LEADERSHIP DECISION TOOL

AI Quick Scope

A leadership decision tool for scoping AI use cases, data readiness, and operating controls.

Purpose

Use this tool before selecting, piloting, integrating, or expanding an AI-enabled workflow. It converts a broad AI idea into a bounded decision: what will the system do, what data may it use, who is accountable, and what must be true before it operates.

The standard is not perfect data. The standard is a defined purpose, trusted data assets, accountable ownership, and operating controls strong enough to make AI repeatable rather than accidental.

Organization	**AI use case or business process**
Executive owner	**Assessment date**
Operational owner	**Risk / privacy reviewer (if applicable)**

Designed for: executive sponsors, operations leaders, data owners, technology leaders, privacy/compliance contacts, and security advisors.

START HERE

How to use AI Quick Scope

1	**Assess one use case at a time.** A customer-service assistant, forecasting workflow, HR screen, and security triage tool have different data, error, and oversight needs.
2	**Score each readiness check.** Mark No, Partial, or Yes. A useful Yes names the system, owner, policy, report, contract, test, or decision record that supports it.
3	**Treat critical guardrails as gates.** Ownership, approved data boundaries, human authority, and operational fallback require deliberate leadership attention before risk scales.
4	**Decide and document the next action.** Use the decision page to pause, run a controlled pilot, proceed with controls, or require executive review.

Scoring model

Score	Meaning	Leadership interpretation	Next action
No	Not defined, not owned, or not evidenced.	Material uncertainty or an unmanaged exposure.	Do not expand until the gap is resolved or formally accepted.
Partial	Exists but is informal, inconsistent, or incomplete.	A limited pilot may be viable when guardrails are explicit.	Assign a remediation owner and due date before pilot expansion.
Yes	Defined, accountable, and supported by evidence.	A control can be relied on for controlled use.	Proceed, but retain monitoring and periodic review.

Critical guardrails

Do not move to production if there is no accountable owner, no approved data boundary, no defined authority for high-impact output, or no safe manual/rollback path. These are decision gates, not items to average away.

DOMAIN 1 OF 6

1. Strategic Scope & Accountable Ownership

Leadership rationale

AI projects often fail in the handoff from a plausible idea to an operating decision. This section ensures leaders can connect the use case to an outcome, specify where the output enters the workflow, and assign authority for benefit, risk, and change.

Framework alignment: NIST AI RMF: GOVERN and MAP; ISO/IEC 42001: organizational context, leadership, and AI-system objectives.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
The business problem and intended outcome are defined. Describe the operational problem before selecting the tool or vendor.	☐	☐	☐	_____ _____
The expected value has a measurable baseline and target. Examples: reduced review time, lower error rate, improved follow-up, or increased service capacity.	☐	☐	☐	_____ _____
The AI output is mapped to a specific decision, task, or workflow step. Record who receives the output, what they can do with it, and what remains out of scope.	☐	☐	☐	_____ _____
An executive owner and an operational owner are assigned. Accountability must sit with named people, not the model, application, or vendor.	☐	☐	☐	_____ _____
The consequence of an incorrect, unavailable, or misused output is understood. Consider customer harm, financial loss, safety, privacy, legal exposure, and reputational impact.	☐	☐	☐	_____ _____

Leadership question: If this use case succeeds, what decision or operational result will improve enough to justify the residual risk and ongoing control cost?

DOMAIN 2 OF 6

2. Data Assets, Decision Rights & Lineage

Leadership rationale

AI inherits the ambiguity of the information it receives. Leaders need a clear account of authoritative sources, who may authorize use, and what data crosses system or vendor boundaries before AI is allowed to shape a decision or workflow.

Framework alignment: NIST CSF 2.0: GOVERN and IDENTIFY; NIST AI RMF: MAP; ISACA data-governance principles on decision rights and accountability.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
Required data sources and data flows are identified. Include systems, databases, shared folders, spreadsheets, APIs, reports, and third parties.	☐	☐	☐	_____
Each critical data source has a business owner and an approved use purpose. The owner is accountable for meaning, access approval, quality expectations, and change notification.	☐	☐	☐	_____
The authoritative source of truth is known for key records. Do not build AI workflows on duplicates or extracts without knowing which record is authoritative.	☐	☐	☐	_____
Sensitive, regulated, and restricted information is located and labeled. Include personal, health, financial, legal, employee, customer, and confidential business data.	☐	☐	☐	_____
Third-party data access, outbound data transfer, and retention paths are documented. Know what leaves the organization, where it goes, and who can access or retain it.	☐	☐	☐	_____

Leadership question: Can we explain, without guessing, which data the system will use, who owns it, and where it travels?

DOMAIN 3 OF 6

3. Data Quality, Context & Fitness for Use

Leadership rationale

A technically connected source can still be unfit for the decision at hand. This section tests whether information is current, complete, consistently defined, and accompanied by known limitations so the AI output is not trusted beyond the evidence that supports it.

Framework alignment: NIST AI RMF: MAP and MEASURE; NIST AI RMF Generative AI Profile (NIST AI 600-1); ISO/IEC 42001 performance evaluation concepts.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
Data is current enough for the intended operating cadence. Some work tolerates monthly data; fraud, safety, security, or customer-support work may not.	☐	☐	☐	_____ _____ _____ _____
Data is sufficiently complete and representative for the decision. Missing, biased, or unevenly captured records can make an output appear more reliable than it is.	☐	☐	☐	_____ _____ _____ _____
Key terms, labels, formats, and business definitions are consistent. Align naming, units, date formats, statuses, and category definitions across the scope.	☐	☐	☐	_____ _____ _____ _____
Known data limitations are documented and visible to reviewers. Record gaps, stale sources, excluded populations, unresolved duplicates, and confidence limits.	☐	☐	☐	_____ _____ _____ _____
There is a process to correct inaccurate source data and prevent recurrence. Name who corrects the record, how the correction propagates, and how recurrent issues are handled.	☐	☐	☐	_____ _____ _____ _____

Leadership question: What decision would become unsafe or materially misleading if the underlying data were stale, incomplete, or misunderstood?

DOMAIN 4 OF 6

4. Data Protection, Privacy & Vendor Boundaries

Leadership rationale

AI systems introduce new routes for data access, transformation, storage, and possible secondary use. This section puts security, privacy, and contract controls around those routes before sensitive or business-critical information is exposed.

Framework alignment: NIST CSF 2.0: GOVERN, PROTECT, and DETECT; NIST AI 600-1; ISO/IEC 42001 governance and data-management controls.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
Approved and prohibited data classes for this use case are defined. Not every document, prompt, email, file, or record is appropriate for AI use.	☐	☐	☐	_____ _____
Access follows least privilege and is tied to a business need. Review identities, roles, shared accounts, service accounts, exports, and administrative access.	☐	☐	☐	_____ _____
Vendor terms address data use, retention, training, subprocessors, and deletion. Confirm whether the provider may retain or use organizational data beyond the agreed service.	☐	☐	☐	_____ _____
Privacy, legal, and regulatory requirements are evaluated for the actual use case. Escalate handling of personal data, health data, financial data, employment data, and regulated decisions.	☐	☐	☐	_____ _____
Meaningful access and activity logging exists for important data and AI interactions. The organization should be able to reconstruct material access, changes, prompts, and outputs where appropriate.	☐	☐	☐	_____ _____

Leadership question: Would we be comfortable explaining to a customer, regulator, or board exactly what data is shared with the AI provider and under what protections?

DOMAIN 5 OF 6

5. Governance, Human Authority & Accountability

Leadership rationale

The technical system cannot hold organizational authority. Leadership must define who approves use, who may rely on output, who can override it, and how people report or escalate unsafe, harmful, or incorrect behavior.

Framework alignment: NIST AI RMF: GOVERN and MANAGE; NIST CSF 2.0: GOVERN; ISACA guidance on accountability across the AI lifecycle; ISO/IEC 42001 management-system approach.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
The use case has a risk tier and documented approval path. Define who may approve, pause, expand, restrict, or retire the workflow.	☐	☐	☐	_____ _____
Human review, override authority, and decision rights are explicit. High-impact outputs require accountable human authority, not a nominal checkbox.	☐	☐	☐	_____ _____
Employees know which AI tools are approved and which are prohibited. Reduce shadow AI by making the permitted path usable and clear.	☐	☐	☐	_____ _____
Employees know what data must never be entered or uploaded. Policy must translate into practical behavior for ordinary work, not only legal language.	☐	☐	☐	_____ _____
There is an incident, correction, and escalation process for harmful or unexpected output. Treat failures as operational events with ownership, evidence, response, and lessons learned.	☐	☐	☐	_____ _____

Leadership question: When the system is wrong, unsafe, or challenged, who has the authority to stop it and who owns the outcome?

DOMAIN 6 OF 6

6. Validation, Operations & Resilience

Leadership rationale

Readiness is not proven at launch. Performance can change as data, users, vendors, threats, and business conditions change. This section confirms the organization can test before release, observe behavior in use, fall back safely, and improve controls from evidence.

Framework alignment: NIST AI RMF: MEASURE and MANAGE; NIST CSF 2.0: DETECT, RESPOND, and RECOVER; ISO/IEC 42001 monitoring and continual improvement.

Readiness check	No	Partial	Yes	Evidence, owner, or decision record
Success, quality, and safety metrics are defined before use. Measure outcome quality, error rate, time saved, user acceptance, customer impact, and cost where relevant.	☐	☐	☐	_____ ____ _____ ___
The workflow is tested with normal, edge-case, and misuse scenarios. Use representative test data and evaluate failure modes before customer-facing or consequential use.	☐	☐	☐	_____ ____ _____ ___
Monitoring covers output quality, access, data changes, and material incidents. Decide what signals trigger review, who receives them, and how evidence is retained.	☐	☐	☐	_____ ____ _____ ___
A documented manual process, rollback, or safe-stop path exists. The organization must be able to operate safely when the tool, data connection, or vendor fails.	☐	☐	☐	_____ ____ _____ ___
A review cadence is defined for drift, model/vendor changes, control effectiveness, and lessons learned. Tie review frequency to the use case's impact, pace of change, and residual risk.	☐	☐	☐	_____ ____ _____ ___

Leadership question: If this AI workflow fails tomorrow, can the business detect it, contain the impact, and continue the underlying service safely?

DECISION & ACTION

Executive Readiness Decision

Do not average away a critical gap. Use the following decision logic after the six domains are completed and evidence has been reviewed.

Decision	Minimum condition	Leadership action
Pause	A critical guardrail is No, or risk cannot be understood.	Do not deploy. Define the purpose, owner, approved data boundary, and minimum controls first.
Controlled pilot	Critical guardrails are met; remaining items are Partial with named owners.	Limit users, data, duration, and authority. Add human review, logging, tests, and a remediation date.
Proceed with controls	Controls are Yes with evidence and residual risks are accepted.	Authorize controlled operation; monitor quality, access, incidents, and business impact on a defined cadence.
Executive review required	Sensitive/regulated data, high-impact decisions, public-facing automation, or material vendor uncertainty.	Obtain documented executive, privacy, legal, and security review appropriate to the use case.

Required decision record

Decision selected	_____
Residual risks accepted by	_____
Pilot / deployment boundary	_____
Next review date	_____

DECISION & ACTION

30-Day Action Plan

Prioritize actions that remove decision gates, reduce material exposure, or create the evidence leadership needs to make a controlled deployment decision. Keep the plan narrow: a small number of accountable actions is stronger than a long unowned list.

Priority	30-day action	Accountable owner	Due date
1			
2			
3			
4			
5			

DECISION & ACTION

Escalation Criteria & Executive Sign-Off

If any condition below applies, pause expansion until the appropriate executive, privacy, legal, and security authorities have reviewed the use case. Record the decision and any accepted residual risk in the use-case file.

Escalate for executive review when

- The use case handles sensitive or regulated data without a documented data boundary or vendor safeguard.
- AI output can materially affect customers, employees, pricing, legal position, health, financial outcomes, or safety without accountable human authority.
- The organization cannot explain the source data, decision logic, operational fallback, or evidence needed to investigate a material failure.

Leadership sign-off

Executive sponsor: _____ Date: _____

Operational owner: _____ Next review: _____

Decision notes

_____

_____

_____

_____

_____

EVIDENCE BASE

Framework References & Implementation Notes

This tool is a practical scoping aid, not a substitute for legal, privacy, regulatory, clinical, or cybersecurity advice. It does not certify compliance. Its framework references support a common language for decisions, evidence, and control design.

How the references are used

NIST AI RMF provides the core leadership logic: govern the system, map the context and impacts, measure risks and performance, and manage the response. NIST CSF 2.0 provides complementary cybersecurity outcomes. ISO/IEC 42001 and ISACA resources reinforce management accountability and lifecycle governance.

[1] NIST AI RMF 1.0. Artificial Intelligence Risk Management Framework (NIST AI 100-1, 2023): Govern, Map, Measure, and Manage. Open official source

[2] NIST Generative AI Profile. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1, 2024). Open official source

[3] NIST CSF 2.0. The Cybersecurity Framework 2.0 (NIST CSWP 29, 2024): Govern, Identify, Protect, Detect, Respond, and Recover. Open official source

[4] ISO/IEC 42001:2023. Information technology — Artificial intelligence — Management system. Official standard overview. Open official source

[5] ISACA. Governing AI Across Its Lifecycle: A Framework for Risk Practitioners (2026). Open official source

Reference check date: August 2026. Frameworks and implementation guidance evolve; organizations should validate current obligations for their jurisdiction, sector, and use case.

Next step

Use the completed AI Quick Scope to identify the few gaps that materially constrain safe, valuable adoption. PWNSENTINEL can facilitate an executive review, map critical data assets, evaluate governance and implementation controls, and define a practical readiness roadmap.

PWNSENTINEL | Govern AI with confidence from executive strategy to production deployment.