

Google Cloud

Data Agents Guidebook



Table of contents

01

What are data agents?

03

Definitions, core capabilities, and key components that differentiate data agents from standard AI models.

02

Why the hype around agents?

05

Technological advancements driving the rise of agents, and the evolution of multi-agent workflows.

03

Google Cloud's data agents

07

Google Cloud's pre-built agents for data engineers, data scientists, business users, developers, and data stewards.

04

Build your own data agents on Google Cloud

11

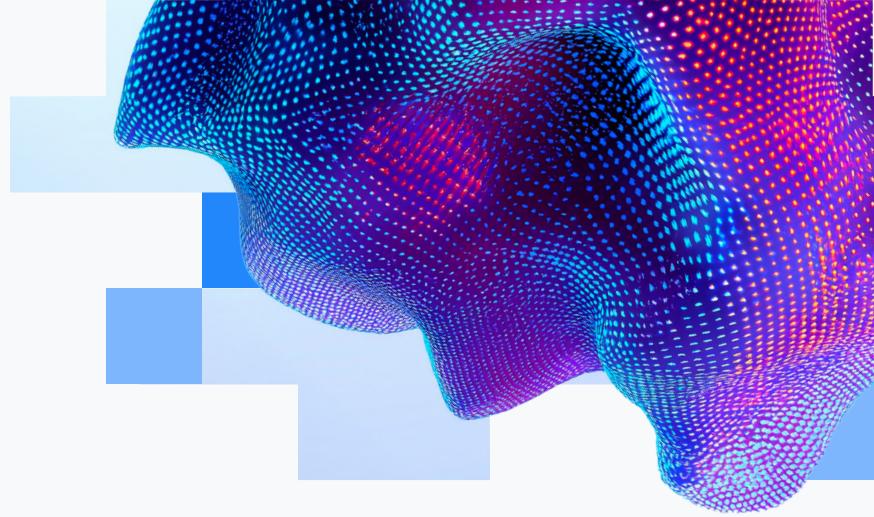
A practical guide to building a custom data agent, from raw data to deployment and optimization.

05

Build on a foundation of AI leadership

24

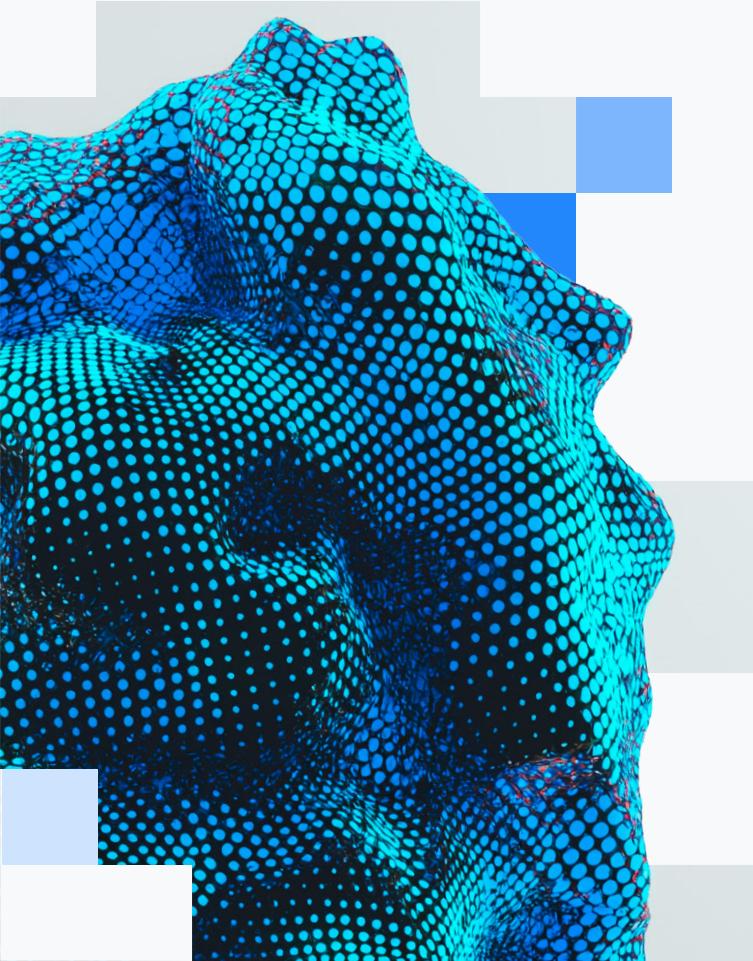
The strategic advantages of Google Cloud for building and scaling agents.





What are data agents?

A data agent is an intelligent, goal-oriented system that is designed to solve complex enterprise data problems, acting as an autonomous partner from start to finish.



Data agents have five core capabilities:

01

Understanding deep context – for example, schema, permissions, or jargon.

02

Gathering and synthesizing information from diverse sources.

03

Acting as goal-oriented problem-solvers with multi-step plans.

04

Planning and executing with alternative approaches.

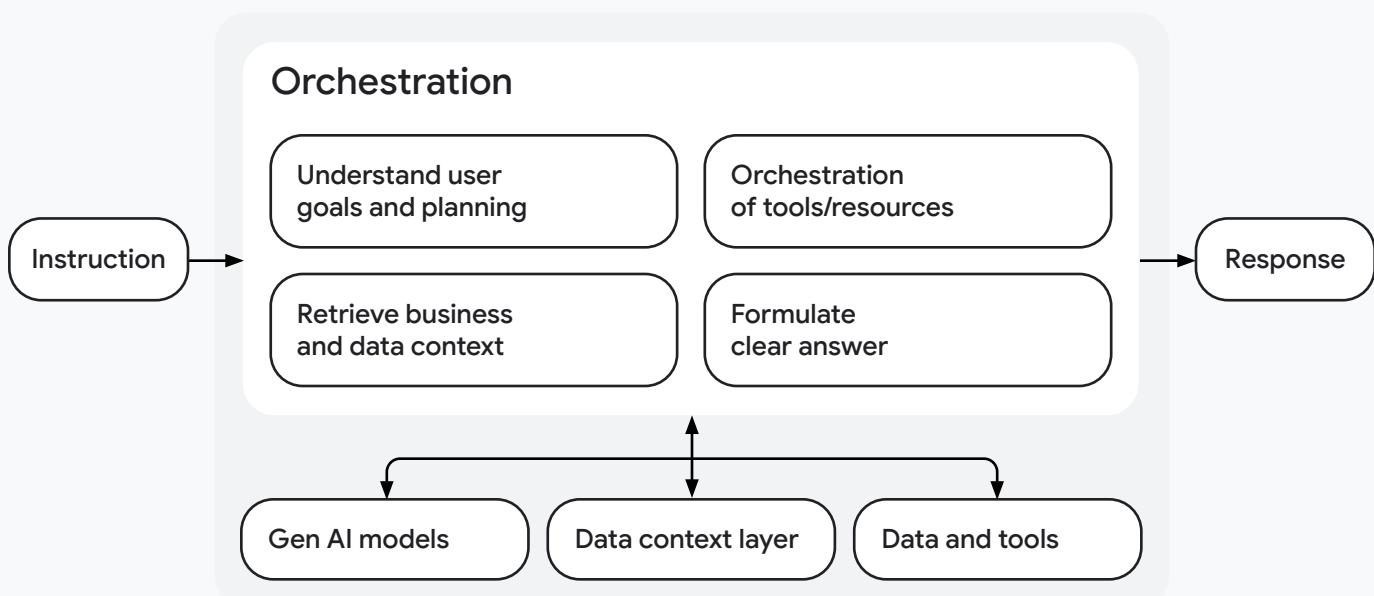
05

Leveraging various tools, including machine learning (ML) models and other agents.



A data agent is a coordinated system with an AI reasoning engine, a data context layer, and an orchestration engine that selects optimal tools for comprehensive outcomes.

For example, a data agent responding to “*Why did sales dip in the Northeast?*” would generate and run queries across sales, marketing, and Customer Relationship Management (CRM) data, then synthesize a summary like: “*The sales dip correlates with a 50% reduction in digital ad spend for the Northeast region,*” with supporting visualizations. This demonstrates the true power of a data agent.





Why the hype around agents?

Organizations are constantly demanding higher levels of automation, efficiency, and data-driven decision-making. With the introduction of **AI agents** – enabled through rapid technological advancements including more powerful foundation models, specialized data capabilities, and sophisticated reasoning engines – these demands can now be met.

Such AI agents have evolved from single-task models to sophisticated systems that are capable of complex reasoning and orchestrating multi-agent workflows to automate entire processes.

Example text	Agent solutions				
<p>Large Language Model (LLM)</p> <p>Reword a sentence to sound more polite</p> <p>Reword this to sound more polite: I don't have time to help you with that.</p> <p>I would love to assist, but unfortunately, I'm fully booked at the moment.</p>	<p>Fine-tuned model</p> <p>Answer a query about insurance policy terms</p> <p>What are the terms and conditions of the XYZ insurance policy.</p> <p>XYZ insurance policy includes...</p>	<p>Retrieval Augmented Generation (RAG)</p> <p>Answer a query about latest market sentiment</p> <p>What is the latest market sentiment based on the uploaded data...</p> <p></p>	<p>Single agent workflow(s)</p> <p>Create an IT ticket for a technical staff to address</p> <p></p> <p>Ticket #INC - 5542 has been created in ServiceNow. Priority set to 'Medium' and assigned to the Network team.</p>	<p>Multi-agent workflow(s)</p> <p>Write a credit memo – no human input beyond initial setup</p> <p></p> <p>Initiate the credit approval process for the new vendor application #882.</p> <p></p> <p>Credit Memo Finalized. Final document saved to ERP system and sent for signature.</p>	

Complex reasoning and adaptability



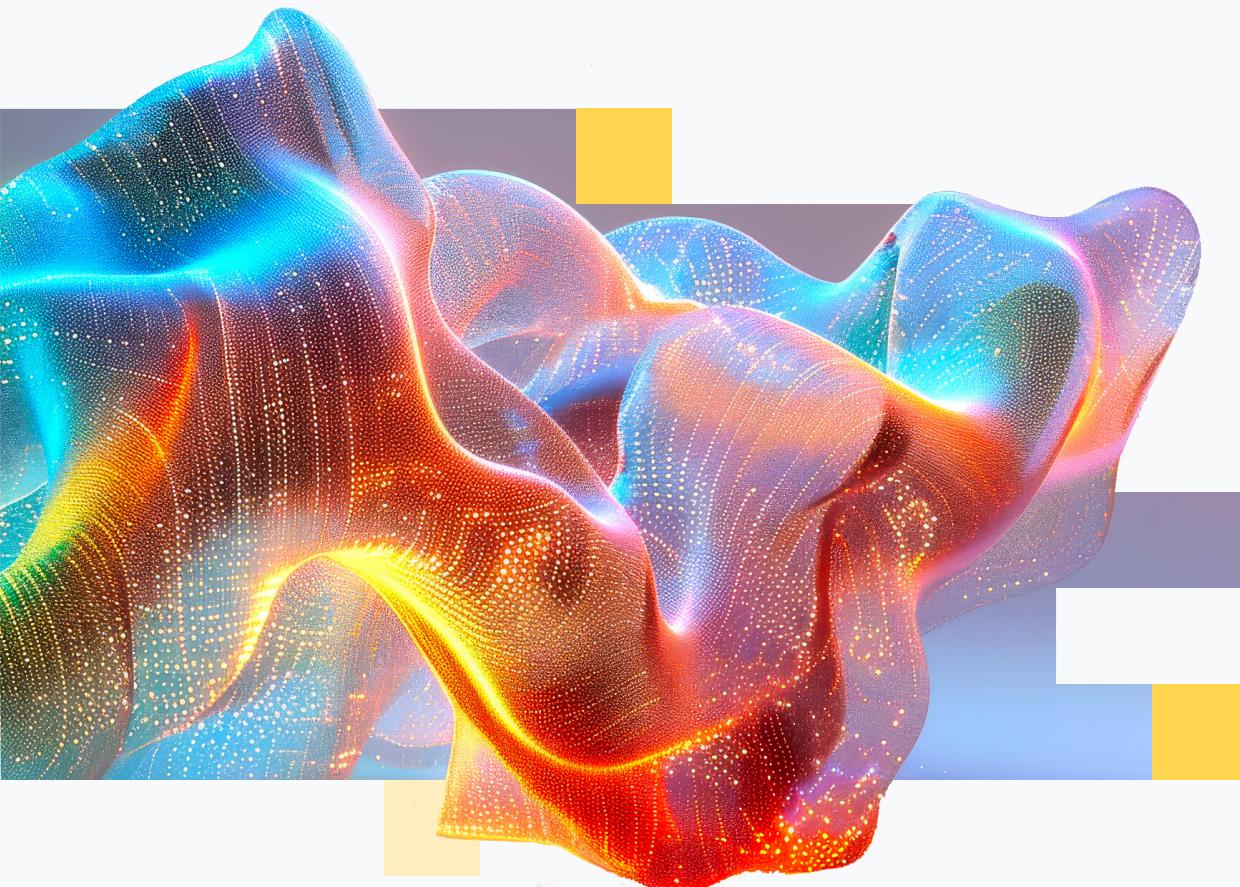
In a June 2025 report, Gartner predicted that by 2027, augmented analytics will become autonomous, managing 20% of business processes with proactive, collaborative, contextual, and continuous benefits.

Source: [Gartner](#)

With data analytics teams increasingly stretched, a shift to greater automation is crucial. Data agents offer a powerful solution by extending teams' capabilities:

They provide self-service conversational and event-driven analytics for **business users**, act as tireless exploratory data analysis assistants for **data analysts**, automate ML lifecycles for **data scientists**, translate natural language to SQL queries and data pipelines for **developers** and **data engineers**.

By handling repetitive tasks, agents shift focus from execution to strategic oversight and innovation.





Google Cloud's data agents

Google Cloud understands that every organization's AI journey is unique. That's why our approach to data agents is built on two core principles: power and choice. Instead of a one-size-fits-all solution, Google Cloud provides a suite of pre-built, ready-to-use agents designed for specific roles that allow your teams to gain value immediately, and offers a robust set of tools for those who need to build fully custom agentic experiences.

In this chapter, we will explore Google's out-of-the-box data agents that are designed to assist and automate workflows for every data user.

Data agents

Data analyst
Query & analyze



Business user
Gain insights & make decisions



Developer
Build intelligent apps



ML engineer
Deploy scalable models



Data scientist
Build models that work



Data engineer
Get clean, useful data





Data Engineering Agent

Data engineers are constantly challenged by the manual effort required to build, fix, and modernize data infrastructure. The toil of authoring complex pipelines, troubleshooting failures, migrating legacy systems, and handling complex data ingestion consumes time that could be spent on high-value innovation.

Google's [Data Engineering Agent](#) is an intelligent partner that automates this entire lifecycle. For authoring, it builds multi-node pipelines and creates code documentation. For troubleshooting, it finds the fix for broken jobs that you direct it toward.

For migrating systems, it translates legacy Spark or on-premises code into standardized, modern pipelines to accelerate the migration process. Finally, for simplifying complex data ingestion it parses industry-specific formats and fetches data from APIs.

By automating these tedious tasks, the agent frees engineers to focus on more strategic work and enables a massive boost in productivity that allows teams to deliver higher-quality data, faster, and at scale.

Sample prompts

"Extract data from the sales.customers table in the us_west_1 region, and load it into the reporting.dim_customers table in BigQuery. Match the schema of the destination table."

"Create a pipeline with the following steps:

Extract data from the ecomm.orders table

Join the extracted data with the marts.customers table on customer_id

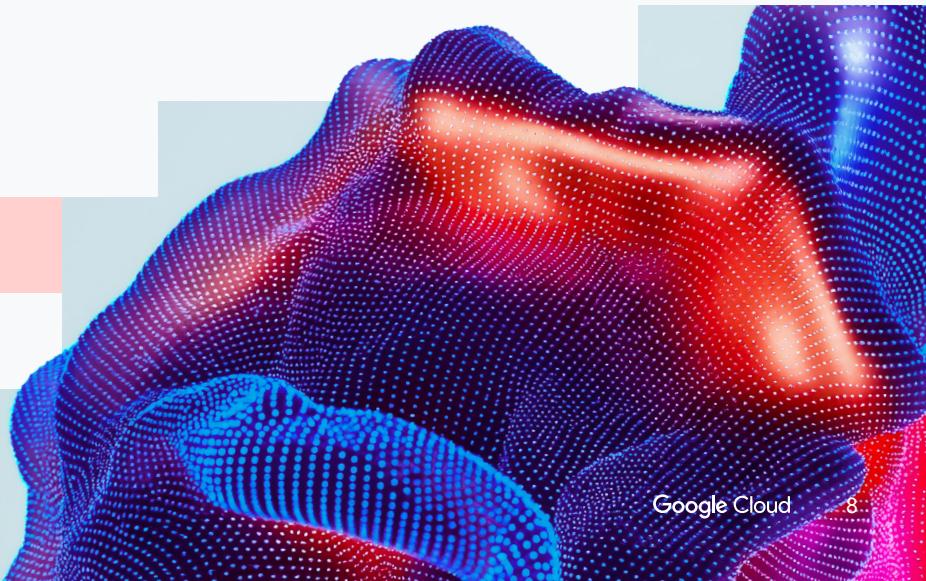
Load the final result into the reporting.customer_orders table."

Demo



How the Data Engineering Agent uses AI to generate SQL and create data pipelines.

[Watch demo](#)





Data Science Agent

Data scientists are often slowed by manual, repetitive tasks that hinder productivity and delay insights. Extensive data wrangling, detailed exploratory analysis, and complex, iterative model development consume valuable time.

Google's [Data Science Agent](#) is integrated directly into Colab Notebooks and streamlines these time-consuming tasks. From a single interface, Data Science Agent offers Gemini-powered AI assistance to automate data exploration, generate visualizations, and write Python or SQL code for you. It also accelerates your path to production with unified access to BigQuery and Vertex AI, and brings powerful tools like pre-trained TimesFM for forecasting and in-built model evaluation directly into your Colab Notebook.

By automating this laborious set up work, the agent allows data scientists to move seamlessly from raw data to production-ready models. This frees data scientists to focus on what matters most: refining model accuracy and delivering high-impact business insights, faster.

Sample prompts

"Create and evaluate a classification model on bigquery-public-data.ml_datasets. census_adult_income using BigQuerySQL. Using SQL, forecast the future traffic of my website for the next month based on bigquery-public-data.google_analytics_sample.ga_sessions_. Then, plot the historical and forecasted values."*

"Create a pandas DataFrame for the data in project_id:dataset.table. Analyze the data for null values, and then graph the distribution of each column using the graph type. Use violin plots for measured values and bar plots for categories."

Demo



How the Data Science Agent helps detect anomalies in datasets

[Watch demo](#)



Conversational Analytics Agent

For business users and data analysts, traditional dashboards often can't answer every question, leading to delays and reliance on technical teams for ad-hoc reports. This friction means that valuable insights remain locked away, slowing down the pace of business.

Google's [Conversational Analytics Agent](#) transforms this experience by allowing you to simply "chat with your data." Powered by Gemini and grounded in Looker's trusted semantic layer, this agent lets users ask questions and have follow-up engagements in natural language that lead to instant and reliable answers, charts, and data explorations. This moves beyond static reports by enabling a dynamic, interactive dialogue with your data.

By democratizing data access, the agent empowers business teams to self-serve trusted insights and frees up data analysts from mundane tasks to focus on high-impact work.

Sample prompts

"What is the trend of product returns over time, segmented by product department?"

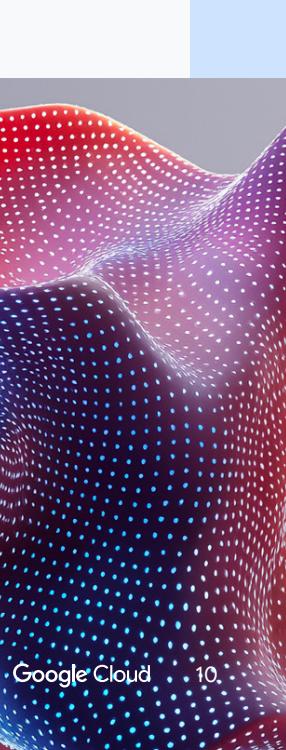
"How has the average price of products in each department changed year over year?"

Demo



Ask questions in plain language and receive answers directly from your BigQuery data

[Watch demo](#)





Build your own data agents on Google Cloud

Every business is unique, so pre-built agents don't always fit an organization's needs. Building a custom data agent offers a powerful, tailored solution, but the path from a promising demo to a production system can be challenging. Development is complex, but productionizing is even harder – with hurdles ranging from ensuring security and managing costs to delivering consistent, reliable results.

Google Cloud demystifies this journey with a structured five-step path. This blueprint lets you build, deploy, and manage robust, scalable agents with confidence, to turn your vision for a custom solution into a production reality.

01

Context

Provide your agent relevant information

02

Setup

Frameworks, tools, models and guardrails

03

Deploy

Run, scale, and monitor your agent

04

Publish

Make agents accessible

05

Analyze

Agent usage, cost, and performance





A quick check – are you ready to build a data agent?

Before diving into development, it's crucial to ensure that your project is set up for success. A great agent requires more than just good technology – it needs a solid organizational and data foundation.

Start by asking a few fundamental questions:

01

Have you identified a clear business use case with measurable outcomes?

02

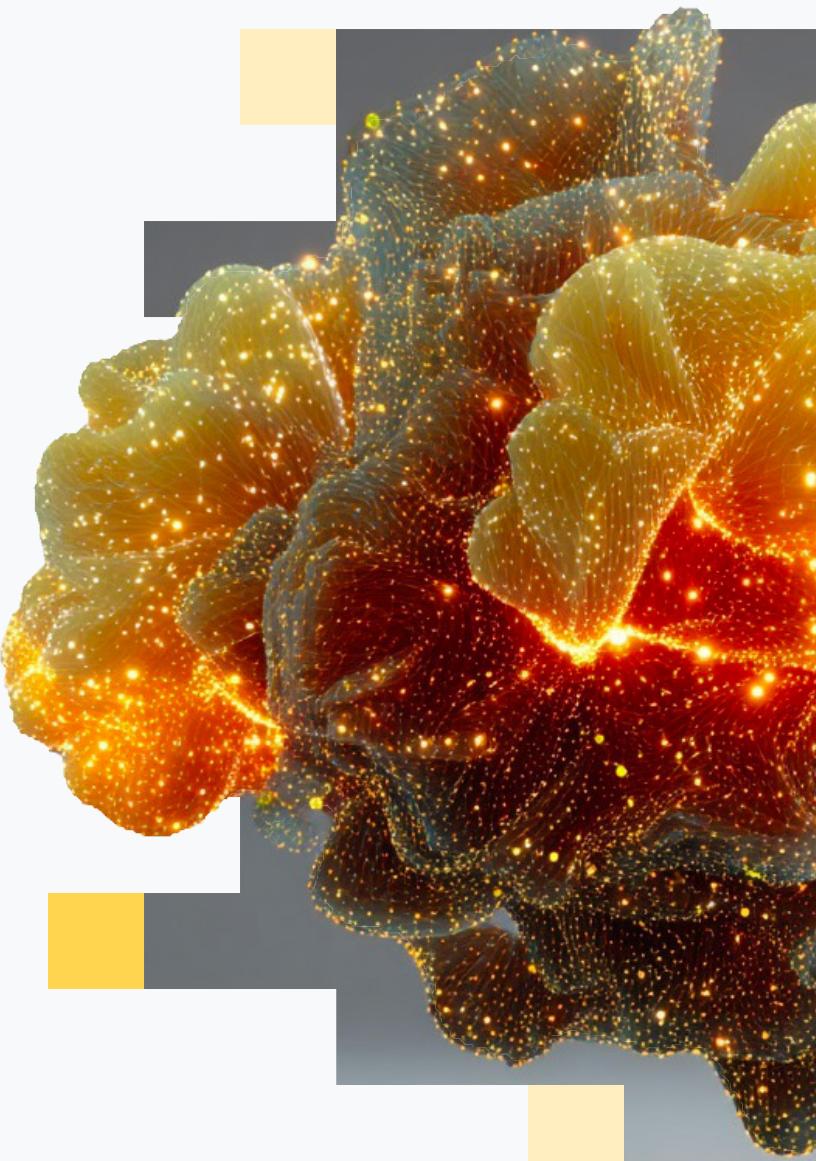
Do you have both business and data leaders ready to sponsor the agent?

03

Do you have the in-house expertise to build an agent, or will you need support?

04

Is your data well-documented and your business context understood? A willingness to invest in preparing your data is often the biggest predictor of success.





Provide rich business context

The success of any data agent hinges on one critical element: context. Without context, even the most powerful AI model is navigating without a map. To achieve accurate, reliable results, you need to create a “business context layer” – an essential bridge that connects your raw data to the agent’s reasoning engine.

Google Cloud provides a suite of tools designed to build this layer with maximum automation.



[Metadata augmenting and cataloging](#): Retrieve metadata for Google Cloud resources – such as BigQuery, Cloud SQL, Spanner, Vertex AI, Pub/Sub, Dataform or Dataproc Metastore – and third-party resources that you bring into Dataplex Universal Catalog, for an instant data catalog. Augment agents with operational data for real time enterprise truth.

[Data discovery](#): Scan Databases such as AlloyDB and Bigtable for structured data, and Cloud Storage buckets for unstructured data to extract and catalog their metadata.

[Semantic layer](#): Manage business-related terminology and definitions across your organization. Create calculations, relationships, and logic tailored to the unique AI algorithms and models. Translates complex data into business terms, allowing data exploration and visualization using natural language.

[Data insights](#): Use AI to generate natural language questions about your data to help uncover patterns, assess data quality, and perform statistical analyses.

[Data profiling](#): Identify common characteristics of column data in your BigQuery tables – for example, typical data values, data distribution, and null counts – which can inform data classification and quality assurance.

[Data quality](#): Define and measure the quality of data in your BigQuery tables by validating data against organizational policies and logging alerts if data doesn't meet quality criteria.

[Data lineage](#): Track how data moves through your systems: where it comes from, where it is passed to, and what transformations are applied to it.



Input

Question

Context

Reasoning Engine



Capabilities

Output

Messages

Detect Schema

Reason and reflect

Forecast

Sample data

Query data

Key driver analysis

Fuzzy search

Generate charts

Detect anomalies

Golden Queries

Provide insights

Period comparison

Data connectors



Looker



BigQuery



AlloyDB



Data and
business context



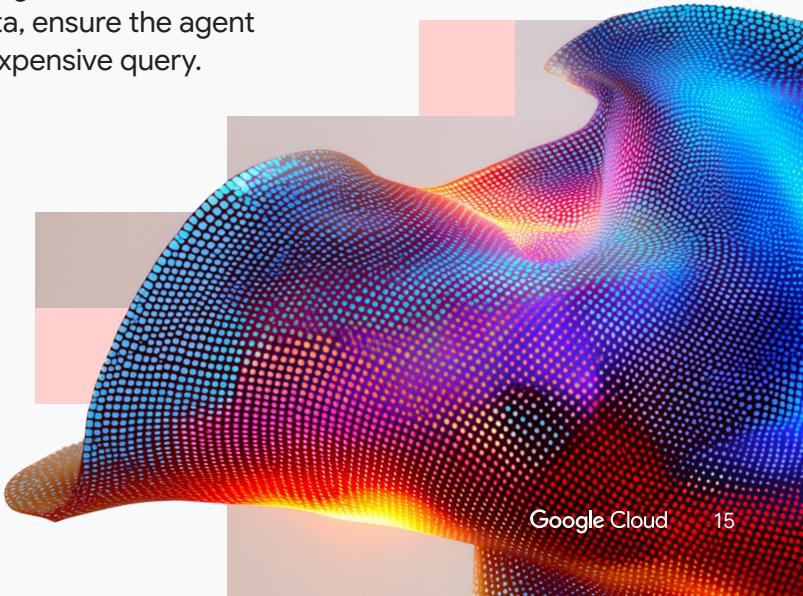
By defining table relationships, providing few-shot query examples, and writing clear instructions, you give the agent the guardrails that it needs to understand your unique business landscape. This automated approach transforms messy, cryptic data into a trusted foundation, which is the single most important step in building an agent that consistently delivers value with impact.

Instructions and information to guide your data agent

Instructions and overall workflow 	Few-shot examples 	BigQuery table context 
<p>What: High-level process and best practices you want the Agent to follow</p> <p>How: Write detailed instructions /prompts in a text-based config file</p>	<p>What: 5-10 SQL examples of typical queries to give the Agent</p> <p>How: BigQuery Dataset Insights / Knowledge Engine and Example Store</p>	<p>What: Get information about the table(s) – schema, column descriptions, metadata, data types etc</p> <p>How: BigQuery commands generated by the model</p>
Use-case specific info 	Table relationships 	Data profiles 
<p>What: Common questions, use-case terminology, important fields</p> <p>How: Manually collected from SMEs and Dataplex Glossary</p>	<p>What: How are tables related – entity relationships, join keys, types etc.</p> <p>How: Pull data from BigQuery Data Insights</p>	<p>What: Data profile results (column stats, value distribution, distinct values, row count, table size, data patterns etc.)</p> <p>How: BigQuery Data Profiles</p>

Once you have provided the right context, also set up guardrails that let you: monitor what the agent is doing with your data, ensure the agent doesn't delete your data, and avoid running a super expensive query.

Deep dive blog 
10 tips to safeguarding your data, and wallet, from agents





Build and configure your data agent

Once your context is established, the next step is to configure your agent's core components. Google Cloud's approach is built on flexibility, centered around APIs and open-source technologies such as:

[Conversational Analytics API](#): Lets you embed natural-language query functionality in your agents, or workflows, all backed by trusted data access and scalable, reliable data modeling. It's the same API that powers the out-of-the-box conversational experiences in Looker and BigQuery.

It allows you to build custom data experiences that provide data, chart, and text answers while leveraging Looker's trusted semantic model for accuracy or providing critical business and data context to agents in BigQuery. You can embed this functionality to create intuitive data experiences, enable complex analysis via natural language, and even orchestrate other conversational analytics agents.

Codelab



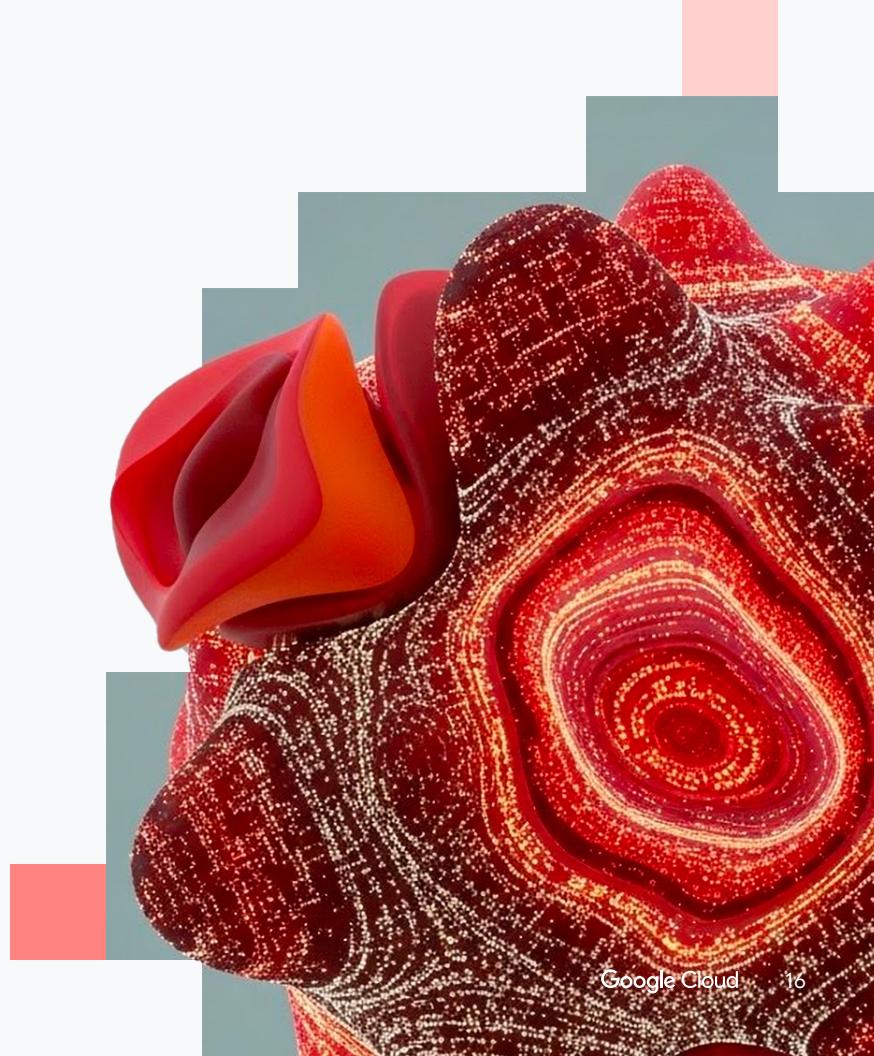
Build a chat app with Conversational Analytics API and Looker

[Explore codelab](#)

[Agent Development Kit \(ADK\)](#): Serves as the foundational framework for building agents, providing the core code structure, lifecycle management, and evaluation tools needed to create reliable agents in Python, Java, or Go.

[Model Context Protocol \(MCP\)](#): Functions as the standard "connector," allowing the agent to securely plug into external data sources and tools – such as databases or enterprise APIs – without needing custom integrations for every new service.

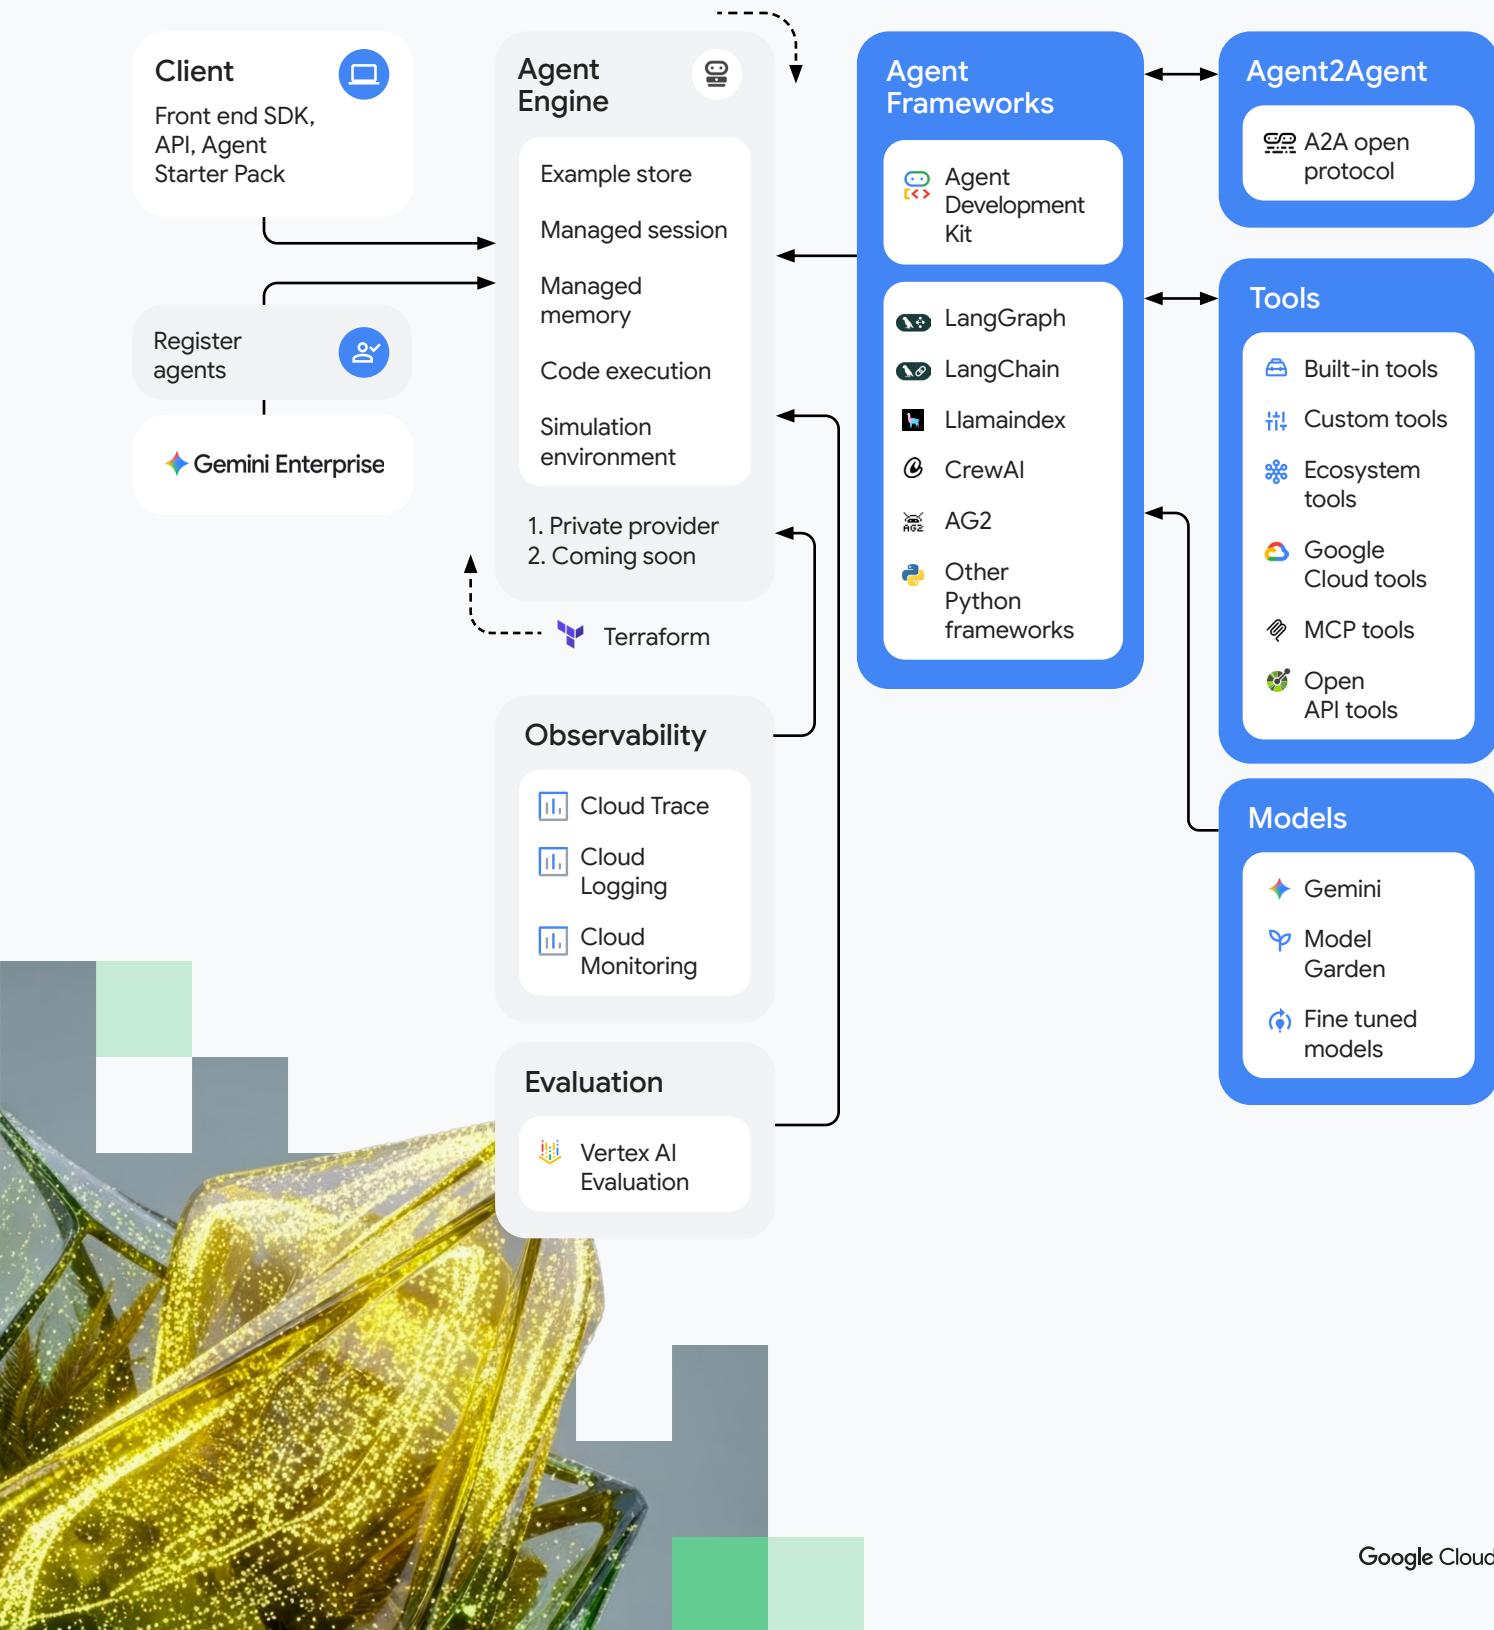
[Agent-to-Agent \(A2A\) Protocol](#): Acts as the networking layer, enabling these individual agents – regardless of how they were built – to discover one another, exchange information, and collaborate on complex, multi-step workflows as a unified team.





This full-stack framework simplifies agent development and is multi-agent by design, allowing you to choose the best Gemini models and tools for your organization's specific needs.

Frameworks, tools, and models to set up your agent





Google Cloud offers a suite of BigQuery tools that streamline how AI agents interact with enterprise data. These tools effectively remove the need for developers to build custom database connectors from scratch and work through ADK and MCP integration methods.

ADK integration

For developers that build agents using Google's ADK, [BigQuery](#) and [Spanner](#) integration tools are available as pre-built "skills" that can be instantly assigned to an agent. This built-in toolset includes ready-to-use functions that enable agents to autonomously: explore data, understand schemas and column definitions, and run queries.

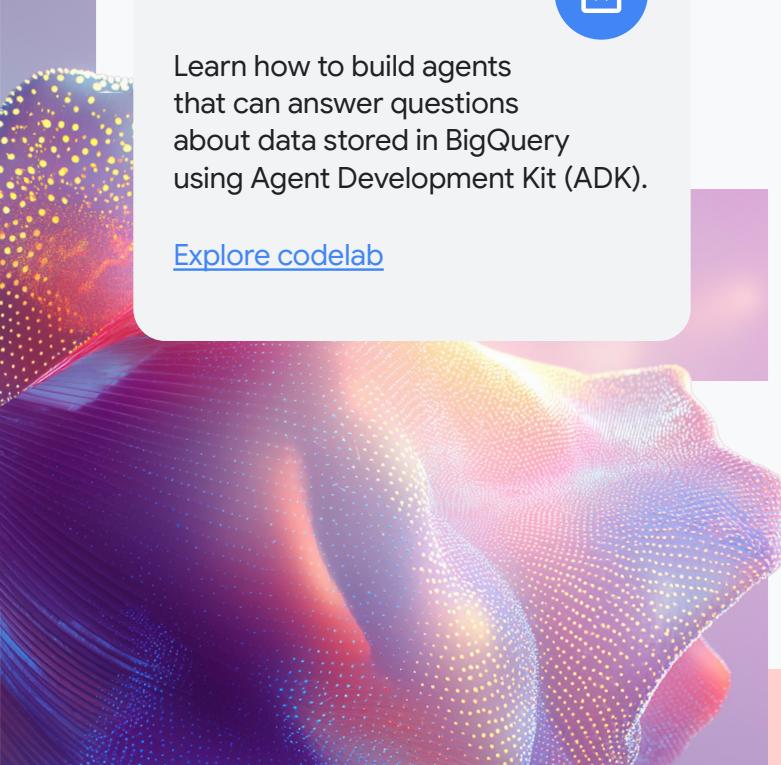
By importing this toolset, an agent gains the ability to "explore" a database and answer natural language questions – for example, "what were the total sales last quarter?" – without the developer writing a single line of SQL generation code.

Codelab



Learn how to build agents that can answer questions about data stored in BigQuery using Agent Development Kit (ADK).

[Explore codelab](#)



MCP integration

[MCP on Google Cloud](#) allows developers to connect an AI agent, or standard MCP client like Gemini CLI or Antigravity, to a globally consistent and enterprise-ready endpoint of Google Cloud services, such as BigQuery and AlloyDB. This enables agents to interpret schemas and execute queries against enterprise data without the security risks or latency associated with moving data into context windows. Thereby, MCP on Google Cloud provides direct access to BigQuery features like forecasting, while ensuring that data remains securely in-place and appropriately governed.

For more flexibility and control, use the [MCP Toolbox](#) – an open-source server that centralizes the hosting and management of toolsets, decoupling agentic applications from direct database interaction. Instead of managing tool logic and authentication themselves, agents act as MCP clients, requesting tools from MCP Toolbox for Databases. It's also available with a variety of IDEs and developer tools including [Gemini CLI](#) and [Antigravity](#), allowing you to securely connect your AI agents to services like [AlloyDB](#), [BigQuery](#), [Spanner](#), [Looker](#) and more.

Codelab



Learn how to make BigQuery datasets available using MCP Toolbox for Databases.

[Explore codelab](#)

Deploy and run data agents at scale

Deploying an agent requires more than just a server – it demands a robust, production-ready environment. [Vertex AI Agent Engine](#) is Google Cloud's preferred runtime for agents, as it is designed specifically for this purpose. As a fully managed service, Vertex AI Agent Engine removes the operational burden by handling security, authentication, and auto-scaling to allow your agent to perform reliably under real-world traffic.

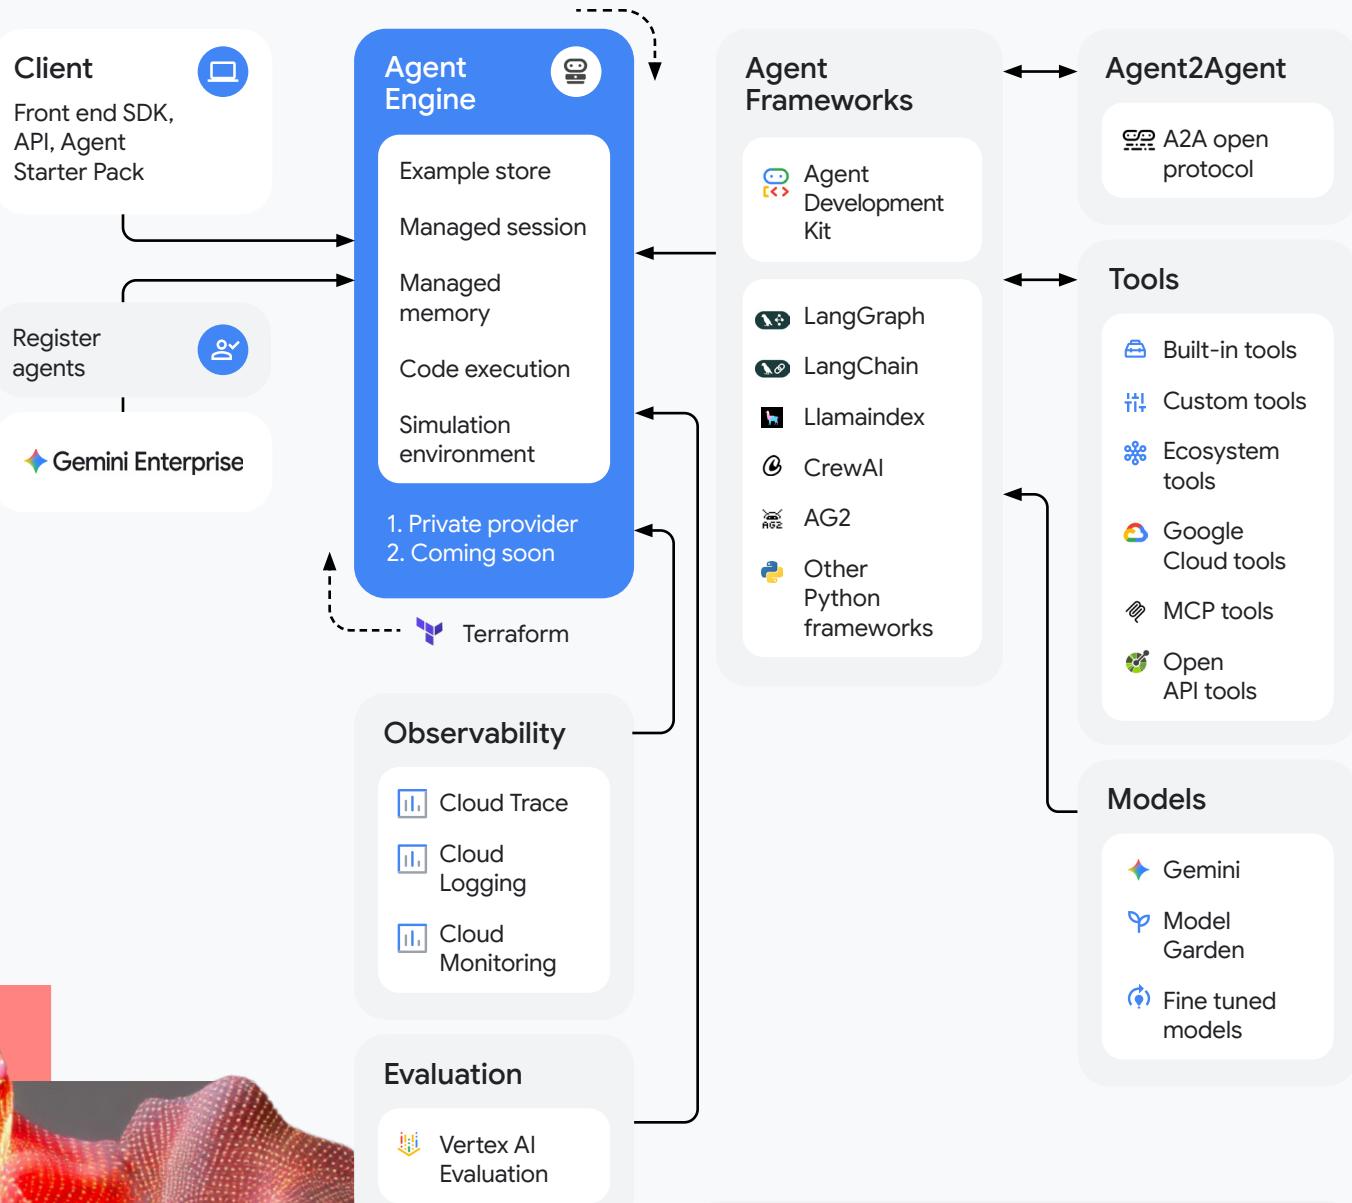
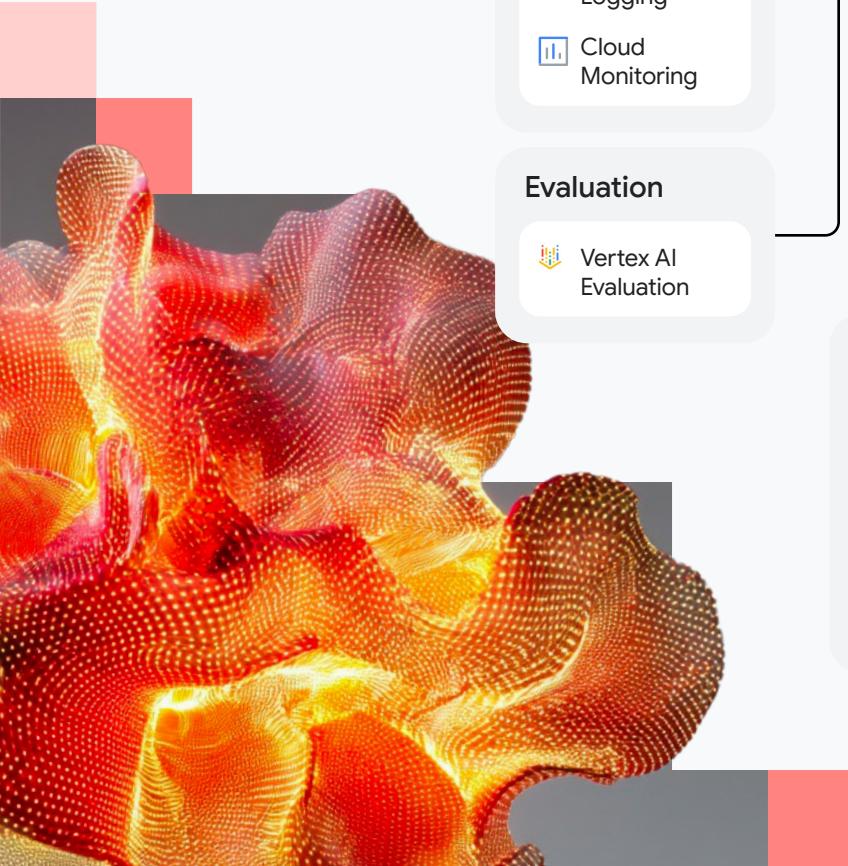
Vertex AI Agent Engine's framework-agnostic design also lets you deploy agents built with ADK, LangChain, or custom frameworks, providing maximum flexibility without the infrastructure overhead.

Although the purpose-built Vertex AI Agent Engine is recommended for most use cases, [Cloud Run](#) offers ultimate flexibility for those requiring precise container control or maximum portability. The decision is clear: if your goal is to reduce time-to-market and leverage a managed, enterprise ecosystem, choose Vertex AI Agent Engine. If strict security policies or custom binaries are your primary concern, choose Cloud Run to gain the control that you need.





Tools to deploy, run, and evaluate your agent



Codelab



Follow our codelab to learn how to build a multi-agent system with ADK, Agent Engine, and AlloyDB.

[Explore codelab](#)

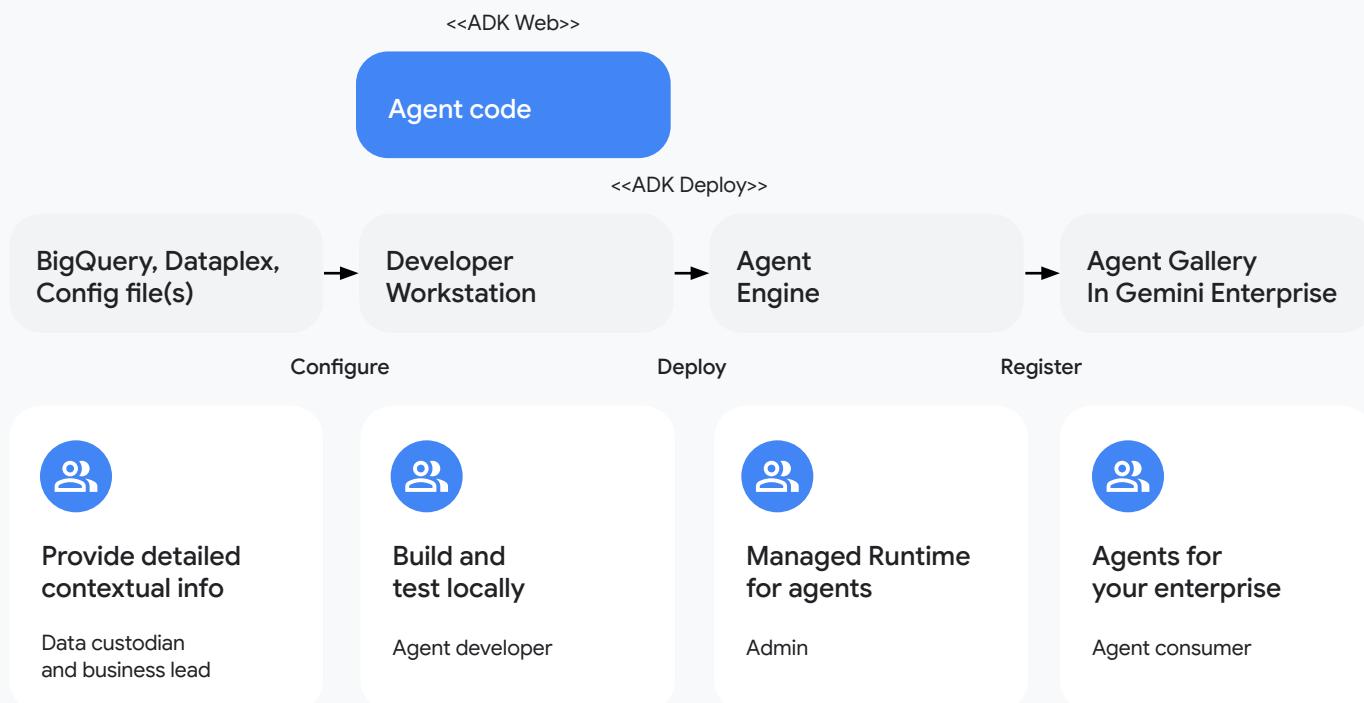
Publish agents to make them accessible

An agent's value is only realized when it's easily accessible. Google Cloud transforms custom agents from isolated projects into discoverable enterprise assets through either [Gemini Enterprise](#) or your own distribution platforms, which act as your internal "app store for agents."

The process is relatively straightforward. After deploying your agent to a managed runtime, like the Vertex AI Agent Engine, use a registration tool to formally publish the agent.

This straightforward step makes your agent's capabilities and endpoint available to the wider organization, turning a custom tool into a reusable service.

Once your agent is published in Gemini Enterprise, it appears in the [Agent Gallery](#). This centralized hub allows business users to browse, search for, and discover available agents, complete with detailed descriptions of their functions, owners, and usage instructions. The gallery provides more than just discovery – it's a governed environment where administrators can manage permissions to ensure that the right users have access to appropriate agents and their underlying data sources. This structured approach fosters a secure, collaborative ecosystem that prevents redundant development efforts and maximizes the impact of every custom agent that you build.



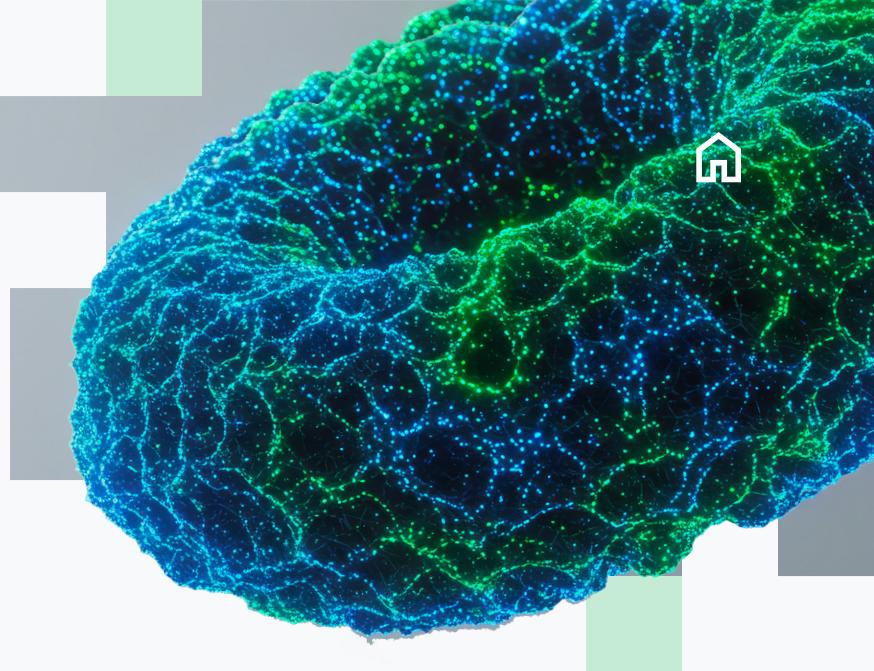
Publish agents in Gemini Enterprise



AgentOps: Analyze and optimize usage, cost, and performance

With the baseline agent now configured, the critical question is “can this agent truly perform in a real-world production environment, and can it successfully navigate the nuanced, complex queries posed by actual users?” Specifically, does it possess the core capability to identify the correct tables from a schema, generate accurate SQL queries and, most importantly, can it resist the critical tendency to hallucinate? A comprehensive evaluation is necessary to validate an agents’ readiness.

Although the Agent Development Kit (ADK) offers an evaluation service, this is currently limited to a few metrics solely for assessing the agent's final response. [Vertex AI GenAI evaluation service](#) empowers you to rigorously assess and understand your AI agents. It includes a powerful set of evaluation metrics that are specifically designed for agents built with different frameworks, and provides in-built agent inference capabilities to streamline the evaluation process.



Building an agent is an iterative process; to improve an agent, you must understand how it's being used. Historically, gathering data on agent usage, cost, and performance required a massive engineering effort to build custom logging pipelines and dashboards.

Google Cloud eliminates this challenge with the [BigQuery agent analytics plugin](#), a powerful, out-of-the-box solution. This single plugin lets you instantly activate a scalable data pipeline to capture, analyze, and visualize agent usage, cost, and performance.

This tool allows you to go from deployment to analysis in under five minutes. You can track costs, understand user behavior, and identify which tools are being used most often. Crucially, gaining access to these insights creates a powerful feedback loop that lets you use real-world data to continuously improve your agent's quality and demonstrate its value to the business.

Codelab



Learn how to evaluate agents using Vertex AI's GenAI Evaluation service.

[Explore codelab](#)

Codelab

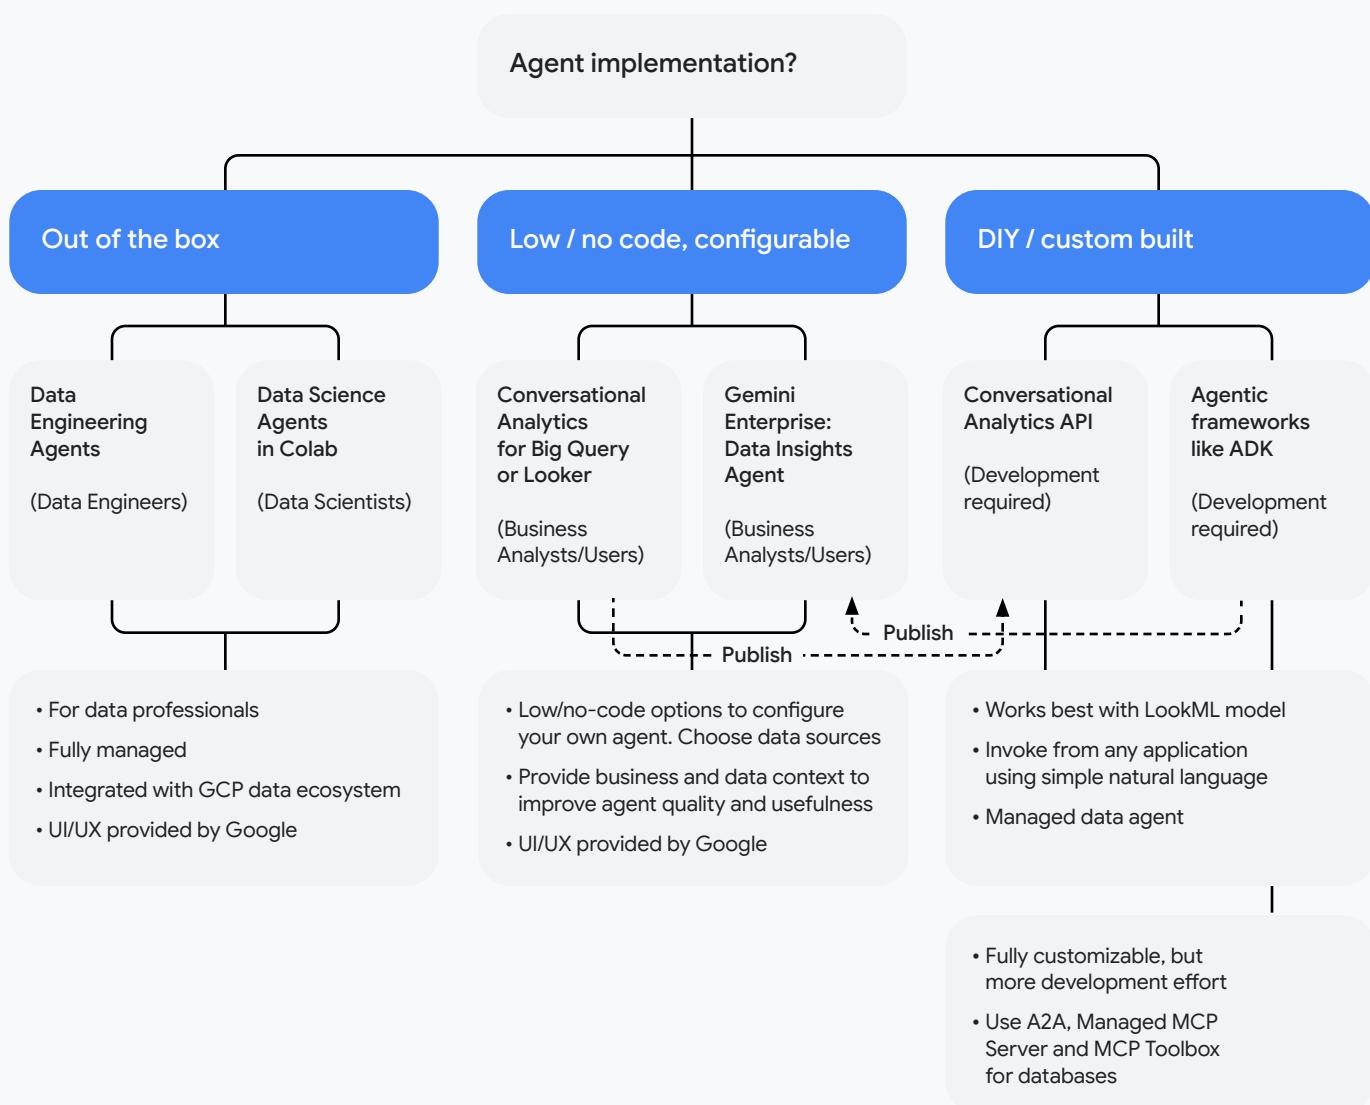


Learn how to enable agent observability to analyze conversation traces and agent tool usage.

[Explore codelab](#)

A summary of Google Cloud's approach to data agents

Google Cloud's approach to data agents



Google Cloud's approach to data agents is built on a philosophy of choice, offering a spectrum of solutions designed to empower every user – from business analysts to expert developers. This ensures that teams can adopt agentic AI in a way that best fits their use case and technical capabilities.



The approach is divided into two main paths:

01

Out-of-the-box, fully managed prebuilt agents: For data professionals who need to accelerate their workflows, Google Cloud offers a suite of pre-built, persona-specific agents. Today, this includes the [Data Engineering Agent](#), and the [Data Science Agent](#).

These are fully managed solutions that are deeply integrated within the Google Cloud ecosystem and designed for rapid time-to-value with minimal development effort. For business users, [Conversational Analytics Agent](#) provides an immediate, intuitive way to interact with data.

02

Custom-built agents: For developers who require maximum flexibility and control, Google Cloud provides foundational building blocks for creating bespoke agentic experiences. Using the [ADK](#), [MCP](#), and the [Conversational Analytics API](#), teams can design fully customizable agents with their choice of models, frameworks, and logic, and embed them directly into their own applications.

This dual approach means that you are never forced into a one-size-fits-all solution. You can start with a managed agent to solve an immediate need and build a custom agent for a unique, strategic opportunity – all on a single, unified platform.





Chapter 05:

Build on a foundation of AI leadership

Google Cloud provides the industry's only fully integrated, AI-optimized stack designed specifically for the era of agentic AI. Unlike fragmented solutions that require stitching together vector databases, inference engines, and foundational models from disparate vendors, Google Cloud offers first-party technology across every layer of the stack. This spans from a secure, purpose-built AI infrastructure up to the data analytics, semantic layers, and ML platforms that power the agents themselves. This end-to-end approach eliminates technical debt, reduces latency, and ensures peak performance, creating an environment for developing and deploying agents that can reason, plan, and act at speed.

At the base, our AI Hypercomputer architecture integrates performance-optimized hardware – including industry-leading TPUs and GPUs – with open software and flexible consumption models. This is the same infrastructure that is used to train and serve Google's most capable models, ensuring that your agents operate on a foundation designed for massive throughput and low-latency reasoning.



Agents are only as good as the data they can access. Google Cloud creates a “data-to-AI” lifecycle that is unequaled in the current market. By unifying BigQuery’s serverless data warehousing and AlloyDB’s high performance database with the semantic modeling of Looker, we provide agents with trusted business definitions. This allows agents to ground their responses in your operational reality – understanding metrics, governance, and real-time facts – rather than relying solely on pre-trained knowledge. This deep data integration reduces hallucinations and ensures that your data agents are enterprise-ready.

These agents are powered by the Gemini family of models, which offer industry-leading multimodal capabilities and massive context windows. This allows your agents to process vast amounts of unstructured data – such as documents, code, video, or audio – alongside structured business data for maximum impact.

You can also easily customize these models using Vertex AI, utilizing advanced reasoning engines that allow agents to break down complex user goals into multi-step workflows.

Google Cloud is backed by over 20 years of leadership in data and AI. This is the same technology that underlies Google Search and YouTube, providing customers with unequaled economies of scale and reliability every day. Google pioneered the Transformer architecture that sparked the generative AI revolution, and continues to lead in this space through DeepMind’s breakthroughs. With Google Cloud, you don’t just gain individual tools; you gain a complete, flexible ecosystem with practical pathways to get started quickly. This combination of a superior, unified stack and decades of proven data and AI expertise makes Google Cloud the ideal place to build, manage, and scale your intelligent agents, accelerating your time to real business value.



Build your AI intelligent apps in the right place

With a comprehensive, flexible, integrated AI platform, from infrastructure to purpose-built solutions, Google Cloud offers a complete ecosystem, not just tools.



Benefit from the Deepest AI expertise

Exemplified by the world's most successful AI product – Google Search – Google provides customers with unmatched scale, research depth, and economies of scale.



Get started, now

Google Cloud provides practical pathways (plus tools, resources, and support) to both get started and accelerate time to value with AI in our 30 day, ramp plan.



Step into the agentic era

In this guide, we've journeyed from defining what a data agent is to exploring the powerful, persona-driven agents available on Google Cloud, even mapping out the steps to building your own agent along the way. Make no mistake, the agentic era is not a distant future – it is here and offering tangible solutions to the immense pressures facing data teams today.

The next step is to move from learning to implementing.

Your agentic transformation starts today.

Get hands-on experience

The best way to gain an understanding of the power of this new paradigm is to experience it firsthand. Google Cloud's pre-built agents for engineering, data science, and analytics are integrated directly within BigQuery Studio, ready to use today. You can start automating pipelines, generating insights, and asking questions of your data in natural language right now to see how intelligent assistance can transform your workflows.

Try BigQuery for free



Try AlloyDB for free



Try Looker for free



Plan your strategy with our experts

Every organization's path to adopting agentic AI is unique. If you're ready to discuss how data agents can address your specific challenges and help you achieve your business goals, our team is here to help. Let's design your roadmap together.

Contact us to schedule a personalized consultation



Contributors

Manoj Gunti,
Sr Product Marketing Manager,
Google Cloud

Geeta Banda,
Head of Outbound Product Management,
Google Cloud

Sean Zinsmeister,
Director of Outbound Product Management,
Google Cloud

Sean Rhee,
Senior Product Manager,
Google Cloud

Ani Jain,
Senior Outbound Product Manager,
Google Cloud