




Steps for Transforming Data Silos with AI Data Assistants

High-Level Guidance




01 - Select a Large Language Model (LLM)

The LLM forms the foundation of knowledge and intelligence used to help derive answers and provide assistance.




02- Configure a Retrieval-Augmented Generation (RAG) Capability

RAG is a popular gen AI implementation pattern and technique that uses enterprise-specific content combined with the intelligence of the LLM to prepare prompts and provide answers.




03- Collect and Consume Existing Assets into the RAG

Existing assets include a wide variety of file formats, catalogs, glossaries, learning material, process descriptions, reports, data dictionaries, and more. During this step, you chunk, embed, and vectorize.



04 - Test the AI Data Assistant

Use a battery of questions and an ad-hoc approach to exploring what the AI Data Assistant understands and can answer.



05 - Tweak the Existing Asset RAG Format and Repeat Testing

Sometimes a bit of tweaking to better represent enterprise-specific content as language can improve the accuracy of responses.