DAOS: Storage Innovations Driven by Intel® Optane™

Liang.zhen@intel.com

Principal Engineer, Intel

Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

No product or component can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/benchmarks .

Intel Advanced Vector Extensions (Intel AVX) provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

DAOS overview

What is DAOS?

- Storage pools globally accessible over the fabric
- Overcomes industry bottlenecks by leveraging Intel® Optane™ Persistent Memory and NVMe SSDs
- Delivers exceptionally high bandwidth and IOPS, meeting the demands of HPC and AI
- Strong distributed consistency (Database like)
- Tightly integrated with Applications
- Can be deployed as either a standalone file system, or a performance tier integrated with existing storage systems

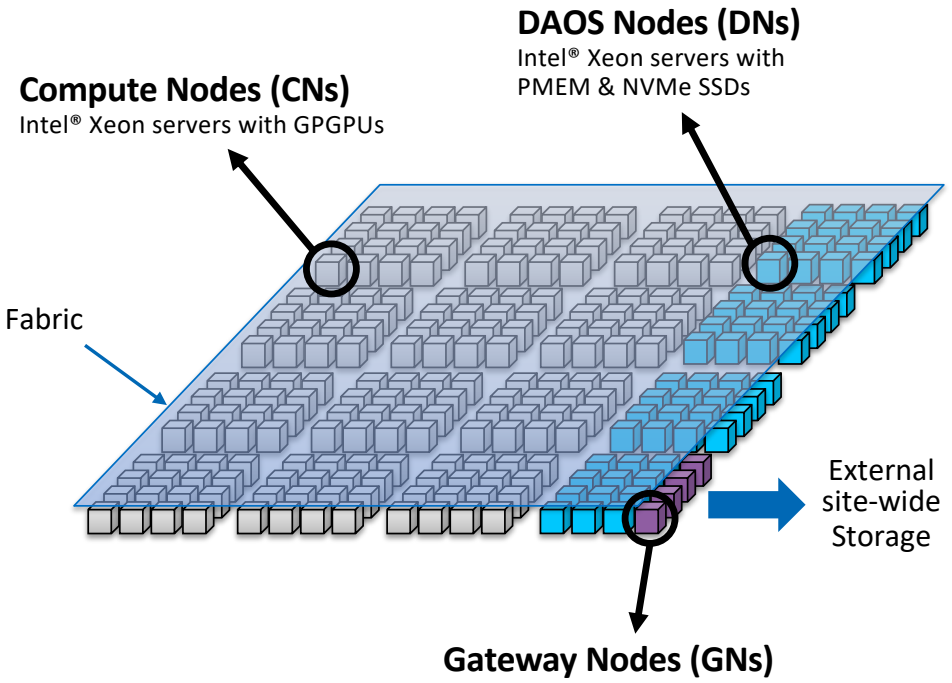

Generational leap forward in storage performance

DAOS architecture

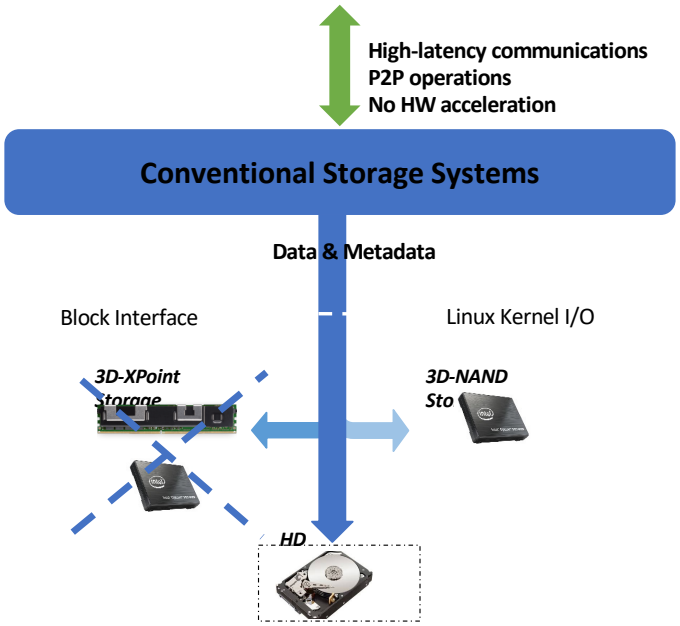

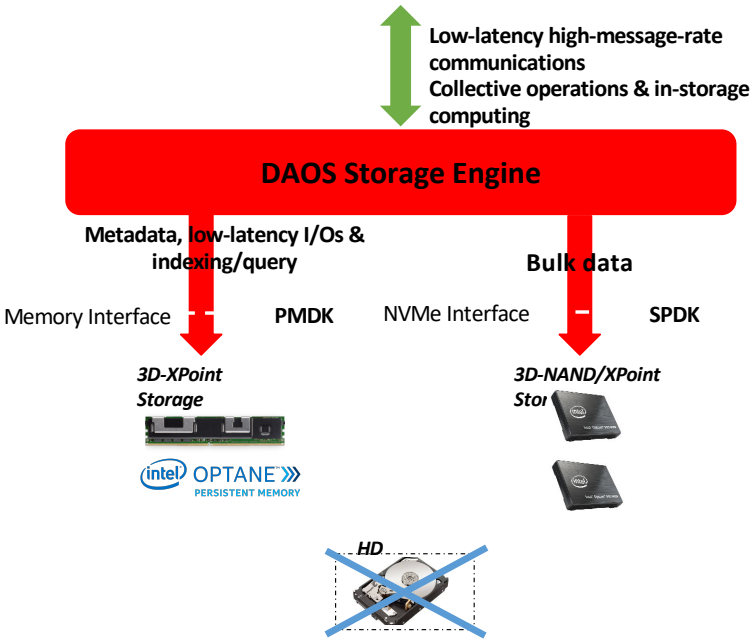

DAOS stack overview

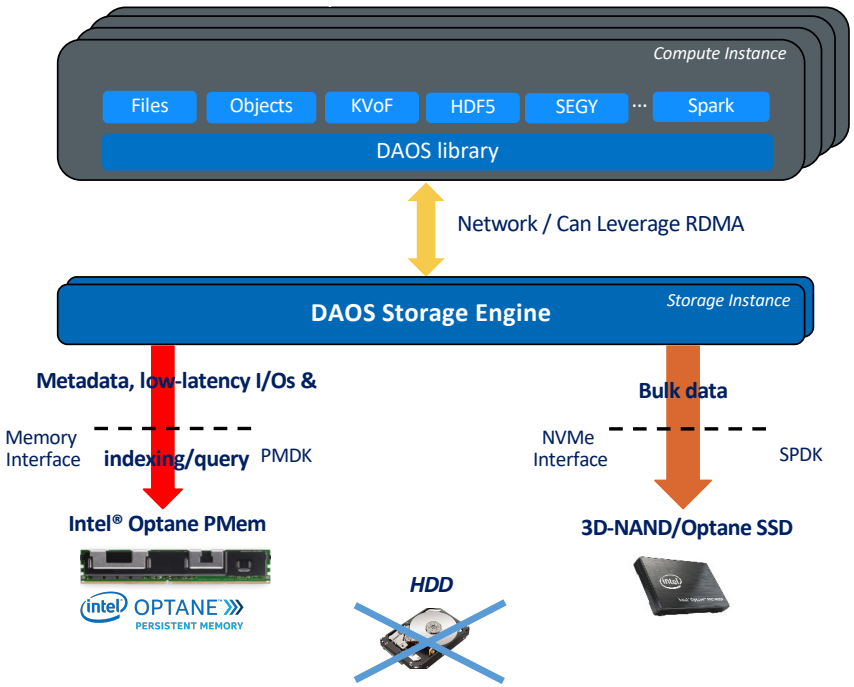

- High **throughput/IOPS @arbitrary** alignment/size (2M IOPS in 1U)
- **Low-latency fine-grained** I/O
- Data access time **orders of magnitude faster** (µs vs ms)
- Highly **scalable**
- Operates in **userspace**
- **Rich** storage semantics
- Support **smooth migration** with support in common frameworks such as Apache Spark*, MPI-IO, HDF, POSIX, etc

DAOS in the Overall Cluster Architecture

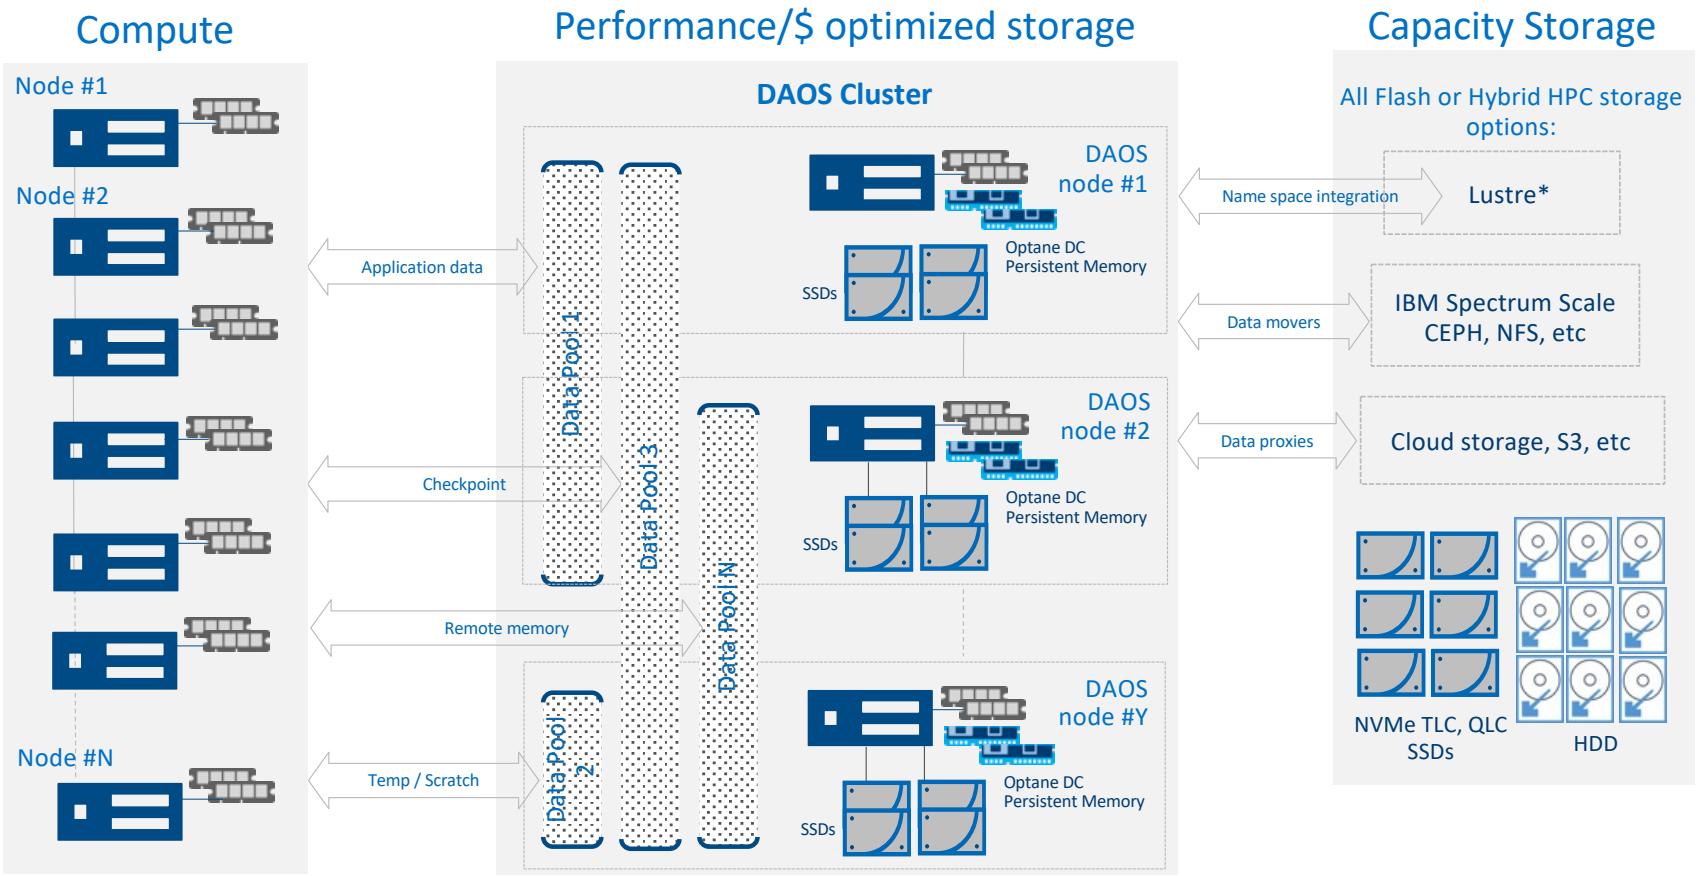

DAOS Node Design

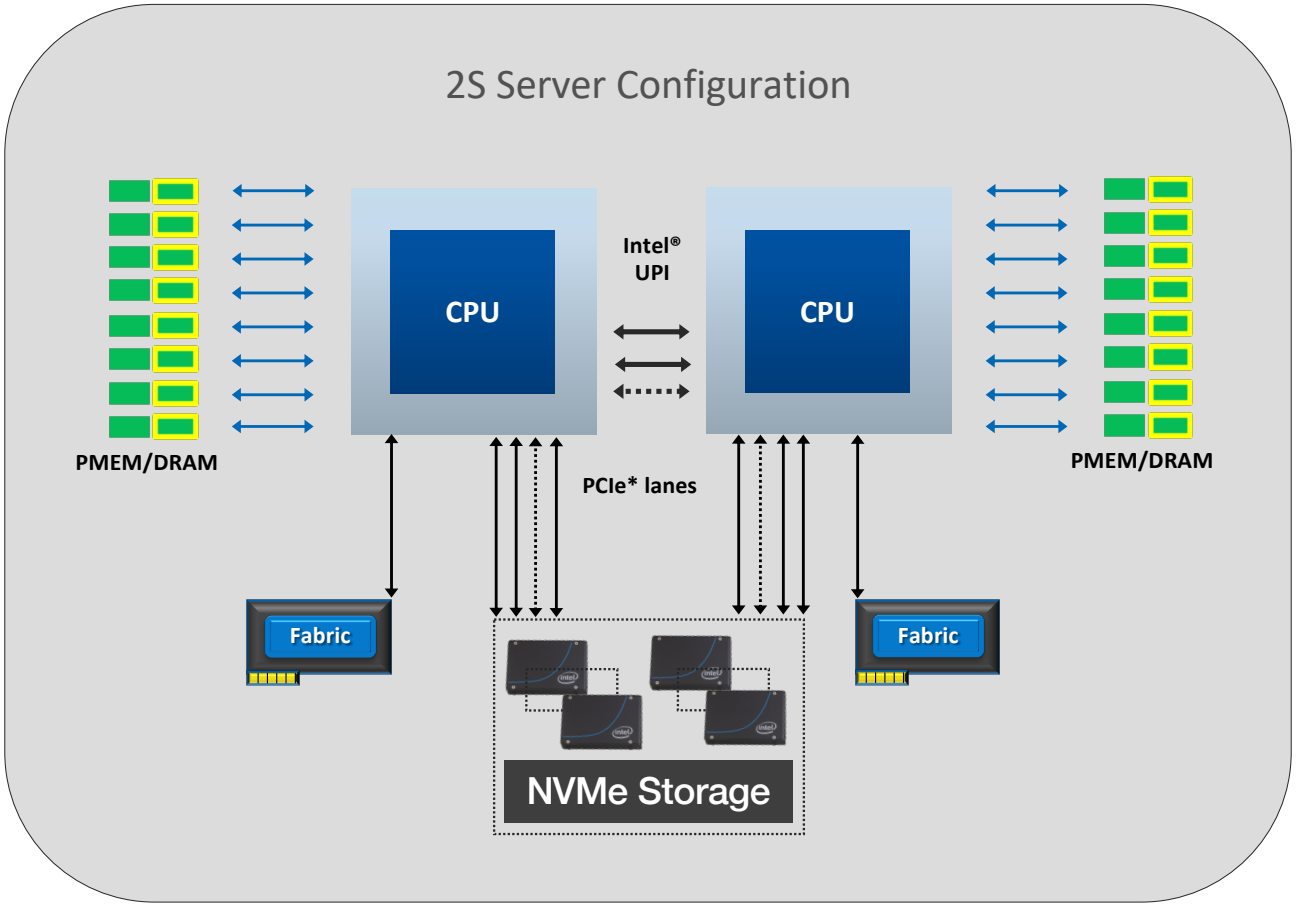

DAOS Node Design

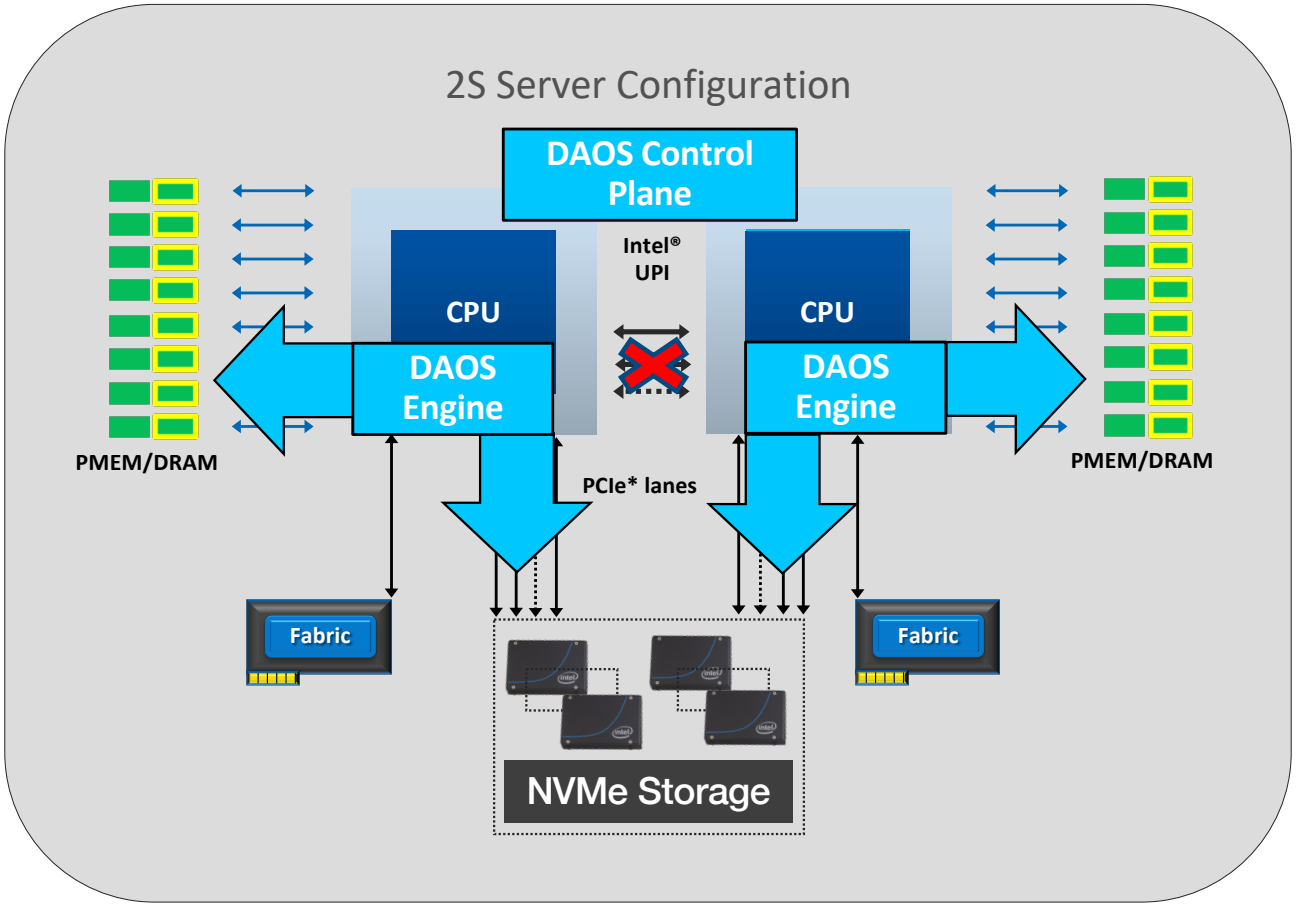

Engine: Media Management

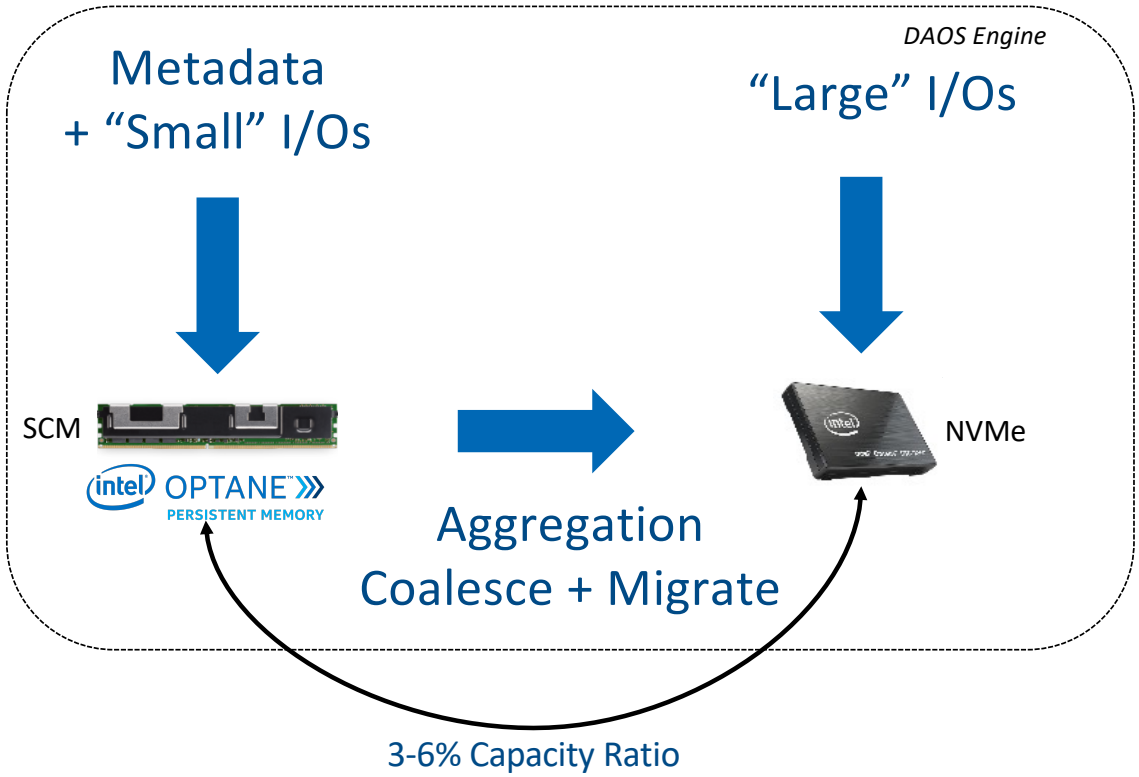

DAOS Data Model: Storage Pooling

Applications
Compute Nodes (CN)

DAOS System
DAOS Nodes (DN)

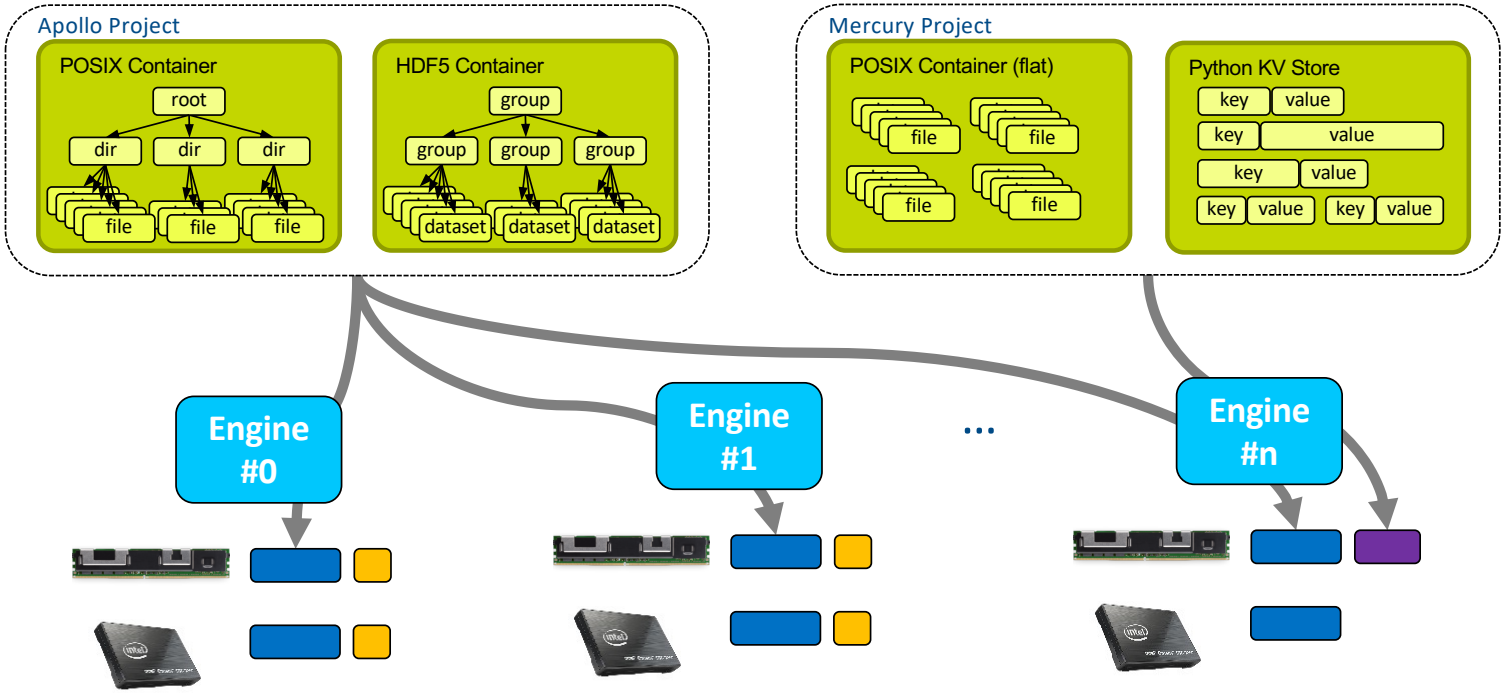

Pool 1	Project Apollo	100PB usable	20TB/s	200M IOPS
Pool 2	Project Gemini	10PB usable	2TB/s	20M IOPS
Pool 3	Project Mercury	30TB usable	80GB/s	2M IOPS

DAOS Data Model: Objects

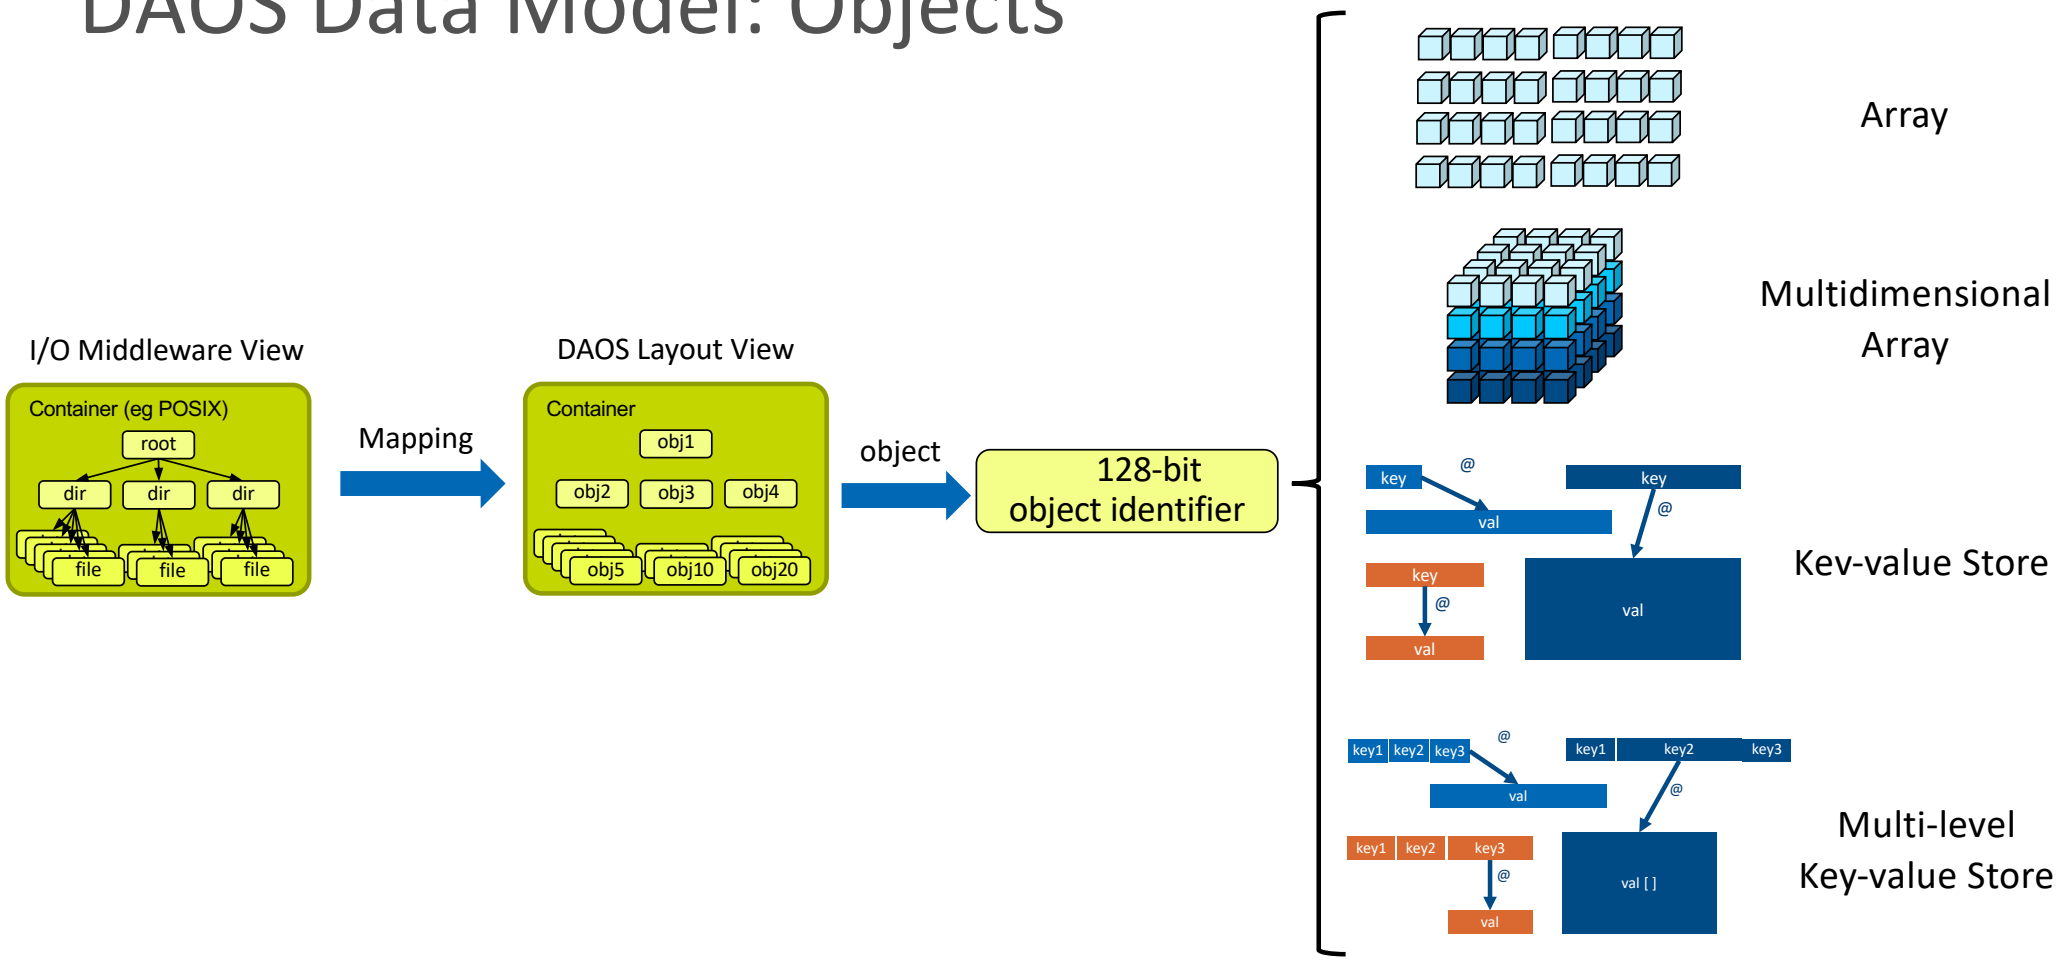

DAOS Data Model: Distribution & Fault Tolerance

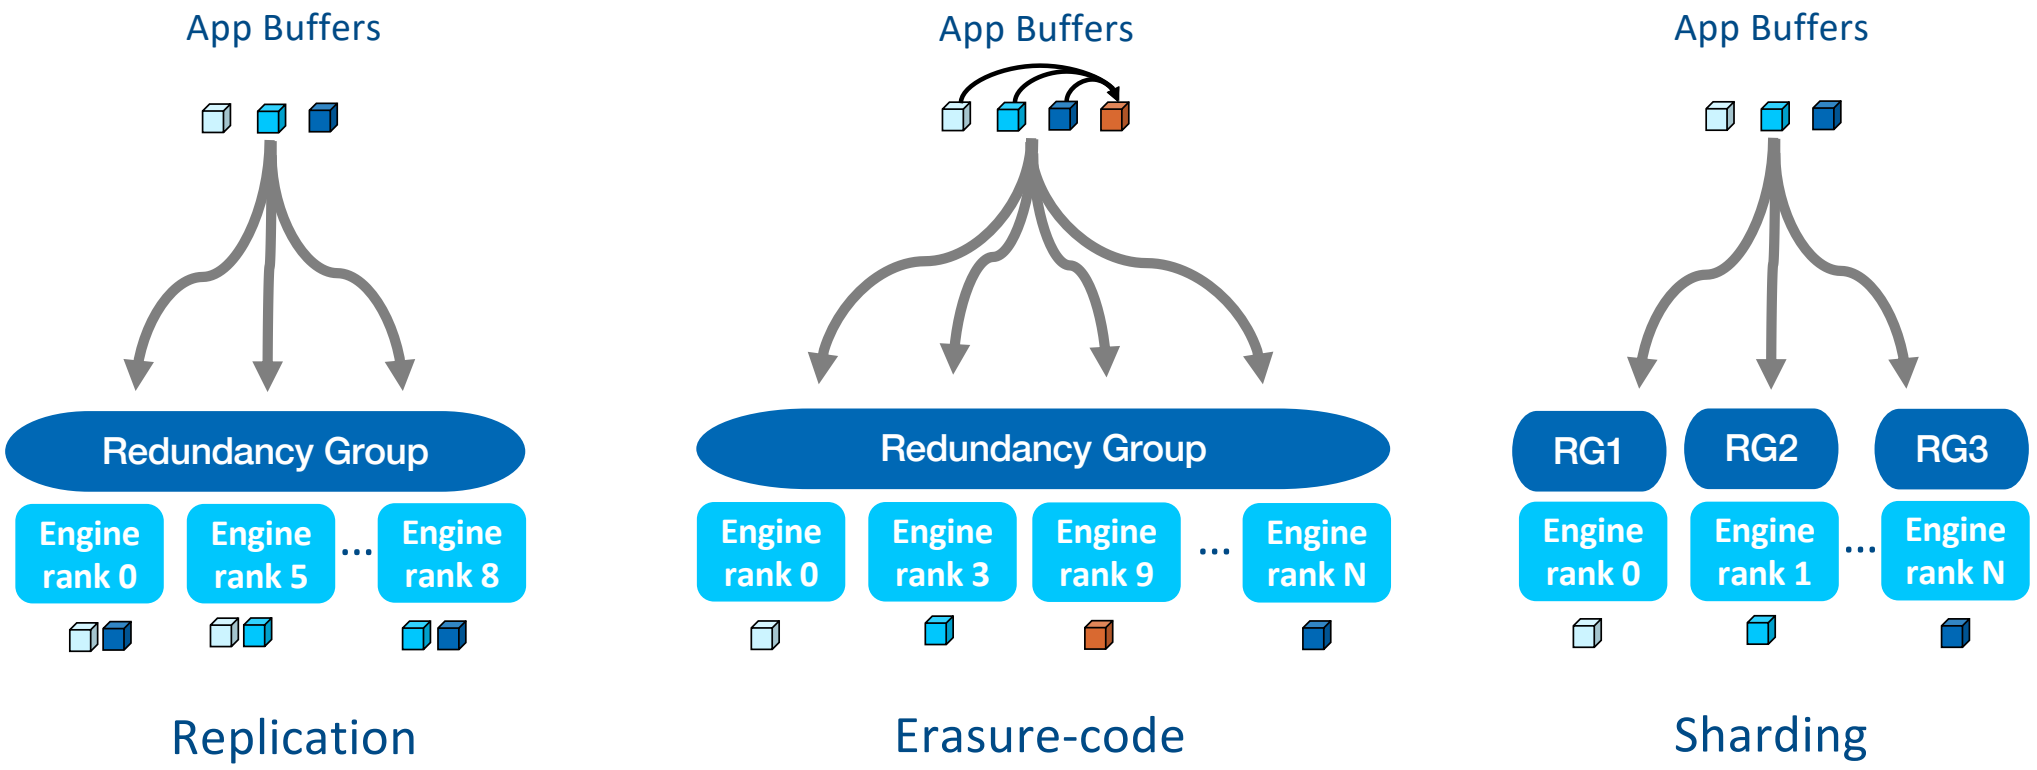

Algorithmic object placement

- No explicit object layout
  - Expensive to maintain
  - Extra round-trip to get the layout
- Object class
  - Data protection method
  - Data distribution requirements
  - Pre-defined attributes, identified by 16-bit integer (class ID)

- Object ID
  - Object class ID + 96-bit ID
- Algorithmic object placement
  - determines where those shards will be stored on the physical system (node, target) based on pool configuration and OID

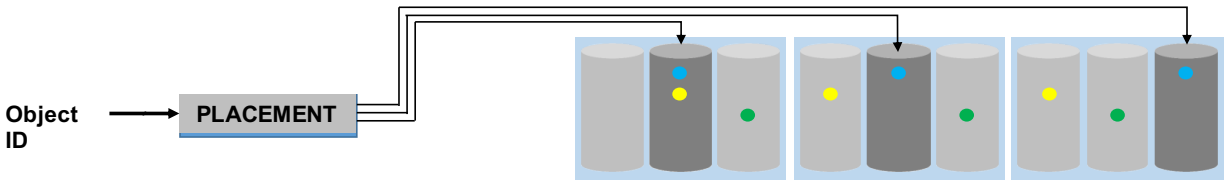

Automated Exclusion & Self-Healing

- Health monitoring
  - Detects failed nodes via SWIM
  - Failed nodes are automatically evicted

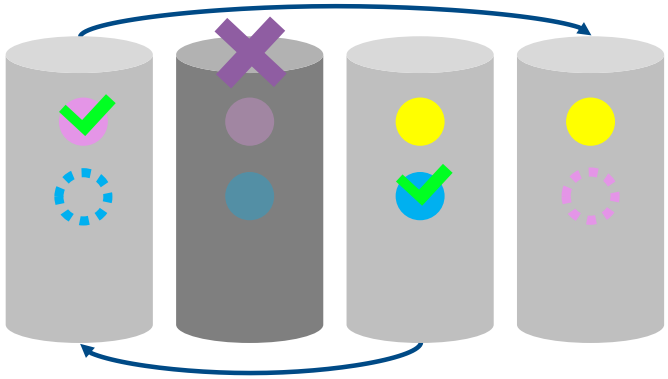

- Failure recovery (“rebuild”)
  - All other surviving nodes are notified of the failure
  - Impacted objects are automatically determined and reconstructed on surviving nodes

Reintegration (“reintegrate”)

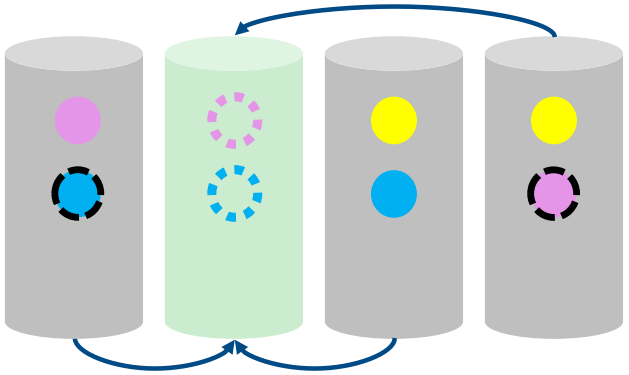

- Restore previously evicted servers/devices to active service
- Data moves back from fallback locations to reintegrated capacity
- Fallback devices service read/write requests until reintegration is complete
- Writes during reintegration also go to reintegrating device to ensure consistency

Extension (“extend”)

- Add new servers/devices to the storage system
- Data rebalances automatically across new capacity

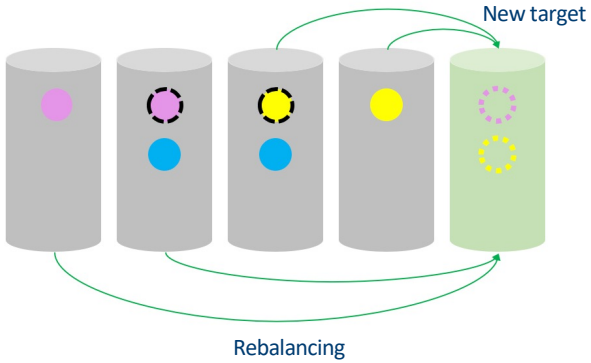

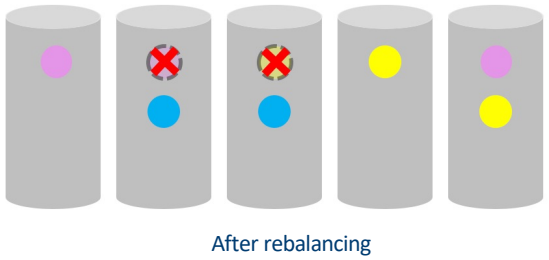

- Original devices service read/write requests until extension is complete
- Writes during extension also go to new devices to ensure consistency

DAOS Software Ecosystem

DAOS Ecosystem

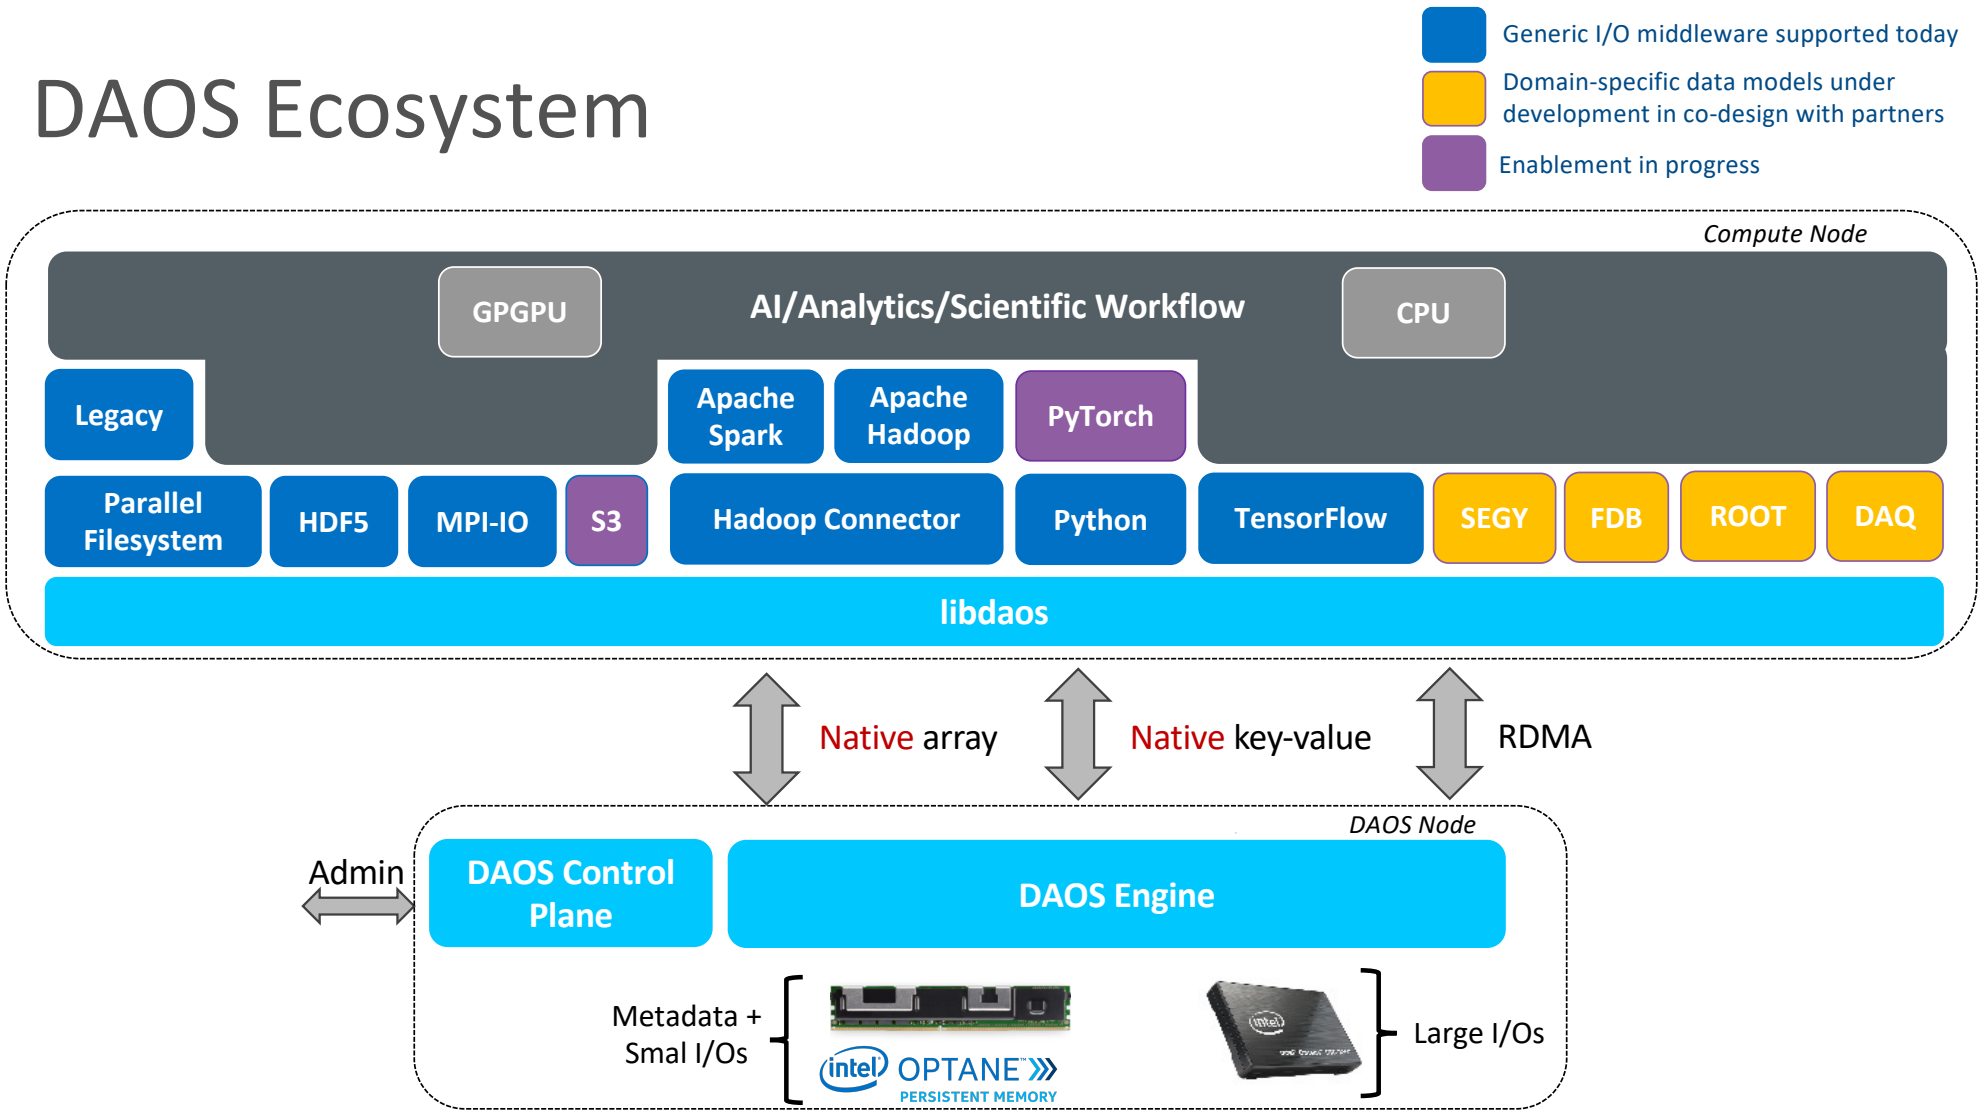

POSIX I/O Support

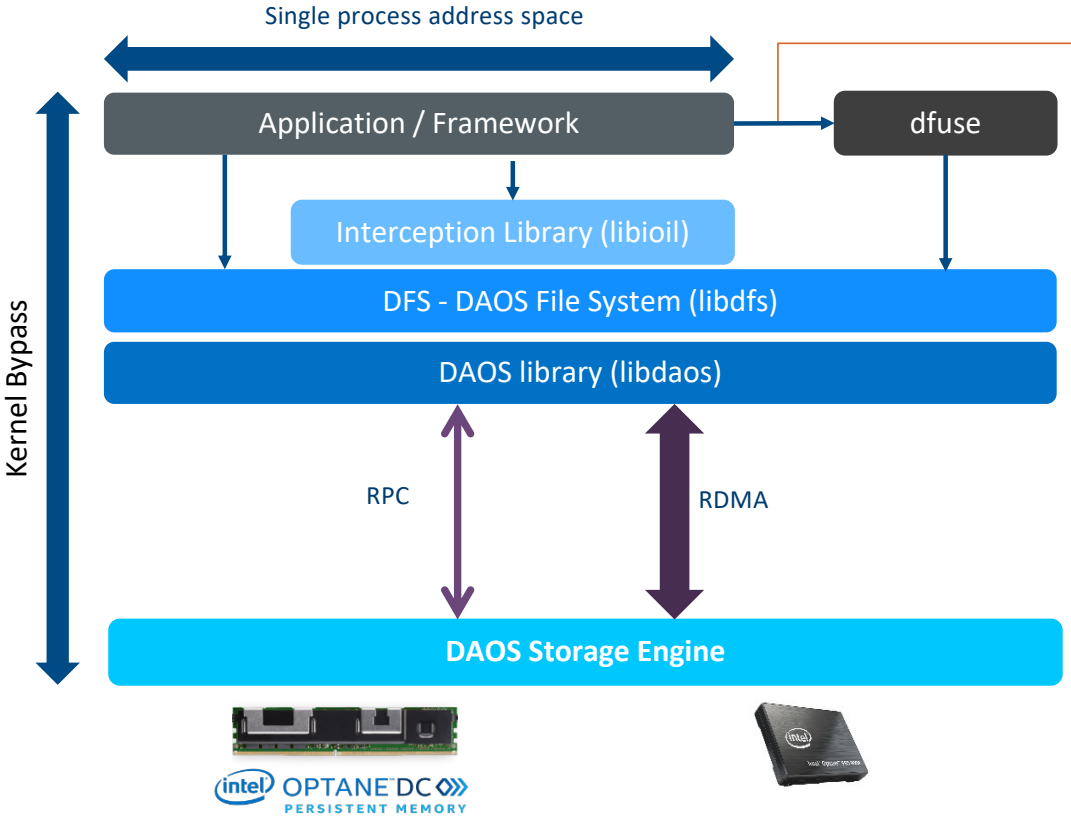

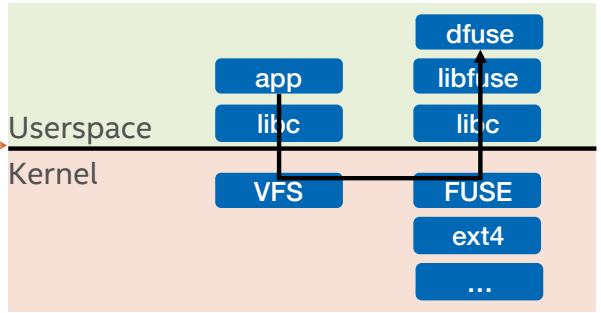

- User space DFS library with an API like POSIX.
  - Requires application changes (new API)
- DFUSE plugin to support POSIX API
  - No application changes
  - Limited performance
- DFUSE + IL
  - No application changes, runtime LD_PRELOAD
  - Good raw data I/O performance, limited metadata performance

DFS API

POSIX	DFS
mkdir(), rmdir()	dfs_mkdir(), dfs_rmdir()
open(), close(), access()	dfs_open(), dfs_release(),dfs_lookup()
pwritev(), preadv()	dfs_read/write()
{set,get,list,remove}xattr()	dfs_{set,get,list,remove}xattr
stat(), fstat()	dfs_stat(),ostat()
readdir() …	dfs_readdir() …

Mostly 1-1 mapping from POSIX API to DFS API.

Instead of File & Directory descriptors, use DFS objects.

All calls need a DFS mount which is usually done on initialization with the pool / container access handles.

PyDAOS Primer

- python module primarily written in C
  - Expose DAOS key-value store objects as a python dictionary
    - Support python iterator, direct assignments, …
    - Bulk insert/retrieve
  - Other data structures are under consideration (see later)
- Python objects allocated by PyDAOS:
  - are **persistent**
    - identified by a string name
  - are immediately **visible** upon creation
    - to any process running on the same or a different node.
  - have a **very low memory footprint** since the actual content is stored remotely
    - This allows to manipulate gigantic datasets way bigger than the amount of memory available on the node

TensorFlow Integration

- Done via TensorFlow-IO (see https://www.tensorflow.org/io)
- Initial integration at the DFS level
  - Use TF-IO filesystem API
  - Compatible with POSIX container
  - Provide full OS bypass
- Development completed & testing underway
  - https://github.com/daos-stack/tensorflow-io-daos/tree/devel
  - https://github.com/daos-stack/tensorflow-io-daos/blob/devel/docs/daos_tf_docs.md

DAOS Performance

DAOS Bandwidth on IO500

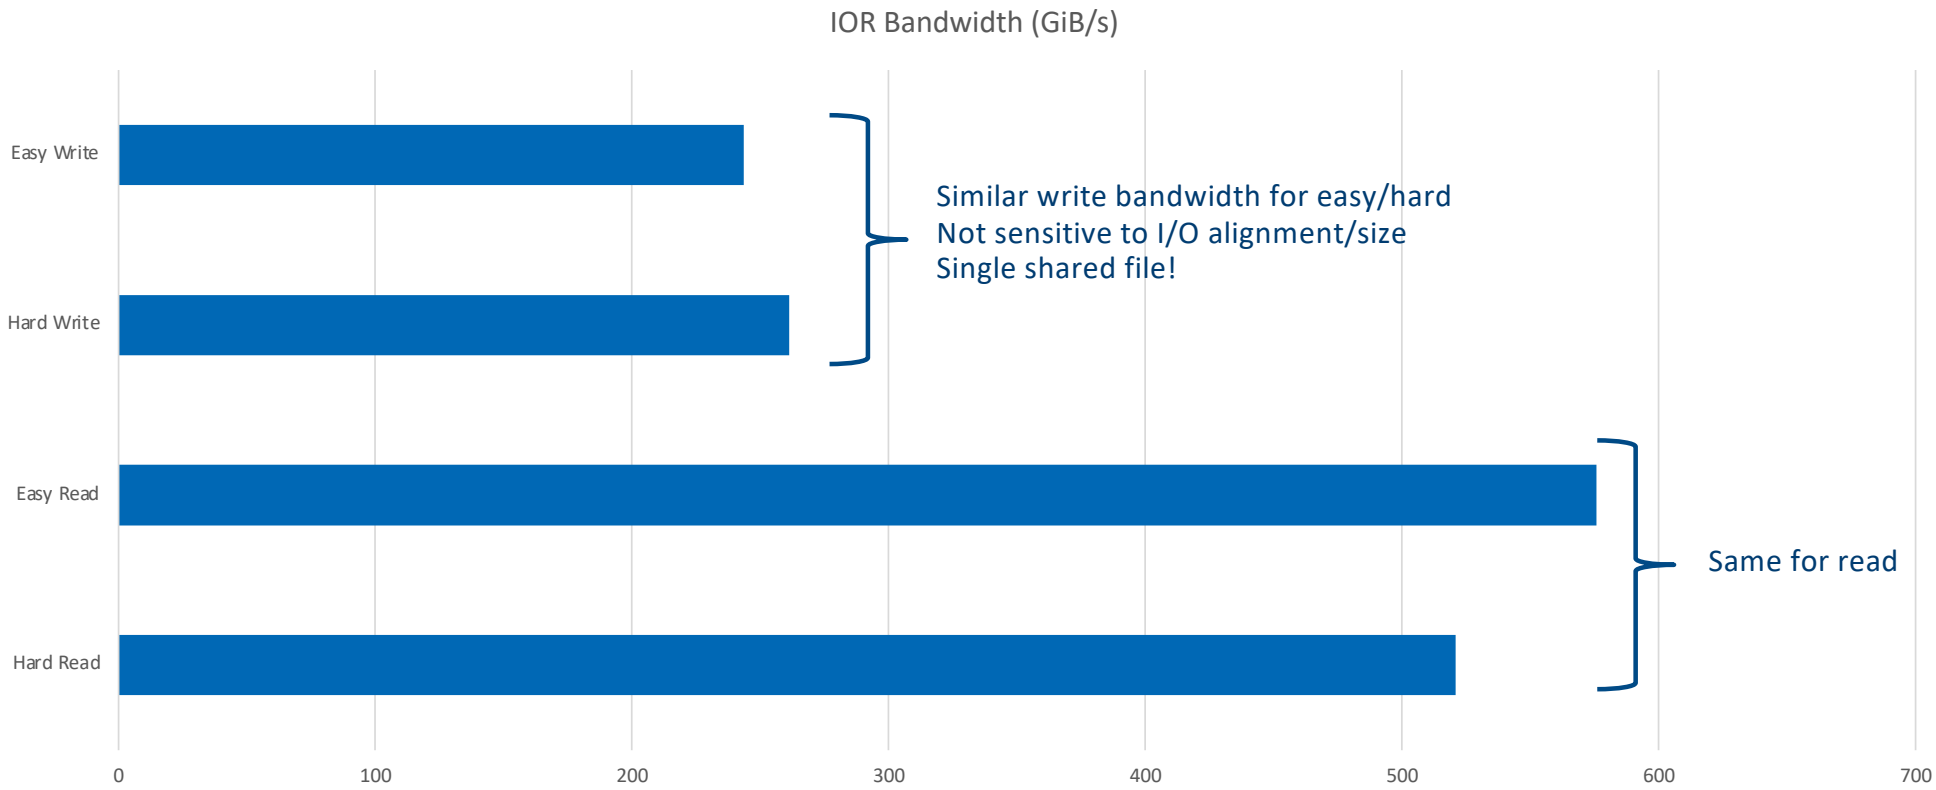

Source: https://io500.org/submissions/view/1

DAOS Metadata Performance on IO500

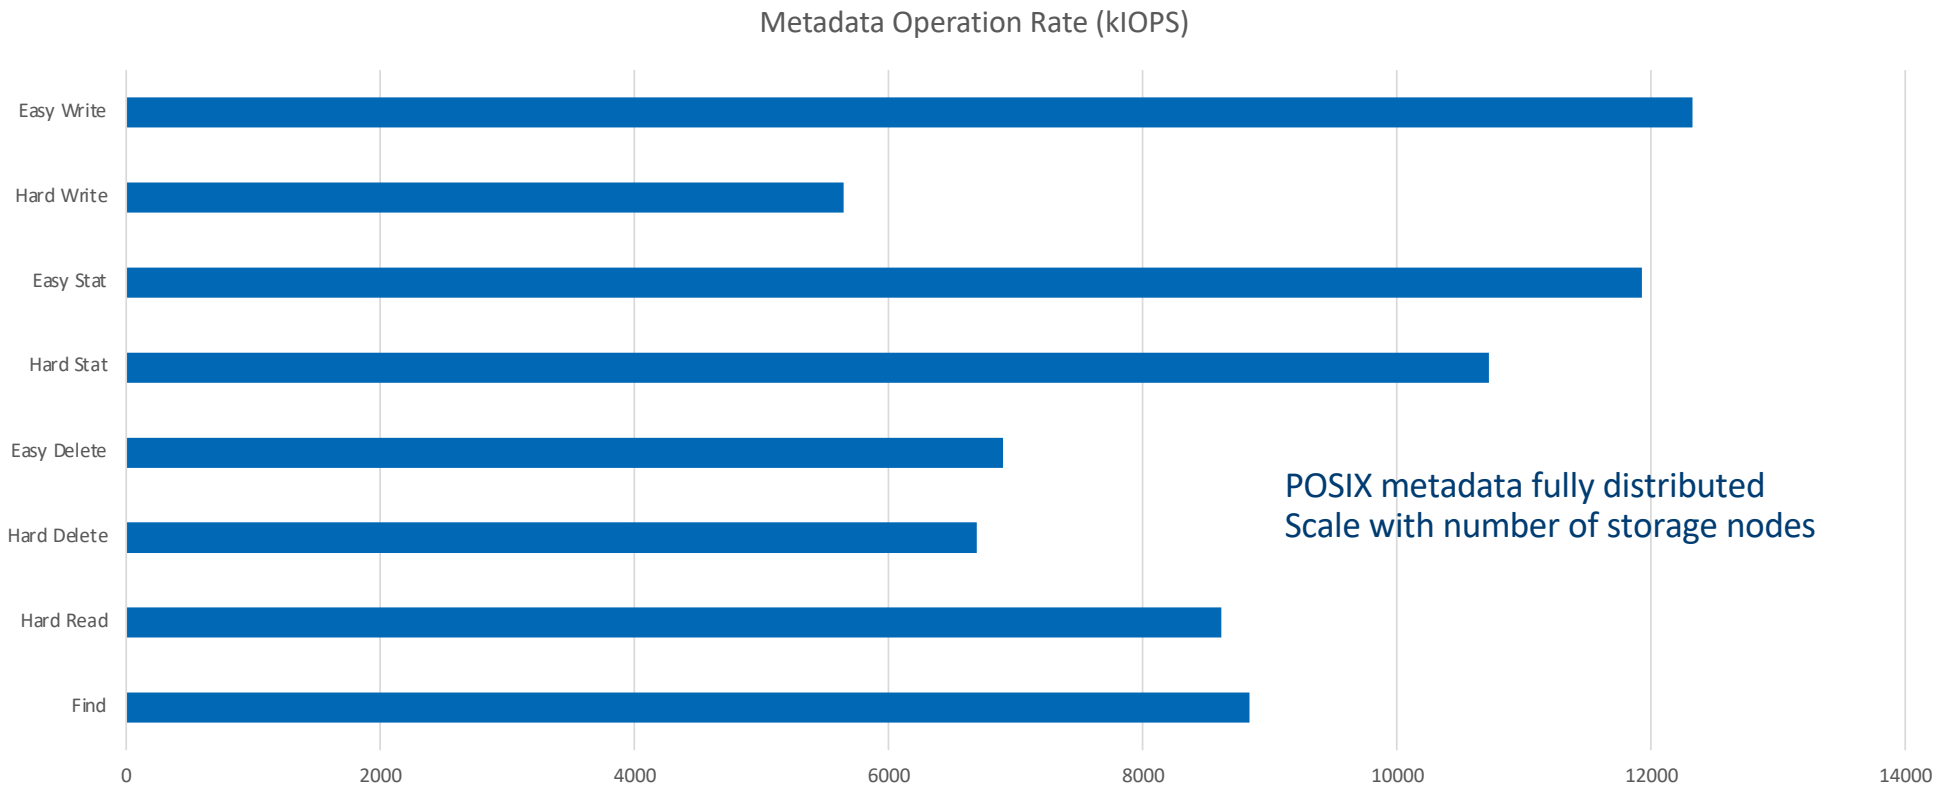

Source: https://io500.org/site/submissions/view/1

Resources

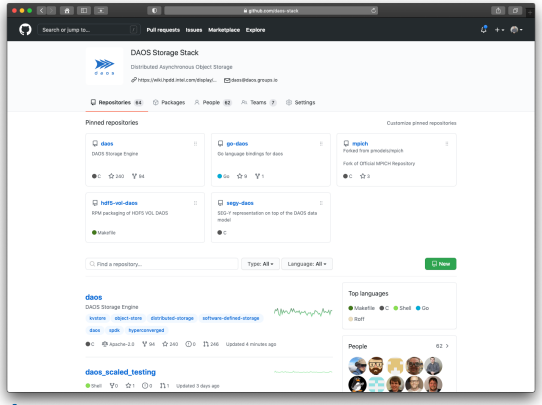

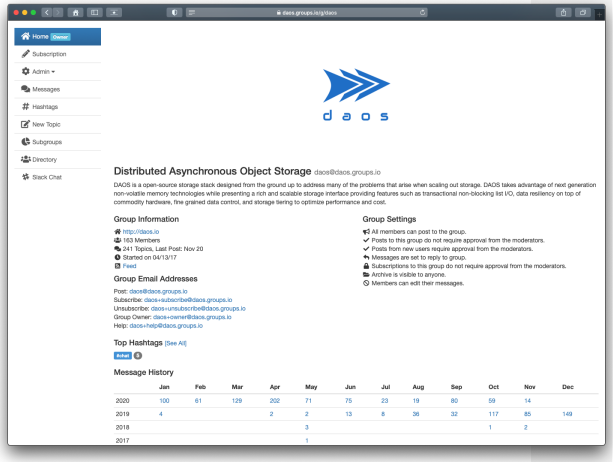

- Community Resources
  - Github: https://github.com/daos-stack/daos
  - Online doc: http://daos.io
  - Mailing list & slack: https://daos.groups.io
  - YouTube channel: http://video.daos.io
- 5th DAOS User Group (DUG'21)
  - Recordings available at http://dug.daos.io
- Intel landing page
  - https://www.intel.com/content/www/us/en/high-performance-computing/daos.html

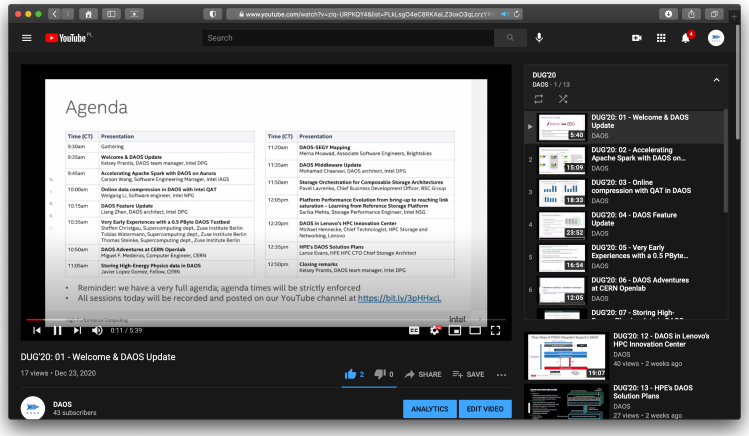

Q & A


intel