

[SEARCH](#)[LOGIN](#)[JOIN SNIA](#)[ABOUT](#) [STANDARDS](#) [EDUCATION](#) [TECHNOLOGY FOCUS](#) [NEWS & EVENTS](#) [RESOURCES](#) [MEMBERSHIP](#)[Educational Library](#)[SNIA Dictionary](#)[What Is...?](#)[White Papers](#)

Home » Education » Educational Library » SkyhookDM: Storage and Management of Tabular Data in Ceph

SkyhookDM: Storage and Management of Tabular Data in Ceph

Author(s)/Presenter(s): Jeff LeFevre, Adjunct Professor of Computer Science and Engineering, UC Santa Cruz University of California, Santa Cruz

Library: Podcast

Content Type: Presentation

Publish Date: Wednesday, September 23, 2020

Event Name: SDC 2020

Event Track: Computational Storage

Abstract:

The Skyhook Data Management project (skyhookdm.com) at UC Santa Cruz brings together two very successful open source projects, the Ceph object storage system, and the Apache Arrow cross-language development platform for in-memory analytics. It introduces a new class of storage objects to provide an Apache Arrow-native data management and storage system for columnar data, inheriting the scale-out and availability properties of Ceph. SkyhookDM enables single-process applications to push relational processing methods into Ceph and thereby scale out across all nodes of a Ceph cluster in terms of both IO and CPU. To highlight the benefits, we will present performance for various physical layouts and

query workloads over example tables of 1 billion rows, as we scale out the number of nodes in a Ceph cluster. In this talk, we first describe how we partition Apache Arrow columnar data into Ceph objects. We consider both horizontal and vertical partitioning (rows vs. columns) of tables. In contrast to objects storing opaque byte streams where the meaning of the data must be interpreted by a higher level application, Apache Arrow data can be partitioned along semantic boundaries such as columns and rows so that relational operators like selection and projection can be performed in objects storing semantically complete data partitions. Next we introduce our SkyhookDM extensions that utilize Ceph's "CLS" plugin infrastructure to execute our methods directly on objects, within the local OSD context. These access methods use the Apache Arrow access library to operate on Arrow data within the context of an individual object and implement relational processing methods, physical data layout changes, and localized indexing of data. Relational processing methods include SELECT, PROJECT, ORDER BY, and GROUP BY with partial aggregations (e.g., local min, max, sum, count, etc.). Physical data layout operations currently supported include transforming objects between row and column layouts, which we plan to extend to co-group columns on objects. Localized indexing is performed as a new object write method and supports index lookups that are beneficial to both point queries and range queries. SkyhookDM is accessed via a user-level C++ library on top of librados. The SkyhookDM library comes with Python bindings and is used in a PostgreSQL Foreign Data Wrapper. The source code is available at github.com/uccross/skyhookdm-ceph-cls under LGPLv2. SkyhookDM is an open source incubator project at the Center for Research in Open Source Software at UC Santa Cruz (cross.ucsc.edu). This work was in part supported by the National Science Foundation under Cooperative Agreement OAC-1836650.

Learning Objectives

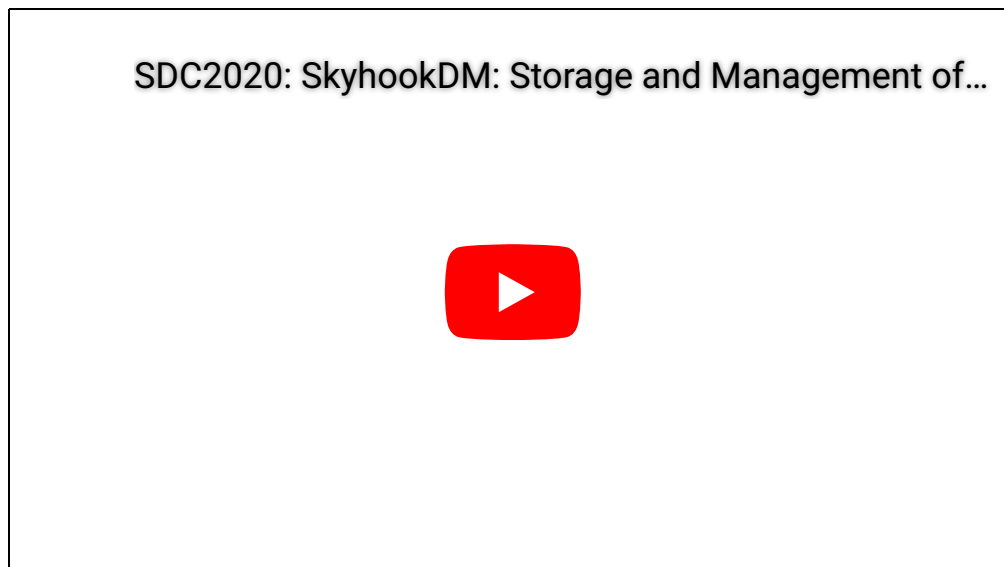
Using Ceph object classes to scale out relational data storage and access, Formatting, storing, and processing relational data directly in Ceph objects using librados, Indexing and physical

data layout of relational data within Ceph objects

[Listen to podcast](#)

[Download PDF](#)

Watch video:



Advancing Storage and Information Technology

Explore, discover, share, and meet other like-minded industry members. Get ahead, stay ahead, and create industry curves. Become a SNIA member today!

SNIA International

[SNIA Brazil](#)

[SNIA India](#)

[SNIA Japan](#)

Discover SNIA

[About](#)

[Standards](#)

[Education](#)

[Technology Focus Areas](#)

[News & Events](#)

[Resources](#)

[Membership](#)

SNIA Blogs

[SNIA Storage Blog](#)

[Cloud Storage Blog](#)

[Networked Storage Blog](#)

[Compute, Memory and Storage Blog](#)

[Storage](#)

[Management Blog](#)

Join

SNIA

Copyright © 2022 SNIA

[Site Map](#) | [Contact Us](#) | [Privacy Policy](#) | Chat provider: LiveChat