

# Accelerating Data-driven Workflows with DAOS

Johann Lombardi, Senior Principal Engineer, AXG, 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.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

# Agenda

- Project History
- DAOS one-slide Introduction
- 4x Architecture Choices
- 1x Future Direction
- Partner Ecosystem
- DAOS on Aurora
- Resources

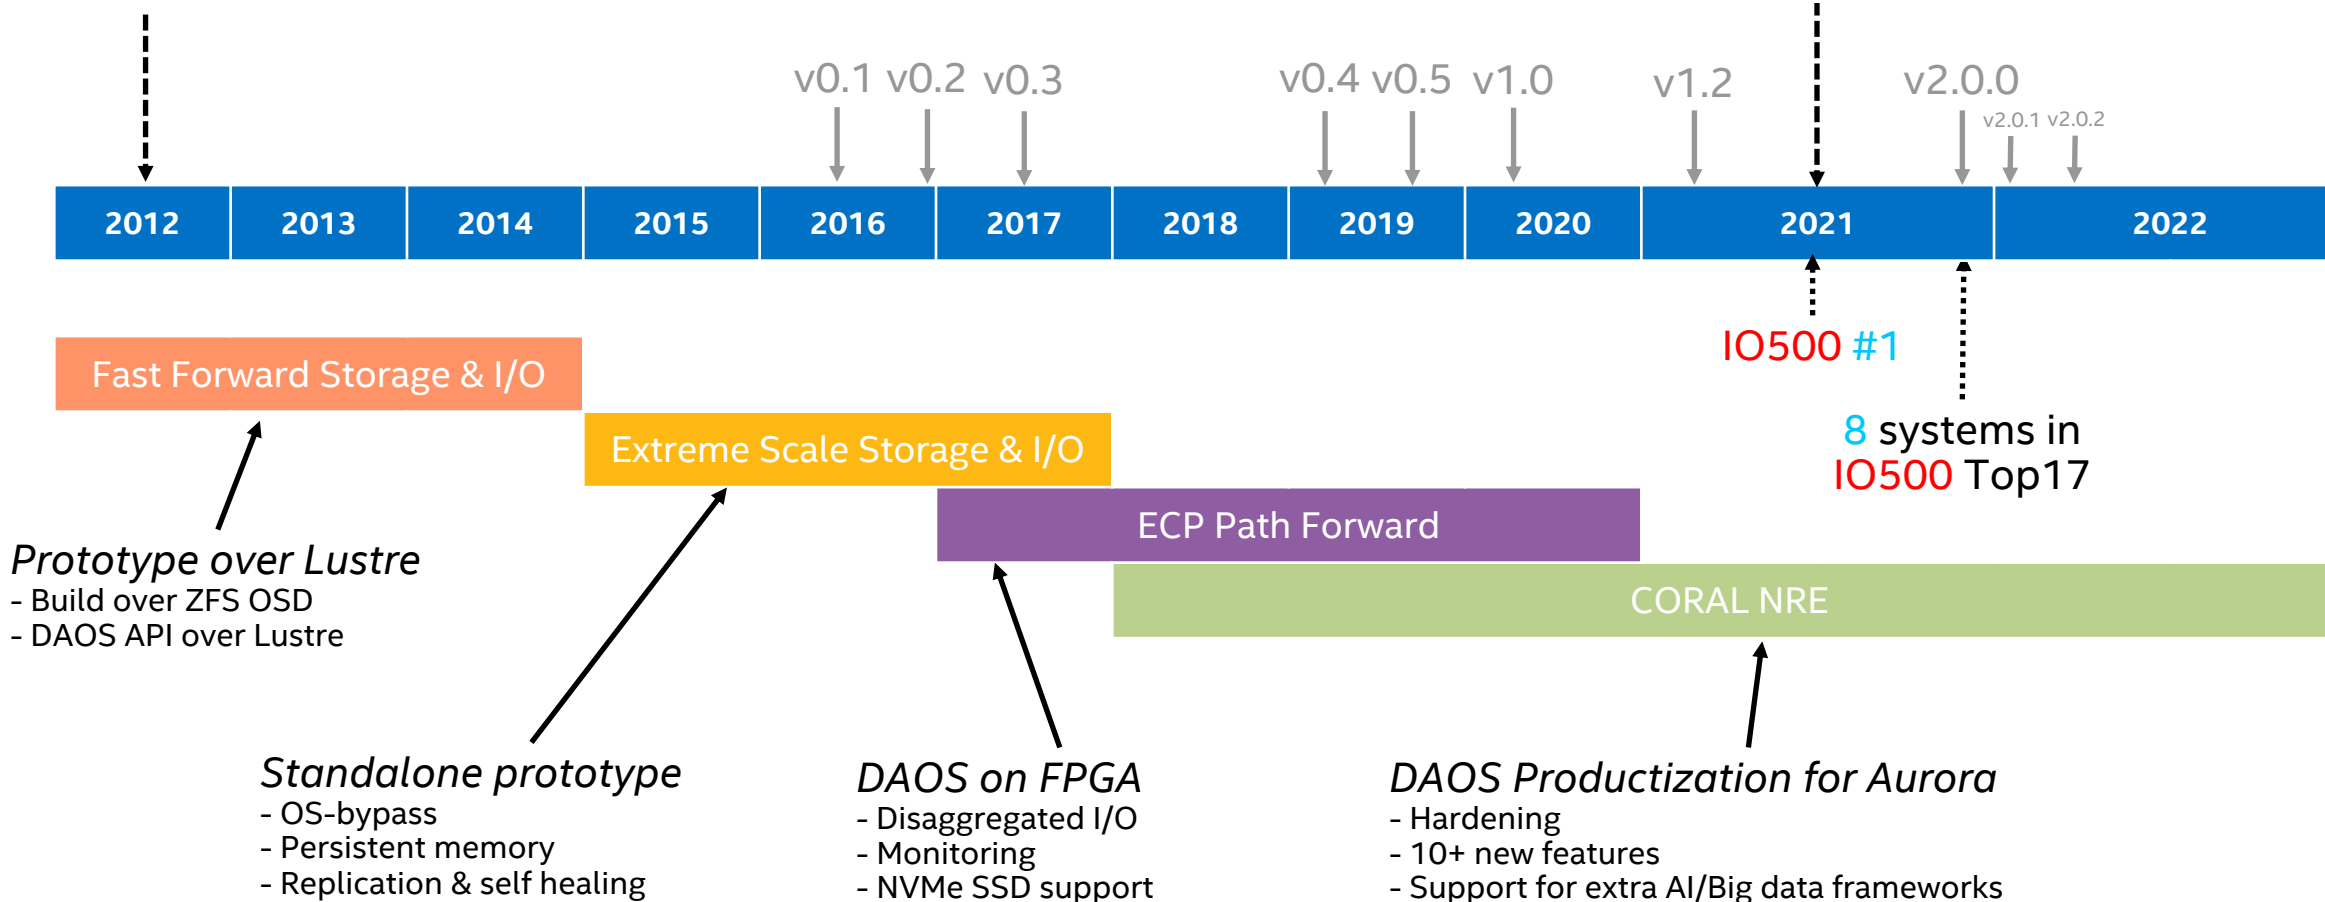
# DAOS History

IO500



Intel **acquires**  
Whamcloud

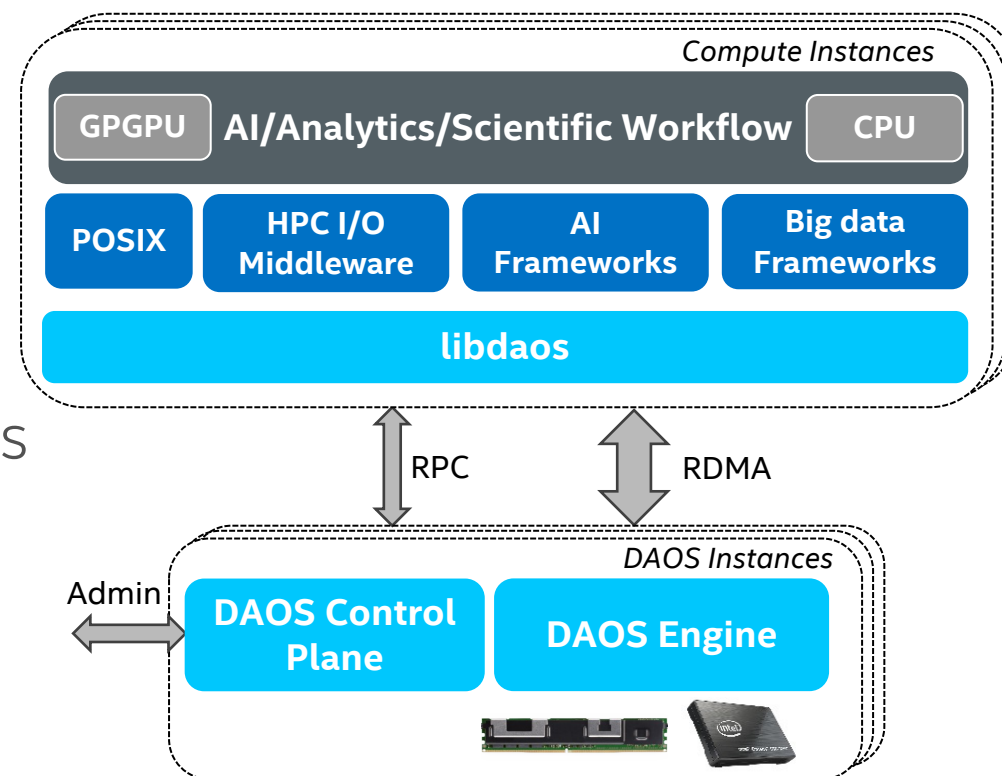
Intel offers  
**L3 Support**



# DAOS: Nextgen Open Storage Platform




- Fully Distributed multi-tenant global namespace
- Platform for innovation
  - Modular API and layering
  - Can leverage latest HW & SW technology
- Built for high performance
  - 10's  $\mu$ s latency, billions of IOPS, TB/s to PB/s
- Full userspace model
  - Run on-prem or in the cloud
- Growing open-source community




# Arch: Storage Management


- New **data model** to unwind 30+y of file-based management
- Introduce notion of dataset: **DAOS container**
  - Basic unit of storage
  - Containers have a type
  - POSIX containers can have billions of files/directories
  - Many parameters
  - ACLs
- Advanced container **query** capabilities




POSIX Namespace




File-per-process




HDF5 « File »



Key-value store



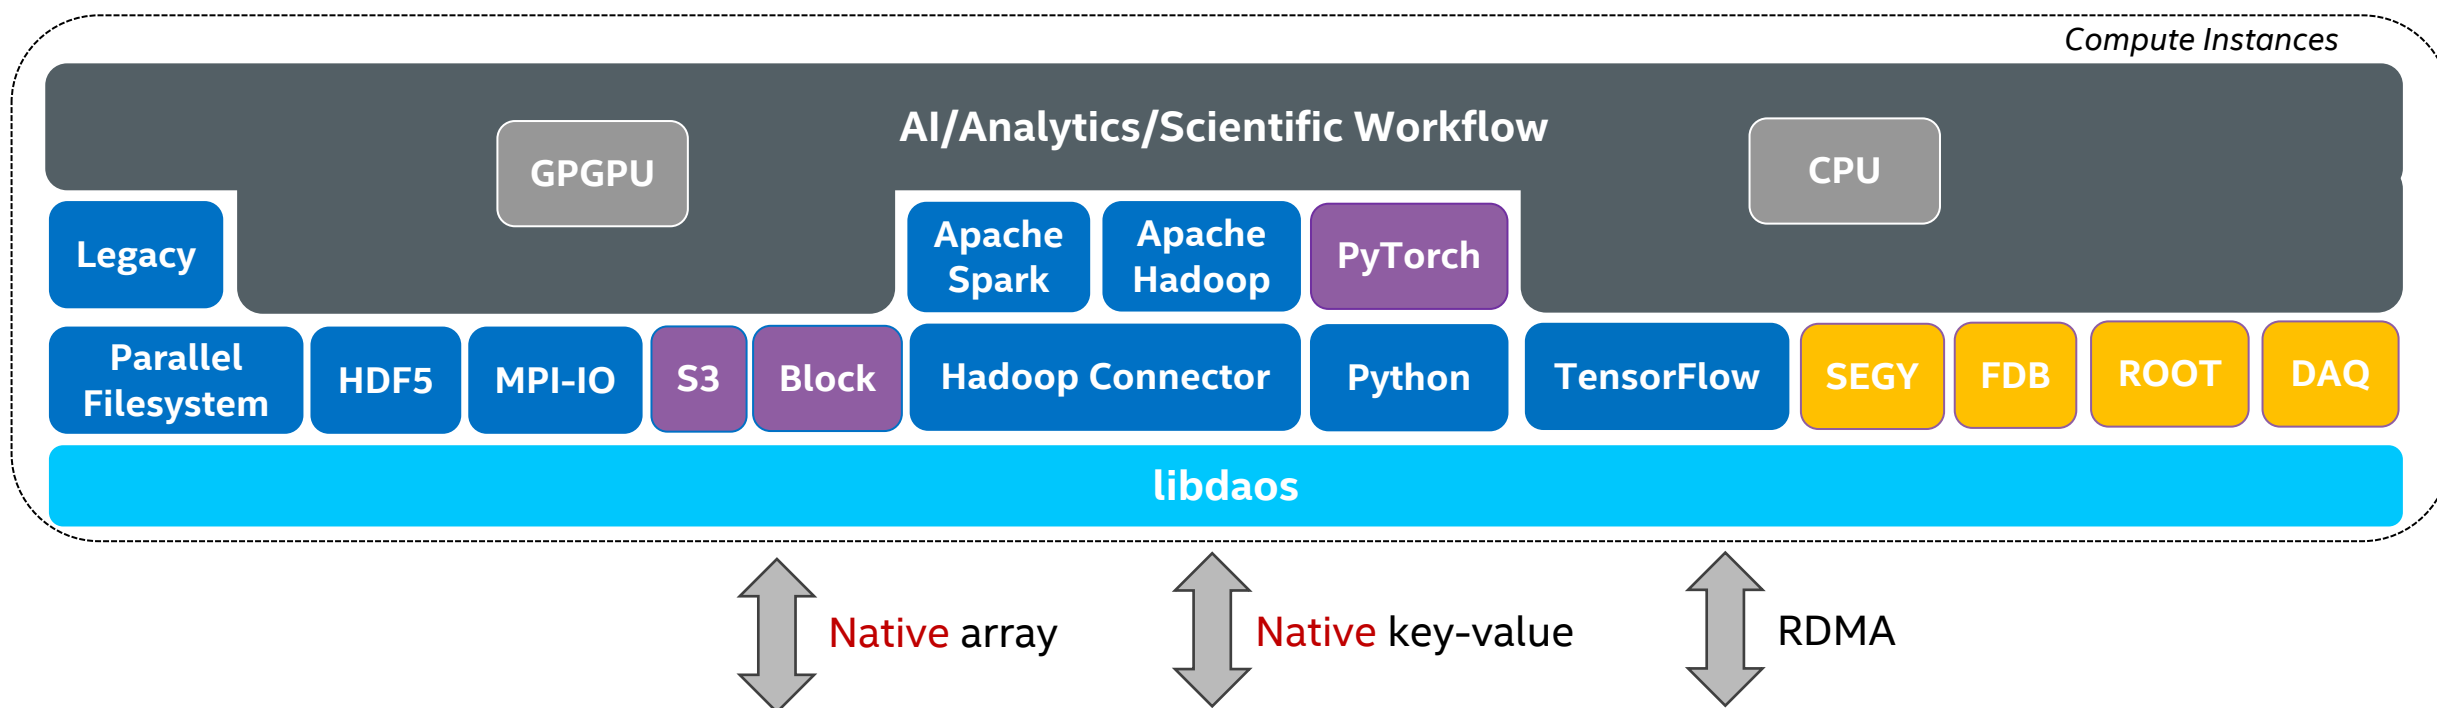
S3 Bucket



Graph

# Arch: Multi-API

- Generic I/O middleware supported today
- Domain-specific data models under development in co-design with partners
- Enablement in progress



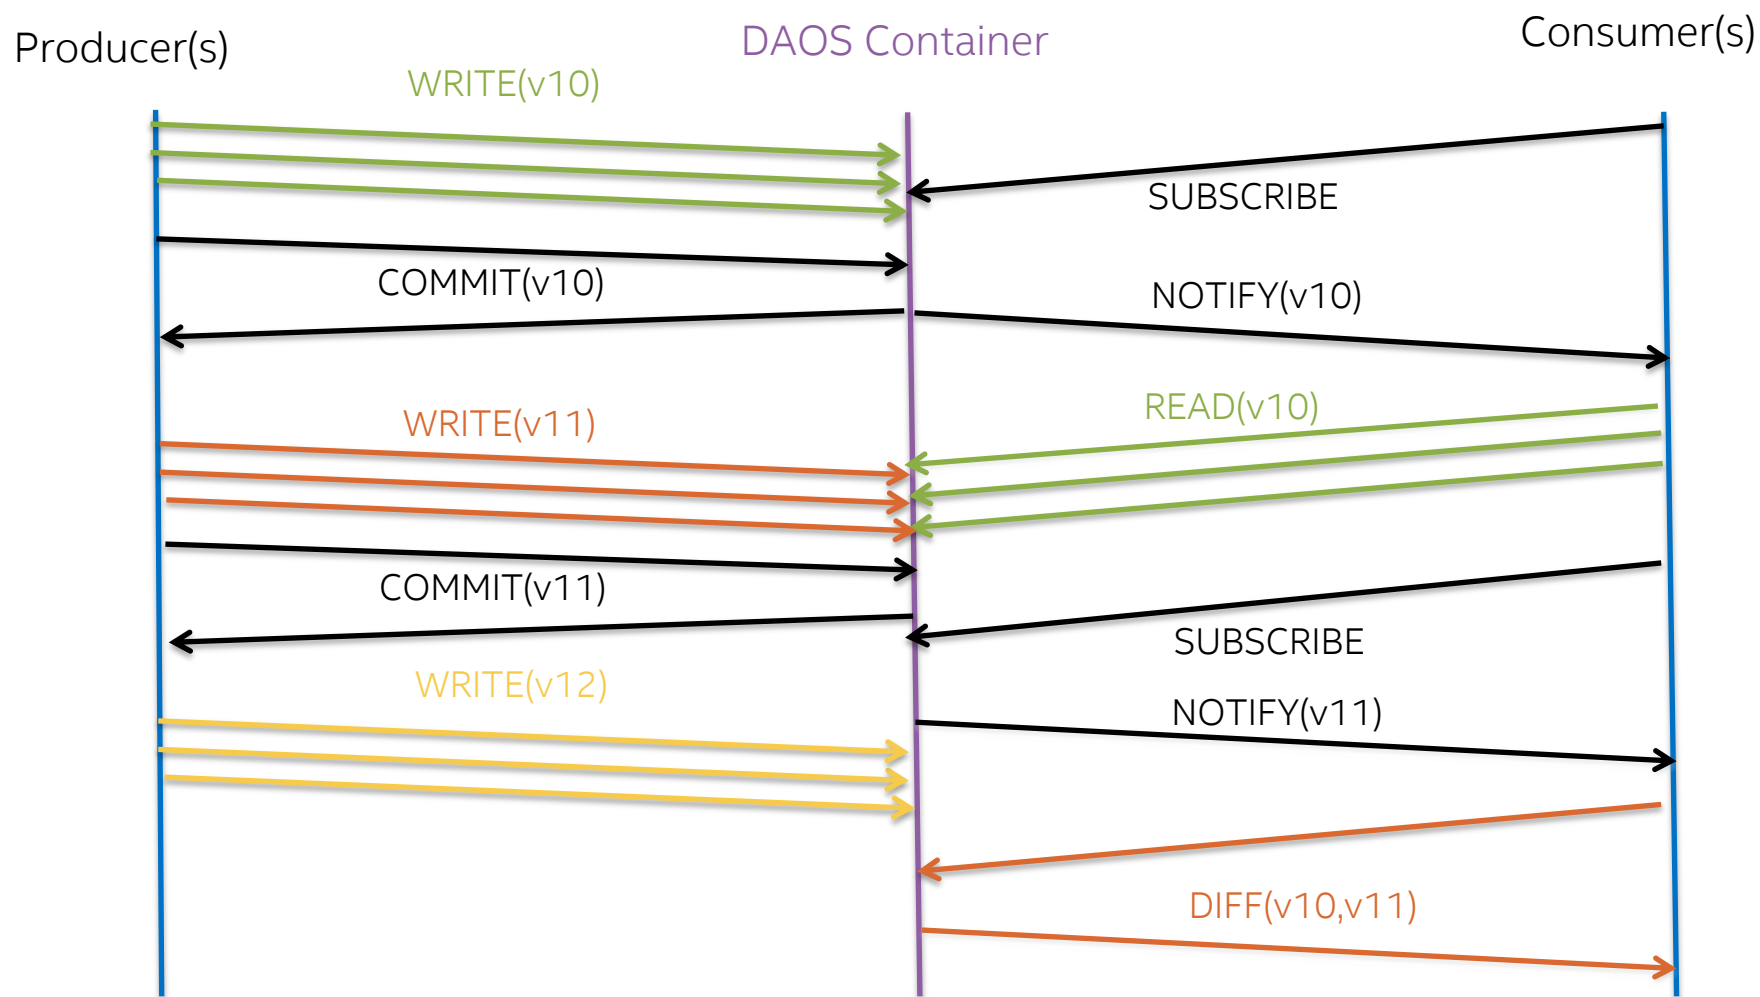
# Arch: AI Acceleration



- Low-**latency** & high **IOPS/BW** reads
  - Can read from any replica
  - No parity check on read (use checksum to validate data)
- Tight **integration** with AI frameworks
  - PyTorch, TensorFlow-IO
- WORM **containers**: optimizations for read-only datasets
  - Aggressive caching
  - Read-optimized layout
  - Size optimizations for immutable files

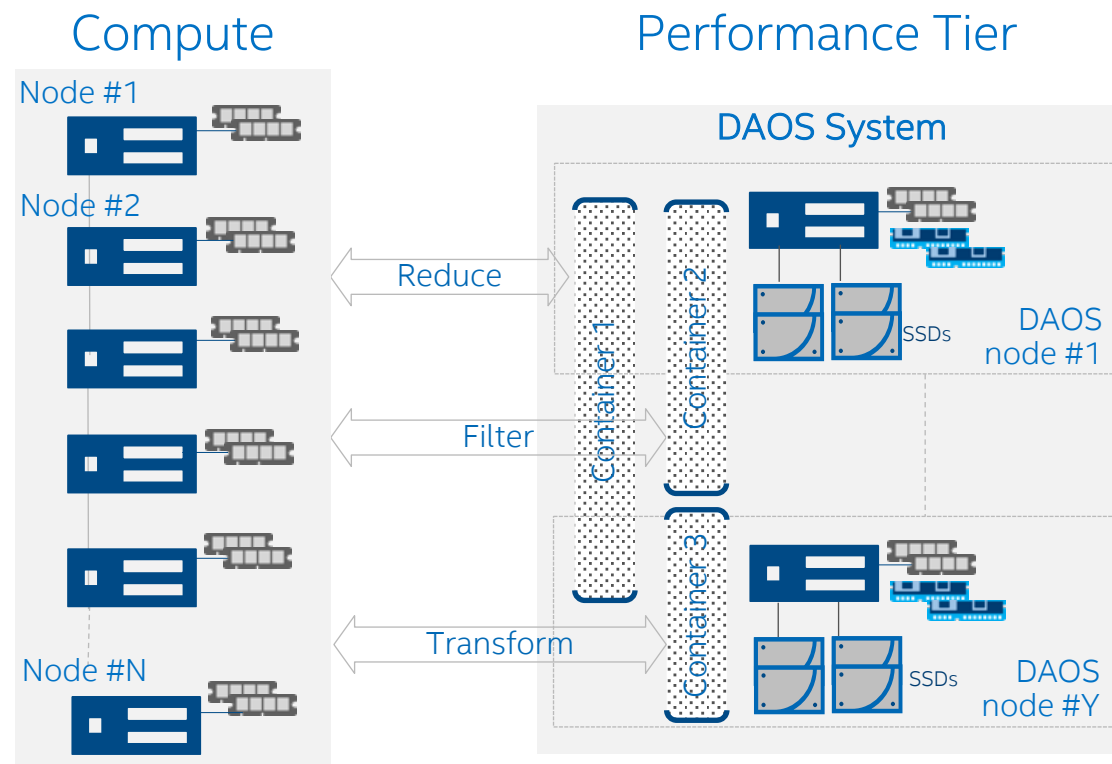


# Arch: Workflow Acceleration



# Direction: Scale-out Active Storage

- DAOS **pipeline** API
  - Offload data-intensive processing to storage
  - Pre-defined or user-defined (ubpf)
- Leverage **HW** acceleration
  - computational storage devices
  - accelerators/smartNICs
- Many **use cases**
  - POSIX find(1)
  - SQL query / MariaDB prototype using predicate push-down
  - In-place data filtering/pre-processing/transformation for AI frameworks
  - Calculate max/min/sum/avg ... or searching for specific pattern/value on metadata/data



# DAOS Partner & Customer Ecosystem

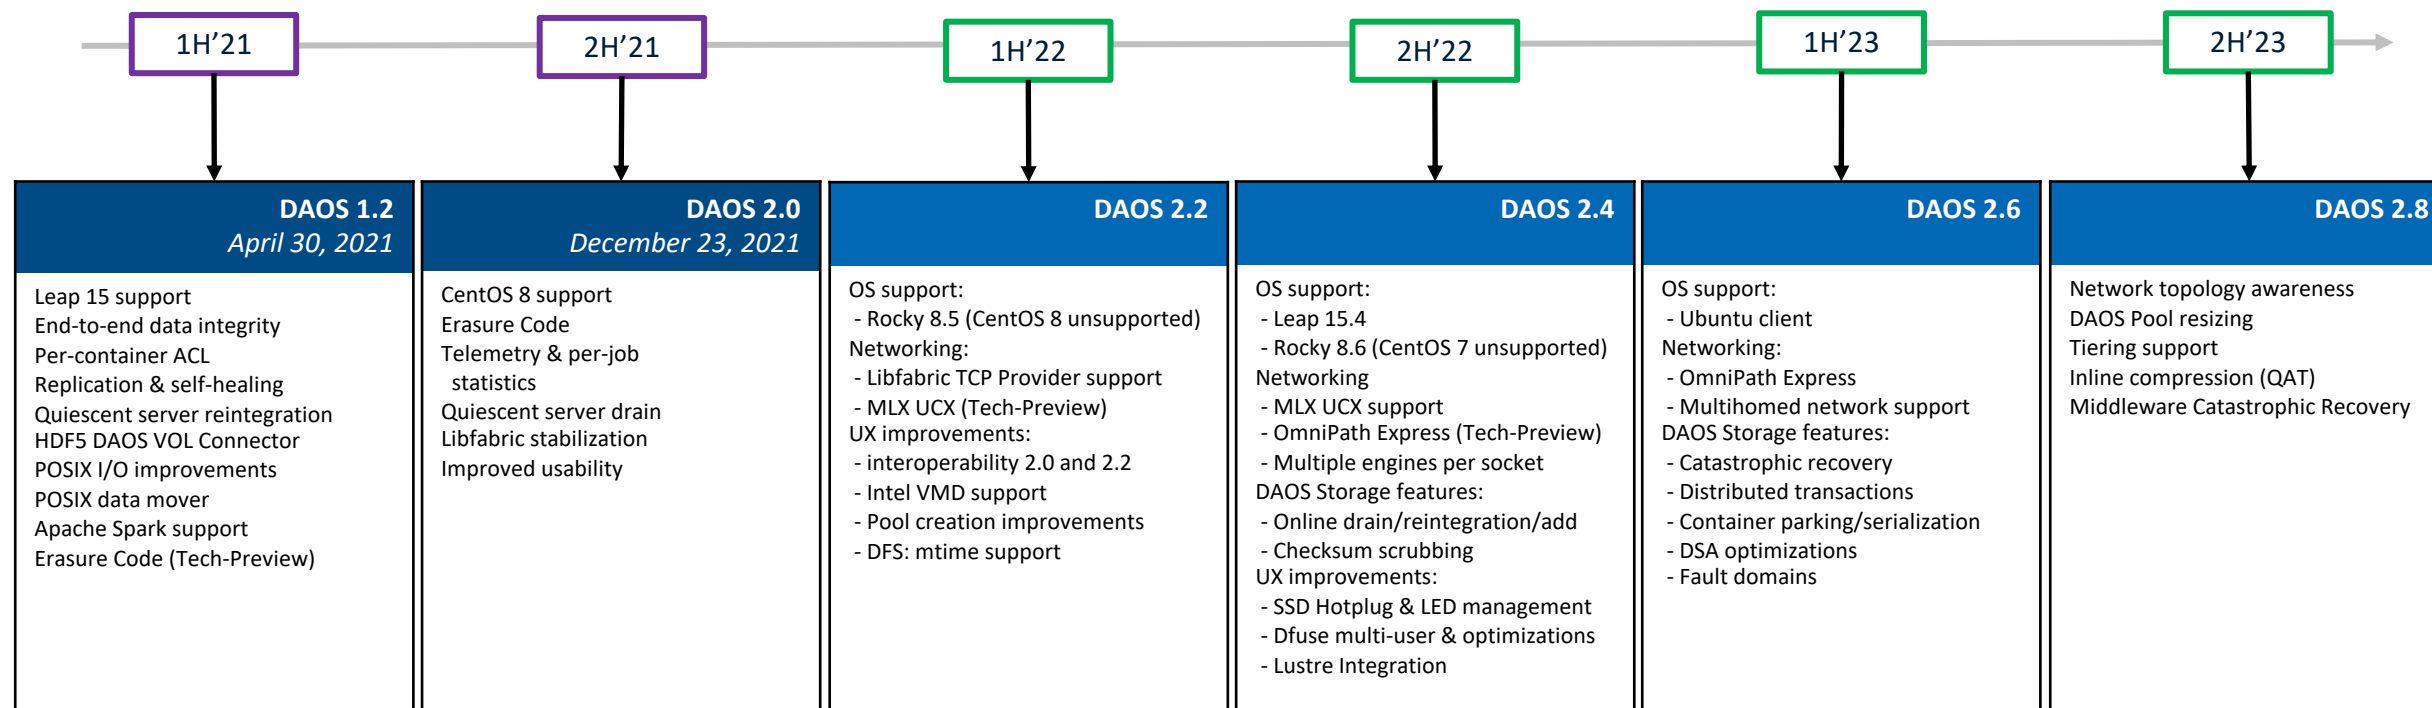
Hardware Partners	       
Reseller Partners	     
Software Development and 3 <sup>rd</sup> party support	     
End Customers	    

# DAOS on Aurora

- 1024x DAOS nodes, each with:
  - 2x Xeon 5320 CPUs
  - 512GB DRAM
  - 8TB Optane Persistent Memory 200
  - 244TB NVMe SSDs
  - 2x HPE Slingshot NIC
- Usable capacity
  - between 220PB and 249PB depending on redundancy level chosen



# DAOS Community Roadmap



**NOTE: All information provided in this roadmap is subject to change without notice.**

# Resources



## ■ Open-source Community

- 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>

## ■ 5<sup>th</sup> DAOS User Group (DUG'21)

- Recordings available at <http://dug.daos.io>

## ■ Upcoming BoFs at ECP Annual meeting & ISC

## ■ Intel landing page

- <https://www.intel.com/content/www/us/en/high-performance-computing/daos.html>

