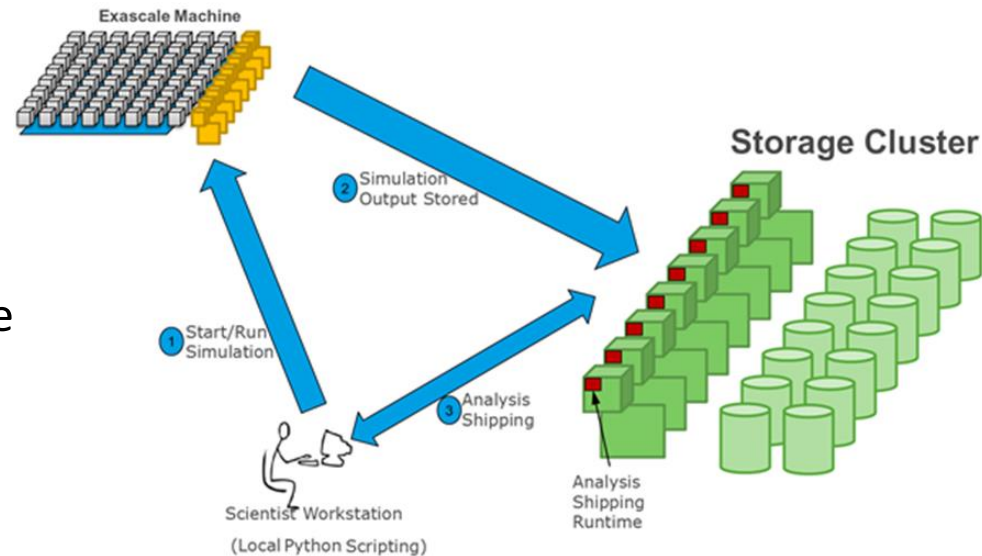


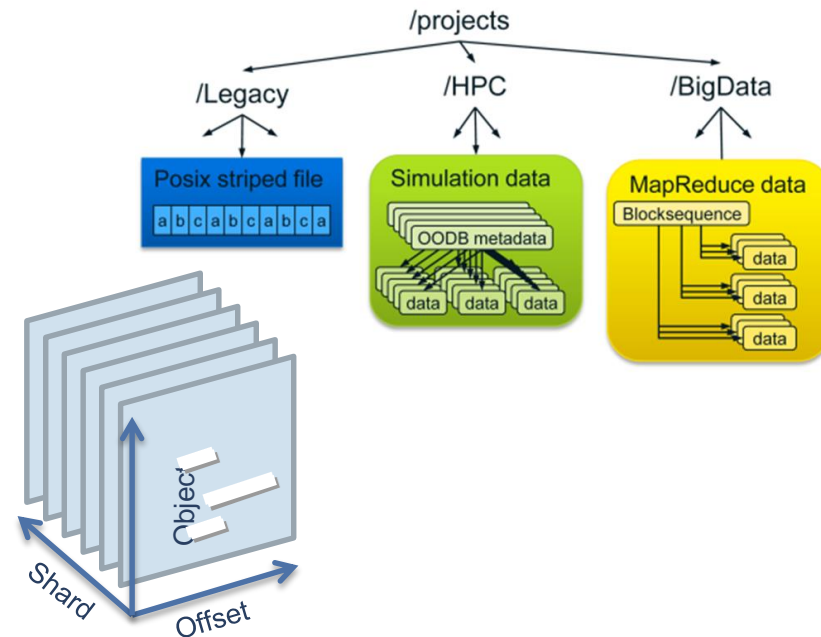
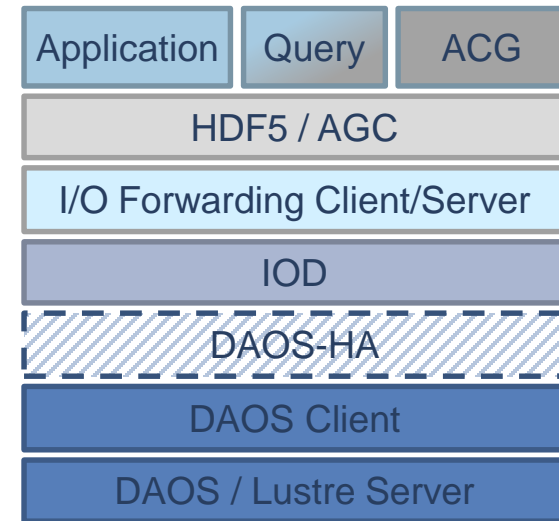
Fast Forward Storage & IO Project

- Intel, EMC, HDF Group, DDN, Cray
- Make Exascale storage a tool of the Scientist
 - Tractable data management
 - Comprehensive interaction
 - Move compute to data or data to compute as appropriate
- Overcome today's IO limits
 - Multi-petabyte datasets
 - Explosive growth of metadata
 - Horizontal scaling & jitter
- Support unprecedented fault tolerance
 - Deterministic interactions with failing hardware & software
 - Fast & scalable recovery
 - Enable multiple redundancy & integrity schemes



Stackable components

- Application I/O APIs
 - Multiple domain-specific API styles & schemas
 - HDF5+extensions & Graph Computation libraries
- I/O forwarding
 - Keeps top level APIs on Compute Nodes when IOD runs on the Burst Buffer
- I/O Dispatcher (IOD)
 - Impedance match application I/O to storage capabilities
 - Semantic resharding
 - Burst Buffer management
- DAOS-HA
 - High-availability scalable object storage
 - Follow-on project from Fast Forward
- DAOS Containers
 - Transactional persistent PGAS
 - Virtualized shared-nothing object storage
 - Unpolluted storage system namespace



Impact

- Performance
 - Commutative storage eliminates ordering requirements
 - Asynchronous APIs mitigate jitter
 - Storage hierarchy supports fine-granularity
- Model Integrity
 - Atomic updates
 - Consistency guarantees for application data / metadata
 - Right-size containers - no need for separate files for separate updates
 - Multiple redundancy schemas
 - Replicate most sensitive data/metadata more widely
 - End-to-end integrity checks / improved recoverability & snapshots
- Higher-level & richer functionality
 - Complex workflows
 - Analysis shipping / Burst Buffers / Distributed Persistent Memory
 - High-level object APIs
 - Object database index/query