

Unlocking the real power of Multi Agent System (MAS) to improve efficiency and effectiveness of your employees

The New Frontier: Why AI Agents Are the Future of Enterprise Work

The conversation around AI is rapidly shifting. We're moving beyond large language models (LLMs) that simply generate text and images. Today, the focus is on building **AI Agents**—designed to reason, plan, use tools, and complete complex, multi-step business tasks with or without constant human intervention. Every C level in almost every organization is focused on improving efficiency and effectiveness of their employees by introducing AI Agents.

Organizations today collect a lot of data, but access to that data remains a key challenge. Customers need insights for critical business decisions instantly, but they are often forced to navigate siloed systems or wait days for an overloaded BI team to build a new report. How can we go from data gathering to on demand data insights?

The Foundation — The "Conversational Analytics" aka Data agent

Thankfully, we now have tools like **Looker Conversational Analytics**.

This technology is a powerful first step. At its core, Looker [Conversational Analytics](#) is an AI-powered feature that integrates with the Looker platform. It allows any user—regardless of technical skill—to ask questions in natural language and get answers directly from the organization's governed, trusted data models (their "single source of truth").

It enables users to have a natural language conversation with their data (whether it's in BigQuery or another source), providing quick, easy, and instant insights, significantly reducing the time spent waiting for new reports.

Every AI Agent out there can actually “answer” this question. However, the key point here is, is the response accurate? Is the data coming from the agent correct? Can we trust AI?

Ensuring AI responses are trustworthy, secure, and free from hallucinations is a critical challenge for organizational leaders today. Without this, building AI agents is merely a science

project at best. This is why it is absolutely essential to make sure AI agents are grounded in your company data that has the business context in addition to the technical context.

Let us assume you are someone from the Marketing team (Sara) in your organization and would like to have data served to you at all times so that you can take data driven decisions quickly.

You are interested in understanding the top 3 products that are driving revenue in each of your regions this month. Conversational Analytics (Data agent) converts the natural language question into SQL query and runs against the data source and provides a response back to the user with pinpoint accuracy.

So how did this happen? The beauty of Conversational Analytics is, when a question is asked, the data agents use the Looker Semantic layer under the hood to convert the natural language questions into SQL queries. In simple terms, the Semantic layer is a fancy SQL generator which has the business context and the technical context to handle user queries. The semantic layer then runs the SQL queries on top of the Big Query (in this example) and retrieves the data for the user.

What are the top 3 products driving revenue in each region this month?

Debug info
Yesterday, 6:55 PM

I'll write a query for the question: "What are the top 3 products driving revenue in each region this month?"

Here's the query result for Top Products By Revenue Per Region This Month.

	Dim Stores Region	Dim Products Product Name	Fact Transactions Total Revenue
1	Gulf Coast	Magellan Outdoors Men's Fishing Shirt	\$5,672.73
2	Gulf Coast	Stanley Quencher H2.0 FlowState	\$990.00
3	Gulf Coast	Yeti Tundra 45 Hard Cooler	\$650.00
4	Midwest	ProForm Carbon TL Treadmill	\$600.00
5	Midwest	Stanley Quencher H2.0 FlowState	\$450.00

The top 3 products driving revenue in each region this month are: **Gulf Coast:**

- Magellan Outdoors Men's Fishing Shirt: \$5,672.73
- Stanley Quencher H2.0 FlowState: \$990.00
- Yeti Tundra 45 Hard Cooler: \$650.00

Midwest:

- ProForm Carbon TL Treadmill: \$600.00
- Stanley Quencher H2.0 FlowState: \$450.00
- Rawlings 11.5 Baseball Glove: \$419.93

Ask a question

The "Aha!" Moment — The Wall of Silos

Conversational Analytics is a powerful tool by itself. It's great for querying a single, well-defined data source.

But it got me thinking: how can we take this to another level?

In the real world, a marketing analyst's job is rarely done with just sales data. To understand *why* a campaign failed or succeeded, they need to combine data from multiple systems.

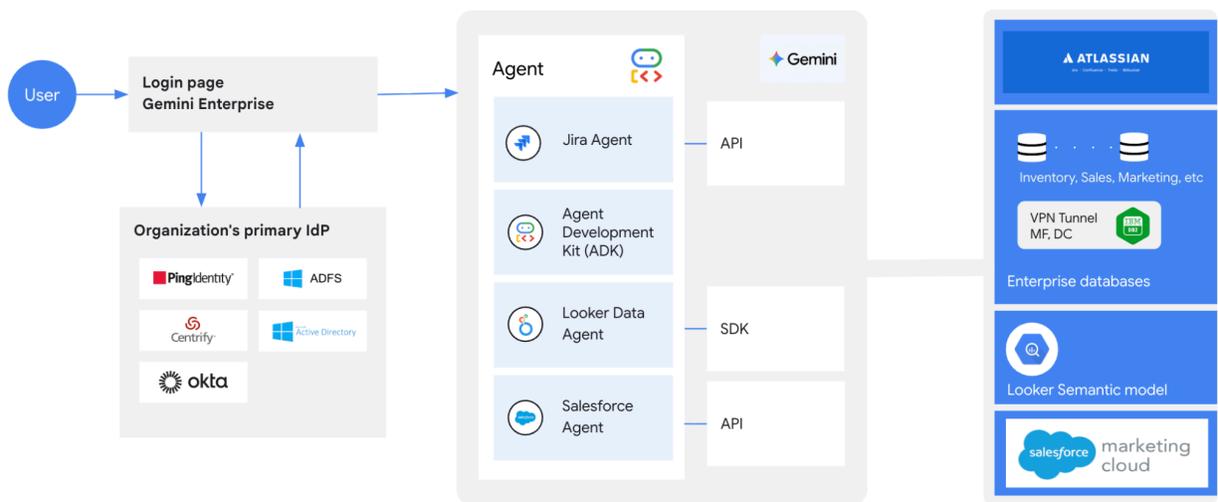
- **Sales & Revenue Data** (from BigQuery, queried via Looker)
- **Campaign Performance Data** (from Salesforce Marketing Cloud)
- **Inventory & Cost Data** (from an internal ERP or database)

This means logging into three different systems, manually merging spreadsheets, and crushing data in Excel just to get a single answer. I felt we could use the power of the **Gemini Enterprise** platform to make the life of the marketing team dramatically more efficient.

The Solution — A Multi-Agent System

A **Multi-Agent System**—a single, intelligent interface built using the Agent Development Kit (ADK) that connects to all these different data sources (Databases, EDW, APIs, Other Cloud platforms) to bring the data closer to the user as opposed to the users going to the data.

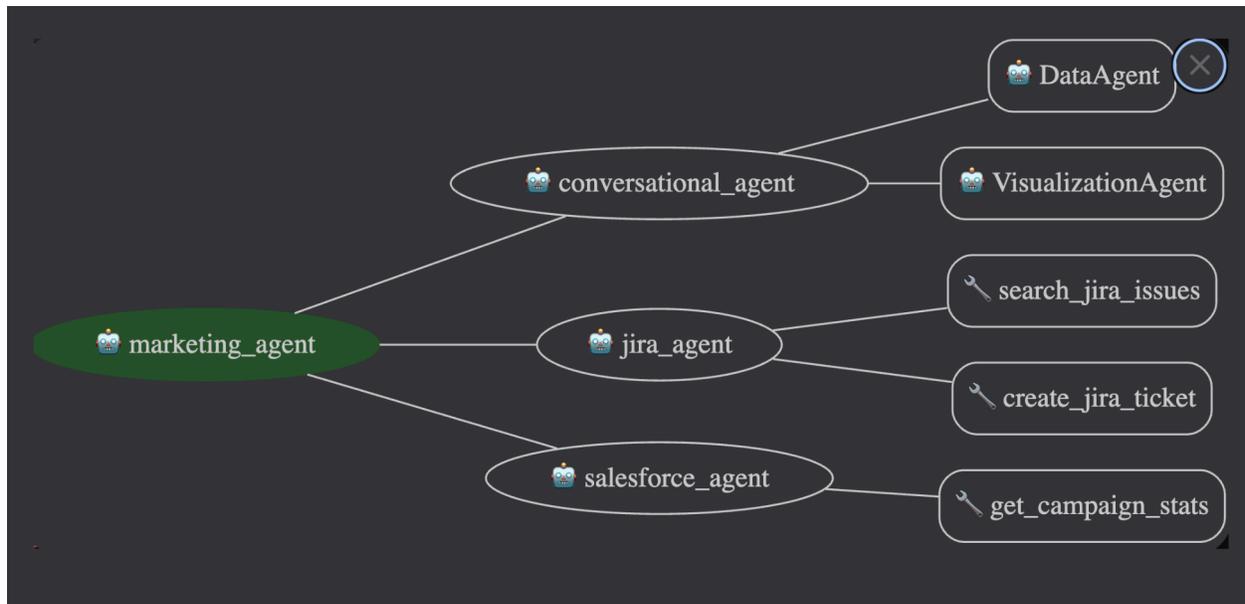
Here's a high level architecture of what we are trying to accomplish for the Marketing person.



Using Agent Development Kit, we are building a marketing agent (root agent) that has three sub agents -

- I. Jira agent
- II. Conversational Analytics agent
- III. Salesforce agent.

Each of these agents has corresponding tools to accomplish their tasks. Let us now see how these agents work together to answer questions for Sara. Let's look at how agents and tools work together for this architecture.



Architecture Deep Dive: The Multi-Agent System in Action

The **Marketing Agent (Root Agent)**, powered by [Gemini Enterprise](#), acts as the orchestrator and the single point of interaction for the user (Sara). When Sara submits a complex, multi-part query—such as "What were the top 3 revenue-driving products this month, cross-reference that with any related open Jira tickets, and summarize the relevant marketing campaign performance from Salesforce Marketing Cloud?"—the Marketing Agent springs into action.

The Marketing Agent's primary job is to parse the user's intent and decompose the complex query into smaller, manageable sub-tasks. It then delegates these tasks to specialized **Sub-Agents**:

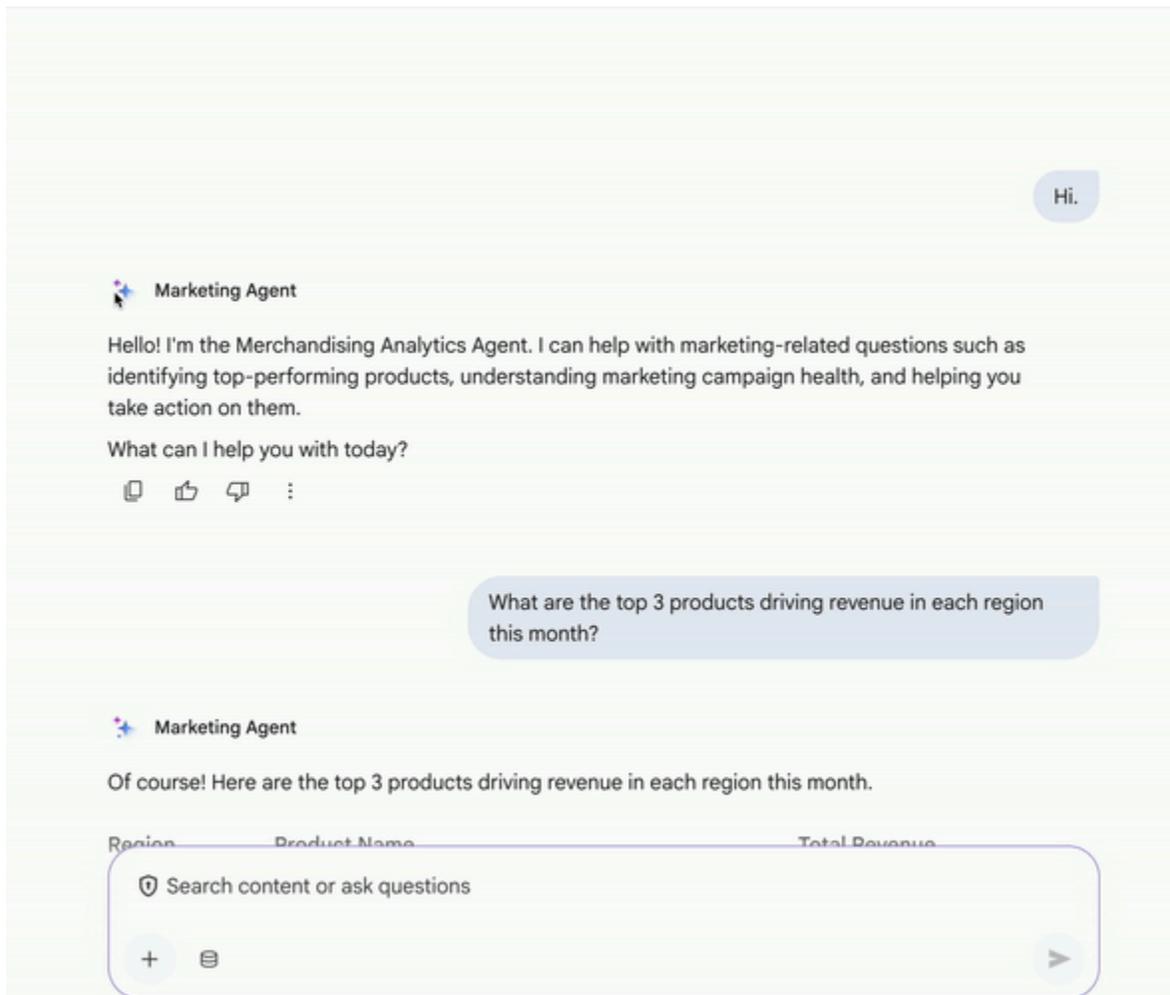
- 1. Conversational Analytics Agent (Data Agent):**
 - **Tool:** Looker Conversational Analytics / Looker Semantic Layer.
 - **Task:** Handles the data retrieval question: "Identify the top 3 revenue-driving products this month."
 - **Mechanism:** Converts the natural language query into optimized SQL using the Looker Semantic Layer, executes it against the governed data source (e.g., BigQuery), and returns the validated, accurate sales and revenue data.
- 2. Jira Agent:**
 - **Tool:** Jira API Connector.
 - **Task:** Retrieves operational context: "Find related open Jira tickets for those top 3 products."

- **Mechanism:** Uses the product names provided by the Data Agent to query the Jira instance via its API, filtering for open issues, bugs, or feature requests associated with those products.

3. **Salesforce Agent:**

- **Tool:** Salesforce Marketing Cloud API Connector.
- **Task:** Gather external performance data: "Summarize the performance of the latest marketing campaigns targeting those products."
- **Mechanism:** Connects to the Salesforce Marketing Cloud API to pull metrics (e.g., click-through rates, conversion rates, spend) for campaigns linked to the top products.

Once all Sub-Agents execute their tasks and return their respective data segments, the Marketing Agent performs the final critical step: **Synthesis and Reasoning**. It aggregates the data from the three siloed sources (revenue data, operational issues, and campaign performance), uses Gemini's reasoning capabilities to connect the dots, and generates a single, coherent, data-driven insight presented back to Sara in natural language. This architecture fundamentally shifts the effort from the user (manual data crushing) to the system (intelligent orchestration).



In Summary: From Silos to Synthesis

What I've demonstrated here are the endless possibilities of what you can achieve with Gemini Enterprise and a multi-agent approach.

We saw the journey from a simple, siloed agent (Looker Conversational Analytics) to an advanced, multi-agent system. This agentic architecture improves the efficiency of a marketing person. Instead of Sara going from one system to another to manually gather data, Gemini Enterprise provides a **single, intelligent platform** for them to be not just efficient, but truly effective. Hope it gave you ideas on how to think about unlocking the true potential of a multi agent system with Gemini Enterprise.

Thanks for reading.