AI Governance Checklist

15 critical questions every organization should answer before giving AI agents access to business systems.

- [ ] **01 Do we know what business systems the AI agent can access?**
  Map every system the agent can reach before it goes live.
- [ ] **02 Have permissions been reviewed and minimized?**
  Apply least privilege and remove access it does not strictly need.
- [ ] **03 Do we understand what data the AI agent can access?**
  Know which records, files, and fields are in scope.
- [ ] **04 Are sensitive or regulated datasets protected?**
  Give PII, financial, and health data extra safeguards.
- [ ] **05 Can all actions performed by the AI agent be audited?**
  Every action should leave a reviewable trail.
- [ ] **06 Is there human oversight for critical actions?**
  High-impact steps should require human approval.
- [ ] **07 Do we know which MCP Servers or integrations are being used?**
  Maintain an inventory of every connection.
- [ ] **08 Have MCP integrations been reviewed for security risks?**
  Assess each integration before trusting it.
- [ ] **09 Can agent permissions be revoked immediately?**
  You need a fast, reliable way to pull access.
- [ ] **10 Do we have monitoring and observability in place?**
  See what the agent is doing in real time.
- [ ] **11 Can the AI agent be isolated or disabled if needed?**
  A clear kill switch limits damage during incidents.
- [ ] **12 Have compliance requirements been evaluated?**
  Confirm the deployment meets your regulatory obligations.
- [ ] **13 Do users understand what the AI agent can and cannot do?**
  Set clear expectations to prevent misuse.
- [ ] **14 Has an independent risk review been performed?**
  An outside view catches blind spots internal teams miss.
- [ ] **15 Would we be comfortable explaining this agent to an auditor or regulator?**
  If not, governance is not ready yet.

Governance Maturity Model

Use these levels to rate where your organization stands today — and where to go next.

1 — Level 1 — Ad Hoc
No formal governance; access granted case by case.

2 — Level 2 — Managed
Basic policies exist but are applied inconsistently.

3 — Level 3 — Governed
Permissions, reviews, and oversight are standardized.

4 — Level 4 — Trusted
Agents are independently assessed and continuously monitored.

5 — Level 5 — Enterprise Ready
Trust, risk, and compliance are measurable and auditable.

Good governance, at a glance

AI Agent → Trust Layer → Governance Controls → Enterprise Systems

Controls: Permissions · Monitoring · Audit Logs · Risk Management · Compliance · Human Oversight