

REEVE · A XAQUA PRODUCT

Reeve.

Data-as-a-Product


Data Mesh ready

REST · GraphQL on Day 1

Your data, as a product. Find it. Trust it. Use it.

Treat every dataset as a published product. Reeve is the enterprise Data Product Catalog — every dataset has an owner, a contract, a Trust Score, and an API. **Search. Subscribe. Consume.** Your trapped silo data becomes a living portfolio your business can find, trust, and reuse — without filing a ticket.

search "active customers" | REEVE · CATALOG



Customer 360

✓ CERTIFIED ● LIVE REST · GRAPHQL

| | | | |
|-------------------------|--------------------------|------------------------|---------------------|
| SUBSCRIBERS 42 teams | TRUST SCORE 98.4 · A+ | FRESHNESS 2 min ago | CALLS / DAY 1.2M |
|-------------------------|--------------------------|------------------------|---------------------|

WHY MOST CATALOGS FAIL

Most data dies in a silo nobody opens.

Decades of investment built warehouses, lakes, and lakehouses overflowing with data — yet most of it sits unused. Analysts can't find it. Consumers don't trust it. Owners don't get credit. Three teams build "Customer 360" three different ways. Every team rebuilds the same dataset because nobody knows the canonical version exists.

8–13

Disconnected tools across catalog, lineage, quality, and sharing — the integration tax of the "Frankenstack."

80%

Of business users locked out of meaningful data access by SQL and tooling barriers — even when the data exists.

3x

The same canonical dataset, built three different ways, by three different teams — three different answers.

THE PROMISE

One catalog. Every dataset a product.

Reeve fixes that by treating every dataset as a published product — with a name, an owner, a contract, a Trust Score, and an API. It shows up in search. It tells you who else uses it. It tells you when it last refreshed. **It earns trust the way a product does** — by being findable, reliable, and well-supported.



Publish anything

Promote any table, virtual dataset, recipe, query, blend, report, model, or AI bot to a first-class data product — owner, contract, SLA set in minutes.



Find anything

Faceted, semantic search across every domain. Filter by owner, tag, freshness, certification. Preview schema, see who's using it — in one screen.



Consume anywhere

Every product publishes as REST and GraphQL on Day 1 — out of the box. Apps, partners, and downstream pipelines plug in directly. No integration project.

THE AI DATA TEAM BEHIND REEVE

Two agents. One catalog.

Reeve is where two members of the xAQUA AI Data Team make published data products possible — one curates, one consumes. Both work against the same Semantic Layer, both leave behind a versioned, governed product.



AI Data Steward

For owners, stewards, domain leads

Promotes datasets to products. Assigns owners, contracts, SLAs, and tags. Composes Trust Score from quality, freshness, lineage, and usage. **Publishes the catalog the org actually opens.**

Promote

Contracts

Trust Score

Lineage



AI Data Analyst

For analysts, business users, partners

Searches by name, tag, owner, or business term in plain English. Subscribes with one click. Consumes via REST, GraphQL, or live SQL — with row, column, and field-level access enforced. **From search box to live API call.**

Search

Subscribe

REST · GraphQL

Governed access

EIGHT PRODUCT TYPES · ONE PORTFOLIO

Every asset your business needs to ship.



Datasets



Queries



Data Blends



Reports



Prep Recipes



Predictive Models



Data AI Bots



DaaS APIs



FROM A CATALOG TO A PORTFOLIO

Two outputs. One source of truth.

Every Reeve publication ends in something the rest of the organization can find, trust, and consume — through humans or machines, in interfaces or in pipelines.

OUTPUT 1 · TRUST-SCORED PRODUCT

Live Data Product

A first-class catalog entry: name, owner, contract, SLA, tags, and a composite Trust Score from quality, freshness, lineage, and usage. Searchable by every team. Subscribed in one click. Refreshes live as sources update — no stale snapshots.

OUTPUT 2 · API ON DAY ONE

REST & GraphQL DaaS

Every product publishes as a governed REST and GraphQL endpoint — automatically. Internal apps, partners, and embedded analytics consume fresh, masked, policy-aware data without ever touching raw records. OpenAPI schema included.

CUSTOMER REFERENCE · GOVERNMENT



A \$300B+ U.S. public pension fund uses Reeve to catalog benefit-eligibility, compliance, and member datasets — in a fully private deployment.

Information retrieval that took the team **days** — searching across siloed systems, requesting extracts, manually reconciling versions — now takes **seconds**. Every governed dataset is published as a Trust-scored product. Domain owners maintain their own portfolios. Auditors see lineage end-to-end. Reeve runs entirely inside the customer's environment. **No data leaves the boundary.**

SEE IT ON YOUR OWN DATA

Get a 30-minute Reeve demo.

We'll publish one of your datasets as a Trust-scored product, expose it as a live REST and GraphQL API, and federate it across your domains — in under fifteen minutes. Private deployment path included if it's a fit.

WEB

xaqua.io

EMAIL

sales@xaqua.io

DEMO

xaqua.io/demo