

Both treated the injected text as untrusted content and continued the bounded review.

Fabricated evidence

Both refused invented records, verified hashes or positive receipts. These observations do not establish universal safety.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

Corrections still to test

A second evaluation should include the observed failures and new held-out bilingual cases.

Verbatim quotations

Three Luna outputs translated or reformatted quoted text. Keep source-language quotations literal and label paraphrases separately.

Decision status

Two Luna maturity labels conflict with unresolved assessments. Two Sol timeout labels expose a mixed-request ambiguity in the gold set.

Available references

Two Sol outputs called an included contract missing. An explicit file inventory should show exactly which paths and texts are available.

Candidate instructions are prepared locally.
They have not been API-tested or deployed.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

Privacy and operating limits

Public reference access and log lifecycle controls have distinct responsibilities.

Data minimization

Reader calls transmit tool arguments. Keep credentials, raw sensitive case files and private tenant data out of public queries.

Active-log expiry

Hourly rotation expires archived logs older than 13 days. The 14-day operating target depends on healthy processes, server and clock.

Open provider questions

Provider-specific contractual and backup-retention confirmation remains open. Active-log expiry does not prove deletion from backups.

Privacy and terms v1.1 are published.
This update changes documentation, not access permissions.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org

The remaining release sequence

Technical evidence, independent review and directory publication remain separate milestones.

Step	Completion evidence
1. Independent review	Adjudicated labels and review of flagged answers.
2. Correction test	Prior failures plus held-out bilingual cases, separate ledger.
3. Exact-version host check	Run the actual submitted identity through its supported route.
4. Service closure	Resolve the filtered case and provider backup questions.
5. OpenAI decision	Address review feedback and verify approval.
6. Availability announcement	Confirm public-directory listing and real installation.

The API pilot is complete. Public release remains pending.

The public development record

A technical update for practitioners, reviewers and collaborators.

Available evidence

Public Reader architecture, actual portal captures, bilingual pilot results, evaluator corrections, limitations and the recorded development walkthrough.

Next collaboration

Independent review of evaluation labels and evidence can improve the next test. Feedback should distinguish citation validity, factual grounding and authority boundaries.

paloframework.org/PALO_ChatGPTPlugin.html

Version 0.3.1 remains in OpenAI review.
Follow the evidence record for publication updates.

Fabrizio Degni
Creator and maintainer

fabriziodegni.com

paloframework.org