AUTOMATIZE · BUSINESS INTELLIGENCE

Qrly.

Ask. Build. Share. Self-service business intelligence, on your own infrastructure.

Connect a database. Ask a question. Build a report, a chart, a dashboard. That is the whole system — until you need more, at which point Qrly also carries a full spreadsheet engine, four cache and replication tiers, an approval-gated AI agent workforce, and a productised BCBS 239 risk-data governance module that no other BI platform ships.

~40 data sources | Self-hosted | White-label | BCBS 239 | Agent-ready

40
CONNECTION TYPES

13
VISUALISATION TYPES

500
WORKSHEET FUNCTIONS

14
PRINCIPLES GOVERNED

BANKING · INSURANCE · REGULATED ENTERPRISE · EMBEDDED ANALYTICS

PRODUCT BROCHURE 2026.1

THE PROBLEM

Self-service analytics stops being self-service the moment it matters.

The tool is bought so business users can answer their own questions. Six months later every meaningful number still routes through two data engineers, the important reporting has migrated back into spreadsheets nobody can trace, and the AI features that were meant to help are the ones the risk function will not approve — because nobody can say what the model was told, or what it changed.

Engineer bottleneck
Every new question is a ticket. The tool's promise dies in the queue.

Shadow spreadsheets
The real reporting lives in files with hand-typed columns beside a stale export.

Unauditable AI
A model that writes queries and saves content, with no record of the prompt or the change.

Governance as a project
Data lineage, critical data elements and quality rules assembled by hand in documents.

THE SOLUTION

A BI platform that takes governance as seriously as it takes charts.

Query without SQL
A point-and-click definition compiled into dialect-correct SQL for around forty database types — with drill-down, calculated fields, and a converter that turns hand-written SQL back into the visual builder.

Spreadsheet, properly
Not an export button — a real grid with around five hundred worksheet functions, four conditional-formatting families, and refresh-safe blocks that keep your hand-written columns beside their record when the query reorders.

Governed by construction
Every query, every AI prompt, every agent tool call and every schema change is written to an audit table. Schema authoring is off by default at three independent levels.

THE DIFFERENTIATOR

BCBS 239, in the product.

All fourteen Basel Committee principles catalogued in code, with periodic self-assessments and a sign-off chain, a critical data element register, data quality rules scored across six dimensions, element-to-element lineage, a risk report calendar with attestation, a remediation issue register — and a printable dossier for the regulator.

AI assistance in this module writes to **separate audit codes from human actions**, so a reviewer can always tell what was machine-drafted.

HOW A QUESTION BECOMES AN ANSWER

1. **Connect**
   Point Qrly at a database. It reads the schema, caches it, and refreshes nightly.
2. **Ask**
   Build visually, write SQL, or describe what you want in your own words and get a query back you can read and edit.
3. **Shape**
   Thirteen visualisation types, a semantic model layer with rollups and period comparison, and drill-down to underlying rows.
4. **Share**
   Dashboards with cascading filters, scheduled deliveries, alerts, public links, embedded panels, or an embedded assistant.

CONNECT, QUERY, PRESENT

The analytics surface.

1 Data connectivity

- Around forty connection types — PostgreSQL, MySQL, SQL Server, Snowflake, BigQuery, Redshift, Databricks, Trino, ClickHouse, DuckDB, Synapse, Vertica, DB2, Sybase and more
- Managed-cloud variants for the major providers recognised natively
- Per-connection pooling, query timeout, row limits, a concurrency governor with a wait queue, and telemetry on both
- Nightly schema introspection with cached columns and foreign keys; a data dictionary and an entity-relationship view
- **Freshness probes** so the interface can say “data as of” without re-running the query
- SSH tunnels, TLS modes, and a guard against reaching internal hosts
- Every credential encrypted at rest

2 Query & visualise

- A visual builder covering tables, joins, filters, aggregations, grouping, sorting and limits — no SQL required
- Fourteen filter operators plus date helpers for week, month, year and rolling windows
- A schema-aware SQL editor with completion, hover, formatting and template variables
- **Calculated fields** usable anywhere a column is, sanitised against anything that is not a read
- SQL back into the visual builder, so a hand-written query is not a dead end
- **Drill-down** in two modes — to underlying rows, or narrowed to the selected group — on visual and native queries alike
- Thirteen visualisation types including pivot, funnel, gauge, progress and both regional and point maps
- A semantic model layer with dimensions, measures, subtotals, time hierarchies, period-over-period comparison, top-N within group and pivoting
- Interactive per-question filters and cascading filters that drive many cards with differing field names
- Streaming results, chunked exports, and output as JSON, CSV, newline-delimited JSON or Parquet
- **Verified** as a trust signal, flippable only by curators and logged every time

3 Dashboards & delivery

- Twelve-column dashboards with drag and resize, dashboard-level filters and auto-refresh
- **Live dashboards** pushed on an interval, or driven by a change notification from the source database itself
- A morning briefing that diffs yesterday against today into a banner
- Reports for side-by-side comparison or multi-card layout, with revocable public share links
- Report packages bundling many reports into a single workbook, on demand or on a schedule
- **Report bursting** — one page per value of a chosen filter
- Alerts on rows returned, goal reached or below goal, by e-mail, chat, webhook or SMS
- Scheduled deliveries as PDF or CSV, with optional written commentary

4 The spreadsheet

- A hand-written formula engine and virtualised grid — around **500 worksheet functions**, cross-sheet references, whole-column ranges, lazy evaluation and cycle detection
- Excel's two editing modes, click-and-drag reference picking with range outlining, fill-down with reference rewriting, and row and column reordering that relocates every reference
- Four ways to bind data: a saved question, a visual query, ad-hoc SQL, or a natural-language request resolved to SQL *once* — after which refreshes cost nothing
- **Refresh-safe blocks**: new source columns are appended, removed ones are cleared rather than shifted, your own rows and columns are never written into, and key columns return each record to its prior row
- Conditional formatting by cell rule, colour scale, data bar and icon set
- Export keeping formulas as formulas

> **Why the spreadsheet matters commercially.** The reporting that actually runs a finance or risk function is in spreadsheets, and it is there because analysts need a column of their own beside the data. Qrly does not ask them to give that up — it makes the data underneath it live, traceable and governed.

INTELLIGENCE, PERFORMANCE, GOVERNANCE

The three layers underneath.

5 Agentic BI, approval-gated

- **Four personas** — Analyst (read-only), Composer (drafts questions and dashboards), Modeler (builds semantic models), Investigator (root-cause work)
- A reasoning loop over ten tools, alternating thought, action and observation until an answer or a budget is reached
- **Every write is a proposal a human must approve.** Approval is required by default; the proposal appears as a card with accept and reject
- An audit table storing the **exact prompt, message and injected context** sent to the provider — users can audit what the model was told
- Row-level data access for AI is **off by default**: the model sees schema only unless a connection explicitly opts in
- Per-organisation token, session and step budgets; learned preferences, reusable prompt templates, persona handoff with lineage
- Seven provider families, one of them fully self-hosted
- A built-in **Model Context Protocol server** exposing eight tools, authenticated with existing API tokens — so an external agent works inside Qrly's own permission model

6 Proactive intelligence

- An insights feed detecting trend changes, milestones, new records, stale data, correlations, seasonality and anomalies, with severity and de-duplication
- Anomaly detection over a rolling window, excluding the point being judged
- **Change detection** — a metric that moved is badged, and one click hands it to the Investigator
- Investigations with findings and an executive summary, triggered manually, automatically above a severity threshold, or weekly
- Notification preferences by insight type, severity, channel and digest frequency

7 Four performance tiers

- **In-memory** — encrypted, byte-budgeted, with an optional columnar store
- **Persistent** — the application database, a time-series store, or a paired file per connection; survives a restart
- **Materialised tables** — a question's result written into a database *you own* as a real, queryable table, so other tools can point at it. Four write strategies, drift detection, and reads transparently rewritten to the cached table with fallback to source
- **Replication** — raw source tables copied into a target, with full reload, watermark, snapshot-delta and **change data capture** for three major databases, plus foreign-key-consistent table groups

8 Governance & lineage

- A typed lineage graph across connections, tables, questions, dashboards, alerts and deliveries — with a printable view and a downloadable report
- The **BCBS 239 module**: fourteen principles, scored self-assessments with a draft, submit and approve chain, a critical data element register with owner, risk category and materiality, seven data quality rule types scored across six dimensions, element-to-element lineage, a risk report register with sign-off, and a remediation issue log
- Schema authoring with import, dialect compilation and a plan preview — gated by three independent switches, all off by default, with every executed statement recorded verbatim
- A query audit log capturing actor, query type, full SQL, rows, duration, status and cache state, with retention and an auditor export

Two embed paths

A dashboard SDK with signed embeds, locked or editable parameters and automatic resizing — and a dependency-free embeddable assistant widget scoped to named connections, questions and tools.

A database API too

An auto-generated table and function API compatible with a widely used PostgreSQL REST convention — bulk insert, remote procedure calls, preference headers and transaction control.

Documented in the product

240 help articles across twelve categories with pre-built search, including forty troubleshooting runbooks and nine administrative operational runbooks.

WHO BUYS IT

Where “who changed this number, and on whose authority” is a real question.

Banking & capital markets
The sharpest fit. A productised BCBS 239 implementation aimed at chief risk officers, heads of risk data and regulatory reporting teams at supervised institutions.

Insurance
Risk data aggregation, critical data element ownership, quality scoring and an attestation calendar for periodic regulatory reporting.

Internal audit & compliance
An auditor is a first-class user: audit export endpoints, permission evidence, prompt records, and a printable governance dossier.

Pharmaceutical & life sciences
Full query audit with retention, verified-content trust signals, and AI that proposes rather than writes.

Energy & utilities
Operational reporting across mixed estates, with replication and materialisation reducing load on systems that cannot take analytical queries.

Public sector
Self-hosted, no external asset fetch, and a fully on-premise assistant option where a hosted model is not procurable.

Software vendors & ISVs
A genuine embedded-analytics play: signed embeds with locked parameters, a scoped assistant widget, per-tenant branding and custom domains.

Managed service providers
One deployment serving many client accounts, each with its own branding, domain, backups and administration.

Data platform teams
Change data capture, materialisation into owned tables and an auto-generated database API — Qrly is a replication and API layer as well as a reporting tool.

Legacy enterprise estates
Dialect support reaching databases modern tools ignore — and API code samples in eleven languages including COBOL, Fortran and RPG.

Sovereignty-bound organisations
Defence and financial infrastructure: a single artefact, no content delivery network, self-hosted fonts and libraries, and a local model option.

Multilingual mid-market
Five languages at complete coverage, edited in the product, with machine translation for adding more.

WHO USES IT

Platform superuser · **Organisation administrator** · **Curator / analytics lead** · **Analyst** · **Business viewer** · **Data engineer** · **Risk & compliance officer** · **External auditor** · **Integrator** · **AI agent**

Two permission ladders, granted by group. Content access runs from view through explore and curate to manage; data access distinguishes *restricted* — saved questions only, no raw SQL — from full. A user's effective level is the most permissive across their groups.

AGAINST THE ALTERNATIVES

CAPABILITY	QRLY	TYPICAL BI
Self-hosted, no telemetry	•	—
Full spreadsheet engine	•	—
Change data capture	•	—
AI writes require approval	•	—
Prompt audit table	•	—
BCBS 239 module	•	—
Agent protocol server	•	rare

TECHNICAL SPECIFICATION

One artefact. No content delivery network. Air-gap ready.

Runtime	Java 21 on Spring Boot 4 with virtual threads; native-image hints included
Metadata store	PostgreSQL 14 or later, 86 tables, schema owned entirely by versioned migrations
Front end	Server-rendered with vanilla ES6 — no package manager, no bundler, no build step
Assets	Around 71 MB vendored, fonts self-hosted, **zero content-delivery-network references**
API	91 REST controllers with OpenAPI, an agent protocol server, a database-compatible table and function API, live dashboard channels and streaming results
Auth	Signed tokens with a denylist; LDAP and Active Directory; single sign-on per organisation via the major identity providers
Tenancy	Tenant, organisation and user — with optional database row-level security as a third enforcement layer
Scale	~157 000 lines of application code, 84 entities, 731 test files, architecture tests

SECURITY & GOVERNANCE

- Group-based access on two ladders — content and data — with the most permissive grant winning
- Tenant isolation in three layers: scoped columns, request context, and optional database row-level policies with per-request role switching
- **AES-256-GCM at rest** for connection passwords, identity secrets and integration tokens; the in-memory query cache stores encrypted bytes
- Passwords hashed with bcrypt at cost twelve; account lockout on repeated failures
- API tokens stored hashed, with scopes, per-token rate limits, expiry and rotate-all
- Injection defence through sanitisers, a parser and prepared statements, on read-only connections; a five-pass sanitiser on white-label content
- Resource guards throughout: timeouts, row caps, a concurrency governor, per-user daily budgets, agent step and token budgets, abort thresholds
- Control mapping to a recognised trust services framework, with an explicit statement of what the product does *not* provide, plus a documented data-subject deletion procedure

> Multi-factor authentication and SAML are not implemented in the current release. Single sign-on is available through OpenID Connect and through LDAP or Active Directory.

COMMERCIAL

Licensed per tenant. Not per viewer.

Sold per tenant with unlimited viewers — in BI the value is everybody having the number.

REVENUE	PER YEAR
Up to €500k	**€900**
€1M	**€1,500**
€5M	**€2,500**
€25M	**€5,500**

REVENUE	PER YEAR
€100M	**€13,000**
€500M	**€37,000**
€1B	**€52,000**

Annual licence on your own turnover, ex VAT. The schedule is marginal like income tax — each slice of turnover is charged at its own rate, so there are no cliffs — with a €900 minimum and any yearly increase capped at 25%. Every tranche carries unlimited users, projects and organisations inside one tenant, self-hosting and support that rises with the tranche. Hosting and hardware are yours, AI runs on your own provider keys and is billed by them, and consultancy, setup, migration and training are quoted separately. Illustrative: a per-viewer licence at €12 a month bills four hundred readers €57,600 a year; at €5M turnover this is €2,500.

NEXT STEP

Connect one database.

Give us read access to a single schema. We will build the dashboard your team actually asks for every Monday, and hand you the deployment.

Sales tim@automatize.eu
Partners luc@automatize.eu
Web automatize.eu

BCBS 239 refers to the Basel Committee on Banking Supervision's *Principles for effective risk data aggregation and risk reporting*; Qrly supports compliance work but does not itself constitute or guarantee compliance. Specifications may change.