Deloitte.

Beyond Pilots: Transforming IT for AI at Scale

A Practical Guide for Technology Leaders

Executive Summary

Artificial Intelligence (AI) has moved beyond the proof-of-concept stage. Across industries, AI pilots are giving way to enterprise-wide deployment. Scaling AI effectively is now a strategic imperative: organizations that can operationalize AI across the enterprise will unlock greater efficiency, innovation, and durable competitive advantage in an increasingly digital economy. Those that fail to scale risk being left behind as AI becomes a core driver of business performance and differentiation.

CIOs are discovering that the challenge is no longer how to start with AI. It's how to **embed it in their organization, architect it to scale, govern it, and sustain measurable value over time.**

Harvard Business Review, drawing on recent research from the MIT Media Lab, reports that the vast majority of GenAI investment "approximately 95%" has not yet produced measurable financial returns. HBR explains that while individual employees are experiencing productivity gains from GenAI tools, those improvements are not translating into visible P&L impact, and organizations continue to face challenges scaling GenAI across the enterprise. The analysis also notes that most experimentation budgets still flow toward sales and marketing initiatives, even though operational and back office transformations typically yield stronger returns.[1]

To move forward, organizations need to focus on **how** AI can enable customer and business value and remove **technical, operational, and organizational barriers** that stand between AI investment and business impact.

Across four chapters, this white paper provides practical and actionable insights for technology leaders along two parallel value tracks: First, AI that improves IT performance and operating efficiency, and second, AI that enables direct customer and business value.

- **Chapter 1: Building an AI-Ready Technology Organization**
- **Chapter 2: Strengthening the Technology Core**
- **Chapter 3: Establishing AI Governance and Risk Management**
- **Chapter 4: Defining and Measuring AI Value**

In 2026 and beyond, AI leadership will be measured not by the number of pilots, but by the resilience, trust, and sustainable advantage AI brings to the business.

Chapter 1: Building an AI-Ready Technology Organization

The emergence of AI, in particular Generative AI and Agentic AI, marks a fundamental transformation in technology management and delivery. CIOs are now required to move beyond their traditional roles as technology stewards and become architects of hybrid human-AI enterprises. In these new models, autonomous systems and AI agents coordinate complex, multi-step processes through close human-agent collaboration. While this will reduce costs and improve efficiency, it fundamentally disrupts all aspects of IT – from the underlying technology infrastructure to people and knowledge work. CIOs and IT leaders must start with a refresh of the IT Strategy.

Refreshing the Enterprise IT Strategy

As we will explore throughout this paper, scaling AI is at the core of enterprise technology - from legacy systems integration to data architecture, and from governance to fundamental shifts in work. While standalone AI strategies may be appropriate for early-stage, low-maturity organizations, IT leaders looking to scale must embed AI into their enterprise IT strategy to deliver long-term value.

Embedding AI into the enterprise IT strategy ensures that investments in platforms, data, applications, and talent evolve together, enabling AI to become a multiplier rather than a parallel effort. This approach links AI directly to business outcomes - improved productivity, enhanced customer experience, modernized core systems - and avoids duplicate spending on siloed tooling or one-off models. It also ensures that risk, governance, and regulatory considerations are addressed consistently across the technology landscape.

Redesigning Work: Skills, Culture, and Human Value

A recent study by Anthropic analyzed over 4 million real-world AI interactions and found that 57% augmented human work, while 43% focused on automation. This suggests that AI is more often used to enhance human capabilities rather than replace them.[2] Another study showed that the impact of AI depends less on a person's past experience and more on how effectively they evaluate AI recommendations.[3] For technology teams, this means both junior and senior staff can leverage AI to augment their skills and improve performance.

As AI agents execute tasks at greater speed and scale, human workers must increasingly address challenges that require uniquely human skills. For IT professionals, these shifts demand a balance of technical fluency in AI tools and soft skills in governance, change management, and human-AI collaboration. CIOs need to be a key part of helping redefine how work across the organization is done leveraging AI. This means enabling people to concentrate on higher-value activities that require judgment, creativity, and empathy, while AI agents manage routine tasks that are homogenous and historically manual-intensive.

To develop human-AI collaboration, leaders should:

1. **Foster a willingness to collaborate with Generative AI** by helping employees see it as a partner and a tool to achieve their goals.
2. **Offer learning opportunities** that build a basic understanding of AI and its potential to improve work.
3. **Communicate effectively to build trust**, and provide strategies and support that empower employees to leverage AI.
4. **Create a safe space** for employees to explore Generative AI and integrate it into existing workflows, encouraging adaptability and openness to new processes.
5. **Establish clear ethical guidelines** for the use of Generative AI to ensure compliance with legal and regulatory requirements.
6. **Raise employee awareness of the potential risks** and ethical considerations, and take appropriate steps to mitigate them.

Education and experience have long been the cornerstones for developing expertise and preserving institutional knowledge. Through years of apprenticeship, training, and collaboration, employees gained this knowledge and built expertise.

But Generative AI is commoditizing knowledge. The future of expertise lies in the ability to translate, contextualize, and challenge knowledge in real-world situations where human experience and business context make all the difference.[4]

The traditional career ladder was built based on linear knowledge accumulation. This pathway is no longer sufficient. Organizations that do not redesign their apprenticeship and talent development pathways risk losing both their future talent and their competitive advantage. As discussed in the Deloitte article **Reshaping expertise: How GenAI is changing knowledge work**, to remain relevant, technology organizations must upskill their existing workforce while attracting AI-savvy talent, ensuring employees can effectively collaborate with AI tools in daily workflows.

Transforming the Technology Operating Model

Artificial Intelligence is reshaping the technology landscape, introducing new skills, roles, and ways of working. Traditional operating models are not equipped to manage the complexity of augmented human-AI collaboration, increased automation, and accelerated innovation cycles. To remain competitive and unlock enterprise value, the Technology Operating Model (TOM) must evolve as well.

The first step is assessing the technology organization's maturity, as priorities vary by stage. Once the baseline is clear, a structured capability framework can guide transformation. Deloitte's Technology Capability Model provides this framework by defining the end-to-end capabilities needed to deliver world-class technology services. It offers a consistent lens to evaluate effectiveness, identify gaps, and prioritize investments for secure, scalable AI adoption.

Deloitte's Technology Capability Model

The capability model is organized into the six technology domains shown here.

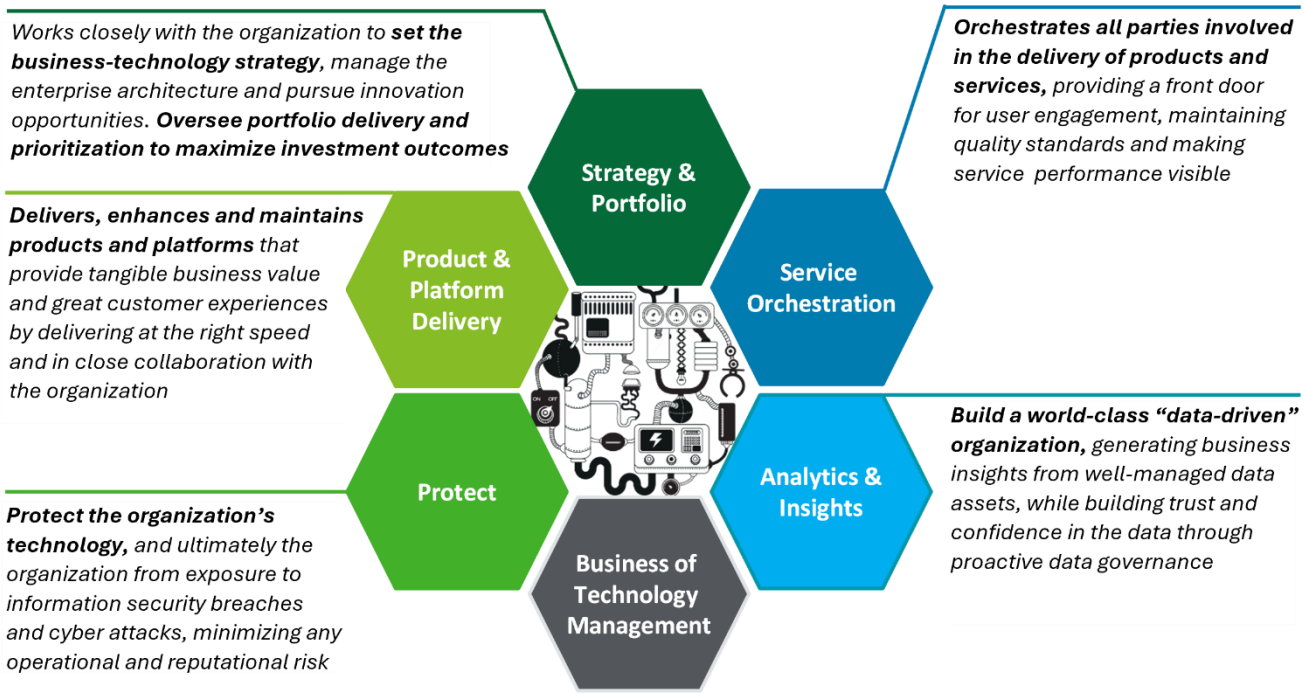

Technology leaders must focus on the critical capabilities that enable scalable and effective AI integration. To progress beyond the foundational stage and unlock greater value, prioritize the following capabilities:

Domain	Strategic Priorities
Strategy & Portfolio	• Refresh the technology strategy to support AI-enabled business priorities and required AI investments in talent, applications, tools, and underlying data and infrastructure.
Product & Platform Delivery	• Establish architecture principles and standardise metadata models. • Leverage AI capabilities of major SaaS platforms to accelerate adoption and benefits (e.g., Salesforce, ServiceNow, SAP, Oracle). • Build access to deep engineering expertise to assemble complex, integrated solutions such as agentic AI. • Leverage tools and capabilities for AI-assisted software development (e.g. code generation, refactoring, or modernization, AI-assisted QA and testing, documentation).
Protect	• Update security policies to address AI-specific risks, such as model drift, data poisoning, and prompt or data bias. • Invest in AI monitoring capabilities to ensure policy compliance.
Business Technology Management	• Launch AI training programs, certifications, and firsthand labs for IT and business teams. • Establish financial processes such as FinOps to align AI investments with strategic objectives and define metrics to measure financial performance, total cost of ownership (TCO), and strategic impact of AI initiatives.
Analytics & Insights	• Establish robust data management practices and analytics strategies to ensure AI initiatives are supported by high-quality, well-governed data. • Invest in data cleansing and improve accessibility to maximize AI accuracy and impact. • Adopt modern data platforms that enable efficient collection, storage, and standardization, supported by consistent data standards and metadata management.
Service Orchestration	• Evaluate and onboard strategic suppliers with proven AI capabilities, and ensure vendor agreements cover data sharing, intellectual property, and AI model ownership to support responsible integration and clarify ownership of intellectual property.

Ultimately, a deliberate and structured approach to evolving the Technology Operating Model and investing in the right foundational capabilities today will enable organizations to scale AI effectively, drive sustainable business value, and remain competitive in an increasingly digital future.

Spotlight: Lemonade - Scaling Performance Through Human-AI Teaming[5]

Lemonade demonstrates how intentionally designed technology architecture, governance, and human-AI collaboration can unlock sustainable enterprise value. Founded in 2015, the company built an AI-native operating model from day one. Its digital agents "Maya" for onboarding and customer interactions, and "Jim" for autonomous claims processing are embedded directly into core workflows rather than layered on top of legacy processes.

Business Impact

Lemonade's approach shows measurable results from AI-first operations:

- **Speed and Efficiency:** Claims processed in seconds through AI-driven automation.
- **Scalable Performance:** Growth without proportional headcount increases.
- **Enhanced Experience:** Improved customer satisfaction through seamless digital interactions.
- **Continuous Improvement:** Built-in learning loops for ongoing optimization.

Human-AI Teaming by Design

AI agents handle high-volume, repeatable processes with speed and accuracy, while people focus on activities that require judgment, empathy, and innovation. This clear delineation of roles accelerates performance while maintaining transparency and trust.

Lemonade illustrates how enterprises can scale AI beyond pilots by embedding agentic capabilities into platforms and workflows, supported by strong governance and human-AI collaboration.

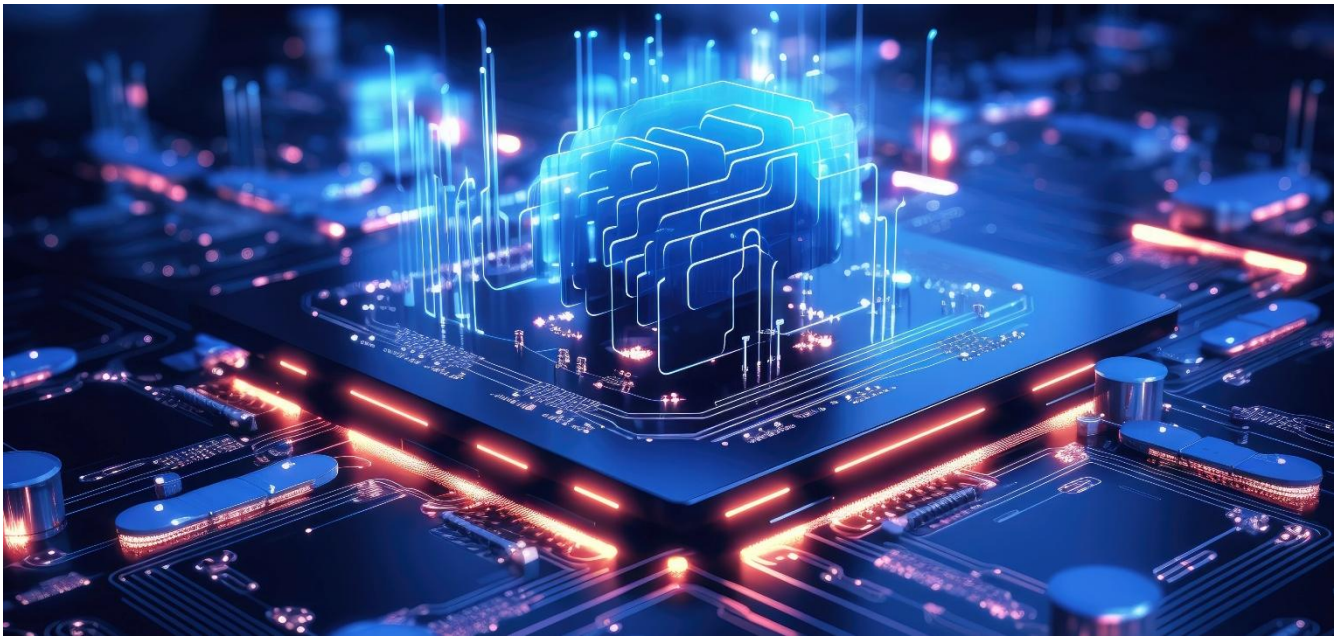

Chapter 2: Strengthening the Technology Core

AI only scales when the technology core makes it easier for the business to adopt and operate AI across functions. This requires shared foundations, for example, data, integration, security, and delivery practices, which reduce time to deploy and improve reliability. If these foundations are weak, AI initiatives multiply in silos, driving duplicate spend and inconsistent risk management. For this reason, decisions about the technology core should be grounded in wider enterprise IT strategy, legacy system, and data environment modernization, ensuring investments support business priorities and can be reused across multiple value streams.

Modernizing Enterprise architecture for the AI Era

Enterprise Architecture has long provided a stable framework for aligning technology with business goals. But in the AI era, the pace of change is challenging that stability. Architecture can no longer be a static blueprint; it must evolve into a capability that responds quickly to changing business needs. This is not just about adding AI to existing systems; it is about rethinking architecture as a driver of growth, resilience, and innovation.

Traditional frameworks, such as The Open Group Architecture Framework (TOGAF) which divides architecture into Business, Data, Application, and Technology domains, still matter. But how they are applied must evolve to meet the demands of AI-driven transformation. This modernization enables architecture to directly support business outcomes such as speed, innovation, and sustainable competitive advantage, while protecting core business processes.

Evolving Business Architecture

Artificial Intelligence (AI) is fundamentally reshaping business architecture and accelerating the pace of change across industries. Organizations are moving beyond traditional process-centric models toward adaptive data-driven architectures that enable agility and innovation.

Key shifts include:

- **Real-Time Decision-Making:** AI enables dynamic, data-informed decisions at scale, requiring architectures that support continuous learning and rapid response.
- **Human–AI Collaboration:** Operating models must define how people and AI can meaningfully collaborate, redefining roles, workflows, and governance.
- **Automation of Core and Support Functions:** Functions such as finance, HR, and supply chain are increasingly automated, transforming value chains and resource allocation.
- **Hyper-Personalization:** AI-driven insights allow businesses to deliver tailored experiences, demanding architectures that support real-time analytics and customer-centric design.
- **Emergence of New Capabilities:** Capabilities like large language model (LLM) management, AI governance, and ethical compliance are becoming essential layers within enterprise architecture.

These shifts require business architects to revisit and renew foundational artifacts including business capability maps, value streams, business process models, and conceptual data models.

Turning Data into a Strategic Advantage

AI requires access to large and diverse datasets with tight integration to enable real-time processing. It uses text, images, audio, video, and other inputs that do not fit into traditional models. In modern businesses, AI runs constantly, powering real-time personalized experiences, autonomous systems, and adaptive decision-making. Tesla's sensor streams, Spotify's listening data, and YouTube's recommendation signals are examples of data that directly drives business outcomes.

The need for diverse data with AI introduces additional complexity to a data landscape that is often fragmented and suffering from poor data quality. These issues slow down AI adoption and increase data-related risks for the organization.

For CIOs, the opportunity is twofold: (i) expand the scope of data architecture to include such new data models and (ii) accelerate governance to match the pace of change, so use cases can scale safely.

To support new types of data, build a flexible data system that can handle text, images, audio, and video all in one place. Start by creating a data lake house which combines the benefits of affordable, flexible storage with organized data management, to reduce the time it takes for teams to find, trust, and reuse data across use cases. Add tools that help the system understand connections and context like special databases for search and a knowledge graph that links related information, making it easier for AI to make sense of complex data and answer questions, such that AI responses are grounded in enterprise knowledge.

At the same time, speed up your data oversight by ensuring processes like DataOps and MLOps are part of your daily operations to reduce defects, rework, and compliance surprises. This means setting up automatic checks for things like data changes, bias, and the history of how data and models are used. Put these checks right into your development process so they happen all the time, not just as occasional reviews.

When data is managed as a strategic asset that is curated, governed, and continuously improved, it becomes the foundation for AI-enabled transformation and measurable business value.

To learn more, please see the Deloitte report **Four Data and Model Quality Challenges tied to Generative AI**[6] , which outlines practical approaches that technology leaders can take to remove barriers and capture value from their data.

Designing AI-Ready Applications for Agility and Trust

In the past, application architecture assumed that a specific input would provide a predictable output. AI-powered applications change that to be more dynamic. Chatbots respond differently to similar prompts, recommendation engines reshuffle results in real time, and fraud detection models adapt as new patterns emerge.

For CIOs, this means rethinking application design for uncertainty and resilience. Requirements should shift from fixed outputs to performance thresholds and risk tolerance (e.g., precision levels, harmful output limits). To maintain trust, applications need built-in monitoring, feedback loops, and fallback logic that allow rapid intervention when AI behavior deviates from expectations.

Equally critical is enabling governed agility. CIOs can accelerate adoption by providing secure, self-service platforms for business users to prototype AI solutions using approved models and prompt catalogs. Modular architectures and low-code capabilities ensure these solutions can scale and industrialize without compromising compliance or security.

Engineering Technology Foundations for Scale and Resilience

Technology infrastructure has traditionally been the backbone of an organization's applications and data. If it was stable and met business performance needs, it operated quietly in the background and out of the limelight. AI has thrust technology into the foreground as a strategic enabler. Training and serving AI models require specialized components such as GPUs, TPUs, vector databases, and orchestration frameworks that did not exist in traditional stacks. These are not incremental upgrades; they are new categories of technology adopted at a pace that outstrips conventional planning cycles. In addition, massive datasets require high throughput storage while training and inference requires low-latency networks.

CIOs and business leaders now need to treat infrastructure decisions as business decisions. Long-term bets on platforms now carry higher risk as specialization accelerates and the pace of change increases. The mandate is clear: build flexible, scalable foundations through hybrid and multi-cloud strategies, embed security and lifecycle management into AI platforms, and leverage ecosystem partnerships to stay ahead.

In practice, this means:

- Rethinking computing strategies, including on-prem, cloud, and hybrid approaches to balance cost, performance, scalability, and regulatory needs;
- Potentially investing in AI-optimized hardware, such as GPUs, high-speed networking, or high-performance storage solutions;
- Accessing deep engineering expertise to assemble AI solutions – there is no one-stop-shop for a platform that meets all needs; and,
- Closely monitoring spend and resource utilization to protect against spiraling costs.

Technology architecture is not just a technical layer, it is a strategic differentiator, directly enabling AI-driven business outcomes.

Finally, Enterprise Architecture can no longer be a back-office function. It is a front-line capability that determines whether AI becomes a source of sustainable advantage or a collection of disconnected pilots. Architecture should evolve from static design to dynamic orchestration which means continuously aligning AI adoption with business strategy, enabling speed and scalability, and embedding trust and resilience at every layer.

Modernizing Legacy Tech

Modernizing legacy technology has been an ongoing, incremental process for years as CIOs reduce technical debt and enable new business capabilities. Like the data challenges noted previously, many legacy solutions are built for binary "yes or no" resolutions and do not lend themselves well to the probabilistic, real-time nature of generative AI outputs, which can demand distributed computing capabilities (e.g. GPUs), efficient storage for large datasets, modern APIs for integration with cloud-native AI platforms, and up to date security models – not to mention different skills sets and governance models.

As outlined in the Deloitte report Three Ways to Approach Legacy Tech Modernization with AI[7], CIOs may be asking: Can artificial intelligence help us accelerate it? Additionally, as CIOs experiment with how AI can evolve operations within the IT function, this can provide further evidence or examples to the rest of the organization of how to leverage AI to transform their operations.

At this stage, three approaches are emerging for tech estate modernization with AI:

1. **Rethinking tech processes** to incorporate generative AI, agents, and reasoning models to help remove technical debt. AI coding tools, for example, can transform how applications are developed and delivered by reducing operational cost and speeding capability delivery. GenAI organizes unstructured data and generates actionable insights, directly supporting better decision-making.
2. **Reengineering the digital core** across a wide range of intelligent technologies that can reshape databases, applications, and platform possibilities. This type of modernization approach covers modern architecture, containers, and the necessary integrations while leveraging AI to help make it as affordable as possible through automated code remediation and data migration.
3. **Reimagining business** capability development with agentic AI. Rather than automating existing processes with AI, businesses might reimagine processes and then incorporate AI. These enterprises have amplified the impact with AI by digitally transforming the front and back office. Some enterprises are just beginning to experiment in this space; for example, using large language models to respond to level one and level two customer calls with minimal human assistance or employing digital twins to simulate the potential impacts of pulling the plug on a legacy system.

Evolving IT Operations

To build and manage AI solutions that are resilient, scalable, and reliably meet business and customer demands, IT operations must adopt new capabilities and tools to manage and govern AI solutions across their lifecycle. Capabilities such as AI Engineering and ModelOps are crucial for efficiency, regulatory compliance, and to maintain trust and transparency in AI systems. Furthermore, while commercial generative AI platforms, custom developed AI solutions, and autonomous agentic AI systems each offer increased business opportunities, they also elevate complexity and risks, requiring enhanced governance and cyber security approaches.

While AI is increasing demand for IT operations teams, AI-enabled operational platforms and tools are driving significant efficiencies to reduce costs and risks while improving service availability, customer satisfaction, and business agility. AIOps uses real-time and historic operations data to predict incidents, identify root cause, and automate remediation workflows, allowing IT operations to shift from a reactive to a proactive model. A large U.S. food processing company, for example, using Generative AI enabled workflows has reduced system-generated incident volumes by 70%, increased first-time resolution by 33%, and improved mean-time-to-resolution by 36%.

Spotlight: AI at Enterprise Scale – Walmart's Platform-Led Transformation

Walmart illustrates how a large, globally distributed enterprise can move beyond pilots to an AI-ready technology organization at scale.

The Challenge: Moving Beyond Fragmented AI Experimentation

Operating across thousands of stores, digital channels, suppliers, and partners, Walmart faced increasing fragmentation as AI adoption accelerated. Moving experiments into production required significant bespoke engineering and existing platforms were not designed to support AI systems at scale. Without a stronger technology foundation, the realization of business value risked being constrained by fragmentation, elevated risks, and accumulating technical debt.

The Solution: Element as the Core of Enterprise AI Transformation

To address these challenges, Walmart invested in Element, a standardized enterprise AI platform embedded into its technology core. Element industrializes AI delivery by providing shared infrastructure, containerized deployment, and MLOps[8] capabilities that support the full lifecycle from experimentation through production. The platform also enables agentic AI through centralized orchestration and observability, supporting standardized "super agents" as governed, scalable entry points to enterprise AI capabilities.[9] As a result, AI shifted from bespoke implementations to a repeatable, enterprise-grade capability.

Business Impact

By treating AI as core enterprise infrastructure rather than a disconnected collection of tools and solutions, Walmart reduced time to production, improved reuse across domains, and strengthened governance as AI adoption expanded. This platform-led approach created the technical conditions for AI to scale reliably, securely, and cost-effectively across a complex enterprise environment.

Walmart's experience underscores that CIOs who invest in resilient AI platforms and standardized engineering foundations can turn AI into durable, repeatable enterprise capabilities rather than a continuous stream of pilots.

Chapter 3: Establishing AI Governance and Risk Management

To operate AI, organizations need a clear strategy driven by AI leaders, robust governance frameworks, and the appropriate technology to successfully manage and scale the AI strategy. The resulting AI framework in the organization should reference and be complemented by global regulatory frameworks and compliance requirements (e.g., EU AI Act, NIST, AI RMF, Canadian directives, and emerging Asian standards). Learn more about how to operationalize AI governance in Deloitte's report "**Operationalizing AI Governance**".

What is Trustworthy AI?

The rapid adoption of artificial intelligence, particularly Generative AI, has underscored the importance of trust in AI systems. Deloitte's **Trustworthy AI framework** provides a comprehensive approach to building, deploying, and operating AI in a manner that is responsible, ethical, and aligned with organizational values. The framework is designed to help organizations navigate the complex landscape of AI risks, compliance requirements, and societal expectations, ensuring that AI delivers value while maintaining stakeholder trust.[10]

But what makes AI trustworthy? At its core, the framework encompasses seven dimensions: **Privacy, Transparency and Explainability, Fairness, Responsibility, Accountability, Robustness and Reliability, and Safety and Security**. These collectively address the risks associated with AI and guide organizations in establishing governance structures that promote responsible AI adoption. The importance of each dimension is outlined below.

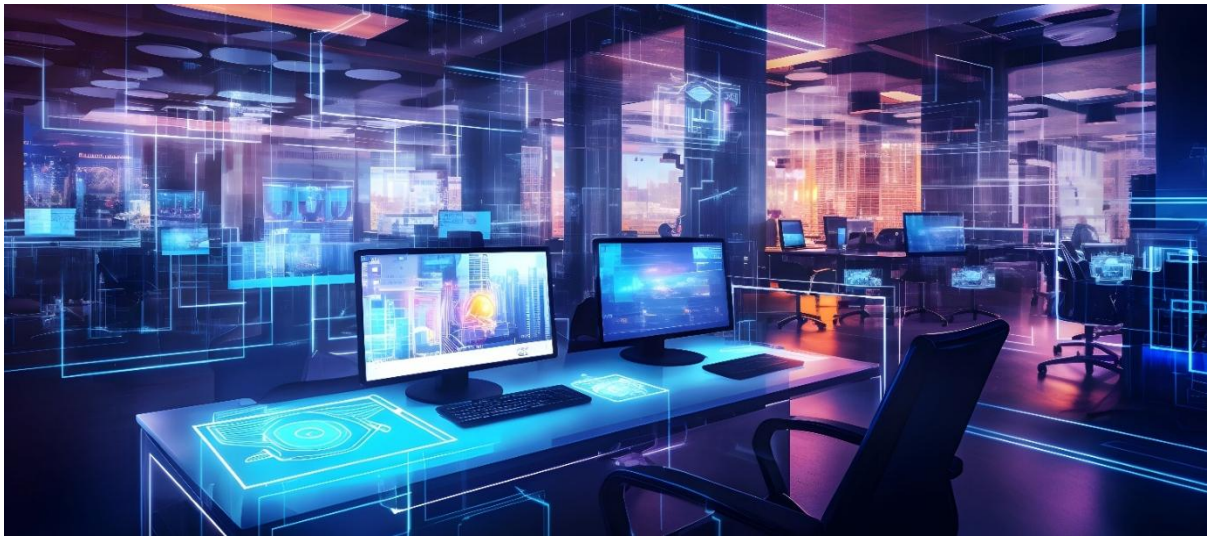

Deloitte's Trustworthy AI Framework

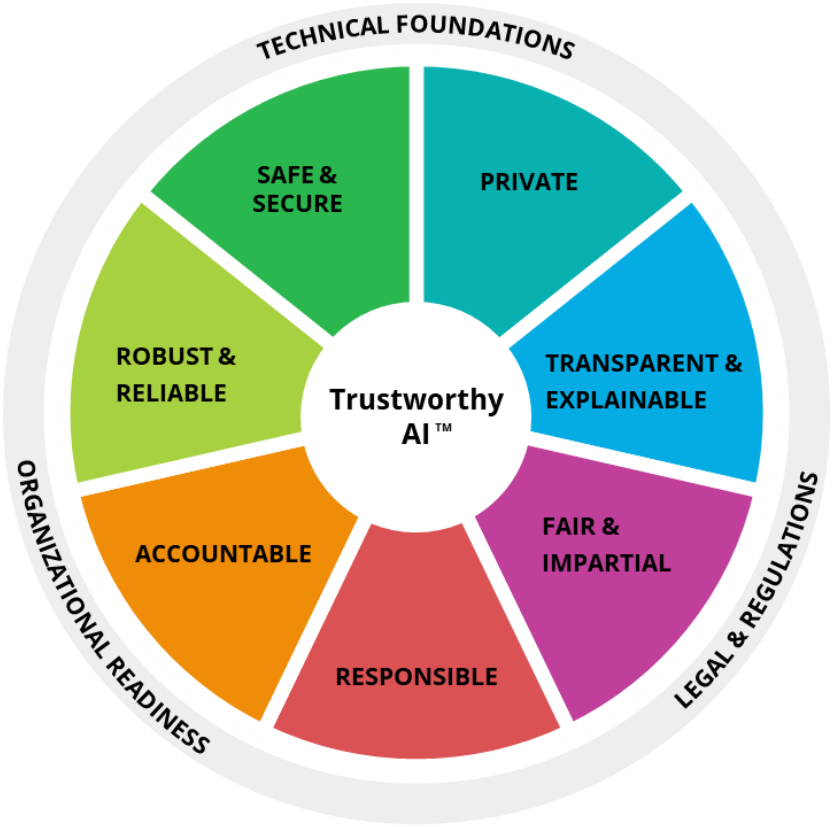

Privacy: Responsible handling of personal and sensitive data is critical throughout the AI lifecycle. With evolving regulations like GDPR and CCPA, strong privacy practices protect organizations from compliance risks and build trust with users.

Transparency and Explainability: Clear, explainable AI decisions make systems easier to audit and govern, which is critical for compliance and stakeholder confidence.

Fairness: Fair AI prevents bias and discrimination, protecting organizations from reputational harm and regulatory penalties. It also promotes trust and acceptance among stakeholders.

Responsibility: Responsible AI ensures appropriate human oversight and intentional use of AI systems. Clear boundaries and use of guardrails help organizations mitigate unintended consequences while aligning AI outcomes with organizational values and societal expectations.

Accountability: Clear ownership and oversight of AI systems enable organizations to respond effectively to incidents and maintain compliance. Defined accountability structures and escalation protocols strengthen governance and build confidence in AI-driven decisions.

Robustness and Reliability: AI systems must perform consistently across different scenarios and deliver dependable output over time. Robust and reliable AI reduces the risk of failures or errors, ensuring operational continuity and maintaining stakeholder confidence.

Safety and Security: AI systems must be safeguarded against adversarial attacks, misuse, and unintended harm. Ensuring safety and security is essential as AI becomes integral to business operations, helping prevent vulnerabilities and protect organizational assets.

Spotlight: Hallucinations

When AI hallucinates, it undermines trust in enterprise systems and can lead to incorrect decisions, insecure code, or non-compliant outputs. Unlike traditional software defects, hallucinations are unpredictable and can silently propagate across documentation, workflows, and knowledge bases at scale. CIOs must view hallucinations as an enterprise-level reliability and risk challenge, requiring strong guardrails, curated knowledge sources, and human-in-the-loop oversight. Ensuring AI-generated output is governed, validated, and auditable is essential to maintaining confidence and safeguarding business operations.

So What? Practical Implications for CIOs

Trustworthy AI provides actionable guidance for organizations seeking to drive value from AI while managing risks. For CIOs, this means embedding trustworthiness into AI initiatives to enhance organizational resilience, support regulatory compliance, and foster a culture of responsible innovation.

Some practical suggestions include:

- **Establishing an AI governance board** to oversee ethical and risk-related aspects of AI projects.
- **Conducting regular bias and fairness audits** of AI models to identify and mitigate potential issues.
- **Implementing transparent documentation practices** to ensure AI decision-making can be explained to stakeholders.
- **Building cross-functional teams** that include legal, compliance, and data privacy experts to address the full spectrum of AI risks.
- **Investing in employee training** to build awareness of AI ethics and governance best practices.

Embracing Trustworthy AI: Examples

Deloitte's thought leadership highlights several examples where organizations can operationalize the Trustworthy AI framework:

- **Financial services firms** can adopt AI fairness audits to ensure credit scoring algorithms do not inadvertently discriminate against historically underrepresented groups.
- **Healthcare organizations** can leverage explainable AI to improve transparency in diagnostic tools, supporting both regulatory compliance and patient trust.

- **Public sector agencies** can establish cross-functional ethics committees to oversee the deployment of AI in sensitive areas such as social services and law enforcement.

In conclusion, the Trustworthy AI framework offers CIOs a structured approach to embedding governance, ethics, and risk management into every stage of the AI lifecycle. By considering practical steps such as establishing governance boards, integrating AI risk into enterprise frameworks, and fostering a culture of responsible AI, organizations may be better positioned to harness AI's benefits while maintaining trust and compliance. Real-world examples demonstrate that these principles are not just theoretical, they are already being put into practice across industries to drive sustainable, trustworthy AI adoption.

Chapter 4: Defining and Measuring AI Value

Throughout this paper we have outlined steps technology leaders can take to scale AI from experimentation and proof-of-concepts to full production. But scaling AI is only part of the challenge. To ensure initiatives deliver value, organizations need to define realistic business and technology objectives supported by quantitative metrics. These metrics should cover growth, profitability, customer experience, productivity, and operational efficiency. Ongoing measurement ensures AI initiatives deliver the desired impact and remain aligned to business objectives over time. Examples of such metrics are outlined in the table below.

Metrics for Evaluating AI

Growth

Sales Growth
Conversion Rate, Average Order Value (AOV), Customer Acquisition

Market Expansion
Market Penetration, Customer Acquisition Rate

Financial Performance
Revenue Growth, ROI

Profitability

Profitability & Margin Improvement
SG&A, COGS/M as % of Revenue, ROA

Stakeholder Value
EPS, EBITDA, Revenue CAGR

Product, Portfolio & Pipeline Performance
Profitability, P&Ls, Margins

Experience

Customer / Employee Satisfaction
NPS, CSAT, eNPS, Customer Acquisition/Retention Rate

Brand Loyalty & Customer Value
CLTV, Foot Traffic, AOV

Operational Efficiency
Resource Efficiency, Time to Market

Technology

Feature/Capability Adoption
DAU / MAU, tasks per use, % adoption

Cycle Time Reduction, Dev Velocity and Throughput
Incident triage time, percentage of code generated by AI, defect density of AI-generated code

Model / Use Case Effectiveness (Gen AI)
Task Success Rate, Hallucination Rate, Toxicity/PII Leakage Rate

Managing Costs in the Era of Generative AI

As generative AI technologies advance at a rapid pace, CIOs face the dual challenge of enabling innovation while maintaining fiscal discipline. The pressure to invest in innovative models and infrastructure can lead to escalating costs, making it essential to adopt strategic approaches for cost management.

To avoid uncontrolled costs, CIOs should:

- **Adopt a Modular Approach by** building technology architecture as a series of flexible, interchangeable modules. This enables incremental upgrades rather than wholesale replacements, minimizing sunk costs and ensuring compatibility with future innovations.
- **Consider Hybrid Cloud and On-Premises solutions** to scale resources up or down as needed while balancing capital investments and operational expenditures.
- **Ensure that AI expenditures have a clear Return on Investment (ROI) and measurable business outcomes.**
- **Prioritize Low-barrier High-Impact Use Cases** that drive efficiencies and savings that can be reinvested in innovation to unlock more transformational and differentiating opportunities.
- **Foster a Cost-Aware** Culture by encouraging engineering and development teams to consider the financial implications of AI solutions such as licensing, infrastructure, LLM cost-per-token, volatile costs, and GPU scarcity.
- **Budget for Exploratory Learning and Iteration** to build employee knowledge and capabilities that will reduce risking for larger, transformational AI deployments.

Charting a Course Forward

Technology organizations play a pivotal role in unlocking the full potential of artificial intelligence. As AI moves from experimentation to scaled enterprise deployment, technology leaders must ensure their organizations have the right **talent, technology foundations, operating model**, and **governance** structures to drive sustainable value.

As a technology leader, you are responsible for charting a clear, value-driven course forward. This means shifting from isolated AI enthusiasm to deliberate, enterprise-grade execution. The path ahead requires strengthening foundational capabilities and reshaping how people, processes, and platforms work together in a hybrid human–AI environment. Without clear leadership, shadow AI will proliferate, preventing organizations from fully capitalizing on AI driven opportunities and introducing risks across data privacy, regulatory compliance, intellectual property, and the reliability of AI outputs.

To move confidently beyond pilots and build durable competitive advantage, consider the following priorities:

1. **Refresh your Technology Strategy with AI at the core.**
   Avoid standalone AI strategies. Ensure AI investments are embedded within the enterprise technology strategy, aligning architecture, talent, and governance with business priorities.

2. **Assess your end-to-end technology, data, and operations landscape.**
   Identify barriers to scaling AI, including data fragmentation, legacy interfaces, API complexity, and the maturity of operational processes.

3. **Evolve your technology operating model for hybrid human-AI collaboration.**
   Modernize roles, workflows, and team structures to support new ways of working. Ensure teams have the skills and processes to deliver AI solutions at scale.

4. **Leverage embedded AI capabilities while maintaining architectural discipline.**
   Use AI features within existing platforms where the ROI is compelling, while recognizing that no single platform will meet every need. Sound architecture, engineering rigor, and clear governance remain essential to enable flexibility and long-term scalability.

5. **Invest in talent development and organizational change management.**
   Scaling AI depends as much on people as it does on technology. Develop employees' fluency in AI, reshape training pathways, and build a culture that supports responsible AI adoption and continuous learning. Early wins should be used strategically to build momentum and organizational confidence.

By taking these steps, technology leaders can ensure their AI initiatives are not only technically sound but also strategically aligned, ethically responsible, and positioned to deliver sustained business value. The organizations that succeed will be those that treat AI not as a collection of tools but as a catalyst for re-architecting their technology foundations, elevating their people, and reshaping how work gets done in the digital era.

Contact Us

David Noseworthy
Partner,
Technology Strategy & Transformation
Partner, Deloitte Canada
dnoseworthy@deloitte.ca

Bruce Derraugh
Partner, National Practice Leader,
Technology Strategy & Transformation
Partner, Deloitte Canada
bderraugh@deloitte.ca

Authors

Nathan Boey
Manager
Technology Strategy & Transformation

Hailey Moon
Business Analyst
Technology Strategy & Transformation

David Noseworthy
Partner
Technology Strategy & Transformation

Cedric Rossi
Manager
Technology Strategy & Transformation

Pranav Saigaonkar
Consultant
Technology Strategy & Transformation

Contributors

Jonas Yovoyan
Ryan Zuro
Jean-Francois Haeck

Endnotes

[1] Nathan Furr, Andrew Shipilov, ***"Beware the AI Experimentation Trap,"*** *Harvard Business Review*, August 28, 2025.

[2] Nestor Maslej et al., "***The AI Index 2025 Annual Report,"*** AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2025, Chapter 4, page 246.

[3] Nestor Maslej et al., "***The AI Index 2025 Annual Report,"*** AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2025, Chapter 4, page 269.

[4] "***Reshaping expertise: How GenAI is Changing Knowledge Work,"*** Deloitte Canada, n.d., https://www.deloitte.com/ca/en/services/consulting/perspectives/reshaping-expertise-using-gen-ai.html.

[5] Steve Elliott, Meagan Lettau, Karen Pastakia, Nic Scoble-Williams, Stuart Scotis, "***Work Reworked: What it Takes to Win in the Age of Agentic AI,"*** Deloitte, n.d., https://www.deloitte.com/content/dam/assets-shared/docs/deloitte-human-agentic-ai.pdf.

[6] Ashish Verma, Prakul Sharma, Parth Patwari, Alfons Buxo Ferrer, Diana Kearns-Manolatos, and Ahmed Alibage, ***"Four Data and Model Quality Challenges Tied to Generative AI,"*** *Deloitte Insights*, February 6, 2025, https://www.deloitte.com/us/en/insights/topics/digital-transformation/data-integrity-in-ai-engineering.html.

[7] Tim Smith, Faruk Muratovic, Bill Briggs, and Diana Kearns-Manolatos, ***"Three Ways to Approach Legacy Tech Modernization with AI,"*** *Deloitte Insights*, June 6, 2025, https://www.deloitte.com/us/en/insights/topics/digital-transformation/legacy-system-modernization.html?id=us:2el:3dp:wsjspon:awa:WSJCIO:2025:WSJFY26.

[8] ***"Walmart's Element: A machine learning platform like no other,"*** *Walmart*, March 14, 2024, https://tech.walmart.com/content/walmart-global-tech/en_us/blog/post/walmarts-element-a-machine-learning-platform-like-no-other.html.

[9] Sravana Karnati, ***"From models to agents: A new era of intelligent systems at Walmart,"*** *Walmart*, August 29, 2025, https://tech.walmart.com/content/walmart-global-tech/en_us/blog/post/wibey-announcement.html.

[10] "***Trustworthy AI",*** Deloitte Canada, n.d., https://www.deloitte.com/ca/en/services/consulting/services/ai-and-data/trustworthy-ai.html.

About Deloitte Canada

At Deloitte, our Purpose is to make an impact that matters. We exist to inspire and help our people, organizations, communities, and countries to thrive. Our work underpins a prosperous society where people can find meaning and opportunity. It builds consumer and business confidence, empowers organizations to find imaginative ways of deploying capital, enables fair, trusted, and functioning social and economic institutions, and allows our friends, families, and communities to enjoy the quality of life that comes with a sustainable future. And as the largest Canadian-owned and operated professional services firm in our country, we are proud to work alongside our clients to make a positive impact for all Canadians.

Deloitte provides industry-leading consulting, audit and assurance, tax, advisory and managed services to nearly 90% of the Fortune Global 500® and thousands of private companies. We bring together world-class capabilities, insights, and services to address clients' most complex business challenges.