

RangerIO — Verified User Guide

RangerIO current release (0.16.0) · Reviewed October 17, 2026

1. Welcome to RangerIO
2. Why RangerIO Exists
3. How RangerIO Works
4. System Requirements & Installation
5. Bringing Data In
6. The Dossier
7. Privacy Scan
8. Safe Handoff
9. Asking Questions
10. Connecting to Your AI Tools
11. Reports & Exports
12. Use-Case Workflows
13. What RangerIO Doesn't Do
14. Technical FAQ

Chapter 1 · Welcome to RangerIO

Guide scope — RangerIO current release. Last reviewed October 17, 2026. This guide includes only behavior verified in the current product.

RangerIO is a desktop data-intelligence and sanitization tool for working with files that may contain sensitive information. It imports supported files, scans them for potential PII, provides review and correction tools, answers questions about them, builds a reusable file profile called an **Intelligence Dossier** when you ask, and prepares sanitized or tokenized context for the AI provider you choose.

In everyday terms, RangerIO helps you open a work file, understand what is in it, check it for sensitive information, ask questions about it, and prepare an output for review or sharing.

What do you need to do?

Need	Start	Workflow
Bring in a spreadsheet, document, or group of files	My Data → + Add Files	Import one or more files
Check whether a file contains personal or sensitive identifiers	My Data → select the file → Preview Data	Review and correct findings
Correct something RangerIO marked incorrectly	PII review	Choose PII , Not PII , or Change type , then Apply
Create a copy with detected sensitive values hidden	Document preview, Report , or PII Inventory	Export a masked document or a masked data export
Ask for a total, list, comparison, or summary	Explore , or Chat beside a file on My Data	Start a data question
Turn one or more files into a draft report	Report → AI Assisted	Generate an AI-assisted report
Connect your cloud AI provider or your own local model server	Home screen → + Add Model Provider , or Models	Connect an AI model
Use RangerIO from Claude Desktop	Settings → Claude Desktop Connector	Configure Claude Desktop
Repeat a review for a new client, month, or reporting period	My Data → Actions	Repeated engagement workflow

A good first workflow

Use when: you want to learn the product with one representative file before importing an entire engagement.

Before you start: choose a supported file that you are authorized to use. If possible, begin with a copy rather than your only copy of the file.

- 1 Make sure an AI model is connected (see **Connect an AI model**).
- 2 Open **My Data**.
- 3 Choose **Actions** → **New Project** and create a clearly named project.
- 4 Choose **+ Add Files** → **Files** and pick one file.
- 5 Check the file's type, size, rows, and columns, then choose **Start Import**.
- 6 Wait for the green tick. New files go to **Library**; use **Move to project** to put the file in your project.
- 7 Select the file and choose **Preview Data** to open **PII review**.
- 8 Review the findings. Correct any with **PII**, **Not PII**, or **Change type**, then **Apply**.
- 9 Choose **Chat** beside the file to ask about it, or continue to **Report**.

Done when: the file appears in the correct project, its detail view is available, and you understand both its readiness state and its privacy-scan state.

If it fails: open **Events** in the header, read the failed or interrupted task, and use the recovery steps in “During import.” Do not keep importing duplicate copies while an earlier task is still active.

RangerIO is not only a chat interface. The data engine works before a model is involved:

- 1 You import a file.
- 2 RangerIO reads its structure and indexes it.
- 3 RangerIO scans it for personal and sensitive values.
- 4 You review the findings.
- 5 You ask questions, build an Intelligence Dossier, create a report, or export a masked copy.
- 6 Your AI provider is used only when the task needs a model, and personal details are replaced before anything is sent.

RangerIO Enterprise

RangerIO Enterprise is a licensing tier, not a separate product. It runs the same data engine, privacy workflows, and export paths described in this guide, and adds two things for organizations: **centralized management** for multi-seat deployments, and **fully supported database and REST API data sources** alongside file imports.

- Centralized management for deployment, configuration, and policy across many workstations.

- SQL database sources as a supported, fully backed import path.
- REST API sources as a supported, fully backed import path.
- Full support for the organizations running these connected sources.

Enterprise — RangerIO Plus is planned for Q1 2027. Enterprise licensing is expected to begin shipping in Q3 2027 and is available on request today. Database and REST API sources are not usable in the current desktop product. For more information on our Enterprise version, please contact info@rangerio.com.

What you can do

- Create projects and project groups.
- Import individual files, or a Cowork Folder that RangerIO watches for new and changed files.
- Follow import tasks and recover an interrupted import.
- Preview data and build an Intelligence Dossier for a file.
- Review PII findings in tables and documents.
- Mark missed sensitive values, dismiss false positives, and change the PII type.
- Ask questions about one or more ready files.
- Review grounding, coverage, data-quality caveats, and what was replaced before a cloud request.
- Export data, reports, and masked copies.
- Turn on the local MCP server and connect Claude Desktop.

What needs an AI model

RangerIO does not include an AI model. You connect your own: a cloud provider with your API key, or a local model server you run, such as Ollama, LM Studio, or vLLM. Questions on **Explore**, the Intelligence Dossier, and AI-assisted reports need a connected model. Import, privacy review, and standard exports do not.

Before relying on an output

A completed import, a “Ready” status, a PII scan, a masked export, or an AI answer is not a certification. Review consequential outputs against the source material and your organization's policy.

Terms you will see

Product term	Plain-language meaning
Project	A workspace that groups the files for one client, matter, review, or period. It is not the same as a folder on your computer.
Source	RangerIO's working record for a file you imported. A workbook's sheets belong to its file.
Intelligence Dossier	A reusable profile of a file that RangerIO builds when you ask: structure, statistics, quality signals, sensitive-data findings, and other available context.
PII	Personally identifiable information: values that may identify or relate to a person. The legal definition depends on the jurisdiction and context.
Severity	How risky the findings that were detected may be.
Coverage	Whether the privacy scan finished sufficiently. It does not mean the file is clean.
Ready	The inputs needed for a particular action are available. It is not a guarantee that the source or result is correct.
AI provider	The cloud service or local model server you connect for tasks that need an AI model.
MCP connector	The optional connection that lets Claude Desktop request approved RangerIO tools.

Chapter 2 · Why RangerIO Exists

Professional work often starts with files that should not be pasted directly into a general-purpose AI chat: contracts, employee files, accounting workbooks, research exports, operational records, and other client material.

RangerIO separates **preparing and understanding the source** from **choosing an AI engine**.

The practical problem

Without a preparation layer, a user has to decide between:

- doing the work manually;
- sharing more source content than intended;
- manually removing identifiers;
- rebuilding the same context in every new AI conversation; or
- using an AI answer without a reusable, inspectable source profile.

RangerIO addresses this by building reusable local intelligence first.

The operating model

Local preparation

The data engine imports, profiles, indexes, scans, and stores working intelligence locally.

Explicit model choice

The user can use:

- local lookups and exact calculations;
- a cloud AI provider you configure with your own API key;
- your own local model server; or
- Claude Desktop, through the connector you turn on.

Review before disclosure

The user reviews scan coverage, findings, masking, and the organizational rules that apply to the data.

The product boundary

RangerIO reduces preparation work and provides safety controls. It does not make an otherwise prohibited disclosure permitted, replace legal review, or guarantee that every sensitive value will be detected.

Chapter 3 · How RangerIO Works

End-to-end workflow

- 1 **Select files**
- 2 **Check their structure**
- 3 **Import and scan for personal data**
- 4 **Review and correct findings**
- 5 **Ask, report, or mask**
- 6 **Export locally or use your AI provider**

Step 1 — Select and inspect

For file imports, RangerIO performs a light structure scan before the full import:

- tables show their type, size, rows, and columns;
- files that cannot be read show an error before the import proceeds;
- you can remove a file before importing.

This structure scan is not the final PII scan.

Step 2 — Import and scan

The import creates the working source, reads its schema and structure, indexes it for search, and scans it for personal data. Imports run in the background, so you can keep working. A green tick on **My Data** means the file is ready for questions.

The **Intelligence Dossier**, a richer profile of the file, is not built during import. You build it when you need it from the file's details (see **The Dossier**).

Step 3 — Privacy coverage

Import completion and PII-scan completion are separate states. A file may be imported while its privacy scan is still pending or incomplete. Questions wait until the scan is done.

Step 4 — Review readiness and correct findings

Open the file's details and **PII review** after processing. Review structure, data-quality signals, scan coverage, and findings. Correct false positives or missed values and apply your decisions; RangerIO rescans the file to match. A refresh re-imports the file; a rescan refreshes detection and coverage.

Step 5 — Route the task

Task	Local data engine	Model needed?
Read schema or row count	Yes	Usually no
Run a local analytical query	Yes	Not always
Produce an interpretive answer	Yes, for context	Yes
Review or correct PII	Yes	Scan components may use models, but user review remains required
Export data	Yes	No for standard data export
Generate an AI-assisted report	Yes, for source context	Yes
Use an MCP host	Local connector plus host	Yes, in the external host

Local and external model paths

Path	Data-handling note
Cloud AI provider you configure	Personal details are replaced before sending, and every provider call is checked before it leaves your computer
Your own local model server	Runs where you host it; the server must be reachable from your computer
Claude Desktop through the connector	Tool results return through Claude Desktop and stay sanitized or tokenized

Important distinction: in-app cloud vs. MCP

For an in-app cloud call, RangerIO can mask outbound values and rehydrate supported response values locally.

For an external MCP host, public rehydration is blocked because an MCP tool result is sent back through that host's conversation context.

Chapter 4 · System Requirements & Installation

Current desktop builds

- **macOS:** Apple Silicon, arm64
- **Windows:** x64

Startup requirements

Component	Minimum	Recommended
RAM	8 GB	16 GB
Processor	4 CPU cores	8 CPU cores
Free storage	5 GB on an SSD	10 GB
GPU	Not required	Not required

Actual working capacity depends on file size, file complexity, and other applications using memory.

Install on Windows

- 1 Download the current Windows x64 installer.
- 2 Run the installer. It downloads the application package while it installs.
- 3 Accept the license and choose the install folder.
- 4 On the last page, choose **Run RangerIO**. You can tick **Delete the downloaded installation files** to free about 3.4 GB; future updates then download the full package.
- 5 Complete the in-app setup.

Install on macOS

- 1 Download the current Apple Silicon disk image (.dmg).
- 2 Open the disk image.
- 3 Move RangerIO to Applications.
- 4 Start RangerIO.
- 5 Complete any macOS security prompt and the in-app setup.

First launch

RangerIO starts its local data engine, checks your computer, and asks you to connect an AI model. You can choose **Skip for now** and connect one later. On the last step you can also turn on the Claude Desktop connector.

Connect an AI model

Import, privacy review, and standard exports do not need an AI model. Questions on **Explore**, the Intelligence Dossier, and AI-assisted reports do. RangerIO does not include a model: you connect your own.

Which model option fits your situation?

Your situation	Use	What to know
You only need to import, review privacy findings, or export standard data	Continue without a model	RangerIO tells you when a feature needs one
Your organization permits an external AI service and you have its API key	+ Add Model Provider on the home screen	Personal details are replaced before anything is sent; the provider's terms apply to what it receives
Your organization runs Ollama, LM Studio, vLLM, AnythingLLM, or another compatible local server	Models → Llama / vLLM / Network Model	The server must be running and reachable from your computer; RangerIO does not install or manage it

Connect a cloud provider

- 1 On the home screen, choose **+ Add Model Provider** in the **Active AI Model** panel. (Or open **Models** from the  menu and choose **Cloud / API**.)
- 2 Choose the provider: **OpenAI**, **Anthropic**, **Fireworks AI**, or **OpenAI-compatible**. The **API Endpoint** is filled in for known providers.
- 3 Paste the **API Key**.
- 4 Choose the model from the list, or enter a model ID confirmed by your provider.
- 5 Enter a **Profile name**.
- 6 Choose **Test Connection**, then **Save & Connect**.
- 7 Confirm the model shows as active in the **Active AI Model** panel.

Use a large, current model. Small models can fail on long files with a “context too long” error. Provider credentials, charges, retention, and contractual terms belong to that provider.

Connect a local model server

- 1 Start the server and load at least one model in it.
- 2 Open **Models** from the  menu.
- 3 Choose **Llama / vLLM / Network Model**.
- 4 Under **Local LLM servers**, choose **Connect** beside the running server.

- 5 If no server is detected, choose **Configure** or enter a custom endpoint, then choose **Test & Connect**.
- 6 Confirm the server appears as **Active** under **Saved Configurations**.

Switch, edit, or remove a saved model

- 1 On the home screen, open **Models and Providers**.
- 2 Under **Saved Models**, choose **Make Active** on the model you want to use. **Check now** confirms it still answers.
- 3 Use **Choose model** to pick a different model from the same provider.
- 4 Use **Edit** to change the endpoint, model, name, or key, then choose **Save Changes**.
- 5 Use **Remove** only when the saved model is no longer needed.

Removing a saved model does not delete imported files.

Working-data location

RangerIO uses a local data directory under the signed-in user's home directory. The current engine uses `RangerIO_local` as its local working directory. Imported working data, indexes, metadata, registries, and audit information can be present there.

Updates

macOS: the app checks for updates and asks before downloading; choose **Update & Restart** to install. **Windows:** automatic updates are off. Choose **Check for Updates** in the  menu, then **Open RangerIO Releases** to download the latest signed installer.

Troubleshooting startup

If RangerIO cannot start a required component:

- 1 Read the message RangerIO shows.
- 2 Close memory-heavy applications if the system is low on RAM.
- 3 Retry the action.
- 4 If a model is needed, check that your AI provider answers: on the home screen, open **Models and Providers** and choose **Check now**.
- 5 If the issue persists, open the  menu and choose **Generate Support Bundle** (or **Export Support Snapshot** at the top of **Settings**).
- 6 Save the snapshot, check what it contains, and email it to info@rangerio.com.

Chapter 5 · Bringing Data In

Projects and project groups

A **project** organizes related files for an engagement or analysis. A **project group** organizes multiple projects.

Create a project

- 1 Open **My Data**.
- 2 Open **Actions**.
- 3 Choose **New Project**.
- 4 Enter a project name.
- 5 Choose **Create**.

Create a project group

- 1 Open **My Data**.
- 2 Open **Actions**.
- 3 Choose **New Project Group**.
- 4 Enter a **Group Name**, choose a color, and add projects to the group.
- 5 Choose **Create**.

Verified import formats

Tabular and structured

- CSV
- TSV
- XLS
- XLSX
- XLSM
- JSON
- JSONL
- Parquet

Documents and text

- PDF
- DOCX
- PPTX
- HTML

- TXT

Needs an extra tool — legacy .doc and .rtf files open only if Pandoc is installed on your computer; RangerIO does not include it. Otherwise, save them as DOCX first. Scanned PDFs (pictures of text) need to be imported again with OCR, as the message explains.

Import one or more files

- 1 Open **My Data**.
- 2 Choose **+ Add Files**. The **Import Data** dialog opens.
- 3 Choose **Files** and pick one or more files.
- 4 Check each file's type, size, rows, and columns. Remove anything picked by mistake with the x.
- 5 Optionally open **Advanced Options** to set the **Domain** (default **Auto-detect**), a **Reporting Profile**, and **Key Notes**, for example "Q4 audit data; focus on vendor transactions above \$10,000." Most imports can leave these as they are.
- 6 Choose **Start Import**.
- 7 Follow progress in the import card or in **Events** in the header.
- 8 New files go to the **Library** project. To file them elsewhere, choose **Move to project** on the file, or select files and choose **Move selected**.

The Domain, Reporting Profile, and Key Notes fields guide later intelligence and reporting. They do not change the file's data.

Drag files onto My Data

My Data accepts dragged files and folders. They open the import dialog with the files already listed. You still review them and start the import.

Import a Cowork Folder

A **Cowork Folder** is a folder RangerIO watches. Its files are imported into their own project, and RangerIO checks the folder for new and changed files and asks you before applying anything. Subfolders are not imported.

- 1 Open **My Data**.
- 2 Choose **+ Add Files → Cowork Folder**.
- 3 Choose the folder.
- 4 Select the **File Types** to include.
- 5 Choose **Scan Directory**.
- 6 Review the discovered files. Use **Select All**, **Deselect All**, or select individual files.
- 7 Choose **Continue with N files**.

- 8 Follow the import tasks. The files appear in a new project named **Directory:** followed by the folder name.

When files in the folder change, the project shows **N changes**. Open the project, review the list, and choose **Review & Sync** to apply the changes or **Dismiss** to ignore them.

Enterprise — The desktop app imports files and folders. Database and REST API sources are part of **RangerIO Enterprise** licensing, together with centralized management and full support. Enterprise licensing is expected to begin shipping in Q3 2027 and is available on request today. Contact info@rangerio.com.

During import

The task system reports current work and can show partial, interrupted, paused, failed, or completed states.

If an import is interrupted after a restart:

- 1 Look for the banner on **My Data**.
- 2 Choose **Resume** (or **Continue preparation**) to continue when resources permit.
- 3 Choose **Dismiss** to stop the unfinished work instead.
- 4 If Resume is refused because memory is low, close other applications and retry.

What appears after import

A file's row and details can show:

- file name and format;
- project placement;
- row or document information;
- a green tick when the file is ready for questions;
- privacy severity and scan coverage;
- data-quality and structure warnings;
- Intelligence Dossier status;
- **Preview Data**, **Chat**, **Refresh**, **Move to project**, and **Delete** actions.

The **My Data** toolbar also supports:

- search;
- **Filter** by status or severity, and **Sort by**;
- selecting one or more files;
- **Use in Chat** when two or more ready files are selected;
- **Refresh selected** and **Move selected**; and
- **Delete**.

Source status

Use these concepts carefully:

- **Ready** (green tick): the file is ready for questions.
- **Needs attention**: a warning, incomplete process, or uncertain structure requires review.
- **Failed**: the import did not complete. The message says why and what to do.
- **PII scan required**: the privacy scan has not completed.
- **unverified** after a severity: PII findings may be real, but the scan did not finish fully.

A source can be analytically ready while its privacy coverage still needs attention.

Excel workbooks

A workbook is shown as one file with its sheets beneath it. Users should:

- 1 Open the workbook on **My Data**.
- 2 Review each sheet's rows, columns, quality, and PII information.
- 3 Ask about the specific sheet when a question must be limited to it.
- 4 Verify which file and sheet an answer or report used.

Refresh a source

- 1 Open **My Data**.
- 2 Select the file.
- 3 Choose **Refresh** on the file, or select files and choose **Refresh selected**.
- 4 Follow the task in **Events**.
- 5 Review structure and PII coverage again.

A refresh re-imports the file. If it has an Intelligence Dossier, rebuild it afterwards.

Delete a source

- 1 Select the file or files.
- 2 Choose **Delete** (for a workbook, **Delete workbook**).
- 3 Confirm the exact files.
- 4 Verify that the correct project remains.

Deletion removes RangerIO's working records for the selected files. It does not delete the original file on your computer.

Chapter 6 · The Dossier

What a dossier is

The **Intelligence Dossier** is RangerIO's reusable structured understanding of a file. It is not built during import: you build it when you need it. Reports need it, and questions and the Claude Desktop connector use it when it exists.

A dossier is not a second copy of the original document for human reading. It is a structured intelligence layer.

What can appear in a dossier

Depending on the source and completed processing:

- source name, format, size, and timestamps;
- row, column, page, paragraph, or document-structure information;
- schema and inferred types;
- per-column statistics such as missingness, uniqueness, common values, and numeric summaries where applicable;
- data-quality and parsing-confidence signals;
- detected entities;
- PII findings and scan coverage;
- searchable chunks and retrieval metadata;
- domain and reporting-profile context;
- workbook/sheet attribution; and
- provenance used by questions and reports.

Not every field is available for every format.

Build and open a dossier

- 1 Open **My Data** and select the file.
- 2 Wait until the privacy scan is complete and an AI model is connected.
- 3 Choose **Build Intelligence Dossier** in the file's details.
- 4 Wait for **Intelligence Dossier · built** followed by the time.
- 5 Choose **Preview Dossier** to read it.

Building a dossier uses your AI model. Personal details are replaced before anything is sent. If a dossier cannot be built, RangerIO says why (for example, the privacy scan is incomplete or no model is connected) and offers **Retry**.

You can also build dossiers during import: in **Advanced Options**, turn on **Build Dossier Intelligence**. It is off by default.

When the dossier changes

A dossier changes only when you build or rebuild it. If the file changes afterwards, RangerIO marks the dossier as being from an earlier version of the file; choose **Rebuild** to refresh it. Review its warnings again before relying on it or exporting.

Data-quality signals

RangerIO can surface structure confidence, missingness, anomalous values, normalization issues, and other data-quality context depending on the source and operation.

For analytical answers, caveats can be operation-aware. For example, a total over a materially incomplete column can carry a missingness warning.

A missing warning is not proof that the data is clean.

Chapter 7 • Privacy Scan

What the Privacy Scan does

RangerIO scans supported structured and document content for potential sensitive values. Detection can combine machine-validatable patterns, model-detected entities, prior user marks, and stored corrections.

Findings are reviewed in the **PII review** window. Each finding is labeled **Critical** or **Non-critical**.

Coverage and severity are different

- **Severity** describes the risk level of findings that exist.
- **Coverage** describes whether scanning completed sufficiently.
- **PII scan required** means the file has not been scanned. It is not a severity.
- **unverified** after a severity means findings may exist but the scan is incomplete.

Zero findings on a file that says **PII scan required** does not mean “No PII.”

Review and correct findings

The same steps work for tables and documents.

- 1 Open **My Data** and select the file.
- 2 Choose **Preview Data**. The **PII review** window opens. (From a file preview, choose **Privacy**.)
- 3 Start with **Needs Your Attention**, which lists the PII types still to review, or open **Browse detections** to see every finding.
- 4 Use **Next finding** and **Previous finding** to move between findings, and **Next match** and **Previous match** to step through each occurrence.
- 5 Select a finding, cells, a column, or text, then choose **PII**, **Not PII**, or **Change type**.
- 6 If the scan missed something, select it and choose **Mark Selection as PII**, or use **Add PII Values** to enter a value.
- 7 Choose **Review decisions**, check the list, then **Apply**. Apply saves all your decisions together.
- 8 RangerIO rescans the file so protection matches your decisions. You can keep browsing while it finishes.

Show PII highlights and **Hide PII highlights** turn the marks on and off. A saved **Not PII** decision can be undone with **Restore finding**.

Correction scope

A decision applies to every matching occurrence. The decision card shows how many (**Exact matches in this document**, or in this sheet). Check the count before you apply.

Re-scan and export gate

Masked document export uses the last complete scan. After you apply decisions:

- 1 Wait for the rescan to finish.
- 2 If the scan did not finish, your decisions are still saved. Use the scan action RangerIO shows (**Rescan** or **Run the scan again**).
- 3 In the document preview, wait for **Scanned [time] · export ready** before using **Export masked**.

When the scan is out of date, the document preview shows **Masked export paused until re-scan**. Choose **Re-scan now**.

If the scan found nothing

A completed scan with no findings is not a guarantee. Review the file and use **Mark Selection as PII** for anything missed.

What automated detection cannot promise

- masking of health terms inside documents, which are not detected in this release;
- complete recall for every locale and identifier;
- perfect distinction between real and template data;
- correct classification of every organization, date, account reference, or short name;
- a legal conclusion about whether a value is personal information; or
- permission to disclose the resulting file.

Human review remains required.

Chapter 8 · Safe Handoff

Purpose

Safe Handoff is the set of controls used when an explicitly configured cloud or connector workflow receives RangerIO context.

In-app cloud-provider flow

- 1 **Read privacy state.** RangerIO checks the selected files' scan coverage and current registry.
- 2 **Build task context.** Only the context the question or report needs is assembled.
- 3 **Replace detected values.** Registered values are replaced with stable stand-ins.
- 4 **Check outbound content.** RangerIO checks the request before every provider call. If it cannot confirm the request is sanitized, the cloud call is blocked.
- 5 **Send to your provider.** You configured this provider and are responsible for its account and terms.
- 6 **Receive the response.** RangerIO puts the original values back locally.
- 7 **Show what was sent.** Under the answer, **Cloud request** lists what was replaced before sending, and **View lineage** opens the full record.

Stable stand-ins

Stable stand-ins let repeated values remain linkable during an engagement. This can preserve relationships such as repeated counterparties without sending the original detected value.

The mapping is maintained in a local registry scoped to the engagement/project workflow.

Fail-closed behavior

If RangerIO cannot validate a cloud-bound context, the expected outcome is a blocked cloud request with an explanation—not a silent downgrade.

The user can then:

- rescan the file;
- correct findings;
- select fewer or different files; or
- use your own local model server.

External MCP flow

An external MCP host is different:

- 1 The host asks RangerIO to run a tool.
- 2 RangerIO authorizes the tool and builds a response.
- 3 PII-aware scrubbing and egress controls apply.
- 4 The tool result returns to the external host.
- 5 That result becomes part of the host's conversation context.

Because of step 5, RangerIO's public MCP surface blocks rehydration. The external host receives sanitized or tokenized content.

What MCP tools may return

Depending on the tool and scope:

- project and source metadata;
- dossier summaries;
- schema and quality information;
- PII summaries;
- sanitized row previews;
- sanitized analytical query results; and
- tokenized exports when the user explicitly enables that permission.

Review requirement

Masked context can still expose structure, categories, counts, excerpts, or undetected values. The user must decide whether the task and external provider are permitted.

Chapter 9 · Asking Questions

Start a data question

- 1 Open **Explore**.
- 2 Choose + (**Select a data source**). The **Select files to query** panel opens.
- 3 Search for and tick one or more ready files. They appear as chips above the question box.
- 4 Preview a selected file when needed.
- 5 Type the question.
- 6 Type @ to reference a column by name where useful.
- 7 Press Enter.
- 8 Follow progress, and choose **Cancel** if the question should stop.
- 9 Review the answer and the signals under it.

On **My Data**, you can also choose **Chat** beside a ready file, or select two or more files and choose **Use in Chat**.

How RangerIO routes a question

RangerIO uses multiple answer paths:

- 1 **Direct lookup.** Schema, counts, known fields, and precomputed facts can be returned without an AI model.
- 2 **Exact local calculation.** Totals, counts, and averages that map to a safe query are worked out exactly on your computer, using every row, and nothing is sent to your AI provider. When a question does not map cleanly, RangerIO passes it to your model instead of guessing.
- 3 **Model-assisted answer.** Interpretive questions use your model with context from the files. Smaller tables are sent in full; very large ones as a sample, and the answer says so.

The user does not need to choose the internal route.

Example direct questions

- “How many rows are in this source?”
- “Which columns are present?”
- “What is the inferred type of Invoice Date?”

Example analytical questions

- “Show the top 10 vendors by total amount.”
- “Count transactions by status.”
- “List records over \$10,000.”
- “Compare missingness in Q3 and Q4.”

Example interpretive questions

- “Summarize the main patterns in this source.”
- “What should I review first?”
- “Explain the largest differences between these two files.”
- “Draft a cautious summary that includes the data-quality caveats.”

Multiple sources

When selecting multiple sources:

- confirm that each source is ready;
- ask a question that identifies the desired comparison or join;
- review source attribution;
- reduce the number of sources if coverage or model limits become an issue; and
- do not assume all cross-source relationships are valid without checking keys and source structure.

Evidence and honesty signals

An answer can display:

- the files used;
- a coverage notice when only part of a table or document was used;
- how the answer was produced, for example **Answered using masked data**;
- a grounding badge for checked figures, for example **Numbers traced to your data**;
- data-quality caveats;
- a request to clarify when the question is unclear;
- a warning when a response was shortened;
- **Cloud request**, listing what was replaced before sending, with **View lineage**; and
- **Validate metrics**, an optional review that asks the model to check the answer against your data. Its status, the model used, and the time taken show in the bar at the bottom.

If a cloud request is blocked for privacy

- 1 Read the message. It lists the blocked files.
- 2 Choose **Re-scan now**.
- 3 Wait for **Scan complete**, then ask again.
- 4 If the scan finds problems, review and correct them first.
- 5 If RangerIO warns before sending that a file is not ready, you can choose **Ask without it**.

If an answer is incomplete

- ask a narrower question;
- select fewer sources;
- name the exact field or operation;
- review missingness and caveats;
- refresh or re-scan the source;
- add a related source only when it is genuinely needed; or
- calculate consequential numbers independently.

Limits of external-client answers

RangerIO can validate and sanitize its tool output. It cannot control or validate new prose that an external AI host generates after receiving that tool output.

Chapter 10 · Connecting to Your AI Tools

What the connector is

RangerIO includes a local MCP server that lets an MCP-compatible host request approved RangerIO tools. The server is **off by default**.

The guided setup is for Claude Desktop. Other MCP hosts may work but are not covered by this guide.

Enable the MCP server

- 1 Open **Settings**.
- 2 Under **General**, expand **Claude Desktop Connector**.
- 3 Turn on **MCP server**.
- 4 Wait for the local connector service to start.
- 5 Choose + **Add Connector**.

Configure Claude Desktop

- 1 Make sure Claude Desktop is installed.
- 2 If RangerIO says **Claude Desktop is open**, close it.
- 3 Read “We'll add RangerIO to your Claude Desktop connectors” and choose **Connect**.
- 4 Open Claude Desktop and pick **RangerIO** from the connector menu.
- 5 Return to RangerIO and confirm the status changes to **Connected** after activity.

Connector status

- **Not configured:** the connector has not been installed.
- **Pending:** configured and waiting for the first connection.
- **Connected:** RangerIO has observed connector activity.
- **Idle:** configured without current activity.
- **Error:** the sidecar or configuration needs attention.

Use **Re-test** to check the connector and **Disconnect** to remove the configured connection.

Trust settings

The MCP connector is conservative by default.

Allow tokenized export from stdio agents is off by default. When off, an agent cannot request a tokenized project bundle outside RangerIO's data directory.

When enabled:

- a trusted agent can request that export;
- detected PII is replaced with tokenized/synthetic equivalents;
- the resulting bundle still leaves RangerIO's data directory; and
- the setting should be disabled when it is no longer needed.

Disable the server

Turning the MCP server off stops the local connector service and unregisters the managed desktop connection. Use this when connector access is not required.

Chapter 11 · Reports & Exports

A. Export standard data

- 1 Open **Report** from the  menu.
- 2 Choose the file under **Select Source**.
- 3 Leave **AI Assisted** off.
- 4 Review the preview.
- 5 Choose a **Format**.
- 6 Choose an option under **Select Masking Strategy**.
- 7 Choose **Export** followed by the format, for example **Export CSV**.
- 8 Follow progress, then choose **Download**.
- 9 Inspect the exported content before sharing it.

Format availability varies by file; the formats shown for the selected file are the ones available.

Selected file or workflow	Formats offered
Table or spreadsheet	CSV, Excel, JSON, Parquet, TXT
Document	TXT, DOCX
AI-assisted report	DOCX, PDF, XLSX

B. Generate an AI-assisted report

- 1 Make sure each file has an Intelligence Dossier (see **The Dossier**).
- 2 Open **Report**.
- 3 Choose one or more ready files under **Select Source**.
- 4 Turn on **AI Assisted**.
- 5 Enter a **Report Title**.
- 6 Choose a **Report Template**, or keep **Custom** and pick the **Sections**.
- 7 Choose the **Output Format**: DOCX, PDF, or XLSX.
- 8 Under **Report masking**, choose **Mask sensitive values** or **No masking — original values**. This applies to the downloaded report; requests to your AI provider stay protected either way.
- 9 Add **Discoveries** when offered.
- 10 Choose **Generate Report** (for a custom report, **Generate DOCX Report**, **Generate PDF Report**, or **Generate XLSX Report**). The report runs in the background.

- 11 For a template report, review the sections and their variants, then choose **Download** for all variants or **Export Report** for the selected ones. A custom report opens a save dialog when it is ready.
- 12 Read any **Report notes** and sections that could not be generated.

A generated report must still be reviewed for factual and professional accuracy.

If a report cannot be generated

If a file's Intelligence Dossier is not ready, generation stops with a message that names the file and what to do, and offers **Retry**. If some sections could not be generated, the report says so and lists **Report notes** in the document.

C. Export a masked document from the preview

- 1 On **Explore**, select the document, turn on **Preview**, and choose **Open full preview**.
- 2 Complete the **PII review** and **Apply** your decisions (see **Privacy Scan**).
- 3 Wait for **Scanned [time] · export ready**.
- 4 Choose **Export masked**. RangerIO saves a masked DOCX copy.
- 5 Follow progress.
- 6 Open the exported copy and inspect every sensitive section before sharing it.

D. Export from PII Inventory

For supported sources:

- 1 Open **PII Inventory** from the  menu.
- 2 Expand the file.
- 3 Review findings by type.
- 4 Tick the PII types to mask.
- 5 Choose the **Strategy: Partial, Redact, or Hash**.
- 6 Choose **Export Masked**.
- 7 Review the resulting copy.

Masking strategies

Select Masking Strategy on the **Report** page offers these options when the selected file supports them. When the file has critical findings, **Partial** is selected by default.

- **None — export as-is:** does not mask detected PII. Do not use it for a sensitive file you will share.
- **Partial:** shows the first and last characters, for example J***n.
- **Full:** replaces all characters of the detected value.

- **Redact:** removes the detected value from the output.
- **Tokenize:** replaces detected values with stand-ins that cannot be reversed.

Availability and behavior depend on the source and export path.

Audit trail

RangerIO records local audit information for protected operations, including connector and export activity where supported. Audit records are technical evidence of what RangerIO recorded; they are not a regulatory certification.

Chapter 12 · Use-Case Workflows

Check a document and create a masked review copy

Use when: you have a PDF or DOCX file that may contain names, contact details, account information, or other identifiers, and you need a copy with reviewed findings hidden.

Before you start: confirm that the document is a supported format and that you are authorized to process it.

- 1 Import the document on **My Data** and move it to the correct project.
- 2 Select the document and choose **Preview Data** to open **PII review**.
- 3 Start with **Needs Your Attention**.
- 4 Use **Next match** and **Previous match** to inspect every occurrence.
- 5 Choose **PII**, **Not PII**, or **Change type** where appropriate.
- 6 Select any missed sensitive text and choose **Mark Selection as PII**.
- 7 Choose **Review decisions**, then **Apply**, and wait for the rescan to finish.
- 8 On **Explore**, select the document, turn on **Preview**, choose **Open full preview**, and wait for **Scanned [time] · export ready**.
- 9 Choose **Export masked**.
- 10 Open the exported copy and compare sensitive passages with the original.

Done when: the exported copy opens successfully and you have manually checked the areas that contain or could contain sensitive information.

If it fails: if the preview says **Masked export paused until re-scan**, choose **Re-scan now**. If the rescan does not finish, your decisions are kept; run the scan again, and do not use an older masked export as if it included the new decisions.

Ask a spreadsheet question and verify the answer

Use when: you need a total, count, ranked list, exception list, or period breakdown from a spreadsheet without writing a formula or query.

Before you start: identify the sheet, fields, filters, and time period that should be included. A precise business question produces a result that is easier to verify.

- 1 Import the workbook.
- 2 Open it on **My Data** and review each sheet's columns, missing values, and warnings.
- 3 Choose **Chat** beside the file, then **Yes** to open **Explore**.
- 4 Ask one precise question, such as "Show the top 10 vendors by total Amount for Invoice Date in Q4."
- 5 Type @ to pick a column when the name could be ambiguous.
- 6 Review the files used, grounding, coverage, and caveats shown with the answer.
- 7 For important figures, choose **Validate metrics** under the answer.
- 8 Compare material totals or example rows with the original workbook.
- 9 Ask a narrower follow-up if the answer mixed periods, sheets, or columns.

Done when: the question names the intended scope, RangerIO attributes the result to the correct source, and consequential figures have been checked against the file.

If it fails: name the exact sheet and column, use the words in your file, reduce the number of requested operations, and read any missing-value or structure warnings.

Compare two files or reporting periods

Use when: you want to compare last month with this month, an original file with a revised file, or two related source exports.

Before you start: confirm that the files represent comparable populations and that their key fields have the same meaning. RangerIO cannot decide that two similarly named columns are business-equivalent.

- 1 Put both files in the same project when they belong to the same engagement.
- 2 Import both files.
- 3 Review each file's structure separately and wait for both green ticks.
- 4 Select both ready files on **My Data**.
- 5 Choose **Use in Chat**.
- 6 Ask for a named comparison, for example: "Compare total Amount and record count by Status between March and April."
- 7 Review which source each number came from.
- 8 Read missingness and data-quality caveats for both files.
- 9 Verify the largest differences against the underlying records.

Done when: the answer identifies both sources, uses the intended fields, and the significant differences have been checked.

If it fails: compare one measure at a time, state which source represents which period, and avoid asking for a join until you have confirmed a reliable matching key.

Prepare a draft report for professional review

Use when: you need a structured first draft based on one or more imported sources.

Before you start: finish reviewing the files, confirm privacy coverage, connect a permitted AI model, and build an Intelligence Dossier for each file. A report is model-assisted and must be reviewed.

- 1 Open **Report**.
- 2 Choose the files under **Select Source**.
- 3 Turn on **AI Assisted**.
- 4 Enter a clear **Report Title**.
- 5 Choose the **Report Template** or **Sections** that fit the assignment.
- 6 Choose DOCX, PDF, or XLSX as the **Output Format**.
- 7 Choose **Report masking**.
- 8 Choose **Generate Report**.
- 9 Continue other work while the report runs in the background.
- 10 Open the completed report preview, or save the custom report.
- 11 Read every section and any **Report notes**.
- 12 For a template report, choose the variants, then **Download** or **Export Report**.
- 13 Check material facts, calculations, names, and conclusions before circulation.

Done when: the report file opens, its source scope is correct, its warnings have been addressed, and a qualified person has reviewed the draft.

If it fails: confirm the files are ready and have an Intelligence Dossier, confirm a model is connected, reduce the number of files, or remove a section whose inputs are unavailable.

Review a folder of engagement files

Use when: a client or colleague gave you a folder of files and you want to choose which ones belong in RangerIO.

Before you start: remove known backups or irrelevant files, or plan to deselect them after the scan. RangerIO imports the folder into its own project and keeps watching it for new and changed files. Subfolders are not imported.

- 1 Open **My Data**.
- 2 Choose **+ Add Files → Cowork Folder**.
- 3 Choose the folder.
- 4 Select the **File Types** to include.
- 5 Choose **Scan Directory**.
- 6 Review the discovered files.
- 7 Deselect backups, duplicates, irrelevant files, and unsupported material.
- 8 Choose **Continue with N files**.
- 9 Follow the import tasks in **Events**.
- 10 Review failed or interrupted files separately rather than assuming the entire group succeeded.
- 11 When the folder changes later, open its project and choose **Review & Sync** or **Dismiss**.

Done when: the intended files appear in the folder's **Directory:** project and each file's task state is accounted for.

If it fails: narrow the file types, import a problem file on its own with **Files**, and use **Resume** or **Dismiss** on the **My Data** banner for interrupted imports.

Use RangerIO from Claude Desktop

Use when: you prefer to ask questions in Claude Desktop and want Claude to request approved RangerIO information without manually attaching the original source file.

Before you start: import and review the file in RangerIO, turn on the MCP server, and complete the guided Claude Desktop connector setup. Your organization must permit the external host.

- 1 Confirm the source is ready in RangerIO.
- 2 Review its privacy state.
- 3 Confirm the connector is **Connected** in **Settings** → **Claude Desktop Connector**.
- 4 Open Claude Desktop.
- 5 Ask a question that clearly identifies the RangerIO project or source needed.
- 6 Review any tool request shown by the host.
- 7 Review the returned answer knowing that RangerIO sanitizes its tool output but does not control Claude Desktop's final prose or retention behavior.

Done when: Claude Desktop used the intended RangerIO source and the returned result contains only information you are permitted to place in that external conversation.

If it fails: return to RangerIO, choose **Re-test**, confirm the MCP server is on, and relaunch Claude Desktop if the connector wizard instructs you to do so.

Individual professional workflow

- 1 Create a project for the engagement.
- 2 Import supported files.
- 3 Review structure warnings.
- 4 Review the privacy scan.
- 5 Correct findings and apply your decisions.
- 6 Ask local or model-assisted questions.
- 7 Verify important answers.
- 8 Export the dossier, report, or masked copy appropriate to the task.
- 9 Retain or delete the local project according to organizational policy.

Accounting and audit workflow

- 1 Create a project for the client and period.
- 2 Import ledgers, workbooks, and supporting documents in verified formats.
- 3 Review sheet structure and header confidence.
- 4 Inspect missingness, unusual values, and parsing warnings.
- 5 Review PII across the sources.
- 6 Ask precise questions such as totals, counts, status breakdowns, and exceptions.
- 7 Read grounding and data-quality caveats.
- 8 Verify material totals against the underlying rows.
- 9 Export data or a report for professional review.

Legal and HR workflow

- 1 Create a project for the matter or HR review.
- 2 Import supported DOCX, PDF, spreadsheet, or text sources.
- 3 Open each document with **Preview Data**.
- 4 Start with **Needs Your Attention** in **PII review**.
- 5 Step through each occurrence.
- 6 Choose **PII**, **Not PII**, or **Change type**, and mark missed values.
- 7 **Apply**, and wait for the rescan to finish.
- 8 Ask questions with a permitted AI provider or your own local model server.
- 9 Export a masked copy only after reviewing the export itself.

Healthcare and research workflow

- 1 Create a project for the authorized dataset or study workflow.
- 2 Import supported tabular exports and documents.
- 3 Review the source's structure and scan coverage.
- 4 Inspect names, dates, identifiers, locations, and other potential re-identification fields.
- 5 Mark specialized identifiers that the scan missed. Health terms inside documents are not detected in this release: mark them yourself or leave those files out.
- 6 **Apply**, and wait for the rescan to finish.
- 7 Where policy requires it, use exact totals and counts (worked out on your computer) or your own local model server.
- 8 If an external provider is permitted, review sanitized context and organizational rules first.
- 9 Verify all analytical results against the source.

Repeated engagement workflow

- 1 Start a new project for the new period or client.
- 2 Import the current files.
- 3 Confirm that source structure has not changed unexpectedly.
- 4 Review new scan findings rather than assuming prior classifications apply.
- 5 Ask repeatable questions using the same field names and operations.
- 6 Export the current output and record the review.

Chapter 13 · What RangerIO Doesn't Do

It is not a regulatory certification

RangerIO provides technical controls and review tools. It does not make an organization HIPAA-compliant, SOC 2-certified, PCI DSS-compliant, or legally authorized to disclose data.

It is not DLP or endpoint security

RangerIO does not monitor all filesystem, email, browser, or network activity and does not enforce organization-wide data-loss policy.

It is not a cloud warehouse

The desktop app is workstation-oriented. It is not a centralized multi-user warehouse or enterprise identity/access-control plane.

It is not a real-time streaming pipeline

RangerIO imports files and can watch a Cowork Folder for new and changed files, which you review before they are applied. It is not a continuous event-stream pipeline.

It does not import from databases or REST APIs

The desktop app imports files and folders. Database and REST API sources are part of RangerIO Enterprise licensing, with full support and centralized management, and are expected to begin shipping in Q3 2027. For more information on our Enterprise version, please contact info@rangerio.com.

It does not certify a source as free of PII

A complete scan can miss values or produce false positives. A user must review the source and output.

It does not guarantee correct AI output

Models can be incorrect, incomplete, or misleading. Grounding and caveat controls reduce risk but do not eliminate it.

It does not control an external AI host

RangerIO controls its own tool response and egress checks. The host's final prose, retention, account settings, and provider behavior remain outside RangerIO.

It does not include an AI model

Questions, the Intelligence Dossier, and AI-assisted reports use the cloud provider or local model server you connect.

Chapter 14 · Technical FAQ

Which release does this guide cover?

Current release (0.16.0), reviewed October 17, 2026. Unverified and roadmap features are omitted.

What is the architecture?

RangerIO is an Electron desktop application with a local backend. The backend handles ingestion, profiling, PII processing, analytics, retrieval, exports, connector tools, and calls to your AI provider.

Current technical components include:

- FastAPI for the local backend;
- DuckDB for staging and analytical queries;
- SQLite for application metadata and registries;
- LanceDB for vector retrieval;
- fastembed through ONNX for embeddings; and
- GLiNER2 through ONNX for entity extraction.

Ports and internal process topology are implementation details and should not be presented as stable user configuration.

Can RangerIO use a local model?

Yes, if you run your own local model server, such as Ollama, LM Studio, vLLM, or AnythingLLM, and connect it in **Models** → **Llama / vLLM / Network Model**. RangerIO does not include a model.

What runs locally?

On your computer:

- file and folder import;
- structure analysis;
- privacy scanning, review data, and corrections;
- indexing and retrieval;
- exact totals, counts, and averages;
- dossier storage;
- standard exports; and
- the local connector service.

Requests to your cloud AI provider and Claude Desktop conversations are external paths.

Where is working data stored?

The desktop engine uses a local RangerIO data directory under the signed-in user's home directory. The current working directory is named `RangerIO_local`. It can contain imported working data, metadata, indexes, registries, logs, and audit data.

Which formats are verified for import?

CSV, TSV, XLS, XLSX, XLSM, JSON, JSONL, Parquet, PDF, DOCX, PPTX, HTML, and TXT.

Legacy .doc and .rtf open only if Pandoc is installed on your computer; otherwise save them as DOCX first.

Why can a file be imported while it says PII scan required?

Import and the privacy scan are separate stages. A file can exist while its scan is still pending, interrupted, or incomplete. Questions wait until the scan is done.

What does “unverified” mean?

The displayed findings may be valid, but the scan did not finish fully. Rescan the file and review it.

Do corrections survive a re-scan?

Yes. Your decisions are saved when you choose **Apply**, and the rescan that follows uses them.

Why does masked export pause after a correction?

The masked export relies on the last complete scan. The rescan that follows **Apply** reconciles the whole document and unlocks export. If the preview shows **Masked export paused until re-scan**, choose **Re-scan now**.

Is a masked export guaranteed safe?

No. Registry-first masking and backstop checks reduce risk, but the user must inspect the exported file.

How are questions answered?

RangerIO answers exact totals, counts, and averages with a local query using every row, and sends interpretive questions to your model with context from the files. When a question does not map cleanly to a query, it goes to the model instead of being guessed.

What do grounding and caveats mean?

Grounding shows that supported claims—especially numeric claims—were checked against source/query facts. Caveats describe relevant data limitations such as missingness or anomaly coverage.

Neither is a guarantee that the entire answer is correct.

Can an in-app cloud answer restore masked values?

Yes. RangerIO puts supported values back locally after the provider response returns. Under the answer, **Cloud request** lists what was replaced before sending.

Can an external MCP client receive restored sensitive values?

RangerIO's public MCP surface blocks rehydration because tool results return through the external host.

Does MCP expose raw rows?

MCP tools can return sanitized previews and sanitized analytical results. PII-aware scrubbing and outbound controls apply, but the result still requires review.

Is the MCP server always running?

No. It is off by default and starts when you turn on **MCP server** in **Settings**, or the Claude Desktop option during first-run setup.

Does RangerIO send telemetry?

MCP telemetry is off by default. When explicitly enabled, it writes to a local rotating log. Update checks on macOS, and any AI provider or connector you enable, create network traffic for those functions. No usage data is sent to RangerIO servers.

Does RangerIO require the internet?

Local ingestion and local analysis can operate without a cloud provider. Downloads, updates, configured cloud providers, and external connector hosts require connectivity.

Does RangerIO check for updates?

On macOS, yes: it checks and asks before downloading. On Windows, automatic updates are off; **Check for Updates** in the  menu links to the latest signed installer.

What happens when an import is interrupted?

The banner on **My Data** offers **Resume** and **Dismiss**. Resume is user-driven and can be refused when the system lacks sufficient memory.

Does deleting a source delete the original file?

No. It deletes RangerIO's working records for that source, not the original external file.

Are database, API, or SharePoint imports supported?

Not in the current desktop product. SQL database and REST API sources are part of **RangerIO Enterprise** licensing, with full support, and are expected to begin shipping in Q3 2027. SharePoint is not a documented source. For more information on our Enterprise version, please contact info@rangerio.com.

Is there an enterprise or centrally managed version?

Yes. RangerIO Enterprise is a licensing tier over the same data engine. It adds centralized management for multi-seat deployments and fully supported database and REST API data sources in addition to file imports. Enterprise licensing is expected to begin shipping in Q3 2027 and is available on request today.

For more information on our Enterprise version, please contact info@rangerio.com.

Is RangerIO a compliance program?

No. It is one technical component in a broader organizational process.