

Agentic AI Readiness Assessment (AARA™)

Strategy & Business Alignment

1. Our organization has a documented AI or Agentic AI strategy

A written plan defining where, why, and how autonomous workflows will be deployed.

Scale: 1 (Lowest) to 5 (Highest)

2. Leadership understands where autonomous workflows create value

Executives can name the processes where agents reduce cost, risk, or cycle time.

Scale: 1 (Lowest) to 5 (Highest)

3. We can identify our top three agentic AI opportunities

Concrete, high-impact candidates for supervised or autonomous execution.

Scale: 1 (Lowest) to 5 (Highest)

4. AI investments are prioritized using business outcomes

Funding follows defined success metrics, not hype or one-off requests.

Scale: 1 (Lowest) to 5 (Highest)

Data & Knowledge Readiness

5. Critical business knowledge is centralized and organized

Key documents, policies, and data live in known, governed locations.

Scale: 1 (Lowest) to 5 (Highest)

6. Data owners and stewardship are clearly identified

Someone is accountable for the accuracy and access of each domain.

Scale: 1 (Lowest) to 5 (Highest)

7. Information is searchable and retrievable on demand

People and systems can find the right content quickly and reliably.

Scale: 1 (Lowest) to 5 (Highest)

8. We would trust an AI agent to act using our current knowledge assets

The agentic multiplier: would you let an agent act on what it can retrieve today?

Scale: 1 (Lowest) to 5 (Highest)

Governance, Risk & Trust

9. We have defined approval processes for deploying AI systems

A known path from idea to production with the right sign-offs.

Scale: 1 (Lowest) to 5 (Highest)

10. AI outputs are reviewable and auditable

You can trace what an AI system did, why, and on what basis.

Scale: 1 (Lowest) to 5 (Highest)

11. Human override and escalation mechanisms exist

A person can stop, correct, or escalate an agent at any time.

Scale: 1 (Lowest) to 5 (Highest)

12. We define which decisions agents may execute autonomously

Agentic-specific: explicit boundaries on what an agent can do without a human.

Scale: 1 (Lowest) to 5 (Highest)

Workforce & Operating Model

13. Employees understand AI capabilities and limitations

People know what agents do well, and where humans must stay in the loop.

Scale: 1 (Lowest) to 5 (Highest)

14. Leaders actively model AI usage

Executives use AI visibly, signaling it is safe and expected.

Scale: 1 (Lowest) to 5 (Highest)

15. Teams receive practical, role-specific AI training

Training is hands-on and tied to real workflows, not theory.

Scale: 1 (Lowest) to 5 (Highest)

16. Employees view AI as augmentation, not a threat

Culture treats agents as teammates that remove drudgery, not replacements.

Scale: 1 (Lowest) to 5 (Highest)

Agentic Architecture & Technology

17. Core systems expose APIs that agents can call

Your key platforms can be driven programmatically, not just by humans.

Scale: 1 (Lowest) to 5 (Highest)

18. Identity and access management is automated

Machine and agent identities are provisioned and controlled, not manual.

Scale: 1 (Lowest) to 5 (Highest)

19. AI interactions are logged and monitored

Every agent action is observable for debugging, safety, and audit.

Scale: 1 (Lowest) to 5 (Highest)

20. We support AgentOps, LLMOps, observability, and tool orchestration

Agentic-specific: the operational backbone for running agents reliably.

Scale: 1 (Lowest) to 5 (Highest)

AI Portfolio & Execution Maturity

21. We measure AI business outcomes and ROI

You know what your AI work returns, not just that it exists.

Scale: 1 (Lowest) to 5 (Highest)

22. AI solutions are reused across teams

Successful patterns become shared assets, not one-off builds.

Scale: 1 (Lowest) to 5 (Highest)

23. AI projects move from pilot to production

Promising pilots actually ship and stay in operation.

Scale: 1 (Lowest) to 5 (Highest)

24. We can scale a successful agent across the organization

A win in one team can be repeated everywhere without starting over.

Scale: 1 (Lowest) to 5 (Highest)