TFCompanion Release Notes, Version 6.5

Release date: April 22, 2026

Scope: Model refinement, Production Batch, Library Manager, property dialog cleanup, database groups, FFT performance, FFT chart annotation, and chart stability updates.

Summary

This update combines workflow improvements in model refinement and Production Batch, library import/export usability updates, database group administration, and a major FFT optimization and chart-annotation pass. Several related UI and chart-stability issues were also corrected, while integrating these features.

Model Refinement

1. Added a staged current-fit refinement option in the Refine Model workflow. This path keeps the current model structure, enables calculated parameters gradually, hands accepted stages to Marquardt, and backs off parameters that destabilize the fit.

2. Layer selection in the refinement dialog was broadened so it is no longer limited to HOscillator layers when the current model contains other parameterized materials with calculated parameters.

3. The refinement dialog now streams intermediate candidate results into the results table while the search is running instead of waiting until the full process finishes.

4. The dialog now supports stopping an active refinement run cleanly by forwarding a stop request into the optimization control path.

5. The dialog was made more usable on shorter screens by allowing the full panel to scroll when needed.

6. Refinement focus now defaults to the full spectrum instead of preselecting a local wavelength region automatically. Local focus can still be re-enabled by selecting a refinement area explicitly.

Production Batch and Calibration Export

1. Normal Production Batch export with Apply Calibration now stays on the regular batch export path even when a template is loaded. This corrected the case where a loaded template incorrectly forced the template-comparison export path and produced no output.

2. Repeated file headers in calibrated batch export were removed. Export now produces one header, then all rows, then the statistics block.

3. Template-based matching for correlation/export was aligned better with the partial-name loading behavior so name mismatches are handled more consistently.

Technical note: the original issue was primarily an export-routing problem, not strictly a partial-template-load problem.

Library Manager and Import/Export Workflow

1. The Library Manager now remembers the last used directory during the session across import/export browsing actions.

2. Multi-selection export was enabled, so multiple materials, filmstacks, and recipes can be exported together.

3. Binary zip export was simplified so the zip file name is retained while browsing to a target directory instead of forcing the user to re-enter the name.

4. Absolute zip export paths are now handled correctly when saving outside the default database directory.

5. Imported records are automatically tagged with the imported group in the supported import paths.

Property Dialog Cleanup

The Remove trend checkbox was removed from the Property Calculation dialog because the FFT helper path was already effectively performing the trend correction unconditionally. The underlying XML property was left in place for backward compatibility.

Database Groups

1. Records now support arbitrary custom group names in addition to the built-in groups: standard, development, production, and imported.

2. Custom groups are persisted per database using lightweight store metadata, so newly created custom groups survive restart even if they are temporarily empty.

3. Legacy records remain backward-compatible. Missing or blank group values continue to default to standard.

4. Group filters were propagated to the main sidebar, Production Batch, MultiExperiment, and Material/Layer sidebars.

5. The side-panel Set Group popup now refreshes dynamically, so custom groups can be assigned from the popup as well as selected in the filter dropdown.

Technical note: the custom group discovery originally depended on scanning full records. That was replaced with a store-level catalog to avoid heavy startup scans and memory pressure.

Group Management Tools

A new Manage Record Groups dialog was added to the Tools menu.

Supported operations:

- Rename a group globally within a selected database.

- Merge one group into another group.

- Delete an empty group from the catalog.

- Delete a non-empty group by reassigning its records to standard.

- View custom groups together with their record counts.

Technical note: the dialog now loads asynchronously and shows a scanning message while the selected database is being analyzed.

FFT Performance

The FFT processing path was optimized significantly.

Key changes:

- Removed redundant FFT setup/windowing work by applying the window before loading data into the FFT path.

- Replaced Marquardt-based polynomial trend fitting with direct least-squares regression for FFT trend removal.

- Preserved the previous code paths as comments where appropriate for rollback/reference during validation.

Observed result on repeated runs of the same data: the full FFT processing path decreased from roughly 35 to 40 ms into the single-digit millisecond range after warm-up, with repeated hot-path runs around 6 to 8 ms.

Technical note: the first run remains slower because of JVM class loading and JIT warm-up. Later runs represent the more realistic steady-state throughput.

FFT Chart Annotation

FFT thickness PSD charts can now show selected FFT peaks directly on the chart with thickness labels.

Behavior:

- Annotation is controlled by the calculation property FFT_annotate in the FFT Options panel.

- Labels use the corrected thickness values associated with the selected FFT peaks.

- Annotation now uses only the FFT peaks explicitly selected by the FFT/thickness assignment logic.

- Label font size was increased for readability.

- Marker drawing and placement were adjusted to avoid duplicate labels and to reduce overlap with chart edges and the legend.

Chart Stability

Repeated FFT chart refreshes previously triggered intermittent Chart repaint exceptions in axes, grid, scatter series, connected series, and tooltip rendering. The refresh path was hardened by rebuilding charts more safely and by guarding transient repaint-time null states in the affected chart objects.

Other UI and Logging Fixes

- Side-panel group loading was serialized to avoid concurrent background DB reads that could trigger index/store read exceptions.

- The Material2DChart scaled-axis debug printout was commented out to reduce console noise.

- Several intermediate chart, dialog, and selector behaviors were adjusted during integration to keep the UI responsive while database scans or FFT chart rebuilds are running.

Compatibility

- Legacy databases remain readable.

- Legacy records without explicit group metadata still default to standard.

- Pre-catalog custom groups that no longer exist on any record cannot be rediscovered automatically. New custom groups created with this version are preserved in the database-level catalog.

Material Import and Parametric Models

1. Imported parameterized materials now load more reliably from the Layer/View Import Material path, including text exports such as Sellmeier and Sellmeier2 definitions.

2. Material charts no longer freeze when imported parameterized models produce non-finite optical-constant values near resonances. Invalid points are skipped safely during chart sampling and axis-range calculation.

3. Sellmeier-family evaluation was corrected to use lambda^2 - Li^2 denominators, fixing wrong dispersion from imported Sellmeier material data.

4. The material UI and parameter tables now support imported Sellmeier materials with more than 5 oscillators instead of failing with combo-box index errors or fixed-name limits.

5. EMA interface creation in place now copies tabular materials safely by falling back from same-class copy constructors to generic Material copy constructors and clone().

Material-Model Review Fixes

1. Cody-Lorentz-Urbach was corrected to use the proper Gt denominator sign, and the Urbach invalid-value guard now detects NaN correctly.

2. Tauc-Lorentz deserialization now restores a missing or invalid C parameter correctly.

3. Tauc-Lorentz-Drude now uses a consistent flat parameter order for display/access and mutation.

FFT Peak Refinement and Thickness Confidence Limits

1. FFT peak assignment no longer loses the selected secondary peak because next-maximum index propagation was corrected.

2. Approximation fitting now uses signed residuals for Marquardt-based peak fitting, improving numerical behavior during Gaussian peak refinement.

3. Gaussian peak refinement remains the default path, with parabolic interpolation used only as a fallback when a formal Gaussian confidence limit is unavailable.

4. Thickness confidence limits from FFT peak refinement are now carried into the assigned filmstack thickness parameters more consistently, and the fallback path derives its confidence limit from the local quadratic curvature.

UI Stability and Cleanup

1. Multiple startup and material-dialog tab panes now use scroll-tab layout together with safer dynamic-tab updates to avoid intermittent JGoodies PlasticTabbedPaneUI ArrayIndexOutOfBoundsException crashes during startup and in the Layer View dialog.

2. FilmstackUtils cleanup removed redundant doubleValue() calls left over from pre-autounboxing code, making the FFT utility code easier to read and maintain.