Expert Report

D2A2
DIGITAL DATA ANALYTICS
ARTIFICIAL INTELLIGENCE

D2A2 Insights

2026 Digital, Data, Analytics, and AI (D2A2) Predictions and Prescriptions






Foreword

As we enter 2026, the convergence of digital, data, analytics, and AI are accelerating into a new phase—one where intelligence becomes embedded into every layer of the enterprise. Last year's D2A2 Insights report highlighted the rise of multimodal interfaces, the shift toward data-centric AI, and the growing importance of responsible innovation. Those themes continue, but they now serve as the foundation for a more adaptive, accountable, and strategically focused AI landscape.

This year's predictions reflect a decisive shift toward systems that optimize and govern themselves. AI-native data architectures will redefine modern engineering, enabling pipelines that continuously learn and self-correct. Data observability will become the backbone of AI governance, transforming trust and resilience into measurable operational capabilities. At the strategic level, data products and mesh-based ecosystems will take center stage, creating interconnected environments where data behaves as a dynamic, reusable asset.

Also anticipated is the resurgence of classical AI techniques—revitalized by modern computing capabilities—to deliver efficient, interpretable models for specialized tasks. At the frontier, reasoning LLMs will push AI beyond pattern matching toward cognitive capabilities such as justification, planning, and complex decision-making. Hybrid multimodal interfaces will continue the trajectory we noted last year, bringing voice, vision, text, and gesture together into seamless, context-aware interactions.

Enterprises will also pivot toward smaller, domain-focused language models that deliver faster performance, lower costs, and tighter alignment with business needs. Local and offline AI—powered by edge computing and efficient architectures—will reach critical mass, bringing privacy-preserving intelligence to any environment. Meanwhile, AI governance will evolve from principles to quantifiable standards, and new ROI frameworks will introduce long-needed discipline into AI investment decisions.

Taken together, these developments signal an era where AI becomes more precise, more explainable, and more deeply woven into the fabric of digital strategy. The prescriptions that follow are designed to help organizations harness these shifts—ensuring that their data and AI capabilities are not only innovative, but resilient, ethical, and built for the cognitive future ahead.

Regards and thank you
Raef Lawson

Executive Director
Profitability Analytics Center of Excellence (PACE)
W: www.profitability-analytics.org

Introduction

We are standing at the threshold of a new economic revolution, driven by D2A2 (Digital, Data, Analytics and AI) technologies. Data is no longer just a byproduct of transactions and interactions; it has become the very foundation of economic growth. The increasing digitization of every industry sector has resulted in the explosion of data to further fuel innovation, improve decision-making, and unlock new opportunities for both businesses and individuals.

With this backdrop, this research work **D2A2 Insights: 2026 Digital, Data, Analytics, and AI (D2A2) Predictions and Prescriptions**, not only provides the 12 predictions, but also gives a **prescriptive recommendation and the required capabilities** for organizations to prepare for these changes in 2026 and beyond. Each of the **12 D2A2 Insights**, has three main sections.

1. Background i.e., the **context** for the prediction
2. The **rationale** for the prediction
3. Finally, the **prescriptive recommendation** on how to respond and capitalize on this prediction so that organization can plan and prepare.

This research report is a teamwork from some of the finest D2A2 experts in the industry. Firstly, I thank the entire D2A2 Analyst team i.e. **Arun Marar, PhD**, **V Balasubramanian, PhD, MBA**, **Sumi Singh, PhD, Tobias Zwingmann, MSc,** and **Valentina Tomov, PhD** for collaborating with me. Secondly, special thanks to **CFO.University, Arcalea, Aztra, Exasol**, and **Qualdo** for supporting this research. Last, but not least, I thank **Raef Lawson**, CEO of PACE for writing a fantastic foreword for this report.

I hope you will find the D2A2 Insights report useful for your organization. For any questions or clarifications, do not hesitate to contact me. Have a great 2026!

Regards and thank you

Prashanth Southekal, PhD., MBA, MS, ICD.D
Founder and Managing Principal
Data for Business Performance (DBP) Institute and D2A2.ai
W: www.d2a2.ai

Contributors: The Analyst Team

The **2026 D2A2 Insights** report is researched, analyzed, and prepared by the following Analysts.

Arun Marar, Ph.D. is an expert at implementing AI/ML models utilizing a wide range of technologies such as supervised and unsupervised learning, text analytics, natural language processing (NLP) and deep learning. Arun has consulted with clients all around the world across domains such as Banking and Financial Services, Retail, Healthcare and Transportation Logistics. Arun is also an expert in optimization, and this has helped him aid companies in moving up the value chain of analytics from predictive analytics to prescriptive analytics. Dr. Marar has a PhD in Operations Research from Princeton University, and he can be reached at amaar@d2a2.ai.

Prashanth Southekal, PhD, MBA, ICD.D is a Consultant, Author, and Professor. He has worked and consulted for over 80 organizations including P&G, GE, Shell, Apple, and SAP. He is the author of three books — “Data for Business Performance”, "Analytics Best Practices” and “Data Quality” — and writes regularly on data, analytics, and ML in Forbes and CFO University. His second book, ANALYTICS BEST PRACTICES was ranked #1 analytics book of all time in May 2022 by BookAuthority. Apart from his consulting and advisory pursuits, he has trained over 4,500 professionals worldwide in Data, Analytics, and AI. Dr. Southekal is also an Adjunct Professor of Data and Analytics at IE Business School (Madrid, Spain). He has a PhD from ESC Lille and MBA from Kellogg School of Management. Dr. Southekal can be reached at psouthekal@dbp-institute.com and psouthekal@d2a2.ai.

Dr. Sumi Singh is the Founder and CAIO of GAIL180, a consulting and innovation lab pioneering AI-First transformation. She holds a PhD in Computer Science and brings more than 15 years of experience across enterprise AI, product development, research, and academia. Dr. Singh partners with CXOs, boards, and innovation leaders to accelerate the adoption, commercialization, and scale of AI across industries. A recognized thought leader, technologist, and educator, she leads GAIL180 in guiding Fortune 500 companies, universities, and SMBs to shape and execute AI-First strategies that deliver measurable business impact. Her email id is sumi@gail180.com.

Dr. V. "Bala" Balasubramanian is a Strategic Advisor for the Healthcare and Life Sciences (HLS) industry. Till recently, he was Senior Vice President and Global Head for HLS at Orion Innovation. Bala was the Co-Founder and CEO of Cabeus for nine years, a niche products and services firm specializing in regulatory for Life Sciences. Over his 36 years of experience, Bala has implemented innovative information management solutions for firms such as Bristol-Myers Squibb, Aventis, Roche, Merrill Lynch, AT&T, Bell Atlantic and IBM. Bala is associated with Rutgers University, TiE NJ, and DIA, a professional association for the Life Sciences industry. Bala is a co-author of one of the leading industry references titled "Achieving Excellence in Regulatory Information Management", available on Amazon. He has a PhD and MBA from Rutgers University. Bala can be reached at vbala@vbinsightsllc.com.

Tobias Zwingmann is Managing Partner of the German AI advisory firm RAPYD.AI and on a mission to help companies implement machine learning and AI faster while delivering meaningful business value. He is the author of the 3 books “AI-Powered Business Intelligence”, “Augmented Analytics” and “The Profitable AI Advantage”. He has more than 15 years of professional experience working in a corporate setting where his responsibilities included building data science use cases and developing an enterprise-wide data strategy. Mr. Zwingmann can be reached at tobias@rapyd.ai and tzwingmann@d2a2.ai.

Valentina Tomov supports research and thought leadership at D2A2 focused on the ethical impacts of AI. She draws on her background analyzing corporate governance and ESG reporting. Valentina has led global education and leadership programs across sectors, designing peer-learning initiatives and cross-cultural collaborations that inform her human-centered approach to AI. She is a former Fulbright student with a PhD in Social Anthropology. Her education and work across academia, nonprofits, and industry informs her human-centered approach to AI governance. You can reach her at vtomov@d2a2.ai.

D2A2 Insight 1:
Headless Architectures Will Drive Modern Systems

Background

In recent years, the digital landscape has undergone a significant transformation, driven by the demand for a more flexible, scalable, and user-centric systems. Traditional monolithic architectures, where the front-end and back-end are tightly coupled, have given way to more modular approaches. Headless architectures, which separate the front-end (the user interface) from the back end (the databases and content management system), are gaining traction for developing digital solutions across industries, functions, platforms and devices.

Rationale for Prediction

The rapid evolution of digital experiences has driven the need for systems that are agile, scalable, and adaptable. Headless architectures offer a streamlined approach by decoupling content and data from presentation. This decoupling allows businesses to better manage and distribute data and content across a variety of channels. As the demand for personalized, omnichannel experiences increases, headless solutions offer the flexibility and speed. Additionally, advancements in APIs and microservices are accelerating the adoption of headless architectures, making them the ideal foundation for future-proof digital systems.

Prescriptive Recommendation

As businesses seek to create more personalized and seamless digital experiences, headless architectures will be the driving force behind the next generation of digital systems. Below are three recommendations for organizations that are seeking to implement headless architecture.

1. **Invest in API-first Development:**
   For organizations looking to adopt a headless architecture, starting with an API-first approach is essential. This means designing and developing all back-end systems with APIs at the core, allowing for easy integration with different front-end platforms. APIs provide the flexibility to deliver content seamlessly across various platforms and devices, ensuring a consistent user experience.

2. **Embrace Modular, Microservices-Based Approaches**
   To truly maximize the benefits of headless architectures, businesses should also adopt a microservices-based approach for their back-end systems. Microservices enable the modularization of business logic, making it easier to scale and manage different components of a digital ecosystem independently. By breaking down monolithic applications into smaller, manageable services, organizations can build systems that are more resilient, cost-effective, and adaptable.

3. **Focus on Omnichannel User Experience**
   As headless architectures enable the delivery of content and data across multiple touchpoints, businesses must prioritize an omnichannel approach to user experience. Companies should leverage the flexibility of headless systems to tailor content dynamically, optimize for different form factors or screen sizes across various devices, and integrate personalized recommendations.

D2A2 Insight 2:
AI-Native Solutions Will Redefine Data Foundations

Background

AI-native data solutions represent a shift from traditional, schema-first data architectures. These AI solutions will interpret, enrich, and orchestrate data and content rather than relying on extensive human-driven modeling. Core concepts include semantic data layers (using embeddings, vector search, and knowledge graphs to encode meaning), AI-driven data quality (continuous learning-based monitoring and remediation), and autonomous data engineering (AI-generated and self-healing data pipelines).While traditional data modeling using ER modeling techniques will become obsolete in the near future, AI will be the driving force behind modern data engineering, powering intelligent, and self-optimizing/healing data pipelines

Rationale for Prediction

Enterprises today face exponential growth in content and data, overwhelming traditional approaches to data acquisition, curation, integration, and engineering. At the same time, success with LLMs (Large Language Models) and advanced analytics hinges on the semantic richness and interoperability of the underlying data and content. Vendors across cloud, data management, and observability markets are rapidly embedding AI-native features such as vectorization, agentic automation, and semantic catalogs, making them unavoidable components of next-generation architectures. Finally, business expectations for rapid time-to-value require data teams to move beyond manual engineering work towards oversight of automated intelligence. These forces are collectively making AI-native data foundations inevitable.

Prescriptive Recommendation

Organizations can prepare for this change by adopting the following strategies.

1. **Prioritize Semantic Enrichment as a Core Capability**
   Organizations should move beyond metadata capture for building semantic layers that encode relationships and context. Firms should invest in data pipelines, vector indices, and knowledge graph extensions and ensure all new data products and AI solutions support semantic retrieval so that AI agents and LLM applications can operate with richer understanding and less manual transformation.

2. **Operationalize AI-Driven Data Quality and Observability**
   Deploy AI agents that continuously profile data, detect anomalies, and recommend or execute remediation steps. Use these capabilities to reduce human effort in data reconciliation ad remediation activities. Build governance processes that focus on validating AI-driven decisions, setting escalation triggers, and ensuring transparency. Update roles, skills, and team structures to leverage automation deeply, while enabling AI agents to handle repetitive integration and transformation tasks.

3. **Introduce Autonomous Data Engineering in Targeted Domains**
   Begin by automating routine or high-change workloads such as ingestion of new sources, document generation, schema change detection, and lineage mapping. Evaluate platforms capable of generating pipelines, optimizing SQL, or auto-suggesting transformations. Establish a human-in-the-loop review model to maintain trust while accelerating delivery.

D2A2 Insight 3: **Data Observability Will Be The Heart Of AI Governance**

Background

Data observability is the ability of the enterprise to understand the health, quality, and behaviour of its data through continuous monitoring of data pipelines, data quality, lineage, and system performance offering an end-to-end, real-time and automated view of data dependencies and anomalies. AI governance encompasses the frameworks, controls, and processes that ensure data and AI systems are ethical, compliant, reliable, and aligned with organisational and regulatory expectations. As AI systems become increasingly embedded in critical decision processes, data quality is no longer a downstream operational issue; it is a strategic governance concern. Poor or drifting data directly translates into poor or drifting AI behavior. This makes data observability a foundational enabler of trustworthy AI.

Rationale for Prediction

AI governance has long focused on model-centric controls i.e. validation, explainability, monitoring, and documentation. But most of the issues stem more from data issues than from model flaws. Regulators are responding and the EU AI Act, NIST AI RMF, and emerging UK/EU/US guidelines emphasise continuous oversight, traceability, and demonstrable control of data - training and operational. As a result, AI governance is shifting from model oversight to data-supply-chain oversight. In this regard, data observability becomes essential infrastructure where detecting data drift, tracking lineage for compliance, validating training datasets, and uncovering system-level risks are all becoming the core of AI governance.

Prescriptive Recommendation

So, how can organizations prepare for this prediction? Below are three prescriptive recommendations for implementation.

1. **Embed Data Observability into the AI Governance Operating Model**
   Treat data observability as a governance control, not an engineering or technical function. Incorporate data observability checkpoints covering data freshness, completeness, drift signals, and lineage validation. Define clear thresholds for acceptable data behaviour and align these with AI governance requirements. Also, ensure that observability outputs automatically trigger governance actions, such as model review, rollback, or human validation.

2. **Establish Data Contracts for Unified Data Management**
   Create end-to-end, cross-platform data contracts between engineering, analytics, and AI teams that covers the full data lifecycle—from acquisition to integration to transformation through to model training and inference. Ensure that all changes to upstream data sources are visible, tested, and assessed for regulatory and operational impact. Such transparency strengthens governance decisions, accelerates incident resolution, and reduces model risk.

3. **Operationalise AI Risk Management**
   Implement continuous observability across data pipelines, with alerting configured for anomalies that influence model accuracy, fairness, or stability. Integrate observability platforms with model monitoring, MLOps, and incident-management tools to enable coordinated response. Develop playbooks that map specific data issues such as schema changes, null surges, drift, unusual distribution shifts to predefined mitigations and communication routes. Last, but not least, train teams to respond rapidly, ensuring AI systems remain compliant, reliable, and aligned with business intent.

D2A2 Insight 4:
Data Products And Data Mesh Will Be Part Of The Data Strategy

Background

Data products are well-managed, and reusable data assets which provide business-ready, trusted, and easily accessible data streams for improved business performance. By treating data as a product, organizations improve data clarity, confidence, and usability, and reduce operational friction and manual data reconciliation resulting in better decision making and greater business alignment. A data mesh ecosystem is a decentralized data architecture pattern that organizes data ownership and management by business domains (e.g., marketing, sales, customer service). It empowers domain teams to own and manage their data as data products, including data ingestion, transformation, governance, and serving. This decentralization breaks down data silos, fosters collaboration, and supports scalable, self-service data management with robust yet federated governance.

Rationale for Prediction

Together, data products and data mesh are important for enterprise data strategy because they shift the focus from centralized, reactive data management to a proactive, product-oriented, and domain-driven approach. This improves scalability, data quality, governance, and self-service, enabling organizations to unlock more value from their data while accelerating decision-making and reducing reliance on central data teams. They help enterprises move from fragmented, siloed data to a cohesive, transparent, and agile data ecosystem thereby aligning technology efforts with business outcomes for sustainable growth and innovation.

Prescriptive Recommendation

The recommendations below will enable a scalable, decentralized, product-focused data ecosystem that will align data initiatives with business outcomes, reduce bottlenecks, and improve data trustworthiness and accessibility across the enterprise.

1. **Strategy and Alignment**
   Secure strong buy-in and alignment with the business to ensure stakeholder support and clarity on goals, especially on data sharing across the enterprise. Define data domains clearly based on organizational structure, expertise, and business value streams to ensure domain ownership and efficient collaboration.

2. **Data Product Development and Implementation**
   Build and test data products and data mesh architectures by selecting manageable pilot projects with clear business value, well-documented pipelines, and domain teams. Treat data as a product by identifying owners, users, business and technical metadata, and service level objectives. Develop a self-service data platform to provide domain teams with the tools and capabilities to independently build, maintain, and serve data products as MVP (Minimal Viable Product).

3. **Governance and Continuous Improvement**
   Establish federated governance to balance domain autonomy with enterprise-wide standards for security, compliance, data quality, and observability. Use agile, iterative methods to launch, scale, and evolve the architecture, leveraging feedback, metrics, and continuous improvement from each phase of data product deployment.

D2A2 Insight 5:
Classical AI Will Complement Generative AI

Background

In the current landscape dominated by Generative AI (GenAI) and Large Language Models (LLMs), there is a widespread perception that these technologies can address most of the AI-related challenges. But AI is more than GenAI. While GenAI and LLMs have revolutionized content creation, Classical AI, the kind that predicts and optimizes business processes based on structured data, is being grossly neglected these days.

Rationale for Prediction

The hype surrounding GenAI has led many in the C-suite to prioritize launching GenAI initiatives. However, the notion that GenAI could replace Classical AI for a wide range of business problems has now been dispelled. As a result, we are witnessing a renewed focus on Classical AI—particularly for critical business use cases to automate and optimize business processes. Overall, across industries, Classical AI is indispensable for structured data analysis, risk prediction, and operational optimization, turning insights into actionable strategies that enhance operational performance, minimize risk, and improve decision-making.

Prescriptive Recommendation

The following strategies can help every organization realize the potential of Classical AI along with complementing Generative AI solutions:

1. **Prioritizing High-Value AI Use Cases and Performance Measurement.** Identify areas where Classical AI can enhance revenue, reduce costs, or mitigate risk. The selected use-cases should be based on factors such as data quality, resource availability, organizational skills, data maturity, business priorities and more to deliver measurable business value. Track performance metrics for the Classical AI paradigm such as ROI, engagement, adoption, efficiency and accuracy.

2. **Applying Separation of Concerns Across Classical AI and Generative AI.** Separation of concerns (SoC) is a proven and recommended Enterprise Architecture pattern. Use Classical AI to handle data-driven decision-making, including risk assessment, pattern recognition, and operational optimization. Generative AI can be adopted to manage language, creativity, personalization, and communication, transforming decisions into engaging outputs.

3. **Operationalizing Insights Through a Dual-Layer AI Workflow.** Implement the insights workflow as a structured process used to turn insights into decisions and business outcomes. Classical AI layer offers insights on actions to take, when, and for whom the action is applicable to. The Generative AI layer acts as the execution layer, communicating those insights into making decisions for improved business outcomes.

D2A2 Insight 6:
Agentic Commerce Will Be The Cornerstone Of Agentic AI Workflows

Background

Recently, we marked the third anniversary of ChatGPT's initial release. In that short span, we have witnessed the rapid rise of Generative AI and Large Language Models (LLMs). This widespread adoption has evolved into what is now known as Agentic AI i.e. an ecosystem of collaborative AI agents capable of autonomously executing complex business workflows. Agentic AI now spans a wide spectrum of tasks, from customer-centric processes to invoice processing to the analysis of large and complex datasets.

While the GenAI success rates are relatively low, we are observing a similar pattern with Agentic AI workflows. However, there is one notable exception is Agentic Commerce. Agentic Commerce has given rise to the phrase "Buy For Me" reflecting its core mission: enabling a new era of online shopping in which AI agents autonomously act on behalf of customers by discovering products, comparing options, negotiating prices, and completing purchases, often through a conversational interface.

Rationale for Prediction

Agentic Commerce isn't just a new interface or tool. It is a potential paradigm shift in how commerce works. Because AI agents can reduce friction, personalize for you, act on your behalf, and scale into complex workflows, many believe it will quickly become the default mode of online shopping. These agents are more than simple chatbots or recommendation engines. They have memory, reasoning, access to tools/APIs, and can make multi-step decisions based on budget, preferences, delivery and time automatically.

Prescriptive Recommendation

The following strategies can help organizations position themselves for success in the era of Agentic Commerce:

1. **Understand Your Customer**
   Agentic Commerce rests on understanding the customer. Companies must develop capabilities to understand the customer well—one whose interactions are mediated by AI agents and non-AI agents. Mastering the workflows of Agentic Commerce will be critical to achieving success related to customer engagement, personalization, and fulfillment.

2. **Enhance Product Data and Infrastructure**
   Ensure product metadata such as features, inventory levels, pricing, delivery windows, images, and reviews, is organized in machine-readable formats like APIs or structured markup (e.g., Schema.org). This represents the new form of SEO (Search Engine Optimization) for AI agents, often referred to as Answer Engine Optimization (AEO). Implement open APIs that allow AI agents to seamlessly retrieve product information, check stock, and initiate transactions. Adopt API-first checkout solutions that enable agents to securely complete purchases using tokenized payments or virtual cards, ensuring both speed and safety.

3. **Automate Operations While Prioritizing Trust and Oversight**
   AI agents can streamline internal operations—forecasting demand, managing inventory, triggering replenishment, and optimizing logistics and delivery routes based on demand patterns and competitor activity. At the same time, retailers must protect customer data and ensure trust through transparency, strong security, and clear human oversight. This includes setting permissions for high-impact decisions, preventing unauthorized purchases, and deploying secure protocols and fraud-detection systems.

D2A2 Insight 7:
Reasoning Llms Usher In The Age Of Cognitive AI

Background

The first wave of GenAI emphasized content generation from text, images, code, video, and audio. The second wave introduced multimodal perception integrating modalities. In less than half a decade, enterprises are entering a third transformational phase: large language models (LLMs) with explicit reasoning capabilities. Unlike prior AI systems that responded reactively to prompts, these new agents can decompose goals, plan multi-step actions, use applications and APIs, self-correct through reflection loops, validate results for accuracy, and maintain short- and long-term memory. This evolution positions AI as a collaborative problem-solving system capable of navigating complexity across diverse business functions.

Rationale for Prediction

AI systems are shifting from reactive assistants to reasoning-first agents capable of planning, validating, and executing multi-step tasks. Frontier models from OpenAI, Anthropic, Google, DeepSeek, and emerging reasoning startups now show major gains in deliberate reasoning, mathematical accuracy, and structured decision-making, using multi-step inference, self-checking, and task decomposition.

Enterprises are also adopting architectures that connect LLMs with RPA tools, SaaS APIs, rules engines, knowledge graphs, and vector stores, allowing agents not just to interpret information but to act across complex workflows. This shift mirrors the broader trend in AI system design going beyond modality for action-oriented systems, where diverse model types are becoming interoperable components.

Prescriptive Recommendation

Below are the three techniques enterprises can adopt to implement reasoning LLMs in the age of Cognitive AI.

1. **Build co-reasoning systems**
   Re-architect workflows so the goal isn't just automation but AI and humans reasoning together. The shift is from "AI executes a step" to "AI co-reasons with humans." For example, the AI can first draft a reasoning map, humans refine constraints, preferences, or risk tolerance, and then the AI recalibrates and re-plans based on that input.

2. **Develop LLMOps and Reasoning drift detection**
   Enterprises must build robust LLMOps pipelines covering model/prompt versioning, reasoning-quality evaluation, monitoring, staging, and controlled rollouts, so changes in an agent's reasoning are traceable, testable, and safely deployed. They also need human-feedback loops, override logging, and automated regression tests to continuously refine agent behavior and prevent reasoning drift as systems scale.

3. **Build a Reasoning Clean-Room for Safe Simulation and Stress Testing**
   Enterprises need a clean-room environment where agents can execute full reasoning cycles i.e. planning, tool calls, memory updates, and edge-case handling, against synthetic, masked, or historical datasets with zero production access. This clean room must support adversarial scenarios, multi-step simulations, regression tests, and reasoning failure analysis so enterprises can detect brittleness, unsafe actions, and logic drift long before deployment.

D2A2 Insight 8:
Hybrid Interfaces Enable True Multimodal Intelligence

Background

Multimodal models expand well beyond image–text pairing, incorporating video, audio, and temporal context to enable sequence reasoning rather than static interpretation. Today we have video-text systems performing video QA, action recognition, and cross-modal captioning, while industry prototypes use real-time multimodal inputs i.e. camera, audio, and sometimes gesture, to build context-aware agents that perceive environment, speech, and visual cues simultaneously. At the same time, advances in embeddings and transformer architectures—such as cross-modal attention, temporal–spatial fusion, and unified latent spaces for vision, audio, and text, are making true hybrid interfaces (voice + vision + gesture + text) increasingly practical. Given this momentum, by 2026 one can expect early real-world deployment of hybrid multimodal agents in enterprise settings, particularly in domains requiring environmental awareness and multi-sensory input, including field operations, customer service, remote collaboration, and physical-asset management.

Rationale for Prediction

By 2026, hybrid multimodal AI agents—systems that can listen, see, track motion and gesture, interpret context, and respond through voice, text, visuals, or gestures—will move from research prototypes to enterprise-ready tools in select high-value verticals. Operating in real-world environments, these agents will blend modalities to deliver seamless, context-aware interactions that feel richer and more natural. This shift will reshape training in fields such as medicine, sales, and education, and transform how context-intensive domains approach hybrid human–machine collaboration.

Prescriptive Recommendation

Below are three strategies to implement multimodal solutions.

1. **Design Interactions Around Situations, Not Single Modalities**
   Build workflows around real-world contexts—hands-busy tasks, spatial navigation, customer-facing interactions, or safety-critical checks—rather than anchoring experiences to voice-only or text-only interfaces. Hybrid multimodal agents work best when they can combine cues, using visual context to clarify voice commands, or gestures to guide complex actions.

2. **Invest in Infrastructure**
   To fully leverage hybrid interfaces, enterprises should upgrade data pipelines to include video streams, sensor inputs, gesture signals, and audio metadata, processed through unified embeddings or vector stores. This foundation enables agents to jointly reason across modalities and ensures real-world data and content is available for insight derivation and decision-making.

3. **Build Governance for Multimodal Inputs**
   Because multimodal agents operate in dynamic physical environments, organizations must implement guardrails that manage permissions, protect sensitive visual/audio data, and define when human oversight is required. Clear governance reduces operational risk and ensures responsible deployment of context-aware AI systems.

D2A2 Insight 9:
Enterprises Will Turn To Small Language Models (SLMs)

Background

Since the release of ChatGPT, enterprises have experimented with Large Language Models with mixed results, mostly sub-optimal in terms of addressing business challenges and use cases. The accompanying challenges around infrastructure costs, vast amounts of energy consumption to extract or summarize large amounts of business-related knowledge, combined with issues such as hallucinations, information loss, lack of accuracy and precision are leading enterprises to rethink the use of LLMs. A recent study by NVIDIA has revealed that Small Language Models (SLMs), rather than their larger counterparts LLMs, could become the backbone of the next generation of intelligent enterprises. A recent HBR article says "the age of bigger is better may be giving way to smaller is smarter" [1].

Rationale for Prediction

SLMs are compact GenAI models designed for natural language processing tasks, emphasizing efficiency, faster response times, lower computational costs, and better suitability for specialized or resource-constrained environments without needing enormous infrastructure. SLMs use fewer parameters, millions to a few billions, compared to LLMs. Since SLMs are trained on a narrower range of specific data, SLMs are often more reliable for domain-specific agentic tasks compared to the general, broad scope of LLMs. Their growing importance stems from technical and practical advantages such as customizability, scalability, privacy, accessibility, and reduced energy costs. For example, Microsoft has already made SLMs available such as the Phi model family on Azure AI for businesses seeking efficient and cost-effective AI solutions.

Prescriptive Recommendation

Following are some recommendations for implementing SLMs pragmatically and securely.

1. **Apply SLMs to High Value Use-cases**
   SLMs can streamline and automate sector-specific workflows such as customer service interactions, IT helpdesk responses, expense approvals, and procurement processes. By tailoring SLMs to the language and rules of each function, organizations can reduce operational costs while maintaining accuracy, consistency, and efficiency.

2. **Strengthen Data Privacy, Compliance, and Real-Time Edge Performance**
   SLMs can be deployed locally to enhance privacy and regulatory compliance - particularly valuable in regulated sectors like Healthcare, Banking and Insurance. Their smaller footprint enables training on proprietary datasets with minimal cloud exposure. At the same time, their low-latency performance makes them ideal for real-time or edge use cases such as voice assistants, autonomous vehicles, wearable health devices, and smart factories.

3. **Enable Specialized Knowledge, Rapid Experimentation, and Scalable Adoption**
   SLMs excel at fine-tuning on narrow domain knowledge, improving reliability in fields like agriculture, law, medical diagnostics, finance, and supply chain management where general LLMs may be less precise. They also offer a cost-effective way to prototype, experiment, and embed AI within existing workflows or legacy systems. This supports safe, incremental AI adoption and helps organizations scale value quickly without major infrastructure overhauls.

References

1. A. Kumar, et al., The Case for Using Small Language Models, Harvard Business Review, September 8, 2025.

D2A2 Insight 10:
Local And Offline AI Workflows Reach Critical Mass

Background

Enterprise AI today is heavily dependent on cloud infrastructure. Organizations send data to external servers, face unpredictable consumption-based costs, and must accept shifting model behavior. The scale is immense; OpenAI alone processes billions of tokens per minute. This cloud-first model became standard because running advanced AI required specialized hardware and expertise that only hyperscalers could provide.

But the environment is changing. Open-weight models now rival proprietary systems; powerful consumer-grade hardware can run them efficiently; and mature tooling allows non-experts to deploy local AI stacks. At the same time, concerns about cloud reliance are growing. EU data-protection rules complicate processing sensitive data off-premises, major outages highlight the fragility of centralized infrastructure, and many organizations are finding cloud AI costs volatile and unexpectedly high.

Rationale for Prediction

Local and offline AI workflows are reaching critical mass because on-premises deployment now offers stronger economic, regulatory, and operational advantages for many enterprise use cases. For high-volume, repetitive tasks such as document classification or data extraction, cloud API pricing becomes prohibitively expensive, while local models provide predictable, lower-cost processing. In sectors like healthcare, legal, finance, and government, strict privacy and compliance requirements make on-prem AI essential, eliminating data-residency risks and avoiding third-party processing altogether. Operational resilience is another driver: cloud outages, rate-limit changes, and shifting model behavior can disrupt critical workflows, whereas local infrastructure ensures stable, self-reliant AI performance.

Prescriptive Recommendation

Organizations should act now to develop local AI capability before competitive pressure makes this transformation urgent rather than strategic. The following actions position enterprises to capitalize on the shift toward sovereign AI infrastructure.

1. **Start with the right local workloads to scale infrastructure**
   Identify early use cases that combine high-volume, sensitive data and moderate latency tolerance such as document processing, internal content generation, data extraction, and knowledge-base enrichment. These deliver fast ROI while building internal capability. Match infrastructure to workload: use a workstation-class GPU system for exploration and evaluate appliances or small data-center deployments only when production needs justify it. Avoid both overbuilding and underpowering.

2. **Build the automation layer that turns local AI into operational value**
   Local AI becomes useful only when embedded in business processes. Invest in workflow automation, orchestration, and integration tools so AI output flows directly into operations. This transforms local AI infrastructure from a lab experiment into a revenue- or efficiency-producing capability.

3. **Develop organizational expertise and a hybrid, multi-tier strategy**
   Assign staff to own local AI operations, provide structured training, and establish communities of practice. Combine this with a clear hybrid architecture strategy that allocates workloads between local and cloud environments based on sensitivity, economics, latency, and capability needs. Plan for the emerging three-tier model i.e. edge devices, departmental servers, and data-center deployments, each with its own governance. This preparation ensures scalable, efficient adoption as needs grow.

D2A2 Insight 11:
AI Governance Will Evolve Towards Measurable Accountability

Background

Responsible AI has shifted from ideals to board-level and investor priorities. By 2026, stakeholders will demand measurable proof of fairness, trust, and compliance, pushing AI governance into an audit era like ESG. Leading organizations will operationalize ethical frameworks, unify trust functions, and adopt continuous monitoring to demonstrate trustworthy AI performance. Those that make Responsible AI quantifiable and accountable will gain trust, regulatory confidence, and competitive advantage by proving their systems are consistently trustworthy in real-world use.

Rationale for Prediction

Regulatory pressure is accelerating the shift from principles-based AI governance to evidence-based accountability. The EU AI Act, U.S. executive directives, and emerging global standards all require ongoing monitoring, documentation, and proof of human oversight, not just responsible-AI statements. At the same time, enterprises are deploying AI deeper into customer journeys, operational workflows, and decision systems, creating new obligations to track drift, fairness, explainability, and data quality across the entire model lifecycle to preserve trustworthy system behavior.

This rising operational complexity is pushing organizations toward more structured, measurable governance frameworks. And it's not just about risk: the connection between governance and business outcomes is now empirically visible. According to a Gartner survey [1], "organizations that perform regular audits and assessments of AI system performance and compliance are over three times more likely to achieve high GenAI business value." This validates a growing industry reality: quantified governance is becoming a driver of ROI, not a drag on innovation. As leaders enter 2026, measurable accountability will differentiate organizations that can scale AI in trustworthy, auditable ways from those that cannot.

Prescriptive Recommendation

1. **Establish Unified AI Governance and Trust KPIs**
   Effective AI governance begins with consistent and measurable indicators of trustworthy performance. Organizations should quantify key KPIs such as fairness stability, model drift, explainability scores, and human-oversight coverage. These metrics must be tied to business outcomes—customer impact, operational efficiency, and risk reduction—ensuring governance becomes a strategic driver, not just a compliance exercise. Centralized dashboards showing inventories, risk tiers, lineage, and bias indicators help leaders continuously validate whether systems remain trustworthy at scale.

2. **Embed Compliance & Trust into Architecture and Delivery**
   Trustworthy-by-design should be the baseline architecture principle. AI applications must align with regulatory frameworks such as the EU AI Act, mapping each system to its corresponding risk tier, and required controls. Governance, privacy, data lineage, and AI assurance should be consolidated into a unified "Data & AI Trust Command Center" to avoid operational silos. Fairness testing, explainability assessments, and human oversight processes should be treated as non-negotiable requirements that must be met before any model is deployed or scaled.

3. **Strengthen Cross-Functional Accountability & Evidence Readiness**
   Sustainable AI governance requires cross-functional responsibility for maintaining trustworthy AI systems. Joint KPIs should be created across data, security, risk, product, and compliance teams to reduce fragmentation and improve alignment. Quantitative metrics should be complemented by qualitative cultural and anthropological assessments to capture local expectations of what "trustworthy" means in different markets. Finally, organizations must build a robust evidence trail that encompasses model cards, audit logs, oversight reports to demonstrate measurable trust to investors, partners, and regulators and to prepare for the emerging landscape of responsible AI ratings.

References

1. Gartner. "Gartner Survey Finds Regular AI System Assessments Triple the Likelihood of High GenAI Value." Gartner, Nov. 4, 2025.

D2A2 Insight 12:
ROI Frameworks Will Discipline AI Investments

Background

Enterprise AI investment has entered a crisis of accountability. Global enterprise AI spending is exceeding $600 billion in 2025. Yet, most of this investment does not see tangible P&L impact. The measurement problem is structural. Traditional IT ROI frameworks assume value materializes within 7-12 months; but only a small fraction of organizations achieved AI payback in under a year. Attribution proves nearly impossible when AI augments human work rather than replacing discrete processes. And consumption-based pricing creates nonlinear cost scaling that conventional budgeting cannot model. This has created organizational dysfunction. Boards demand ROI justification using frameworks designed for ERP implementations. Business units respond with inflated claims or abandon measurement entirely. This results in significant capital committed to AI with no shared language for evaluating success and viability.

Rationale for Prediction

AI ROI frameworks will bring more discipline to combat the current chaos that is unsustainable. One practical consensus on differentiated evaluation is already emerging, which is that AI investments fall into two categories requiring different measurement approaches:

1. Productivity AI (ChatGPT, Copilot, coding assistants) warrants qualitative and adoption metrics. This is pragmatic: productivity AI often doesn't generate measurable P&L impact. Sales teams write faster emails but still send the same number. Developers code better but follow the same release cycle. Hard ROI calculation produces false precision. What matters here is whether employees find tools useful and adoption is growing.

2. Engineered AI (process automation, system redesign, custom applications) require hard financial discipline. These investments consume significant capital and should deliver measurable impact. This demands defined cost caps, explicit value targets, and clear accountability as organizations discover prototype-to-production cost gaps with spiralling API costs.

Prescriptive Recommendation

1. **Establish separate evaluation tracks.**
   Productivity AI should be measured on adoption, user satisfaction, and perceived value—not forced into ROI calculations it cannot support. Engineered AI solutions require traditional financial discipline such as explicit value targets, clear attribution methodology, and accountability.

2. **Implement cost caps for engineered AI.**
   Establish hard spending limits per project and time. Deploy anomaly detection for usage spikes. Create experiment sandboxes with pre-approved spend zones. Require total cost of ownership analysis before approval. Build progressive budget gates: small initial allocation to prove feasibility, expanded funding contingent on progress.

3. **Accept appropriate timelines and build the governance structures.**
   Productivity AI should show adoption traction within 3-6 months. If employees aren't using tools after reasonable enablement, reconsider. Engineered solutions, however, could require multiple years for full value. Define intermediate milestones and track "trending ROI" indicators before financial returns materialize. Centralize evaluation methodology to prevent each business unit from reinventing approaches and enable pattern recognition across the portfolio.

Required D2A2 Capabilities

Fundamentally, the purpose of data and analytics is to offer insights to **know or to act** for the organization. If the purpose of analytics is to just **know**, ask: *Why do you want to know? How much do you want to know* - the time horizon? *What is the value of knowing and not knowing?* But if the purpose of insights is to act, ask what are the events that are adversely affecting the organization's ability to plan, respond, and control from the predictive analytics insight? Also, spare a thought on who will own or consume the derived insights? How can you realize the change?

In this backdrop, here are the 12 important capabilities required in effectively adopting the ten DBP insights in the organization. The below capabilities are specific to D2A2 and assumes that organizations have good capabilities on problem solving, collaboration/teamwork, customer service, talent management, and more.

- Office of the CDAO (Chief Data and Analytics Officer)
- Enterprise Performance Management (EPM)
- Digital Solutions
- Master Data Management (MDM)
- Data Integration and Engineering
- Canonical Platforms (DWH/Data Lake/Data Lakehouse)
- Semantic Layer and Data Catalog
- BI/AI/ML Platforms
- Data Literacy and Change Management
- Data Governance
- Data Products Mindset
- Decision making using data under Uncertainty and Risk

Criteria to Identify D2A2 Projects

So, how do you implement these D2A2 predictions? Selecting the D2A2 Projects for tangible business impact is based on **five key criteria**.

1. The D2A2 Projects selected should be practical and relevant for the business. The D2A2 Projects selected should measure and improve/optimize business objectives and performance. Specifically, the problem definition should be linked to 3 key business objectives: enhancing revenue, reducing costs and/or minimizing risks.

2. The project selected should be based on questions and factors that are under the influence and control of the business. The business question (dependent variable) should be contextual, and the data sourced (independent variables) should be relevant for deriving the answer to the question.

3. The D2A2 Projects should apply appropriate concepts and techniques such as EDA (Exploratory Descriptive Analytics), association, segmentation, prediction, Prescriptive Analytics, and so on to measure and improve/optimize business performance. The prerequisite for this is availability of data and its ease of integration. Fundamentally Data and Analytics help in leveraging the past performance to drive future results and performance. In addition, the project should apply quality data supported by realistic assumptions.

4. The D2A2 Projects should drive change and address implementation as well as change management aspects. Specific questions include:
    a. Whose problem (person or persona) are you solving?
    b. Who owns the Data and Analytics solution?
    c. Availability of resources and capabilities/skills
    d. What are the risks of implementing this solution? Is the Data and Analytics solution ethical?
    e. What solution and governance components are in play in the Data and Analytics Projects?

5. Finally, the D2A2 Projects should deliver quantifiable business impact measured using KPIs such as ROIC (Return on Invested Capital), FCF (Free Cash Flow), Growth, CX (Customer Experience), etc.

Review of 2025 D2A2 Insights

Since 2020, we are making predictions and recommending prescriptions for the following year. At the beginning of last year, we made the following 10 D2A2 predictions and prescriptions for 2025.

1. 3Ms (Mapping, Modeling and Measurement) will fuel Data Engineering Renaissance.
2. Data and AI Governance will guide Trustworthy D2A2 Solutions.
3. Seamless AI Environment will Integrate Data & Content across Digital Products.
4. Data Products will bridge Analytics and AI Solutions.
5. Cloud-Native Applications will transition to AI-Native.
6. Generative AI will enable Zero Technical-Debt Application Modernization.
7. Agentic AI will not Instantly boost Business Productivity.
8. AI Initiatives will be Strongly Evaluated against Use Cases and ROI.
9. True Agile Development Methodologies will drive in-house development.
10. There will be Greater demand for AI-trained IT and business professionals.

So, how did these insights perform? Based on the analysis of market adoption trends of our predictions, we assigned ratings on the following 1 -5 scale

- 5 = Fully realized
- 4 = Largely realized
- 3 = Partially realized
- 2 = Early traction
- 1 = Largely unrealized

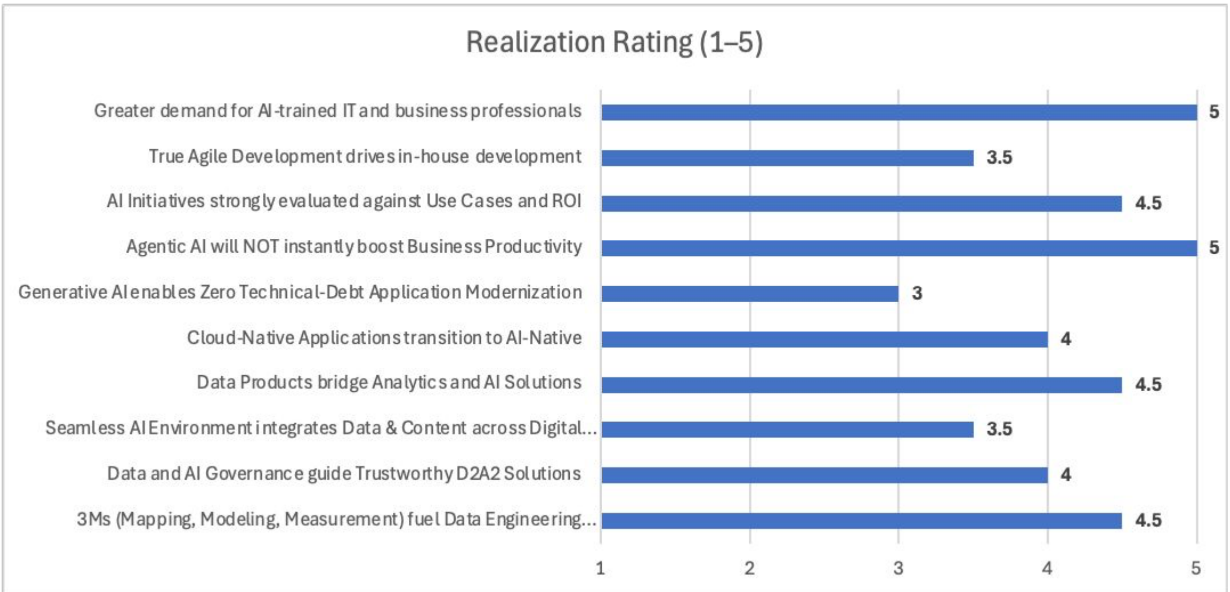

Figure 1: Review of 2025 D2A2 Insights

Overall, our 10 D2A2 predictions and prescriptions for 2025 were not hype driven. They were execution-aware and value-focused, which is why most of them materialized strongly in 2025.

#	2025 Prediction	Realization Rating (1–5)	Industry Reality in 2025
1	**3Ms (Mapping, Modeling, Measurement) fuel Data Engineering Renaissance**	4.5	Semantic layers, data contracts, observability, and outcome-driven measurement became central to modern data engineering.
2	**Data and AI Governance guide Trustworthy D2A2 Solutions**	4.0	Governance expanded to include AI risk, model oversight, provenance, and regulatory compliance, though maturity varies.
3	**Seamless AI Environment integrates Data & Content across Digital Products**	3.5	RAG pipelines and vector databases advanced integration gained traction, but tooling and architecture fragmentation remains.
4	**Data Products bridge Analytics and AI Solutions**	4.5	"Data as a Product" became operational, enabling shared foundations for BI, analytics, and AI use cases.
5	**Cloud-Native Applications transition to AI-Native**	4.0	New applications are increasingly AI-first, with embedded copilots and inference pipelines as core components.
6	**Generative AI enables Zero Technical-Debt Application Modernization**	3.0	GenAI accelerated modernization and refactoring, but true zero technical debt remains aspirational.
7	**Agentic AI will NOT instantly boost Business Productivity**	5.0	Early agentic systems faced reliability, governance, and orchestration challenges; productivity gains required redesign.
8	**AI Initiatives strongly evaluated against Use Cases and ROI**	4.5	Enterprises shifted from experimentation to value engineering, with stronger CFO and business scrutiny.
9	**True Agile Development drives in-house development**	3.5	In-house product and platform teams grew, though skills gaps and vendor dependence persist.
10	**Greater demand for AI-trained IT and business professionals**	5.0	AI literacy became essential across roles, driving sustained upskilling and talent demand.

Conclusion

By 2026, D2A2 strategies will center on more flexible, intelligent, and accountable systems. Headless architectures and AI-native platforms will modernize data foundations, while data observability and formal governance frameworks ensure trust, quality, and measurable accountability. Organizations will adopt data products and data mesh patterns to scale data sharing and value creation. AI will become more sophisticated and diversified: classical AI will work alongside generative models, reasoning-capable LLMs will drive cognitive automation, and hybrid multimodal interfaces will enable more natural human-AI interaction. Agentic workflows—especially in commerce—will automate complex tasks end-to-end. Enterprises will shift toward smaller, efficient language models and embrace local or offline AI to gain control, reduce cost, and improve privacy. Finally, ROI-driven frameworks will guide AI investment decisions, ensuring that innovation aligns with clear business outcomes.

In today's volatile, uncertain, complex, and ambiguous (VUCA) global market, the D2A2 landscape also offers enterprises unprecedented opportunities to derive business value. However, realizing value from D2A2 solutions requires deliberate, sustained effort. Enterprises must invest in acquiring high-quality data, adopting AI/ML models, building the right talent pool, re-engineering business processes and models, and strengthening data maturity, governance, and literacy. At its core, generating meaningful results from data and analytics is a change-management challenge. Ultimately, only the ability to translate data into action is what creates value.

Sponsors & Partners

Appendix

D2A2 Acronyms

- AEO – Answer Engine Optimization
- AI – Artificial Intelligence
- CDO – Chief Data Officer
- CDW – Cloud Data Warehouses
- CLIP – Contrastive Language-Image Pretraining
- D2A2 – Digital, Data, Analytics and AI
- DLC – Data (and Analytics) Lifecycle.
- EDW – Enterprise Datawarehouse
- EDG – Enterprise Data Governance
- EPM – Enterprise Performance Management
- GenAI – Generative AI
- IT – Information Technology
- KPI – Key Performance Indicator
- LLM – Large Language Models
- ML – Machine Learning
- MVP – Minimal Viable Product
- ROI – Return on Investment
- RBAC – Role Based Access Control
- SEO – Search Engine Optimization
- SLM – Small Language Model
- SoC – Separation of concerns
- TAVI – Text, Audio, Video and Images.
- TCO – Total Cost of Ownership
- VUCA – Volatile, Uncertain, Complex and Ambiguous

D2A2 Glossary

- → Algorithm: A mathematical formula or statistical process used to perform analysis of data.
- → API (Application Program Interface): A set of programming standards and instructions for accessing or building web-based software applications.
- → Artificial Intelligence: The ability of a machine to apply information gained from previous experience accurately to new situations in a way that a human would.
- → Big Data: Big data refers to datasets whose size is beyond the ability of typical database software tools to capture, store, manage and analyze. Big data sets are characterized by 3Vs i.e., volume, velocity, and variety.
- → Business Intelligence: The general term used for the identification, extraction, and analysis of multi-dimensional data.
- → Change Management: Change management is the discipline that guides how we prepare, equip and support individuals to successfully adopt change to drive organizational success and outcomes.
- → Cloud Computing: A distributed computing system hosted and running on remote servers and accessible from anywhere on the internet.
- → Dashboard: Dashboard is a visual snapshot of business performance using KPIs to help users make smarter, data-driven decisions.
- → Data Analytics: The process of answering business questions using data. Businesses typically use the three types of analytics: Descriptive, Predictive and Prescriptive Analytics.
- → Data Analysis Workspace: It is a flexible browser-based interface that allows you to quickly build analyses and share insights. Using the drag-and-drop interface, you can craft your analysis, add visualizations to bring data to life, curate a dataset, share and schedule projects with anyone in your organization.
- → Data Architecture: It is the mechanism in which data is collected, and how it is stored, arranged, integrated, and used in data systems and in organizations.
- → Data Governance: A set of processes or rules that ensure data integrity and that data management best practices are met.
- → Data Integration: The process of combining data from different sources and presenting it in a single view.
- → Data Lake: A large repository of enterprise-wide data in raw format – structured and unstructured data

- Data Mart: The access layer of a data warehouse used to provide data to users.
- Data Mining: It is finding meaningful patterns and deriving insights in large sets of data using sophisticated pattern recognition techniques. To derive meaningful patterns, data miners use statistics, machine learning algorithms, and artificial intelligence techniques.
- DataOps: DataOps as defined by Gartner is a collaborative data management practice focused on improving the communication, integration, and automation of data flows between data managers and data consumers across an organization.
- Data Product: A data product is the application of data for improving business performance; it is usually an output of the data science activity.
- Data Science: A discipline that incorporates statistics, data visualization, computer programming, data mining, machine learning and database engineering to solve complex problems.
- Data Warehouse: A repository for enterprise-wide data but in a structured format after cleaning and integrating with other sources. Data warehouses are typically used for conventional data (but not exclusively).
- Database: A digital collection of data and the structure around which the data is organized. The data is typically entered into and accessed via a database management system.
- Deep learning: Deep learning is a function of AI that imitates the human brain by learning from how it structures and processes information to make decisions. Instead of relying on an algorithm that can only perform one specific task, this subset of machine learning can learn from unstructured data without supervision.
- Descriptive Analytics: Condensing big numbers into smaller pieces of information. This is like summarizing the data story. Rather than listing every single number and detail, there is a general thrust and narrative.
- ETL (Extract, Transform and Load): The process of extracting raw data, transforming by cleaning/enriching the data to make it fit operational needs and loading into the appropriate repository for the system’s use.
- Hallucination: Hallucination refers to an incorrect response from an AI system, or false information in an output that is presented as information.
- Hypothesis. A hypothesis is an assumption, an idea, or a gut feeling that is proposed for the validation so that it can be tested to see if it might be true.

→ Insight. It is the understanding of a specific cause and effect within a specific context. In this book, the terms insight and information are used interchangeably.
→ KPI. A Key Performance Indicator (KPI) is a measurable value that demonstrates how effectively the entity is achieving key objectives or targets.
→ Machine Learning: A method of designing systems that can learn, adjust and improve based on the data fed to them. Using statistical algorithms that are fed to these machines, they learn and continually zero in on "correct" behavior and insights and they keep improving as more data flows through the system.
→ Metadata. Any data used to describe other data — for example, a data file's size or date of creation.
→ ModelOps: ModelOps as defined by Gartner is on the governance and life cycle management of a wide range of operationalized artificial intelligence (AI) and decision models, including machine learning, knowledge graphs, rules, optimization, linguistic and agent-based models.
→ Online analytical processing (OLAP). The process of analyzing multidimensional data using three operations: consolidation (the aggregation of available), drill-down (the ability for users to see the underlying details), and slice and dice (the ability for users to select subsets and view them from different perspectives). OLAP systems are used in BI reports.
→ Online transactional processing (OLTP). The process of providing users with access to large amounts of transactional data in a way that they can derive meaning from it. OLTP systems are used in Transactional reports.
→ Predictive Analytics: Using statistical functions on one or more data sets to predict trends or future events.
→ Prescriptive Analytics: Prescriptive analytics builds on predictive analytics by including actions and make data-driven decisions by looking at the impacts of various actions.
→ Prompt: A prompt is an input that a user feeds to an AI system to get a desired result or output.
→ RBAC (Role Based Access Control): RBAC is an access-control mechanism defined around factors such as authority, responsibility, and competency.

- → Semantic Layer: The semantic layer is the representation of data that helps different business end-users discover and access the right data efficiently, effectively, and effortlessly using common business terms.
- → SQL (Structured Query Language): A programming language for retrieving data from a relational database.
- → Stable Diffusion: It is a latent text-to-image diffusion model capable of generating photo-realistic images given any text input, cultivates autonomous freedom to produce incredible imagery, empowers billions of people to create stunning art within seconds.
- → Systems of Insight (SoI). It is the system used to perform data analysis from the data that is combined from the SoR or transactional systems.
- → System of Record (SoR). The authoritative data source for a data element. To ensure data integrity in the enterprise, there must be one — and only one — system of record for a data element.
- → Structured Data: Data that is organized according to a predetermined structure.
- → Unstructured Data: Data that has no identifiable structure, such as text, audio, video and images.
- → Vector Databases: A vector database stores, manages, and searches data as high-dimensional numerical vectors called embeddings. These vectors are mathematical representations of unstructured data like text, images, or audio, capturing their semantic meaning.

D2A2 (Digital, Data Analytics, and Artificial Intelligence) is a thought leadership firm designed to empower D2A2 leaders navigating the complex world of digital transformation, data analytics, and AI.

For more information, visit www.d2a2.ai

D2A2

DIGITAL DATA ANALYTICS
ARTIFICIAL INTELLIGENCE