

AI-READINESS SELF-ASSESSMENT

AI-Readiness Checklist

A Guide to Auditing Your Business for AI

A practical, fill-in audit tool for leaders who want an honest read on whether their organization is actually ready to build with AI — across data, infrastructure, security, people, use cases, and how you operate. Work through it with your team, check what's true today, and score each domain.



Readiness is built domain by domain — not bought in a single tool.

6 DOMAINS • 44 CHECKPOINTS

01
Data
Foundation

02
Infrastructure
& Compute

03
Security &
Governance

04
Talent
& Skills

05
Use-Case
& Value

06
Operating
Model

"Not AI. AI for Humans."

DataExos — cloud + data + AI
www.dataexos.com • 1-800-460-2261

How to use this workbook

1

Assemble the right people

Bring together the voices that actually know the answers — data, IT/engineering, security, and a business owner. Readiness is cross-functional; a solo guess isn't an audit.

2

Check what's true today

For each domain, tick only the statements that describe your organization *right now* — not what you intend to do. Empty boxes are your backlog, not a failure.

3

Score & interpret

Rate each domain 1–5 on the maturity scale, carry the six scores to the results page, total them, and read your readiness tier and recommended next moves.

Maturity scoring key

RATE EACH DOMAIN 1–5

1

Ad hoc. No defined approach. Handled reactively, informally, or not at all.

2

Aware. The need is recognized and early, informal efforts are underway — but nothing is standardized.

3

Defined. A documented approach exists and is understood, though it's applied inconsistently across teams.

4

Managed. The approach is applied consistently, owned by someone accountable, and measured against targets.

5

Optimized. Continuously improved and embedded in how the organization operates day to day.

Be candid. This tool is only useful if you score what *is*, not what you hope. There is no passing grade to perform for — an honest low score on a domain is exactly the signal that tells you where to invest first. Human judgment stays in the loop throughout: **"Not AI. AI for Humans."**

Data Foundation & Infrastructure



DOMAIN 01

Data Foundation

AI runs on data. Before models, ask whether the data that would feed them is known, trusted, governed, and reachable.

- ☐ We maintain an inventory of the data sources that would feed AI use cases — systems, owners, and refresh cadence are documented.
- ☐ Critical datasets have clear definitions and lineage, and each has a named data owner accountable for it.
- ☐ We measure data quality — completeness, accuracy, timeliness — and track it against thresholds rather than assuming it's fine.
- ☐ Sensitive and regulated data is classified and tagged (PII, PHI, financial), so we know what we hold and where.
- ☐ Data flows through governed pipelines, not manual exports and one-off spreadsheets, and can be consolidated across silos without heavy reconciliation.
- ☐ Retention, consent, and usage rights are documented — including whether data may be used to train or fine-tune models.

Domain maturity

Circle one · 1–5

AD HOC

1

2

3

4

5

OPTIMIZED



DOMAIN 02

Infrastructure & Compute

AI workloads have different appetites than typical apps. Assess whether your compute, environments, and costs can carry training and inference at scale.

- ☐ We understand our compute profile — training vs. inference, CPU vs. GPU/accelerator — for the workloads we're considering.
- ☐ We have a path to scalable, elastic compute (cloud, hybrid, or on-prem) that can flex up and down with demand.
- ☐ Cloud and accelerator spend is monitored with budgets and cost controls, so AI workloads don't surprise us on the invoice.
- ☐ Environments are reproducible and provisioned as code, with parity across development, test, and production.
- ☐ We have a model-serving strategy — managed API, self-hosted, or hybrid — and know the latency, throughput, and storage it demands.
- ☐ Data-residency and region requirements are accounted for in how infrastructure is designed and where workloads run.

Domain maturity

Circle one · 1–5

AD HOC

1

2

3

4

5

OPTIMIZED

Security, Governance & People



DOMAIN 03

Security & Governance

Trust is a precondition, not a feature. Assess whether AI decisions are governed, risks are assessed, and humans stay accountable for high-stakes output.

- ☐ **An accountable owner or committee governs AI** — policies, acceptable use, and approval of new use cases sit with someone specific.
- ☐ **AI and model risks are assessed before deployment** — bias, privacy, security, and potential for misuse are reviewed, not discovered later.
- ☐ **Access to data, models, and AI tools follows least privilege** and is reviewed periodically; data is encrypted in transit and at rest with managed keys.
- ☐ **Human-in-the-loop review is defined for high-impact output** — a person, not a model, owns consequential decisions. **"Not AI. AI for Humans."**
- ☐ **AI inputs and outputs are logged and monitored** for auditability, incident response, and the ability to explain what a system did.
- ☐ **Our approach is aligned with recognized AI-governance and security frameworks** (compliance-aware — an alignment target, not a certification claim).
- ☐ **Third-party and vendor AI is governed** — data handling, model provenance, and terms are reviewed before we route data through someone else's model.

Domain maturity

Circle one · 1–5

AD HOC

1

2

3

4

5

OPTIMIZED



DOMAIN 04

Talent & Skills

Tools don't deliver AI — people do. Assess whether the skills, understanding, and ownership exist to build, adopt, and sustain it.

- ☐ **Leadership shares a realistic view of AI** — what it can do, what it can't, and where the risk lives — not just enthusiasm.
- ☐ **We can access the skills to build and integrate AI** — data engineering, ML/MLOps, and solution design — in-house or through partners.
- ☐ **Business teams can spot and scope AI opportunities** in their own domains, rather than waiting for a central team to imagine them.
- ☐ **Staff are trained on responsible, effective use of AI tools**, including where to be skeptical of model output.
- ☐ **Roles and responsibilities for AI work are defined** — initiatives have owners, not just "someone's side project."
- ☐ **We plan for capability gaps deliberately** — hiring, partnering, or upskilling — and manage the change AI brings to workflows and roles.

Domain maturity

Circle one · 1–5

AD HOC

1

2

3

4

5

OPTIMIZED

Value & the Operating Model



DOMAIN 05

Use-Case & Value

AI earns its place by solving real problems. Assess whether use cases are chosen for value and feasibility — and whether that value is actually measured.

- ☐ We have an inventory of candidate use cases tied to real business problems, not to hype or a tool someone wants to try.
- ☐ Use cases are prioritized by value, feasibility, and risk — a clear-eyed shortlist, not a wish list ranked by novelty.
- ☐ Each priority use case has a success metric and a baseline to measure against, defined before the build starts.
- ☐ We can estimate total cost — build, run, and maintain — before committing, and we account for the cost of the model being wrong.
- ☐ We can pilot quickly, then scale or kill on evidence, rather than letting a stalled proof-of-concept linger indefinitely.
- ☐ Value realization is tracked after deployment, not just celebrated at launch — we know whether it actually paid off.

Domain maturity

Circle one · 1–5

AD · HOC

1

2

3

4

5

OPTIMIZED



DOMAIN 06

Operating Model & MLOps

Getting to production is the start, not the finish. Assess whether AI systems are versioned, monitored, owned, and safe to change once they're live.

- ☐ AI systems have a defined lifecycle — build, validate, deploy, monitor, retire — rather than being launched and forgotten.
- ☐ Deployments are versioned and repeatable — CI/CD for models and prompts, so releases are controlled, not hand-crafted.
- ☐ Models in production are monitored for drift and degradation, and there's a defined path to retrain, update, or roll back.
- ☐ AI incidents have an escalation and response path, so a bad output has a clear owner and a clear next step.
- ☐ Systems are documented, so running AI doesn't depend on one person's memory — knowledge survives a departure.
- ☐ Ownership for running AI is clear — who's accountable, who's on call, and who signs off on changes is not ambiguous.

Domain maturity

Circle one · 1–5

AD · HOC

1

2

3

4

5

OPTIMIZED

Score & interpret

DOMAIN	SCORE (1-5)
01 Data Foundation	
02 Infrastructure & Compute	
03 Security & Governance	
04 Talent & Skills	
05 Use-Case & Value	
06 Operating Model & MLOps	
Total readiness score	/ 30

6-13 TIER 1	Foundational — start with the basics The building blocks aren't in place yet. Focus on data quality and ownership, a minimal governance stance, and one clear, low-risk pilot to build momentum and learn.
14-19 TIER 2	Developing — prove value, tighten control You have real building blocks. Tighten governance and security, put human-in-the-loop review where it matters, and prove measurable value on one or two use cases before scaling.
20-25 TIER 3	Operational — ready to scale The foundations hold. Invest in MLOps, monitoring, and repeatable delivery so you can run a growing portfolio of AI reliably rather than heroically.
26-30 TIER 4	AI-Native — optimize for advantage AI is embedded in how you operate. Shift focus to portfolio-level governance, continuous optimization, and turning capability into durable competitive advantage.

This scoring is a directional self-assessment to guide prioritization and conversation — not a benchmark, audit, or guarantee of outcomes. Your lowest-scoring domain is usually the most valuable place to look next.

WHERE TO GO FROM HERE

Turn the gaps into a plan.

DataExos is a cloud, data, and AI partner built for organizations that want to adopt AI without handing over judgment. If this workbook surfaced gaps, that's the starting point — we help you close them in the right order, with humans firmly in the loop.

Cloud
Scalable, cost-aware infrastructure for AI workloads.

Data
Governed, quality data foundations you can build on.

AI
Use cases delivered with security and human oversight.

www.dataexos.com

1-800-460-2261

"Not AI. AI for Humans."

© 2026 DataExos, LLC — a Delaware limited liability company

AI-Readiness Checklist · 06