

# RNTuple: Towards First-Class Support for HPC Data Centers

Giovanna Lazzari Miotto<sup>1,2</sup>, Javier Lopez-Gomez<sup>2</sup>

<sup>1</sup> Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil

<sup>2</sup> EP-SFT, CERN, Geneva, Switzerland

E-mail: <sup>1</sup>glmiotto@inf.ufrgs.br, <sup>2</sup>javier.lopez.gomez@cern.ch

**Abstract.** Future HENP experiments, e.g., at the HL-LHC, will increase the volume of generated data by an order of magnitude compared to LHC Run 1 and Run 2. In order to sustain the expected analysis throughput, ROOT's RNTuple I/O subsystem has been engineered to overcome the bottlenecks of the TTree I/O subsystem, focusing also on a compact data format, asynchronous and parallel requests, and a layered architecture that allows supporting distributed filesystem-less storage systems, e.g., HPC-oriented object stores. In a previous publication, we introduced and evaluated the RNTuple's native backend for Intel DAOS. Since its first prototype, we carried out several improvements, both on RNTuple and its DAOS backend, aiming to saturate the physical link, such as support for vector writes and an improved RNTuple-to-DAOS mapping. In parallel, the latest developments allow for better integration between RNTuple and ROOT's storage-agnostic, declarative interface to write HENP analyses, RDataFrame. In this work, we contribute with the following: *(i)* a redesign of the RNTuple DAOS backend, including a mechanism for efficient population of the object store based on existing data; and *(ii)* an experimental evaluation on a single-node platform, showing a significant increase in the analysis throughput for typical HENP workflows.

## 1. Introduction

In the last few decades, the ROOT TTree [1] I/O has been used to efficiently store more than one exabyte of high-energy physics data and has become the *de facto* standard format for HENP. Its on-disk columnar layout enables efficient reading of a subset of the columns, a common use case in HENP analyses. However, future experiments, such as those ushered by the HL-LHC, are expected to increase the size of datasets by at least one order of magnitude, which makes high-bandwidth low-latency NVMe/SCM memory and distributed object stores especially relevant. TTree cannot fully exploit the capabilities of the aforementioned storage systems. RNTuple [2] is the new columnar I/O subsystem for ROOT that addresses TTree's shortcomings. In particular, it targets high performance on low-latency NVMe devices, asynchronous and concurrent bulk I/O, native support for object stores, e.g., DAOS, and robust, user-friendly interfaces.

Our contribution in this paper can be summarized as follows:

- We propose a matured, RNTuple–DAOS backend that delivers significantly improved read/write throughput. We recognize the importance of write throughput given that data must be imported into the object store to be processed in HPC facilities.
- We provide an experimental evaluation that demonstrates the performance improvements of the proposed techniques in a realistic, end-to-end analysis.


## 2. Background

In this section, we introduce concepts that are key to both ROOT’s RNTuple I/O subsystem and Intel DAOS, the object store targeted in this work.

### 2.1. ROOT’s RNTuple I/O subsystem

In RNTuple, data are stored column-wise on disk, similarly to TTree and Apache Parquet [3]. RNTuple consists of the actual stored data and the metadata structures that describe how the former will be interpreted. User data are organized into pages and clusters. A page contains an array of contiguous values of a fundamental data type (e.g., `Int32`) for a given column, while clusters hold all the pages containing data for a range of rows in the dataset. Pages for the same column and cluster denote a page group. On-disk pages are usually compressed. Columns of a complex C++ type, e.g., `std::vector<float>`, are broken into columns of fundamental data types. Metadata encompasses the header, the page lists, and the footer which contain, respectively, the schema description, the location of pages and clusters, and summary information and other metadata. For an illustrative example of the on-disk layout, see Figure 1.

RNTuple’s design decouples data representation from raw storage of byte ranges, enabling the implementation of backends for different storage systems, such as POSIX files or object stores.



**Figure 1.** RNTuple’s on-disk layout. Pages store data for the structure of matching color.

### 2.2. DAOS

DAOS [4] is part of the Intel exascale storage stack, and provides a fault-tolerant distributed object store targeting high bandwidth, low latency, and high IOPS provided by NVMe and SCM devices. DAOS completely bypasses the operating system’s virtual filesystem and block I/O layers. On the one hand, this forgoes kernel I/O coalescing and buffering, which is mostly relevant for spinning disks. On the other hand, the POSIX strong consistency model may limit the scalability of parallel filesystems. A DAOS system comprises multiple servers running a Linux daemon that exports local NVMe/SCM storage. RDMA is used to copy data from servers to clients where available, e.g., over InfiniBand.

DAOS’ storage space can be partitioned into pools, which are further divided into container namespaces, which hold objects. An object is a key–value store that is identified by a 128-bit object ID (`oid`). DAOS splits its keys into two parts: the distribution key (`dkey`) and the attribute key (`akey`); in particular, the former impacts data locality by determining the physical device in which the data are stored.

## 3. Efficient storage of HENP data in DAOS

In previous publications [5], we introduced a prototype of a DAOS backend for RNTuple. However, this initial version did not fully exploit the data link. In the following paragraphs, we describe several improvements carried out to improve the read throughput and shorten the data ingestion stage, i.e., the time required to import a dataset into the object store.

Orthogonal to the changes described below, and given the relatively high cost of creating an operation queue, we reworked the backend to maintain a single, persistent queue per accessed ntuple. Previously, queues were created and destroyed for each vector read/write.

### 3.1. Co-locality-based mapping function

Guided by typical HENP analysis patterns, RNTuple data are designed to be accessed columnar-wise, on a page group basis. However, the migration to an object store paradigm can disband related pages. Although previous mappings between RNTuple and DAOS did not exploit it (see Eqn. 1), the locality-controlling *dkey* imparts a way to retain columnar semantics by crafting a mapping function that keeps page groups together (see Eqn. 2). This unveils an opportunity to read data efficiently in parallel, provided that neighboring pages’ requests are coalesced in advance. In the equations below,  $Cluster_i$  and  $Column_j$  denote, respectively, the cluster and column identifiers, whereas  $Page_k$  is a strictly increasing integer that uniquely identifies a page.

$$\frac{\phi: \langle Cluster_i, Column_j, Page_k \rangle \rightarrow \langle oid, dkey, akey \rangle : \mathbb{N}^3}{\phi_{object-per-page}(Cluster_i, Column_j, Page_k) \mapsto \langle Page_k, \alpha_{dkey}, \alpha_{akey} \rangle} \quad (1)$$

$$\phi_{locality-driven}(Cluster_i, Column_j, Page_k) \mapsto \langle Cluster_i, Column_j, Page_k \rangle \quad (2)$$

### 3.2. Request coalescing

In order to benefit from the parallel reads made possible by the locality-driven mapping in Section 3.1, requests for coinciding pages in DAOS servers must be made in the same call. The DAOS’s client API allows calls with an arbitrarily long batch of requests for elements sharing the *oid* and *dkey*. By deferring I/O calls until a cluster’s requests have been coalesced by  $\langle oid, dkey \rangle$ , we realize the potential for much faster throughput and minimize operation queue overhead.

### 3.3. Vector writes

Though applicable to both reading and writing, the aforementioned improvements could not be fully leveraged by the I/O subsystem without support for vector writes, i.e., performing multiple update operations at once. Such batched transfers can improve throughput and amortize the overhead costs associated with the network fabric, and thus shorten the data import stage.

RNTuple was already able to buffer pages and defer writing until the entire cluster was committed, but pages were still individually issued to the underlying backend. Now, the interface has been extended to trigger vector writes in storage backends that may benefit from it.

### 3.4. Scatter-gather concatenation

The block size can have a significant impact on data transfer throughput. Yet, the size of a page, RNTuple’s data chunk unit, is fixed when the dataset is generated. The default page size (64 KiB) is reasonable for a local filesystem; however, DAOS may benefit from larger chunks to offset network overhead, especially over RDMA [6]. To circumvent this limitation, we introduce a scatter-gather-based mechanism to splice many pages within the same page group that are, then, transferred as a single unit.

## 4. Evaluation

In this section, we present the results of the experimental evaluation after applying the improvements described in Section 3. The experiments were all conducted with the hardware and software environments described below.

**Hardware platform.** Hewlett-Packard Enterprise’s Delphi cluster, comprising two servers and six client nodes interconnected via InfiniBand fabric. A single client node was used.

*Server nodes.* 4× Intel Xeon Gold 6240M CPU (18 physical cores), 2.60 GHz, 24.75 MB L3 cache, 185 GB DDR4 RAM, SMT enabled. These nodes are equipped with a Mellanox MT28908 ConnectX-6 InfiniBand adapter.

*Client node.* 2× Intel Xeon E5-2640 v3 CPU (8 physical cores), 2.60 GHz, 20 MB L3 cache, 131 GB DDR4 RAM, SMT enabled. High-speed interconnection is available through a Mellanox MT27800 ConnectX-5 InfiniBand adapter.

**Software.** The operating system is Red Hat Enterprise Linux 8.4 (kernel 4.18.0-305), employing `daos 2.2.0` (`ofi+verbs` provider), and `libfabric 1.15.1`.

**Test cases.** To evaluate the proposed method, we used the LHCb Run 1 open data "B meson decays to three hadrons" dataset comprising 26 columns, of which 18 are read for an analysis that culminates in a histogram for the B mass spectrum. The original dataset contains 8.5 million events and amounts to an uncompressed size of 1.5 GB. To simulate the use case for HPC clusters, the dataset is artificially replicated 10-fold to a total of 15 GB.

We measured both (i) the transfer rate for importing the entire dataset into DAOS, and (ii) the end-to-end analysis throughput, i.e., from storage to histogram, for the following scenarios:

**Baseline:** experimental backend that predates the proposed improvements. It employs the naïve object-per-page mapping between RNTuple pages and DAOS objects.

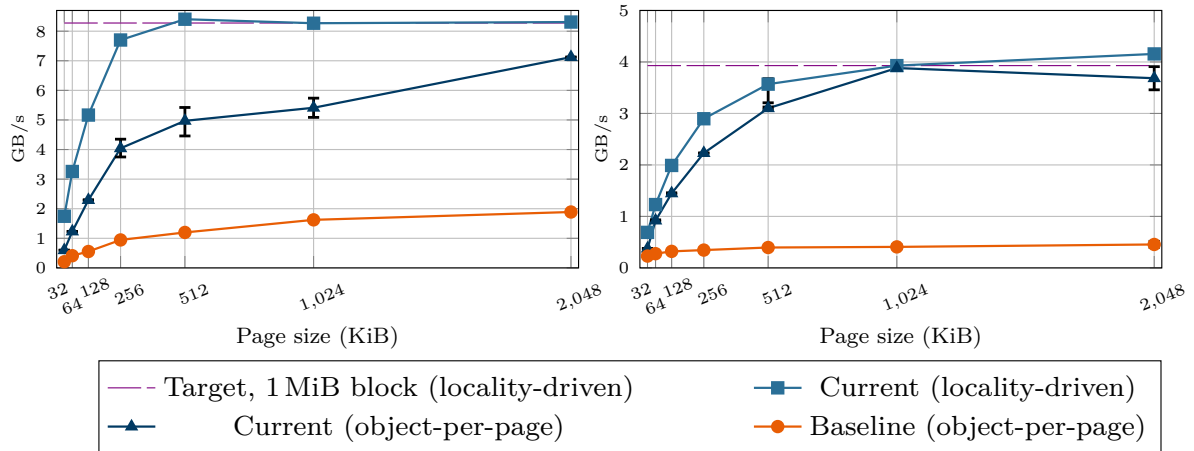
**Current (object-per-page):** with the improvements proposed in Section 3 and the previous, object-per-page mapping.

**Current (locality-driven):** proposed improvements, mapping that exploits co-locality.

**Target:** same as above, but the optimization described in Section 3.4 has been enabled and using a target transfer size of 1 MiB.

(2.a): write throughput (no compression)


(2.b): read throughput (no compression)



**Figure 2.** LHCb analysis throughput across different mappings, with `OC_SX` objects and 50 MB clusters.

Figure 2.a and 2.b show, respectively, the measured write throughput in GB/s while importing data and the read throughput during the analysis, both averaged over five runs. These plots compare performance w.r.t. various page sizes in each scenario described. The best results are achieved by using locality-driven mapping together with page sizes of around 1 MiB or more. The horizontal dashed line corresponds to splicing pages up to a target size of 1 MiB.

Figure 3 shows a comparison of the read/write throughput attained by both versions of the backend, before and after the improvements, for a dataset stored either uncompressed or using `zstd`. The results demonstrate, respectively for write and read scenarios, 4.5 to 9× and 3 to 10× better performance with the proposed method across all page sizes.



**Figure 3.** Throughput comparison for native 64 KiB pages with the baseline and with the current backend, with 1 MiB concatenation enabled.

## 5. Conclusion and future work

The results shown in the experimental evaluation demonstrate that, with the latest proposed developments, RNTuple’s DAOS backend is ready for high-throughput analyses in HPC data centers. In particular, we measured, respectively, a 9 and 4.3× improvement in the write and read throughputs for the default 64 KiB page size, which directly translates into much faster data ingestion and end-to-end analysis. By fully leveraging scatter-gather lists, the backend can target transfer sizes associated with higher throughput independently from the page size.

In future work, we plan to evaluate the backend in multi-process and distributed analyses with ROOT’s Distributed RDataFrame to saturate the link layer’s capacity. Moreover, significant development efforts are underway to extend native support for the S3 object store.

## Acknowledgments

This work benefited from support by the CERN Strategic R&D Programme on Technologies for Future Experiments CERN-OPEN-2018-006 and the Intel–CERN openlab collaboration. Access to the hardware for the experimental evaluation was provided by Hewlett-Packard Enterprise.

## References

- [1] Rene Brun and Fons Rademakers. “ROOT—An object oriented data analysis framework”. In: *Nuclear instruments and methods in physics research section A: accelerators, spectrometers, detectors and associated equipment* 389.1-2 (1997), pp. 81–86.
- [2] Blomer, Jakob et al. “Evolution of the ROOT Tree I/O”. In: *EPJ Web Conf.* 245 (2020), p. 02030. DOI: 10.1051/epjconf/202024502030.
- [3] Deepak Vohra. “Apache Parquet”. In: *Practical Hadoop Ecosystem: A Definitive Guide to Hadoop-Related Frameworks and Tools*. Apress, 2016, pp. 325–335. DOI: 10.1007/978-1-4842-2199-0\_8.
- [4] Zhen Liang et al. “DAOS: A Scale-Out High Performance Storage Stack for Storage Class Memory”. In: *Supercomputing Frontiers: 6th Asian Conference, SCFA 2020, Singapore, February 24–27, 2020, Proceedings*, pp. 40–54. DOI: 10.1007/978-3-030-48842-0\_3.
- [5] Javier Lopez-Gomez and Jakob Blomer. “Exploring Object Stores for High-Energy Physics Data Storage”. In: *EPJ Web of Conferences* 251 (2021). Ed. by C. Biscarat et al., p. 02066. DOI: 10.1051/epjconf/202125102066.
- [6] Philip Werner Frey and Gustavo Alonso. “Minimizing the Hidden Cost of RDMA”. In: *2009 29th IEEE International Conference on Distributed Computing Systems*. 2009, pp. 553–560. DOI: 10.1109/ICDCS.2009.32.