

AI Agent Governance Checklist

PRE-DEPLOYMENT & AUTHORIZATION REVIEW

Solutions by Jewel
AI Governance Advisory
solutionsbyjewel.com

Use this checklist before deploying, expanding, or authorizing any AI agent capable of accessing systems, using tools, making decisions, recommending actions, or executing actions. Confirm each item below prior to authorization.

1 AGENT IDENTITY

- ☐ Agent has a unique identity
- ☐ Model/version is documented
- ☐ Agent environment is documented
- ☐ Agent owner is identified
- ☐ Agent activity can be attributed to the agent identity

2 BUSINESS PURPOSE

- ☐ Business purpose is documented
- ☐ Intended outcome is defined
- ☐ Authorized use case is documented
- ☐ Business owner is identified
- ☐ Out-of-scope uses are defined

3 AUTHORITY & ACCESS

- ☐ Permissions are explicitly defined
- ☐ Access is limited to required resources
- ☐ Privileged access is controlled
- ☐ Credentials are managed securely
- ☐ Access can be revoked
- ☐ Authorization is traceable

4 DECISION BOUNDARIES

- ☐ Authorized actions are defined
- ☐ Prohibited actions are defined
- ☐ Human approval requirements are defined
- ☐ Escalation conditions are defined
- ☐ Financial or consequential actions have appropriate controls

5 TOOLS & CONNECTED SYSTEMS

- ☐ Tools available to the agent are inventoried
- ☐ External systems are identified
- ☐ APIs and integrations are documented
- ☐ Data flows are understood
- ☐ Tool permissions match the approved use case

6 SECURITY & RISK

- ☐ Security risk has been assessed
- ☐ Privacy risk has been assessed
- ☐ Data sensitivity has been evaluated
- ☐ Prompt injection / manipulation risks have been considered
- ☐ Agent misuse scenarios have been considered
- ☐ Failure modes are documented

7 HUMAN OVERSIGHT

- ☐ Accountable human owner is identified
- ☐ Review requirements are defined
- ☐ Escalation path exists
- ☐ Human intervention is possible
- ☐ Override or shutdown capability exists where appropriate

8 EVIDENCE & AUDITABILITY

- ☐ Agent decisions/actions are logged
- ☐ Relevant outputs are retained
- ☐ Authorization evidence is retained
- ☐ Exceptions are documented
- ☐ Evidence can be retrieved for review

9 MONITORING

- ☐ Agent activity is monitored
- ☐ Security events are monitored
- ☐ Permission changes are monitored
- ☐ Material behavior changes are identified
- ☐ Alerts and escalation thresholds are defined

10 REASSESSMENT

- ☐ Reassessment trigger is defined
- ☐ Reassessment date is established
- ☐ Model changes trigger review where appropriate
- ☐ Tool or permission changes trigger review where appropriate
- ☐ Business-purpose changes trigger review
- ☐ Material incidents trigger review

AI AGENT GOVERNANCE STATUS

☐ Ready for Deployment

☐ Ready With Controls

☐ Additional Assessment Required

☐ Executive Review Required

☐ Not Authorized

Accountable Owner

Governance Reviewer

Date

Reassessment Date

Decision Record Reference

GOVERNANCE PRINCIPLE

An AI agent should have identifiable authority, defined boundaries, accountable ownership, usable evidence, and a mechanism for continuous oversight and reassessment.

Notes