DR. AHMAD KHOKHAR

AI Readiness Checklist for CEOs

A board-friendly readiness scan for government, healthcare, and enterprise leaders.

Executive briefing for leaders building secure, governed, production-grade AI systems.

AI Architecture | Private LLMs | Healthcare AI | Smart Cities | Government AI

AI Readiness Checklist for CEOs

AI readiness is not measured by how many tools an organization has tested. It is measured by whether leadership can connect AI to outcomes, data, workflows, governance, security, adoption, and accountability.

Author's perspective

Dr. Ahmad Khokhar recommends that leaders evaluate readiness before vendor selection. The wrong sequence creates scattered pilots. The right sequence creates a reusable AI architecture that can scale across departments.

Executive readiness

- A senior owner is accountable for AI outcomes, not only procurement or experimentation.
- Priority use cases are ranked by value, urgency, risk, and feasibility.
- The board or executive team understands where AI may assist, recommend, automate, or only summarize.
- The organization has defined what success looks like: time saved, quality improved, access expanded, risk reduced, or oversight strengthened.
- Budgeting accounts for integration, governance, evaluation, training, and support, not only software licensing.

Readiness scoring model

Use this as a leadership diagnostic. Low scores do not mean AI is impossible; they show where architecture work must happen before scale.

Dimension	Score 1: weak	Score 5: production ready
Strategy	AI interest is generic and tool-led.	AI priorities are tied to clear institutional outcomes and executive ownership.
Data	Sources are scattered, unclassified, and poorly governed.	Data owners, sensitivity, quality, lineage, and permissions are known.
Workflow	Processes are informal or undocumented.	Workflow owners, decision points, bottlenecks, and review paths are mapped.
Governance	No clear rules for use, review, escalation, or logging.	Policies, approval paths, audit logs, evaluation, and risk controls exist.
Security	Sensitive data boundaries are unclear.	Access control, hosting model, privacy, and incident response are defined.

Dimension	Score 1: weak	Score 5: production ready
Adoption	Users are expected to adapt after deployment.	Training, change management, and feedback loops are planned from the start.

Critical questions before vendor selection

- What workflow will the AI system improve, and who owns that workflow?
- Which data sources are required, and who has authority to approve their use?
- What data may never be sent to public models or external vendors?
- What output quality is required before the system can affect real operations?
- Which decisions require human review regardless of model confidence?
- How will the organization audit, investigate, and correct AI mistakes?
- What deployment model fits the risk profile: public cloud, private cloud, hybrid, sovereign cloud, or on-prem?
- Which internal team will operate the system after the pilot?

CEO dashboard for AI readiness

Metric	Why it matters	Example target
Use case clarity	Prevents generic experimentation.	Top five workflows ranked by value and risk.
Data availability	Determines whether AI can produce reliable outputs.	Critical sources identified with owners and access rules.
Review coverage	Keeps accountability intact.	All high-impact workflows mapped to human review points.
Evaluation cadence	Detects failure, drift, and misuse.	Monthly quality review for live systems.
Auditability	Supports trust, compliance, and incident review.	All sensitive queries and outputs logged.
Adoption health	Shows whether AI is changing work, not just producing demos.	Usage, override, satisfaction, and cycle-time metrics tracked.

What a readiness audit should produce

- An AI opportunity map by department, workflow, value, risk, and readiness.
- A data and integration map showing systems, documents, owners, and access boundaries.
- A governance model defining review, escalation, logging, evaluation, and approval rules.
- A recommended first pilot with measurable outcomes and controlled risk.
- A 90-day implementation roadmap and a 12-month institutional AI architecture roadmap.

The purpose of readiness work is not to slow AI adoption. It is to prevent expensive fragmentation and make the first serious deployment strong enough to become reusable infrastructure.

Executive discussion guide

Use these questions to move the conversation from interest in AI to a serious operating decision. They are designed for boards, founders, ministers, hospital executives, CXOs, program leaders, and technical teams that need a shared view of readiness and risk.

- What institutional outcome will improve if this AI system succeeds, and how will that improvement be measured?
- Which data sources, documents, systems, and permissions are required for the workflow to operate safely?
- Where does AI assist, where does it recommend, where can it automate, and where must it stop for human review?
- Who owns the final decision when AI output influences a citizen, patient, customer, employee, budget, safety, or compliance outcome?
- What evidence, citations, logs, monitoring, and evaluation will be available when leadership needs to audit the system?
- Which deployment model fits the data sensitivity, latency, cost, resilience, and governance requirements?

Leadership lens	What to verify	Evidence of maturity
Value	The use case has measurable operational, clinical, financial, service, or oversight value.	Baseline metrics and target outcomes are documented.
Risk	Sensitive decisions, data exposure, safety impact, and reputational risk are understood.	Risk register, review rules, and escalation paths exist.
Governance	Policy is translated into daily operating controls.	Role matrix, audit logs, approval flows, and model evaluation cadence exist.
Scale	The first deployment can become reusable institutional capability.	Reusable data, retrieval, model, workflow, and monitoring services are planned.

Dr. Ahmad Khokhar's recommended leadership discipline is simple: do not approve AI scale until the organization can explain value, data, workflow, governance, deployment, and human accountability in one coherent architecture.

Recommended next step

A readiness review should happen before major AI procurement or large-scale implementation. It gives leadership a defensible roadmap for where AI can create value, what must be governed, and how to scale responsibly.

For confidential institutional discussions, project details should be scoped under appropriate confidentiality expectations. Sensitive government, healthcare, security, or enterprise matters can be summarized at the architecture-pattern level before deeper review.

Contact: drk@drkhokhar.com | drkhokhar.com