

DATASHEET

Speicherkraft AI Vault Application

Local AI (LLM) directly in M-Files – fully on-premises



The Speicherkraft AI Vault Application integrates a freely configurable, locally operated AI connection (LLM) directly into M-Files. All AI calls run exclusively server-side – **no document ever leaves your infrastructure**. Users ask questions about documents, receive AI-powered metadata suggestions and benefit from rule-based background prefill – right in the metadata card of the New Desktop Client and in M-Files Web.

— HIGHLIGHTS

- **AI assistant in the metadata card:** Free-text questions about the selected object; metadata and document content are included as context server-side, answers are formatted and list the context sources.
- **Cross-document questions:** Select multiple objects and compare or summarize them in a single question.
- **Metadata suggestions with one-click apply:** Generate suggestions per rule, tick the values you want and apply them with one click – suggestions are never saved without confirmation.
- **Background prefill:** Rule-based and asynchronous via a persistent task queue with retry on errors; filter conditions, target fields and minimum confidence per rule.
- **Image analysis (vision models):** Multimodal models read image content directly – with automatic fallback to a vision profile if the default profile is not a vision model.
- **Flexible provider profiles:** Multiple profiles in parallel; Ollama as reference, Azure OpenAI and Google Gemini prepared. Timeout, temperature and model per profile.
- **Dashboard in M-Files Admin:** Metrics on configuration, interactive usage, background processing and the last error; failed tasks can be re-run.

In focus: Linked documents checked automatically

Rules can **include linked objects in the AI evaluation**. Example incoming invoice: when the invoice is read, the linked **purchase order** and **delivery note** are read along with it – line items, quantities and amounts can be cross-checked directly. The number of linked objects and the context size are configurable per rule.

— UP AND RUNNING IN THREE STEPS

1

Install

Deploy a single installation package (.mfappx) in M-Files Admin – the “AI Assistant” client tab is installed automatically.

2


Create a provider profile

Enter the endpoint URL and model of your local LLM backend (e.g. Ollama); multiple profiles can run in parallel.

3

Define rules

Set target fields, filter conditions and, optionally, linked objects – suggestions and prefill work immediately.



AI assistant in the M-Files client: metadata suggestions for an incoming invoice – invoice number, supplier, amounts and due date are detected and applied with one click.

— TECHNICAL DATA

Target system	M-Files On-Premises (current server version)
Clients	New Desktop Client and M-Files Web – “AI Assistant” tab in the metadata card
Deployment	One installation package (.mfappx): server application (VAF) incl. automatically installed UIX client
AI providers	Local AI Server / Ollama (reference); Azure OpenAI and Google Gemini prepared; multiple profiles in parallel, API keys stored encrypted
Supported content	PDF (text/OCR layer), Office (.docx/.xlsx/.pptx and .doc/.xls/.ppt), text files (txt, csv, xml, etc.)
Image analysis	PNG, JPG/JPEG, GIF, BMP, WEBP, TIFF – up to 4 images (max. 8 MB each) per object
Data privacy	All LLM calls run exclusively server-side; the client never calls an AI backend directly
Administration	Configuration entirely in M-Files Admin; system prompts adjustable without redeployment; configurable logging
Recommendation	GPU for inference (CPU-only possible but significantly slower); a GPU is practically required for image analysis

Contact & further information

IT-Speicherkraft Vertriebs- und Dienstleistungs GmbH · office@Speicherkraft.com · www.Speicherkraft.com