

SEMANTIQ · A XAQUA PRODUCT

# SemantiQ.

AI Data Catalog

Federated Semantic Layer

Governed by Design

## Auto-discovered. AI-enriched. **Steward-confirmed** Data Catalog — for the AI Era.

The catalog AI drafts and your stewards confirm. SemantiQ connects to every data source you have — across catalogs, sources, files, and virtual assets — discovers the schema, enriches it with business meaning, and becomes the common language behind every plain-English question. **Your data never moves.** The output is one federated semantic layer — read by every xAQUA agent and every developer integration.

### YOUR ENTERPRISE CATALOGS

  
**Customer 360**  
412 assets

  
**HR Master**  
287 assets

  
**Finance & GL**  
523 assets

  
**Operations**  
386 assets

  
**Compliance**  
239 assets

↓ FEDERATED ↓

### SemantiQ Federated Search Layer


BM25 keyword + semantic vector hybrid · One query, all catalogs · Zero data movement


+ BM25

+ SEMANTIC

= HYBRID

↓ TWO ACCESS MODES ↓

  
**Ask Zyra**  
Conversational · Built for users

  
**Hybrid Search API**  
Programmatic · Built for developers

WHY TRADITIONAL CATALOGS DIE

# Buy a **traditional catalog**. Hire a team to populate it. **Watch it go stale.**

Traditional catalogs are **documentation-first**. Every business name, every description, every PII tag, every glossary entry — typed by hand, by your team, on top of their day jobs. Six months in, half is wrong. A year in, half doesn't exist. The data team writes definitions; the business team disagrees; the catalog you bought to standardize answers becomes the source of three different ones.

**27%**

Average coverage of legacy catalogs after 11 months — the documentation backlog never shrinks.

**8–13**

Tools in a typical enterprise catalog stack (catalog, glossary, lineage, quality, MDM) — the “Frankenstack.”

**40%**

Of data budgets consumed by integration and licensing across that fragmented stack.

THE INVERSION

# SemantIQ inverts the model. **Curation scales. Transcription doesn't.**

That's not a team failure. It's a catalog failure. Documentation-first catalogs scale with human effort — and human effort doesn't scale. SemantIQ flips it: AI does the writing; your stewards do the reviewing. The catalog stays current because curation scales while transcription doesn't.



### AI drafts. Stewards confirm.

Every business name, description, and classification drafted by AI from schema metadata — never raw data. **Confidence-scored. Lockable. Auditable.**



### One language. Every catalog.

Federated semantic layer across HR, Sales, Finance, Operations, Compliance. **BM25 + semantic vector hybrid** — one query, all catalogs.



### Strict governance, by design.

No direct edits. Every change flows through a change request, reviewed by the catalog owner or steward, **accepted or rejected with reasoning logged.**

WHAT SEMANTIQ DOES

# Catalog. Glossary. Lineage. Quality. Governance. All native.

One product. **Five pillars. Eleven capabilities.** A single shared semantic model — every module reads and writes through the same governed layer.

<p><b>P 01</b></p> <h3>Connect &amp; Discover</h3> <ul style="list-style-type: none"><li>Bound Sources</li><li>File Assets</li><li>Schema Sync</li></ul>	<p><b>P 02</b></p> <h3>Enrich &amp; Classify</h3> <ul style="list-style-type: none"><li>AI Enrichment</li><li>Quality &amp; Metrics</li><li>Business Glossary</li></ul>	<p><b>P 03</b></p> <h3>Relate &amp; Reason</h3> <ul style="list-style-type: none"><li>Cross-Source Joins</li><li>Analytical Semantics</li></ul>	<p><b>P 04</b></p> <h3>Govern &amp; Control</h3> <ul style="list-style-type: none"><li>Certification</li><li>Change Mgmt</li></ul>	<p><b>P 05</b></p> <h3>Conversational Intelligence</h3> <ul style="list-style-type: none"><li>AI Insights · Zyra</li></ul>
--	---	---	--	--

HOW IT WORKS

## Six steps. Mostly automatic. Always under steward control.

<p><b>1</b></p> <h3>Connect</h3> <p>Database, file, or SaaS source. No data movement.</p>	<p><b>2</b></p> <h3>Sync</h3> <p>Schema introspection. Every change logged.</p>	<p><b>3</b></p> <h3>Register</h3> <p>Permanent UUIDs. Renames tracked. Soft deletes.</p>	<p><b>4</b></p> <h3>Enrich</h3> <p>AI drafts. Stewards review via change request.</p>	<p><b>5</b></p> <h3>Relate</h3> <p>Cross-source affinity scan. One-click confirm.</p>	<p><b>6</b></p> <h3>Power</h3> <p>Every xAQUA agent reads the catalog.</p>
---	---	--	---	---	--

**One catalog. One language. Every product.**

\* ARCHITECTURAL NOTE

**Zero raw data ever leaves your environment.** The LLM sees schema metadata only — never customer rows. Enforced at the service layer. The same model can run inside your private deployment if you need it to.

HOW YOU ACCESS THE CATALOG

# Two access modes. One semantic layer.

The same federated semantic layer that powers the conversational experience also powers your applications. Search across every catalog you have access to — through Zyra in plain English, or through the Hybrid Search API.

Z

## Ask Zyra

Conversational · Built for users

Plain English. Catalog-grounded. Cross-source and cross-catalog by default. Ask *"find all tables with PII data across the enterprise"* — Zyra returns one answer that spans HR, Sales, Finance, Operations, Compliance. With root-cause, lineage, and a suggested next action.

Plain English

Cross-catalog

Root-cause analysis

Self-improving

{ }

## Hybrid Search API

Programmatic · Built for developers

BM25 + semantic vectors. One endpoint. Returns ranked matches across every catalog you have access to — with metadata, sensitivity, and confidence scores. Built for app developers and data engineers integrating SemantIQ into their own products and pipelines.

BM25 + Vector

RRF Fusion

Federated scope

JSON / REST

CUSTOMER REFERENCE · GOVERNMENT



## A \$300B+ U.S. public pension fund uses SemantIQ for benefit eligibility, compliance reporting, and data cataloging in a fully private deployment.

Information retrieval that took the team **days** — searching across siloed systems, preparing extracts, and manually reconciling files — now takes **seconds**. SemantIQ runs entirely inside the customer's environment. No data leaves. The federated semantic layer is the foundation; every other xAQUA agent reads from it.

SEE IT ON YOUR OWN DATA

## Get a 30-minute SemantIQ demo.

We'll connect a real source from your stack — show you the AI enrichment, the federated search, and the strict change-management workflow — and outline a private deployment path if it's a fit.

WEB

[xaqua.io](https://xaqua.io)

EMAIL

[sales@xaqua.io](mailto:sales@xaqua.io)

DEMO

[xaqua.io/demo](https://xaqua.io/demo)