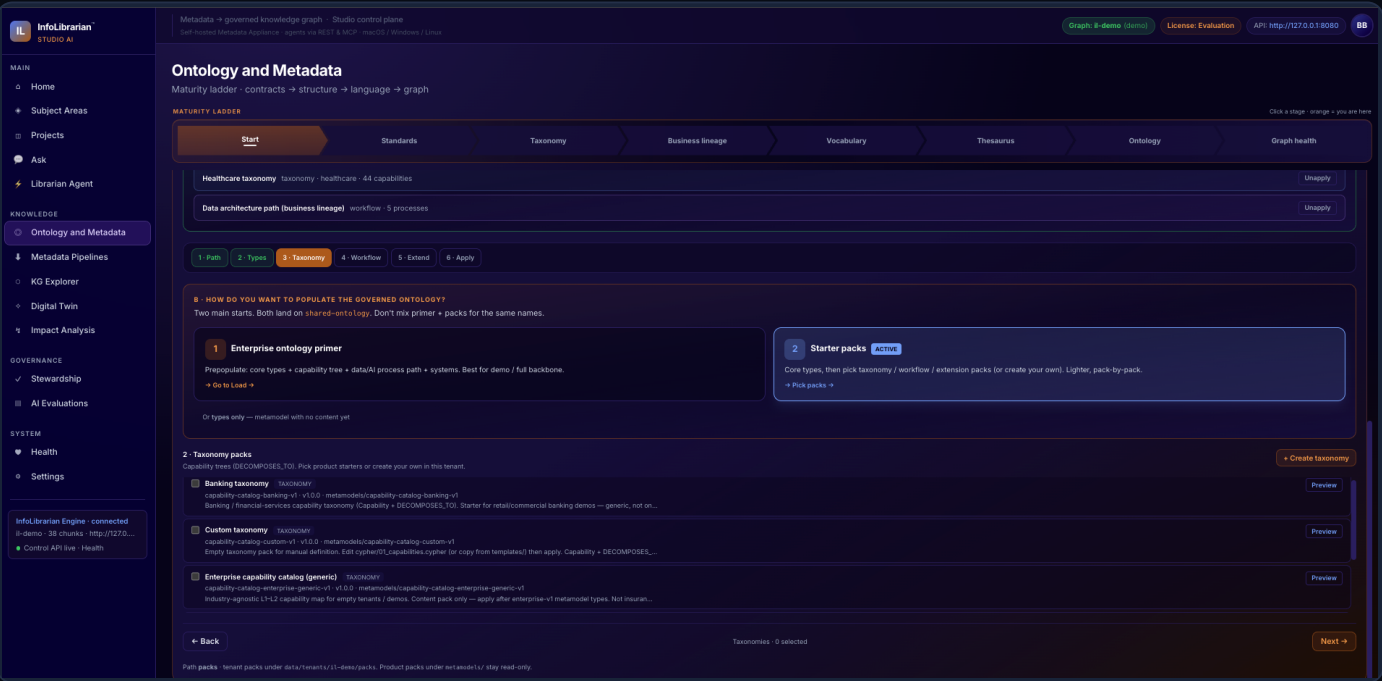


KNOWLEDGE GRAPH ENGINE · PRODUCT BRIEF FOR THE C-SUITE
 LICENSE ONLY · RUNS IN YOUR ENVIRONMENT

Enterprise metadata into a knowledge graph.

Our enterprise metadata heritage became an AI-native knowledge graph platform. One subject area at a time. Low human-in-the-loop. Studio as the control plane. Agents consume the same governed graph your people already trust.

CFO · CIO · CTO · CDO · EA · Risk



Studio · Ontology and Metadata · start here · sneak peek, not the whole store

WEEKS

not years

to a governed graph on one subject area

LOW HITL

stewardship

AI proposes. Experts approve exceptions.

100%

air-gapped

or BYO frontier models. You choose.

ONE

control plane

Studio + agents + chat on the same graph.

The AI gap is not a model gap.

Flagship models already reason fluently. What they do not have is how your enterprise is defined and how it runs: capabilities, processes, systems, terms, lineage, and the documents that surround them. Without that map, every copilot is a lottery.



CFO

Spend without P&L

Token firehoses and multi-year catalog programs do not show up as operating change. You need weeks-to-value and scored answers, not another platform tax.



CIO / CTO

Tech debt + wrappers

Agent platforms bolted onto ungoverned estates recreate the last decade of integration debt — with a larger blast radius.



CDO / DATA

Heavy stewardship

Classic catalogs onboard Data Citizens through years of manual curation. Exceptions pile up. Time-to-trust never arrives.



EA / PROCESS

The real context is stranded

Capability models, process definitions, and business metamodels live in decks and modeling tools. Agents never see them.




RISK / AUDIT

No explainability

Similarity search cannot say who owns it, which version is in force, or why the answer is allowed.

A knowledge graph is a contract, not an index.

It encodes the objects the business already argues about, the relationships RAG cannot infer reliably, and the trail back to evidence. Documents are footnotes. The graph is the map.



Not a search index — types, relationships, evidence, and a signature.

ENTITIES

Capability, process, system, term, contract, product — one name, many source keys.

YEARS OF MODELING

Enterprise ontologies die in committee. A corporate ontology for everything is a stall.

RELATIONSHIPS

Decomposes-to, flows-to, maps-to, depends-on, supports. Multi-hop paths, not word overlap.

UNCONSTRAINED EXTRACT

Let the model invent types and you get label pollution, not a platform.

EVIDENCE

Documents, models, and glossaries attached to those things — cited, not stuffed into a window.

HEAVY HITL CATALOGS

Data Citizens, multi-year onboarding, and stewardship as a standing army.

GOVERNANCE

A bounded ontology. Types and rules. Humans approve what AI proposes.

WRONG LAYER

A lakehouse semantic layer names metrics. It does not know how the business actually runs.

One package. Graph, Studio, agents.

No graph-vendor shopping list. The platform ships the store, the engine, the control plane, and the agent surface. Bring your own models — on-prem, air-gapped, or approved frontier endpoints.



05 CONSUMPTION · SAME GRAPH

Studio AI (people) · MCP / APIs (any agent or chatbot)

Browse, contribute, steward, ontologist — one GUI. Not a second portal product.

One subject area. Then the next.

A corporate ontology for the entire firm is how programs die. InfoLibrarian compounds quality by scope: define the area, apply a starter ontology, load models and documents, steward the exceptions, ask questions, expose agents.



Quality compounds by ring — not a big-bang corporate ontology.

1

Define the subject area

Claims. Member. Payments. Media rights. One bounded corpus of high-value sources.

2

Apply a governed ontology

Enterprise spine plus a vertical starter — Insurance, Healthcare, Travel, Banking, Media — then customize.

3

Ingest the real context

EA capabilities and processes, modeling-tool exports, glossaries generated or uploaded, unstructured decks and policies.

4

AI gap-close + HITL

Missing links are proposed and scored. Stewards review exceptions — not every row.

5

Ask, measure, expose

Hybrid retrieval with citations. Deterministic mode when audit requires it. MCP to any agent.

TIME TO MARKET

Weeks to a working graph on one area. Months to a compound estate. Not a multi-year catalog transformation.

Traditional metadata is a standing army.

Data Citizens, glossary wars, and multi-year onboarding are the old operating model. InfoLibrarian is AI-native: the machine does the volume. People do the exceptions. That is stewardship — not a second job for the enterprise.



THEN — data-governance monolith, every row. NOW — one steward, low-confidence exceptions only.

THEN

Heavy stewardship

- Manual harvest of every term
- Years to onboard a domain
- Catalog as a project, not a product
- Agents never get a trusted map

NOW

Low HITL, AI-powered

- AI extract, map, and gap-close
- Scored exception queues
- Steward accepts / rejects
- Same graph for people and agents



TECH DEBT

Unstructured + manual = more debt.

Every year of paper catalogs, invented labels, and agent wrappers on an ungoverned estate is interest. A bounded ontology with scored exceptions stops the bleed. You do not hire your way out of an index.

STEWARDSHIP

Automation proposes. Your experts certify.

Gap-close finds relationships that are true but never written down. Each proposal is scored for fast review. Nothing skips a human on the authority edge. That is how you get speed without gambling the catalog. Roles — browser, contributor, steward, ontologist, admin — apply at catalog level. Portal and Studio are one GUI.

Do not start from a blank schema.

Vertical starters you customize — not a one-size corporate ontology, and not a pile of invented labels. The real context is already in EA: capabilities, processes, metamodels. Load them. Then add documents and structured metadata.

ONTOLOGY

Insurance

+ Insurance

- Policy
- Member
- Claim
 - Payment / producer

ONTOLOGY

Healthcare

+ Healthcare

- Member
- Encounter
- Care pathway
 - Claim

ONTOLOGY

Travel

+ Travel

- Inventory
- Reservation
- Itinerary
 - Fulfillment

ONTOLOGY

Banking

+ Banking

- Lending
- Payments
- Risk
 - Counterparty

TAXONOMY

Media

+ Media

- Brand
 - Franchise
 - Title
- Rights

STARTER

Your vertical

+ Your L1

- Define L2
- Define L2
 - Then load EA
 - Then documents

THE REAL CONTEXT · BUSINESS LINEAGE

Enterprise architecture belongs on the graph.

Business capability models (taxonomy). Business process definitions (lineage of work). Modeling-tool exports. Glossaries generated by harvest or uploaded. Pipelines in Studio map structured metadata in days — not a custom integration program.



If EA stays in PowerPoint while you embed tickets, you rebuilt search. Architecture, data metadata, and documents — one catalog.

Ask. Cite. Score. Then trust it.

Studio is the product surface: hybrid Ask with citations, RAGAS-style judge on every answer, deterministic mode when the auditor is in the room. Live product. Crystal clear.

InfoLibrarian™
STUDIO AI

MAIN

- Home
- Subject Areas
- Projects
- Ask
- Librarian Agent

KNOWLEDGE

- Ontology and Metadate
- Metadate Pipelines
- KG Explorer
- Digital Twin
- Impact Analysis

GOVERNANCE

- Stewardship
- AI Evaluations

SYSTEM

- Health
- Settings

InfoLibrarian Engine · connected

- 8-demo · 38 chunks · http://127.0.0.1:8080
- Control API live · Health

Metadata → governed knowledge graph · Studio control plane

- Self-hosted Metadata Appliances · agents via REST & MCP · macOS / Windows / Linux

Graph: 8-demo (demo) · License: Evaluation · API: http://127.0.0.1:8080 · 88

Ask / Retrieve

Show Ask options · Mode: subject-area · GPU

Conversation

- Hybrid
- Expand

A managed-care PCP check failed. Ops wants to blame member demographics or the ASO capitation print flag. What affiliation field should they inspect first, and what does that field actually mean?

1. Inspect the "ACCEPTING PATIENTS CODE" field first:

- The context suggests that when PCP validation fails, it's important to check the "ACCEPTING PATIENTS CODE" on the affiliation before assuming a demographics mismatch. This is because the accepting patients code indicates if a provider is accepting new patients and is used in PCP validation for managed care.

2. The "ACCEPTING PATIENTS CODE" field means:

- The "ACCEPTING PATIENTS CODE" is a business term that indicates whether a provider is currently accepting new patients. It's crucial for PCP validation in managed care processing, as it helps determine if the issue might be related to patient acceptance policies rather than demographic mismatches.

Citations: 5 — open tab

70803ms total · 8 stages — Timings

Judge avg 1.00 — open tab

Walk the loop — authorization batch date from claim-line processing, to the Erwin logical attribute, to the VendorOut column. What are the two confirmations an analyst must make, and which MDA close-loop in the pharmacy guide is the same pattern for the approved drug deductible?

1. First Confirmation: The analyst must confirm that the AUTHORIZATION BATCH DATE from claim-line processing maps to the logical attribute in Erwin.

2. Second Confirmation: The analyst must verify that this logical attribute then maps to the VendorOut column.

The MDA close-loop for the approved drug deductible is the same pattern as described above, which involves:

- Business term (APPROVED DRUG DEDUCTIBLE AMOUNT)
- Erwin attribute
- Physical column via MAPS_TO

Citations: 5 — open tab

84670ms total · 8 stages — Timings

Judge avg 0.90 — open tab

Ask a question about the graph...

Send

Diagnostics for selected answer

BUILD: Timings · Stats · Trace · Debugging

EVIDENCE: Citations · Relevant docs · Chunks · Context · Bridge graph · Graph role · Topics · Graph query · RAGAS / Judge

LLM-as-judge (RAGAS-style) scores for this answer — 0-1 scale. A small judge model often saturates near 1.0 on fluent extractive answers; treat a flat 1.00 as "didn't penalize," not perfect retrieval. Not stored in the graph.
Retrieval pressure: 12 chunks · 5 bridge · 9 entities · top score 0.982. Graph support is the bridge/entity counts, not the RAGAS average.

OVERALL AVERAGE 0.90
Mean of the four metrics below — strong

FAITHFULNESS 0.75
Grounded
Does the answer stick to the retrieved context, or invent claims not supported by sources?
"The answer aligns with the retrieved context and does not introduce unsupported facts."

ANSWER RELEVANCY 1.00
On-topic
How well the answer addresses the question asked (vs. drifting off-topic).
"The answer directly addresses all parts of the question without adding off-topic content."

CONTEXT PRECISION 0.85
None
How much of the retrieved context was useful for the answer (less noise).
"Most of the retrieved context was relevant to answering the question, with only a few chunks being slightly noisy."

CONTEXT RECALL 1.00
Coverage
Whether enough supporting context was retrieved for a complete answer.
"The answer covers all key points needed for a practitioner to act on this question."

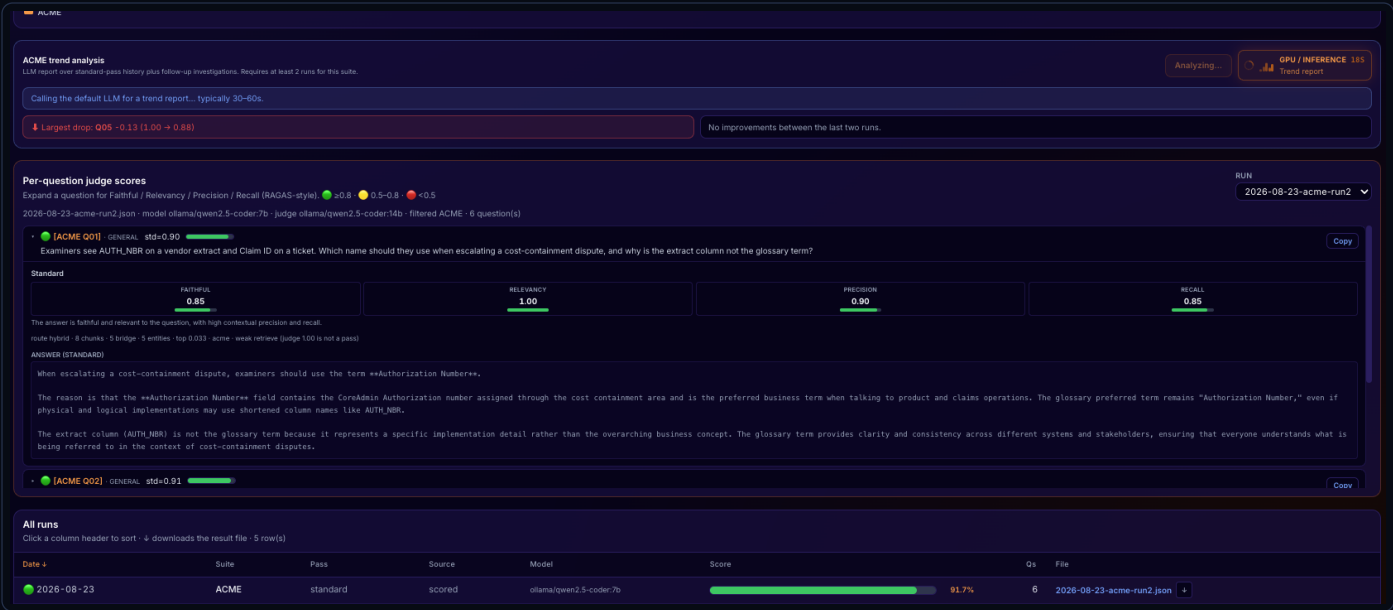
OVERALL RATIONALE
The answer is comprehensive and specific, aligning well with the retrieved context while fully addressing the question.

InfoLibrarian™ · infolibcorp.com/platform · License only · Confidential

8 / 13

Trust is a number. Trend is a story.

Gold questions. Judged runs. Faithfulness, relevancy, precision, recall — per question and across the suite. This is the board pack: not a vibe, a regression you can show.



Studio · AI Evaluations · ACME suite, per-question RAGAS-style judge · live product

AI-GENERATED TREND ANALYSIS

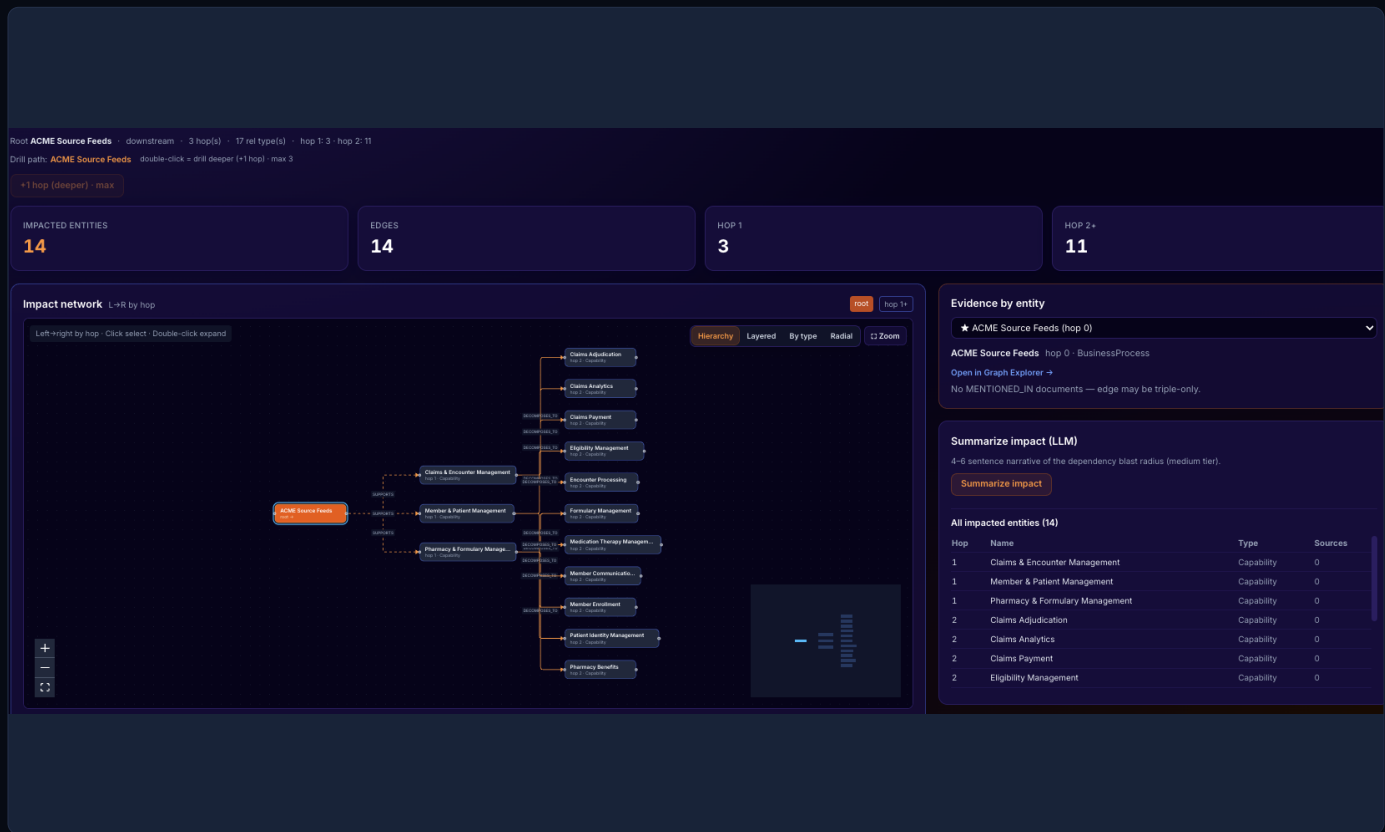
The suite writes the trend report — largest drop, no-improvement, next investigation. GPU/inference, not a consultant appendix.

LOW HITL · QUESTIONS + GROUND TRUTH

Gold questions and ground truth are generated with humans on the exceptions — not a standing labeling army. Same low-HITL rule as stewardship.

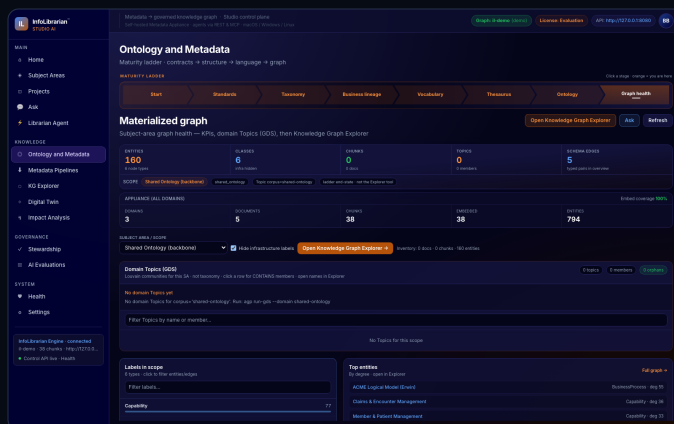
The control plane, in the product.

Impact as a path. The ontology as a contract. Graph health you can ladder. Neighborhoods, not a blob. Same GUI for browser, contributor, steward, ontologist.

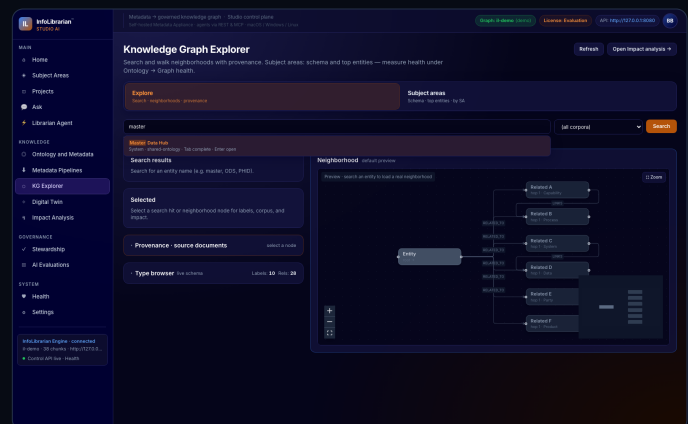


Impact · ACME Source Feeds · hops, not a report

Ontology contract · type-level schema, layered



Graph health · ontology ladder, materialized



KG Explorer · neighborhoods, not a blob

People and agents on the same graph.

MCP and open APIs mean any chatbot or agent framework consumes governed context — not a private prompt pile. Studio is how operators see what the agent sees.

MCP → ANY AGENT

The librarian surface is the same catalog your analysts Ask.
No second truth for machines.

100% AIR-GAPPED

Metadata, content, graph, and inference stay inside the
boundary when policy requires it.

BRING YOUR OWN MODELS

On-prem local models, private endpoints, or approved frontier
APIs. The platform does not own your meter.

DETERMINISTIC MODE

When the answer must be auditable: graph and structured
metadata, not a creative completion.

IMPACT ANALYSIS


What breaks downstream if this changes — capabilities,
processes, systems, terms.

BUNDLED TRIAL


Studio + engine as a package. Point at your environment. Not
a GitHub science project.

Every answer is a number and a path.

Provenance from the first parse. Citations on every Ask. Built-in evaluation suites and a dashboard — not a slide that says “we are responsible.”



Example from the live Ask surface (overall 0.90) — the board pack is this product, not a consulting appendix.



Hybrid retrieval you can tune — from fully deterministic (audit) to graph-augmented synthesis.

Gap-close proposals scored so stewards spend minutes, not weeks. Lineage and impact as first-class product.

Catalog-level roles: browser, contributor, steward, ontologist, administrator. AI evaluations as a board surface.

THE TEST THAT MATTERS

Can two agents, a chatbot, and a human analyst give the same answer to the same as-of question, with the same citations? If not, you have prompts. If yes, you have a platform.

Weeks. Not years.


You already have enterprise metadata, architecture, and documents. InfoLibrarian turns them into a knowledge graph people and agents can use — without a standing catalog army and without sending the estate to someone else’s model host.



Time to solution is the CIO metric. A catalog program is a calendar. A subject-area graph is a decision.

TIME TO A GOVERNED GRAPH · ONE SUBJECT AREA

Same destination. Different clocks.



- 1 Pick the area

Wrong answer has a cost
- 2 Apply the spine

Vertical starter + EA
- 3 Load the context

Models, glossary, docs
- 4 Steward exceptions

Minutes. Not every row.
- 5 Ask · score · MCP

Then clone the pattern

License only. Runs in your environment.

Platform and Knowledge Graph Engine are InfoLibrarian commercial IP. Classic product line is heritage; the next-generation platform is AI-native and shipping. Pilot, production, or enterprise — same software, you size the room.