

[WHITEPAPER]

AI in GoodData

Architecture, Management,
and AI Experiences for
Analytics

Contents

Executive Summary	01
[01] Why AI Changes Analytics	01
[02] The Critical Challenge: Why Analytics Is Hard for AI	02
[03] Why Generic AI Breaks in Production Analysis	03
[04] Why GoodData Is Different	03
[05] The Technical Foundation: A Bottom-Up View	04
[06] The Reasoning Layer: Skills and Agents	05
[07] From Interactive AI to Agentic Workflows	06
[08] Multi-Tenant AI at Production Scale	07
[09] Embedded AI for Customer-Facing Products	07
[10] Management and Scale: The AI Control Plane	08
[11] Open Integration and Interoperability	08
[12] Enterprise Trust, Governance, and Observability	09
[13] How Customers Apply GoodData AI	10
Conclusion	10

Executive Summary

AI is changing analytics from a system people navigate to a system that answers, explains, and acts. But in production, AI only works when it is grounded in governed metrics, business context, permissions, and operational control.

GoodData approaches AI as a native part of the analytics system itself. Rather than adding a generic assistant on top of dashboards, GoodData connects AI directly to the semantic layer, analytics logic, and execution model that already power production analytics. This ensures that AI can operate on trusted definitions, respect user and tenant boundaries, and produce outputs consistent with how the business actually measures performance.

WHY IT MATTERS

Analytics is not just a language problem. In real deployments, AI must correctly interpret governed metrics, apply business rules, leverage organizational knowledge, and remain manageable across thousands of environments. GoodData is built for this reality.

The result is an AI-native analytics platform that supports interactive assistants, embedded copilots, reusable skills, configurable agents, and agentic workflows — all operating on the same governed foundation.

[01]

Why AI Changes Analytics

For years, analytics has been organized around dashboards, filters, reports, and exploration tools. That model is not disappearing, but it is no longer the only interface that matters.

Users increasingly expect systems to answer questions directly, explain what changed, identify drivers, recommend next actions, and automate parts of the analytical workflow. Product teams increasingly want to embed AI-native analytics into customer-facing applications. External agents and AI applications increasingly need governed access to analytical capabilities through APIs and open protocols.

This shift raises the bar for what an analytics platform must do. It is no longer enough to visualize data well. The platform must also support AI reasoning on top of governed semantics, enforce permissions in every interaction, operate reliably across large multi-tenant environments, and expose analytics as composable capabilities that can participate in broader AI systems.

This is where generic AI approaches begin to fail.

The Critical Challenge: Why Analytics Is Hard for AI

It is easy to make AI look useful in a demo. A chart can be summarized. A trend can be described. A question can be answered in fluent language.

The difficulties begin when AI is expected to work in production.

In production analytics, AI must use the organization's actual definitions of "revenue," "retention," "pipeline," or "active customer," not the nearest matching field in a database. It must respect permissions, tenant boundaries, entitlement rules, business policies, and product-specific constraints. It must understand that the same word may carry different governed meanings in different contexts. It must reason from both structured metrics and unstructured organizational knowledge. And it must do all of this consistently, not just persuasively. A fluent answer can still be wrong in ways that matter to the business.

GoodData addresses this challenge by grounding AI across three dimensions:

1.

Data:

AI must operate on a reliable analytical execution layer with governed access to structured data.

2.

Decisions:

AI must reason through business rules, prior analytical signals, and governed logic rather than rediscovering meaning from raw tables.

3.

Domain Knowledge:

AI must use documentation, organizational language, policy, and historical context so outputs reflect how the business actually works.

This grounding model is what turns AI from an impressive interface into a production capability.

Why Generic AI Breaks in Production Analytics

Generic AI layers often fail in production analytics for predictable reasons:

- They reason from raw schema instead of business meaning.
- They retrieve plausible but ungoverned fields instead of governed metrics.
- They ignore permissions complexity.
- They lack tenant awareness.
- They treat company documentation as optional rather than necessary context.
- They may produce answers that sound coherent while being analytically inconsistent or operationally unsafe.

In simple environments, these weaknesses may be tolerable. In embedded analytics, enterprise deployments, and multi-tenant products, they are not.

THE GOODDATA PRINCIPLES

GoodData is designed to solve exactly these problems. **AI does not bypass the analytics platform. It works through it.**

Why GoodData Is Different

GoodData's approach to AI is shaped by a set of structural choices that make the system usable in real analytical environments.

■ Governed semantic foundation

GoodData connects AI directly to the semantic layer that defines metrics, attributes, relationships, and business logic. This allows AI to reason in business terms rather than technical structures and ensures outputs remain aligned with the same definitions used elsewhere in the platform.

■ AI-ready as-code the foundation

GoodData is built on an as-code and API-first architecture. Analytics assets — from logical data models to metrics — can be versioned, reused, and deployed through standard engineering workflows. A platform that exposes analytics through code is naturally well-suited for AI-assisted development and fast delivery into production.

■ Multi-tenant by design

GoodData is built for environments where one platform must serve many customer organizations, each with its own users, rules, context, and analytical experience. AI for one company is a feature; AI across thousands of customer environments is an operating model.

■ Operational control, not just model access

Once AI moves beyond a demo, teams need to control what agents can do, what knowledge they can use, how behavior is configured, and how changes are rolled out. GoodData includes a centralized operating model for AI management, governance, and observability.

■ Open AI ecosystem

Organizations do not want to lock themselves into a single model or interface pattern. GoodData supports flexible deployment, bring-your-own LLM approaches, and integration through APIs, MCP, and A2A, allowing customers to evolve their AI strategy without replacing their analytics foundation.

[05]

The Technical Foundation: A Bottom-Up View

To understand how GoodData AI works, it helps to start with the platform layer underneath everything else.

Infrastructure and performance

At the base of the system is an **open analytical infrastructure** layer responsible for connectivity, federation, execution, and performance. This includes data connectivity and query execution capabilities that provide the speed and reliability AI requires. If the execution layer is weak or retrieval is slow, the AI layer becomes unreliable regardless of model quality.

Semantics and builder enrichment

The **semantic layer** is where data becomes analytically usable. Metrics, facts, attributes, and relationships define business meaning. AI operates through this semantic system rather than against raw structures.

Builders enrich this foundation through governed object design, naming, descriptions, and catalog management. The quality of AI output depends heavily on the quality of the semantic layer underneath it. Clear definitions and well-maintained analytical objects improve both precision and trust.

Context management: the grounding layer

Semantics define meaning, but semantics alone are not enough. AI also needs context.

GoodData's **context management** layer brings together governed analytical objects, durable memory, and searchable organizational knowledge so agents can reason from the right supporting information.

This includes:

Analytics Catalog	A governed source of metrics, objects, definitions, and descriptions that communicate business intent.
AI Memory	A structured key-value store for durable context, user preferences, organization-wide terms, and memory policies.
AI Knowledge	A search-first knowledge layer that transforms internal documents such as PDFs, markdown, and text files into searchable knowledge for grounded responses.

Together, these layers allow AI to combine structured analytical logic with relevant business context.


[06]

The Reasoning Layer: Skills and Agents

GoodData separates analytical capability from reasoning control. This makes it possible to support many AI experiences without rebuilding the system for each one.

Modular skills

Skills are reusable analytical building blocks — opinionated recipes built on top of platform tools. A skill represents a predefined path from input to output. These capabilities can include:



Because skills are modular, new capabilities can be added to the platform without redesigning the whole experience.

Configurable agents

Agents are the configurable reasoning units that guide how skills are used. An agent receives an assignment, determines the necessary steps, and activates the required skills.

They can support read-only diagnostic and explanation scenarios, or write-capable builder scenarios that help create and modify analytical assets. This model supports both conversational and operational AI while preserving governance and administrative control.

Decision-based governance and control

Agent behavior can be shaped by captured decisions and explicit rules. If a workflow needs to change, proposals can be reviewed by humans before taking effect, helping the system improve through supervised oversight rather than uncontrolled drift.

[07]

From Interactive AI to Agentic Workflows

AI in analytics often begins with interactive experiences such as assistants and copilots. These help users explore data, understand results, and work more naturally with analytics in the moment.

The next step is extending AI beyond a single interaction. In production environments, AI also needs to support repeatable analytical processes: detecting issues, surfacing recommendations, and contributing to workflows that run on a schedule, on a trigger, or as part of a broader business process.

CONTROL DOES NOT DISAPPEAR

As AI becomes more operational, organizations need clear governance, traceability, and human oversight to ensure that recommendations and actions remain trustworthy.

Assistants and copilots

Assistants support conversational exploration and explanation. Copilots provide contextual help inside specific work moments, such as modeling, dashboard building, or interpretation workflows.

Operational AI and workflows

Moving beyond manual prompts, workflows can run on a schedule or trigger to detect problems and surface recommendations automatically. Through the Agent Execution API, configured agents can be executed programmatically with structured input and output, allowing analytical AI to participate in external scripts, products, and operational systems.

Governance and human oversight

Operational AI should not mean uncontrolled autonomy. In real deployments, recommendations, changes, and actions often need review, traceability, and clear approval paths. GoodData is designed for this governed middle ground: more automation where it helps, with control where it matters.

[08]

Multi-Tenant AI at Production Scale

AI for a single internal team is one problem. AI across thousands of customer environments is another.

In multi-tenant deployments, AI must not only answer accurately — it must do so while preserving strict tenant isolation, respecting tenant-specific context and rules, handling different user groups and capability tiers, and remaining centrally manageable across many deployments.

GoodData architecture supports tiered AI delivery, where different capabilities can be assigned to specific users or groups. Upgrading a customer's AI experience can happen through configuration and reassignment rather than code changes or redeployments. This makes large-scale rollout more manageable and more commercially practical.


[09]

Embedded AI for Customer-Facing Products

GoodData AI is not limited to internal analytics teams.

A major use case is embedding governed AI analytics into customer-facing software products. In this model, product and engineering teams can deliver AI-powered analytical experiences directly inside their own applications while retaining control over the user experience, product logic, and workflow design.

GoodData provides the governed analytical intelligence underneath:



This is what allows AI analytics to move beyond internal BI and become part of the product itself.

[10]

Management and Scale: The AI Control Plane

Once AI moves into production, model quality is only part of the problem. The larger challenge is operational management.

Organizations need a central place to control which agents exist, which skills they can use, what knowledge sources they can access, how behavior is configured, which user groups receive which capabilities, and how changes are rolled out over time.

AI HUB

GoodData provides this through a centralized **AI management** layer — the AI Hub — which serves as the control plane for operational AI. It gives administrators and platform teams a deliberate operating model for rollout, change management, capability assignment, and governance across complex environments.

Without that control layer, AI becomes a collection of exceptions. With it, AI becomes something teams can run deliberately.

[11]

Open Integration and Interoperability

GoodData does not confine AI to a single interface or channel. The platform supports integration through APIs, MCP, and A2A.

MCP	Exposes analytics capabilities as callable tools that external agents and AI applications can use directly.
A2A	Supports agent-to-agent communication when work spans multiple systems.
APIs	Provide the broader integration surface that allows governed analytical capabilities to participate in products, workflows, and operational environments.

The value is not just that GoodData speaks these protocols. The value is that what sits behind them is governed semantics, controlled execution, and multi-tenant-aware access to business meaning.

[12]

Enterprise Trust, Governance, and Observability

For AI to be trusted in analytics, it must be governed end to end.

■ Policy enforcement and guardrails

Administrators must control what tools, skills, and knowledge each agent can access. Governance needs to apply across input, reasoning, and output so AI behavior remains deliberate and bounded.

■ Traceability and explainability

Teams need visibility into how an answer or recommendation was produced. This includes step-by-step traces, the specific skills invoked, and the underlying inputs used during execution.

■ Auditability

Administrative actions such as creating agents, changing skill allowlists, or modifying knowledge configurations should be recorded as auditable events for compliance and operational accountability.

■ Reproducibility and context clarity

Summaries and findings should be stamped with the relevant filter and drill context at the time of the request, so analytical outputs remain clear, reproducible, and unambiguous.

■ Safe rollout at scale

AI capabilities need to be deployable gradually across environments, groups, and customers. Governance is not only about restricting AI — it is also about introducing it safely and deliberately.

How Customers Apply GoodData AI

Interactive analytics	Customers use GoodData AI to let business users ask questions in natural language, receive governed answers, understand KPI changes, and get narrative explanations grounded in trusted metrics and business definitions.
Analytical reasoning and investigation	Customers apply GoodData AI to detect anomalies, investigate performance changes, identify likely drivers, explore trends, run what-if analysis, and support forecasting and guided analytical reasoning.
AI-assisted creation	Analytics builders and product teams use GoodData AI to accelerate the creation of dashboards, metrics, insights, and analytical workspaces through guided, agent-assisted workflows.
Operational and workflow automation	Customers use GoodData AI to run analytical workflows on a schedule or trigger, detect issues automatically, surface recommendations, and return structured outputs that can feed broader operational processes.
Embedded AI experiences	Product and engineering teams embed GoodData AI into customer-facing applications to deliver governed analytical experiences directly inside their own products while maintaining control over the interface and workflow.
Open agent ecosystem integration	Customers expose GoodData capabilities to external agents and AI systems via MCP, coordinate work across systems via A2A, and integrate governed analytics into broader applications and workflows via APIs.

Conclusion

AI will not transform analytics by becoming a better chatbot on top of dashboards, but by becoming a governed execution layer that can answer, explain, recommend, and act on top of trusted business meaning.

By grounding AI in the semantic layer, combining structured analytics with organizational context, exposing capabilities through modular skills and configurable agents, and managing the whole system through a centralized control plane, GoodData makes AI usable in the environments where analytics actually matters: production systems, embedded products, and multi-tenant platforms at scale.