Conversational Analytics in Looker

Google Cloud

Today's Agenda

01	What is Looker?
02	What is Conversational Analytics?
03	LookML: The Foundation of Accuracy
04	Agents and Instructions: The "Brain" of the Chat
05	How It All Works Together
06	The "Ask "Lifecycle: What Happens Behind the Scenes?
07	Feedback & Iteration: Improving the "Data Brain"
08	Quantifiable Improvement

Bytecode IO

What is Looker?

Serve up real-time reports and dashboards that inspire more in-depth analysis.

Modern BI & analytics

Infuse relevant data with your existing tools for an enhanced experience and more effective results.

Integrated insights

Super-charge operational workflows with complete, near-real time data.

Data-driven workflows

Infuse trusted metrics directly into Gemini and LLM workflows to ground AI reasoning in a single source of truth.

Integrated AI Insights

Looker

API-first and cloud native for integrating into existing workflows

Semantic modeling layer for enterprise-wide governance

In-database architecture for access to real-time data

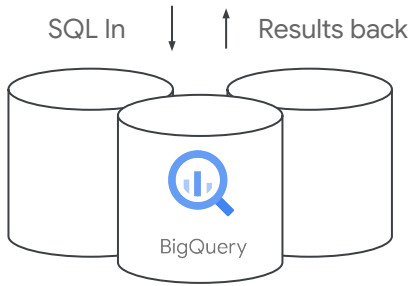

What is Conversational Analytics?

Conversational Analytics is a **Gemini-powered** interface that allows users to “chat with their data.”

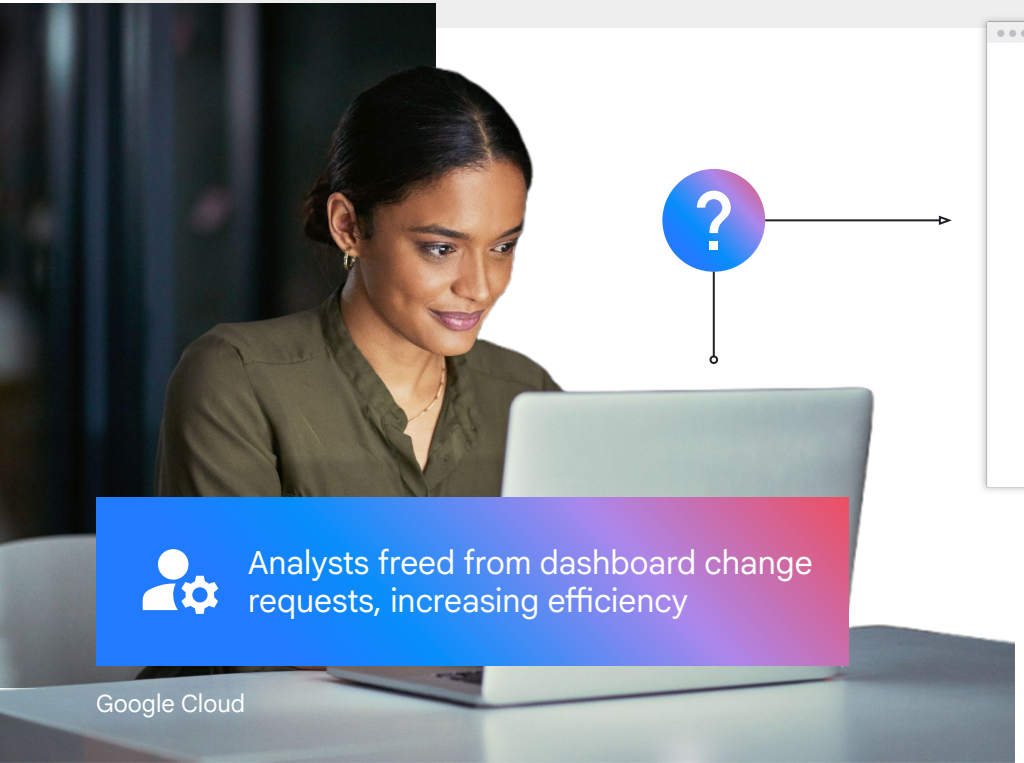

Analysts freed from dashboard change requests, increasing efficiency

From “What?” to “What next?”

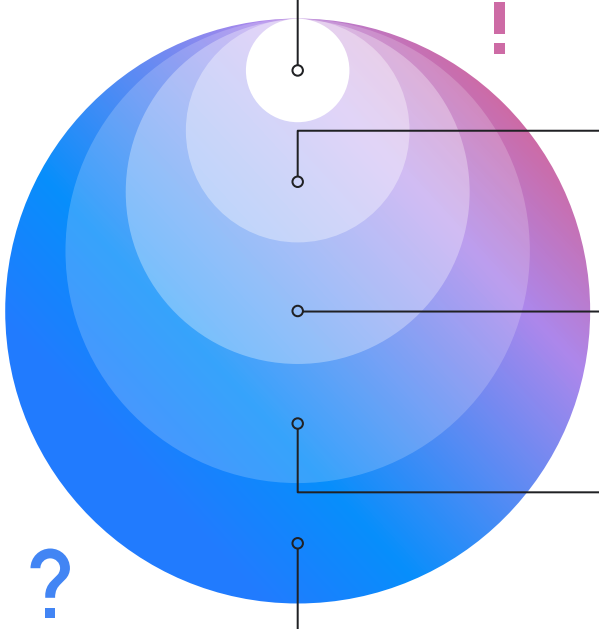

“What?”

Sales figures, inventory levels, metrics

“How does this compare?”

Benchmarking, vs time, vs groups

“Why?”

Root cause analysis, correlations

“What will happen?”

Forecasting, trend prediction

“What should we do?”

AI-powered recommendations

Bytecode IO

Conversational Analytics on Looker Dashboards

Chat with your Dashboards

Empower dashboard viewers to go beyond static reports and engage in dynamic conversations with their data, instantly getting answers to their questions directly within the dashboard.

Unlock Deeper Understanding

Enable viewers to easily ask "why" behind the data displayed in dashboard tiles, generate summaries, and uncover hidden trends, fostering a more intuitive understanding.

Drive Data-Driven Decisions

Equip viewers with the ability to explore beyond pre-set filters, integrate more context, and ultimately make better, faster decisions based on a better understanding of their dashboard data.

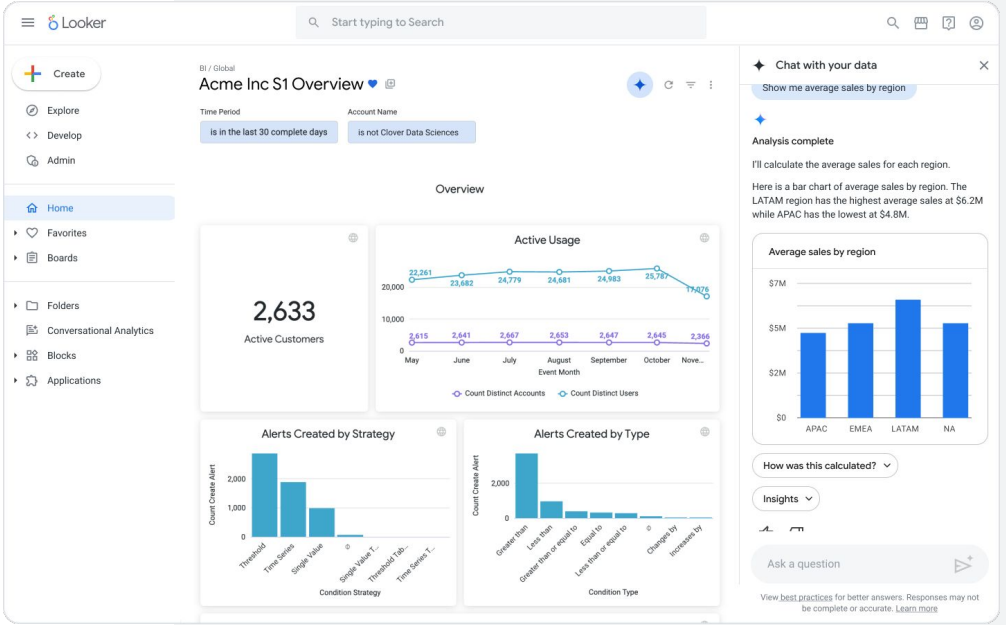

Demo

Semantic Layer powers conversational analytics

Looker significantly reduces the load on a Large Language Model

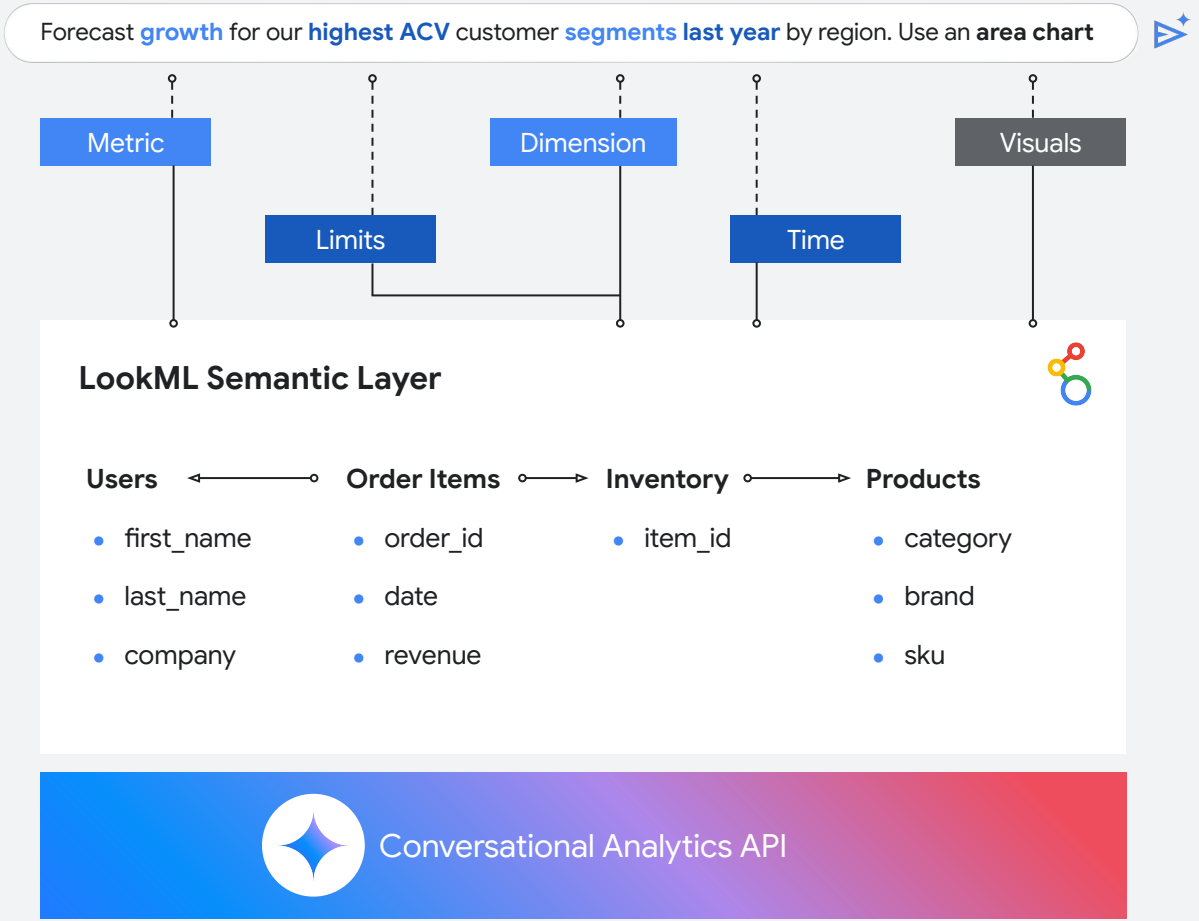

LookML: The Foundation of Accuracy

Explore

Define the scope: Use description, label, group_label, and QuickStarts to frame the starting point.

Join

Focus on relationships: Use view_label to ensure joined tables appear intuitive in the conversational flow.

Field

Precision mapping: Utilize labels, group_labels, group_item_labels, descriptions, synonyms, and suggestions for high-fidelity intent mapping.

Readme

Strategic intent: Describe the primary use case for each Explore in the Readme to help the Agent understand the data's core purpose.

Future-Proofing Tip While Conversational Analytics does not leverage every parameter today, robust LookML metadata is essential for human usability and ensuring your model is "AI-Ready" as capabilities evolve.

Bytecode IO

LookML: The Foundation of Accuracy

Curation Checklist

If your Explore is curated for users, it's ready for the AI Agent.

- **Limit & Hide:** Suppress technical "plumbing" like system IDs and raw foreign keys.
- **No Duplicates:** Ensure every logical concept has one (and only one) clear field.
- **Zero Ambiguity:** Avoid generic names like "Date"; use "Order Date" or "Return Date."
- **Define Units:** Add units (e.g., USD, kg) to descriptions for measurement clarity.
- **Discrete Values:** List exact possible values in descriptions for categorical fields.
- **Example Values:** Provide real data samples in descriptions to "teach" the Agent.
- **User Semantics:** Labels must reflect business meetings, not SQL columns.

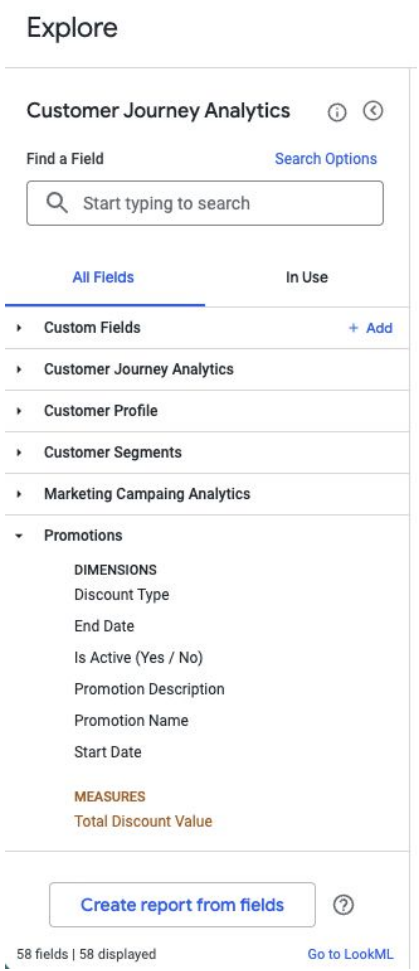

Bytecode IO

Labels & Synonyms

Labels should be clear, concise, and reflect business terms.

Synonyms don't need to be exhaustive, but they should catch the common patterns business users may use to refer to the field.

```
dimension: specialty {
  group_label: "Prescriber Details"
  label: "Medical Specialty"
  description: "The primary clinical area of expertise (e.g., Cardiology, Endocrinology)."
  synonyms: ["Area of Medicine", "Doctor Type", "HCP Specialty"]
  type: string
  sql: ${TABLE}.SPECIALTY ;;
}
```

Bytecode IO

Descriptions

Effective descriptions act as the Agent's primary "Cheat Sheet" for intent mapping.

Minimalist Formula

[What it is]. [Logic/Constraint]. [Unit/Format].

- **Revenue:** "Total revenue from completed transactions, excluding returns and shipping fees, in USD."
- **Status:** "The current state of a membership based on the most recent successful billing event."

Maximalist Formula

[Flags]. [What it is]. [Purpose]. [Guidance].

- **Audited Financials:** "[AUDITED]. Total revenue after returns. Use for board-level reporting. Do not confuse with 'Gross Revenue'."
- **ML Churn Score:** "[BETA / ML-DERIVED]. Likelihood (0-1) of cancellation in 30 days. Updated every Monday."

No Technical Jargon: Never spell out SQL calculations or DB column names; the AI already understands the code.

Avoid the Noise: If descriptions are too wordy, the AI gets "lost." Keep language concise and high-signal.

Dual Benefit: Clean, business-semantic language is the "Golden Rule" for both human users and LLMs.

LookML: The Foundation of Accuracy

Optimizing Suggestion Performance

Fuzzy search and database sampling use Looker's suggestion API. Without optimization, retrieving data values can be slow and incur unnecessary load.

Limit Joins

Set full_suggestions: no to query only the base view containing the filter field. This skips unnecessary Explore joins, significantly reducing query complexity and execution time.

Static Values

Use the suggestions parameter to provide a hard-coded list of values. This eliminates database requests entirely for fixed categories or status fields, providing instant search feedback.

Redirect Queries

Point to efficient dimensions using suggest_explore and suggest_dimension. Instead of searching large fact tables, redirect the agent to a smaller, indexed summary table.

Key LookML Takeaways

Keep these takeaways in mind when defining LookML for Explores used as data sources for Conversational Analytics:

Clear Labels

Reflect how business users talk. Avoid `"amt_usd_curr"`; use `"Amount (USD)"`.

Seamless Mapping

Use `synonyms` and `descriptions` to help the agent map user questions to the correct fields.

Centralized Logic

Define calculations as LookML dimensions or measures for a single source of truth.

Streamline Context

Hide technical or internal fields like IDs to reduce noise and improve accuracy.

Search Optimization

Define hardcoded values in `suggestions` or use `suggest_explore`, or disable full suggestions using `full_suggestions`.

Query Efficiency

Use aggregate awareness, caching, and efficient join logic to optimize performance.

Demo

Agents and Instructions: The “Brain” of the Chat

An Agent isn’t just a search bar; it’s a **specialized assistant** designed for a specific data set.

It lives at the intersection of:

- **The Data Source:** A curated Looker Explore.
- **The Guardrails:** Logic and filters that keep the conversation on track.
- **The Persona:** A personality that speaks your department’s "dialect."

Agents and Instructions: The “Brain” of the Chat

The Power of Instructions: Instructions are the “system prompts” you write to guide Gemini.

Category	The Instruction (Abbreviated)	The Business Value
Semantic Routing	Differentiate "Net" vs "Gross" Revenue based on context.	Ensures "one version of the truth" across departments.
Logic Mapping	Define complex segments like "Leaky Bucket" or "At-Risk."	Translates high-level strategy into precise queries.
Proactive Insight	Force the agent to suggest root-cause follow-up queries.	Transitions the AI from a "Data Fetcher" to a "Strategic Partner."

Agents and Instructions: The “Brain” of the Chat

Anatomy of Agent Instructions

How we translate business intent into technical precision.

1. Persona & Tone
2. Business Context & Terminology
3. Default Query Behaviors
4. Field Mapping & LookML Guidance
5. Golden Queries (Examples)
6. Guardrails & Limitations

Agents and Instructions: The “Brain” of the Chat

Tone & Persona

Since the standard CA agent is already professional, we only customize the persona when we want to drive a specific behavior. While the data is the priority, the persona dictates how the agent handles failure or follow-up.

- **Define the Role:** "You are an Executive Chief of Staff. Your tone is incredibly brief and bottom-line-oriented. Avoid all technical jargon. If a user asks a question, provide the number first, followed by a single sentence on the 'so what' (the business impact)."
- **Handling "I Don't Know":** Instruct the agent on how to handle out-of-scope questions.
    - *Example: "If a user asks about HR data, politely inform them that this agent only has access to Sales and Inventory data."*

Agents and Instructions: The “Brain” of the Chat

Business Context & Terminology

The agent needs to speak your company’s dialect. This section is for bridging the "Semantic Gap"—the distance between how a human asks a question and how the data is structured.

A. Define Composite Concepts and “Standard View” Instincts

Map a single business term to a curated set of measures and dimensions that together answer a "big picture" question.

- **Term:** "Account Health"
- **Instruction:** "When a user asks about 'Account Health', provide the following metrics: orders.total_revenue, order_items.average_margin, users.days_since_last_order, and support_tickets.open_ticket_count. Present these as a table grouped by accounts.name."

B. Map Seasonal & Campaign Logic

The agent does not know your internal marketing calendar. Explicitly map campaign names to hardcoded date and product filters.

- **Term:** "'The Summer Sale'"
- **Instruction:** "'The Summer Sale' refers to the period between June 1, 2024, and August 31, 2024, filtered for the 'Seasonal' product category."

Agents and Instructions: The “Brain” of the Chat

Use LookML & Filter Syntax

The agent is most reliable when you speak its "native" language. Don't just describe a field; reference it by its LookML path.

- **Reference Fields Directly:** Use the view_name.field_name format to eliminate ambiguity.
  - *Example*: "For 'Sales', use the measure order_items.total_sale_price."
- **Leverage Standard Filter Expressions:** Use Looker’s built-in filter syntax (e.g., "last 90 days", "is not null") within your instructions so the agent can map them directly to the query.
  - *Example*: "Unless specified otherwise, always filter orders.status to 'complete'."

Agents and Instructions: The “Brain” of the Chat

Set Default Behaviors

Users are often vague. Instructions should act as a "safety net" for ambiguous questions.

- **Default Timeframes:** Decide what happens when a user asks "How many orders?" without a date.
  - *Best Practice:* "If no date range is provided, default to the last 30 days using orders.created_date."
- **Default Aggregations:** If a user asks for "performance," tell the agent which metric represents performance for that Explore (e.g., Conversion Rate vs. Gross Margin).

Agents and Instructions: The “Brain” of the Chat

Implement "Golden Queries"

Golden Queries are pairs of a Natural Language Question and a Verified LookML Query. They are the most powerful way to "train" the agent on complex patterns.

Question	LookML Component Goal
"Who are my top 5 customers this month?"	{ "model": "advanced_ecomm", "explore": "orders", "fields": ["users.name", "order_items.total_revenue"], "filters": {"users.state": "California", "orders.created_date": "this month" }, "sorts": ["order_items.total_revenue desc"], "limit": "5" }
"What is the attachment rate for shoes?"	{ "model": "advanced_ecomm", "explore": "sales", "fields": ["sales.attachment_rate"], "filters": {"product.category": "Shoes"}

Agents and Instructions: The “Brain” of the Chat

Optimize Instruction Efficiency

Avoid Redundancy: Keep the Prompt Lean

If LookML is the dictionary, the Agent has already read it.

Trust the Foundation:

The agent automatically ingests LookML labels, descriptions, and view logic. Do not repeat basic field definitions that already exist in your .lkml files.

Focus on Complexity:

Use your "instruction budget" for logic LookML cannot easily handle: multi-field composite metrics, nuanced persona goals, and departmental dialects.

Maintain Prompt Hygiene:

Handle simple aliases or synonyms at the field level in LookML using descriptions or tags. This keeps the instruction prompt focused and efficient.

Pro-Tip: High signal-to-noise ratio in your instructions maximizes the **reliability** of intent mapping and ensures the most **precise** execution of your trusted LookML logic.

Bytecode IO

Demo

Looker's Agentic Architecture

Powering Conversational Analytics

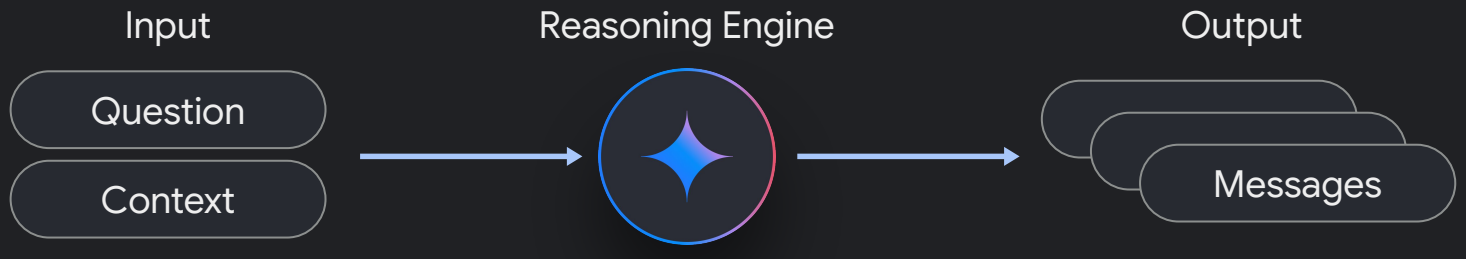

Capabilities

Data and business context

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

The “Ask” Lifecycle: What Happens Behind the Scenes?

What happens when a user asks: “What was my best selling product in Q3?”

1	2	3	4	5
Intent Recognition Users Question	**Semantic Mapping** Gemini + LookML+ Agent Instructions	**Query Generalization & Fuzzy Search** Generates Looker Query (JSON-based)	**Execution & Visualization** Looker Engine runs query: Gemini selects chart	**Narrative Summary** Natural Language Explanation

Bytecode IO

Demo

Over reliance on small teams of data experts has created an analytics backlog

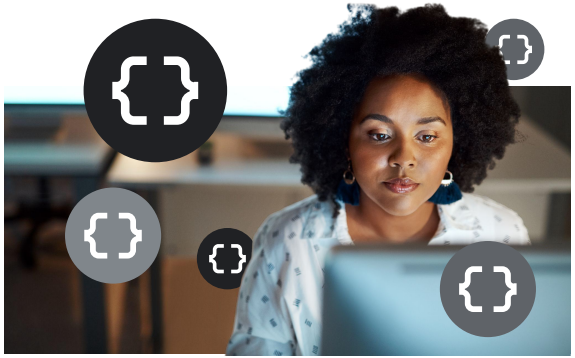

Growing number of data requests →

Limited number of data experts →

Insufficient resources and frequent delays

Quantifiable Improvement

The “Deflection” Metric: Track the reduction in “Ad-hoc” SQL requests to the BI team.

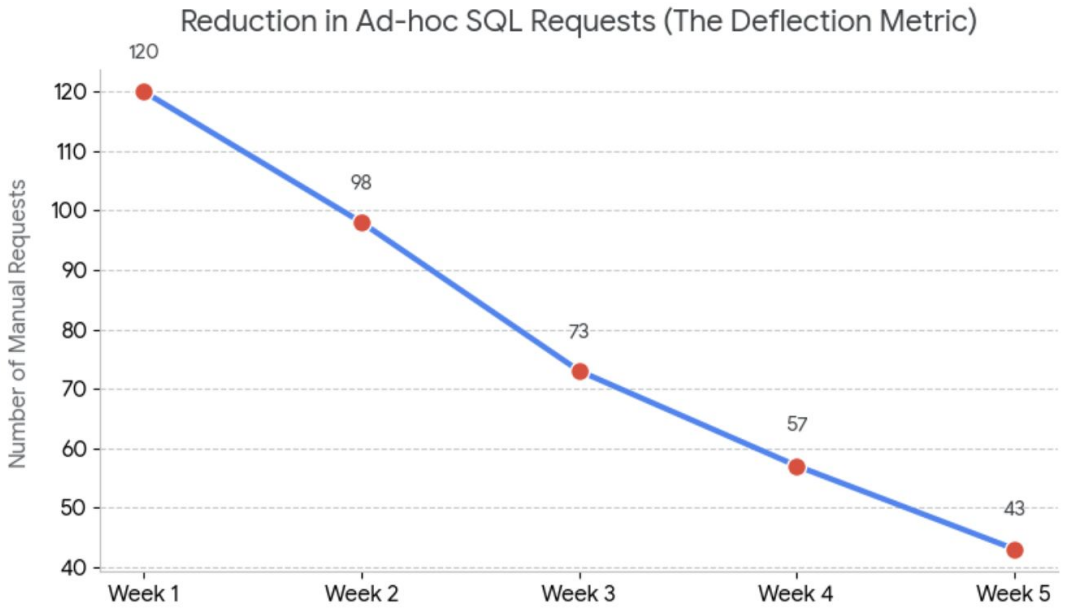

Bytecode IO

Quantifiable Improvement: Evaluation Framework

A structured scorecard to move from baseline hallucinations to 100% governed precision.

1. Test

Run the **"Golden Dataset"** of 50 real-world queries. Record the Returned Explore Link to see exactly how the AI mapped intent.

2. Diagnose

Categorize failures: **LookML Context** (Needs better labels/descriptions) vs. **Agent Instructions** (Needs specific logic/acronyms).

3. Refine

Apply fixes and conduct Round 2 testing. Quantify the delta in accuracy to ensure the "Data Brain" is learning correctly.

20-50	+35%	LookML Changes	Zero
GOLDEN QUESTIONS	ACCURACY LIFT	Trigger	HALLUCINATION GOAL

Bytecode IO

Context: LookML or Agent Instructions

The Developer Loop

As a developer, continue to test and update the Agent Instructions or the LookML Descriptions to close the gap. If the agent didn't understand "Q3", you update the instructions to define the fiscal calendar.

1. **LookML:**
    - Context intended for all users (human, AI)
    - Field-specific metadata, including: synonyms, descriptions, hidden fields
2. **Agent Instructions:**
    - Persona, terms, behavior, golden queries (until golden queries UI is GA)
    - Behavior may range from defaults under conditions to how the agent interacts with the user

Don't duplicate information the agent already knows: how to create a query, explaining Looker concepts, listing accessible Explores, or basic date filtering.

Context: LookML or Agent Instructions

Below is a tried and tested framework for Agent Evaluation with proven results.

1. **Build the test set.**
    a. Build a list of questions the agent will realistically encounter (make sure to include multi-step).
2. **Record baseline performance.**
    a. Test CA with each question (multiple times each question)
3. **Optimization.**
    a. Determine where gaps lie for wrong answers.
    b. If model context, modify LookML descriptions, labels synonyms, etc..
    c. If handling rules, modify agent instructions.
    d. It is possible both LookML and agent instruction changes are necessary.
4. **Retest and rescore.**
    a. Track changes and agent versions
5. **Repeat until acceptance threshold is reached.**

Demo

Conversational Analytics & AI Services

2026 Bytecode IO Program Offerings

It’s time we **unlock self-service** for everyone

Bytecode IO

Conversational Analytics Services

Unlock high quality, consistent conversational analytics for business users to self-serve

- Building trust and confidence in AI
- Repeatable development framework
- Combine LookML and Agent Instructions
- Internal launch templates & plan

Enable your technical teams to build trusted agent powered experiences and help set the right expectations with your users to maximize value.

"It's more than just turning on a feature, it's about building a **repeatable foundation** for AI powered workflow **that people can trust**"

Bytecode IO

Question	Baseline	Round 1	Round 2	Round 3
1	25.00%	100.00%	100.00%	100.00%
2	14.30%	85.70%	100.00%	100.00%
3	25.00%	87.50%	100.00%	100.00%
4	14.30%	85.70%	85.70%	85.70%
5	50.00%	100.00%	100.00%	100.00%
6	0.00%	40.00%	100.00%	100.00%
7	44.40%	100.00%	100.00%	100.00%
8	66.70%	100.00%	100.00%	100.00%
9	44.40%	88.90%	77.80%	77.80%
10	71.40%	100.00%	100.00%	100.00%
11	60.00%	80.00%	80.00%	100.00%
12	80.00%	80.00%	80.00%	100.00%
13	12.50%	75.00%	87.50%	100.00%
14	22.20%	77.80%	100.00%	100.00%
15	33.30%	88.90%	88.90%	100.00%
16	12.50%	62.50%	75.00%	100.00%
17	11.10%	55.60%	100.00%	100.00%
18	50.00%	62.50%	100.00%	100.00%
19	62.50%	62.50%	62.50%	100.00%
	0.00%	26.30%	57.90%	100.00%

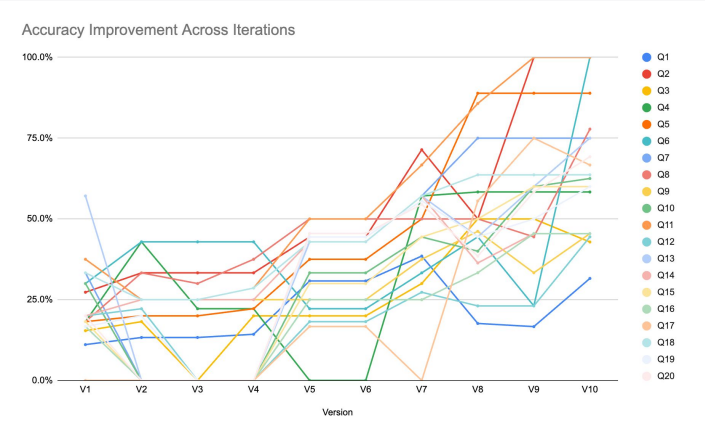

Proven Development Framework Delivers Results

The Bytecode IO development framework enables the organization to quantitatively improve response quality and consistency through iterations of LookML context maximization and agents tailored to business use cases.

The evaluation framework is the critical process required to delivering iterative improvement

Bytecode IO

Conversational Analytic Fast Starts

Blueprint

Assess AI readiness and develop single use case with a small team looking to test and validate Conversational Analytics for company use cases

2 weeks

$10,000

Fast Start

Multiple iterations of tuning to deliver 75%+ accuracy and consistency for a complex use case plus planning for launch and beta user onboarding

4 weeks

$20,000

Value Engine

Developer AI training program to maximize Looker for agentic workflow including delivery of a production use case to a broad audience. Build the repeatable engine for AI

10 weeks

$50,000

Bytecode IO

Conversational Analytics Blueprint

Assess readiness and co-develop an agent experience

Work with the Bytecode IO team to assess AI use cases, workshop how Conversational Analytics fits into analytics strategy and co-develop a use case in Looker for end user pilot.

Assess use cases and map into ROI

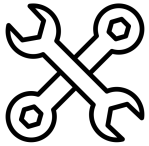

Optimise LookML and build data agent

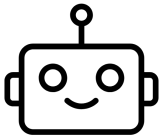

Deliver first CA agent experience

Develop roadmap of agentic value

Bytecode transforms your data into your success engine by aligning Google Cloud's most advanced capabilities with your specific business objectives. We focus on delivering governed, high-impact outcomes ensuring your Analytics, Data and AI investments translate into a measurable business value.

Investment

- 2 weeks
- $10K

Activities

- AI readiness workshop
- Use case to ROI mapping
- LookML and agent workshop
- Co-development sessions

Impact

- Your team understands where CA fits into analytics strategy
- Build your first LookML and agent optimized CA experience
- Ready to build and launch internally

Related Services

- Analytics Foundation Fast Start
- Data Foundation Fast Start
- Agentic Foundation Fast Start

SOW Template

Google partner funding available.

Bytecode IO

Conversational Analytics Fast Start

Assess readiness and deliver conversational analytics

Work with the Bytecode IO team to assess AI use cases, develop a production ready CA experience and implement the foundational development processes to delivery quality AI powered workflows

Assess use cases and map into ROI

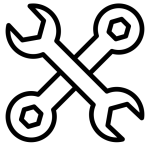

Optimise LookML and build data agent

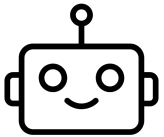

Deliver first CA agent experience

Develop roadmap of agentic value

Bytecode transforms your data into your success engine by aligning Google Cloud's most advanced capabilities with your specific business objectives. We focus on delivering governed, high-impact outcomes ensuring your Analytics, Data and AI investments translate into a measurable business value.

Investment

- 4 weeks
- $20K

Activities

- AI readiness workshop
- Use case to ROI mapping
- LookML and agent workshop
- Evaluation framework co-dev
- Production ready CA experience

Impact

- Your team understands where AI fits into analytics strategy
- Launch your first LookML and agent optimized CA experience
- Ready to build and launch internally

Related Services

- Data Foundation Fast Start
- Agentic Foundation Fast Start

Google partner funding available.

Bytecode IO

Looker Powered AI Value Engine

Developing agentic workflows and unlocking maximum value

Work with the Bytecode IO team establish a repeatable development and launch process to deliver AI powered experiences to end users across the enterprise

Intake framework to assess, plan and build to max ROI

Optimize Looker, Vertex and ADK to deliver AI workflows

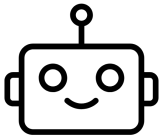

Deliver a Looker powered agent workflow to users

Develop roadmap of agentic value

Bytecode transforms your data into your success engine by aligning Google Cloud's most advanced capabilities with your specific business objectives. We focus on delivering governed, high-impact outcomes ensuring your Analytics, Data and AI investments translate into a measurable business value.

Investment

- 8 -10 weeks
- $50K

Activities

- AI readiness workshop
- Agent solution architecture
- LookML and AI workshop
- Agent development & deployment
- Production ready agent capability

Impact

- Your team understands where AI fits into analytics strategy
- Launch an AI powered workflow on Google Cloud Platform
- Roadmap to deliver maximum ROI agent workflows

Related Services

- Data Foundation Fast Start
- Agentic Foundation Fast Start

Google partner funding available.

Bytecode IO

Thank you

Governing the AI Experience

Granular Permissions for Looker Conversational Analytics

gemini_in_looker

The foundational permission. Enables the core Gemini UI and AI capabilities across the Looker platform for the user.

chat_with_agent

Interaction access. Grants permission to converse with pre-configured agents within the Looker Agent Builder interface.

chat_with_explore

Ad-hoc analysis. Allows users to "Ask Gemini" directly inside any Explore to generate real-time governed queries.

save_agents

Allows users to build, save, and modify their own custom agent configurations and instructions.

publish_agent_externally

Grants authority to deploy agents into external applications via the ADK or embedded contexts.

admin_agents

Administrative rights to view, manage, and delete all agents across the Looker instance.

Agent sharing: permissioned within Looker

Governed Consistency

Share pre-trained agents to ensure all teams query the single source of truth, eliminating data debates and driving uniform decisions.

Enhanced Collaboration and Reuse

Teams easily create and share agents (with groups or users) pre-configured with approved business context (e.g., Finance), accelerating data adoption and knowledge sharing.

Instant Insights

Deploying pre-trained agents drastically reduces setup time, enabling faster, actionable insights across specialized business units.

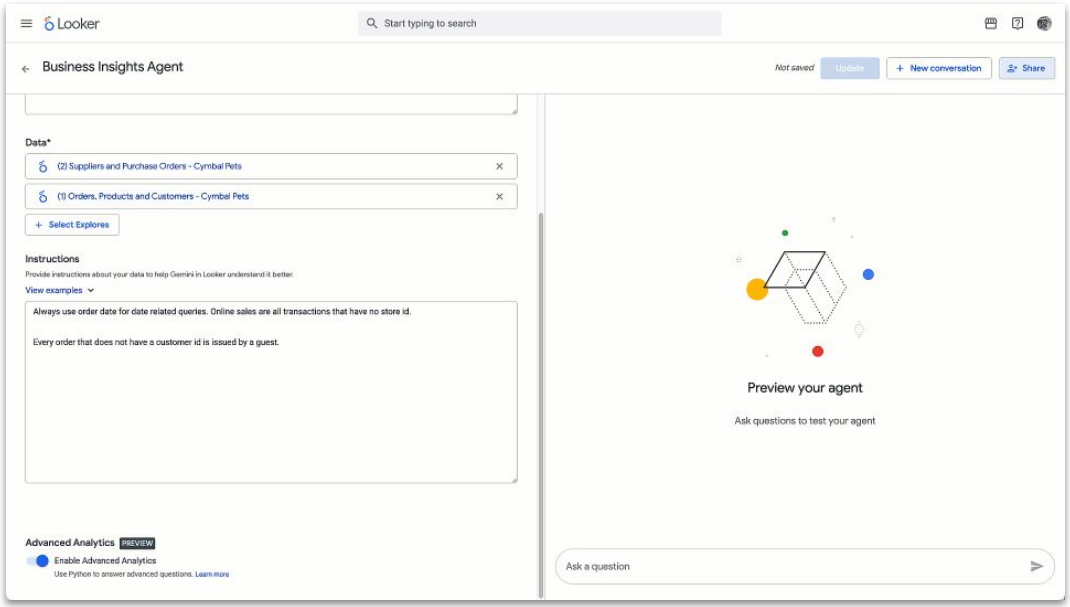

Multi-explore support

Agents will allow up to 5 Looker Explores to choose from.

Enables cross-domain analysis

Users can ask questions that seamlessly integrate insights from up to five distinct Looker Explores, spanning multiple business areas.

Deepens self-service exploration & delivers comprehensive answers

Boosts user autonomy by combining relevant Explores for complex investigations, leading to richer, more holistic insights and uncovering hidden correlations.