Turn Every File Into Structured, Searchable Intelligence

Agentic AI Orchestration Platform At Scale

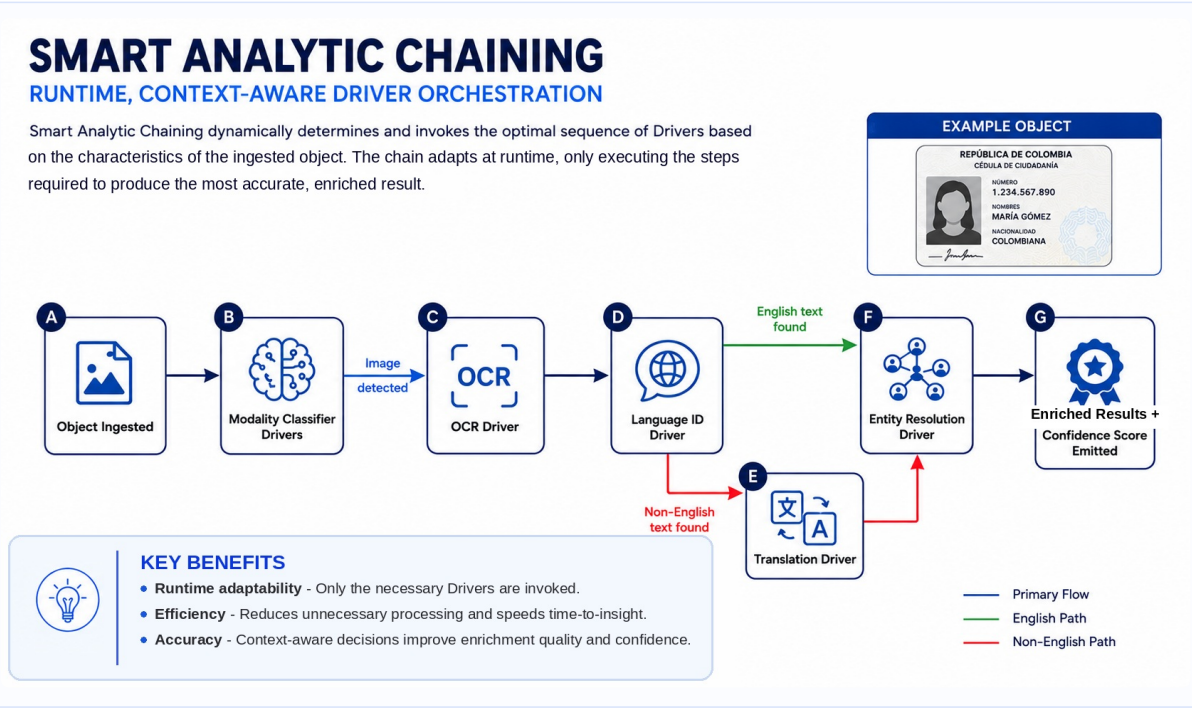
BeaconIQ® ingests files of any type, including PDFs, scans, images, audio, and more, and runs each one through a modular AI pipeline that extracts content, enriches it with meaning, and makes it instantly searchable. Every step is chained together dynamically at runtime, so each object only travels through the processing it actually needs. Built entirely on open-source technology and deployable fully air-gapped, BeaconIQ® enables an accelerated AI pipeline deployment and scales into a production-grade Kubernetes platform.

Modular
Every step is its own container. Swap a model or add a capability without touching anything else.

Scalable
Workers and compute scale independently, provisioning CPU or GPU only where the workload needs it.

Extensible
New file types, models, and analytic packs drop in as new steps, and no existing worker is ever modified.

Configurable
Prompts, models, and thresholds are tuned from a web interface and take effect on the next file.



A Prebuilt AI Orchestration Pipeline, From Raw File to Insight

Every file flows through the same five phases, and which steps actually run depends on what the file is and what is contained in the file.

Extract
Identifies true file type and validates every incoming object before any processing begins.

Parse
OCR, speech-to-text, and translation turn any format and language into clean, structured text.

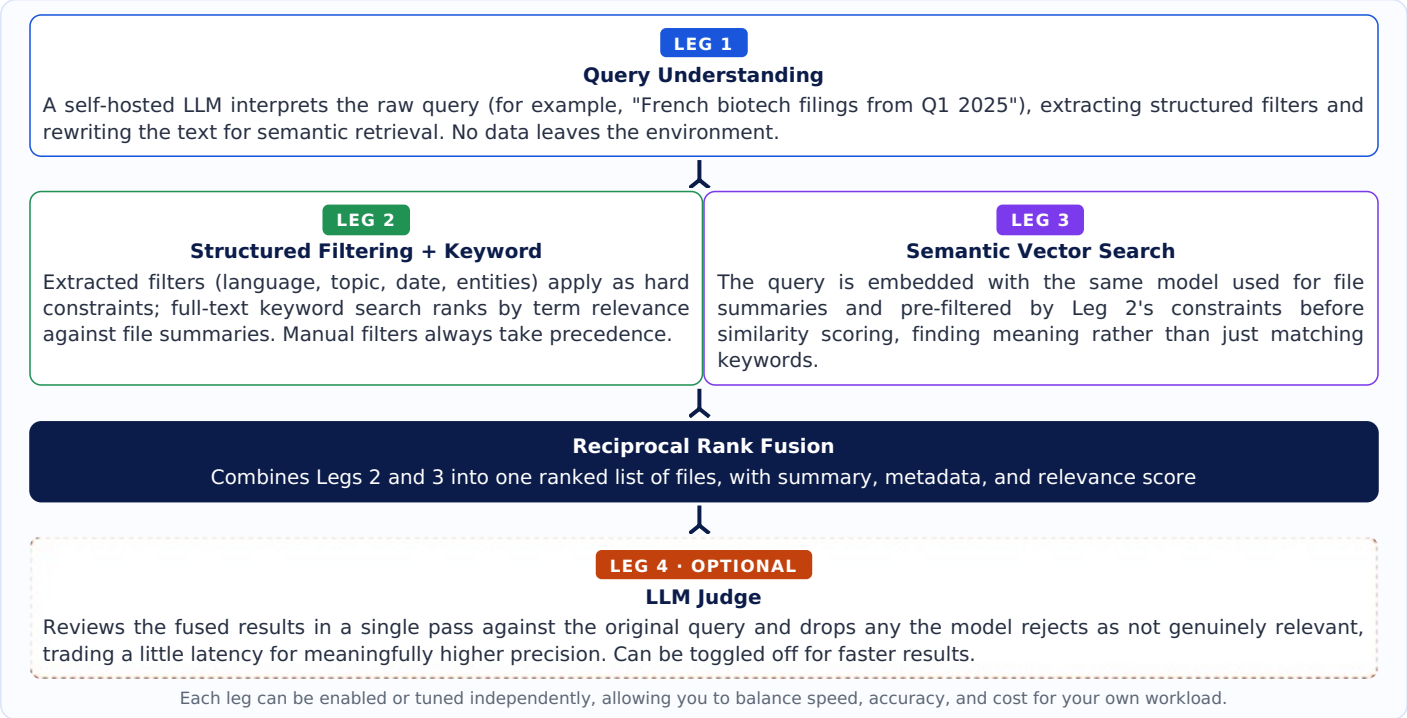
Chunk
Segments extracted text into coherent units sized for downstream embedding and retrieval.

Embed
Text, image, and audio embeddings power semantic search and similarity across every modality.

Enrich
Named entities, topics, objects, captions, and summaries extract the meaning inside every file.

Search That Understands Intent, Not Just Keywords

A natural-language query flows through four complementary legs, query understanding, structured filtering, semantic search, and an optional LLM judge pass that is fused into a single ranked result, entirely on self-hosted infrastructure.



Each leg can be enabled or tuned independently, allowing you to balance speed, accuracy, and cost for your own workload.

Built on a Foundation You Can Trust

- **Open-source only.** Every model and tool in the pipeline is open-source, with no proprietary black boxes and no per-call licensing.
- **Air-gap ready.** No third-party SaaS dependency anywhere in the pipeline or search layer. Runs fully disconnected if required.
- **Secrets never touch config.** Credentials are injected at container startup and fully IAM-audited.
- **Full audit trail.** Every file records exactly which configuration version processed it, with one-click rollback.

Designed to Grow

- **One container per model.** Replacing or upgrading a model is a single image swap, and nothing else in the system is touched.
- **Configuration without code.** Prompts, thresholds, and model parameters are tuned live from the web interface.
- **Right-sized compute.** Lightweight steps run on CPU; heavy vision and language models get GPU only when needed.
- **Highly adaptable.** Designed to easily on-board new models without requiring any code.

Let BeaconIQ® light the way to what matters most.

See BeaconIQ® Run on Your Own Files

Request a walkthrough of BeaconIQ® and chart your migration path to production.

CONTACT US
ai@iconsultcorp.com