هيئة الاتصالات والفضاء والتقنية
Communications, Space & Technology Commission

Awareness Guide for Artificial Intelligence Adoption in Technology Companies

Version 1

Jul 2026

cst.gov.sa

Executive Guide

This guide is issued by the Communications, Space and Technology Commission (CST) to support the practical and responsible adoption of Artificial Intelligence (AI) by technology companies operating in the Kingdom of Saudi Arabia, aligned with CST's mandate to regulate and enable the telecommunications and information technology sector. This guide is non-binding; companies remain responsible for compliance with applicable laws and regulations.

AI is no longer an experimental technology; it is a foundational capability reshaping productivity, service delivery, cost structures, product design, workforce allocation, and competitive positioning across digital markets. While experimentation is widespread, sustained and scalable value realization remains concentrated in organizations adopting structured operating models rather than isolated pilots.

For companies operating in the Kingdom, this shift should be assessed not only through global benchmarks, but also through local regulatory alignment, data governance expectations, sovereign and cross-border data considerations, Arabic-language performance in customer-facing use cases, sector-specific realities, and public trust expectations.

Against this backdrop, AI should be treated as a structural enterprise capability rather than a collection of tools. Adoption begins with organizational readiness across five interdependent dimensions: context, data, infrastructure, talent, and culture. This assessment determines which initiatives are appropriate, how aggressively to scale, and which foundational gaps should be addressed first, while AI literacy serves as a core enabler across leadership and operational roles.

Building on this foundation, AI application spans internal operations, customer-facing solutions, and AI agents. Internal use cases typically offer the fastest and lowest-risk path to value, while customer-facing AI enables stronger differentiation but requires higher

maturity due to visibility, trust, and regulatory sensitivity. AI agents represent a more advanced execution model and should begin with constrained authority, human oversight, and gradual expansion based on demonstrated reliability.

In parallel, adoption should be approached as a full-stack capability, where interdependent decisions materially affect cost, scalability, vendor concentration risk, and strategic flexibility. Companies should address intellectual property considerations, manage vendor lock-in through modular architecture and contractual clarity, and internalize inference economics. Sustainable advantage rarely comes from model access alone, but from proprietary context, embedded workflows, customer signals, and feedback loops.

On this basis, the guide introduces an AI-first operating principle in which strategy is defined before tool selection. This requires clarifying where AI creates measurable advantage, where proprietary data confers defensibility, what levels of autonomy are acceptable, and whether capability should be built, partnered, or acquired. Execution should then proceed through phased scaling: readiness validation, controlled pilots, enterprise integration, and continuous optimization.

As adoption progresses, governance is framed as a risk-based operating model that begins with disciplined use-case selection rather than post-deployment controls. Poorly governed AI systems can introduce legal, financial, and reputational risk. Governance should be proportionate to use-case risk, considering impact, autonomy, data sensitivity, and customer visibility. Responsible AI use requires clear problem definition, measurable objectives, compliance with applicable legal and regulatory frameworks, robust data governance, proportionate transparency, appropriate safeguards, and controlled experimentation with phased deployment.

Table of Content

Table of Figures

01 Introduction

This guide is issued by the Communications, Space and Technology Commission (CST) pursuant to its mandate to regulate and develop the telecommunications and information technology sector in the Kingdom of Saudi Arabia, in accordance with the Telecommunications and Information Technology Act and its bylaws, issued pursuant to Council of Ministers Resolution No. (592) dated 01/11/1443H and approved by Royal Decree No. (M/106) dated 02/11/1443H, as well as Article Seven of Cabinet Resolution No. (292) dated 27/04/1441H. Further to CST's amended Statute under Council of Ministers Resolution No. (133) dated 21/05/1424H, which entrusted CST with regulatory responsibilities for the information technology sector, and pursuant to its amended Statute under Council of Ministers Resolution No. (430) dated 23/06/1446H, which entrusted CST with promoting efforts that enhance the adoption of emerging digital technologies and digital transformation in the Kingdom, enabling and developing technology and innovation in sectors related to CST's mandate, maximizing their benefits to support the adoption of modern and innovative business models, and stimulating and developing emerging technologies, CST issues this guide for technology companies operating in the Kingdom, ranging from startups to large enterprises, within the scope of the Information Technology and Emerging Technologies Sector Classification issued by CST.

Within this mandate, and consistent with CST strategy to foster advanced and sustainable technology markets, artificial intelligence (AI) is recognized as a foundational capability with increasing impact on the internal operations of technology companies and on the solutions and applications they deliver to the market. Supporting effective and responsible adoption of AI is therefore integral to market development, competitiveness, and long-term value creation.

This guide is intended for technology companies operating in Saudi Arabia, from startups to large enterprises, within the scope of the Information Technology and Emerging Technologies (IT/ET) Sector Classification issued by CST. It addresses the practical adoption of AI across internal operations, customer-facing solutions and applications, and AI-enabled execution models, including AI agents, where relevant.

The guide is written primarily for executive leadership, strategy owners, and senior decision-makers within technology companies, with supporting relevance for technology leaders, product owners, and functional managers responsible for translating AI from strategic intent into operational reality. It is intended to support informed decision-making, structured implementation, and responsible scaling of AI across enterprise contexts.

This guide is non-binding in nature and does not replace or supersede any laws, regulations, rules, instructions, standards, or regulatory decisions issued by CST or other competent authorities. It should not be construed as a legal reference or as creating any legal obligations or rights. Each company remains fully responsible for interpreting, assessing, and applying the content of this guide in a manner consistent with its own context, priorities, and strategic direction.

02 The Strategic Case for AI Adoption

AI has seen rapid expansion in the scope of its use and applications, evolving from limited use cases into a widely adopted General-Purpose Technology (GPT) that is reshaping productivity, service delivery, and competitive positioning across digital markets. Like earlier GPTs such as electricity and the internet, AI is pervasive across sectors, continuously improving, and capable of enabling entirely new products, services, and business models rather than delivering isolated efficiency gains. International assessments estimate that AI will contribute between USD 13–15 trillion to global economic output by 2030, driven primarily by productivity gains, process optimization, and data-enabled innovation, with productivity improvements alone accounting for nearly 60 percent of the projected impact [UNDP, 2022; UNCTAD, 2023]. Recent enterprise surveys also show rapid growth in generative AI adoption, alongside a widening gap between experimentation and scaled value creation, reinforcing the need for structured operating models rather than pilot-driven adoption [McKinsey & Company, 2024]. This economic potential reflects AI's flexibility, as it can be applied across a wide range of functions and industries, from routine automation to complex decision-making, without requiring fundamental redesign for each use case. This pattern is consistent with broader evidence that most companies remain early in scaling AI reliably, even as capability and investment accelerate [Stanford University HAI, 2024].

In parallel, organizational adoption has expanded rapidly across sectors, but most companies remain early in translating experimentation into sustained and scalable operational impact. [UNESCO, 2023].

This disparity highlights that the value of AI as a GPT is not realized through experimentation alone, but through structured adoption aligned with organizational readiness and long-term strategic intent. Companies that develop early operational competence in AI benefit from cumulative learning effects, improved data utilization, and workforce capability development that are difficult for late adopters to replicate. Conversely, delayed or unstructured adoption increases exposure to competitive pressure, rising customer expectations, talent constraints, and cost disadvantages as AI-enabled capabilities become baseline market expectations. Accordingly, companies are advised to treat AI not merely as a component of digital transformation, but as a cross-cutting, foundational capability that reshapes how strategy, operations, products, and decision-making are designed and executed, requiring deliberate prioritization, sequencing, and governance at the enterprise level rather than isolated or tactical deployment.

Having established the strategic rationale for AI adoption, the next question is whether the organization is sufficiently prepared to adopt AI in a disciplined and sustainable manner.

03 Organizational Readiness

Effective adoption of AI depends on organizational readiness across five interdependent dimensions: context readiness, data readiness, infrastructure readiness, talent readiness, and organizational culture. Together, these dimensions determine whether a company is positioned to move beyond experimentation toward sustainable AI adoption. This guide treats readiness as a practical diagnostic rather than a theoretical concept, recognizing that weakness in any single dimension constrains overall outcomes regardless of strengths elsewhere, a principle consistently emphasized in international policy guidance [ITU, 2022; UNESCO, 2023].

Companies should assess readiness through a two-step approach:

First, a qualitative assessment:
using the five dimensions to identify structural imbalances, such as strong technical capability applied to an ill-defined problem, or a market where AI adds limited incremental value.

Second, a quantitative, indicative assessment:
using the Self-Assessment Checklist provided in the appendix.

This produces a readiness score per dimension and an overall readiness range, which serves as a decision aid rather than a benchmark or regulatory classification.

The output is a clear readiness profile, not a pass/fail judgment. That profile should directly inform:

- **Which categories of AI use cases are appropriate at the current stage,**
- **How initiatives should be sequenced, and**
- **What foundational gaps must be addressed before scaling adoption.**

The five organizational readiness dimensions answer distinct questions:

- **Context readiness:** whether the use case is justified within its market, customer, regulatory, and operational context. In the Saudi context, this also encompasses alignment with national data governance frameworks, sector regulator coordination where applicable, sensitivity considerations for AI in government-linked contracts, and Arabic NLP capability adequacy for customer-facing use cases.

- **Data readiness:** whether the required data exists, meets quality thresholds, and can be lawfully used.

- **Infrastructure readiness:** whether compute, integration, security, and hosting can support AI workloads.

• **Talent readiness:** whether the required technical capability, domain expertise, and AI literacy are available.

• **Cultural readiness:** whether leadership commitment, governance clarity, and change management are sufficient to support adoption.

As shown in Figure 3.1, organizational readiness is structured across five interdependent dimensions that together determine the feasibility and sequencing of AI adoption.

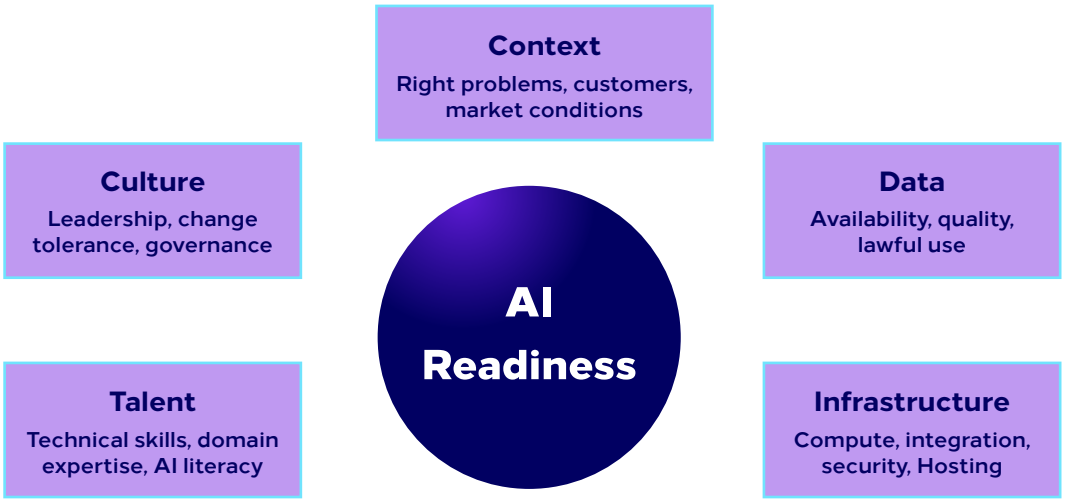

Figure 3.1: Five Readiness Dimensions Framework

This framing makes it explicit that AI readiness is not only about internal capability, but also about choosing the right problems, for the right customers, in the right market conditions. It prevents technically feasible but strategically unsound AI initiatives, and anchors adoption decisions in business reality rather than technological possibility alone.

Companies are advised to align AI ambition with demonstrated readiness, recognizing that premature deployment beyond current capability often results in failed initiatives that undermine confidence and delay future adoption more than a disciplined, phased approach would. The readiness dimensions used in this guide are aligned with widely adopted AI maturity models that assess companies across data, technology, people, and governance dimensions, and are intended to support prioritization and sequencing rather than produce a single compliance score [Gartner, 2023].

3.1. AI Literacy and Organizational Understanding

AI literacy represents a foundational enabler of effective and responsible AI adoption. AI literacy should be developed as a structured capability across roles, with differentiated requirements for leadership, technical teams, and general staff, supported by training programs, tracking mechanisms, and alignment with relevant competency frameworks. It should extend beyond engineering teams to leaders, product, commercial, legal, compliance, and operational roles. Without sufficient literacy, companies risk unrealistic expectations, poor interpretation of outputs, and ineffective governance. Companies should therefore treat AI literacy as an ongoing capability-building requirement rather than a one-time training exercise. Once organization readiness has been assessed, the next step is to determine how AI adoption should be structured across enterprise strategy, sourcing, implementation, and scaling. AI-First Enterprise Strategy and Implementation Model.

04 AI-First Enterprise Strategy and Implementation Model

AI adoption should be grounded in a single enterprise framework that aligns strategy with execution. An AI-first approach is not a side initiative or a technology upgrade; it is an operating principle that assumes AI is the default layer in designing products, workflows, and services. This does not mean automating everything blindly. It means starting with a simple question: can this task be executed reliably, economically, and within acceptable risk by an AI system? If the answer is yes, AI becomes the primary execution layer, while human roles shift toward oversight, judgment, and higher-value decisions.

Sustainable advantage, however, does not come from choosing the most advanced AI model. Frontier models rapidly commoditize. What compounds is proprietary context: domain-specific data, embedded workflows, customer signals, and structured feedback loops that continuously improve performance. Organizations that capture real usage data, measure outcomes, integrate corrections, and refine systems over time build intelligence that competitors cannot replicate. The true moat in AI adoption is not model access, but ownership of context and discipline in turning it into operational advantage.

4.1. Strategic Foundation: From Intent to Direction

Before selecting use-cases or investing in tools, companies should define a clear AI-first enterprise posture that answers structural questions, such as:

- Where does AI create measurable advantage across the enterprise? This includes internal productivity, cost structure redesign, product differentiation, new revenue streams, risk mitigation, and strategic defensibility.
- Which processes should be AI-led by default, and which require human execution due to complexity, regulatory sensitivity, or ethical constraints?
- What level of autonomy, data exposure, and decision authority is acceptable given the company's market role and regulatory environment?
- Where should the company build proprietary AI capability, where should it rely on partners, and where should it acquire capability?

An effective AI-first strategy aligns AI with business objectives, competitive positioning, proprietary data assets, cost structure realities, and governance expectations. Without this alignment, adoption fragments into disconnected pilots driven by tools rather than value.

4.2. AI Adoption Decision Framework Under the "AI-First" Principle

The following framework sets out steps to help determine the strategic direction before beginning detailed execution, after confirming the suitability and feasibility of AI compared with other alternatives:

Step 1: Enterprise Objective
Define whether AI is primarily intended to redesign internal cost structures and productivity, differentiate products and services, enable new business models, or support a combination of these objectives.

Step 2: Data Advantage Assessment
Assess whether the company possesses proprietary, high-quality, domain-specific data that is appropriately available and governed, and that can support defensible AI capabilities.

Step 3: Human vs Agent Allocation Principle
Determine which processes can be executed reliably by AI within risk tolerance and which should remain primarily human-led with AI support.

Step 4: Control vs Speed Trade-off
Decide whether long-term differentiation, speed of deployment, or a balanced approach should drive sourcing choices, while taking into account applicable legal and regulatory requirements.

Step 5: Risk and Exposure Mapping
Assess whether the AI system affects customer rights, pricing, safety, sensitive data, or mission-critical services, and calibrate oversight accordingly. The output of this framework is not a list of tools.

It is a defined enterprise direction that determines:

- Has the suitability and feasibility of the AI option been established compared with other? alternatives?
- Which AI capabilities are core?
- Which execution model to adopt?
- How aggressively to scale?
- How can AI support the development of human capabilities and the reallocation of roles and tasks in ways that enhance productivity?
- And how AI integrates into long-term growth strategy?

4.3. Execution Architecture: From Strategy to Scale

Strategic clarity should be translated into execution discipline. AI initiatives typically fail not due to weak ambition but due to weak sequencing, sourcing, and cost management. Execution should therefore integrate capability sourcing, phased scaling, and value tracking within one coherent model.

4.4. Capability Sourcing: Build, Partner, Acquire

AI capability may be developed through build, partner, or acquire pathways. Build approaches offer greater control and support strategic differentiation. Partner approaches support faster deployment and access to specialized capability. Acquire approaches provide rapid access to products, talent, or market position. In practice, many technology companies will adopt a hybrid model depending on the importance of control, speed, and strategic value.

As illustrated in Figure 4.1, organizations typically evaluate capability sourcing across build, partner, and acquire options, each with distinct trade-offs in control, speed, and strategic value.

Figure 4.1: Build-Partner-Acquire Decision Framework

4.5. Phased Scaling Model

AI adoption should proceed through phased scaling. Discovery and structural readiness establish priorities and assess feasibility. Pilot and validation test technical performance, economics, and risk in controlled environments. Enterprise integration expands validated use cases and strengthens monitoring and cost observability. Continuous optimization improves performance, cost efficiency, and operating model maturity over time.

As shown in Figure 4.2, adoption progresses through four stages discovery, pilot, scale, and optimization, each with defined objectives and outputs.

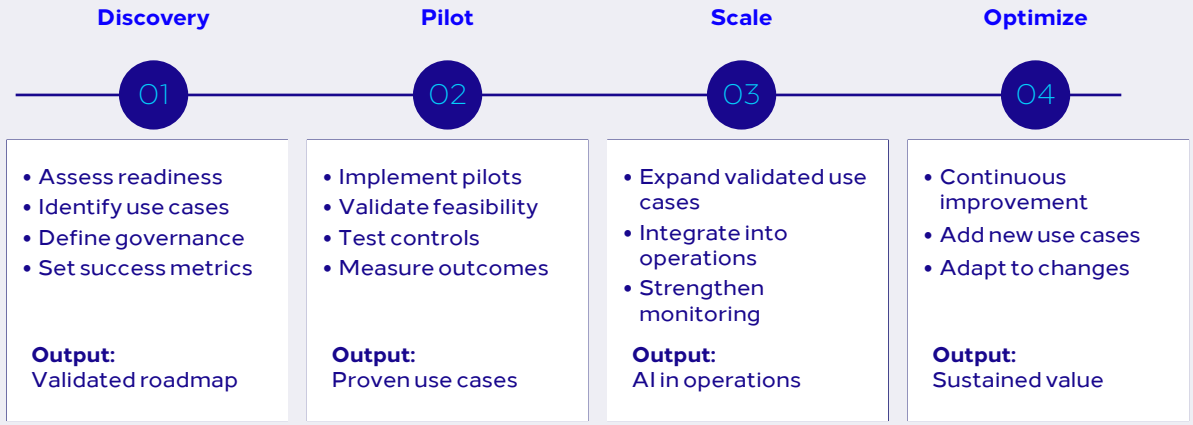

Figure 4.2: Phased Adoption Timeline

4.6. Workforce Reallocation and Organizational Design

As AI adoption expands, tasks that can be executed reliably by AI may shift toward agent-led or AI-supported execution. Human roles should increasingly focus on oversight, escalation, innovation, relationship-driven work, and higher-value judgment. Talent strategy should therefore prioritize AI literacy, architecture capability, and risk management rather than routine execution alone.

4.7. Investment Discipline and Cost Structure Management

AI introduces cost structures that differ from traditional IT, including talent, infrastructure, development, inference, monitoring, and governance costs. Inference costs in particular can materially affect margins at scale. Companies should therefore apply usage forecasting, cost-per-transaction tracking, value measurement, and capital allocation discipline from the outset. AI initiatives should be managed as a portfolio, with redesign or retirement of underperforming initiatives where necessary.

With strategy and implementation principles in place, the next step is to identify where AI can be applied most effectively across internal operations, customer-facing solutions and applications, and agent-based execution models.

05 AI Application Domains

This section translates the strategic and organizational foundations set out earlier in the guide into practical domains of AI application for technology companies. It addresses AI as a full-stack capability, internal operations, customer-facing solutions and applications, and AI agents. The examples presented are illustrative rather than limiting and are intended to support structured prioritization and sequencing.

5.1. AI as a Full Stack: Considerations

AI adoption should be understood through the lens of a full, integrated AI stack, rather than as isolated tools, models, or point solutions. This stack spans multiple interdependent layers, each of which may represent a source of competitive advantage, cost exposure, or service opportunity for technology companies (Stanford HAI, 2024; Gartner, 2023; NVIDIA, 2023):

Solutions and services

industry-specific offerings, managed services, system integration, and advisory capabilities built on top of AI.

Applications and agents

user-facing systems, AI assistants, and autonomous or semi-autonomous agents.

Models

large language models (LLMs), domain-specific models, and both open-source and proprietary alternatives.

Platforms and infrastructure services

cloud, edge, orchestration, security, observability, and data pipelines.

Compute and chipsets

GPUs (Graphics Processing Units), TPUs (Tensor Processing Units), and other accelerators required for model training and large-scale inference.

Infrastructure and facilities

data centers, power availability, cooling, and physical hosting environments that support AI workloads.

As illustrated in Figure 5.1, the AI stack spans multiple interconnected layers, where decisions at one layer can materially affect performance, cost, scalability, and vendor dependency across the system.

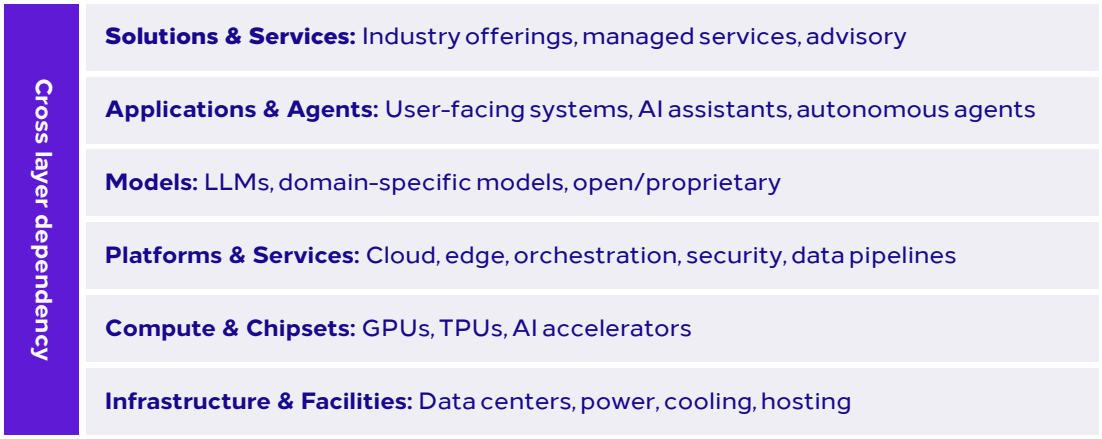

Figure 5.1: AI Stack

These layers are no longer loosely coupled. Cross-layer dependency has become a dominant planning constraint, as choices made at one layer materially affect feasibility, operating cost, performance, and time-to-market for AI-enabled solutions and applications [Stanford University HAI, 2024]. As a result, AI strategy, when applicable, should account not only for technical capability, but also for structural, commercial, and governance implications that emerge across the full stack.

Sustainable advantage depends on proprietary data, embedded workflows, and feedback loops rather than model access alone. It emerges from proprietary data, embedded workflows, domain-specific feedback loops, and continuous refinement based on real usage

Intellectual Property and Ownership Across the Stack

Intellectual property considerations arise across the AI stack and should be addressed explicitly in adoption decisions. These include ownership of outputs, rights over adapted models, licensing obligations, and constraints related to data usage. Companies should therefore treat IP not as a legal afterthought, but as a design and governance consideration.

Vendor Lock-In and Dependency Management

The AI stack introduces multiple forms of vendor dependency across models, infrastructure, orchestration, and deployment patterns. Effective dependency management requires architectural and contractual choices that preserve optionality, including modular design, abstraction layers, support for multiple providers where appropriate, and periodic reassessment of vendor concentration risk.

5.2. Training, Inference, and Token-Based Economics

Economic considerations within the AI stack are increasingly dominated by inference economics, rather than by one-time development or deployment costs. Unlike traditional software systems, many AI-enabled services operate under usage-based and token-based pricing models, where costs scale dynamically with consumption [McKinsey & Company, 2023; Gartner, 2024].

Training Economics

Model training and fine-tuning typically involve high upfront costs driven by compute intensity, data preparation, experimentation cycles, and specialized talent. While these costs are often episodic and predictable within defined scopes, they can escalate rapidly when companies pursue custom models, frequent retraining, or large-scale domain adaptation. Training decisions therefore carry long-term implications for cost structure, capability ownership, and time-to-market [Stanford University HAI, 2024].

Inference Economics

Inference represents a recurring and scalable cost, directly tied to user adoption, interaction frequency, context size, and response length. Each prompt, response, agent interaction, or autonomous action consumes tokens and compute resources. As AI capabilities are embedded into customer-facing services, internal workflows, or always-on agents, inference costs become a central determinant of operating margins and service viability.

Token-Based Pricing Implications

Token-based pricing introduces a structural shift in how digital services are costed and monetized. Longer prompts, larger context windows, RAG, multi-step reasoning, and agent-based workflows can significantly increase token consumption per transaction. Failure to manage inference economics can materially erode margins and undermine scalability. Without deliberate design, services that appear successful from a usage perspective may become economically unsustainable at scale.

As a result, companies should explicitly incorporate token economics into product and service design. This includes forecasting usage patterns, optimizing prompt structure, managing context size, selecting appropriate models for different tasks, and aligning customer pricing with underlying inference costs. Cost observability, usage monitoring, and continuous optimization are no longer optional; they are foundational capabilities for any company operating AI at scale.

Failure to account for training and inference economics early can result in structurally mispriced offerings, margin erosion, or forced redesign after market entry. Conversely, companies that internalize these dynamics can design AI services that scale sustainably, balance performance with cost, and maintain strategic control over their AI-enabled value propositions.

5.3. Internal Operations

Internal operations typically represent the fastest and lowest-risk path to AI value. They provide controlled environments, clearer success metrics, and lower external exposure, while allowing companies to build practical capability before moving to more visible deployments.

AI-Enabled Interaction with Enterprise Systems and Process Automation

AI-enabled interaction layers may be applied across enterprise systems such as ERP, CRM, HR, and finance platforms, allowing users to express requests in natural language while execution remains subject to existing authorization and control frameworks. This can reduce administrative friction and improve internal process efficiency where logging, auditability, and permission controls remain in place.

Software Development, and Architecture, Testing

AI-assisted software development can support code generation, documentation, design artifacts, and testing activities, but outputs should remain subject to formal review and validation by accountable teams. AI-enabled testing may also improve coverage, prioritization, and cycle times when integrated into existing engineering and CI/CD practices.

Communication, Content, and Marketing Enablement

AI supports document processing, content creation, and marketing activities by extracting and structuring information from documents, generating drafts for reports, proposals, and campaigns, and assisting with internal and external communications. In marketing contexts, AI-enabled platforms support content ideation, campaign personalization, and performance analysis. Effective use requires defined scopes, human review workflows, and safeguards to prevent quality, brand, or credibility risks [UNCTAD, 2023].

AI-assisted correspondence tools generalize beyond email to include internal messaging, notifications, customer acknowledgments, and standardized communications. These tools improve consistency and efficiency but should operate under clear policies governing confidentiality, tone, and approval, with accountability remaining fully with the company [OECD, 2023; Gartner, 2024].

IT Service Operations

AI may support IT service operations through ticket analysis, troubleshooting, access management, and service inquiry handling. It may also strengthen cybersecurity monitoring by identifying anomalies, correlating signals, and supporting triage. In both cases, deployment should begin in assistive or monitoring modes before expanding to more automated action.

Functional Job Groups and Knowledge Work

AI increasingly supports functional roles across finance, legal, customer service, procurement, human resources, and internal consulting through drafting, summarization, analysis, and routine inquiry handling. Productivity gains are strongest where AI supports first drafts, insights, or recommendations while humans retain responsibility for validation, judgment, and final decisions.

5.4. Customer-Facing Solutions and Applications

Customer-facing AI applications reflect the most visible and strategically significant opportunities, but also require higher maturity due to direct user impact, trust considerations, and regulatory sensitivity. As a GPT, AI enables companies not only to enhance existing products, but to redefine service offerings and create entirely new customer value propositions across markets. For conversational interfaces such as chatbots, companies should ensure clear user disclosure, defined escalation to human support, and appropriate handling of personal data in line with regulatory requirements.

In this context, the use cases discussed in this section are presented by way of example and not limitation. They are intended to illustrate representative application patterns rather than define an exhaustive or prescriptive list of AI applications. Technology companies are encouraged to assess AI opportunities through the lens of the specific sectors they serve, the problems faced by their customers, and the value gaps that exist within those sectors.

Adopting AI from a sector-oriented perspective allows companies to tailor capabilities to domain-specific needs and, in many cases, to create differentiated niche markets where specialized knowledge, data, and workflows provide defensible advantage. Such an approach supports deeper customer relevance, clearer value propositions, and more sustainable competitive positioning than generic, horizontal AI offerings.

Accordingly, Companies are encouraged to approach customer-facing AI strategically by embedding AI into existing products through a phased roadmap and, where appropriate, by using domain-specific models, retrieval-augmented generation (RAG), or other approaches that strengthen relevance, defensibility, and customer value. Capability development may involve build, partner, or acquire choices depending on the importance of speed, control, and strategic differentiation.

AI-Enabled Client Engagement, Acquisition, and Growth

AI may support client engagement and growth through qualification, demand forecasting, proposal support, inquiry handling, and account analysis. Value is strongest where AI augments expert judgment and is embedded into existing commercial and delivery workflows with clear accountability.

Personalization, Intelligent Search, and Experience Optimization

AI-driven personalization may improve discovery, support, and user experience by adapting content, recommendations, or workflows to customer behavior and context. Such use should remain proportionate, transparent, and aligned with customer expectations and applicable privacy requirements.

Pricing, Offers, and Service Optimization

AI may support pricing, offer design, and service optimization in contexts where demand, usage, and customer needs vary across different segments. When applying AI in pricing-related contexts, companies should consider its operating costs, governance requirements, and implications for customer trust, while ensuring sufficient transparency and consistency with consumer protection requirements and other relevant regulations.

Anticipating and Reducing Service Issues

Predictive and preventive services use operational signals to anticipate issues before they affect customers. This may improve reliability and reduce reactive support when supported by sound data, clear intervention thresholds, and appropriate human oversight.

AI may also be introduced as a modular value-added capability within existing offerings, allowing customers to adopt enhanced functionality without requiring uniform adoption across the full service. This can support gradual uptake, clearer monetization pathways, and controlled expansion of AI-enabled features.

5.5. AI Agents

AI agents represent a more advanced execution model that can support multi-step reasoning, planning, and controlled interaction with enterprise systems or workflows. They may be applied across internal operations and customer-facing contexts where tasks span multiple stages or systems. Premature autonomy in agent systems can introduce disproportionate risk.

Because agents may operate with greater autonomy and system access than conventional AI applications, deployment should begin with constrained authority, human oversight, comprehensive logging, and gradual expansion based on demonstrated reliability and acceptable risk.

06 Governance as a Risk-Based Operating Model

Effective AI governance begins before deployment, not after. It starts with disciplined use case selection and continues through proportionate risk management as systems move into production. For technology companies, governance is therefore not primarily about adding approval layers or slowing delivery, but about ensuring that AI initiatives are selected, sequenced, and controlled in a way that maximizes value while managing risk in proportion to impact.

Not all AI opportunities merit pursuit. Companies should prioritize AI initiatives based on expected value, implementation feasibility, and strategic alignment, rather than technological enthusiasm. International practice distinguishes between narrowly scoped, assistive AI use cases and more complex, autonomous applications, emphasizing deliberate selection and sequencing to avoid misaligned or premature deployment [ISO/IEC TR 24030].

A commonly used decision tool at this stage is the impact–effort matrix, which plots candidate use cases along two dimensions:

- **Expected impact**, including financial return, strategic value, and operational benefit; and
- **Implementation effort**, including cost, technical complexity, time, integration burden, and risk.

High-impact, low-effort use cases represent early wins that should be prioritized, as they build organizational confidence and capability while funding larger initiatives. High-impact, high-effort use cases represent strategic bets that require phased execution, executive sponsorship, and stronger governance controls. Low-impact use cases should be deferred regardless of effort, while high-effort, low-impact initiatives should be explicitly rejected, as interesting technology alone does not constitute a business case.

Evaluation of candidate use cases could be grounded in realistic assumptions rather than optimistic projections. Core evaluation criteria may include:

- **Quantified value and clear success metrics,**
- **Feasibility assessment covering data availability**, technical complexity, integration requirements, and talent needs,
- **Risk assessment** addressing the consequences of failure, misuse, or unexpected behavior,
- **Strategic fit with organizational** direction rather than distraction from priorities.

Figure 6.1: Impact-Effort Prioritization Matrix

Additional factors such as operational sustainability, regulatory and compliance exposure, user and stakeholder acceptance, and long-term scalability often determine whether an AI application delivers lasting value beyond initial deployment. Companies should also be explicit about dependencies, recognizing that advanced customer-facing applications frequently rely on internal AI maturity as a prerequisite. Pursuing high-exposure use cases before mastering internal AI operations is a common cause of avoidable failure.

Once a use case passes this selection discipline, governance becomes the mechanism that enables safe execution and scaling. AI risks differ significantly depending on how and where systems are used. Internal productivity applications such as AI-assisted coding, document drafting, or analytics typically carry limited external impact and can be governed through lightweight controls embedded in existing engineering and operational processes. In contrast, customer-facing applications, automated decision-making, and AI agents with access to systems or workflows introduce higher levels of risk related to trust, security, fairness, and service continuity.

A practical governance approach is therefore anchored in use-case-level risk assessment rather than uniform controls. Before deployment, each AI initiative should be assessed across a focused set of risk dimensions, including:

- **Potential impact** of incorrect or harmful outputs,
- **Degree of autonomy** granted to the system,
- **Sensitivity of** the data involved,
- **Exposure to** misuse or security threats, and
- **Level of visibility** to customers or external stakeholders.

The outcome of this assessment determines the depth of governance required, not the technology itself. One structured way to operationalize this approach is to align governance controls with recognized AI risk management functions, such as governing, mapping, measuring, and managing AI risks, so that risk treatment is explicit, repeatable, and proportionate across the AI portfolio [NIST, 2023].

As illustrated in Figure 6.2, governance intensity should scale with risk level, with progressively stronger controls applied to higher-risk AI applications.

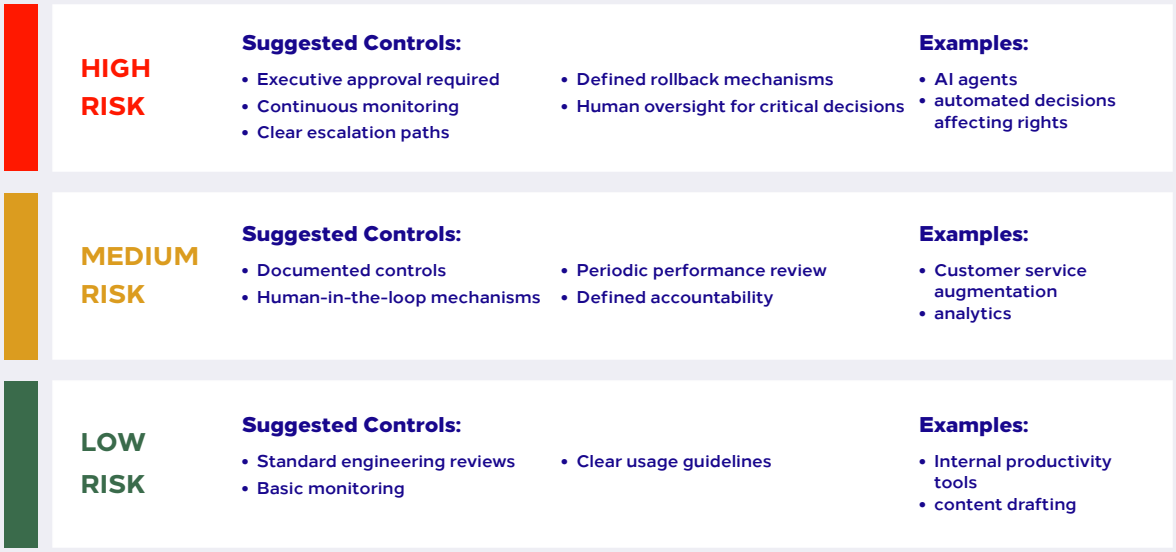

Figure 6.2: Governance Intensity by Risk Level

Low-risk use cases may proceed with standard engineering reviews, basic monitoring, and clear usage guidelines. Medium-risk applications typically require documented controls, defined human-in-the-loop mechanisms, and periodic performance review. High-risk applications, such as AI agents executing actions across systems or AI-driven decisions affecting customer rights, require explicit executive approval, continuous monitoring, clear escalation paths, and well-defined rollback mechanisms. Oversight mechanisms may range from human-in-the-loop to human-on-the-loop and fully autonomous systems with defined boundaries, with appropriate override and deactivation capabilities for higher-risk applications.

Ultimately, governance effectiveness depends on clear ownership. Accountability should remain explicit and aligned with organizational roles. Executive leadership sets strategic direction and risk tolerance. Operational functions define standards, guardrails, and monitoring mechanisms. Delivery teams remain accountable for system behavior in production, regardless of whether outputs are generated by humans or AI. The use of artificial intelligence does not transfer responsibility away from the company.

When governance is framed in this way, starting with disciplined use case selection and extending through risk-proportionate controls, it becomes an enabler of confident AI adoption, allowing companies to scale capabilities while maintaining trust, resilience, and control as systems evolve. While this section addresses how AI initiatives should be selected, prioritized, and governed, responsible adoption also depends on the practical safeguards applied during design, deployment, and operation. These safeguards are addressed in the following section.

6.1. Responsible AI Use: Operational Safeguards and Risk Awareness

Responsible AI use requires more than governance structure alone. It also requires deliberate consideration of regulatory, operational, ethical, and organizational factors to ensure that AI use is justified, proportionate, and aligned with strategic objectives. AI should be applied where it meaningfully improves outcomes, not simply because it is technically feasible. Introducing AI without a clearly defined problem, measurable objective, or value hypothesis can lead to inefficient use of resources and unnecessary complexity, a risk consistently highlighted in international guidance on AI deployment and digital transformation [UNESCO, 2023; OECD, 2022].

A foundational consideration is compliance with applicable national laws, regulations, and regulatory frameworks governing technology, data, cybersecurity, privacy, and sector-specific activities. AI systems rely on data availability, quality, and lawful use; where relevant data is unavailable, of insufficient quality, or cannot be used for the intended purpose, AI initiatives are unlikely to succeed and may expose companies to regulatory or operational risk. Effective AI adoption therefore depends on well-governed data foundations, clear data provenance, and alignment with national regulatory expectations [ISO/IEC TR 24030; ITU, 2022]. Particular attention should be given to the use of external AI tools, ensuring that sensitive, confidential, or personal data is not exposed without appropriate contractual, technical, and policy safeguards.

Companies should also prioritize transparency and coordination with customers and stakeholders, particularly when AI influences decisions, recommendations, or automated actions that affect user experience, pricing, service delivery, or rights. Clear communication about the role of AI, the nature of automated decision-making, and the boundaries of system behavior is essential to maintaining trust and managing expectations. Lack of transparency can undermine adoption even when technical performance is strong.

AI systems introduce security risks that extend beyond traditional software vulnerabilities and require model-specific controls. These include prompt injection attacks that manipulate system behavior through malicious inputs, data poisoning during training or fine-tuning processes, model extraction or distillation attempts that replicate proprietary capabilities, adversarial inputs that degrade reliability, and agent-level risks where autonomous systems interact with enterprise tools or sensitive data. Unlike conventional applications, AI systems can be influenced indirectly through content, queries, or training artifacts, making input validation, output filtering, dataset governance, access segmentation, monitoring, red-teaming, and rollback mechanisms essential safeguards. Generative AI systems introduce additional risks including hallucination, content authenticity concerns, and intellectual property implications, which require verification processes, clear labeling, and appropriate safeguards. Organizations should treat model integrity as a core governance concern, ensuring that AI deployments incorporate layered controls proportionate to system autonomy, data sensitivity, and operational exposure.

The use of AI to process confidential, highly confidential, or personal data warrants heightened scrutiny due to potential implications for privacy, security, and trust. Companies should assess whether AI adds sufficient value in these cases to justify the additional risk, and should implement safeguards such as data minimization, purpose limitation, controlled access, and continuous monitoring to ensure responsible use [UNESCO, 2023; UNCTAD, 2023]. In regulated or sensitive environments, not all technically feasible AI applications are appropriate for deployment.

Finally, effective AI adoption depends on experimentation, testing, and gradual scaling. Companies are advised to validate AI systems through controlled pilots, testing under realistic conditions, and staged deployment rather than immediate large-scale rollout. A phased approach supports learning, risk reduction, and capability building while allowing companies to adapt to rapid technological change. Continuous monitoring of performance, emerging risks, and evolving best practices is essential, as AI technologies and regulatory expectations continue to develop.

Taken together, these considerations position AI governance not as a barrier to innovation, but as an enabling framework that helps companies deploy AI confidently, responsibly, and in alignment with national priorities and long-term value creation.

07 Appendix

A. References

- Gartner, How to Prioritize AI Use Cases for Business Value, Gartner Research, Stamford, CT, USA, 2023.
- Gartner, AI Maturity Model Toolkit, Gartner Research, Stamford, CT, USA, 2023.
- Gartner, Building AI Literacy Across the Enterprise, Stamford, CT, USA: Gartner Research, 2023.
- Gartner (2023) Building an AI Strategy That Delivers Business Value. Stamford, CT: Gartner Research.
- Gartner, How to Build an AI Roadmap, Gartner Research, Stamford, CT, USA, 2023.
- Gartner, How Generative AI Will Change Knowledge Work, Gartner Research, Stamford, CT, USA, 2024.
- Gartner, AI in Sales and Client Engagement for Technology Providers, Stamford, CT, USA: Gartner Research, 2024.
- Gartner, "The Real Cost of Generative AI: Understanding Inference Economics," Stamford, CT, USA, Gartner Research, 2024.
- Gartner, Stop, Start, Continue for 2025: AI Technique Heat Map, Gartner Research, Stamford, CT, USA, 2024.
- International Organization for Standardization and International Electrotechnical Commission, ISO/IEC TR 24030:2021 – Information Technology – Artificial Intelligence (AI) – Use Cases, Geneva, Switzerland, 2021.
- International Organization for Standardization and International Electrotechnical Commission, ISO/IEC TR 24030:2023 – Information Technology – Artificial Intelligence (AI) – Use Cases (2nd ed.), Geneva, Switzerland, 2023.
- International Telecommunication Union, AI for Good: Enabling AI Readiness and Adoption in ICT and Digital Services, Geneva, Switzerland, 2022.
- International Telecommunication Union, AI in Cybersecurity and Network Operations, Geneva, Switzerland, 2022.
- International Telecommunication Union, Guidelines for AI Governance and Digital Transformation, Geneva, Switzerland, 2022.
- International Telecommunication Union, Artificial Intelligence in Software Development, Testing, and Operations, Geneva, Switzerland, 2023.
- International Telecommunication Union, Towards Autonomous and Intelligent Networks and Services, Geneva, Switzerland, 2024.
- International Telecommunication Union, AI Systems, Risk, and Trust in Critical Digital Infrastructure, Geneva, Switzerland, 2024.
- McKinsey & Company, The State of AI in 2024: Generative AI's Breakout Year, McKinsey Global Institute, New York, NY, USA, 2024.

- McKinsey & Company, The Economic Potential of Generative AI in B2B Sales and Client Operations, New York, NY, USA: McKinsey Global Institute, 2023.
- Microsoft, The Economic Potential of Generative AI: The Next Productivity Frontier, Microsoft Work Trend Index Special Report, Redmond, WA, USA, 2023.
- National Institute of Standards and Technology (NIST), Artificial Intelligence Risk Management Framework (AI RMF 1.0), Gaithersburg, MD, USA, 2023.
- NVIDIA, "Understanding GPU Economics for AI Training and Inference," Santa Clara, CA, USA, NVIDIA Technical White Paper, 2023.
- NVIDIA, NVIDIA AI Enterprise Architecture and AI Infrastructure Overview. Santa Clara, CA, 2023
- Organisation for Economic Co-operation and Development (OECD), "AI, Compute, and Market Concentration," Paris, France, OECD Publishing, 2023.
- Organisation for Economic Co-operation and Development, Algorithmic Pricing and Competition Policy, Paris, France, 2021.
- Organisation for Economic Co-operation and Development, Consumer Policy and Algorithmic Pricing, Paris, France, 2021.
- Organisation for Economic Co-operation and Development, AI Value Creation, Trust, and Adoption in Digital Markets, Paris, France, 2022.
- Organisation for Economic Co-operation and Development, Framework for the Classification of AI Systems, Paris, France, 2022.
- Organisation for Economic Co-operation and Development, Generative AI and the Future of Work, Paris, France, 2023.
- Stanford University Human-Centered Artificial Intelligence Institute, AI Index Report 2024, Stanford, CA, USA, 2024.
- United Nations Conference on Trade and Development, Digital Economy Report 2023: Value Creation and Capture, New York and Geneva, United Nations, 2023.
- United Nations Conference on Trade and Development, Digital Economy Report: Trust, Data, and Artificial Intelligence, New York and Geneva, United Nations, 2023.
- United Nations Development Programme, Artificial Intelligence and the Future of Work and Productivity, New York, NY, USA, 2022.
- United Nations Educational, Scientific and Cultural Organization, Guidance on AI Ethics and Readiness for Public and Private Organizations, Paris, France, 2023.
- United Nations Educational, Scientific and Cultural Organization, Guidance on Ethical and Responsible Artificial Intelligence, Paris, France, 2023.
- United Nations Industrial Development Organization, Industrial Applications of Artificial Intelligence and Quality Automation, Vienna, Austria, 2022.
- World Economic Forum, Harnessing Artificial Intelligence for the Next Wave of Digital Transformation, Geneva, Switzerland, 2023.

08 Appendix

B. Self-Assessment Checklist

This checklist is a self-diagnostic tool designed to support structured AI readiness assessment. It represents the quantitative component of the two-step readiness approach outlined earlier and complements the qualitative assessment across the five readiness dimensions. It is provided for guidance only and is not a compliance, audit, scoring, or certification exercise.

Each statement should be marked only if it is currently met in practice, not merely planned or aspirational. The preferred outcome is to meet all applicable criteria; however, relevance depends on context, sector, and intended use cases. The checklist should be used to identify readiness gaps, support prioritization, and guide sequencing of AI initiatives, rather than to label maturity or measure performance.

Strategic Readiness
AI objectives are documented and tied to business outcomes
A designated executive sponsor is formally accountable for AI initiatives
Budget is allocated for AI investment
Competitor AI activity is understood

Data Readiness
Data relevant to priority use cases is inventoried
Data quality is assessed with remediation plans for gaps
Legal basis for intended data uses is confirmed
Cross-border data transfer constraints are understood
Data is classified according to confidentiality

Infrastructure Readiness
Compute capacity for AI workloads is available or accessible
Data residency requirements can be met
Integration architecture supports AI system connectivity

Talent Readiness
Existing AI-related skills across the company have been identified and documented
Skill gap closure plan exists (hire, train, or partner)
Leadership understands AI sufficiently for investment decisions
Staff training plan exists for AI system adoption

Organizational Readiness
Management can identify AI opportunities in their domains
Leadership visibly supports AI adoption through actions
Expectations about timelines and outcomes are realistic
Change management capacity is adequate
Governance roles and policies are established

Per-Initiative Readiness
Problem is clearly defined with measurable success criteria
Value is quantified with realistic assumptions
Data availability and quality are confirmed
Technical feasibility is validated
Risks are assessed
Resources are allocated
Stakeholders are identified and engaged

هيئة الاتصالات والفضاء والتقنية

Communications, Space &
Technology Commission