

COLROWS • DEPLOYMENT GUIDE


How Colrows works with your data stack

A plain-language guide for business and IT leaders

Colrows connects to the systems you already have: databases, data platforms, catalogs, BI tools and documents. It learns what your business terms mean, then answers questions in plain English with one approved definition behind every number. Your data stays where it is, and your current tools stay in place.

[Shared SaaS](#)[Dedicated SaaS](#)[Private cloud](#)[On-premise](#)engage@colrows.com

1 How Colrows fits into your data estate




Grey sources hold the data that queries run on. Colrows reads the other sources only to learn business meaning.

- **Connectors** read your sources with read-only accounts. They collect table structures, catalog glossaries, BI calculations and query history, not your business records.
- **Document intelligence** reads policies, manuals, wiki pages and scanned documents. For scans and images, it uses optical character recognition (OCR) to extract the text. It splits each file into sections and converts each section into a vector embedding, a numeric record of its meaning. Colrows can then find the right paragraph by meaning, not only by keyword, and cite the page in its answer.
- **The semantic layer** is a knowledge graph of your business terms, metrics, entities and rules, stored in Neo4j. Document sections, embeddings and term vectors are stored in MongoDB, which also serves as the vector store. The vectors connect the words people use, such as "sales", "turnover" or "revenue", to the correct approved definition. Both stores run inside your deployment.
- **AI agents** propose new definitions, validate them and detect changes in your sources. People approve. The agents never promote a definition on their own.
- **The governed query engine** applies access rules, writes the database query and records an audit trail for each answer.
- **Your AI model** reads the wording of a question. It does not decide what a metric means or who can see what. You choose the model, and it can run inside your network.

2 Example deployment: Tableau, Snowflake and Oracle

This illustrative example shows a typical company. It uses Tableau dashboards on Snowflake, keeps finance data in Oracle, keeps a glossary in Colibra and keeps finance policies as PDF files.



Step 1: Connect the data

IT creates read-only accounts for Snowflake and Oracle. Colrows imports the Tableau workbooks and the Colibra glossary, and the finance team uploads its policy PDFs, including older scanned copies. Users sign in with the company login (Microsoft Entra ID in this example). No data is exported or moved.

Step 2: Build the semantic layer

Colrows scans tables, columns and query history, reads the Tableau calculated fields, reads the scanned pages with OCR and turns all policy documents into searchable vectors. It then proposes each business term with its evidence. For "net revenue", the evidence looks like this:

Source	What Colrows found
Tableau calculated field	Net revenue = gross sales minus returns
Colibra glossary	Net revenue is reported after discounts
Finance policy PDF, page 4	Revenue counts when the order ships
Finance owner (decision)	Approves one definition: gross sales minus returns and discounts, counted at shipment

The approved definition becomes version 1. If finance changes it later, Colrows creates version 2 and keeps the history, so past answers stay explainable.

Step 3: Ask and get answers

A regional manager asks "What was net revenue by region last quarter?" in Microsoft Teams. Colrows checks the manager's role, limits the data to that region, writes one read-only query and sends it to Snowflake. The answer shows the number, the definition used and a citation to the policy page. The

Tableau dashboard connects to Colrows and uses the same definition, so the dashboard and the chat answer agree.

3 Deployment options

Option	Where Colrows runs	Best fit for
Shared SaaS	Colrows cloud, shared with other customers	Fast start and evaluations
Dedicated SaaS	Colrows cloud, only for you	Isolation without running servers
Private cloud	Your own AWS, Azure or Google Cloud account	Data and keys stay in your cloud
On-premise	Your own data centre, air-gapped if needed	Regulated and sovereign workloads

4 The deployment plan

A private-cloud or on-premise deployment usually takes 4 to 8 weeks. SaaS starts faster because there is nothing to install. Phases 3 to 5 often run at the same time.

Phase	What happens	Your team's part	Typical time
1. Discovery	Agree on use cases, sources, security needs and go-live date	Business owner, IT and security	Before start
2. Install	Colrows engineers install and secure the platform in your environment	IT provides servers or a cloud account	2 to 3 weeks (not for SaaS)
3. User sign-in	Connect your company login and user groups	IT	1 to 2 business days
4. Connect sources	Read-only access to databases and platforms; import catalogs and BI logic; upload documents	Data team gives access	3 to 5 business days
5. Build definitions	Colrows proposes business terms and metrics with evidence	Business and data owners approve or correct them	2 to 3 weeks
6. Go live	Users ask questions in chat, Microsoft Teams or their BI tool	Users give feedback	2 weeks close support, then standard support

5 Common questions

Does our data leave our systems? No. Colrows reads table structures and metadata, not your records, to build definitions. To answer a question, it sends a read-only query and shows the result. Uploaded documents are the one thing you send to Colrows, and they stay inside your deployment.

Do we need to replace our BI tool or database? No. They stay. Any tool that connects to a database can connect to Colrows, and if you have no BI tool, users start with the Colrows chat or Microsoft Teams.

Can the AI make up a number? No. The AI reads the question, but every number comes from an approved definition and a query on your data. If a question is unclear, Colrows says so instead of guessing. A chatbot that writes queries directly is much less reliable: in the Spider 2.0 benchmark of enterprise data tasks, the best AI agent solved only 21.3 percent of tasks ([Lei et al., 2024](#)).

Next step: book a discovery workshop at engage@colrows.com. We confirm your sources, the connector for each one, the security review and a go-live date.