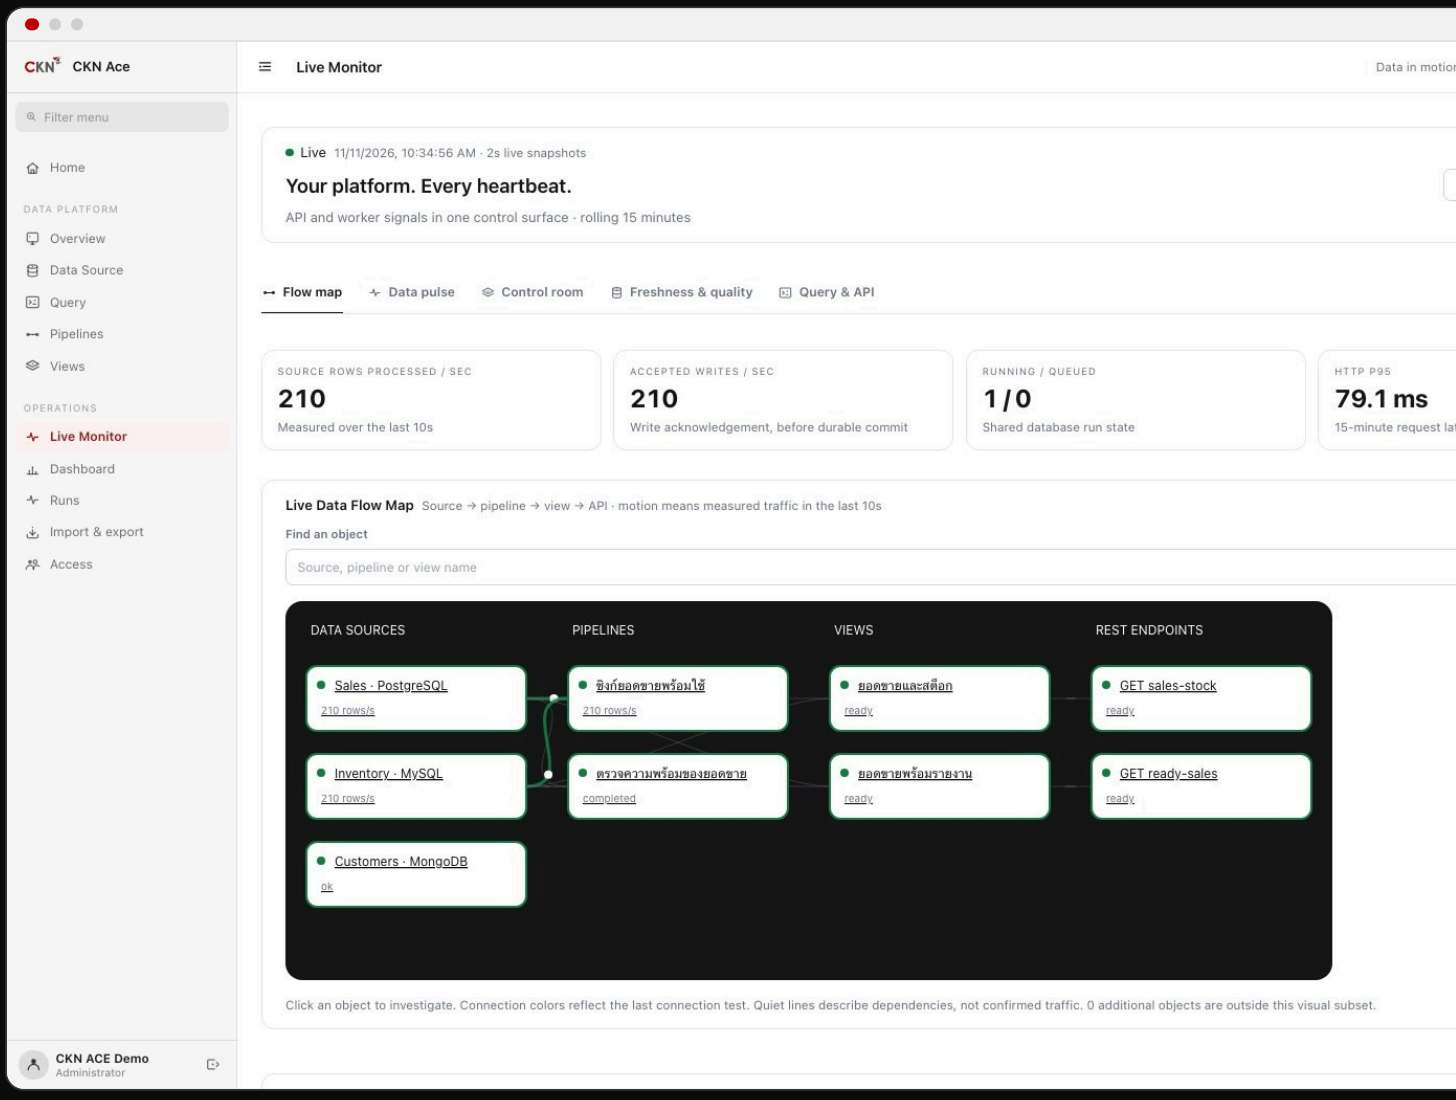
CKN ACE

DATA PLATFORM

Connect data. See it work.
Deliver value.

One data platform for connecting databases, querying across systems, running ETL, monitoring live, and publishing data as a REST API.



WHY CKN ACE

One question shouldn't take three systems

Sales, stock and customer data usually live in different databases. Teams pull each one out and merge them in a spreadsheet. CKN ACE connects every source in one place, answers questions across systems straight away, and moves data where it is needed, with control.

Arabica coffee: reorder now?

Sales · PostgreSQL	Warehouse · MySQL	Customers · MongoDB
2,501	**7**	**?**
units sold	units in stock	who reorders

Before CKN ACE: open each system, export files, merge by hand.

WITH CKN ACE · ONE SQL STATEMENT

Product	Sold	Stock	Status
Thermal tumbler	2,501	48	In stock
Arabica coffee	2,501	7	Restock
Coffee maker	2,498	3	Restock

Real result from the demo dataset: 3 rows in 177 ms

01 CONNECT

Connect

SQL databases, MongoDB and files, with a schema catalog

02 QUERY

Query across

One SQL statement over many sources, on DuckDB

03 ETL

Extract, transform, load

Guided pipelines or a drag-and-drop canvas

04 MONITOR

See it live

Flow map, volume, job status, data quality

05 DELIVER

Deliver

Views as REST APIs, a Trigger API and webhooks

9
source types (4 SQL · MongoDB · 4 file formats)

17
transform operators in pipelines

2 s
Live Monitor refresh over WebSocket

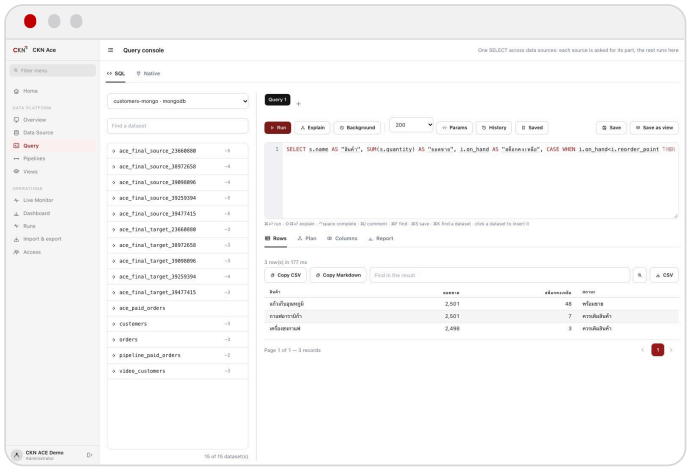
01 CONNECT · 02 QUERY

Connect every source. Ask in one query.

Supported data sources

PostgreSQL · MySQL · Oracle · SQL Server · MongoDB · Excel · CSV · JSON · NDJSON

- **SAP** via exported files: bring SAP data in as Excel or CSV and use it in queries and pipelines right away. For a direct SAP connection, contact our team.
- Set each connection to **read** or **read-write**. Pipelines never write to a read-only source.
- Connection credentials are encrypted with AES-256-GCM and never returned in any response.
- Catalog: column search, structure and sample rows, column profiles (nulls, distinct, min/max), ER diagram, schema-drift detection, and DDL export per SQL dialect.
- Refresh the catalog as a background job or on a cron schedule.



Query console: one SQL statement joins sales (PostgreSQL) with stock (MySQL)

Federated Query

Across databases

Splits the SQL statement, pushes filters down to each source, and combines the results in DuckDB.

Native mode

Runs on a single connection with bound-parameter SQL or a Mongo filter. On PostgreSQL, MySQL and Oracle it runs in a read-only transaction.

Reusable

Explain, history, saved queries with $name parameters, and scheduled runs.

Background runs

Follow progress as it runs. Results are kept as Parquet for 24 hours.

Export

Download CSV, Parquet or JSON, or deliver results into a table (upsert/replace) or a webhook.

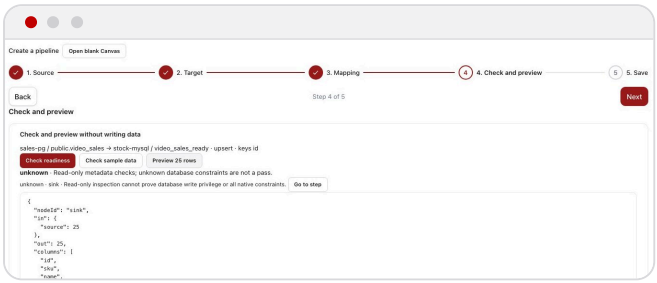
Safe by default

Limits on time, rows and concurrent queries, adjustable for your organization.

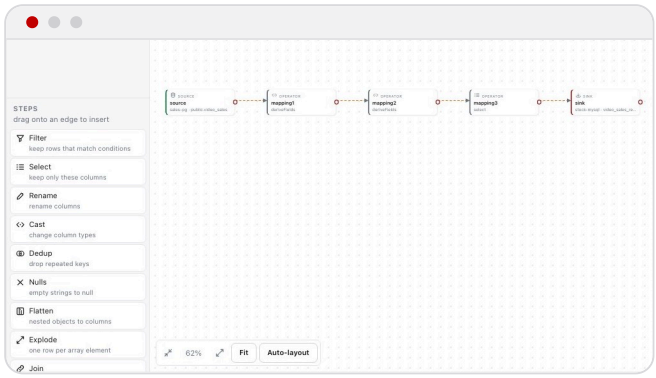
03 ETL PIPELINES

Extract, transform, load. All in one place.

Build pipelines two ways: a guided setup for everyday data moves, or a drag-and-drop canvas for complex flows. Both let you preview the data before anything is written.



Guided setup: source → target → column mapping → check and preview → save



Canvas: drag datasets and transform steps into a chain

17 transform operators

`filter` `select` `rename` `cast` `flatten` `explode` `nulls` `dedup` `hashJoin` `tagSource` `union` `deriveFields` `lookup` `validate` `route` `window` `windowJoin`

Every condition is structured JSON. The platform never runs code typed in by users.

Check before you run

A readiness check reports blockers, warnings and unknowns without writing to the target. Preview data before loading.

Run your way

Stream, memory or hybrid mode. Insert or upsert by key. Event-time windows supported.

Scheduling

Cron with a time zone. Scheduled runs act as the recorded owner.

Incremental

Resume from the last position with a cursor, plus lookback, soft-delete markers and isolated backfill.

Versions

Immutable revision history, compare changes, and restore as a new version.

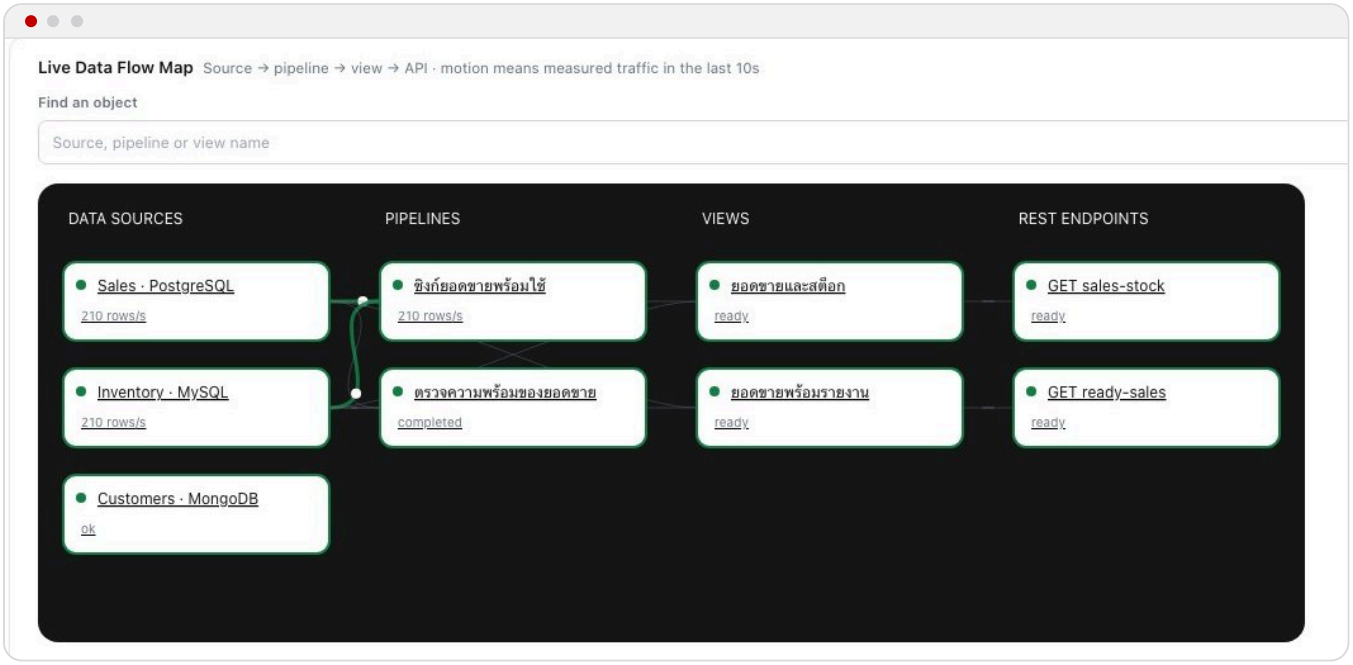
No lost rows

Rows that fail a rule are quarantined, then replayed after a fix or discarded, with retry and recovery.

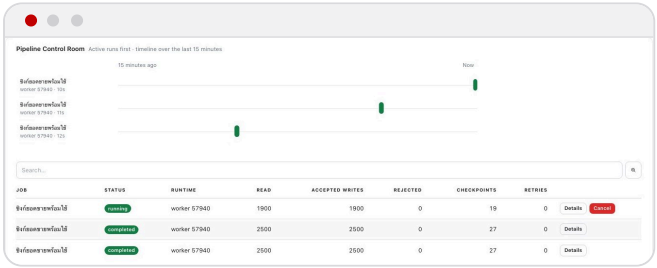
04 LIVE MONITOR · DATA QUALITY

Know every job. Trust every dataset.

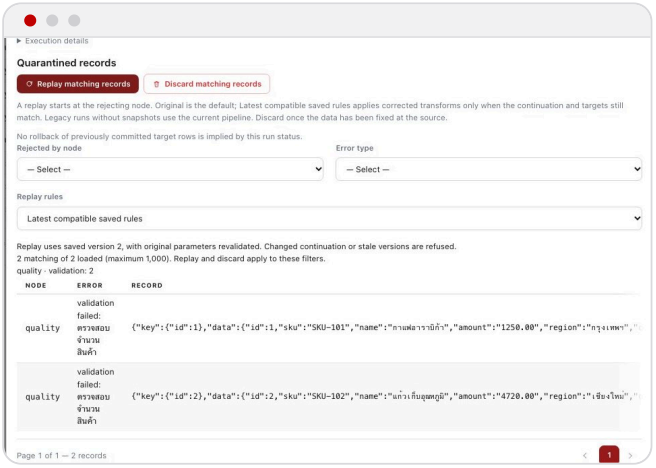
Live Monitor refreshes every 2 seconds with data paths, volume, job status and data freshness. It never stores SQL text or business data.



Flow map: sources → pipelines → views → REST endpoints, with measured traffic over the last 10 seconds



Control room: 15-minute timeline, job status, and cancel across processes



Quarantine: rejected rows filtered by step and error kind, then replayed

5 dashboards

Flow map · Data pulse · Control room · Freshness & quality · Query & API

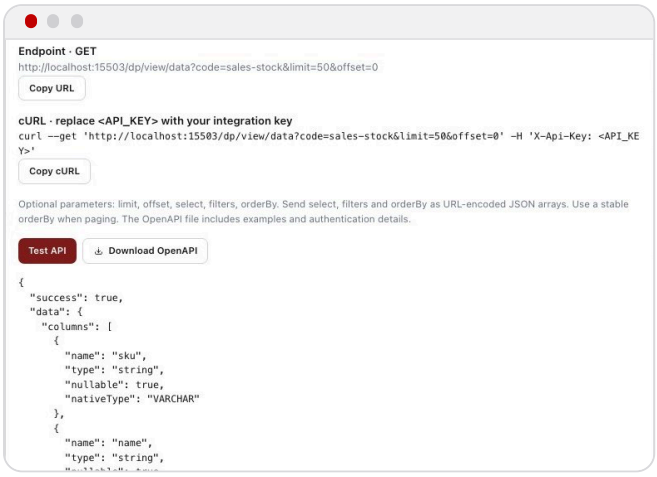
Data freshness

Compares each view's data age with its target and shows it in colour at a glance.

Notifications

In-app history and webhooks when a run completes, fails, is partial or cancelled, rejects rows, or stalls.

Deliver data, securely



Generate REST API: endpoint, cURL example, test button and OpenAPI 3.1 file

Views → REST API

- Views combine several sources and serve them **live**, **cached** (up to 24 h) or **materialized** (refreshed on a schedule or continuously).
- GET `/dp/view/data` supports limit, offset, select, filters and orderBy.
- Download OpenAPI 3.1 so developers can integrate right away.

Connect other systems

- **Trigger API** starts pipelines from outside systems, with an Idempotency-Key and typed parameters.
- **Signed webhooks** receive data, signed with HMAC-SHA256, with key rotation.
- **MCP server** lets AI agents search columns, explain and run SQL (dp:read required).

Security and access control

Built-in OAuth 2.1

Authorization code + PKCE, refresh tokens, client credentials, RS256 tokens, discovery and JWKS.

Fine-grained access

Scopes dp:read · dp:write · dp:admin, plus read/write grants per connection.

Request protection

CSRF, idempotency, rate limits, security headers and an access audit log.

Secrets kept safe

Credentials encrypted with AES-256-GCM, passwords hashed with scrypt, API keys stored as SHA-256.

Outbound control

Blocks dangerous internal addresses and supports an allow-list for webhooks.

Your data stays put

Business data stays in your source databases. CKN ACE keeps only its configuration in PostgreSQL.

CKN ACE

Book a demo
with your own data

We connect a sample of your data sources, build your first pipeline, and publish it as an API, all in one meeting.

Book a CKN ACE demo

WEB www.ckn.ooo/products/ace.html

TEL +66 65 518 5509

EMAIL team@innobiz.studio

Scan for details