Gartner

Gartner Insights

2026 Planning Guide for Analytics and Artificial Intelligence

Sumit Agarwal, Georgia O'Callaghan, Christopher Long, Wilco van Ginkel, Maryam Hassanlou, Zain Khan, Cuneyd Kaya, Raúl Arrabales

2026 Planning Guide for Analytics and Artificial Intelligence

13 October 2025 - ID G00837539 - 30 min read
By: Sumit Agarwal, Georgia O'Callaghan, Christopher Long, Wilco van Ginkel, Maryam Hassanlou, Zain Khan, Cuneyd Kaya, Raúl Arrabales
Initiatives: Analytics and Artificial Intelligence for Technical Professionals; Architect, Implement and Scale Data and Analytics Solutions; Generative AI Resource Center

> In 2026, organizations must maximize business value while accelerating AI and analytics implementations. To achieve this, data and analytics technical professionals should utilize all available data, design modular and flexible architectures, and implement robust governance to ensure trust.

Overview

Key Findings

- The convergence of analytics and AI enables organizations to analyze structured and unstructured data through conversational interfaces. While this approach is transformative, it depends on the maturity of AI agents, large language model (LLM) technologies and semantic architectures to deliver reliable analytical outcomes.
- More organizations are deploying generative AI (GenAI) use cases in production. Most implementations focus on productivity gains, but face challenges with user adoption and achieving measurable returns on investment.
- Heightened risk perception and regulatory scrutiny around GenAI have increased organizational concerns and delayed approvals and implementations. Risks stem from the technology's black-box nature and concerns around safety, data privacy and reliability.
- The pressure to leverage data for competitive advantage and better decision making, alongside rapidly evolving technology and products, has exposed gaps in organizational skills, solution adoption and risk awareness.

Recommendations

- Utilize structured and unstructured enterprise data to create a comprehensive view, leveraging AI agents and embedded models to drive agile decision making. Use an analytics mesh approach to ensure consistent analytics development, deployment and governance.
- Evolve beyond AI assistants to implement AI-driven workflows that deliver greater business impact and value. Carefully prioritize and select use cases to integrate within a business process workflow, while maximizing results.
- Strengthen AI governance by embedding controls directly into the architecture. Involve the architecture review board (ARB) in creating and maintaining a framework for evaluating architecture decisions and assessing AI technologies and products.
- Implement data and AI literacy programs, backed by effective change management, to boost adoption of AI-driven analytics and solutions. Harness the momentum of technology trends to motivate upskilling across the organization.

Analytics and Artificial Intelligence Trends

Analytics and AI capabilities are converging. They are transforming organizations' abilities and agility to extract value from data and make decisions. AI is now embedded in all major technology platforms, reshaping business operations. In 2026, AI will play a critical role as enterprises demand faster, smarter decisions and operational excellence.

According to the 2025 Gartner AI Survey: CIO and Technology Leader View, 82% agree that the pace of change within their organizations is accelerating rapidly, reflecting the speed of AI innovation. [1] New GenAI models and AI agent frameworks are launching continuously, and significant investments from technology leaders signal that disruption will only intensify.

This convergence presents several challenges and strategic opportunities for enterprises:

- Enterprise data is often messy and multistructured, while most analytics tools are designed for well-managed structured data.
- AI tools like Microsoft Copilot can deliver quick productivity gains, but actual business transformation occurs when AI powers entire workflows, not just isolated tasks.

- The AI landscape is expanding, along with expectations and budgets. Without disciplined, targeted investment, organizations risk wasting resources and missing out on ROI.

To stay ahead, organizations must upskill and reskill their teams; prioritize piloting new AI products and frameworks; and architect solutions for scale, flexibility and reliability.

As shown in Figure 1, these objectives align with key architectural and technical trends that Gartner believes will shape analytics and AI implementations in 2026 (click links below to jump to trends):

- Analytics and AI solutions will require a holistic view of structured and unstructured data.
- The transition from AI assistants to AI-powered workflows will demand a composite of AI techniques.
- Growing AI investments will require a strong focus on AI governance and AI literacy.

Figure 1: 2026 Key Trends in Analytics and Artificial Intelligence

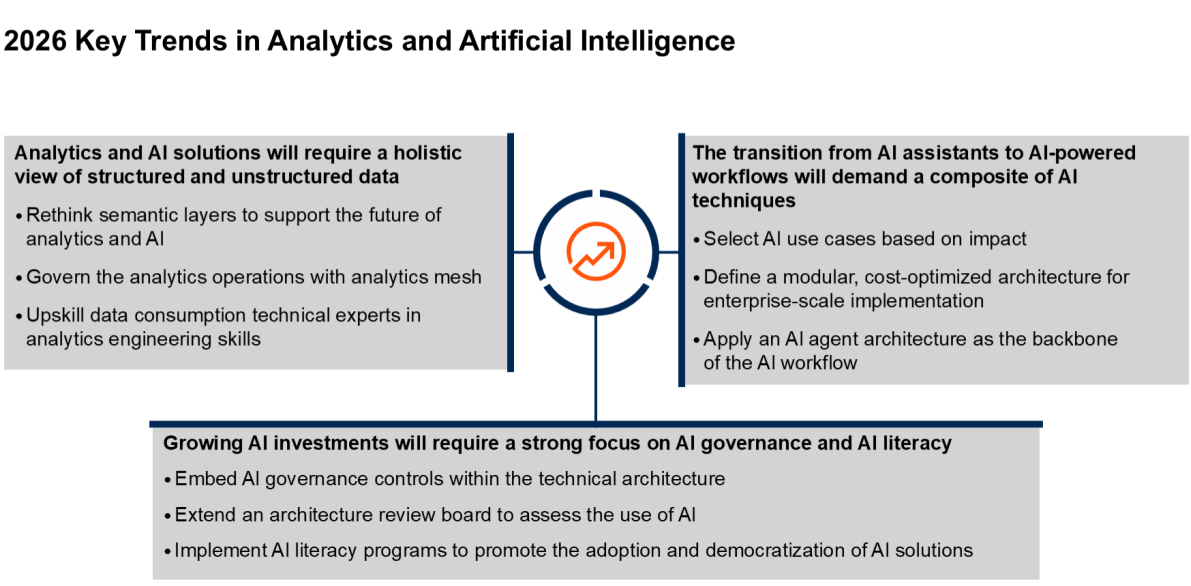

Source: Gartner
837539

Analytics and AI Solutions Will Require a Holistic View of Structured and Unstructured Data

Back to top

Enterprise data management strategies and technologies have traditionally focused on structured data. GenAI technologies, especially those using LLMs, now offer the potential to unlock significant value from unstructured data. Despite this advancement in AI capabilities, structured and unstructured data often remains siloed.

However, as technology continues to evolve, the ability to combine these data types and provide conversational interfaces is truly transformational. This shift offers organizations a new avenue from purely relying on interactive dashboards to prompt-based, question-driven analysis. Data analysts and business leaders can leverage this approach for agile, question-based decision making. It is also essential to recognize that the technology and solutions offering such a capability are still evolving and immature.

To implement solutions that enable access and analysis of multistructured data, analytics and AI professionals — particularly technical architects, data scientists and analytics specialists — should take the following steps (click links to jump to sections):

- Rethink semantic layers to support the future of analytics and AI.
- Govern the analytics operations with analytics mesh.
- Upskill data consumption technical experts in analytics engineering skills.

Planning Considerations

Back to top | Back to planning considerations list

Rethink Semantic Layers to Support the Future of Analytics and AI

A semantic layer is an abstraction layer that delivers a consistent and unified view of data from multiple sources. It translates complex data structures and technical terminology into familiar business terms and concepts. Traditionally, this has involved mapping various structured data sources to create a business-friendly perspective, hiding complex table joins and exposing fields using business language. As the solutions advance, the semantic layer expands to become a composite. Figure 2 illustrates a scenario where a knowledge graph supports the semantic layer to connect and model both structured and unstructured data.

Figure 2: Semantic Layer Powered by Knowledge Graph

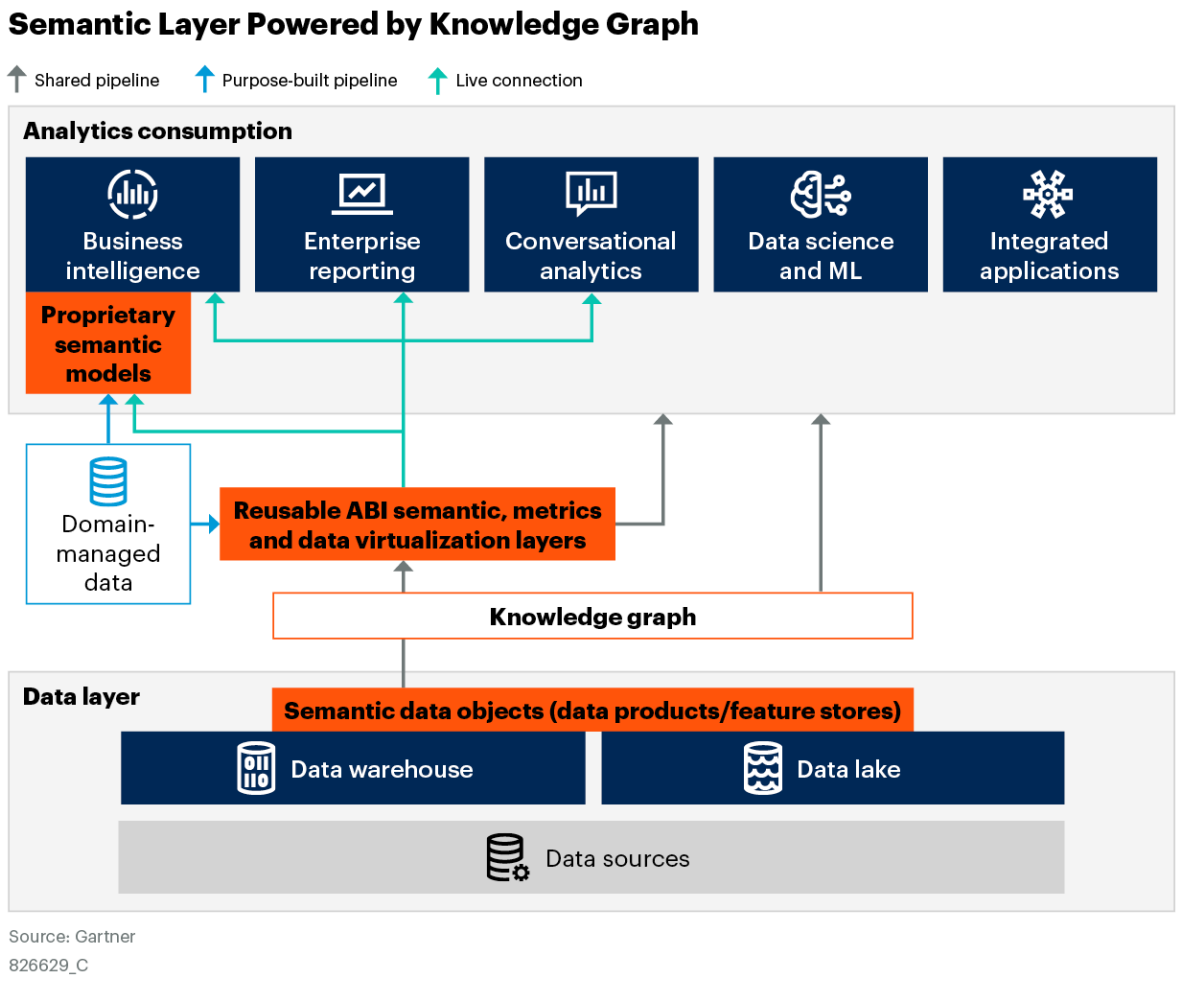

Source: Gartner
826629_C

Gartner

Historically, data management teams have overseen well-structured data sources, focusing primarily on structured data. However, the rise of GenAI and LLMs has transformed how organizations analyze and interact with unstructured content. Retrieval-augmented generation (RAG) has become the preferred solution for answering user questions from various organizational document types, including PDFs, PowerPoint presentations, Word documents and HTML files. LLMs also enable text-to-SQL capabilities, allowing users to convert natural language questions into SQL queries or generate Python code to access and analyze structured data.

Despite these advancements, an operational divide remains between querying structured data and analyzing unstructured data. The future of analytics lies in bridging this gap to provide a holistic view of all data. This approach aims to democratize data access, empowering users to ask complex, natural language questions spanning structured and unstructured data types.

The following prompt demonstrates a scenario where analyzing both structured and unstructured data is essential:

Provide a summary of customer complaints submitted last month for products that experienced a 10% sales drop compared to the previous month.

An LLM-based solution would break this query into several steps:

- Generate an SQL query to identify customers who submitted product reviews in the last month, filtered by products with a 10% decrease in sales compared to the previous month.
- Use this list of customers to retrieve their reviews.
- Apply a sentiment analysis using an LLM or BERT to identify the negative reviews or complaints.
- Invoke a RAG tool to extract and summarize the customer complaints.
- Combine the SQL and RAG query results to deliver a comprehensive summary.

To perform this analysis, the following components are required:

- **Semantic layer via knowledge graph:** For the initial step, the LLM must identify relevant database fields for product sales, sale dates, customer identifiers and complaint identifiers to generate the SQL query. A robust metadata mapping that links business definitions to database tables and fields is critical. The metadata should clearly define relationships between product and customer tables, which a knowledge graph or other semantic layer tool can manage. The semantic layer may also be extended to provide mapping between various entities in unstructured content and connect them to structured data.
- **AI agent:** The AI agent acts as the orchestrator, planning and sequencing the steps. It leverages an LLM with reasoning capabilities to break down the problem. The LLM prompt should include contextual information to guide its reasoning, or it may need fine-tuning with this mapping. The orchestrator agent can also call specialized agents: an SQL agent for text-to-SQL conversion, a RAG agent for executing RAG workflows, and a summarization agent to generate the final summary.
- **Sentiment analysis model:** This may be an LLM or a smaller BERT model that accepts a text segment as input and provides the sentiment prediction as an output.

The effectiveness of these components depends on the quality and completeness of metadata for all data sources and objects. Business teams must work closely with data management teams to continuously improve metadata, addressing gaps identified by AI teams during query response evaluations.

Related Research

- Rethink Semantic Layers to Support the Future of Analytics and AI
- How to Evaluate the Applicability of Knowledge Graphs for Your Use Cases
- How to Enhance RAG Performance Using Knowledge Graphs
- How GenAI Is Transforming Analytics and Business Intelligence Platforms
- What D&A Professionals Can Expect From Generative AI for Analytics

Govern the Analytics Operations With Analytics Mesh

Back to top | Back to planning considerations list

According to the 2024 Gartner Analytics and AI Engineering Survey, among organizations with a unified team responsible for analytics delivery across development, testing and production stages, only 32% of respondents confirmed that a centralized team performs these analytics functions. [2] For the remaining, analytics are managed jointly by a centralized team and business unit teams, or by single or multiple teams within business teams. This data supports the assertion that enterprises have consistently encouraged analytics to be decentralized and performed within the business teams.

This model has enabled the centralized IT teams to focus on establishing common frameworks, including semantic layers and metrics frameworks, and developing best practices for dashboard development. Despite this, self-service analytics has struggled with a lack of collaboration, reusability, varied delivery platforms, lack of interoperability, and, many times, dashboard sprawl. These challenges require rethinking and refactoring analytics development, deployment and governance approaches.

Analytics mesh combines federated architectures and governance with a product-centric delivery model.

Analytics mesh (see Figure 3) is an organizational model designed to address the challenges of scaling analytics operations across diverse and distributed business units or functions. This approach enables each domain to operate semiautonomously, while still adhering to overarching governance and interoperability standards. As a result, organizations can maintain consistency and compliance across all domains.

Figure 3: Analytics Mesh Framework

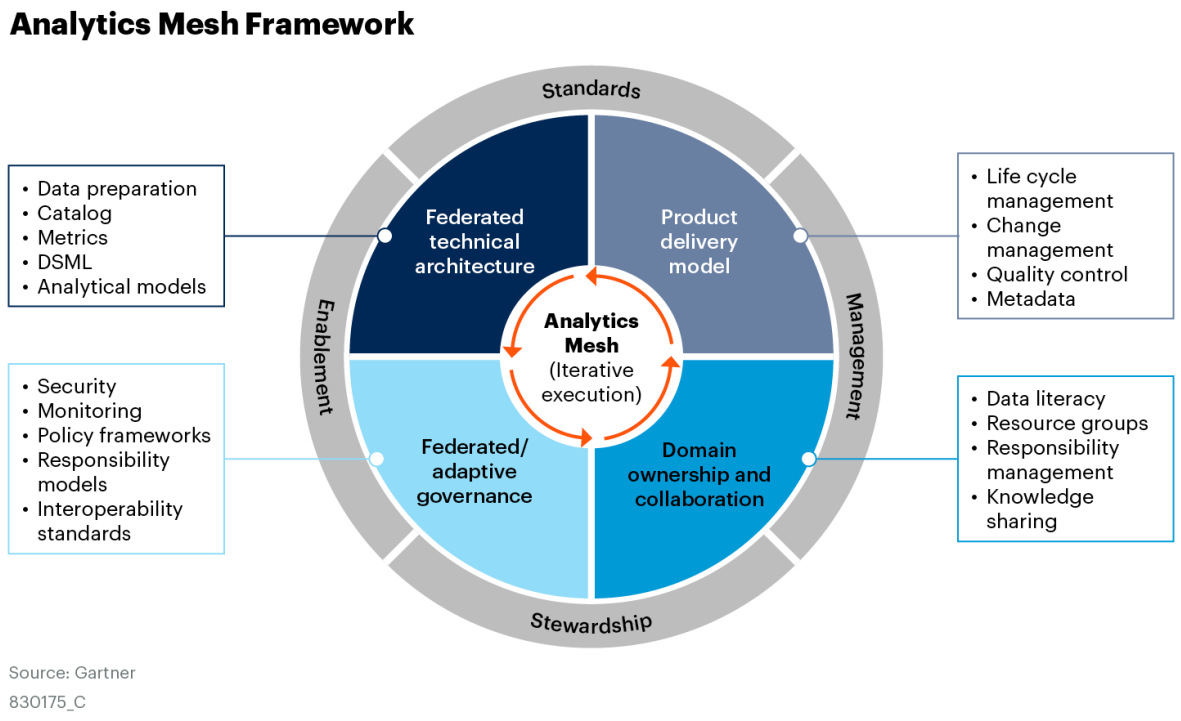

Source: Gartner
830175_C

Gartner

Organizations looking to optimize analytics development and delivery using an analytics mesh approach must focus on four foundational pillars, as illustrated in Figure 3:

- **Federated technical architecture:** Distributed, composable technical analytics capabilities systematically operationalized within business units
- **Federated/adaptive governance:** Coordinated, shared governance model for shared decision control
- **Domain ownership and collaboration:** Business units powered by domain and technical expertise, where consistent training and support are provided, and knowledge is shared

- **Product delivery model:** Ownership shift from artifact delivery to product-centric delivery and support.

While each of the four foundational pillars addresses specific challenges in distributed analytics, their true power emerges from their interconnectedness and synergistic operation as a cohesive framework. Organizations must strive to implement all four pillars concurrently and iteratively for an effective implementation.

Related Research

- Scale and Govern Your Analytics Operations With Analytics Mesh
- How to Govern and Scale Existing Self-Service Analytics Initiatives

Upskill Data Consumption Technical Experts in Analytics Engineering Skills

Back to top | Back to planning considerations list

As analytics platforms connect various structured and unstructured data sources, building a shared understanding of business context and data is crucial. Organizations can achieve this by implementing composite semantic layers, which may include knowledge graphs. Robust semantics will serve as the foundation for modern conversational analytics platforms.

Traditionally, data engineers have focused on building data pipelines to transform data for target architectures. In contrast, composite semantic layers require a different approach. They must provide a metadata map linking business context to technical data structures, demanding a deep understanding of business functions and data.

While data engineers excel at organizing and moving data, developing and maintaining these semantic layers creates a skills and role gap within data engineering teams. To address this, organizations need analytics engineers who blend technical data expertise with business insight to manage semantic layers.

Analytics engineers:

- Specialize in creating and scaling data assets, models and artifacts for analytics consumption, including metrics stores, data marts and semantic layers.

- Apply data engineering best practices — such as code versioning, continuous integration/continuous deployment (CI/CD), environment and artifact management, and orchestration — to analytics development.
- Have a thorough understanding of business metrics, KPIs and terminology.
- Typically work within domain analytics teams alongside BI administrators and self-service analytics developers, but may also be part of centralized data and analytics teams.

To close the skills gap, organizations, especially data and analytics teams, should encourage data engineers and data analysts to expand their understanding of business functions and learn methods for ingesting semantic mappings into knowledge graphs or semantic layer solutions. They should also become familiar with metrics stores and technical approaches for consistently populating KPIs and other measures, building on their expertise in data, SQL and other analysis tools.

Related Research

- Accelerate AI-Infused Analytics With the Analytics Engineer Role
- Establish the Essential Roles for Advanced Analytics and AI Initiatives

The Transition From AI Assistants to AI Workflows Will Demand a Composite of AI Techniques

Actionable, objective insights

Position your IT organization for success. Explore these additional, complementary resources and tools for data leaders and their teams:

Guide

Deliver on Strategic IT Priorities With These Essential Skills

Equip your team with skills to execute your digital strategy and drive enterprise outcomes.

Download Now

Insights

Data Architecture Solutions for Technical Professionals

Access action plans for the key trends for data and analytics architects.

Learn More

Webinar

10 Best Practices for Technical Professionals to Scale GenAI Across the Enterprise

Learn industry best practices from leading adopters to elevate your GenAI initiatives.

Watch Now

Tool

Gartner Data and Analytics Maturity Score for CDAOs

Assess and improve your function's readiness.

Learn More

Already a client?
Get access to even more resources in your client portal. Log In

Connect with us

Get actionable, objective business and technology insights that drive smarter decisions and stronger performance on your mission-critical priorities. Contact us to become a client:

U.S.: 1 855 322 5484

International: +44 (0) 3300 296 946

Become a Client

Learn more about Gartner for Technical Professionals
gartner.com/en/gartner-for-technical-professionals

Stay connected to the latest insights

Attend a Gartner conference
View Conference

Gartner®