


Enterprise AI Platform – Your Knowledge. Your Rules.

From finding to answering to acting –
the sovereign foundation for all your AI applications

With iFinder, you can quickly find company information; with iAssistant, you get permission-aware answers from your own data; with iHub, you complete tasks using generative AI. We are now presenting what has always connected these applications for what it is: the Enterprise AI Platform – permission-aware context for every AI application, enabling employees to find information, chatbots to answer questions, and AI agents to take action. Source-backed, deployable under your control, and open to your own assistants and applications via MCP and APIs – in both businesses and public-sector organizations.




One platform instead of multiple point solutions

AI applications are often developed in isolation – each with its own data connection, permission model, and infrastructure. IntraFind's Enterprise AI Platform provides a shared technological foundation on which a wide range of applications can be built incrementally and expanded as needed.

It handles all the heavy lifting in the background: data connectivity while preserving existing access permissions, ingestion, indexing and preprocessing, linguistic normalization, auto-tagging and metadata enrichment – and, for generative AI, passage chunking, embeddings and vectorization. Authentication, authorization, and the integration of interchangeable language and embedding models are included as well.

You can deliberately start small – for many customers, with iFinder, our permission-aware Enterprise Search. On the same platform, you can then add a chatbot, specialized assistants, agents, or iHub – all the way through to your own business applications. There is no need to start from scratch for every new AI use case.

Integrate once – unlock new use cases over time



ENTERPRISE AI PLATFORM – the foundation you no longer have to build yourself

Data Connectivity • Ingestion • Preprocessing • Indexing • Vectorization • Model Integration • Permissions • APIs • MCP

Expand over time – the platform stays the same

Which application for which use case?


Application	Core task	Benefit to employees
iFinder	Enterprise Search – finds any file, document, or email across your systems in seconds, with permissions enforced.	Where is the information?
iAssistant	RAG-based GenAI chatbot – answers questions using your own data, with source citations down to the exact passage.	What does it say, and what does it mean for me?
iHub	Micro-apps, assistants, and agents – turn knowledge into finished outputs and automate recurring processes.	How do I get the task done now?
Your own assistants & apps	Via MCP and documented APIs, your own assistants, business applications, and data science teams (e.g., Jupyter Notebook) access prepared, permission-aware index data directly.	What do we build next?

Our guiding principle: Your resources should be focused on your domain expertise and business processes – not on repeatedly rebuilding the same AI infrastructure foundations. Data connectivity, ingestion, preprocessing, indexing, vectorization, model integration, permissions, APIs, and MCP are already built in.

This foundation is not new – what is new is that we are presenting it as a platform. For 25 years, we have continuously expanded our data integration layer. iFinder Enterprise Search has been built on it from the very beginning; the platform has always been part of the product, just never its name. What now powers chatbots and agents has matured over 25 years of practical use in business and public administration.

Trusted by Design: sovereign, permission-aware, and secure

Answers are generated exclusively from authorized sources within your organization – with citations down to the exact passage. Only passages the person asking the question is authorized to access are included in the language model's context: two employees, the same question – two answers, each precisely within that person's permissions. This **grounded generation significantly reduces the risk of hallucinations and makes every statement verifiable and auditable** – for business units, internal audit, data protection, and information security.




+ Trusted by design: sovereign • permission-aware • source-backed • auditable

- 1 Data Layer** – content, metadata, and permissions in one controlled layer.
- 2 Permission-aware Passages** – permissions enforced down to passage level.
- 3 Hybrid Retrieval** – combines full-text search, linguistic analysis, semantic search, and vector search.
- 4 Reranking** – the most relevant passages first.
- 5 Grounded Generation** – answers based only on authorized passages, with citations.
- 6 Eval & Audit** – measured, logged, and traceable.

Enterprise AI Platform capabilities

- + **Data connectivity with permission inheritance:** 80+ connectors and hundreds of file formats – access permissions, including inherited permissions, are captured during indexing.
- + **Prepared for search and GenAI:** linguistic processing, auto-tagging, metadata enrichment, passage chunking, embeddings, and vectorization – prepared once and usable across all applications.
- + **Passage-level permission checks:** in results, filters, autocomplete, previews, LLM context, and agent steps – with no intermediate state in which unauthorized content exists.
- + **Open via MCP and API:** the platform serves as the permission-aware retrieval layer for your own assistants and agents – via the Model Context Protocol (MCP) and documented APIs, without losing permissions or source citations.
- + **Model choice without lock-in:** local open-source models, commercial cloud models, or both – use the right model for each use case, with guardrails, logging, and centralized usage and token controls.
- + **Measurable quality:** search relevance, answer quality, and models are evaluated against customer-specific gold standards and continuously optimized in production.

Your Data. Your Infrastructure. Your Choice.



In your own data center, fully air-gapped, in a private or sovereign cloud, in a hybrid setup, as a managed service, or as SaaS – the same platform and the same feature set, container-based with Kubernetes, Helm, and OpenShift. You can switch deployment models later because all variants run on the same foundation.

Benefits of the Enterprise AI Platform

- + **Information at your fingertips:** find knowledge from documents, business applications, and repositories in seconds – instead of searching through folder structures.
- + **Reduced workload:** automate recurring inquiries and standard tasks with assistants and agents – so employees can focus on their core responsibilities.
- + **Consolidation:** consolidate AI initiatives on one platform instead of creating a patchwork of disconnected pilot projects – every new application inherits data connectivity, permissions, and quality assurance.
- + **Auditability:** permissions, source citations, and an audit trail – AI that stands up to internal audit and data protection requirements.
- + **Costs under control:** use the right model for each task and centrally manage usage and token consumption – instead of incurring uncontrolled AI costs.

Secure, sovereign, compliant	
ISO/IEC 27001	IntraFind is ISO/IEC 27001 certified.
EU AI Act	The Enterprise AI Platform complies with Regulation (EU) 2024/1689 (EU AI Act).
Data protection	GDPR-compliant operation.
No model training	Your data is not used to train AI models.
Data residency	You decide where your data is processed and stored.

How do you get started with the Enterprise AI Platform?

Start small and clearly scoped – for example with iFinder Enterprise Search – and expand step by step to chatbots, agents, and your own applications. We will be happy to advise you.

Your path to expansion in 20 minutes: Bring a specific use case – a search task that is currently cumbersome and time-consuming, a question a chatbot should answer from your company's own data, or a process an agent could take over. We will show you how it can be implemented and used on the platform.

[Book a free consultation](#)



IntraFind – A reliable partner for businesses and public-sector organizations

Proven in the market for 25 years, our software is developed and hosted in Germany and delivers the highest levels of security – compliant with the EU AI Act and GDPR – and is trusted by leading organizations across the private and public sectors.

INTRAFIND

IntraFind Software AG
Landsberger Straße 368
80687 München
Germany

+49 89 3090446-0
sales@intrafind.com
www.intrafind.com

IntraFind Inc.
120 Wall Street, 20th Floor
New York, NY 10005
USA

+1 212 584 9724
sales@intrafind.com
www.intrafind.com