



# RNTuple API Review – Discussion of Midterm Findings

---

Jakob Blomer for the ROOT Team

RNTuple Workshop

2 December 2024



- `RNTupleView`
  - The class has been changed to behave more like an `REntry` with a single field, in particular the question of owning or non-owning storage became a runtime decision
  - `#16321` (done)
- `REntry` use in `RNTupleReader`, `RNTupleWriter`
  - Intention for the reader and writer API wrt. `REntry` handling is to be symmetric
  - Both `RNTupleReader::LoadEntry()` and `RNTupleWriter::Fill()` take an optional `REntry` argument.
  - If not provided, they use the default entry of the `RNTupleModel`
  - Related issues: model from on-disk info can now be created without default entry
  - `#16324` (done)



- Page Size Tuning & Memory Consumption on Write
  - Addressed by a new, adaptive algorithm to set page sizes (done).
  - The new algorithm grows the pages as needed, so that dense columns get large pages and sparse columns small ones.
  - Pages still have an absolute limit (default 1MB) and the overall memory budget used for page buffers is limited.
  - Further memory management optimizations [▶ #16753](#) (done) [▶ #16752](#) (done)
  - Aim: single nob for write memory consumption: target cluster size
- Flexible Control of RClusterPool
  - Tracked as issue [▶ #16325](#)



- Indexing
  - Larger scope; work on it has started.
  - A new class, the `RNTupleProcessor` implements iterations of non-trivial joins of `RNTuples` (in contrast to simple/single `RNTuple` iteration of the `RNTupleReader`)
  - Initial version of the `RNTupleProcessor` with support for unaligned friends merged.
  - Full functionality expected in 2025



- `RNTupleParallelWriter`
  - Clear guarantees about the locking around `TFile`
  - New method “`FillNoFlush()`” allows framework to control time of `TFile` access
  - New staged cluster committing allows to set the logical cluster ordering after flushing; facilitates “data barriers” such as lumi block separation
- Smaller items:
  - Point 7 (`RNTupleReader::GetView`): fixed [▶ #16556](#)
  - Point 8 (token use): fixed [▶ #16236](#)
  - Point 9 (token use): fixed [▶ #16557](#)



- `RNTupleModel` & `GetToken()`
  - The frozen state can be explicitly set by the user through `Freeze()` and `Unfreeze()` APIs. Both calls are idempotent.
  - Users can call `Freeze()` and `Unfreeze()`. Note that unfreezing a model will change the model id. As a result, after refreezing, existing `REntries` cannot be used anymore for reading and writing.
  - The model is implicitly frozen when passed to the `RNTupleWriter` / `RNTupleReader` and on committing a changeset for the late model extension (`RNTupleModel::RUpdater::CommitUpdate()`)
  - The model is implicitly unfrozen at the beginning of the `RNTupleUpdater` (`RNTupleModel::RUpdater::BeginUpdate()`).
  - `GetToken()` can be called on any frozen model and on model construction.



- Projected Fields
  - Field projections are stored as projections on-disk.
  - When reading, the user can decide whether the model reconstructed from disk should treat projections as projections, or present them as if they were physical fields (see `RCreateModelOptions`)
  - Note that models with projected fields cannot be used for the `RNTupleReader` (but, e.g., as a source for cloned model for skimming). The restriction on the `RNTupleReader` can be lifted if needed.
- Late Model extension
  - Late model extension will unfreeze the model at the beginning of the transaction and (re-)freeze the model when the extension is committed.
  - As a result, the model ID will change.
  - All existing `REntry` objects and tokens created from the model cannot be used anymore but new entries and tokens need to be retrieved.



- We think that all but the following points of the midterm review have been addressed
- Improvements to the `RClusterPool` may overflow into next year
- The work on indexing and the `RNTupleProcessor` will most likely conclude only in 2025



Many thanks for the thorough and useful feedback!