

## GEMINI ENTERPRISE

### The Benefits of GraphRAG

Explore the technique that significantly enhances accuracy, traceability, and explainability of your AI outputs to reduce hallucinations.



A simple, secure, and trustworthy generative AI experience to accelerate enterprise rapid decision making

# Contents

Introduction	<b>1</b>
Benefits of GraphRAG	<b>2</b>
• Higher Accuracy and Comprehensive Answers	<b>2</b>
• Enhanced Data Understanding and Insights	<b>3</b>
• Governance: Verifiability and Data Security	<b>5</b>
How Gemini Enterprise Works and Benefits Your Business	<b>6</b>
Conclusion	<b>7</b>
About Gemini Enterprise	<b>8</b>
Deployment Options	<b>10</b>

A simple, secure, and trustworthy generative AI experience to accelerate enterprise rapid decision making

Gemini Data's flagship enterprise AI assistant platform – Gemini Enterprise, is built based on **GraphRAG**, a technique that **significantly enhances accuracy, traceability, and explainability to reduce hallucination in AI outputs.**

With Gemini Enterprise, organizations can transform their enterprise data into a semantic layer of understanding, enabling more effective reasoning and retrieval for large language models (LLMs).

### The benefits of GraphRAG fall into three main categories:

#### 1.Higher Accuracy and Comprehensive Answers:

GraphRAG leverages the interconnected structure of knowledge graphs to produce highly accurate, contextually relevant answers tailored to complex, multi-layered queries.

#### 2.Enhanced Data Understanding and Insights:

GraphRAG enables a deeper understanding of data by creating an interconnected network of information that reveals patterns and connections, driving compounding insights and scalable data expansion.

#### 3.Governance: Verifiability and Data Security:

With advanced explainability, traceability, and robust access control mechanisms, GraphRAG supports data governance by ensuring all data interactions are transparent, verifiable, and securely managed.

### DEFINITION

#### RAG (Retrieval Augmented Generation):

RAG is a **technique that enhances the accuracy and reliability of generative AI** models by incorporating facts retrieved from external sources.

**Language models often struggle with knowledge-intensive tasks** due to limitations of their training data. With RAG, an LLM's knowledge is expanded beyond its initial dataset, making the model far more accurate, impactful and relevant to enterprise users.

#### GraphRAG (Graph Retrieval-Augmented Generation):

GraphRAG is a methodology that **integrates knowledge graphs with RAG models to enhance the process of data analysis and natural language generation.** By incorporating a structured knowledge graph, GraphRAG enables AI to produce more accurate and contextually rich responses.

This integration leverages the interconnectedness of data points within the knowledge graph, providing deeper insights and improving the overall performance of the AI in generating relevant and useful outputs.

GraphRAG is particularly valuable in complex scenarios, where understanding the relationships between different pieces of information is crucial for accurate and comprehensive analysis.

By harnessing the combined power of knowledge graphs and generative AI, enterprises can unlock powerful opportunities for knowledge-driven decision-making, innovation, and operational efficiency.

Knowledge graphs, with their structured representation of a domain, enhance the performance of generative AI by providing context, validating outputs, and reducing biases – ensuring alignment with strategic business objectives.

In return, generative AI enriches knowledge graphs by filling knowledge gaps and predicting future trends, thereby increasing the utility, accuracy, and relevance of these graphs.

The synergy between knowledge graphs and generative AI serves as a game-changer for businesses, driving transformative impacts across various organizational functions.

## THE BENEFITS OF GRAPH RAG

### #1: Higher Accuracy and Comprehensive Answers

GraphRAG's most immediate and significant benefit is its ability to produce higher-quality responses. This improvement is evident in Gemini Data's numerous customer use cases, as well as growing academic research.

Microsoft, in their [research blog](#) and [associated research paper](#), identified two key problems with baseline RAG (which relies on vectors):

#### 1. **Baseline RAG struggles to connect the dots.**

*This happens when answering a question requires traversing disparate pieces of information through their shared attributes in order to provide new synthesized insights.*

#### 2. **Baseline RAG performs poorly when being asked to holistically understand** summarized semantic concepts over large data collections or even singular large documents.

Microsoft found that "GraphRAG vastly improves the 'retrieval' portion of RAG, populating the context window with higher relevance content, resulting in better answers and capturing evidence provenance."

For instance, a [study](#) by data catalog company Data.World tested whether an LLM powered with knowledge graphs would produce more accurate answers compared to one powered with SQL databases.

Their results show that compared to SQL databases, knowledge graphs improve the accuracy of LLM responses threefold, tested across 43 business-related questions.

### 3x Accuracy when answering business-related questions

In addition, **GraphRAG is more scalable and cost-effective than baseline RAG.** In a [separate research paper](#), Microsoft found that LLMs powered by GraphRAG required significantly fewer context tokens than those using baseline RAG, while still delivering more comprehensive responses.

Specifically, for "a modest drop in performance," GraphRAG required over 97% fewer context tokens than baseline RAG. Even at maximum performance, GraphRAG required 26-33% fewer tokens.

### Up to 97% fewer context tokens compared to baseline RAG

GraphRAG enables LLMs to deliver answers that are not only more accurate but also highly **relevant and useful**.

In a [recent case study](#) by LinkedIn on GraphRAG's impact within their customer service application, results demonstrated a "significant improvement" in answer quality, along with a 77.6% increase in retrieval accuracy and a 28.6% reduction in resolution time per issue.

### 28.6% reduction in resolution time



## Vector Only approach

Vector Only Search does not have context and it is something like this:

Based on the context provided, the asset managers most vulnerable to lithium shortage are those with significant exposure to companies that rely heavily on lithium batteries, such as electric vehicle manufacturers. Specifically:

- Asset managers with large positions in electric vehicle companies like Tesla, NIO, Li Auto, and Xpeng may face risks if those companies are unable to obtain sufficient lithium supplies for their vehicle batteries. This could negatively impact their production capacity and profitability.
- Asset managers invested in battery producers that rely on lithium, such as Panasonic, LG Chem, and Samsung SDI, may also face risks if those suppliers are unable to source enough lithium to meet demand. This could limit their growth and earnings potential.
- Firms with holdings in lithium mining companies and lithium-based chemical producers could see lower revenues and profits if lithium shortages drive up costs or limit their production volumes. Examples include Albemarle, SQM, Livent, and Ganfeng Lithium.
- More broadly, asset managers overweight in the electric vehicle supply chain may face valuation pressures if lithium shortages negatively impact the growth prospects for the industry.

So in summary, asset managers with the largest exposures to electric vehicle manufacturers, lithium battery producers, lithium miners, and the EV supply chain in general appear most vulnerable to risks from tighter lithium supplies. Their holdings could face growth headwinds if lithium shortages persist.

Done!

## Vector + Graph approach

Vector+Graph has full context like this:

Based on the context provided, the asset managers most vulnerable to lithium shortage are:

- Quantamental Technologies LLC
- DEPRINCE RACE & ZOLLO INC
- Atria Wealth Solutions, Inc.
- CITIGROUP INC
- CHARLES SCHWAB INVESTMENT MANAGEMENT INC
- ROYAL LONDON ASSET MANAGEMENT LTD
- STIFEL FINANCIAL CORP
- TUCKER ASSET MANAGEMENT LLC
- Patriot Financial Group Insurance Agency, LLC
- SeaCrest Wealth Management, LLC

The context indicates that Stanley Black & Decker uses lithium batteries in many of its products. The increased demand for lithium could lead to supply constraints and increased costs. This would negatively impact the profitability of Stanley Black & Decker.

The asset managers listed above all have holdings in Stanley Black & Decker, and therefore may be vulnerable to the risks of lithium shortage impacting that company. The context provides the asset manager name and holdings for Stanley Black & Decker, so I have listed the asset managers with the highest holdings based on the score provided.

Done!

Figure 1: Sample query comparison between one LLM using vectorRAG and another using GraphRAG.

Figure 1 shows a sample query from Neo4j's GenAI workshop, which shows two LLM-generated answers: one using vectorRAG, and one using vector in conjunction with GraphRAG.

Note the difference between describing the characteristics of companies likely to be impacted by a lithium shortage, as opposed to listing specific companies that are likely to be.

For investors or companies looking to adjust their portfolio in response to market changes or supply chain in the aftermath of a natural disaster, the latter response is far more valuable.

**GraphRAG enables AI to output more accurate AND actionable answers.**

## #2: Enhanced Data Understanding and Insights

Knowledge graphs and GraphRAG offers transformative advantages for enterprises by driving exponential growth in data understanding and operational efficiency.

By creating an interconnected network of diverse information sources, organizations can unlock compounding insights that deepen with each new data point.

As new data sources are integrated, each addition compounds previous knowledge, enabling the discovery of patterns that drive **informed, data-driven strategies**.

To illustrate the transformative benefits of GraphRAG and knowledge graphs in real-world applications, we will examine two case studies from *IBM Research* and *Siemens Energy*.

These cases demonstrate how organizations can leverage interconnected data networks to enhance operational efficiency, achieve greater accuracy, and drive continuous innovation.

In their research paper, IBM Research described how they faced the challenge of managing scattered, siloed data across a wide range of research projects, publications, and collaborations.

Traditional methods were unable to effectively connect and leverage this complex information, leading IBM to develop a knowledge graph framework.

By linking structured and unstructured data sources through a unified knowledge graph, IBM created a living network of information that enables **more personalized recommendations and advanced trend analysis**.

In an article from 2020, Amazon Web Services (AWS) described how Siemens Energy, with their global fleet of gas turbines, encountered logistical challenges in managing assets across different environments and contexts.

Tracking inventory, managing spare parts, and meeting specific customer requirements demanded significant manual effort.

Siemens then decided to adopt the metaphactory knowledge graph, transforming their fleet management operations and improving data organization, enhancing part recommendations, and saving thousands of hours in manual processing.

GraphRAG's integration with knowledge graphs in the above cases highlights the exponential growth in data insights that organizations can achieve:

### 1. Network Effect of Data:

Knowledge graphs enable a "network effect" in data, where the value and relevance of the entire network expand as each new data point is added.

This interconnected structure allows organizations to uncover previously hidden relationships and patterns, leading to more informed, data-driven decisions.

For Siemens Energy, the graph structure provides a comprehensive view of turbine performance and part requirements, enhancing visibility across the entire fleet and allowing for smarter maintenance decisions.

For instance, consider a scenario where a turbine in Singapore shows a recurring performance issue, such as overheating in a specific component.

With the knowledge graph, Siemens Energy identifies that similar turbines in Malaysia and Indonesia have experienced the same issue under similar conditions.

By analyzing the network of related data points, the graph highlights a common supplier of the faulty component and a climate-related stress factor.

With this insight, Siemens Energy can:

- Alert all affected sites to **proactively replace** the faulty component, preventing further failures.
- Optimize supply chain logistics by **redirecting available stock** from the nearest regional warehouse to minimize downtime.
- Collaborate with the supplier to **resolve the defect** and prevent future occurrences.

This network effect creates a ripple of value across the organization, where each data point (e.g., a single turbine's performance metric) contributes to a broader understanding, improving maintenance strategies, supplier relationships, and operational efficiency.

### 2. Accelerated Operational Efficiency:

Both IBM and Siemens achieved significant improvements in operational efficiency by deploying knowledge graph solutions.

Using their metaphactory knowledge graph, Siemens **saved 1,500 hours of manual processing in the first year** alone by automating parts catalog management and configuration tracking.

IBM's knowledge graph helped streamline collaboration and minimize redundancies in research workflows, allowing researchers to focus on innovation rather than administrative tasks.

### 3. Scalability and Adaptability:

GraphRAG provides a scalable framework that supports continuous growth and adaptation.

IBM's KGI framework easily integrates with external sources, creating a robust, scalable solution adaptable to various research needs.

Gemini Enterprise too, is compatible with various data sources and data management systems, such as MySQL, Splunk, or Salesforce.

GraphRAG empowers organizations to transform raw data into a dynamic knowledge network, delivering exponential growth in insights and efficiency.

The interconnected structure of knowledge graphs enables continuous improvement and refinement, creating a self-strengthening ecosystem that scales as data needs grow.

By leveraging the compounding insights enabled by GraphRAG, enterprises can achieve faster, more scalable decision-making, unlocking a competitive advantage in a data-driven world.

### #3: Governance: Verifiability and Data Security

As the impact of GenAI decisions grows, so does the need to build trust toward the people who are accountable for their outcomes.

This typically involves being able to audit each decision, backed by a solid and consistent track record of sound judgments. However, that alone is insufficient.

**It is also critical to be able to explain the underlying logic** if and when decisions are scrutinized.

Knowledge graphs enhance governance capabilities by being inherently navigable and queryable, allowing decision-makers to closely examine how the GenAI retrieved its answer.

Figure 2 is a snapshot of a Company X's supply chain data in Gemini Enterprise, showing that the supplier Lukes Micro provides the company with several parts, including Main PCBs.

Figure 2 shows that Alpha Industry is also a Main PCB supplier, but located at a significantly closer distance and delivers orders at a faster time.

Viewers of the graph can then infer based on this information, Company X can potentially source more of their Main PCB parts from Alpha Industry who can supply parts in a more timely manner.



Figure 2: Example snapshot of Company X's supply chain data in Gemini Enterprise

**By storing data using knowledge graphs, enterprises can easily verify the accuracy and relevancy of GenAI responses.**

# Retrieval Augmented Generation Explained - HybridRAG

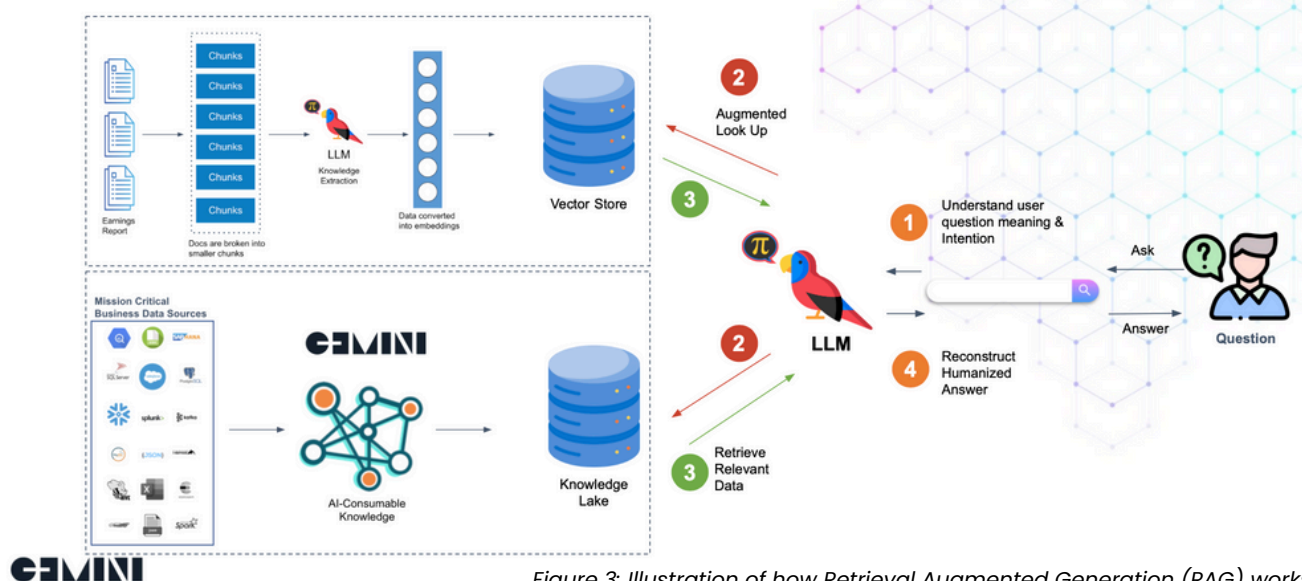


Figure 3: Illustration of how Retrieval Augmented Generation (RAG) works

## HOW GEMINI ENTERPRISE'S GRAPHRAG WORKS

At its current state, LLMs are not able to reliably comprehend and analyze raw data.

Therefore, Gemini Enterprise helps ground the LLM in the data by turning it into AI-consumable semantic knowledge.

**Step 1:** LLM understands the user question's meaning and intention.

**Step 2:** LLM queries the Knowledge Lake, which is filled with AI-Consumable Knowledge derived from user-assigned data sources.

**Step 3:** Knowledge Lake sends relevant data based on LLM-written query.

**Step 4:** LLM uses retrieved relevant data to reconstruct a humanized answer..

## How Gemini Enterprise Benefits Your Business

Gemini Enterprise provides a simple, secure, and trustworthy generative AI experience to accelerate enterprise rapid decision making.

Building on these foundational capabilities, Gemini Enterprise extends its impact across several critical dimensions, including specificity, security, and traceability.

By incorporating these advanced features, Gemini Enterprise enables organizations to harness generative AI with confidence, knowing that each interaction is tailored, secure, and fully traceable.

### Accuracy

With GraphRAG, LLMs can help generate more accurate responses by understanding the connections and dependencies within your data and mapping new links as fresh data becomes available.

### Data Relationships

Leverage the relationships within your data to deliver precise, explicit results.

### Flexible Schema

Easily expand your knowledge base as new information emerges, ensuring continuous growth and adaptability.

### Knowledge Search & Retrieval

Uncover implicit insights and showcase similar responses by searching within the knowledge graph.

### Specificity

By adding a semantic layer of understanding to your business and your data, Gemini Enterprise equips LLMs with information specific to your organization.

This ensures the answers are tailored to your organization's needs.



## Privacy & Security

Gemini Enterprise ensures that data retrieval and usage align with your enterprise security and access control policies, down to the node level.

Here's an illustration of how you can implement fine-grained access controls in a knowledge data store to ensure data security:

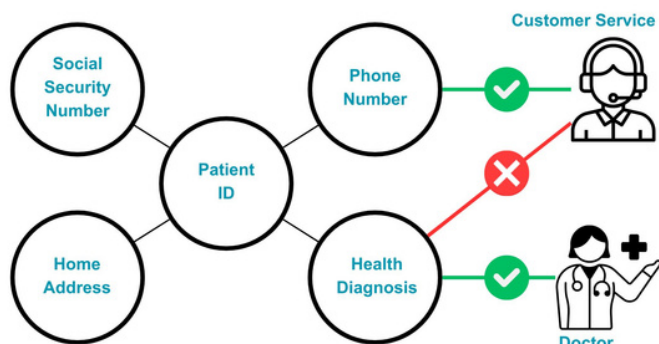


Figure 4: Illustration of how fine-grained access controls can be implemented in a knowledge data store

Gemini Enterprise significantly enhances security and privacy crucial for any organization, especially those in industries such as banking or healthcare, where data access is role-based.

Unlike LLMs or vector databases, which lack robust user-based access controls, an enterprise knowledge graph allows you to govern data access through the database and its underlying model, ensuring that users only see the results they are permitted to view.

### Explainability & Traceability

Gemini Enterprise enables verifiability of each LLM answer by linking each piece of information to its original source.

- Represent data sources as nodes
- Map relationships between results and data source nodes
- Add metadata or annotations

```
Analysis
Total revenue for the month by store location

1 MATCH (sl:StoreLocation)-[:BELONGS_IN]->(d:District)-
  [:WITHIN]->(r:Region)
2 RETURN
3 sl.Name AS StoreLocation,
4 toFloat(sl.`Total_Revenue_For_Month`) AS
  TotalRevenueForMonth
5 ORDER BY TotalRevenueForMonth DESC
6 LIMIT 20
```

Figure 5: An LLM-generated Cypher query retrieving data from the Knowledge Lake based on the user's question.

## CONCLUSION

Semantic knowledge data structure and GraphRAG offer a transformative approach to data management and decision-making, providing unmatched accuracy, scalability, and operational efficiency.

By leveraging these technologies, Gemini Enterprise enables organizations to create an interconnected knowledge network, revealing insights that drive competitive advantage.

### We recommend that organizations:

- Identify and formalize their metadata sources with taxonomies and classifiers.
- Build graph data models by transforming relational schemas and documents into key relationships.
- Use knowledge graphs for enhanced data access, monitoring, and governance.
- Layer AI services on top of this architecture to unlock the full potential of the data.

As data continues to grow as a vital asset, Gemini Enterprise positions businesses to thrive by providing a scalable, adaptable framework that ensures high-quality, data-driven decisions.

Now is the time to explore how Gemini Enterprise can transform your data strategy.

[geminidata.com](https://geminidata.com)

+1 800 549 7888

Amsterdam | Irvine | San Diego  
San Francisco | Taipei | Tokyo

## Contact Us

for a personalized demo to see how our solutions can elevate your approach to AI and data governance.

# About Gemini Enterprise

## Our Mission is to redefine How Data is Used, Analyzed, and Shared

We've built the world's simplest enterprise AI assistant platform to help companies all over the world connect the dots and get answers faster, respond to unusual events, and take the next best action.

## Generative AI for the Enterprise

While AI has long been powerful, it was limited by its complexity, making it accessible only to engineers and data scientists

However, generative AI has transformed this landscape. Now, getting answers from AI is as easy as sending a text to a friend.

**However, businesses need more than just an aggregated, averaged, or statistically probable answer.**

**They need the answer, sourced from their own data, along with a clear understanding of how that answer was reached.**

Achieving this requires data and applications that are optimized for the new world of generative AI.

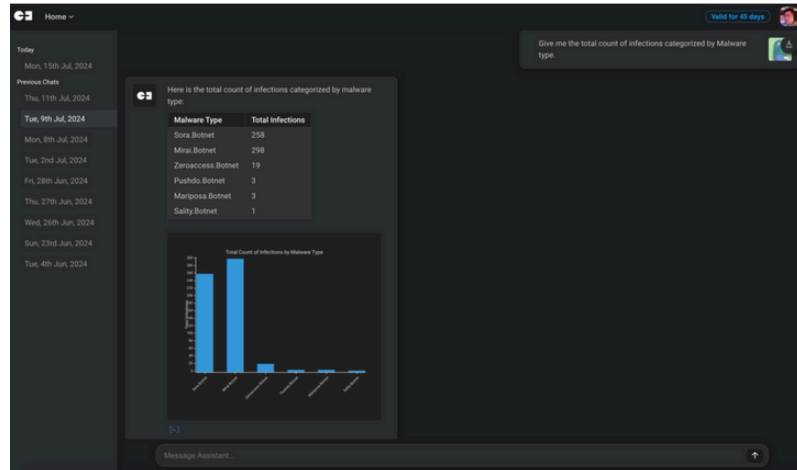
Gemini Data equips every enterprise with the tools necessary to immediately integrate generative AI technologies into their existing infrastructure.

## Modern Enterprise AI Assistant Platform Built for the Future

Gemini Enterprise is engineered to democratize data by enabling technical and business users alike to take full advantage of graph data without needing to understand the underlying algorithms or query languages.

## • Conversational AI Experience Made Easy

Our enterprise AI assistant lets you ask questions through natural, conversational queries and receive comprehensive answers in text, tables, or visual charts.



## • AI Generated Reports

Transform enterprise data into analytical knowledge. Prompt AI to generate reports based on your business needs to speed up decision-making.

Title: Jayson Tatum Shines with Dominant Performance in Recent Game

Introduction:  
In a recent game, Boston Celtics' forward Jayson Tatum exhibited his exceptional skills and proved why he is regarded as one of the brightest stars in the NBA. Tatum's remarkable performance showcased his versatility, scoring prowess, and impact on both ends of the court. Let's delve into the details of his impressive display.

Scoring Prowess:  
Tatum started the game on a high note, making his presence felt with a powerful driving dunk, earning himself two points. Throughout the game, he consistently found ways to contribute to his team's scoring efforts. The talented forward showcased his shooting range, sinking multiple three-point shots from long distance. Tatum's ability to create his own shot and score from various positions on the court was evident, as he tallied an impressive total of 20 points. His offensive contributions were crucial in propelling the Celtics to victory.

Rebounding and Defense:  
Apart from his scoring prowess, Tatum's impact on the defensive end was equally notable. He displayed his tenacity and determination by grabbing rebounds on both ends of the court. With a total of nine rebounds, Tatum showcased his ability to control the boards and provide his team with second-chance opportunities. Additionally, his defensive skills were on full display as he recorded a steal, disrupting the opponents' offensive flow.

Clutch Performances:  
As the game progressed, Tatum's ability to rise to the occasion became evident. In crucial moments, he demonstrated his composure and ability to deliver when it mattered most. Tatum's clutch shooting was exemplified by a remarkable 27-foot three-point shot, which not only extended his team's lead but also energized both his teammates and the crowd. His ability to perform under pressure solidified his status as a go-to player for the Celtics.

Additional prompt (optional)  
Please write a professional sports article featuring the player Jayson Tatum

total 5406 tokens.

Save Copy to Clipboard Regenerate

## About Gemini Data

Our mission is to redefine how data is used, analyzed, and shared. We've built the world's simplest enterprise AI assistant platform to help companies all over the world connect the dots and get answers faster, respond to unusual events, and take the next best action.

© Gemini Data, Inc.

**geminidata.com**

+1 800 549 7888

Amsterdam | Irvine | San Diego  
San Francisco | Taipei | Tokyo

# About Gemini Enterprise

- **Collaboration and Sharing**

Organizations can create multiple data projects, each connected to specific data sources and mapped to a tailored semantic knowledge model, accessible by authorized users or user groups, enhancing collaboration while minimizing the risk of data leaks.

## Connect Your Data to AI Fast

Gemini Enterprise is a modern Enterprise AI Assistant platform that helps organizations start their generative AI journey quickly and efficiently.

A cloud-native, no-code, wizard-driven platform ensures your data and analytics projects to be powered generative by AI in minutes or hours instead of weeks or months.

- **Easily Connect to Enterprise Data**

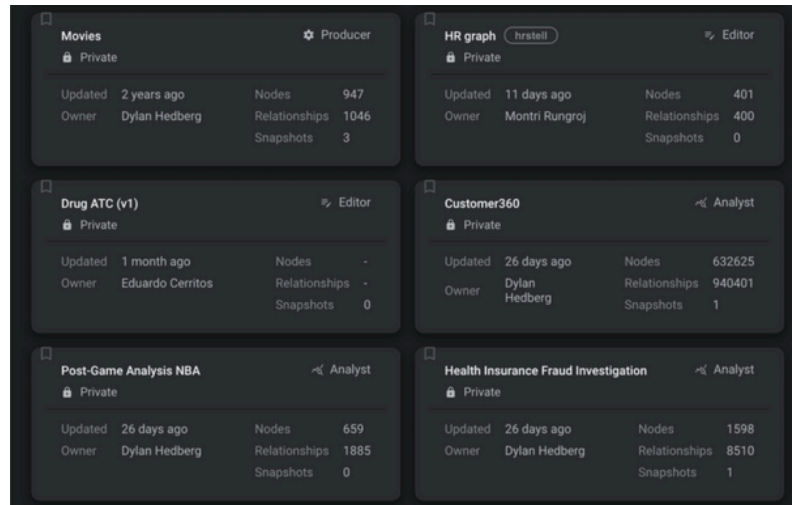
Gemini Enterprise Data Importer makes it simple to setup, connect, and manage various data sources with built-in connectors and a user friendly, guided interface.

- **No-Code**

No specialized IT or query language skills required. Easily import disparate data and create a semantic layer of understanding on top of your enterprise data.

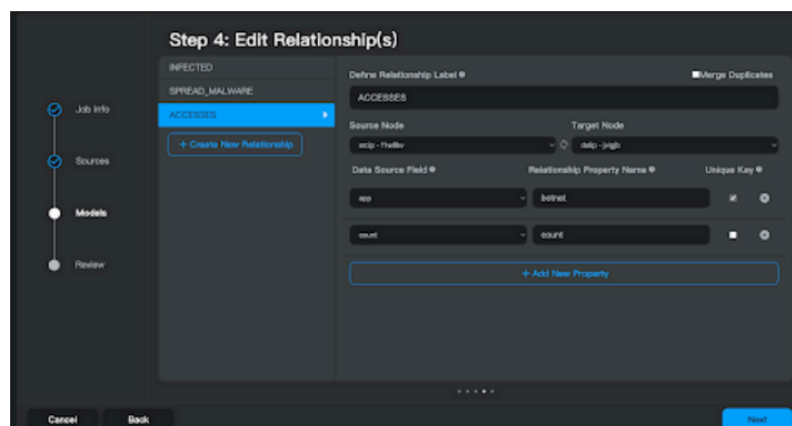
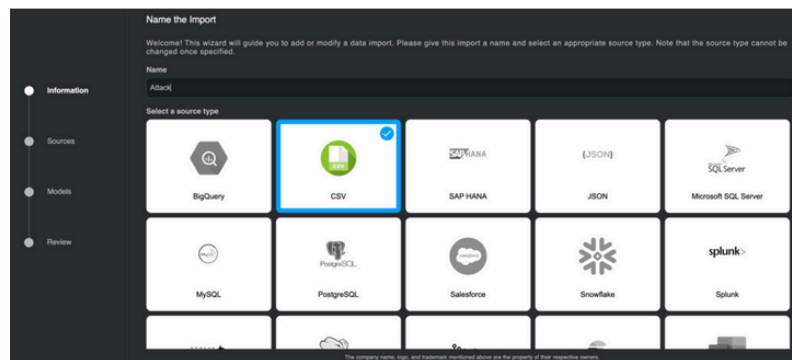
## Interact With Enterprise Knowledge The Smart Way

Gemini Enterprise is powerful and intuitive. Easily interact with data using keyword search, custom calculations, conditional display, and personalized data enrichment.



A screenshot of the Gemini Enterprise data projects dashboard. It displays a grid of project cards, each representing a different data project. Each card shows the project name, its status (Private), the owner, the update date, the number of nodes, relationships, and snapshots. The projects include 'Movies', 'HR graph', 'Drug ATC (v1)', 'Customer360', 'Post-Game Analysis NBA', and 'Health Insurance Fraud Investigation'.

Project Name	Status	Owner	Updated	Nodes	Relationships	Snapshots
Movies	Private	Dylan Hedberg	2 years ago	947	1046	3
HR graph	Private	Montri Rungroj	11 days ago	401	400	0
Drug ATC (v1)	Private	Eduardo Cerritos	1 month ago	-	-	0
Customer360	Private	Dylan Hedberg	26 days ago	632625	940401	1
Post-Game Analysis NBA	Private	Dylan Hedberg	26 days ago	659	1885	0
Health Insurance Fraud Investigation	Private	Dylan Hedberg	26 days ago	1598	8510	1



## About Gemini Data

Our mission is to redefine how data is used, analyzed, and shared. We've built the world's simplest enterprise AI assistant platform to help companies all over the world connect the dots and get answers faster, respond to unusual events, and take the next best action.

© Gemini Data, Inc.

[geminidata.com](https://geminidata.com)

+1 800 549 7888

Amsterdam | Irvine | San Diego  
San Francisco | Taipei | Tokyo

## About Gemini Enterprise

- **Guided Search**

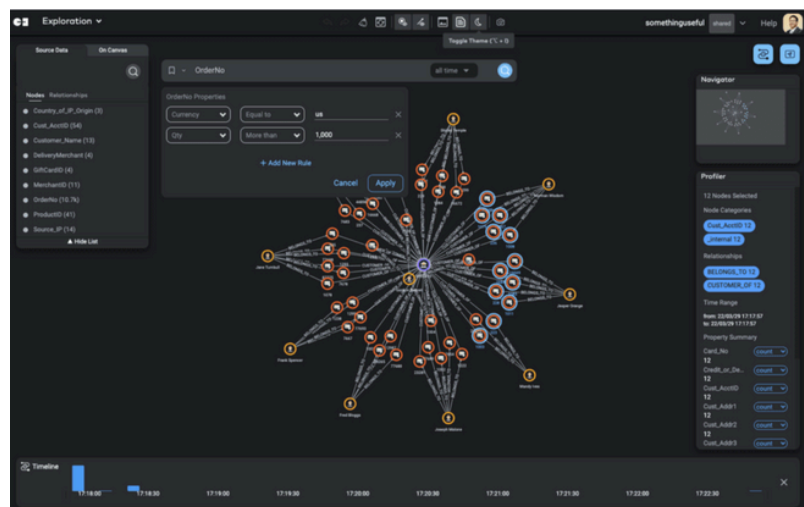
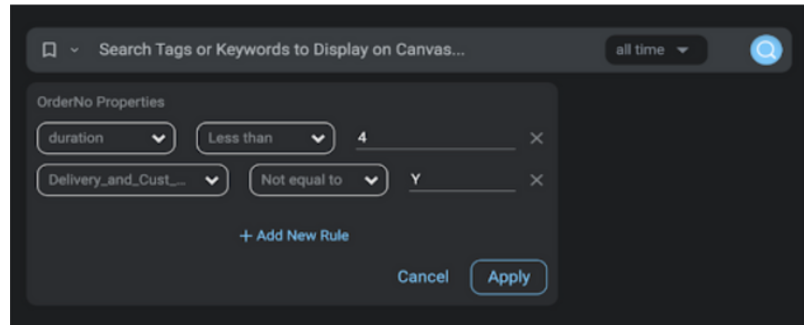
Powerful search allows users to find the proverbial needle in the haystack. With flexible queries and conditional filters, users can precisely curate how enterprise knowledge shows up on the exploration canvas.

- **Flexible Exploration**

Gemini Enterprise Data Explorer gives users all the tools they need to interact with and visualize the enterprise knowledge to uncover hidden clusters, patterns, and outliers.

- **Personalization**

Gemini Enterprise Data Explorer ships with a full suite of styles and customizations such as custom icons, colors and sizes to ensure every piece of enterprise knowledge is presented clearly and compellingly.



## Deployment Options

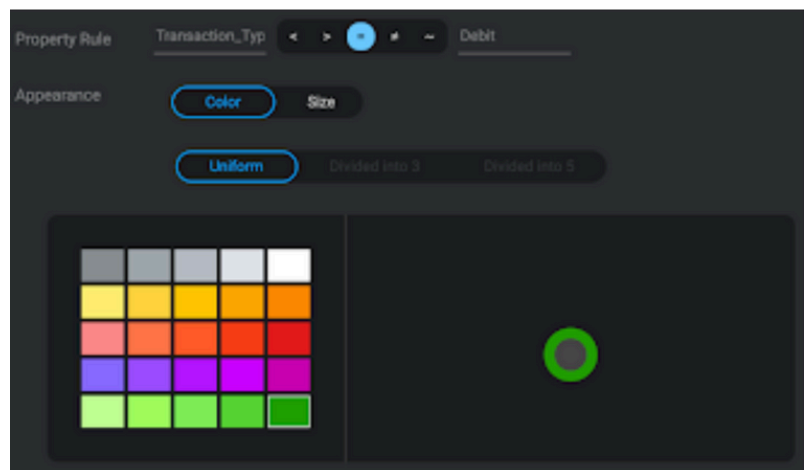
- **SaaS**

Gemini Enterprise is available as a Software as a Service subscription which you could sign up and connect with your enterprise data. No admin or infrastructure required.

- **On-premise**

Gemini Enterprise can be deployed on-premise in your private cloud, virtualized environment or in your own hardware.

We also offer a turn-key appliance for ease of deployment and management at your choice.



## About Gemini Data

Our mission is to redefine how data is used, analyzed, and shared. We've built the world's simplest enterprise AI assistant platform to help companies all over the world connect the dots and get answers faster, respond to unusual events, and take the next best action.

© Gemini Data, Inc.

[geminidata.com](https://geminidata.com)

+1 800 549 7888

Amsterdam | Irvine | San Diego  
San Francisco | Taipei | Tokyo