

AI Practical Adoption Guidance and Best Practices Paper

How to move from “everything is an AI use case” to governed, measurable, and scalable AI-enabled work

Prepared: July 2026

Core message	
AI adoption succeeds when an organization treats AI as a disciplined change in how work is performed, measured, governed, and improved. Almost any process can be described as an AI use case, but only some use cases deserve funding, technical build-out, policy approval, and operational ownership. The practical question is not “Can AI touch this process?” The practical question is “Will AI improve a measurable outcome without creating unacceptable risk, cost, complexity, or control weakness?”	
Adoption principle	Practical meaning
Outcome first	Begin with the mission, customer, control, risk, cycle-time, cost, quality, or decision outcome to improve.
Workflow redesign	Do not simply place an AI chatbot on top of broken processes; redesign roles, evidence, handoffs, approvals, and feedback loops.
Data readiness	Treat trusted data, metadata, lineage, access control, and document quality as adoption prerequisites, not technical afterthoughts.
Human accountability	AI can recommend, draft, triage, classify, and act within limits; accountable humans own policy, judgment, exceptions, and final decisions.
Controlled scale	Move from sandbox to pilot to production only when risk controls, monitoring, audit logs, training, and value metrics are in place.

Table of Contents

- 1. Executive Summary
- 2. The Strategic Problem: Everything Can Look Like an AI Use Case
- 3. What Makes an AI Use Case Worth Doing
- 4. AI Adoption Maturity Model
- 5. Practical AI Adoption Framework
- 6. Best Practices by Adoption Stage
- 7. Data Readiness Guidance
- 8. Resource, Workforce, and Operating Model Guidance
- 9. Technical Environment and Architecture Best Practices
- 10. Process Redesign and Human-in-the-Loop Controls
- 11. Security, Privacy, Legal, and Responsible AI Controls
- 12. Governance, Portfolio Management, and Value Realization
- 13. Recommended Implementation Roadmap
- 14. Practical Checklists and Templates
- 15. Conclusion
- Appendix A. AI Use-Case Scoring Model
- Appendix B. Common AI Use-Case Patterns
- Appendix C. References

1. Executive Summary

Artificial intelligence is now broad enough that nearly any current business process, event, action, transaction, document, meeting, decision, system log, communication, approval, control check, research activity, or customer interaction can be framed as an AI use case. That observation is useful, but it is also dangerous. If every activity is called an AI use case, organizations risk creating long catalogs with weak prioritization, fragmented pilots, unclear ownership, uncontrolled tools, and little measurable value.

The right adoption strategy is to define a practical AI operating model that selects use cases based on business value, data readiness, control fit, workforce capacity, security posture, and repeatable measurement. AI should be used to improve work, not merely to rebrand work. A successful AI use case should make a process faster, more accurate, more traceable, more scalable, more compliant, more insightful, or more responsive. The use case should also be monitorable, explainable enough for its risk level, and owned by a business leader who can change the workflow around it.

Executive recommendation	
Build an AI adoption program around a governed portfolio of high-value workflows, not around isolated tools. Establish an intake and scoring process; create a secure experimentation environment; prioritize use cases with available data and measurable outcomes; require human accountability for consequential decisions; implement security, privacy, and model-risk controls; and scale only after operational, technical, and value evidence is proven.	
Leadership question	Practical answer
Can any business process become an AI use case?	Yes, if it has inputs, work steps, outputs, decisions, or feedback that AI can assist, automate, predict, generate, classify, extract, summarize, reconcile, detect, recommend, or orchestrate.
Should every process become an AI project?	No. Many problems are better solved through process standardization, data cleanup, system configuration, workflow automation, training, or policy clarification.
What separates successful AI adoption from failed experimentation?	Successful adoption has a measurable business outcome, trusted data, workflow redesign, risk controls, human accountability, secure architecture, and sustained ownership.
Where should organizations start?	Start with high-volume, repeatable, information-heavy workflows where AI can reduce cycle time, improve consistency, surface risk, or create better decision support without replacing accountable judgment.

2. The Strategic Problem: Everything Can Look Like an AI Use Case

AI can be mapped onto almost every business activity because most work follows a common pattern: receive information, interpret it, transform it, decide something, communicate the result, and update a record. AI can assist at nearly every point in that pattern. It can read documents, summarize conversations, classify transactions, detect anomalies, draft responses, recommend actions, generate code, compare records, predict outcomes, explain trends, and orchestrate tasks across systems.

This creates a strategic challenge: the number of possible AI use cases can become unlimited, while leadership attention, data-engineering capacity, cybersecurity review, change-management capacity, and funding are limited. A mature organization therefore needs an adoption filter. The goal is to convert unlimited possibility into a prioritized sequence of fundable, controllable, measurable, and scalable AI capabilities.

Activity type	Why it can become an AI use case	Example
Document	AI can extract, summarize, compare, classify, and validate content.	Review policies, invoices, contracts, audit evidence, resumes, claims, case files, or reports.
Transaction	AI can detect outliers, reconcile records, forecast exceptions, and score risk.	Identify duplicate payments, unusual journal entries, suspicious claims, or late approvals.
Meeting or conversation	AI can transcribe, summarize, identify decisions, and create follow-up actions.	Convert interviews, status meetings, or customer calls into structured action logs.
Decision	AI can provide ranked options, evidence summaries, confidence indicators, and risk flags.	Support loan review, audit planning, vendor risk assessment, budget trade-off analysis, or case prioritization.
System event	AI can monitor logs, detect unusual patterns, and trigger alerts.	Identify privilege abuse, failed controls, access anomalies, or operational incidents.
Workflow	AI can coordinate tasks, check completion, route exceptions, and prepare draft outputs.	Guide procure-to-pay, onboarding, audit request management, issue remediation, or grant processing.

Adoption risk
The danger is not that organizations will lack AI ideas. The danger is that they will produce too many loosely defined ideas and too few production-ready capabilities. A practical AI adoption program should require a clear owner, outcome metric,

data source, risk tier, human-control design, production pathway, and funding logic before a use case enters the build pipeline.

3. What Makes an AI Use Case Worth Doing

A strong AI use case has a clear business problem, a repeatable workflow, available data or documents, measurable value, manageable risk, and a defined path into operations. A weak AI use case is a technology idea searching for a problem, lacks owner commitment, has poor data access, cannot be measured, or creates risk that exceeds its value.

Criterion	Strong use case	Weak use case
Business outcome	Improves a named outcome such as cycle time, accuracy, risk detection, cost, compliance, customer response, or decision quality.	Described as “use AI for this area” without a measurable target.
Workflow fit	Targets a repeatable workflow with clear inputs, outputs, exceptions, and accountability.	Targets vague knowledge work with no defined handoffs or decision rights.
Data readiness	Uses accessible, permitted, relevant, and sufficiently reliable data/documents.	Depends on unknown, inaccessible, sensitive, incomplete, or low-quality data.
Risk profile	Risk is identified, tiered, controlled, and proportionate to expected benefit.	Touches rights, safety, money, legal decisions, or sensitive data without control design.
Human role	Defines what AI does, what humans review, what requires approval, and how exceptions escalate.	Assumes the model will independently make decisions without oversight.
Production path	Has sponsor, system owner, security path, support model, and monitoring plan.	Remains a demo, proof-of-concept, or isolated spreadsheet workflow.

Organizations should score use cases against value, feasibility, risk, data readiness, and scalability. A high-value but low-data-readiness use case may still be important, but it should be treated as a data-modernization initiative first. A low-risk drafting assistant may be useful but should not consume the same governance bandwidth as an AI system that recommends eligibility, payment, credit, audit, hiring, safety, or enforcement actions.

4. AI Adoption Maturity Model

Level	Description	Typical behavior	Main risk	Next step
0. Unaware / uncontrolled	AI use is informal or unknown.	Employees use public tools without policy, logging, or data controls.	Data leakage, inconsistent outputs, shadow AI.	Create policy, inventory, and safe-use guidance.
1. Experimenting	Pilots and demos occur in isolated pockets.	Teams test chatbots, document summaries, coding tools, or analytics prototypes.	Pilot sprawl and no path to production.	Create intake, scoring, and secure sandbox.
2. Managed pilots	Use cases are selected, risk-tiered, and evaluated.	Pilots have owners, success metrics, data access, and review checkpoints.	Measurement gaps and weak transition planning.	Build production architecture and operating model.
3. Production integration	AI is embedded into approved workflows and systems.	Users receive AI assistance inside governed applications with logging and monitoring.	Operational dependency without monitoring.	Establish ModelOps/LLMOps, controls, and support.
4. Portfolio scale	AI is managed as an enterprise capability portfolio.	Reusable platforms, approved models, common patterns, and shared services reduce duplication.	Concentration risk and governance bottlenecks.	Optimize performance, cost, and reuse.
5. Adaptive enterprise	AI continuously improves workflows and decisions under governance.	AI-enabled processes feed metrics, lessons learned, and control improvements back into operations.	Over-automation and mission drift.	Continuously validate, audit, retrain, retire, or redesign.

Practical interpretation: maturity is not measured by number of models deployed. It is measured by the organization’s ability to repeatedly convert appropriate use cases into controlled, measurable, adopted, and maintainable capabilities.

5. Practical AI Adoption Framework

A practical AI adoption framework should cover eight dimensions: value, process, data, technology, security, governance, people, and measurement. Each dimension must be addressed before a use case moves from idea to production.

Dimension	Core question	Best practice
Value	What outcome will improve?	Define one primary outcome metric and two to four supporting metrics before building.
Process	How will work change?	Map current state and future state, including human review, exception handling, and decision authority.
Data	What information will AI use?	Document data sources, quality, lineage, permissions, sensitivity, retention, and refresh cycle.

Technology	Where will AI run?	Use approved platforms, secure model gateways, dev/test/prod separation, logging, monitoring, and rollback.
Security	What can go wrong?	Control prompt injection, data leakage, excessive agency, insecure outputs, identity access, and supply-chain risks.
Governance	Who approves and owns it?	Assign business owner, technical owner, data steward, risk reviewer, and operational support owner.
People	Who must change behavior?	Train users, supervisors, reviewers, and support teams on proper use, limits, and escalation.
Measurement	How will success be proven?	Measure baseline, pilot results, production performance, user adoption, quality, risk, and cost.

Minimum viable AI governance
 Every approved AI use case should have: a business owner, purpose statement, risk tier, data inventory, model/tool selection, user group, human-review rule, security review, privacy review when applicable, testing plan, monitoring metric, audit log, value metric, and retirement or reevaluation trigger.

6. Best Practices by Adoption Stage

- Intake:** Capture ideas using a standard use-case form. Require problem statement, target users, source data, current pain point, desired outcome, risk sensitivity, and expected value.
- Triage:** Score ideas using value, feasibility, data readiness, risk, urgency, and scalability. Reject or defer ideas that are really policy, data-quality, workflow, or system-configuration problems.
- Discovery:** Map current process, decisions, documents, systems, stakeholders, errors, cycle times, costs, and control points. Identify whether the AI pattern should be search, summarization, extraction, classification, prediction, anomaly detection, generation, agentic orchestration, or workflow automation.
- Feasibility:** Validate data access, model options, security boundary, integration path, cost estimate, legal/privacy constraints, and human-review requirements. Establish baseline metrics before prototype work begins.
- Prototype:** Build in a secure sandbox with representative data. Test model behavior, failure modes, prompt injection, retrieval quality, hallucinations, accuracy, bias, latency, and cost.
- Pilot:** Use a controlled user group, defined scope, human review, feedback capture, and explicit go/no-go criteria. Compare pilot performance against baseline.
- Production:** Deploy through approved architecture with role-based access, logs, monitoring, incident response, support model, documentation, training, and change control.
- Scale:** Reuse common components, standard prompts, approved connectors, shared RAG pipelines, evaluation datasets, dashboards, and governance templates. Do not rebuild the same pattern repeatedly.
- Monitor and improve:** Track quality, adoption, risk events, drift, cost, false positives, false negatives, user overrides, complaints, exceptions, and realized value. Update or retire the use case when it no longer performs.

7. Data Readiness Guidance

Data readiness is often the difference between AI theater and operational AI. Many AI ideas fail not because the model is weak, but because the data is scattered, inaccessible, inconsistent, low-quality, poorly labeled, undocumented, or not legally permitted for the intended purpose. Data work is not a pre-AI burden; it is part of the AI product.

Data requirement	Guidance	Why it matters
Source inventory	List systems, documents, databases, reports, spreadsheets, APIs, and external sources used by the use case.	Prevents hidden dependencies and supports auditability.
Data quality	Assess completeness, accuracy, timeliness, duplicates, missing fields, inconsistent formats, and known defects.	Model outputs cannot be trusted if inputs are unreliable.
Lineage and provenance	Track where data came from, how it changed, and when it was refreshed.	Supports explainability, troubleshooting, and compliance.
Access and permissions	Apply least privilege and role-based access; prevent the AI layer from exposing data users could not otherwise see.	Reduces insider risk and unauthorized disclosure.
Sensitivity classification	Identify PII, PHI, financial data, controlled unclassified information, classified data, trade secrets, or privileged information.	Determines permitted tools, storage, logging, and sharing.
Ground truth / evaluation data	Create test sets with expected answers, labels,	Allows measurable evaluation instead of

	or expert judgments.	anecdotal demo success.
Metadata and context	Capture definitions, business rules, document type, effective date, source system, owner, and confidence.	Helps AI retrieve and reason over the right context.
Retention and deletion	Define what prompts, outputs, embeddings, logs, and training/evaluation data are stored and for how long.	Supports privacy, records management, and cost control.

RAG versus fine-tuning
 For most enterprise knowledge, policy, document, and reference-answering use cases, start with retrieval-augmented generation rather than model fine-tuning. RAG allows the organization to ground answers in approved content, update knowledge without retraining, control citations, and manage access. Fine-tuning is more appropriate when the goal is to teach a model a stable pattern of style, classification, extraction, or task behavior using high-quality labeled examples.

8. Resource, Workforce, and Operating Model Guidance

AI adoption requires an interdisciplinary operating model. A business team alone may understand the process but lack security, data, and model expertise. A technical team alone may build impressive prototypes that fail in operations. Practical adoption requires a cross-functional team with clear roles.

Role	Primary responsibility
Executive sponsor	Sets priority, resolves barriers, funds the initiative, and holds owners accountable for outcomes.
Business product owner	Owns the process outcome, workflow design, user adoption, and value realization.
Subject matter experts	Define business rules, edge cases, labels, acceptance criteria, and human-review requirements.
Data owner / steward	Approves data use, documents lineage, assesses quality, and manages access.
AI / ML engineer	Builds model, retrieval, evaluation, automation, and integration components.
Application / platform engineer	Connects AI into approved systems, identity, logging, API, and deployment pipelines.
Cybersecurity and privacy reviewers	Assess security boundary, data exposure, adversarial risk, logging, retention, and privacy impact.
Legal, compliance, or policy advisor	Reviews legal authority, obligations, records, intellectual property, procurement, and policy constraints.
Risk / internal audit partner	Reviews control design, evidence, auditability, segregation of duties, and monitoring.
Change-management and training lead	Prepares users, supervisors, job aids, communications, support, and adoption metrics.

Organizations should also establish an AI center of enablement or community of practice. The goal is not to centralize every decision, but to provide reusable patterns, approved tools, evaluation methods, policies, coaching, and shared infrastructure so teams do not repeatedly solve the same problems.

9. Technical Environment and Architecture Best Practices

Practical AI adoption needs a secure technical environment that supports experimentation without creating uncontrolled risk and supports production without making every deployment custom. The architecture should separate sandbox, development, test, and production; enforce identity and access controls; log inputs and outputs according to policy; and provide monitoring for quality, cost, latency, and risk events.

Architecture capability	Best practice
Approved model access	Use a model gateway or approved AI platform to manage models, API keys, rate limits, cost, logging, and provider changes.
Secure sandbox	Give teams a controlled environment for experimentation using approved data or de-identified samples.
RAG and knowledge layer	Use indexed, permission-aware repositories with metadata, chunking, citations, refresh schedules, and retrieval evaluation.
Vector database / search	Store embeddings and retrieval metadata with access controls, versioning, and deletion processes.
Workflow integration	Integrate AI where work already happens: case systems, ticketing, ERP, CRM, document management, BI, collaboration tools, or portals.
Observability	Track prompts, retrieved context, outputs, model version, latency, token cost, user action, override, error, and exception.
Evaluation pipeline	Maintain test prompts, expected answers, labeled examples, red-team cases, and regression tests before release.
Guardrails	Use input validation, output validation, content filters, tool restrictions,

	policy checks, human approvals, and safe-fail behaviors.
Deployment controls	Use change management, code review, security scanning, dependency review, rollback, and production approval gates.

Agentic AI caution
 Agentic AI introduces higher risk because the system can plan, call tools, retrieve data, write outputs, trigger workflows, and sometimes take actions. Treat each agent as a governed digital worker with identity, least privilege, action boundaries, approval gates, audit logs, test cases, and kill-switch procedures.

10. Process Redesign and Human-in-the-Loop Controls

AI is most valuable when the organization redesigns how work is performed, not when it simply asks employees to use a chatbot in parallel with existing work. Practical adoption should map current-state friction and future-state work design. The future state should specify who initiates the AI task, what data is used, what output is produced, who reviews it, what is automatically accepted, what is escalated, and what evidence is stored.

AI role in workflow	Appropriate human control
Drafting and summarization	Human reviews final content before external, official, legal, financial, or policy use.
Classification and routing	Human review for low-confidence, sensitive, novel, or high-impact cases.
Extraction and reconciliation	Human review for exceptions, mismatches, low confidence, and material items.
Prediction and risk scoring	Human decision-maker sees score, explanation, data basis, confidence, and limitations.
Recommendation	Human approves consequential actions and documents rationale when overriding or accepting.
Autonomous action	Allow only for low-risk bounded tasks or with approval gates, rollback, and audit trails.

Human-in-the-loop should not mean vague manual review after the fact. It should be designed as a control: who reviews, when they review, what evidence they see, what standard they apply, what actions they may take, and how the review is logged. For high-risk use cases, the review process should include sampling, second-level review, outcome monitoring, and periodic independent assessment.

11. Security, Privacy, Legal, and Responsible AI Controls

Responsible AI adoption requires controls that match the risk of the use case. A grammar assistant does not require the same control depth as a credit, benefits, enforcement, audit, safety, medical, personnel, payment, cybersecurity, or autonomous-action system. The control model should be risk-tiered, documented, and integrated into the normal system-development lifecycle.

Risk area	Practical control
Prompt injection	Treat external content as untrusted; separate instructions from data; restrict tools; test malicious documents and prompts.
Data leakage	Use approved environments, data-loss prevention, access controls, redaction, retention rules, and no public upload of sensitive data.
Hallucination / confabulation	Use RAG grounding, citations, confidence scoring, refusal behavior, output validation, and human review.
Bias and unfairness	Test for disparate impact, proxy variables, data representativeness, and outcome differences; document mitigation.
Excessive agency	Limit what the model can do; require approvals for external communication, payments, account changes, code deployment, or irreversible actions.
Supply chain	Review model providers, libraries, plugins, connectors, training data claims, data-use terms, and vendor security.
Privacy	Assess PII processing, notice, consent or legal authority, minimization, retention, and purpose limitation.
Intellectual property	Control copyrighted, proprietary, privileged, or confidential input/output handling.
Records and auditability	Preserve prompts, outputs, retrieved sources, approvals, versions, and decisions when required.
Operational resilience	Plan fallback process, outage response, rollback, incident reporting, and continuity procedures.

The control language in this paper is consistent with the general direction of NIST AI RMF, NIST GenAI Profile, OMB federal AI governance requirements, ISO/IEC 42001 management-system concepts, CISA AI security priorities, and OWASP LLM application risk guidance.

12. Governance, Portfolio Management, and Value Realization

AI governance should enable responsible innovation, not stop it. The best governance model gives teams a fast path for low-risk use cases and a deeper review path for high-impact use cases. Governance should be integrated into funding, architecture, cybersecurity, privacy, data management, procurement, and operational performance review.

Governance layer	Responsibilities
AI policy and standards	Define approved and prohibited uses, data rules, human-review expectations, risk tiers, and documentation standards.
AI use-case inventory	Track all AI systems, pilots, owners, data sources, risk tiers, deployment status, and value metrics.
Portfolio board	Prioritize funding, resolve duplication, approve scaling, and retire low-value initiatives.
Risk review	Evaluate safety, rights, privacy, security, legal, model, and operational risks before production.
Architecture review	Ensure solutions use approved patterns, environments, identity, logging, and integration approaches.
Performance review	Monitor value realization, quality, adoption, incidents, cost, and user feedback.
Independent assurance	Periodically assess whether AI controls operate as designed and whether outcomes remain acceptable.

Value realization should be explicit. Each use case should have a baseline and target. Example metrics include hours saved, case backlog reduction, cycle-time reduction, error reduction, false-positive reduction, improved detection rate, avoided cost, improved customer response time, reduced rework, audit exception reduction, compliance timeliness, employee productivity, or revenue growth. In many cases, qualitative outcomes such as improved decision confidence or knowledge retention should be paired with quantitative measures.

13. Recommended Implementation Roadmap

Phase	Timeframe	Key actions	Expected outputs
Phase 0: Mobilize	0-30 days	Name executive sponsor, AI lead, governance team, initial policy, and inventory approach.	AI adoption charter, safe-use guidance, draft intake form, initial risk tiers.
Phase 1: Discover and prioritize	30-90 days	Collect use cases, score portfolio, map data sources, identify quick wins and strategic workflows.	Prioritized portfolio, baseline metrics, data-readiness assessment, pilot candidates.
Phase 2: Build safe foundation	90-180 days	Stand up secure sandbox, approved tools, RAG pattern, evaluation process, logging, and training.	Reusable platform, model gateway, test datasets, security patterns, user training.
Phase 3: Pilot and validate	6-12 months	Pilot selected use cases with defined users, metrics, review controls, and production criteria.	Pilot results, go/no-go decisions, refined workflows, risk-control evidence.
Phase 4: Production scale	12-24 months	Move successful pilots into production systems; establish support, monitoring, value dashboards, and improvement cycles.	Production AI capabilities, adoption metrics, monitored controls, cost/value dashboard.
Phase 5: Enterprise optimization	24+ months	Scale common patterns, automate governance evidence, improve data products, retire weak use cases, and expand higher-value workflows.	Enterprise AI operating model, reusable components, continuous improvement process.

Roadmap discipline
 Do not wait for perfect enterprise data to start. Begin with safe, bounded, measurable use cases while building the data, governance, and technical foundation needed for higher-risk and higher-value AI adoption. Parallel progress is essential: policy, platform, data, pilots, workforce, and measurement should advance together.

14. Practical Checklists and Templates

14.1 AI Use-Case Intake Checklist

- Problem statement: What specific pain point, risk, delay, error, cost, backlog, decision, or quality issue will AI improve?
- Target users: Who will use the AI output, and how will their work change?
- Current process: What are the current inputs, systems, handoffs, decisions, approvals, outputs, and pain points?
- AI task: Is the AI expected to search, summarize, extract, classify, predict, detect, generate, recommend, reconcile, or act?
- Data sources: What systems, documents, repositories, reports, or external data will be used?
- Risk tier: Could the use case affect rights, safety, money, legal status, employment, privacy, security, compliance, or mission operations?
- Human control: Who reviews the AI output, what do they approve, and what exceptions are escalated?

- Value metric: What baseline will be measured and what improvement target is expected?
- Production path: What system, owner, support process, and funding path will sustain the capability?

14.2 Production Readiness Checklist

- Approved business owner and technical owner are documented.
- Risk tier, impact assessment, and required reviews are complete.
- Data sources, lineage, access rules, sensitivity, and retention are documented.
- Model, prompts, retrieval sources, tools, connectors, and architecture are documented.
- Evaluation results meet predefined acceptance criteria.
- Human-review, exception, escalation, and override processes are implemented.
- Logs, monitoring, cost tracking, incident response, and rollback procedures are in place.
- Users are trained on proper use, limitations, prohibited inputs, and escalation.
- Value dashboard and post-deployment review schedule are established.

15. Conclusion

AI adoption is no longer a question of whether organizations can identify use cases. They can identify too many. The harder leadership task is deciding which use cases deserve investment, how they should be governed, what data foundation they require, how processes must change, and how value will be measured after deployment.

A practical AI adoption program should therefore be selective, evidence-driven, and operational. It should begin with real business problems, not model enthusiasm. It should treat data quality, security, privacy, process redesign, workforce adoption, and governance as core delivery work. It should start small enough to control but build reusable patterns that allow enterprise scaling. It should make humans more effective, not less accountable.

The most successful organizations will not be the ones with the longest AI use-case lists. They will be the ones that convert the right use cases into trusted, repeatable, measurable, and improved ways of working. AI should be judged by better outcomes: faster service, stronger controls, reduced rework, improved insight, better compliance, higher productivity, and greater decision confidence. Used this way, AI becomes not a collection of experiments, but a disciplined operating capability.

Appendix A. AI Use-Case Scoring Model

Score area	1 - Low	3 - Medium	5 - High
Business value	Nice-to-have convenience.	Improves a team-level metric.	Improves strategic outcome, risk, cost, mission, compliance, or customer experience.
Workflow repeatability	Ad hoc work with many unique cases.	Moderately repeatable with common patterns.	High-volume repeatable workflow with clear inputs and outputs.
Data readiness	Data inaccessible, unknown, or low quality.	Data exists but needs cleanup or permissions.	Data is accessible, documented, permitted, and reliable.
Technical feasibility	Requires unproven integration or unavailable capability.	Feasible with moderate engineering.	Feasible using approved tools, patterns, and platforms.
Risk manageability	High impact with unclear controls.	Moderate impact with controls needed.	Low to moderate risk with clear controls and human review.
Scalability	One-off local benefit.	Can expand within one function.	Reusable pattern across functions or enterprise workflows.

Suggested formula: Total score = Business value + Workflow repeatability + Data readiness + Technical feasibility + Risk manageability + Scalability. Require separate escalation for any high-impact safety, rights, privacy, legal, financial, cybersecurity, or autonomous-action use case regardless of numeric score.

Appendix B. Common AI Use-Case Patterns

Pattern	Description
Knowledge assistant	Answers questions from approved policies, procedures, contracts, manuals, research, or records.
Document summarizer	Summarizes long documents, emails, meeting transcripts, case files, reports, or legal/regulatory materials.
Data extraction	Extracts fields from invoices, forms, claims, contracts, financial statements, applications, or evidence packages.
Classification and routing	Classifies cases, documents, requests, tickets, emails, or transactions and routes them to the right owner.
Anomaly detection	Finds unusual transactions, events, patterns, users, approvals, vendors, claims, or system behavior.
Reconciliation and matching	Matches records across systems and identifies differences, duplicates, or

	missing evidence.
Predictive scoring	Scores probability of delay, fraud, default, churn, defect, compliance failure, backlog risk, or audit issue.
Decision support	Presents evidence, options, risk flags, and recommended actions while preserving human decision authority.
Draft generation	Drafts reports, emails, memos, code, test scripts, procedures, job aids, or communications for human approval.
Process agent	Coordinates multistep workflows, checks status, gathers evidence, prepares actions, and escalates exceptions.
Monitoring assistant	Continuously watches events, logs, news, regulations, system data, or performance metrics and alerts on changes.
Training and coaching	Provides personalized learning, scenario practice, knowledge checks, and job-specific guidance.

Appendix C. References

- National Institute of Standards and Technology (2023). NIST AI Risk Management Framework 1.0. <https://www.nist.gov/itl/ai-risk-management-framework>
- National Institute of Standards and Technology (2024). NIST AI RMF: Generative Artificial Intelligence Profile, NIST AI 600-1. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
- Office of Management and Budget (2024). OMB Memorandum M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of AI. <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>
- International Organization for Standardization (2023). ISO/IEC 42001:2023 Artificial intelligence - Management system. <https://www.iso.org/standard/42001>
- Open Worldwide Application Security Project (2025). OWASP Top 10 for Large Language Model Applications. <https://owasp.org/www-project-top-10-for-large-language-model-applications/>
- Cybersecurity and Infrastructure Security Agency (2023-2025). CISA Roadmap for Artificial Intelligence. <https://www.cisa.gov/resources-tools/resources/roadmap-ai>
- McKinsey & Company / QuantumBlack (2025). The State of AI: Global Survey 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- McKinsey & Company (2025). Superagency in the workplace: Empowering people to unlock AI full potential at work. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ai-full-potential-at-work>
- Stanford Institute for Human-Centered AI (2025). The 2025 AI Index Report. <https://hai.stanford.edu/ai-index/2025-ai-index-report>