

Workflows

Day-to-day operation of NexChrom: master data, CDS import, methods, calibration, execution, review and certificates.

Qualimetric · Nijmegen, The Netherlands · info@qualimetric.eu · nexchrom.nl · Version 2.36.0 · 22 September 2026

CONTENTS

- | | |
|--|--|
| 1. Scope | 9. Workflow C — Calibration |
| 2. Introduction | 10. Workflow D — External Certificate Import (PDF) |
| 3. Page Orientation | 11. Dashboard |
| 4. Getting Started | 12. Maintenance |
| 5. Master Data Setup | 13. User Management |
| 6. Workflow A — Analysis via Sample | 14. E-mail Alert Notifications |
| 7. Workflow B — Analysis via CDS File | 15. Upload Instrument Documentation |
| 8. Inspecting CDS Files Without Linking (CDS File Browser) | 16. AI Assistant |
| | 17. Troubleshooting and Common Issues |

Scope

This user manual describes how to perform the day-to-day workflows supported by NexChrom as it currently exists. It guides a new user through master-data setup, CDS file import, method configuration, calibration, execution, result review, and certificate generation. It targets the application aligned with database schema DDL v2.9.0.

Introduction

NexChrom is a Streamlit-based application for managing gas chromatography workflows. It runs locally against a SQLite database. All screens are reached from the left-hand navigation once the application is started. Two primary analysis workflows are supported:

- Analysis via Sample — you start from an existing Sample, then link one or more CDS files and run the calculation.
- Analysis via CDS File — you start from CDS files on disk, and a Sample is created for you during import.

Page Orientation

Archiving instead of deleting

- Every master-data table (locations, CDS file formats, users, samples, sample types, instruments, methods, calibration models, certificates, components, aliases) shows **active records only** by default.
- Set **Archive status** to *Archived* to see the **archived records only** (never both at once); set it back to *Active* to return to the active ones. The **Archived At** column (the archive moment, on the tables that record it) appears

in the *Archived* view only; in the *Active* view it would always be empty, so it is hidden there. The filter looks the same as the type filters and is always the rightmost control above the table. A table with extra columns hidden by default (Components: atom counts) has a **Columns** filter of the same kind directly left of it (*Standard / Atom counts*). On Samples the supplier, parent sample and number of linked CDS files, and on Instruments the data folder, CDS format, location, file name filter, documentation folders, notes and created-at are not grid columns at all: tick one instrument and edit them in the **Details of <instrument>** table under the grid — double-click a value (Location is a drop-down of the active locations, CDS format one of the active file formats), the edit turns yellow and is saved with *Save* of the grid or dropped with *Discard*.

- Select one or more rows and press **Archive**: the records get `is_current = 0` (and an archive moment where the table records it). Nothing is deleted — analyses, calibration models, runs and files linked to the record stay where they are, and the record still resolves wherever it is referenced. In the archived view, **Unarchive** reverses it.
- Archived records are not offered in the selection lists of the other pages (instrument on Calibration and Analysis, methods on Execute). On *Analysis via Sample* the sample list has an **Archive status** filter at the far right of the filter row (sample-type filter left, link status and archive status right): *Active* (default) or *Archived*; switching it reloads the sample list and clears the sample and file selection, but leaves the instrument and detector template chosen in section 1 as they are. Historical calibration models stay selectable on Execute for recalculation. CDS files and Analysis records are measurement records and cannot be archived.
- Both actions are written to the Change history. Use **Delete** only for records that were entered by mistake; records referenced by other data cannot be deleted anyway.

Every main CRUD page (Administration, Sample, Analysis, Method, Calibration, Execute, Review, Certificates, Component) exposes an '**About this page**' expander directly below the page title. Expanding it shows a short info banner explaining what the page does and a workflow diagram that highlights the current page within the full NexChrom flow (Component → Sample → CDS files → Certificate → Method → Calibration → Analysis → Execute → Review → Analysis certificate). Some pages (Certificate, Component, Calibration, Administration) only show the part of the flow they participate in. Use this expander whenever you want to confirm where a page sits in the overall workflow before you start entering data.

Filters and buttons

- No filter offers an *All* option: with nothing selected the table shows every record, and clicking the selected option again clears the filter. Category filters (sample type, location, certificate type) let you select several values together; two-way filters (Linked / Not linked, In use / Not in use, Matched / Unmatched, Named / Unnamed, Released / Unreleased, the Review status filters) take one value at a time.
- The controls above a table are **segmented filters** on one row: a type or location filter (*All* plus the values) at the far left, and at the far right *Columns* (*Standard* / the extra columns that are hidden by default, e.g. *Atom counts* on Components) followed by *Archive status* (*Active* / *Archived*), which is always the rightmost control. Clicking the selected option again clears it, which counts as the default.
- A filter that has nothing to switch is shown **disabled** with an explaining help text instead of being hidden — for example the *Peaks* filter of a detector without unidentified peaks in the CDS file browser.
- Wide tables (Users, Certificates, Calculation results on Execute) keep a readable width per column and **scroll horizontally**; the Select column and the identifying column (username, certificate number) stay pinned on the left while you scroll. The registration fields of a certificate (approval status, approved at/by, run, PDF, PDF hash) are read-only: the app fills them when a certificate is generated and approved.
- **Sorting**: click a column header to sort the table (a second click reverses, a third clears). The order stays while you tick rows or save; the date columns of the Certificates table (Issued At, Valid Until, Approved At, Created At, Archived At) sort by date, empty values last.

- **Action buttons** have a material icon, a short label and one fixed width, and sit on one row below the table: **Save** (blue), **Create** (green), **Delete** (red), **Discard** (grey) and **Archive / Unarchive** (amber). Other pages use the same colours: blue for the main action (*Execute, Approve, Link, Register*), green for actions that create something (*Add, Duplicate, Generate certificate*), red for *Reject*, grey for utilities (*Clear, Rename, Restore, Test*). Hover over a button for the detail its short label leaves out.

Feedback and messages

Save, Create, Delete and Discard actions display their success or error message as a **full-width banner** directly below the page title (or directly above the grid). The message is preserved across the automatic page refresh that follows an action, so confirmations no longer flash away. If a save is partially successful, the banner lists every row that was skipped — both new rows that never reached the database and edited rows whose required fields were blanked out — with the sample / component / certificate name for easy identification. Newly created rows now appear at the **top** of their grid, so they are always visible without scrolling.

Getting Started

Launch the Application

- Open a terminal in the project root.
- Activate the virtual environment: `.venv\Scripts\activate` on Windows or `source .venv/bin/activate` on macOS / Linux.
- From the project root, run: `streamlit run app/main.py`.
- Streamlit opens the application in the default browser.

Select the Active Database

- Open the Settings page from the sidebar.
- In the Database tab, enter or browse to the SQLite database file (for example `data_local/database/nexchrom_db_v2.10.0.db`).
- Save the setting. The Dashboard displays the currently connected database at the top of the page.

Verify Master Data Availability

- On the Dashboard, check that Instruments, Components, and Sample Types are populated. Any missing master data blocks downstream workflows.
- If anything is missing, follow the relevant section in Master Data Setup below.

Master Data Setup

Register a Location and an Instrument

- Open the Administration page and create the Location on the Locations tab (name, address, building and other context).
- Open the **Instrument** page (directly below Administration in the sidebar) and create the instrument in the Instruments table. The *Filter on location* control at the top left of the table (*All* plus the Location names from Administration) narrows the table to one location; new rows always stay visible. Select the Location and the CDS File Format that matches how the instrument exports its data.
- Give every physical instrument its **own data folder** (the working convention is `data_local/gc_data/<name>-<manufacturer>-<model>-<serial>/`) and point **data_source_path** at it; the folder name is a human-readable

snapshot, the `data_source_path` link is authoritative. As a fallback for analyzers that must share one folder, **file_name_filter** restricts the listing to file names containing a substring (e.g. the serial number); leave it empty otherwise.

- On the Administration page, File formats tab, verify that the format you need (Agilent ChemStation, Scion CompassCDS, or Thermo Chromeleon*) is present. Add one if needed.

* *Planned: a dedicated Thermo Chromeleon parser is not yet implemented in the current version; the format can already be registered here.*

Record the Scope and the QC Limits (first version)

The **Scope** tab on the Administration page holds the official accreditation scope — the registration header and the CMC lines per component, as on the RvA annex to the declaration of accreditation — next to the internal requirements. Entries are recorded, not yet checked against analyses. Work top-down through its four sections:

- **Step 2 — Enter the CMC lines per section.** Open the expander *Calibration and measurement capabilities* (CMC); it shows the lines of the scope ticked in step 1 (its label — registration · organisation · validity, a status other than active in brackets — stands above the table). Optionally choose the **Section** (*All sections* or one HCS code with its section, e.g. `RM 2 1 Natural gas`); the table reloads for that choice and lists the lines in annex order. *Create* adds a line pre-filled with the selected scope, the chosen section (else the section of the last line), the remark of that section, the location code of the header, the units `cmol/mol` and `%` and the next *Order* number. Per line enter the **Name on annex** (required, exactly as printed), the **Component** it maps to (leave it empty for a group line such as *Heptanes*, *Octanes* or a co-eluting pair, and name the constituents in *Notes*), *Range low* / *Range high* with *Range unit*, *CMC low* (the relative CMC at the low end, or the single value of the line) and *CMC high* (the relative CMC at the high end; leave it empty when the CMC is constant over the range) with *CMC unit*, *Remark* (e.g. `Analysis to ISO 6974`), *Loc.* (only when a line deviates from the header code) and *Order*. *Save* stores the lines. Without a ticked scope the section only shows *Select an accreditation scope in the table above* (*Select checkbox*).
- **Step 3 — Check the Annex view.** Open the expander *Annex view*: it renders the ticked scope as on the RvA annex — the header block (annex to declaration of accreditation, normative document, registration number, organisation, laboratory with code, validity, replaced annex, status with date, register link) and a read-only table *HCS code* / *Measured quantity* / *Range* / *CMC* / *Remarks* / *Location* with one section row per section and one row per active line in order; ranges read as `(0.03 - 0.1) cmol/mol`, CMCs as `12.5 %` or `0.45 % - 0.1 %`, the header notes appear below the table. Compare it with the annex; nothing is edited here — correct the lines in step 2.
- **Step 4 — Record the internal requirements.** Open the expander *Requirements and scope limits*. Every row is one requirement or scope limit for a **Laboratory** (required), optionally narrowed to one **Instrument** and one **Component**: a *Requirement* text, a *Lower limit* / *Upper limit* with *Unit*, a *Reference* (standard or accreditation document) and *Notes*. Use it for what the laboratory is set up to measure beyond — or more strictly than — the official scope.
- Planned: checks of analyses against the scope on Execute / Review, environmental requirements, editing restricted to managers / admins and enforcement of one active scope per laboratory.
- **QC limits from the history of a QC sample.** A QC sample is a cylinder of type *Quality Control* (or *Reference*) that you analyse repeatedly on an instrument — like any sample: link its CDS files, Execute, have the run approved on Review. Once it has been analysed for a while under conditions you judge normal and good, open the **Quality control** page, tab **Control settings**, block *Limits from QC history or certificate*: choose the **Instrument** (only instruments with QC analyses are listed, each with its number of QC samples) and then the **QC sample** analysed on it (with its number of analyses there; a QC sample without analyses is named below the drop-downs but cannot be chosen), basis **History**, the period (default: from the first approved run until today) and, if needed, runs to **exclude** (a known incident — the runs outside the preview bands are listed

below the table). The table shows per component n , mean, s , RSD, min / max and the bands: warning = mean $\pm 2s$, action = mean $\pm 3s$ (Settings). A green banner means the set is **established** (at least 3 months and 13 runs by default); an orange one means it is below the minimum — you can still save it as **provisional** after ticking the confirmation, and replace it later. **Save as limits** (manager / admin) writes the limits and archives the previous set.

- **QC limits from the certificate.** For a QC sample with a certificate (external QC gas, Workflow D) choose basis **Certificate**: every certified component gets warning = value $\pm 1 \cdot U$ and action = value $\pm 2 \cdot U$ (Settings). This is also the fallback when there is no history yet.
- **QC limits by hand.** Tab **QC limits**: Create a row with basis *manual* — *Component*, optional *Instrument* and *QC sample*, *Action low* / *Action high* (and optionally the warning bounds), *Unit*, *Notes* — and *Save*. Rows written by a derivation carry read-only provenance (centre, s , n , period, set by / at, provisional); re-derive instead of editing them. *Archive* takes a limit out of force.
- **Check a QC run.** Execute a QC sample like any other sample. Under the results the expander **QC comparison** shows every component against the limits in force — ● pass, ● warning, ● fail, ● no limit — with the overall verdict; a *preview* is evaluated live and stores nothing, a stored run keeps its verdict (the stored-run message names it) and shows it on the Review page: the QC column in the runs table and the *QC comparison* block in the run detail. A run stored before limits existed has no verdict yet — a reviewer clicks **Run QC check now**; **Re-run QC check** repeats it against the limits in force now. **A failed QC run cannot be approved**: reject it and investigate the analyzer; a warning may be approved with a note. The Dashboard lists failed QC runs of the last 30 days and pending runs with a warning.
- **Westgard rules.** Tab **Control settings**, *Westgard rules*: tick the multirules you apply — 1_2s (warning), 1_3s, 2_2s, R_4s, 4_1s and the N-in-a-row rule (set N, e.g. 8, 10 or 12) — and *Save*. From then on every QC check and the control chart judge a result together with the preceding runs of the same QC sample; the violated rule codes appear in the *Rules* column on Execute and Review and in the chart's hover text. Stored verdicts keep their rules; *Re-run QC check* on Review applies a new selection.
- **Two kinds of QC sample.** An **external QC gas** is a sample of type *Quality Control* with the supplier's certificate imported or registered on the Certificates page (origin external). An **internal QC gas** is a cylinder you analysed repeatedly as an ordinary sample: tab **Control settings**, *Promote a sample to QC* — choose the sample, instrument and period (approved runs only; at least the minimum number of runs of Settings is advised), name the new QC sample and click *Promote to QC*. NexChrom creates the QC sample (parent = the original, same cylinder), an internal QC certificate with per component the mean and $U = k \cdot s$ of those runs (approval *draft* — approve it on Review > Certificates) and records the source runs. From now on link the cylinder's new CDS files to the new QC sample; its history already contains the parent's runs, so you can derive limits at once.
- **Keep an eye on validity.** Tab **Monitoring**, *Expiry monitoring*: what is over the date or approaching it — certificates (*valid until*), released calibration models (newest fit plus the validity months) and QC samples (their certificate) — with the level ● / ⚠ / 🔔 from the reminder windows you set under **Control settings > Reminders & warning** (months before the date per kind, and the model validity). The Dashboard repeats the urgent ones as alerts.
- **Derived limits.** Tab **Control settings**, *Derived limits*: switch the calibrated-range check, the GOF threshold, the expansion factor and the accreditation-scope check on or off and record when they are checked; their values are listed under *Monitoring > Control overview*. With the scope check on, the Review page compares every result with the CMC range of the active accreditation scope of the instrument's laboratory (informative).
- **Accreditation scope at Execute and on the certificate.** The scope check runs on Review by default; tick the moments *execute* and *certificate* under *Derived limits* to see the scope line under the result table on Execute as well and to print an *Accreditation scope* section (with a *reported without accreditation* note for results outside the CMC range) on the analysis certificate.

- **Watch the QC sample.** Tab **Monitoring**, *Control chart*: choose the instrument, then one of the QC samples analysed on it, then the component (the same order as on the other pages; a QC sample without analyses is named below the drop-downs but not selectable) — the results over time against the limits in force (centre, dashed warning band, solid action band; green / orange / red points; hollow = run not approved), the counts below the chart and the *Points* table. *Control overview* lists every limit in force.
- Both pages work like every other master-data table: *Archive status* filter, *Save / Create / Delete / Discard / Archive*, change history (the Annex view is read-only; archived CMC lines drop out of it).

Configure Detector Templates per Instrument

Some GC systems report the same component on more than one detector (for example methane on both a TCD and a FID channel). The detector template records, per instrument, which detector's value should be used, so the Include choice is pre-filled at every import and only exceptions need manual correction.

- Open the **Instrument** page (Instruments tab) and select exactly one instrument in the Instruments table (Select checkbox) — the Detector templates section below follows that selection; instrument documentation lives on the separate Documentation tab.
- An instrument can have multiple named templates (e.g. a different configuration for certain kinds of samples). Pick one in the **Template** selector; the buttons **Create template**, **Rename template**, **Make default** and **Delete template** manage them (the name field serves Create and Rename; the default template cannot be deleted and is the one preselected at Analysis import). An instrument without a template gets a default one named *Standaard* automatically at the first visit of this section or the first import; that creation is recorded in the Change history (DetectorTemplate, field *created*). Both import tabs offer a **Detector template** selector next to the instrument, so any template can be chosen per import.
- **Detectors table** — the registry of the instrument's detectors, with the literal CDS label, a short name and the standardized detector type (TCD / FID / other). Detectors register themselves automatically at the first import; the **Suggest from CDS data** button fills the table from analyses already imported and from the newest CDS files on disk. Assign a detector type where the warning flags one as untyped.
- **Preferred detector per component table** — one row per component naming its preferred detector. The **Suggest from CDS data** button adds one pending row per component seen in the instrument's CDS data; components seen on multiple detectors get one row per detector, and components that already have a row are checked against the data — when the data also shows another detector, the extra row is added and the conflict named in the message. Enable *Highlight duplicate component names*, disable **Include** on the wrong rows, then save — a component can keep only one preferred detector.
- Alternatively, resolve the duplicates once in an import grid and click **Save detector choices as instrument template** there; the template then fills itself from your actual choices.
- From then on, imports pre-fill Include on duplicate component rows automatically; rows from non-preferred detectors are stored with include = 0, so no measurement data is lost.

Maintain Components and Aliases

- Open the Component page.
- On the Components tab, create the chemical components you need to report on (for example methane, nitrogen, CO₂).
- On the **Global component aliases** tab, add every alternative name that might appear in your CDS files for the same component. Aliases are how NexChrom maps variant peak names back to the canonical component; aliases that hold for one instrument only are entered on the Instrument page (*Instrument component aliases*).
- One alias per component exists only once — either global or for one instrument. When you save an alias that already exists for the same component elsewhere, *Save* is refused with a message naming where it exists (*Alias*

X → component already exists as a global component alias / an alias of instrument Y. Remove it there first, then save it here ...); nothing is written until the conflict is resolved.

- On the **Component role catalog** tab, review the catalogue of role types (direct, indirect, fixed) a method can assign to a component; the assignment itself happens per method on the Method page.

Certificate composition

- The *Certificates* table has, next to *Filter on certificate type* (where several types can be selected together, e.g. calibration and qc), the filters **Approval** (Approved / Not approved), **PDF** (With PDF / Without PDF — is a source or frozen PDF on file?) and **Sample** (Linked / Not linked); none of them has an *All* option: nothing selected is no filter, and clicking the selected option again clears it. They combine and new rows always stay visible. *Without PDF* is the quick way to find the external certificates that still need their source document before they can be approved.
- For an external certificate the block **Source document** above the composition grid shows the PDF on file (with the expander *View the PDF on file* — inline viewer and download) or warns that none is on file; choose a PDF there and it is a pending change (yellow line, counted in the *Pending* caption, named in the logout dialog) that the **Save** of this section stores with its SHA-256; *Discard* drops it. Replacing the document of an approved certificate withdraws the approval.
- Select one certificate to see its components. The **Rel. Unc.** column shows the relative expanded uncertainty as a percentage with two significant digits (for example 0.30 %), the same presentation as *Rel. Unc.* on the Execute preview; it is derived from *value* and *uncertainty* and cannot be edited. A caption *Pending: +n / ~n / -n* under the grid counts the new, edited and deleted component rows that are not yet saved.

Maintain Sample Types

- Open the Sample page and select the Sample Types tab.
- Ensure the types you need exist (Sample, Calibration, QC, Reference).

Workflow A — Analysis via Sample

Use this workflow when you already know the sample you want to analyze and want to attach CDS files to it afterwards.

Step 1 — Create the Sample

- Open the Sample page and select the Samples tab. Tick a sample: the read-only *No. of linked CDS files* in the **Details of <sample>** table under the grid shows how many CDS files are linked to it (empty = none yet; the details table also holds the supplier and the parent sample, editable by double-click and saved with *Save*) and the *CDS files* filter (All / Linked / Not linked) next to the sample-type filter separates measured from not-yet-measured samples; linking itself happens on the Analysis page.
- Click Add and fill in the sample details: name, sample type, description, and any identifiers you track.
- Save the record.

Step 2 — Link CDS Files on the Analysis Page

- Open the Analysis page. The tabs are ordered **Analysis via Sample** → **Analysis via CDS File** → **CDS File Browser** (workflow tabs first, audit tab last).
- Select the **Analysis via Sample** tab.
- Section 1 **Select instrument** — pick the instrument that produced the CDS files. This selector sits at the top of the tab, separate from the sample and file sections.

- Section 2 **Select sample** — pick the sample. All samples are shown by default (samples are reusable: a calibration gas calibrates several instruments, and any sample can be re-analysed); the *Link status* column shows how often and on which instrument(s) a sample is already linked. Use the **sample-type filter** (segmented control with *All* plus the SampleType labels) to narrow by type, and the right-aligned **link-status control** (*Linked* | *All* | *Not linked*) to narrow by usage — *Not linked* quickly finds a freshly created sample. The file section has the same control with states *Linked* | *All* | *Available*, where *Available* (default) shows importable files: unlinked ones plus files already linked to the selected sample.
- Section 3 **Select CDS file(s)** — browse to or select one or more CDS files that correspond to injections of this sample. The file table has a **Folder** column next to the file name: the sub-folder under the instrument's data path (empty for files directly in it). Its columns keep a fixed width and the table scrolls horizontally.
- Two files with the same name in different sub-folders are two rows (the rows are keyed by the path under the data folder; the Folder column tells them apart).
- Confirm the import. NexChrom parses each file, creates an Analysis record, and stores the extracted peak responses as AnalysisResponse rows linked to the matched components.
- **Include choices are pending until Save.** Ticking or unticking *Include* in *Measurements from CDS file contents* changes nothing in the database yet: the caption *Pending: ~n include changes* under the table counts them, **Discard** drops them together with the unsaved file selection (*All pending file selection and Include changes discarded.*), and **Save** writes them to the saved measurements and applies them to the files imported in that same Save (on top of the detector template). The result message of Save or Discard appears directly above the Save button.
- A legend under the table explains the **Match** column: *name*, *trivial_name*, *formula*, *iso_code* (the CDS name equals that field of the component), *alias* (Component page), *instrument alias* (this instrument), ..., *suffix X stripped* (recognised after removing the detector suffix), *matched* (imported earlier without a recorded source), *variant of <component>* (a detector variant that is deliberately not imported) and *unmatched*.
- If a peak name does not match a known component or alias, the table *Measurements from CDS file contents* names it above the table. Add the name as an alias — on the Component page (*Global component aliases*) when it holds for every instrument, or on the Instrument page (*Instrument component aliases*) when it is specific to this instrument's method; an instrument alias wins over a global one — then re-select the file. A detector suffix such as *Methane (FID)* needs no alias. If the alias already exists for the same component — globally or on another instrument — Save is refused with a message naming where it exists (*Remove it there first*): one alias per component exists only once, global or for one instrument.

Step 3 — Choose or Define a Method

- Open the Method page. A method is instrument-independent — it meets an instrument only when a calibration model is built with it. The *In use on* column shows on which instruments a method is in use (number of active calibration models); the *Usage* filter (*All* / *In use* / *Not in use*) at the top left separates methods in use from methods still waiting for a calibration.
- Select an existing method that fits your analysis, or create a new one.
- For each component the method reports, assign a role: direct (calibrated), indirect (calculated via RRF), or fixed. For a fixed component enter the concentration in the **Fixed value** column (mol/mol) and, optionally, its expanded uncertainty in **Fixed U**; a fixed role cannot be saved without a value. Execute reports the value as entered (even when the sample shows no peak for it) and includes it in the normalised composition.
- For indirect components, define the RRF rule: reference component and the relative response factor.
- Editing a method that calibration models or calculation runs already reference asks for an explicit confirmation first: roles and RRF rules are read live at (re)calculation time, so in-place changes alter what recalculations of existing runs produce. For semantic changes, use **Duplicate method as new version** (copies the method with

all roles and RRF rules under a date-and-time-stamped name) and edit the copy. All role and RRF saves appear in the Change history.

Step 4 — Ensure a Calibration Model Is Released

- Open the Calibration page.
- Confirm that the instrument you used has a current calibration model covering the components in your method **and that its *Release* column reads *released*** — Execute only offers released models for new runs.
- If there is none, build one using Workflow C (Calibration) below and have it released on the Review page (Workflow C, step 4).

Step 5 — Execute

- Open the Execute page.
- Section 1 **Select Sample** — the list is filtered by a sample-type segmented control (left) and by the **Link status** filter (right). **With *Linked* (the default) the table shows only samples that already have at least one linked CDS file** (those are the samples that can actually feed Execute). Set it to *Not linked* for a read-only view of unlinked samples (the Select column is disabled there so you can inspect but not select them). There is also an `is_current` column on this table.
- Section 2 **Calibration model & method context**:
 - **2.1 Select Calibration Model** — the table is automatically restricted to Calibration Models on the **instrument of the selected sample's CDS files**; an info banner names that instrument. Only **released** models are listed; the *Release status* filter (*Released* / *All*, at the far right above the table) adds draft / rejected / withdrawn models — a run with such a model can only be stored with run type *recalc*, or with a QC / reference sample (the QC check of a new calibration). Any row that would violate the instrument scope is cleared from selection as a backstop; the validation in section 3 flags an instrument mismatch as a hard error and an unreleased model as a warning.
 - **2.2 Combined Method Context** — pick the method via the Select column.
 - **QC coverage** (ordinary samples only) — NexChrom looks for QC runs on the **same instrument** and proposes the newest **passed** one analysed within the QC coverage window (Settings, default 7 days) of your sample's analyses as the **validating QC run**. Keep it, pick another (older runs are listed too) or choose (*none*). A green line confirms the coverage; an orange or red line tells you the QC run is older than the window, has a warning, failed, or is missing — that is a notice, processing continues. The choice is stored with the run, shown in the validation summary, on the Review page and on the certificate.
- Section 3 **Review the validation summary**. All required components must have matching responses from the linked CDS files, and the calibration model's instrument must match the sample's CDS files.
- Section 4 **Execute and results** — organized as three top-level expanders:
 - Expander **Run type** — choose the run type and click Execute; NexChrom computes the results in place (no page jump) and, on saving, creates a CalculationRun with its AnalysisResult rows.
 - Expander **Calculation results** — the per-component table (Mol fraction and Uncertainty with 5 decimals, numeric ISO-code sorting) with the sum of unnormalized mol fractions below it. Uncertainties follow GUM propagation — calibration covariance for direct components, RRF relative standard uncertainties per detector type (ISO 6974-2 Annex B) for indirect components — and are expanded with the coverage factor *k* from Settings. The **Cal. range (mol/mol)** column shows the calibrated range of the fit — the lowest and highest certified concentration of the calibration gases used — and **Range status** tells whether the calculated result lies *within*, *below* or *above* it (indirect components show *ref-below* / *ref-above* when their reference component was extrapolated); flagged rows are red and an EXTRAPOLATION warning names them. The result is still calculated and stored with the flag — it is a review signal, not a block, and the Review page shows the range next to each result so it can be read against your accreditation scope. An

empty range means the model was saved before the range was recorded: re-fit and re-save it on the Calibration page.

- Expander **Preview certificate** — run metrics (Result mode, Sample, Calibration ID, Analysis used) above the certificate-style table with normalized values, GUM display rounding (*Significant digits for uncertainties* control below the table, default 2) and left-aligned relative-uncertainty columns with a % sign; the unnormalized sum appears below this table as well.
- Section **Certificate Follow-up** — renders the analysis certificate as a formatted PDF: sample details, instrument, method, calibration provenance (model, regression method, calibration gases, data age), the GUM-rounded results table with normalized values, an uncertainty statement (expansion factor, ISO 6974-2 RRF basis) and the run provenance. The document carries a status banner (PREVIEW / PENDING REVIEW / APPROVED); it only counts as a released certificate after the run is approved on the Review page. Preview it inline or download it with the  button. For a **stored** run the section also offers **Register certificate in database**: the PDF is frozen to data_local/certificates/generated/ (with a SHA-256 hash) and recorded as a Certificate (origin 'generated', approval status **draft**) with the normalized composition as CertificateComponent rows. A run that already has a registration shows the frozen document with its own download button instead — the live preview above may differ from it when settings changed later. Draft certificates are approved on the Review page (step 7), gated by the reviewer's access role and the rights matrix. You do not have to look the record up there yourself: while the run is still waiting, the blue notice under the PDF carries the link  **Review run n**, and a registered certificate carries  **Review certificate <no>** — each opens the Review page on the right tab with the record already selected.

Step 6 — Review and Approve the Run (second person)

- A **manager or admin** opens the Review page, tab **Runs** (filter *Pending*) — or clicks the run in the Dashboard's *Review queue*, which opens Review on the Runs tab with that run selected. The **Reviewer** card under the introduction shows who is reviewing, that every decision is recorded under that name with the moment and the notes, and the open work (runs, certificates, models); an operator or auditor sees *view only* there with the reason. The banner at the top names the reviewer; an operator or auditor sees the page read-only.
- Tick the run: the summary is a **detail card** — grouped label / value rows (Sample, Calculation, Review, ...) with the statuses as coloured chips (green = approved / released, amber = pending / draft, red = rejected / withdrawn / fail) — followed by the *Calculated results* (Mol fraction, Uncertainty and Cal. min / max shown with 5 decimals, the first column pinned while scrolling) with the calibrated range and, when a result lies outside it, the EXTRAPOLATION warning.
- The card carries an **Earlier decisions** group whenever the record was decided before (moment, user, verdict and note, rebuilt from the Change history) and a chip *previously rejected* / *previously withdrawn* when one of those verdicts was negative — a rejection stays traceable after a later approval.
- Enter *Review notes* and click **Approve** (the analysis is released for a certificate) or **Reject** (not fit for a certificate; re-run after corrections). The decision is stored with your username and the moment; *Review history of this record* shows every earlier verdict. **Reopen** takes a verdict back as long as no approved certificate hangs on the run.

Step 7 — Register and Approve the Certificate

- Back on the Execute page, section *Certificate Follow-up* now offers **Register** (it stays hidden while the run is pending): the PDF is frozen on disk with its SHA-256 hash and a Certificate record (origin *generated*, approval *draft*) is created. You do not have to stay in the same session: the drop-down *Approved runs waiting for registration* lists every approved run without a certificate — pick the run and register it. Without this step no analysis certificate exists; the Certificates page then shows only the imported certificates.

- On the Review page, section **Certificates** (filter *Draft*), tick the certificate (or click it in the Dashboard's *Review queue*): the detail card (header with status chips, the run and its review status, the user who entered an imported certificate (*Entered by*), *Earlier decisions* with the *previously ...* chip where applicable), the component values — with a **Rel. Unc.** column, the relative expanded uncertainty as a percentage with two significant digits, as on the Certificates page — and the frozen PDF (preview and download). **Approve** releases it (approval status *approved*, reviewer and moment recorded, approval basis *approved_run*); **Withdraw** takes a certificate out of use (status *withdrawn*, the PDF stays). Imported, hand-entered and internal QC certificates (Workflow D) appear here as well: approving one **is** the four-eyes check of the values against the source PDF — the source PDF opens next to the component values, you tick *I have checked every component value and uncertainty against the source document* and approve; the user who entered the certificate cannot approve it, and an external certificate without its source PDF on file cannot be approved.
- The Certificates page lists the certificate with its approval status; the approval columns are read-only there.

Workflow B — Analysis via CDS File

Use this workflow when you want to start from CDS files on disk and let NexChrom create the Sample for you during import.

Step 1 — Import CDS Files

- Open the Analysis page and select the **Analysis via CDS File** tab (the middle tab in the new ordering).
- Browse to the CDS files you want to import.
- Choose the matching instrument and confirm the import.
- NexChrom creates a new Sample from the file metadata, registers each file as an Analysis, and extracts AnalysisResponse rows.

Step 2 onward

Continue with Workflow A from Step 3 (Choose or Define a Method).

Inspecting CDS Files Without Linking (CDS File Browser)

Use the **CDS File Browser** tab on the Analysis page (the last of the three tabs) when you want to look at a CDS file on disk without linking it to a sample. The tab is structured as:

- **1. Instrument & folder selection** — pick the instrument whose data folder you want to browse.
- **2. Select CDS file** — single-select grid of files discovered in the instrument's data folder.
- **3. File metadata** — shows the raw metadata fields parsed from the selected file. For Agilent files this now includes *Dilution*, *Multiplier*, *Inj Volume* and *Actual Inj Volume* taken from `Report00.CSV`, plus a warning banner when ChemStation reported warnings on the run.
- **4. Peak results per detector** — shows raw per-detector peak data as it appears in the file (no totals, averages or derived values). Each detector or signal has a **Peaks** filter: *Named* (default) lists the identified peaks, *All (n unknown)* adds the unidentified ones; for a detector without unidentified peaks the filter is disabled and its help text says so.

The values shown in sections 3 and 4 are straight from the file via the shared parser — nothing is recalculated or aggregated, so the browser is safe to use as a quick audit view.

Workflow C — Calibration

Use this workflow to build a new calibration model for an instrument.

Step 1 — Register the Calibration Analyses

- Ensure the calibration samples exist on the Sample page, classified with a Calibration sample type.
- Import the CDS files for your calibration injections on the Analysis page and link them to the calibration samples.

Step 2 — Build the Calibration Model

- Open the Calibration page. The **Instrument** drop-down at the top left of *Calibration models* scopes the whole page: it lists instruments only, as *name (n models)*, and shows manufacturer, type, serial number and location in small type under it; *Archive status* on the right of the same row switches the table to archived models.
- Open the Calibration page.
- Select the instrument.
- Select the calibration analyses to use as input. The table shows one row per sample **and analysis session** (CDS files grouped by analysis date); by default only the most recent session of each sample is selected. Select an older session to deliberately reuse it, or several sessions of one sample to pool them (a warning names the pooled sample). For a model that already holds a fit, a message under the table compares your selection with the input set of the saved fit: *Selection matches the input of the saved fit*, or *Selection differs from the saved fit of this model with Added: / Removed: per sample @ session* — a new regression would then run on other input than the stored one.
- *Calibration input data* lists every calibration point: every component with a certified value in the selected gases is fitted alike, whatever its role in the method (broad fitting). The **Detector** column shows which detector(s) the included replicates come from; a point whose replicates mix detectors is red and a warning names it — fix the Include choices of that file on the Analysis page before trusting the fit.
- **Select a point** in that table to see the replicate peak areas behind its mean: one row per CDS file with its deviation from the mean, z-score and a Grubbs outlier flag, the statistics (n, mean, SD, RSD, u(mean)) and a control chart with the mean and the ± 1 s, ± 2 s and ± 3 s limits of the included replicates. Untick **Include** to leave a replicate out of this model's regression input, or press **Exclude n flagged outlier(s)**; **Restore all** undoes it for this point. These changes are **pending in the session** — the caption *Pending: +n excluded / -n restored* counts them — and are written only when you click **Save** next to them (stored with the calibration model, logged in the Change history, recorded in the fit's input set and copied when you save as a new model); **Discard** reverts them. The input table above and the regression keep using the *saved* exclusions, so **Execute regression is disabled while exclusions are pending** — save or discard them first. The analysis itself and other models are untouched. The Grubbs test is advisory — with few replicates it only catches gross outliers.
- Select the regression method: EIV (errors-in-variables per ISO 6143 — uses uncertainties on both peak area and certified value) is pre-selected; OLS (ordinary least squares), WOLS (weighted by $1/u^2(\text{certified})$) and GLM (robust Huber reweighting) remain available. OLS is kept as an unweighted reference: on data spanning several decades it ignores the low-concentration points and its GOF goes far above the threshold — which is the point of the comparison, not a defect.
- From two calibration points on every fit carries a free intercept — the origin is never forced (D17). The only through-origin fit is the single-point fall-back; at exactly two points the exact line with intercept is offered behind its own checkbox and its intercept is tested (D18).
- Click Execute. NexChrom computes all feasible polynomial orders ($1..\min(3, n-2)$) in one pass and shows the results grid with one row per component.

- In *Calibration results* pick the active polynomial order per component from the dropdown. The Auto column suggests the lowest order whose GOF stays below the **Max allowed GOF** setting (default 2.0); you can accept it or override.
- Any GOF value that exceeds the **Max allowed GOF** threshold (default 2.0, configurable in *Settings* → *SYSTEM* → *Calibration regression thresholds*) is highlighted in **red / bold** in the results grid and in the component-detail metric; a warning banner below the grid lists the offending components. The check is advisory only. A high GOF typically points to inconsistent measurement data (outliers, drift, mis-integrated peaks) — review the input rather than raising the threshold.
- If a component has fewer than three usable calibration points it is skipped with a warning naming the component and the offending sample(s); the other components still fit. Two opt-in checkboxes handle these cases: for a **single** point a 1st-order through-origin fit, for exactly **2 points** the exact line with intercept (b_0 -tested). Such fits are marked as not tested for adequacy in the results and on Execute.
- Select a component in the results grid to open its **regression detail** — the badge says *New regression — not saved*: metrics (method, order, GOF, calibration range, max calibration error and its sample), the coefficients with their uncertainties and $|b|/u(b)$, the covariance and correlation matrices, the fit statistics in groups (Fit, Goodness of fit with the intercept test, Calibrated range, Residual diagnostics: runs test on the residual signs, standardised residuals, Shapiro–Wilk normality), the calibration points with residuals and χ^2 contributions, the calibration curve (hover: peak area, certified value, χ^2 contribution) and the residual plots (y for all methods, x additionally for EIV; hover: per-point relative error). A bow in the residuals of a first-order fit shows up as few runs and a large max $|z|$ — try the next order.
- If a fit *stopped without a verified optimum* (column *Conv* = No with the reason, and a warning above the overview), the EIV iteration was halted by its divergence guard: the coefficients are the last accepted iterate. This nearly always goes together with a very high GOF and means the calibration points do not match their stated uncertainties — check the CDS files linked to the calibration samples and their certificates before saving. *Auto* never proposes such an order when a converged one exists.

Step 3 — Save the Model

- When the fit looks acceptable, save it. Saving into a model that already holds a fit (or that calculation runs reference) is guarded: choose **Save as new model** (recommended; a dated name is suggested) or an explicit overwrite.
- The suggested name for a new model carries date **and time** (models are often re-fitted a few times before one is kept); existing date/time suffixes are stripped first so stamps never stack.
- With *Save as new model*, NexChrom creates the new model as current (*is_current* = 1, *activated_at* set) and archives the previous model (*is_current* = 0, *archived_at* set) with its component fits intact; replicate exclusions travel with the new model.
- A saved model starts as **draft** (columns *Release*, *Released By*, *Released At* in the models table, read-only). Saving new fits into a model that was already released puts it back to *draft* — the flash message says so — because the release covered the previous fits.

Step 4 — QC Check of the Calibration, then Release (second person)

- **Analyse a QC sample in the calibration session.** Run at least one QC / reference sample on the instrument together with the calibration gases (the same day; Settings allow ± 1 day around the first and last calibration analysis). Link its CDS file(s) like any sample.
- **Calculate it with the new model.** On Execute select the QC sample, set the *Release status* filter to *All*, choose the new model and store the run: the QC check runs automatically (see *Check a QC run*). The expander *QC sample in this calibration* on the Calibration page shows the session window, the QC analyses found and the verdict — or what is still missing.

- A **manager or admin** opens the Review page, section **Calibration models** (filter *Draft*), ticks the model (or clicks it in the Dashboard's *Review queue*): the detail card (instrument, method, release status as a chip, *Earlier decisions* with a *previously rejected* / *withdrawn* chip where applicable) and the fits per component (fit type, regression method, polynomial order, GOF with 4 decimals, R^2 with 6 decimals, TSSD with 4 significant digits, Conc. min / max with 6 significant digits, fit moment; the first column stays pinned while scrolling) and the block **QC check of this calibration**: the release is blocked while no QC sample was analysed in the session, the QC sample was not yet calculated with the model, the check has no limits to check against, or the check failed. A warning may be released with a note; a failure means the analyzer did not perform within the QC limits — investigate and recalibrate.
- Under the model card the full-width block **Decisions and notes** holds the waiver reason, the current release notes, the model notes and the numbered decision history (oldest first, the current decision marked), so the card itself stays short.
- **Release** puts the model into use — Execute now offers it for new runs; **Reject** marks it as not fit; **Withdraw** (on a released model) takes it out of use again, runs already calculated keep their results. Notes and the reviewer are stored; the history is shown per model.
- **Changing a model that was decided on.** You do not need to withdraw a model to change it: fit it again on the Calibration page and save (overwrite, or *Save as new model*). New fits put a released, withdrawn or rejected model back to **draft** (the flash and the release notes say so, an earlier QC waiver is cleared) and it appears in the Review queue again for its QC check and release. *Withdraw* is for taking a model out of use, not a step before editing.
- **Release without a QC check**: when no QC sample of the calibration session has a verdict — no QC measurement was made, or it is not calculated with the model yet — an **admin** may still release the model: tick *Release without a QC check (admin only)*, enter the reason (for example a certificate correction with no QC measurement in the session) and click *Release*. The reason is stored on the model (QC *waiver* column on the Calibration page, *Waiver reason* on Review) and in the release notes; the Dashboard shows an orange warning for the model until a QC run with it passes. A **failed** QC check can never be waived, and a manager cannot give the waiver.
- Models that existed before this release step was introduced were released automatically by the migration (released by *migration*); withdraw them here if they should not be in use.
- Reopening a model later: with no fresh results, *Calibration results* lists the stored fits (order, method, GOF, R^2 , calibrated range, save moment, data age). Select one to see its full regression detail — badge *Saved fit* — with the calibration points rebuilt from the recorded input set (a warning tells you when the underlying data changed since the fit was saved; for EIV a re-fit confirms the stored coefficients). Historical models remain available for re-calculation of older samples. Every save and archiving step is recorded in the Change history (AuditLog) on the Calibration page.

Workflow D — External Certificate Import (PDF)

Use this workflow to bring a reference certificate issued externally (VSL, NPL, NMi) into NexChrom so you can link it to samples as a fixed-value source.

Step 1 — Upload the PDF

- Open the Certificates page and select the PDF import tab.
- Upload the certificate PDF.
- NexChrom first attempts to extract text from the PDF layer. If the document is a scanned image, it automatically falls back to OCR.
- It identifies the certificate format (VSL / NPL / NMi) and extracts certified component values.

Step 2 — Review and Save

- Review the extracted values on screen.
- Correct any parsing errors by editing the grid.
- Verify the validity date and cylinder number are correct. A cylinder number that differs from the linked sample's is a **critical** finding: a red notice on the Certificates page and on Review, and a  alert on the Dashboard — a certificate belongs to one cylinder, so correct the link or the number.
- Save. There is no verification step on this tab any more: the certificate is stored as a **draft** under your name (*Entered by*). A second user checks every value against the source PDF and approves it on the Review page (Step 4); until then it is not offered as certified values on Calibration or as a basis on Quality control.
- NexChrom writes a Certificate row and a CertificateComponent row per reported component. The certificate status defaults to active; the approval status is draft.
- On save the PDF itself is stored under `data_local/certificates/imported/<certificate no>.pdf` and its SHA-256 hash is recorded on the certificate (columns *PDF* and *PDF Hash* in the Certificates table), the same way the generated analysis certificates keep their PDF. Uploading a PDF that was already imported under another certificate number shows a warning naming that certificate before you save.

Step 3 — Link to a Sample

- Open the Sample page or the Certificates page and link the imported certificate to the sample that represents the material it certifies.
- Fixed components in future Methods can now reference this certificate as their value source.
- A certificate entered **by hand** in the Certificates table has no PDF yet: tick it, choose the supplier's PDF in *Certificate composition* > **Source document** (a yellow line and the *Pending* caption show it as an unsaved change) and press **Save** in that section — the PDF is then stored with its SHA-256 like an imported PDF; *Discard* drops it. A PDF on file can be opened in the expander *View the PDF on file* (inline viewer and download). Without a source document the certificate cannot be approved.

Step 4 — Check and Approve on Review (four eyes)

- A **different user** than the one who entered the certificate (a manager or admin) opens the Review page, section **Certificates** (filter *Draft*), and ticks the certificate: the detail card names *Entered by*, and the **Source PDF** expander opens next to the component values.
- Compare every component value and uncertainty with the PDF, correct the certificate on the Certificates page if something is wrong (the check then starts again), tick *I have checked every component value and uncertainty against the source document*, enter notes and click **Approve**. The approval is recorded under your name with the moment.
- Two rules block the approval: the source PDF is missing (upload it first) and the four-eyes rule — the reviewer is the user who entered the certificate. Internal QC certificates (created by the QC promotion) follow the same four-eyes approval; their source are the approved runs.
- A later change of the values, the header or the source document of an approved certificate withdraws the approval automatically (the certificate is a draft again, the reason is noted) so the changed certificate is checked again.

Dashboard

The Dashboard is the landing page: the alerts, the work waiting for the second person, the status per instrument and the routes through the application.

- **Panel switcher:** the four sections **Alerts**, **Review queue**, **System status** and **Workflows** are shown one at a time. The bar above them holds four equally wide blocks, the active one filled in its panel colour (red, blue, green, purple), with the number of items small in parentheses (alerts, queue entries, instruments; none behind Workflows); click a block or the ◀ / ▶ arrows and the chosen panel slides in. Alerts opens first; your choice is kept while you are logged in. Every panel opens with a centred coloured title band and a subtitle.
- **Alerts:** six tiles on top (critical, warning, caution, certificates, calibration models, quality control), then every alert as a line sorted by severity — ● critical first, then ● warning, ● caution, ● informational, ● ok — and the legend of the coloured dots at the bottom. Every alert is a link: a certificate alert (over date, expiring, not linked, cylinder mismatch — the mismatch is critical) opens the Certificates page with that certificate ticked and its composition open; a calibration-model alert (without coefficients, expiring, released without a QC check) opens the Calibration page on the model's instrument with the model ticked; a QC-run alert opens Review with that run; a QC-sample alert opens Quality control › Monitoring with instrument and sample preselected. The page switches without a reload.
- **Review queue:** what waits for the second person in four columns — runs pending review (run, sample, instrument, stored date, a QC fail mark), certificates awaiting approval, calibration models awaiting release and methods awaiting release — six per column with a count (... *and n more on the Review page; nothing waiting*) and an *Open Review* link; every entry opens Review with that record ticked, its filter set and only its section expanded.
- **System status:** six tiles (runs pending review, draft certificates, draft models, failed QC runs in 30 days, provisional QC limits, over date / warning), then one line per instrument — traffic light, released model with fit date and age, last QC run with verdict and age, QC limits, pending runs — each a link to the model on Calibration; the legend and the *Expiring / expired* line below.
- **Workflows:** one tab per route — *Analysis via Sample*, *Analysis via CDS files*, *Calibration* and *Execute → analysis certificate* — with the steps in the order you walk them, including the second person's review, approval and release steps (blocks in the Review blue). Every step has a filled arrow button that opens its page and, next to it, a collapsed expander: open it to read what to do on that page (section, button, result, what can block) and to find an *Open <page>* link.
- The Review page itself shows its four parts — **Runs**, **Certificates**, **Calibration models**, **Methods** — as tabs: reviewing runs, certificates, models and methods are separate jobs. Via the menu you arrive on *Runs*; a link from the Dashboard's Review queue opens the tab of that record.
- **Methods need a release too.** A method is complex (component roles, RRF rules, fixed values), so a second person checks it: a new or changed method is a *draft*; on Review › *Methods* a manager or admin who did not make the last change ticks it, checks the roles and RRF rules and presses **Release**. Until then the Calibration page refuses to save a fit on the method and Execute refuses to store a run with it. An **admin** may override that with a reason (tick the override box, enter the reason): the reason is stored on the model or run, logged in the change history and shown as a ● alert on the Dashboard until the method is released. Saving component roles or RRF rules of a released method puts it back to draft. The methods that existed before this step were grandfathered as released.
- At the bottom of the **Workflows** panel (below the three routes) the expander **Complete workflow overview** renders the full NexChrom flow as a fixed grid: the left column holds Component alias (joined to Component by a dotted line without arrowhead), Component, Sample and CDS files; the middle column Certificate (known values), Method, Analysis and Detector template; the right column Analysis certificate, Review, Calibration and Execute; the labelled lines show what feeds what (certified values, components & roles, cal. samples, model, ...).

Maintenance

Database Integrity Checks

- Open the Settings page and select the Maintenance tab.
- Run the integrity check. The system reports anomalies such as orphan foreign keys or missing required values.
- Follow the reported guidance to resolve any findings.

Database Export (Backup)

- From the Maintenance tab, run the database export operation.
- Store the exported copy outside the project directory if you need an off-site backup.

Settings

The Settings page is organised in six tabs (USER SETTINGS, AI ASSISTANT, SYSTEM, MAINTENANCE, DATABASE, SQL). Sections on the SYSTEM and MAINTENANCE tabs are rendered as default-closed expanders so the view stays calm.

- **USER SETTINGS** controls display preferences: table view height, show/hide DB footer, show/hide Import/Export controls, and whether page selections are cleared on navigation.
- **AI ASSISTANT** stores API keys for Gemini, Claude and OpenRouter. Each provider has a **Test API key** button that verifies the key against the provider's listing endpoint without consuming tokens; the last result is shown as  not tested,  OK or  failed in the expander title.
- **SYSTEM** exposes Server configuration, E-mail (SMTP) configuration, the Alert check schedule, **RRF defaults** (relative standard uncertainties of calculated RRFs per detector type, ISO 6974-2 Annex B: FID 2 %, TCD 10 %), **Calibration regression thresholds** (Max allowed GOF, default 2.0, used to flag poor calibration fits in red / bold on the Calibration page and as the threshold of the Auto order proposal; Max calibration data age, default 365 days, 0 = off, used to warn when the analyses behind a stored fit grow older than the limit) and **Uncertainty (expansion factor)** (the app-wide coverage factor k, default 2, used to expand stored certificate/result uncertainties and to convert to standard uncertainties at calibration input — decisions D18/D19/D20).
- **MAINTENANCE** offers the database export, data-integrity checks and record counts — each in its own expander.
- **DATABASE** holds just the active database path.
- **SQL** is the direct SQLite console, moved here from a separate page; it offers SELECT and Script execution modes against the active database.

The path of the settings file is displayed above the Save / Discard buttons at the bottom of the page.

User Management

Manage user accounts on the Administration page, USERS tab.

Create a User

- Open the Administration page and select the USERS tab. The tab holds two sections, each an expander open by default: **Users** and **Access roles & rights** (neither depends on the ticked user: the second is per access role; the minimum alert level for e-mails lives on Settings > System > E-mail (SMTP) configuration) (the Locations and File formats tabs hold one table each and therefore no expander — the rule: one section under a tab is shown bare, two or more sections are expanders).

- Click Create user to add a new row.
- Fill in username (required), full name, email, phone, and address fields.
- Select the user's laboratory Role from the dropdown (Laboratory Manager, Lab Supervisor/Team Leader, Laboratory Technician / Technician II, Lab Assistant/Assistant Technician, Research Scientist/Analyst, or Principal Investigator). This role is informational.
- Select the user's **Access role** (operator, manager, auditor or admin; default operator) — the role the permission system will use; admin always keeps full rights.
- Tick the Email alerts checkbox if this user should receive alert e-mails.
- Click Save to write to the database.
- The minimum alert level ( Critical,  Warning,  Caution) at which subscribed users receive e-mails is set on the Settings page, System > E-mail (SMTP) configuration, next to the SMTP settings; here you only decide *who* receives them (column *Email alerts*).

Roles & rights

- Below the Alerts section, the **Roles & rights** matrices define per access role which pages may be viewed and edited (expanders *View* and *Edit*, one checkbox per page).
- Only an admin can change the matrices; other users see them read-only. The admin row always keeps full rights — changes to it are ignored.
- Change the ticks and click **Save** (the standard blue button; it stays disabled while nothing is pending) to store the matrix — audited, visible in the Change history; **Discard** next to it drops the unsaved ticks. The caption *Pending: ~n rights changes* counts the ticks that differ from the saved matrix.
- The matrix is enforced on every page: without the view right the page is not shown (a lock notice instead), without the edit right the page shows a *View only* notice under its title and every write button (Save, Create, Delete, Archive, Execute, Approve, ...) is disabled; Discard, Clear and read-only actions stay available. Review decisions and the saving of Quality-control limits keep their own, stricter gates on top.

E-mail Alert Notifications

Configure SMTP so the Dashboard can send alert e-mails to subscribed users.

Configure the SMTP Server

- Open the Settings page and select the SYSTEM tab.
- Expand E-mail (SMTP) configuration.
- Enter SMTP host, port, security mode (STARTTLS, SSL or none), username, password and From address.
- The password is stored in the operating system credential store (Windows Credential Manager), not in the settings file.
- Click Send test e-mail to verify delivery to your own user account's e-mail address.

Schedule Alert Checks

- In the same SYSTEM tab, below SMTP configuration, tick "Send e-mail immediately when an alert occurs" if you want each new alert to trigger a message.
- Set the Daily check time for the scheduled digest.
- A yellow notice states that the dispatch scheduler itself is still being wired up; the settings you enter are persisted so they take effect as soon as the scheduler goes live.

Subscribe Users to Alerts

- On the Administration page (USERS tab), tick the Email alerts checkbox for each user who should receive alert messages, and make sure their e-mail address is filled in.
 - On the Settings page (System › E-mail (SMTP) configuration) set the minimum alert level so users only receive e-mails at or above that severity.
-

Upload Instrument Documentation

Provide instrument manuals, specs, and troubleshooting guides for the AI Assistant.

Upload Files

- Open the Instrument page, Documentation tab.
- Select the instrument from the drop-down (it defaults to the instrument selected on the Instruments tab); the expanders Instrument-specific documentation and Shared make/model documentation open.
- If no documentation folder exists yet, click Create documentation folder.
- Use the drag-and-drop area or Browse files to select one or more files (PDF, TXT, MD, DOCX, XLSX).
- Click Upload to save the files to the instrument's documentation folder.
- Documents that apply to every instrument of one make/model belong in the shared section — uploading or deleting there affects all instruments that point to that shared folder.
- To delete a file, select it in the grid and click Delete.

View or Download a Document

- Select exactly one file in either table: a download button appears, plus an inline preview for PDF, TXT and MD files (DOCX and XLSX are download-only).
-

AI Assistant

Use the AI Assistant to ask questions about NexChrom, your instruments, or GC analysis.

Start a Conversation

- Open the AI Assistant page.
- In section AI Configuration, select an AI provider (Google Gemini is free) and a model.
- Optionally select an instrument to load its documentation as context.
- Type your question in the chat box at the bottom of the page.
- The AI responds using its knowledge of the NexChrom database schema, workflows, and any loaded instrument documentation.
- Click Clear conversation to start a new conversation.

Configure API Keys

- Open the Settings page, AI ASSISTANT tab.
- Expand the provider you want to use (each is default-closed).
- Enter the API key:
 - Google Gemini: free key from aistudio.google.com/apikey.
 - Anthropic Claude: paid key from console.anthropic.com.
 - OpenRouter: free key from openrouter.ai/keys.

- Click **Test API key** to verify the key against the provider before saving; the outcome is shown inline and the expander title reflects it with a coloured indicator.
 - Click **Save all settings** to persist the keys.
-

Troubleshooting and Common Issues

A peak in a CDS file was not matched to a component

Add the exact peak name as an alias of the canonical component: on the Component page, tab *Global component aliases*, when the name holds for every instrument, or on the Instrument page, expander *Instrument component aliases* of the selected instrument, when it is specific to that instrument (an instrument alias wins over a global one). Re-select the CDS file (or re-import it) and the response will now be recognized.

Execute is blocked by a validation error

The Execute page shows which component, response, or calibration mapping is missing. Typical causes:

- The selected Method requires a component that is not present in the linked analyses (import more CDS files or change the method).
- The selected CalibrationModel does not cover a required direct component (build a new calibration that includes it).
- A fixed component has no value in the method (enter it in the *Fixed value* column of the component role on the Method page; the validation check names the component).

I deleted a sample and its certificate disappeared from filters

Certificates use ON DELETE SET NULL for sample_id, so the certificate itself is preserved when its sample is deleted — only the link is cleared. Re-link the certificate to a different sample if needed.

I want to rerun an old analysis with the original calibration model

Because calibration models are preserved with is_current flags rather than overwritten, you can still select a historical CalibrationModel on the Execute page. Pick the desired historical model, the original method, and re-run; a new CalculationRun is created and the results are traceable back to the historical model.

The app reports a missing table or column

This typically means the active database was initialized from a different DDL version. Verify the database_version key in the Metadata table (Settings → Database) and either switch databases or re-initialize using data_local/database/init_db_v2.2.0.py against the current DDL.