PRACTICAL GOVERNANCE READINESS GUIDE

Agentic AI Governance Readiness Checklist

A practical, unscored self-review for leaders responsible for AI, Microsoft Copilot, enterprise technology, information governance, security, risk, and operations.

29 questions	**7 governance domains**	**1 action worksheet**

Traditional AI governance asks what people are permitted to do with AI. Agentic AI governance must also address what AI is permitted to do within organizational systems.

Prepared by Patriot Professional Solutions LLC
Version 1.0 | August 2026

Governance-readiness guidance only. This resource is not technical security testing, legal advice, regulatory certification, or a guarantee of AI security or compliance.

How to use this checklist

Use this guide as a facilitated, evidence-first conversation. It is designed to reveal unclear ownership, authority, access, monitoring, and decision rights - not to produce a compliance score.

1	**Assemble the right group**	Include an executive sponsor plus business, IT, security, privacy, risk, information-governance, procurement, and development representatives as relevant.
2	**Rate the actual state**	Select one response for each question. Use current evidence, not intended future state. Record evidence, a required decision, or a next action.
3	**Prioritize patterns**	Pay close attention to Unknown and Not established responses involving authority, sensitive information, privileged access, irreversible actions, or agent shutdown.
4	**Assign action**	Move the most important gaps into the worksheet on page 10. Name an owner, the next decision or action, evidence of completion, and a target date.

Response key

Established	Developing	Not established	Unknown	N/A
Defined, approved, implemented, evidenced, and reviewed.	Partially defined or implemented; evidence or coverage is incomplete.	No documented or consistently operating practice is evident.	The group lacks enough information or evidence to determine the state.	Not applicable, with a documented rationale and accountable owner.

A relationship model for governance

For important use cases, trace ownership, authorized access, decision authority, controls, logging, escalation, human oversight, risk, and documentation across the full chain.

HUMAN	AI	DATA	APPLICATION	IDENTITY	INFRASTRUCTURE

Start with evidence

Useful evidence may include inventories, approval records, policies, architecture diagrams, access reviews, vendor records, data maps, training records, logs, operating procedures, incident exercises, and change approvals. An Unknown response is a useful finding when evidence cannot be produced.

DOMAIN 1 OF 7

Strategy, purpose & accountability

Establish why AI is being used, who is accountable, and how leadership decisions are made.

1	**Has executive leadership approved the organizational purposes, intended outcomes, and risk posture for AI use?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

2	**Is one accountable executive designated for AI governance, with authority to resolve cross-functional decisions and exceptions?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

3	**Are business and technical owners identified for each material AI use case?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

4	**Does a recurring governance forum review AI adoption, unresolved risks, exceptions, and priority decisions?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

DOMAIN 2 OF 7

Inventory, ownership & approved use

Know which AI capabilities exist, why they are used, and whether they are approved.

5	**Does a maintained inventory identify AI tools, models, Copilots, agents, use cases, owners, deployment environments, and connected systems?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

6	**Can teams distinguish approved, pilot, unapproved, and retired AI uses?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

7	**Do approval criteria address intended purpose, users, information handled, capabilities, and risk before deployment?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

8	**Are inventories reviewed when AI features are enabled in existing products or vendors materially change capabilities?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

DOMAIN 3 OF 7

Information & data governance

Understand what information AI can reach and apply the organization's existing information rules.

9	**Is information available to each AI system documented, including Microsoft 365, SharePoint, Teams, repositories, databases, and external sources?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

10	**Are classification, retention, privacy, records, contractual, and information-boundary requirements applied to AI use?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

11	**Are broad, inherited, or public permissions reviewed before AI tools can discover, retrieve, or summarize organizational content?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

12	**Are rules defined for sensitive information, prompts, outputs, training data, retrieval sources, and data sharing with third parties?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

DOMAIN 4 OF 7

Identity, permissions & credentials

Control which identities AI uses, what they can access, and how credentials are governed.

13	**Does each agent or AI workload use an identifiable, accountable identity rather than shared credentials where feasible?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

14	**Are delegated, API, and service-account permissions least-privilege, approved, inventoried, and periodically reviewed?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

15	**Are keys, tokens, and secrets stored, rotated, and revoked under assigned ownership - and kept out of prompts, repositories, and logs?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

16	**Is separation of duties defined between developing, approving, deploying, operating, and reviewing AI systems and agents?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

DOMAIN 5 OF 7

Human oversight, monitoring & escalation

Define where people must approve action and how the organization detects, explains, and stops unwanted activity.

17	**Are human approval points defined for high-impact AI actions such as publishing, communicating, changing records, approving, deleting, executing, or deploying?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

18	**Can logs show which AI or agent acted, when it acted, the identity or credential used, the system or data affected, and approval context?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

19	**Are monitoring ownership, review cadence, anomaly criteria, and documentation responsibilities assigned?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

20	**Are escalation, containment, agent shutdown, stakeholder notification, and evidence-preservation responsibilities documented and exercised?**

[ ] Established	[ ] Developing	[ ] Not established	[ ] Unknown	[ ] N/A

Evidence, decision, or next action: _____

DOMAIN 6 OF 7

Supply chain, development & workforce readiness

Govern the third-party components, generated code, and people that shape AI-enabled operations.

21	**Is there a supplier and component register for third-party models, datasets, libraries, packages, plugins, connectors, APIs, and external AI services?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

22	**Do procurement, architecture, security, and risk reviews consider data use, permissions, provenance, support, change notices, exit needs, and vendor dependencies?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

23	**Are AI-generated code and dependencies reviewed, tested, approved, and documented under secure development and change-management practices before deployment?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

24	**Do users, developers, administrators, owners, and reviewers receive role-based AI training and acknowledge acceptable-use and oversight responsibilities?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

DOMAIN 7 OF 7

Agentic AI security & governance

Govern what autonomous and semi-autonomous AI is permitted to do within organizational systems.

25	**Is every autonomous or semi-autonomous agent documented with its purpose, owner, environment, trigger, tools, connected applications and data, identity, and decision scope?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

26	**For each agent, is action authority explicit - including read, write, modify, delete, execute, publish, communicate, approve, deploy, query, download, upload, and API invocation?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

27	**Are prohibited actions and human-authorization boundaries enforced for material, irreversible, or high-risk actions - not merely stated in policy?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

28	**Can authorized personnel pause or disable an agent and revoke its identity, permissions, and credentials quickly, with named shutdown authority?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

29	**Are agent changes, tool additions, model updates, permission expansions, connector changes, and foreseeable failure modes reviewed before and after deployment?**

[ ] Established | [ ] Developing | [ ] Not established | [ ] Unknown | [ ] N/A

Evidence, decision, or next action: _____

Agentic governance decision

The central question is not only whether an AI output is acceptable. It is whether the agent's identity, information access, tools, action authority, approval boundaries, monitoring, and shutdown path are proportionate to what it can do.

Priority-action worksheet

Select no more than five issues that require near-term leadership attention. Prioritize material authority, access, information, monitoring, and accountability gaps before lower-impact documentation refinements.

#	Gap or required decision	Why it matters	Owner	First action / evidence	Target
1					
2					
3					
4					
5					

Shape a 90-day governance path

Timing	Objective	Decisions, actions, owners, and evidence
0-30 days	Stabilize visibility and decision rights	
31-60 days	Prioritize controls and operating practices	
61-90 days	Implement, verify, and establish review cadence	

Suggested prioritization test

For each issue, consider impact, likelihood, reach, reversibility, sensitivity of information, level of privilege, existing safeguards, dependencies, and the clarity of human accountability. Avoid treating all checklist responses as equal.

From checklist to governed action

The checklist is most valuable when it produces explicit decisions, assigned owners, and verifiable next steps. It is not a maturity certification and should not be used as a substitute for technical testing, legal review, or specialized assurance.

What the response patterns may indicate

■ Multiple Unknown responses may reveal an inventory, evidence, or accountability gap before they reveal a control gap.

■ Not established responses involving privileged access, sensitive information, external communications, destructive actions, or deployments generally warrant early review.

■ Developing responses need a named completion condition: what will be implemented, by whom, by when, and what evidence will demonstrate operation.

■ N/A should include a rationale and owner because tools, agents, integrations, and vendor features change over time.

A practical engagement path

ASSESS	PRIORITIZE	IMPLEMENT	SUSTAIN
Establish current state, evidence, ownership, and significant governance gaps.	Rank decisions and actions by risk, dependency, effort, and leadership importance.	Execute approved policy, process, information, access, or operating improvements.	Review evidence, changes, exceptions, accountability, and roadmap progress over time.

Implementation and sustainment work are separately scoped based on priorities, authority, dependencies, and the capabilities required.

Professional boundaries

PPS's governance-readiness work does not constitute penetration testing, red-team testing, vulnerability scanning, malware analysis, AI model assurance, formal AI safety certification, managed detection and response, Security Operations Center services, incident response, legal advice, regulatory certification, or a guarantee of AI security or compliance. Specialized technical testing or assurance may require a separately scoped, appropriately qualified provider.

When a structured assessment would help

If this checklist reveals unclear ownership, permissions, decision rights, information exposure, agent authority, or oversight, Patriot Professional Solutions can help organize the evidence and decisions through a scoped AI, Copilot & Agentic Governance Readiness assessment.

Explore the readiness assessment | Request a written fit review

Patriot Professional Solutions LLC

Governance, documentation, and technology-readiness support for leaders who need practical decisions, accountable execution, and usable roadmaps.

patriotprofessionalsolutions.com