BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY
Office of Science

Scientific Data Division

THE OHIO STATE UNIVERSITY

Exploring the Proactive Data Containers Runtime System in VAST – A Case Study

Jean Luca Bez
jlbez@lbl.gov
Scientific Data Division
Lawrence Berkeley National Laboratory

Suren Byna
byna.1@osu.edu
Computer Science and Engineering
The Ohio State University

PDSW 2024 @ SC24

Proactive Data Containers

- PDC is an **object-centric** runtime metadata and data management system
  - Designed for transparent, asynchronous, and autonomous data movement
  - Provide data abstractions for **containers**, **objects**, and **regions**
  - `github.com/hpc-io/pdc`
- PDC has support for multiple traditional **storage backends**
- Our goal is to understand PDC's behavior on **novel storage** solutions like **VAST**
  - Enable cross-facility and multi-file-system deployment
  - Preserve the key characteristics of PDC

Proactive Data Containers

Abstractions

- **Container**
  - Groups a collection of objects
- **Object**
  - Describes any byte stream of information
  - Data and metadata
  - It can belong to multiple containers
- **Region**
  - *n*-dimensional arrays can be partitioned
  - Minimal access unit in PDC
  - Parallel access and flexible placement

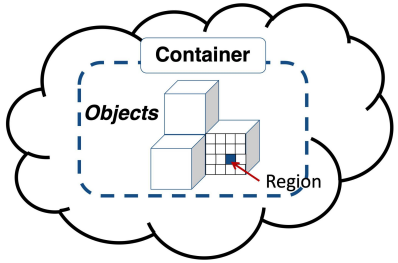

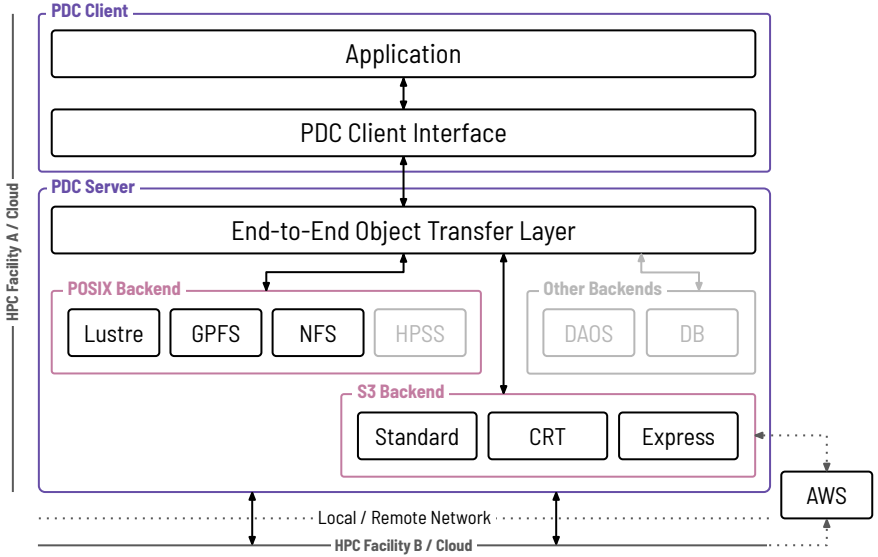

Initial Results

Memory-Intensive Workload

- VAST System:
  - 1 to 8 compute nodes
  - 16 ranks per node
- Microbenchmark:
  - Create and delete a 1K containers / objects
  - 128 ranks and a total of 8 million objects
- Key-results:
  - Maximum of 1.23 × 10^7 objects/second
  - Multiple PDC servers per node is beneficial
  - Investigating high cost of object delete

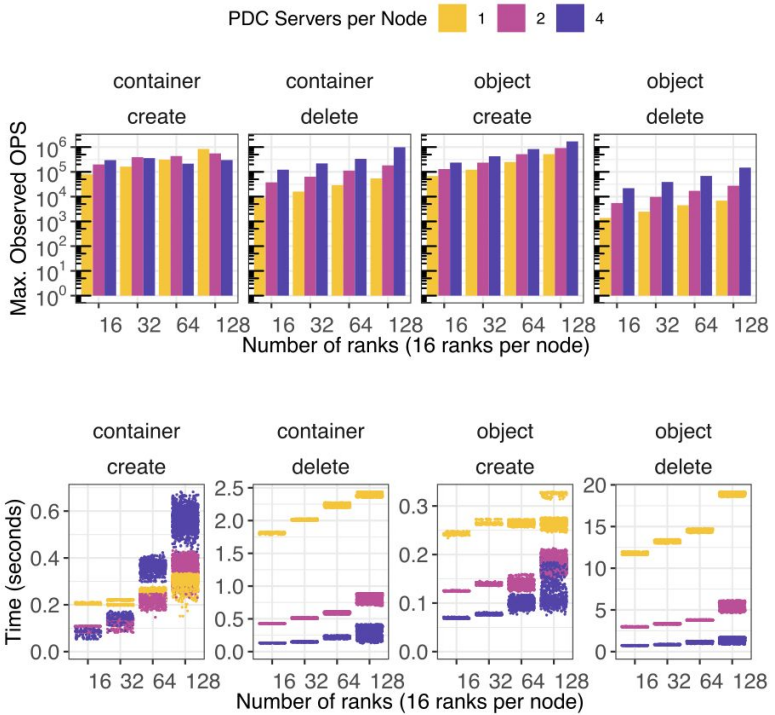

BERKELEY LAB

Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

Office of Science

Scientific Data Division

THE OHIO STATE UNIVERSITY

Exploring the Proactive Data Containers Runtime System in VAST – A Case Study

Jean Luca Bez (jlbez@lbl.gov) and Suren Byna (byna.1@osu.edu)

- PDC is an **object-centric** runtime metadata and data management system
- Ongoing investigation of PDC **performance** with **VAST** as a **backend**
- Next steps:
  - Investigate and optimize for **data**-heavy workloads
  - Cross-site data migration with PDC (e.g., Perlmutter ⟷ VAST-based cluster)