WHITE PAPER

The AI blueprint: A leader's guide for organizational trust and ROI during rapid change

Contents

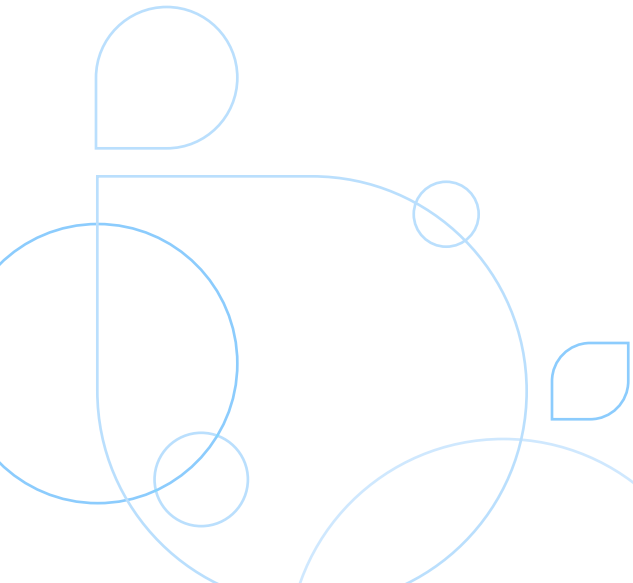

Forward:
Leading with trust and confidence in the AI era

How are you using AI? Can you show value? Are your AI outputs explainable to regulators? You have probably encountered these questions more than once. Executives aren't just curious about AI anymore – they are positioning their organizations for the technology's future. The problem? Most organizations aren't ready and struggle to answer these questions confidently.

AI is rapidly emerging as more than just a technological advancement. It represents a fundamental shift as transformative as the steam engine was to the Industrial Revolution. This profound time in human history is poised to reshape every aspect of life.

Still, only a few organizations are building foundations to scale AI effectively and responsibly. Others are rushing without a plan – missing the infrastructure, governance and cross-functional collaboration to make AI work.

Noting that trust is central to AI success and adoption, **74% of companies struggle to achieve and scale value from their AI initiatives – a sign of the growing trust gap.**[1]

AI success is built on a solid foundation that does not sacrifice control and trust. When trust in AI systems is compromised – whether due to bias, inaccuracy or security vulnerabilities – adoption is hindered, and regulatory and reputational risks escalate. The projected productivity gains and financial returns either fail to materialize or prove unsustainable. Trust is imperative, and leaders are now reviewing a trust index to evaluate the value of AI implementation, operational efficiency, innovation potential and ethical oversight.

Conversely, building trust through transparent and responsible AI practices fosters greater acceptance, encourages wider integration and directly drives higher productivity and return on investment (ROI). Decision makers expect their AI initiatives to deliver at least 3x ROI.[2] Leaders need to prioritize projects that can demonstrate clear ROI, ideally within the first year, as these early successes are crucial for building organizational confidence and securing continued support for broader AI initiatives.

The new AI era has spurred acceleration with organizations of all sizes experimenting in the early stages of large language models (LLMs) to embracing enterprise intelligent decision making. McKinsey research estimates a substantial long-term AI opportunity, projecting US$4.4 trillion in added productivity growth potential from various corporate use cases.[3] Some organizations are already advancing into agentic workflows to improve productivity and accelerate operations. A select few are exploring quantum AI in life sciences, banking and material science.

Despite this widespread recognition and enthusiasm, a notable maturity paradox exists in the current state of AI adoption. While nearly all companies report investing in AI, a mere 16% of leaders classify their organizations as "mature" in AI deployment, defined by AI being fully integrated into workflows and driving substantial business outcomes.[4] This highlights a considerable gap between initial investment and successful, at-scale integration, suggesting that many organizations are still in exploratory or fragmented implementation phases.

> "The AI Trust Index was developed to quantify a critical but often overlooked dimension: how trust influences the success of AI initiatives. It provides organizations with a clear benchmark to assess how trust drives – or stalls – AI adoption and impact. In a landscape flooded with hype, trust is the deciding factor between scalable transformations and stalled pilots."
>
> **Chris Lee Marshall,**
> Research Vice President,
> IDC

Implementing AI extends beyond deploying technologies. It requires strategic alignment and readiness at both cultural and operational levels – connecting technology, processes and people. It requires a new mindset and vision to see AI as a strategic enabler for new business practices and skills development. Therefore, the ownership of AI success must reside at the highest levels of the organization, involving the C-suite and cross-functional leadership, rather than being delegated solely to the IT department.

The pace of AI is relentless, and leaders must act with urgency and clarity. In this paper, we'll explore a proven methodology – the AI blueprint – to navigate emerging AI technology with clarity. One that is structured, repeatable and not an ad hoc experiment. You'll learn how to approach AI at scale – from vision to execution – with trust, confidence and ROI.

> "To navigate the chaos of AI adoption, organizations need a north star – a repeatable, modular AI blueprint built on trust, scale and real-world impact."
>
> **Manisha Khanna**, Global Technology, AI and Emerging Tech, SAS

The agentic AI imperative: A new productivity accelerator

Enterprise leaders are acutely focused on agentic AI, recognizing its emergence as a pivotal driver of future productivity and strategic advantage. The rapid ascent of generative AI and the evolution into agentic AI systems – designed to understand inputs, make decisions and act autonomously or semi-autonomously without direct human intervention to achieve specific goals – have collectively ignited a frenzy. It is propelling an unprecedented surge in enterprise technology investment and strategic priorities.

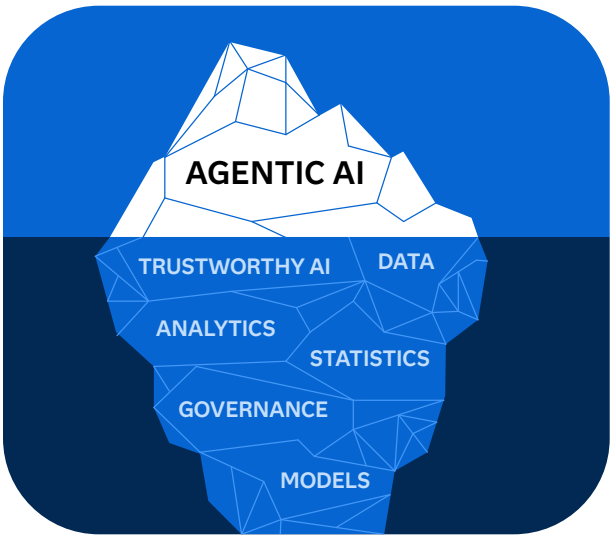

However, this pursuit of the next productivity achievement with agentic AI masks a critical reality. Gartner predicts more than 40% of agentic AI projects will be canceled by the end of 2027, citing escalating costs, unclear business value or inadequate risk controls.[5] Leaders must think about agentic AI metaphorically as an iceberg. While its advanced capabilities represent the visible peak, its true power and sustainability are built upon decades of foundational work in mastering data engineering, statistics, analytics, modeling and the establishment of trustworthy AI systems.

A sound, holistic and achievable AI strategy is critical. A piecemeal approach will fall short of long-term success.

The need for a unified data and AI platform to scale

To realize value from your AI investment, leaders must break down silos between tools, teams and data sources. They must be prepared for emerging technology like agentic AI.

Fragmented data and AI technologies, siloed systems and a lack of seamless data flow are pervasive barriers that prevent AI initiatives from scaling beyond isolated departmental projects – severely limiting their scope, impact and ultimately ROI. With a unified, interoperable data and AI platform, AI models can access the comprehensive data they need to generate accurate results, and AI teams can enable holistic insights across the enterprise.

Furthermore, the rise of edge computing necessitates having a data and AI platform capable of managing and analyzing data in environments with less reliable connectivity, extending AI capabilities beyond centralized data centers and enhancing local decision making and system resilience for distributed operations.

SAS Viya is an end-to-end data and AI platform. With Viya, organizations gain consistency, transparency and speed across the AI life cycle. The platform enables organizations to build, govern and scale AI in an open, modular, cloud-native environment that works across diverse industries and regulatory landscapes.

Balancing AI governance with innovation

Emerging technologies like agentic AI introduce new governance complexities: explainability challenges, autonomous decisions and dynamic learning.

Governance frameworks must be pliable to include advanced monitoring, proactive bias detection and ethical guardrails tailored to these AI systems.

Transparency demands robust documentation, enhanced model cards for transparency and continuous audit trails to provide insight into AI-generated content and agent behaviors. For GenAI, this requires specific monitoring for issues such as loss of context, factual accuracy drift, hallucination or tone alteration (abusive/rude) in model outputs.

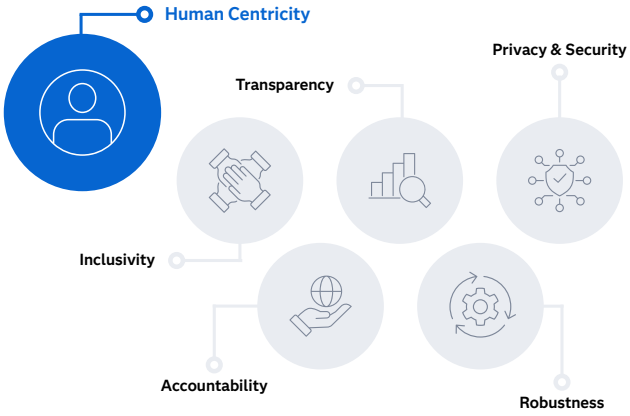

Technology does not exist in a vacuum. Ethical dilemmas can present tensions that should be navigated in the most productive and least harmful ways. Managing challenges in a trustworthy manner requires an unwavering set of principles rooted in proven, effective strategies – from ideation through development and deployment. This includes human centricity, inclusivity, accountability, transparency, robustness, privacy and security.

As organizations adopt agentic AI – where systems act with greater independence and autonomy – these principles become even more critical. They ensure that intelligent agents operate within clear ethical boundaries, remain aligned with human values and are accountable in environments where real-time decisioning and adaptation can amplify both opportunity and risk. The data and AI platform of choice must accommodate this need and, in return, close the trust gap eroding organizational success.

"AI governance is crucial as the impact of AI grows. Ensuring AI systems are trustworthy, explainable and ethical will build organizational trust and value."

Frank van Praat, Partner, Responsible AI, KPMG in the Netherlands

Introducing a future-ready **AI blueprint**

Leading organizations realize that navigating the complexities and opportunities in the AI era demands more than ad-hoc projects – it requires a comprehensive, future-ready AI blueprint. This blueprint serves as a strategic guide, ensuring that AI initiatives are aligned with overarching business objectives and poised for scalable, sustainable impact. This transformative mindset from leadership is crucial, moving beyond mere experimentation to deliberate, enterprise-wide integration. Think of the AI blueprint as a tailored methodology. It's the guiding light across **people, process, technology** and data – bringing clarity to chaos.

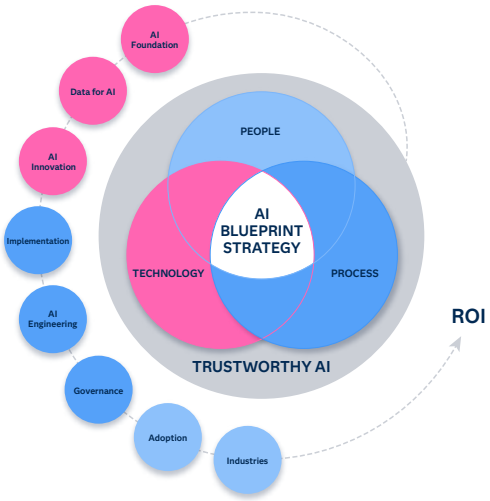

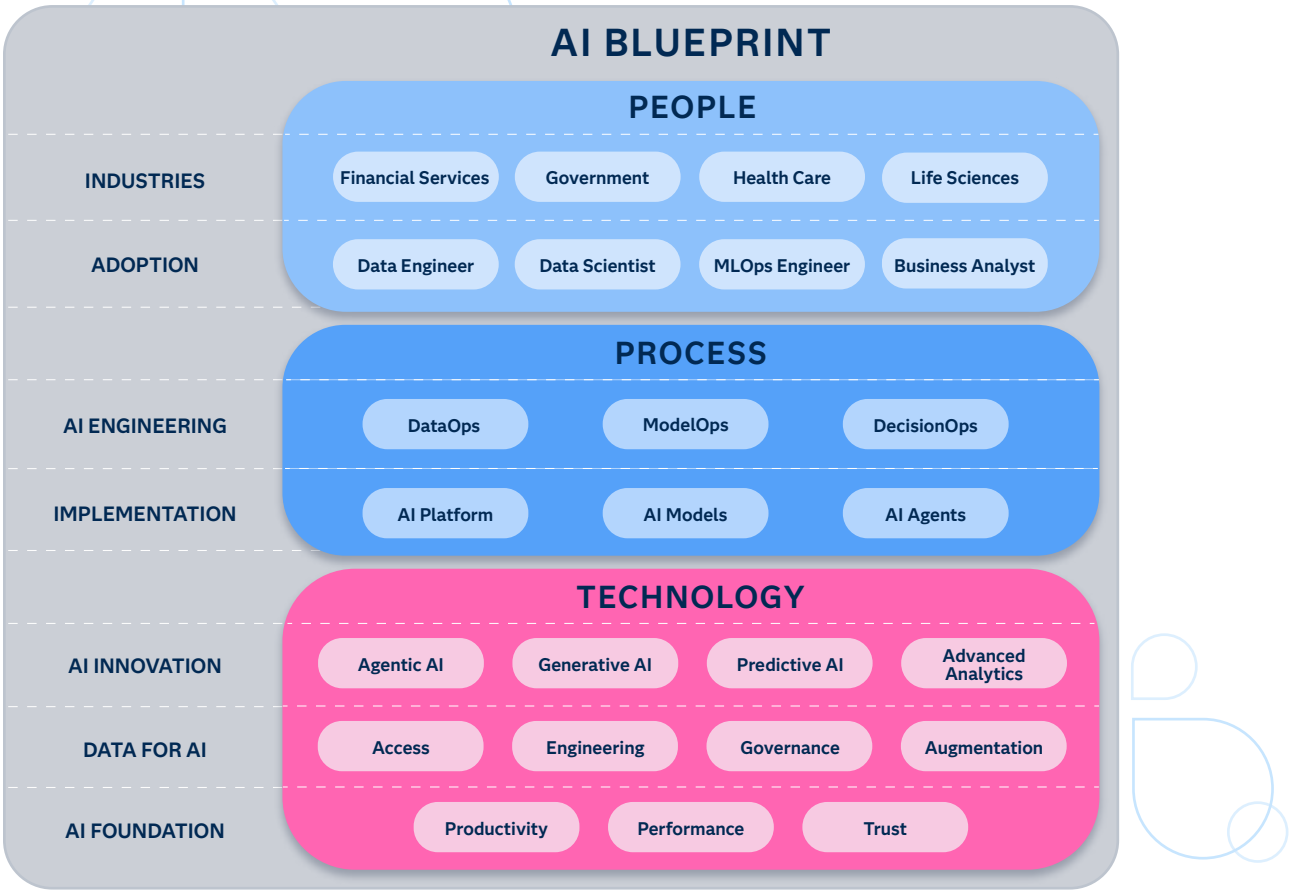

> The AI blueprint transforms vision into execution, cutting through complexity at scale. By embracing emerging technologies, like agentic AI, humans and machines collaborate dynamically, unlocking new levels of productivity and creativity. To stay competitive, organizations must embed this shift directly into their AI strategies."
>
> **Manisha Khanna**, Global Technology, AI and Emerging Tech, SAS

Pillar 1: Designing flexible, scalable technology for tomorrow

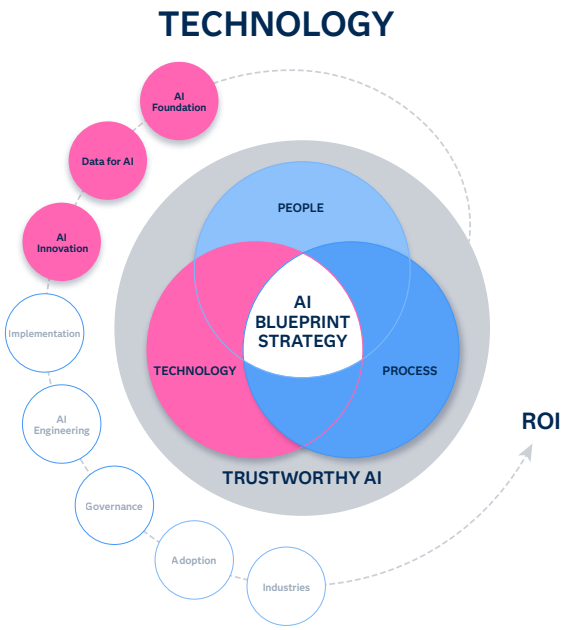

The formidable promise of AI hinges entirely on the quality and integrity of the data that powers it. Organizations must embrace modern data management to ensure data accuracy, trustworthiness and responsible AI. Often, failed AI is directly correlated to bad data. It's not the algorithm.

Trusted data at the foundation

High-quality data is accurate, complete, consistent and timely. It's data you can trust to make critical decisions. Poor quality data without proper data processing can lead to flawed insights, misguided strategies and costly mistakes. Bad data can expose an organization to ethical missteps and regulatory risk. It's important to establish guidelines for responsible AI use and data flow, including detection and mitigation of bias in training data. Trusted data is a nonnegotiable foundation of the AI blueprint and is at the forefront of strategic planning to use AI.

Modern data management

The reliability and utility of AI outputs are directly impacted by the quality and accessibility of this data. A critical step for successful data integration involves addressing existing data silos and fragmentation and establishing robust data governance.

- **Access data.** Simplify access to clean, structured and unstructured data from any source (regardless of its size, format or complexity) for a unified view of all data. Data should be accessible to both business and technical users.
- **Prepare data.** Build robust data pipelines to ingest, transform and serve the highest quality data to analytic and AI models. Embedded profiling, semantic tagging and lineage tools ensure data is AI-ready.
- **Govern data.** Governance at every layer – from profiling to lineage to impact analysis. Use a data management process that is open, scalable and natively tied into the data and AI life cycle, enabling trust end-to-end without complexity with interoperability via open APIs and metadata standards. This enables flexible, future-ready architecture.

> "With tools like SAS Information Catalog within SAS Viya, organizations can automate governance and metadata management, ensuring data is traceable and compliant at every stage. Bias detection capabilities embedded early help flag and correct issues before models are trained, making ethical AI an integral part of data preparation, not an afterthought."
>
> **Dan Soceanu**, Sr. Product Manager, Data and AI, SAS

Augmentation through synthetic data

Emerging technologies like synthetic data are a useful approach for AI innovation. It is artificially generated data based on real data to supplement or fill gaps. Synthetic data can ensure that there's proper representation across all groups.

Synthetic data can also solve data management issues that have challenged organizations for years. It saves time spent acquiring data, preparing data and cleaning data for AI development efforts. Using synthetic data to solve problems is a smart approach to tackle data management challenges and must be a staple of AI strategy.

Moreover, the emergence of agentic AI systems introduces new complexities to data ownership and privacy. The multi-agent interactions mean that data is dynamically modified multiple times during use, moving beyond traditional, static notions of data privacy and control. This necessitates proactive anticipation and preparation for a future where data ownership, provenance and control become even more complex and critical, requiring multi-disciplinary approaches involving legal, compliance, cybersecurity and technical teams.

Unlocking insight through interactive AI

Organizations must approach data to move beyond static reports and interact with data in real time through AI-assisted exploration. This allows organizations to uncover trends, patterns and outliers faster. With intuitive interfaces and natural language querying, asking questions and getting answers becomes accessible to everyone, not just analysts. This democratization of data insights ensures that all team members can engage with and benefit from the data.

Flexible, AI-first infrastructure that scales

With data as the foundation, it's time to advance further into the technology pillar of the AI blueprint. Just like data is critical to the foundation of AI, so is a flexible, AI-first infrastructure. It is built to scale, adapt and support trust at every layer. Flexibility and interoperability are the foundation of long-term AI success.

A critical observation from the current AI landscape is the shift toward more specialized AI platforms and infrastructure. Large-scale AI models, especially GenAI, have demanding computational requirements that cannot be optimally met by traditional cloud infrastructure alone. This necessitates a new class of computing hardware and software specifically designed for parallel processing, massive data throughput and rapid model training, which becomes a competitive differentiator.

The fundamental technology stack for AI comprises several interconnected components: advanced AI datacenter infrastructure (including clusters of Graphics Processing Units (GPUs) with high-bandwidth network connections, machine learning acceleration software, powerful pre-trained AI models, API services for access and the applications where AI outputs are utilized. This layered approach enables efficient development and deployment of complex AI systems.

Cloud infrastructure providers play a pivotal role in establishing the essential groundwork for AI infrastructure, serving as the backbone for large-scale AI operations and democratizing access to powerful AI capabilities.

> SAS Visual Analytics within SAS Viya allows users to visually explore data in real time, spotting trends, patterns and outliers without waiting for IT or specialized analysts. The integrated Visual Analytics copilot takes this further by providing AI-powered natural language assistance, enabling users to ask complex questions conversationally and receive instant, context-rich visualizations – making data discovery faster, more intuitive and accessible to everyone across the organization."
>
> **Manisha Khanna**, Global Technology, AI and Emerging Tech, SAS

> The emergent AI era, driven by the profound capabilities of GenAI and the transformative potential of agentic AI, is fundamentally reshaping the very bedrock of enterprise computing. Leaders now recognize that the demanding computational requirements of large-scale AI models mandate a new class of flexible, scalable and inherently trusted AI-first cloud-ready infrastructure."
>
> **Danielle Greshock**, Worldwide Director, ISV Partner Solutions Architecture, Amazon Web Services (AWS)

Accelerating the right models

To accelerate AI adoption and innovation, AI access must be simplified for the organization to break down infrastructure challenges. Legacy systems need to work with groundbreaking AI technology. Depending on the business question at hand, not all models need to be built from scratch. The right AI strategy balances prebuilt intelligence with the flexibility to customize.

Therefore, it might be appropriate to use ready-made AI models to accelerate time-to-value in regulated and data-intensive sectors. These models are lightweight, prepackaged, industry-specific offerings designed to accelerate AI adoption and are easily implemented across a variety of IT infrastructures.

Deterministic models are becoming even more valuable in the generative era, offering predictable outcomes, robust governance and traceable logic that complements the creativity and fluency of LLMs.

To remain agile, organizations should adopt an open, model-agnostic approach – one that supports integrating LLMs from any provider.

Incorporating digital twins into problem-solving

The data and AI platform must be flexible to accommodate all technology now and in the future. By combining digital twins with AI, organizations can create intelligent systems that enhance productivity, safety and efficiency.

Digital twin technology is transforming industries by enabling real-time simulations, predictive insights and operational efficiencies. Its applications span manufacturing, public sector projects, finance and health care – improving decision making, reducing risk and driving continuous improvement. By creating digital replicas of physical systems, organizations can optimize operations, stress test rare events, reduce waste and plan for the future.

Empowering agents to act with intelligence and integrity

Additionally, architecture should be flexible to accommodate agentic capabilities. Agentic AI brings together automation, decision logic and human judgment. AI agents can drive continuous, real-time decisioning in environments like fraud detection, customer engagement and logistics.

> "In the agentic AI era, strategic advantage lies not in choosing a single provider, but in designing an architecture that can evolve with the market. Flexibility is the safeguard against lock-in, and the engine for continuous innovation."
>
> **Marinela Profi**, Global AI and GenAI Strategy Lead, SAS

Organizations will need to build, test and deploy AI-driven decisions that combine business rules, real-time data and machine learning logic. Using agents that are designed to keep humans-in-the-loop, offering transparency and override options to ensure alignment with ethics, regulations and strategic goals. Technology like agentic AI balances automation with oversight, unlocking productivity while preserving trust.

A flexible data and AI platform, adaptable models and intelligent agents are essential components – but they only unlock real value when paired with productivity-enhancing tools like a copilot. AI copilots emerge as a natural extension of agentic AI: task-specific assistants that augment human capabilities, guide decision making and streamline complex workflows. Copilots accelerate and simplify the AI life cycle by streamlining how users interact with data, build models and deploy insights. To drive meaningful outcomes, copilots must be embedded into day-to-day workflows for data and AI through governed, repeatable processes that scale across the organization. It's not just about implementing AI. It's about operationalizing it with efficiency, consistency and human-in-the-loop intelligence.

> SAS Viya Copilot marks a significant milestone in our collaborative journey with SAS to leverage the power of AI and data analytics. Integrating an AI-driven conversational assistant directly into the SAS Viya platform empowers users to achieve improved levels of productivity and efficiency. This innovation democratizes access to advanced analytics, enabling users to make more informed decisions, streamline their workflows and fully harness the potential of their data."
>
> **Lorraine Bardeen**, Corporate Vice President and CTO, Commercial Solution Areas, Microsoft

Preparing for what's next: Quantum AI

Quantum AI has a lot of potential for innovation gains. Quantum AI brings together quantum computing and AI to develop new algorithms, models and systems that can process complex data, train complicated machine learning models and solve problems considered challenging or impossible with traditional computers. Quantum AI will use hybrid architectures, which include both quantum and traditional, or classical, computing. Multiple industries will benefit from the advanced computations of quantum AI, including life sciences for enhanced drug discovery, financial services for better risk management and manufacturing for improved materials and optimized processes.

Leaders need to keep pace with the speed of change and be ready with architecture that adapts quickly.

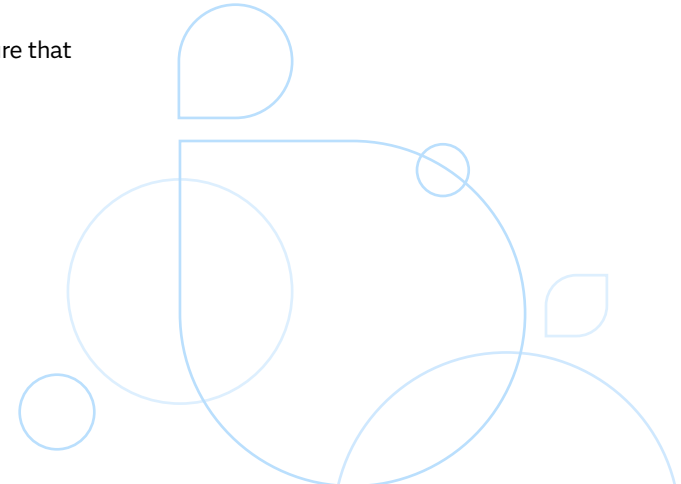

Pillar 2:
Driving results with smart AI processes

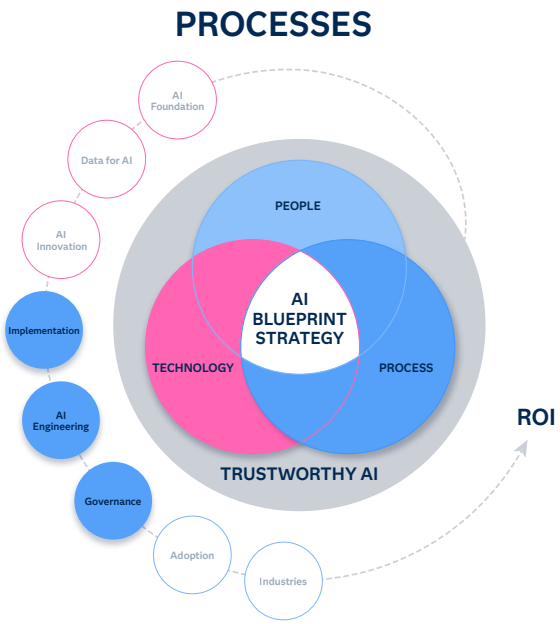

The strategic deployment of AI is not merely about having the right technology; it is fundamentally about implementing smart processes that ensure AI initiatives deliver tangible business results, measure their impact effectively and drive sustained value. AI adoption in regulated industries requires embedding compliance and risk management deeply into processes.

This means that IT and AI teams must not only focus on building models and deploying technology but also rigorously assess and optimize their operational workflows to ensure they meet regulatory standards and enable transparent, auditable AI practices.

Achieving AI-driven innovation without sacrificing governance is the key challenge, and it demands close collaboration between technology, compliance and business teams to design processes that are both agile and compliant.

AI Engineering: Operational discipline for compliance and agility

Processes are imperative to an efficient and trustworthy AI life cycle. AI engineering encompasses XOps disciplines like DataOps, ModelOps and DecisionOps. These processes must be considered for AI implementation for automation and productivity. They provide the structure needed to manage complexity and to scale AI with transparency and trust.

They are especially relevant for agentic AI, as they go beyond simple prompting of GenAI models to encompass the full life cycle management of these advanced AI systems. From data ingestion and model development to deployment, monitoring and governance – the processes ensure both agility and compliance.

They also enable agile development methodologies, characterized by iterative test-and-learn cycles rooted in user feedback. This is crucial for effective AI deployment, allowing for rapid adaptation and refinement.

Critical roles like AI engineering require high AI maturity and deep knowledge of ModelOps. The process helps move newly built AI models through the validation, testing and deployment phases as quickly as possible while delivering quality results for enterprise decisioning. ModelOps helps manage through governance activities that test, validate, schedule, run, monitor and automate decisions at scale. ModelOps also involves architecting AI models and LLMs to be scalable and cost-efficient, managing integration endpoints, performance testing and continuous integration/continuous deployment (CI/CD) capabilities.

Other important processes include DataOps, which focuses on collaboration, integration and automating data management to deliver high-quality, trusted data to AI models. It involves practices like prompt engineering and management for LLMs.

DecisionOps extends these principles to the entire decision-making process, ensuring that AI-driven decisions are managed, governed and optimized consistently. This involves using AI models, AI agents and rules in decisions to support their operationalizing more quickly, consistently and reliably. Automating these monitoring capabilities is essential, along with observability and reporting for application owners to understand user interaction.

For enterprises implementing or building their own models, there is an increasing demand for monitoring for ethical use, abuse prevention and intellectual property (IP) compliance assurance based on the model's provenance and training data.

> *The integrated environment of SAS Viya supports XOps disciplines with compliance-ready features, enabling rapid, yet safe scaling. SAS Viya streamlines DataOps, ModelOps and DecisionOps with built-in audit trails and version control – helping teams maintain compliance while accelerating AI delivery.*

Shared transparency for trust and accountability

With any emerging AI technology, a common concern is whether you can trust the outputs. AI models must be continuously maintained and monitored to ensure optimal performance and be accessible to regulators, compliance teams, business users and data scientists.

Model cards and centralized dashboards offer role-specific insights, bridging gaps and fostering collaboration. Using them in processes to ensure models are securely contained within environments that meet institutional standards for data handling, oversight and auditability.

> With SAS Model Cards, organizations gain a single, accessible source of truth for model documentation, performance and compliance – making governance easier across roles and teams."
>
> **Kristi Boyd**, Trustworthy AI Specialist, Data Ethics Practice, SAS

Smart AI processes to keep pace with rapid change

Smart AI processes embed governance, compliance and operational discipline into every stage of the AI life cycle, turning potential into consistent performance. XOps practices provide the structure needed to manage complexity and to scale AI with transparency and trust. Adaptive governance frameworks ensure AI remains ethical, explainable and aligned with regulatory requirements – even as GenAI and agentic AI evolve.

Effective processes require empowered, skilled people working collaboratively to monitor, maintain and optimize AI outcomes.

> Balancing AI-driven innovation with governance requires end-to-end life cycle management through robust XOps – integrating DataOps, ModelOps and DecisionOps. For agentic AI, these disciplines, combined with AI-first infrastructure, deliver the agility, observability and control needed for real-time decisioning."
>
> **Mary Osborne,** Platform Product Strategy, SAS

> "AI Engineering, including ModelOps, is the power behind enterprise AI. It helps turn ideas and experiments into real-world results; getting models and LLMs from development to production quickly, at scale, and with the necessary governance to create actual beneficial results."
>
> **Colby Hoke**, SAS Model Manager and Open Source Integration, SAS

Pillar 3:
Enabling people to power AI success

It's a fallacy that AI will replace humans. In fact, it's people who activate, govern and evolve AI capabilities. AI is increasingly viewed as an intelligent assistant and copilot designed to augment human creativity by automating repetitive tasks and providing new insights, allowing individuals to focus on higher-value work. People are not just participants in the AI life cycle – they are orchestrators, overseers and amplifiers of AI's strategic value, as previously stated about human-in-the-loop.

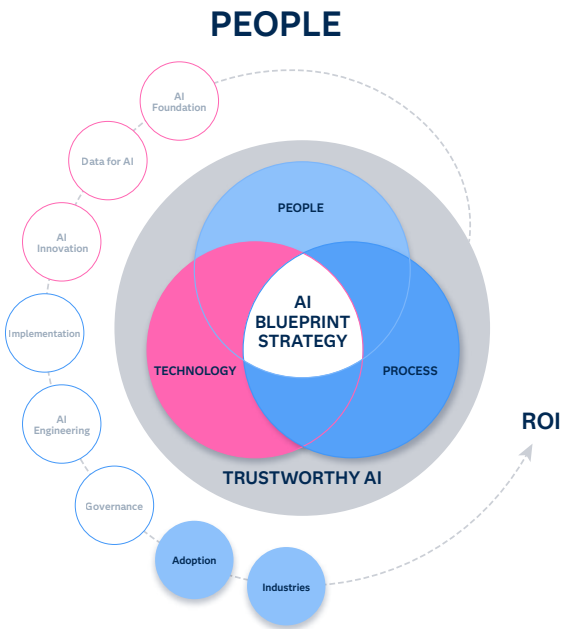

Gartner predicts that by 2026, more than 100 million people will engage with robocolleagues, signaling a widespread shift towards human-AI collaboration.[6]

As agentic AI systems take on more autonomous roles, the human element becomes even more critical in guiding outcomes, enforcing ethics and adapting AI to real-world complexity.

Key roles in the AI ecosystem

People are shifting from passive users to active supervisors and cocreators of intelligent systems. The integration of AI necessitates a fundamental transformation of the workforce, involving not just new skill sets but also the reshaping of roles, the establishment of new organizational structures, and new ways of working.

There are foundational roles in a full data and AI life cycle. These roles remain critical – even more so – in the era of GenAI and agentic AI. But their scope and impact are evolving to meet the demands of greater autonomy, scale and real-time decisioning:

- **Data engineer.** Prepares and manages data pipelines. Supports more complex and multimodal pipelines (e.g., text, image and structured data).
- **Data scientist.** Develops models and drives innovation. Their role shifts from just building models to designing decisions with intelligent, explainable agents.
- **MLOps engineer.** Automates deployment, monitoring and retraining of models. Their role expands to AI engineering – managing and orchestrating AI agents.
- **Business analyst.** Bridges AI insights into strategic business decisions. Analysts help define the business logic and constraints that guide generative and agentic systems.

Emerging technologies have ushered in new roles to support GenAI and agentic AI innovations, including prompt engineers, AI ethicists and AI risk leads. Human roles shape how AI learns, adapts and delivers value – especially as agents act more autonomously.

> "SAS Viya offers a shared platform where roles converge – empowering data engineers, scientists and decision makers to work together across the AI life cycle in a unified, governed environment."
>
> **Manisha Khanna**, Global Technology, AI and Emerging Tech, SAS

Breaking down barriers between technical and non-technical talent

As we have learned, technology, processes and people make AI work. The foundation of technology and processes allows people to collaborate seamlessly across the data and AI life cycle. Collaboration is essential from day one with both technical and non-technical roles. One, centralized data and AI platform eliminates black boxes allowing for transparent, explainable outputs across departments and with regulators.

Breaking down barriers in collaboration allows people to more confidently explore innovative ways to solve problems while maintaining ethical and operational guardrails.

With SAS Viya, transparency is built in – not bolted on. From day one, technical and non-technical teams can collaborate on shared workflows, using explainable models, audit trails, and role-based access to ensure clarity, trust and alignment across departments.

People-centered AI: Skills, stewardship and strategic oversight

AI is the new data literacy. AI-first organizations invest in upskilling, demystifying AI for non-technical teams and empowering responsible innovation. The rapid evolution of AI creates an urgent need for targeted skill development, addressing the anticipated talent shortage in critical new roles through proactive reskilling and upskilling programs.

An adaptable talent strategy allows roles to evolve and their AI capabilities to grow. This requires flexible and innovative workforce strategies for creating top talent, which could include reskilling, partnerships with universities and other novel methods to develop essential skills and continuous learning. All departments within the organization need to embrace this school of thought.

> *The Commonwealth AI Consortium, together with SAS, is offering tailored learning for students, professors and independent learners, enabling access to learning resources and SAS software to build AI skills for the future.*[7]

Closing the AI talent gap

- **Upskill from within.** Invest in targeted training for existing employees – especially data analysts, developers and business users.
- **Build cross-functional AI literacy.** Create blended teams that combine domain experts with technical talent.
- **Recruit for potential, not perfection.** Look beyond traditional AI degrees and experience – hire for adaptability, systems thinking and curiosity.

Humans remain essential to sustaining AI. They will continue to apply domain expertise to determine when AI models or agents no longer reflect current business conditions, regulations or ethical norms. This human involvement can trigger model retraining or intervention when automated systems miss the nuance.

Humans also reinforce accountability through regular checkpoints, escalation protocols and policy alignment to ensure autonomy never comes at the cost of control.

While employees are often more ready for AI than leaders imagine, effective change management – encompassing frequent communication, identifying champions, comprehensive training and incentivizing experimentation – is crucial for maximizing adoption and ensuring technology delivers real impact and ROI.

AI literacy is critical for preparing organizations for the future of work and ensuring that leaders can navigate transformative changes brought about by rapidly changing technology. Organizations will benefit from training employees quickly with AI skills.

Empowering people for scalable AI growth

People are critical enablers who bring the AI blueprint to life and are at the heart of AI adoption. AI is not a replacement for humans but an augmentation tool. Without active human involvement in data curation, model evaluation, feedback loops and strategic oversight, AI systems cannot achieve reliable, beneficial and trustworthy outcomes. Organizations need to invest in roles, encourage collaboration and use an infrastructure that facilitates sustained innovation.

Furthermore, proactive workforce transformation extends beyond just training. The strategic focus should be on redesigning work to uplift unique human capabilities – creativity, critical thinking and emotional intelligence – that AI cannot replicate. It should not be about solely automating tasks. This long-term, iterative process requires sustained leadership commitment and a willingness to rethink traditional organizational paradigms.

> "Human feedback loops are supported in SAS Viya with automated retraining triggers, alerts and governance features – ensuring AI systems stay ethical, compliant and effective over time."
>
> **Kristi Boyd**, Trustworthy AI Specialist, Data Ethics Practice, SAS

Solving real-world problems with AI

Now that we have established the methodology for successful AI, we can examine how the AI blueprint translates into real work.

Identifying and prioritizing high-value AI use cases for measurable impact

The ultimate success of an AI strategy is intrinsically linked to its ability to deliver measurable business value tailored to the specific industry or domain.

This requires a sharp focus on projects that can demonstrate clear ROI within the first year, which is crucial for building organizational confidence and securing continued support for broader AI initiatives. The feasibility of a use case is more important than its perceived business value, ensuring practical implementation and avoiding costly, unachievable projects. When selecting AI use cases, line-of-business stakeholders should clearly articulate the expected tangible business benefits, asking critical questions such as:

- What specific problem is the business trying to solve?
- Who is the primary consumer of this technology?
- And, crucially, how will the impact and value of technology be measured, monitored and maintained?

> *SAS solutions are built with domain expertise at the core – helping organizations tackle industry-specific challenges with confidence, compliance and clarity.*

Furthermore, organizations with extensive AI experience measure success by business metrics (e.g., business growth, customer success, cost-efficiency) rather than solely by project volume or tasks completed.

The business objectives for AI continue to evolve as organizations improve their maturity levels on AI. The highest level of value from AI is achieved when organizations focus on customer experience, market share and resilience as their primary AI goals compared to those that prioritize cost reduction. A broader transformation and shift in perception of using AI for enabling long-term business transformation rather than short-term savings can enable organizations to use AI more strategically.

The AI blueprint in practice

Financial Services: The National Bank of Greece

The National Bank of Greece (NBG) was the first commercial bank in the country and now boasts 180 years of ongoing operations with 5.5 million retail customers. The bank plays a vital role in supporting the Greek economy, contributing to national development and promoting social transformation.

NBG is overcoming the challenge of becoming more agile and modern while improving productivity.

Key takeaways

- Established a Business Analytics Center of Excellence to integrate analytics and AI into daily operations.
- Reduced complexity of processes with the agility, speed and efficiency using SAS Viya on Azure.
- Attracted new talent in data science by allowing developers to integrate with open source.

AI blueprint applicability: Technology, process and people

"SAS Viya is cloud-ready, and the transfer of operations to the cloud is intertwined with our strategy."

Pantelis Maraveas, Assistant General Manager of Retail Segments, Bank Analytics and Liability Products, National Bank of Greece

Government and Public Sector: Malta Tax and Customs Administration

The Malta Tax and Customs Administration (MTCA) launched a modernization initiative to improve taxpayer compliance and morale. SAS Viya was integrated to provide a real-time view of the taxpayer to automatically detect flag compliance issues. MTCA has improved public trust through more timely, more accurate data-driven decisions.

Key takeaways

- Instituted a fully integrated system of fiscal and customs administration that provides an extensive, in-depth view of taxpayer compliance.
- Utilized modern data management to review all data, including third-party data automatically integrated via data-sharing agreements with domestic and international partners and institutions.
- Empowered staff to work more collaboratively and efficiently with a user-friendly interface, enabling all users to generate reports and democratize data, ingest data from any authorized source and perform sophisticated analyses.

AI blueprint applicability: Technology, process and people

"We aim to maximize the data at our disposal to enhance compliance, enforce regulations and optimize resources to be more effective and impactful in our outcomes."

Joseph Caruana, Commissioner for Tax and Customs, Malta Tax and Customs Administration

"The initiative underway foresees a modernization of tax legislation, a true digital transformation encompassing new processes, new skills and advanced analytics systems, with a forward-looking perspective that includes adoption of artificial intelligence-based systems."

Pierre Vella, Project Leader within the Reform Team, Malta Tax and Customs Administration

Health Care and Life Sciences

A large health plan aimed to modernize their value-based analytics and AI environment. They received billions of claims each year and needed the proper tools to integrate and analyze data to generate insights, control costs and improve quality. Different models were built to optimize the effectiveness and costs of house calls, report outcomes of clinical interventions and predict the risk of readmission.

Key takeaways

- Integrated and prepared a large volume and variety of data.
- Used machine learning models, including voice detection, image recognition and clinical workflow mapping, to predict health outcomes, improve clinical recommendations and optimize costs.
- Reduced readmission rates due to operational and clinical performance improvements that were informed by predictive models of patient risk <30 days discharge.
- SAS Viya reduced analytical processing time by 97%, which added financial value in quality program participation.

AI blueprint applicability: Technology, process and people

Manufacturing: Georgia-Pacific

Georgia-Pacific relies on SAS advanced analytics and AI solutions powered by SAS Viya on Amazon Web Services (AWS). From plywood to corrugated boxes to paper napkins, Georgia-Pacific runs more than 30,000 machine learning models to calculate the optimal production settings based on current business needs.

Key takeaways

- Trains models in the cloud and then deploys them instantly at the edge.
- Uses real-time data from 85,000 vibration sensors to intervene early when models predict increased likelihood of process anomalies.
- Experiments with digital twins to run realistic simulations to find the best solutions.

AI blueprint applicability: Technology, process and people

> “The advanced analytics enabled by SAS allows us to find the optimal balance of speed and quality to maximize profitability. We’re constantly pushing the envelope of what’s possible with analytics and AI.”
>
> **Roshan Shah**, Vice President AI & Products, Georgia-Pacific

SAS Viya – Enabling enterprise-ready AI

VALUE

- Increased accuracy
- Confidence
- Higher productivity
- Efficiency
- Flexibility
- Cost savings

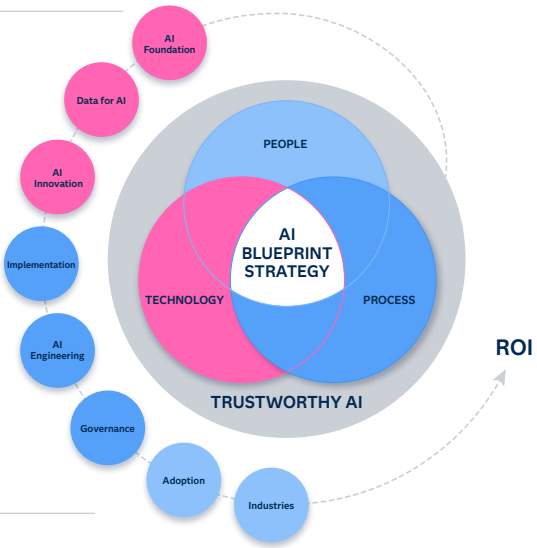

SAS Viya: The platform that operationalizes the AI blueprint

Bringing it all together, the AI blueprint with SAS Viya is designed to support data, technology, processes and most importantly – people.

It covers the entire AI life cycle – from data preparation and model development to monitoring and governance. It helps organizations realize value from AI faster, with trusted and explainable outputs that comply with domestic and international regulatory requirements. It's a strategic guide that prioritizes productivity, performance and trust.

At every step of the blueprint, Viya provides the scalability and flexibility needed to unlock new levels of productivity and innovation. Viya delivers significant productivity gains to organizations using AI and improves performance while uniquely lowering the total cost of ownership in the cloud. The platform enables trustworthy decisions that are explainable and transparent.

Agentic AI presents a lot of opportunities but also an exceptional amount of risk if the technology is not incorporated with governance and responsibility. SAS Viya has embedded governance and explainability at every step for trusted decisions. The platform guarantees that AI systems not only deliver accurate outcomes but also adhere to ethical standards, maintain data privacy and align with the organization's values and regulatory obligations.

> SAS' approach to agentic AI strikes the critical balance between autonomous decision making and ethical governance. Its intelligent agents represent not just technological advancement but a pragmatic framework for responsible enterprise AI adoption – precisely what organizations need as they navigate this rapidly evolving landscape to gain a competitive advantage."
>
> **Nick Patience**, Vice President and Practice Lead, Artificial Intelligence Platforms, The Futurum Group

Notable research

A benchmarking study conducted by The Futurum Group revealed that the SAS Viya data and AI platform helps users execute the life cycle of collecting data, building models and deploying decisions 4.6x faster than selected competitors, helping to increase innovation, speed up decision making and drive revenue growth.[8] The Futurum Group analysis compared Viya to a leading commercial environment and non-commercial open source environments, including Jupyter Notebook with MLFlow and Python Libraries.

The Futurum Group conducted another benchmarking study that found SAS Viya was 30x more scalable and 86% more cost-effective than commercial and open source alternatives.[9]

The **SAS Viya Value Calculator** was created based on research to calculate potential ROI. The calculation factors on-site visitor data inputs, research conducted by The Futurum Group and SAS benchmarks and domain expertise.

> ***SAS is an Emerging Leader in the 2025 Gartner Emerging Market Quadrant for Generative AI Engineering***
>
> *SAS is the only vendor named a leader in each of these 2024 AI analyst validations:*
>
> - *IDC MarketScape: Worldwide Machine Learning Operations Platforms*
> - *2024 The Forrester Wave™: AI / ML Platforms, Q3 2024*
> - *IDC MarketScape: Worldwide Decision Intelligence 2024*

*SAS has been recognized for innovative technology and social innovation projects by **American Business Awards, Fast Company, PRNews, SiliconANGLE Media** and more.*

Learn more at sas.com/viya.

Appendix

Glossary of terms

Artificial intelligence (AI) – AI covers technologies such as machine learning, deep learning and generative AI for natural language processing, voice/speech recognition, content generation, image/video analysis and time series analysis.

AI agent – An AI agent is a software program that uses machine learning and/or generative AI to set goals, make decisions, and take actions autonomously. Note that unlike a simple chatbot or AI assistant that just responds to prompts, an AI agent can perceive its environment, reason about what to do and act independently.

Agentic AI – An agentic AI system involves multiple AI agents that collaborate, learn from each other, make decisions, and take actions

AI governance – A unified, holistic AI governance approach to aggregating, orchestrating and monitoring AI systems, models and agents.

AI model life cycle – The AI model life cycle refers to the series of stages that an AI model goes through from conception to deployment, monitoring, governance and beyond.

AI trust gap – The AI trust gap is the disconnect between an AI system's actual reliability or performance and the level of confidence users or stakeholders have in adopting and relying on it.

AI trust index – An AI trust index is a measurement an organization can use to evaluate the value of AI implementation, operational productivity, innovation potential and ethical oversight.

AI engineering – A discipline focused on the practical application of AI techniques to design, develop and deploy intelligent systems, inclusive of operational disciplines like DataOps, ModelOps and DecisionOps, which are crucial for the AI life cycle.

AI human-in-the-loop – The practice of integrating human input into the AI life cycle, where humans and machines remain connected to ensure ethical and transparent outputs.

Copilot – AI assistants that work alongside human users, helping them with tasks, boosting productivity and enhancing creativity.

Data and AI literacy – Data and AI literacy refers to skills in the modern workforce where users of technology can read, understand, create and communicate using AI with ethical awareness.

Digital twins – A virtual representation of a physical object, system or process. The technology uses real-time data, simulations, machine learning and reasoning.

Edge computing – A distributed computing framework that brings computation and data storage closer to data sources.

Generative AI (GenAI) – GenAI is a subset of AI technologies used to create new content, such as text, images, videos, audio and synthetic data, based on patterns learned from large amounts of existing data.

Machine learning (ML) – ML is another subset of AI technologies that focuses on the design, development, deployment and maintenance of machine learning models.

Model cards – Model cards serve to highlight indicators like accuracy, fairness and model drift, which is the decay of model performance as conditions change. They are used by stakeholders across the data and AI life cycle to promote transparency.

Traditional AI – Includes all AI technologies other than GenAI.

Trustworthy AI – Trustworthy AI is AI that is fair, explainable, transparent, secure and safe. It is used across AI systems to mitigate risk and harm to people and organizations.

Quantum AI – A combination of quantum computing and AI to create algorithms and models that leverage quantum properties, e.g., superposition and entanglement, to process information and solve complex problems faster and/or more accurately than traditional computers.

References

1. **Boston Consulting Group (BCG)**, ***Where's the Value in AI?***, October 24, 2024

2. **IDC Research**, ***The AI Gold Rush in Asia Pacific***, October 23, 2024

3. **McKinsey**, ***The Economic Potential of Generative AI: The Next Productivity Frontier***, June 14, 2023

4. **SAS and Coleman Parkes Research**, ***The Race to Success With Generative AI***, April 2024

5. **Gartner**, ***Gartner Predicts Over 40% of Agentic AI Projects Will be Cancelled by the End of 2027***, June 25, 2025

6. **Gartner**, ***Gartner Experts Answer the Top Generative AI Questions for Your Enterprise***, 2025

7. **The Commonwealth**, ***The Commonwealth Secretariat and SAS join forces to secure an equitable digital future for young people***, August 13, 2024

8. **The Futurum Group**, ***Faster AI & Analytics: SAS® Viya® Outperforms the Competition***, May 2023

9. **The Futurum Group**, ***From Data to Decision: Increasing AI Productivity with SAS® Viya®***, September 2024

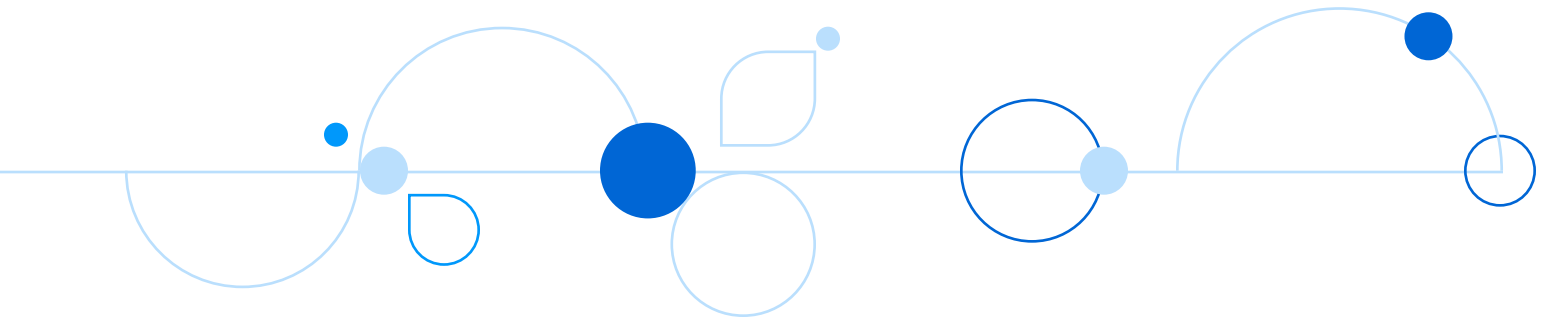

For more information, please visit sas.com/viya.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.
Other brand and product names are trademarks of their respective companies. Copyright © 2025, SAS Institute Inc. All rights reserved. W399500.0825