

Retrieval-augmented generation (RAG): How data becomes generative AI

The most challenging part of building a viable AI model is accessing the large volumes of data necessary to train AI, feeding it to the model and ensuring that it stays up to date.

Modern-day enterprise data originates in applications, databases, file storage systems and event streams. This data needs to be extracted and loaded to a central data repository.

Once data lands in a data lake or data warehouse, it must undergo additional processing to be made useful for generative AI.


One approach is to catalog and assign semantic, real-world meaning to data, constructing a referenceable knowledge graph that illustrates the relationships between entities, events and concepts.

The other, more pertinent approach is to embed raw data – text, images, video and other media – as numerical representations called vectors in a vector database.

Vectorized data can be attached to a request made by a user, called a prompt, and sent to a foundation model as additional context, providing more accurate and relevant answers.

This architecture, in which a retrieval model augments prompts with additional context, is called retrieval-augmented generation (RAG).

RAG enables organizations to leverage the power of off-the-shelf generative AI for their proprietary data, saving money and time.



How Fivetran fuels your GenAI workloads

Fivetran is the global leader in data movement, helping enterprises use their data to power everything from AI applications and ML models to predictive analytics and operational workloads. The Fivetran platform reliably and securely centralizes data from hundreds of SaaS applications and databases into any cloud destination — whether deployed on-premises, in the cloud or in a hybrid approach.

Key capabilities & features

How Fivetran prepares teams for generative AI



Data integration services:

- 500+ no code connectors into all major destinations (warehouses, data lakes, etc)
- Ingests text-rich data and proprietary data held in databases that are key for AI model training
- Best-in-class change data capture and 99.9% uptime, ensuring data is accurate, secure and always available



Pipeline flexibility:

- Fivetran's governed approach to data lakes helps organizations store rich, voluminous AI-ready data at scale
- Platform extensibility (transformation orchestration, REST API)
- Cloud-, deployment- and destination-agnostic so data teams can build a future-proof solution



Data governance:

- Highest levels of compliance and security (CCPA, GDPR, HIPAA, ISO, SOC2)
- Centralized security for standardized governance (RBAC, Teams, SCIM, metadata)
- Integrations with data governance and observability tools, ensuring sensitive data used to train AI is never exposed

Thousands of companies trust Fivetran to move business-critical data



"Fivetran has revolutionized our approach to data, enabling AI/ML and GenAI initiatives on employee performance and providing managers with faster insights. These advancements were unimaginable just two years ago."

Sandro Frattura, Analytics Engineering Manager,
HubSpot People Operations