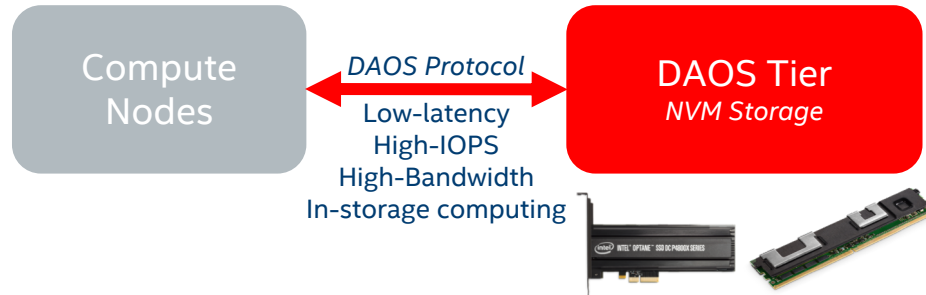




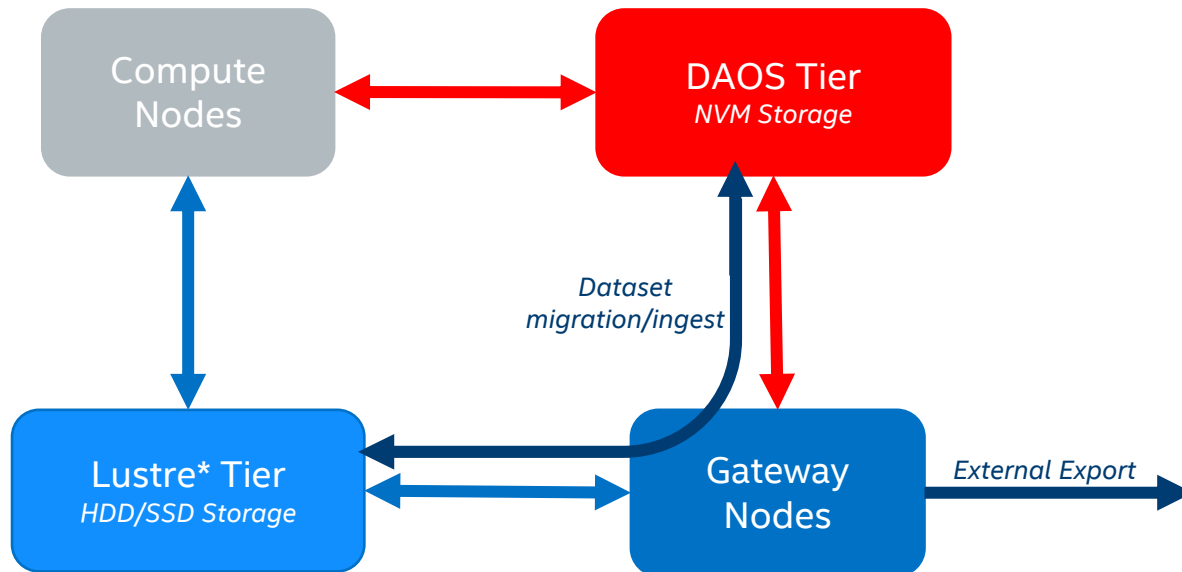
THE OPEN-SOURCE DAOS STORAGE ENGINE

Johann Lombardi, Principal Engineer, Intel

STORAGE ARCHITECTURE



STORAGE ARCHITECTURE



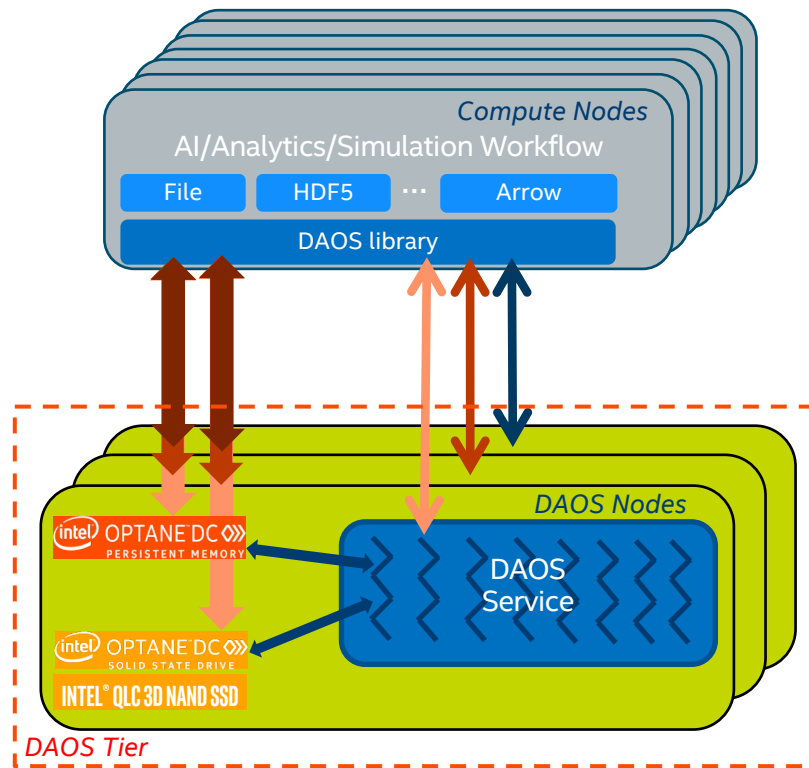
DAOS TIER ANATOMY

DAOS Tier

- Globally accessible from any compute nodes
- Large capacity (100's PB)

DAOS Nodes

- COTS Xeon servers running the DAOS service
- RNIC attached for communications
 - Support multiple RNICs per server to sustain backend storage IOPS/bandwidth
- Mix of storage technologies attached
 - Optane DC Persistent Memory (DCPM)
 - NVMe SSD (*NAND, Optane SSDs)



DAOS VALUE PROPOSITION

Unprecedented storage **performance**

- up to 100's TB/s
- up to billions of IOPS

Enable new **revolutionary** storage models


- Generic key-value/array object containers support multiple high-level I/O & database APIs
- New storage paradigm providing advanced features like native index/query, dataset snapshot, asynchronous I/O, ...

Provide more **resilient** storage

- Distributed transaction exported through the API to guarantee data model consistency
- Ultra-fast storage rebuild that is fault-domain aware and guarantee end-to-end data integrity

Deliver **open-source** storage platform for AI, Big Data and Simulation with built-in accelerator support

- Native support for key-value store, Spark, HDF5, TensorFlow, files, blocks ...



DATA ANALYSIS

Fast data retrieval

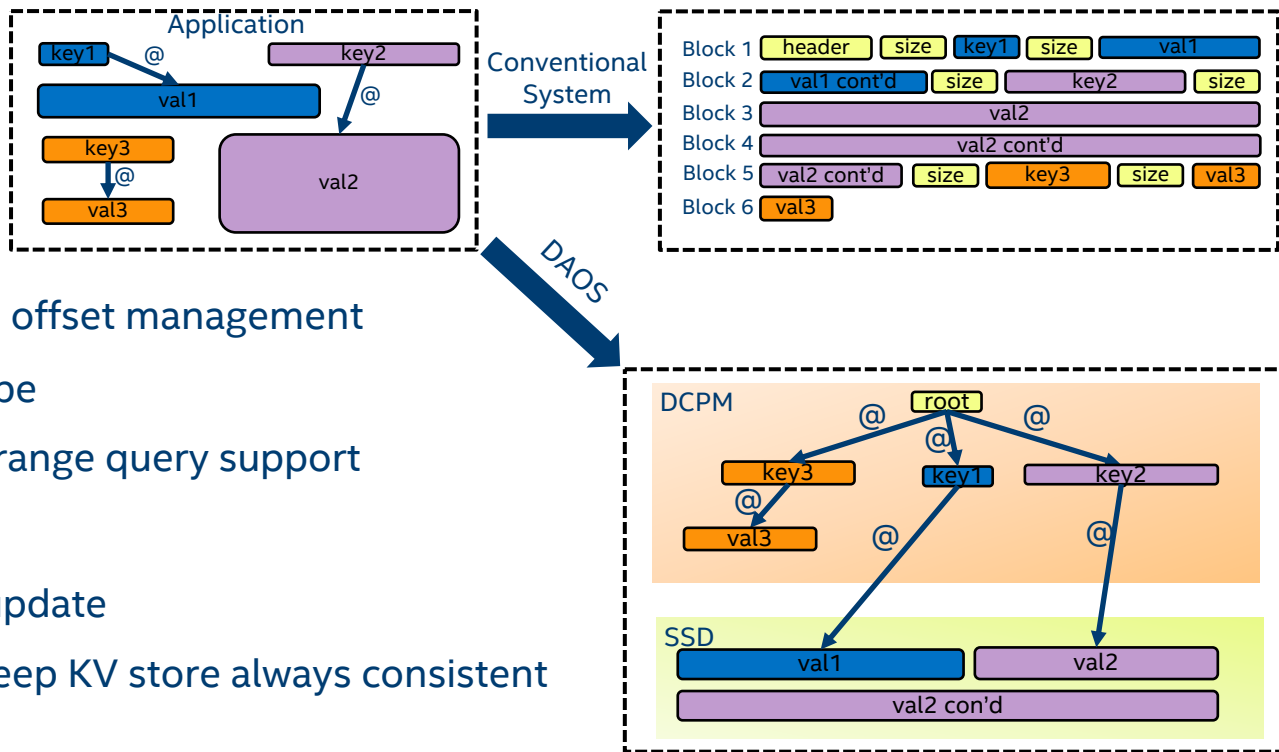
- Avoid file serialization and offset management
- Keys can be of any size/type
- Keys can be ordered with range query support

Scalable insert

- Allow concurrent access/update
- Distributed transactions keep KV store always consistent

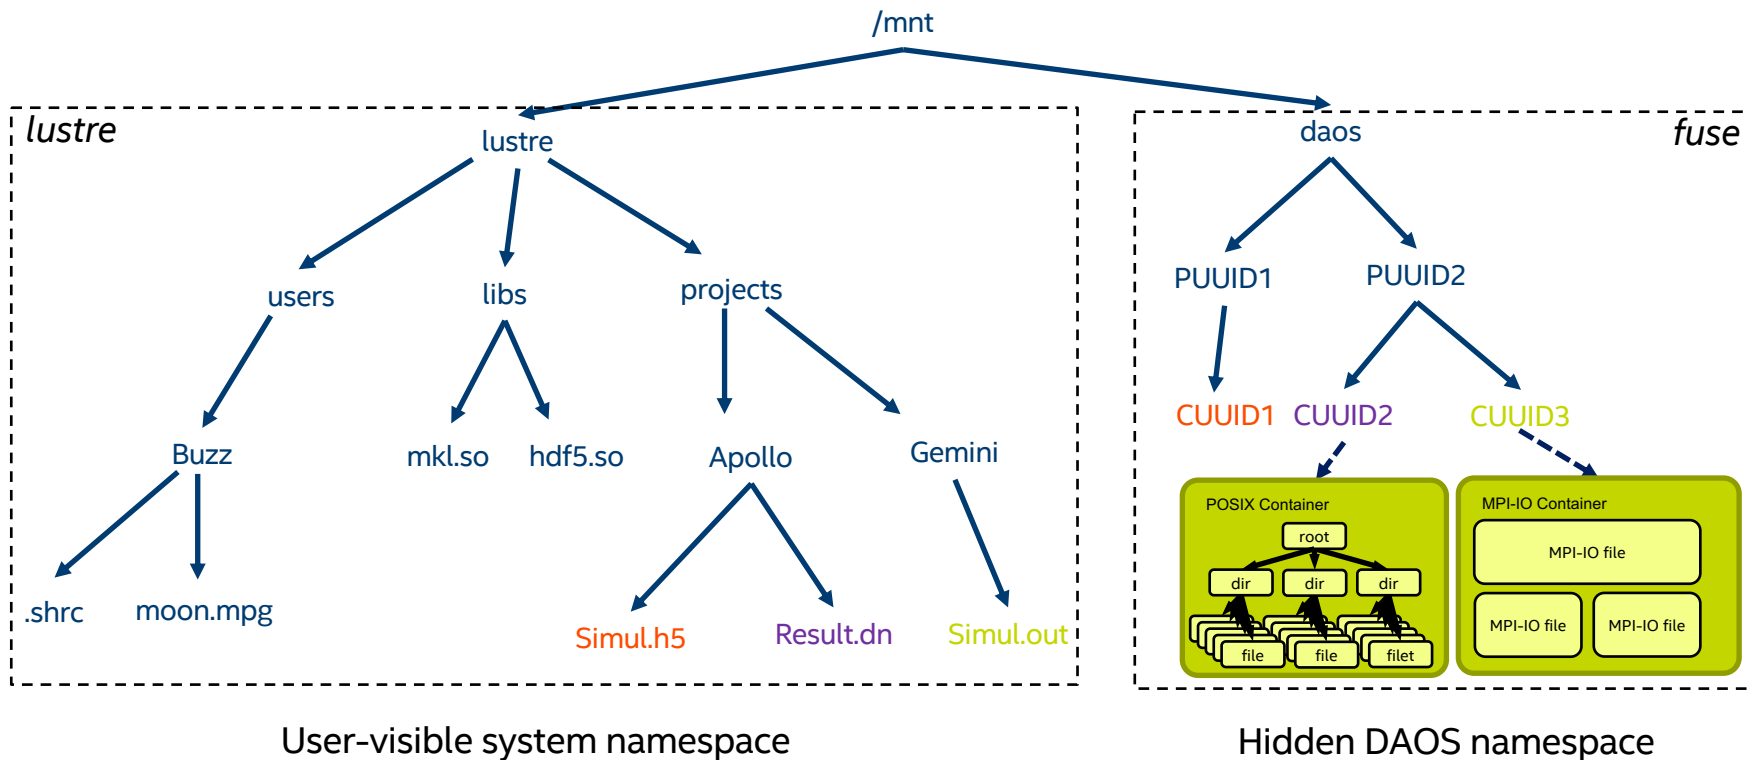
Data indexing

- Query & custom index
- Data provenance



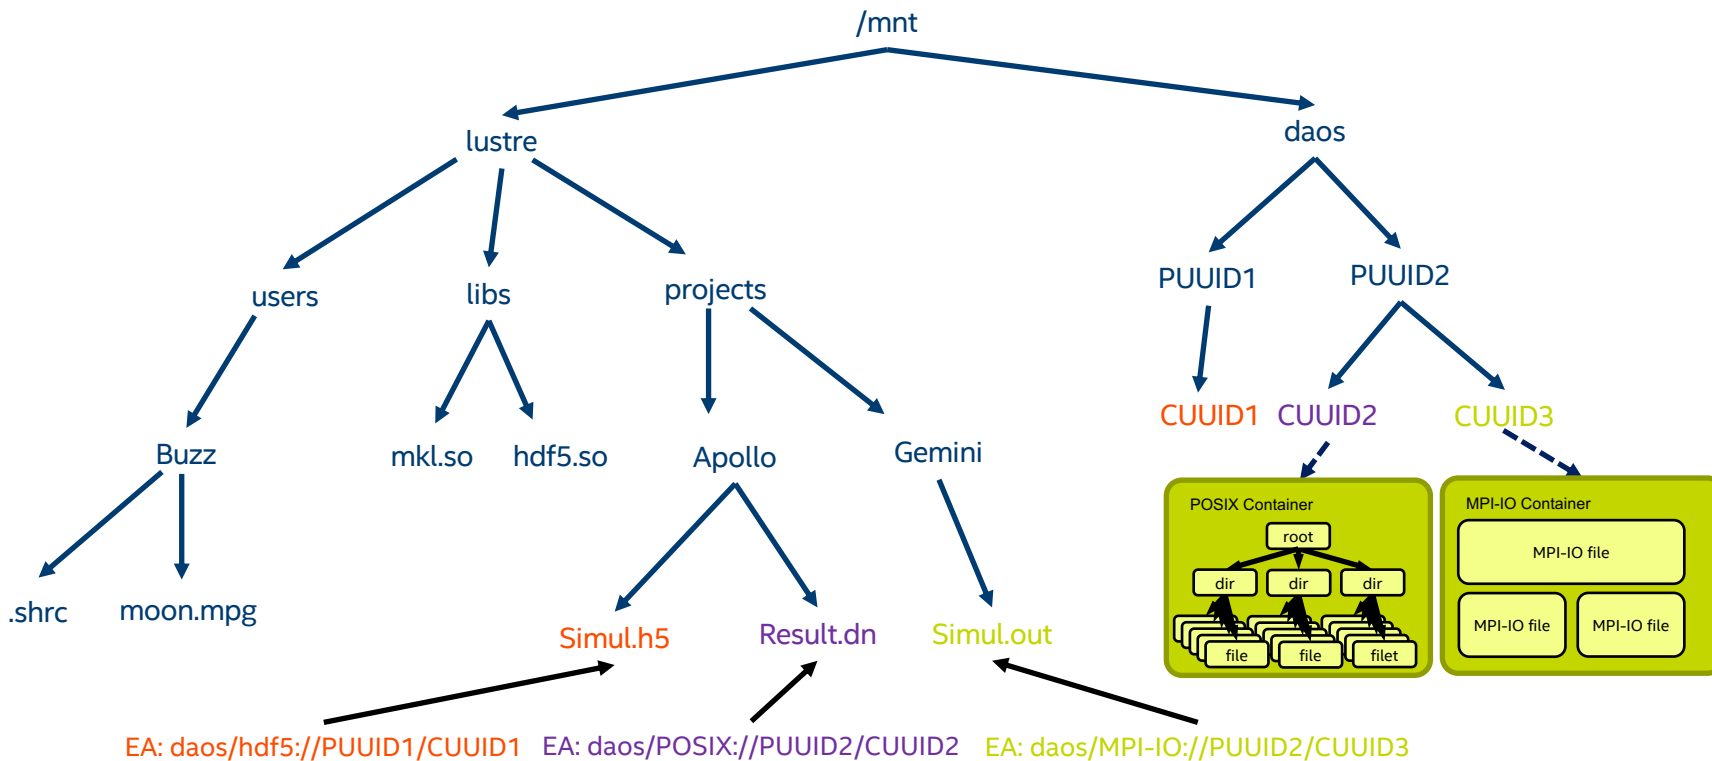
DAOS/LUSTRE INTEGRATION

Regular PFS directories & files
HDF5 Container
DAOS POSIX Container
DAOS MPI-IO Container



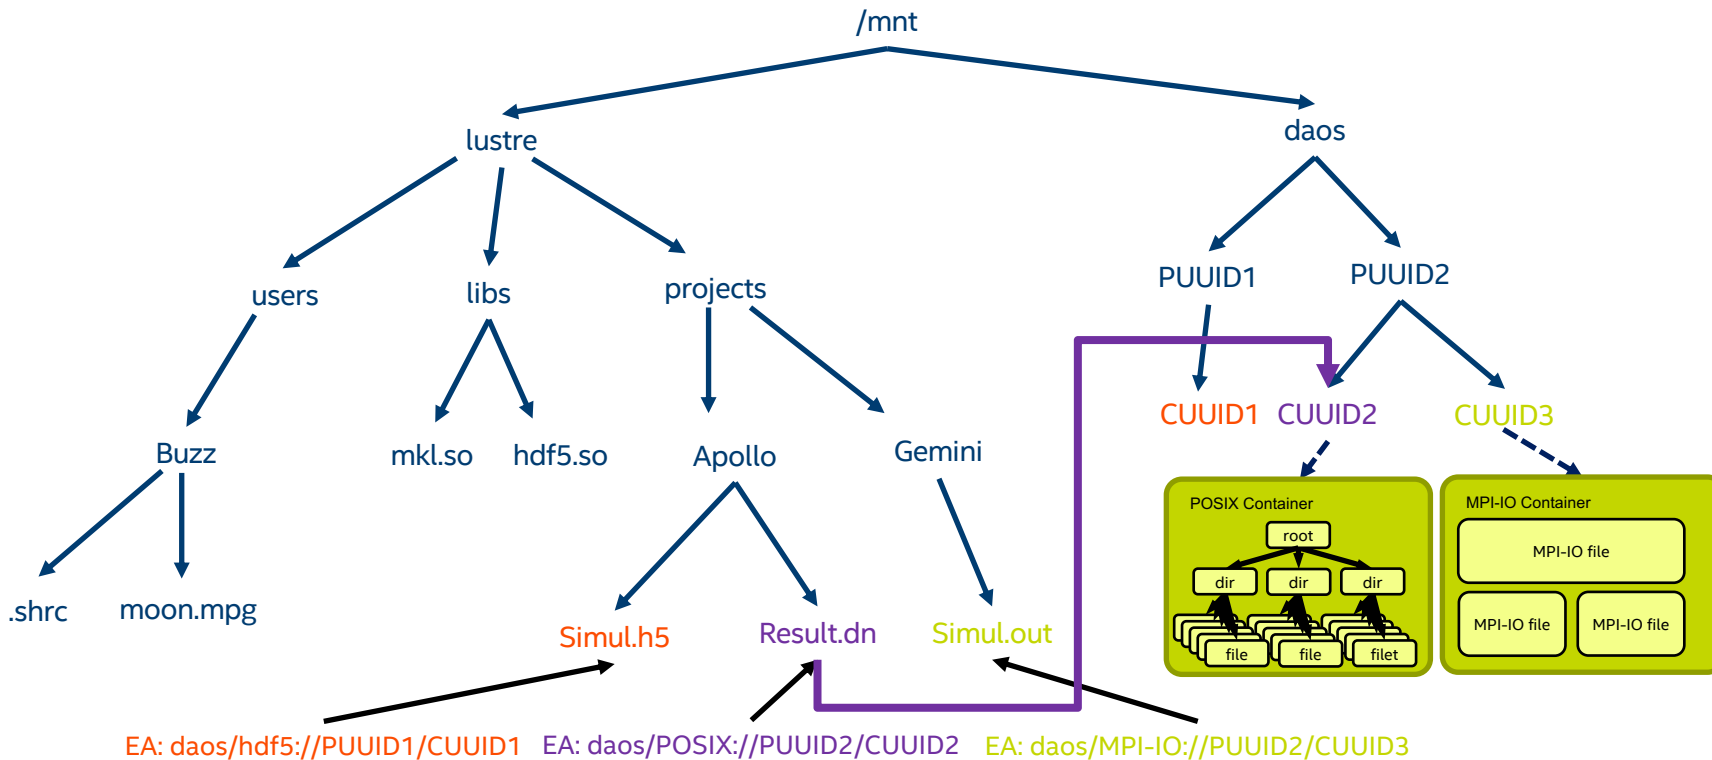
DAOS/LUSTRE INTEGRATION

Regular PFS directories & files
HDF5 Container
DAOS POSIX Container
DAOS MPI-IO Container



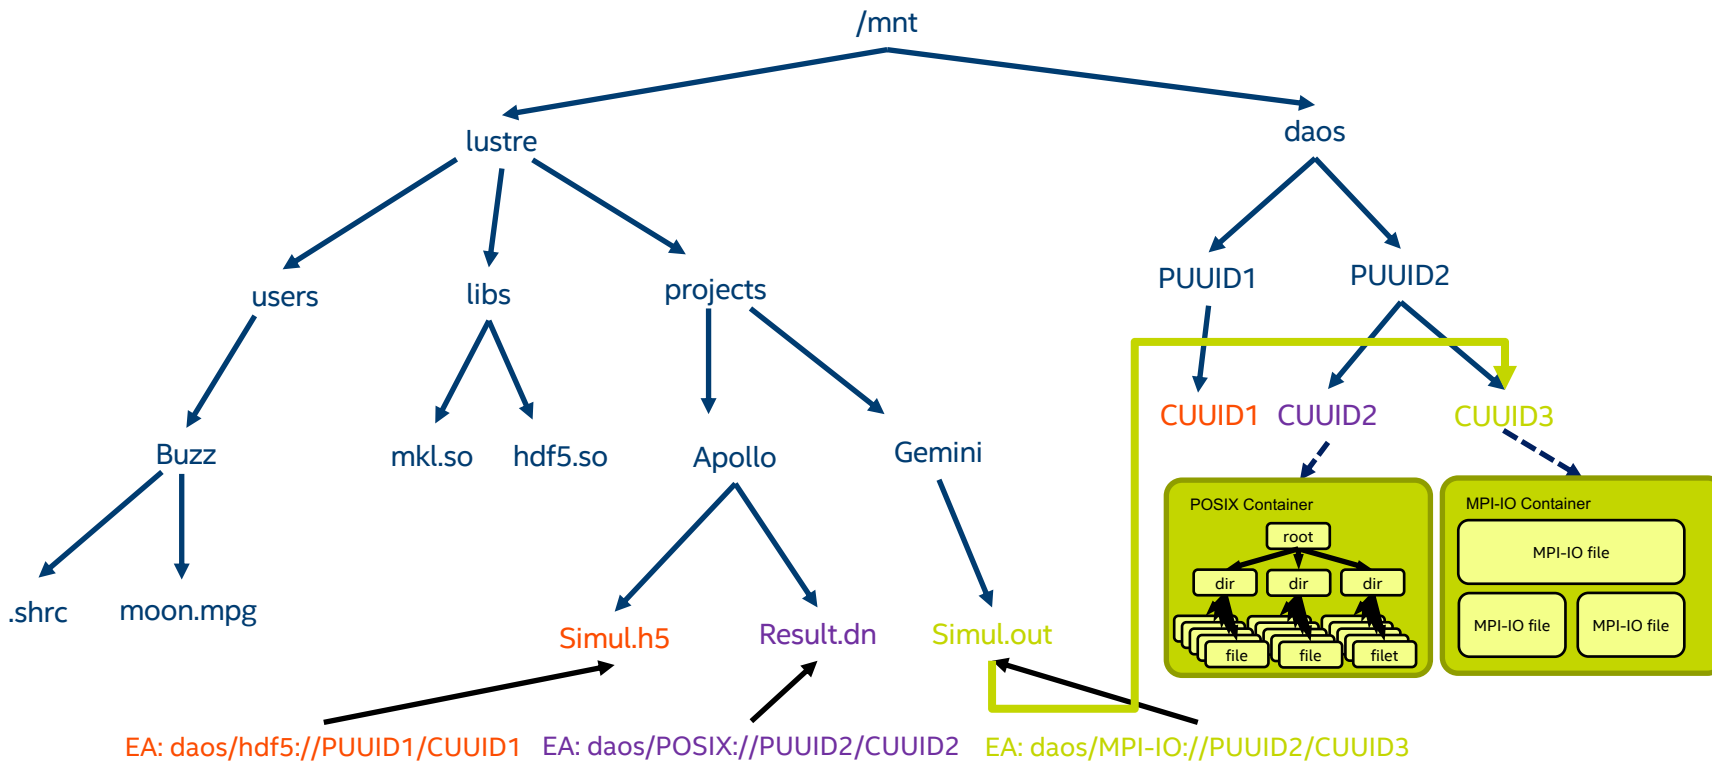
DAOS/LUSTRE INTEGRATION

Regular PFS directories & files
HDF5 Container
DAOS POSIX Container
DAOS MPI-IO Container

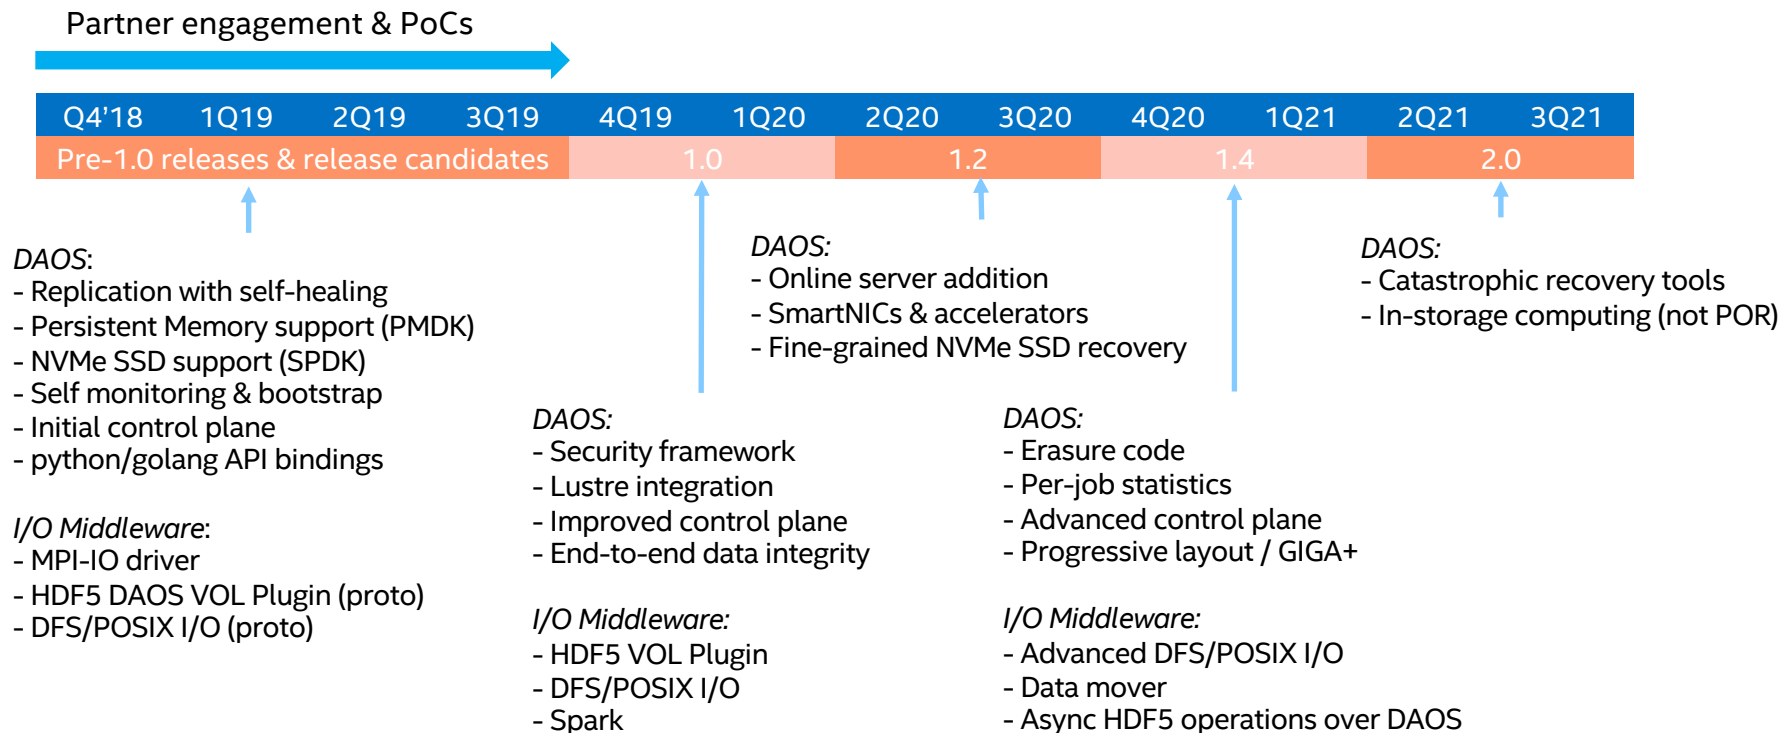


DAOS/LUSTRE INTEGRATION

Regular PFS directories & files
HDF5 Container
DAOS POSIX Container
DAOS MPI-IO Container



DAOS COMMUNITY ROADMAP



SUMMARY



Apache 2.0 License for **both** client & server

Source code **publicly** available on GitHub

- <https://github.com/daos-stack/daos>
- Contributions are welcomed and do not require a CLA

Community **mailing list** on Groups.io

- daos@daos.groups.io or <https://daos.groups.io/g/daos>

Wiki

- <http://daos.io> or <https://wiki.hpdd.intel.com>

Bug tracker

- <https://jira.hpdd.intel.com>

