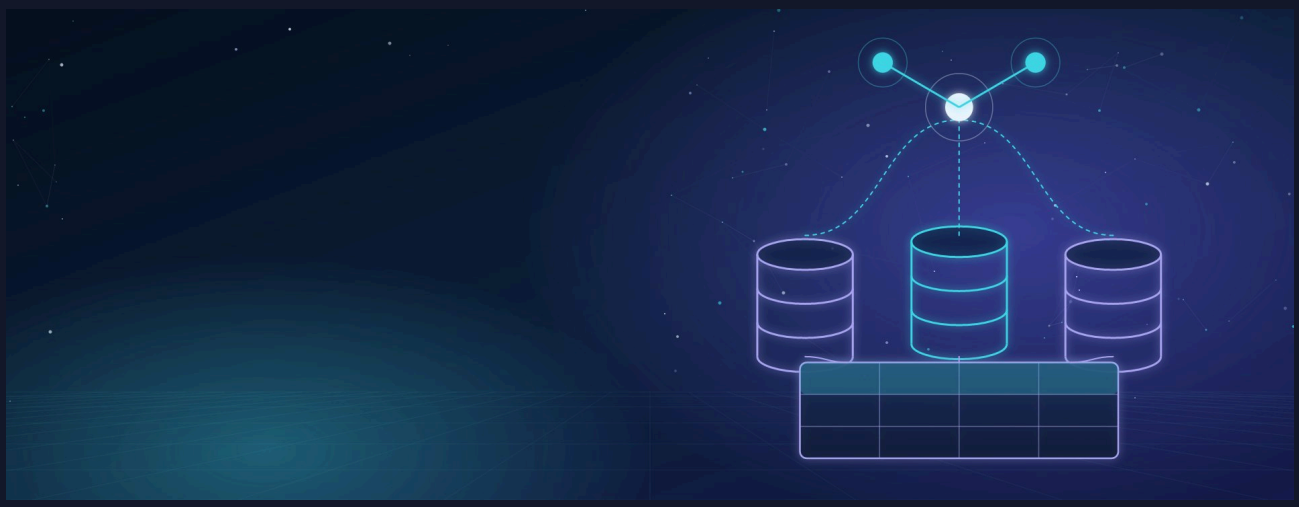
DATAEPIX
Enterprise AI, built on data you can trust.

FREE CHECKLIST

AI Readiness & Governance Checklist

7 steps to AI agents and knowledge assistants built on data you can trust



For business and technology leaders planning AI agents, RAG knowledge assistants, or AI-ready data platforms.

BEFORE YOU START

Who this checklist is for, and how to use it

This checklist is for business and technology leaders who want AI that people can rely on: agents that complete real tasks, assistants that answer from company knowledge, and the data foundation both depend on.

It covers seven steps, from choosing a use case to running AI in production. Most of them are about data: whether it is complete, current, clearly defined, and reachable with the right permissions. That is where most AI projects succeed or stall.

How to use it

- Pick one AI use case you are planning or already running. Score that use case, not the whole company.
- Work through it with the business owner, a data owner, and whoever will build or run the AI. It takes 10 to 20 minutes.
- Score each check: **Yes = 2**, **Partly = 1**, **No = 0**. Total your score on the last page.
- Agree on the two or three gaps to close first, with an owner and a date for each.

What you will get

A clear view of whether a use case is ready to launch, ready for a focused pilot, or needs data work first, and a short list of next actions you can take in the next 30 days.

This checklist is general guidance, not legal advice. Frameworks referenced: NIST AI Risk Management Framework (AI RMF 1.0) and ISO/IEC 42001 (AI management systems).

THE CHECKLIST

Steps 1 and 2: Start with the use case and the data

STEP 1

Choose a use case worth the effort

AI pays off when it is tied to one process, one owner, and one number you can measure.

Check	Yes (2)	Partly (1)	No (0)
A named business owner and a clear process or decision the AI will support.			
A measurable target (time saved, accuracy, cost, or revenue) and today's baseline.			
A defined scope: which users, which tasks, and what is out of scope.			
A decision on autonomy: will the AI suggest, draft, or act on its own?			

STEP 2

Know the data it depends on

Most AI shortfalls trace back to the data, not the model.

Check	Yes (2)	Partly (1)	No (0)
Every source the use case needs is listed, with a named owner.			
Quality is checked: completeness, duplicates, freshness, and known errors.			
Documents are current, with outdated and draft versions removed or labelled.			
You know how much of the needed data the AI can actually reach today.			

Steps 3 and 4: Shared definitions and the right access

STEP 3

Agree on what the data means

An assistant cannot give a reliable answer when two systems disagree on a definition.

Check	Yes (2)	Partly (1)	No (0)
Key business terms (customer, order, active, revenue) have one agreed definition.			
Definitions live in a glossary or semantic layer that people and AI both use.			
Where sources conflict, a declared source of truth decides.			
Changes to definitions have an owner and a review step.			

STEP 4

Set the right access

AI should see what the task needs, and nothing more.

Check	Yes (2)	Partly (1)	No (0)
The AI runs under its own service identity, not a shared or personal account.			
Search and retrieval respect the user's existing permissions.			
Personal and confidential data is identified, and excluded or masked where it is not needed.			
Agents get only the access the task needs, with no standing delete or payment rights.			

Steps 5 and 6: Grounded answers, tested before launch

STEP 5

Ground answers and actions in approved sources

Your organization is accountable for what its AI says and does.

Check	Yes (2)	Partly (1)	No (0)
Answers come only from approved, current sources, with citations users can check.			
The assistant says it does not know, or hands off to a person, when sources do not cover a question.			
High-impact or hard-to-reverse actions need human approval.			
Every answer and action is logged with its sources and inputs.			

STEP 6

Test before you launch

A short, honest test beats a confident demo.

Check	Yes (2)	Partly (1)	No (0)
A test set of real questions and tasks, with expected answers agreed by the business owner.			
An accuracy target the use case must reach before go-live.			
Sensitive topics (pricing, policy, legal, medical, financial) are tested explicitly.			
A pilot group and a rollback plan are in place.			

Step 7: Run, measure, and improve

STEP 7

Run, measure, and improve

AI in production needs an owner, a budget, and a feedback loop.

Check	Yes (2)	Partly (1)	No (0)
A named owner tracks accuracy, usage, cost, and user feedback after launch.			
The data and documents the AI relies on are refreshed on a schedule.			
An AI use policy and inventory cover every AI tool in use, aligned to the NIST AI RMF or ISO/IEC 42001.			
There is a clear process to report, review, and fix wrong answers or actions.			

Your score

Total (out of 56)	What it means
45 to 56	Ready to launch and scale. Keep testing and monitoring as usage grows.
28 to 44	Ready for a focused pilot. Close the gaps in Steps 2 to 5 before a wider rollout.
0 to 27	Start with the data foundation and a scoped plan before building.

My score: _____ Top 3 gaps to close: _____

NEXT STEP

Turn your score into a plan

If the checklist surfaced gaps, the fastest way forward is a short, fixed-scope assessment of one use case.

AI Strategy & Opportunity Sprint: USD 7,500

- Review your top AI opportunities and the data each one depends on.
- Score readiness, risk, and expected return for each use case.
- Leave with a prioritized roadmap and a scoped first project.

When you are ready to build, Data Epix launch packages for AI agents, knowledge assistants (RAG), and AI-ready data start from USD 19,500, each with a clear scope, timeline, and deliverables.

Talk to an AI solutions architect

Book a free AI strategy call at **dataepix.com/contact-us** or call **+1 (619) 377-8090**. We reply within one business day.

Data Epix
11440 W. Bernardo Court, San Diego, CA 92127

DATAEPIX
Enterprise AI, built on data you can trust.