

FileOrbis AI Enabler & AI Governance




Overview

FileOrbis AI Enabler and AI Governance turn the knowledge in your file servers into secure, permission-aware and verified AI answers. AI Enabler keeps that knowledge in a continuously synchronized index, tagged with access permissions and protected against sensitive-data exposure. AI Governance answers every question through a chain of specialized AI agents that check each other's work — an approach known as Agentic RAG.

The result: every user gets an answer that truly addresses their question, built only from data they are authorized to see.

How it works

Answers come only from authorized, protected data — and are verified



AI Enabler and AI Governance flow · two processes, two loops

The left side keeps your data continuously prepared; the right side answers every question only from this permission-aware knowledge base and searches again when an answer falls short.

AI Enabler: Makes your data AI-ready

AI Enabler turns your file servers into an always-current knowledge base that AI can use safely. The process is fully automated and runs continuously.

1

Continuous sync and indexing

Content from all file servers is ingested into the knowledge base. Changes are detected in real time and reflected in the index, so AI does not work with stale data.

2

Permission mapping (permission-aware indexing)

Every content chunk carries its source file's access permissions (ACLs) as metadata. Your file-server permission model is preserved one-to-one.

3

Sensitive data protection

Personal and confidential data is discovered and, according to your policy, excluded from the index, masked or anonymized.

4

Continuous governance loop

Steps 1–3 re-run automatically for every new, modified or deleted file and every permission change.



AI Governance: Accurate, authorized, verified answers

Instead of single-pass RAG, AI Governance uses a chain of specialized agents that check each other's work. Each agent focuses on one job; the result is answers that are more accurate and easier to audit.

1

Query understanding

The first agent interprets the user's question and pins down what is actually being asked.

2

Query expansion

A search agent generates multiple search queries to find every piece of information the answer needs, including documents that use different wording.

3

Permission-aware retrieval

Only chunks the user is authorized to access are retrieved; the most relevant are selected and sent to the LLM together with the question. The LLM never sees data the user cannot access.

4

Answer evaluation

An evaluation agent checks the generated answer against the question. If it falls short, the system loops back to step 2 and searches again; if it is sufficient, the answer goes to the user.

Why FileOrbis



Permissions carry over to AI

If a user cannot open a file on the file server, they cannot reach its content through AI either.



Always up to date

The knowledge base stays in sync with your file servers.



Sensitive data under control

Personal and confidential data is protected before it ever reaches AI, supporting your GDPR and other data-protection compliance efforts.



More accurate answers

The multi-agent verification loop reduces the risk of incomplete or off-topic answers.

Ready to Scale AI without Compromising Security or Budget?

Meet our experts to discover how FileOrbis protects your enterprise data while reducing operational costs

[Request a Demo](#)

[Talk to an Expert](#)




Enterprise content security and governance platform for unstructured data.



FileOrbis Recognized as a Strong Performer in Gartner® Peer Insights™ 'Voice of the Customer' for Document Management

Gartner.
Peer Insights™



Gartner.
Peer Insights™



111 Reviews
as of 30 July 2025.

95% Willingness to Recommend

4.9 ★ Overall Rating

For more information, request a demo.



fileorbis.com



@fileorbis



info@fileorbis.com