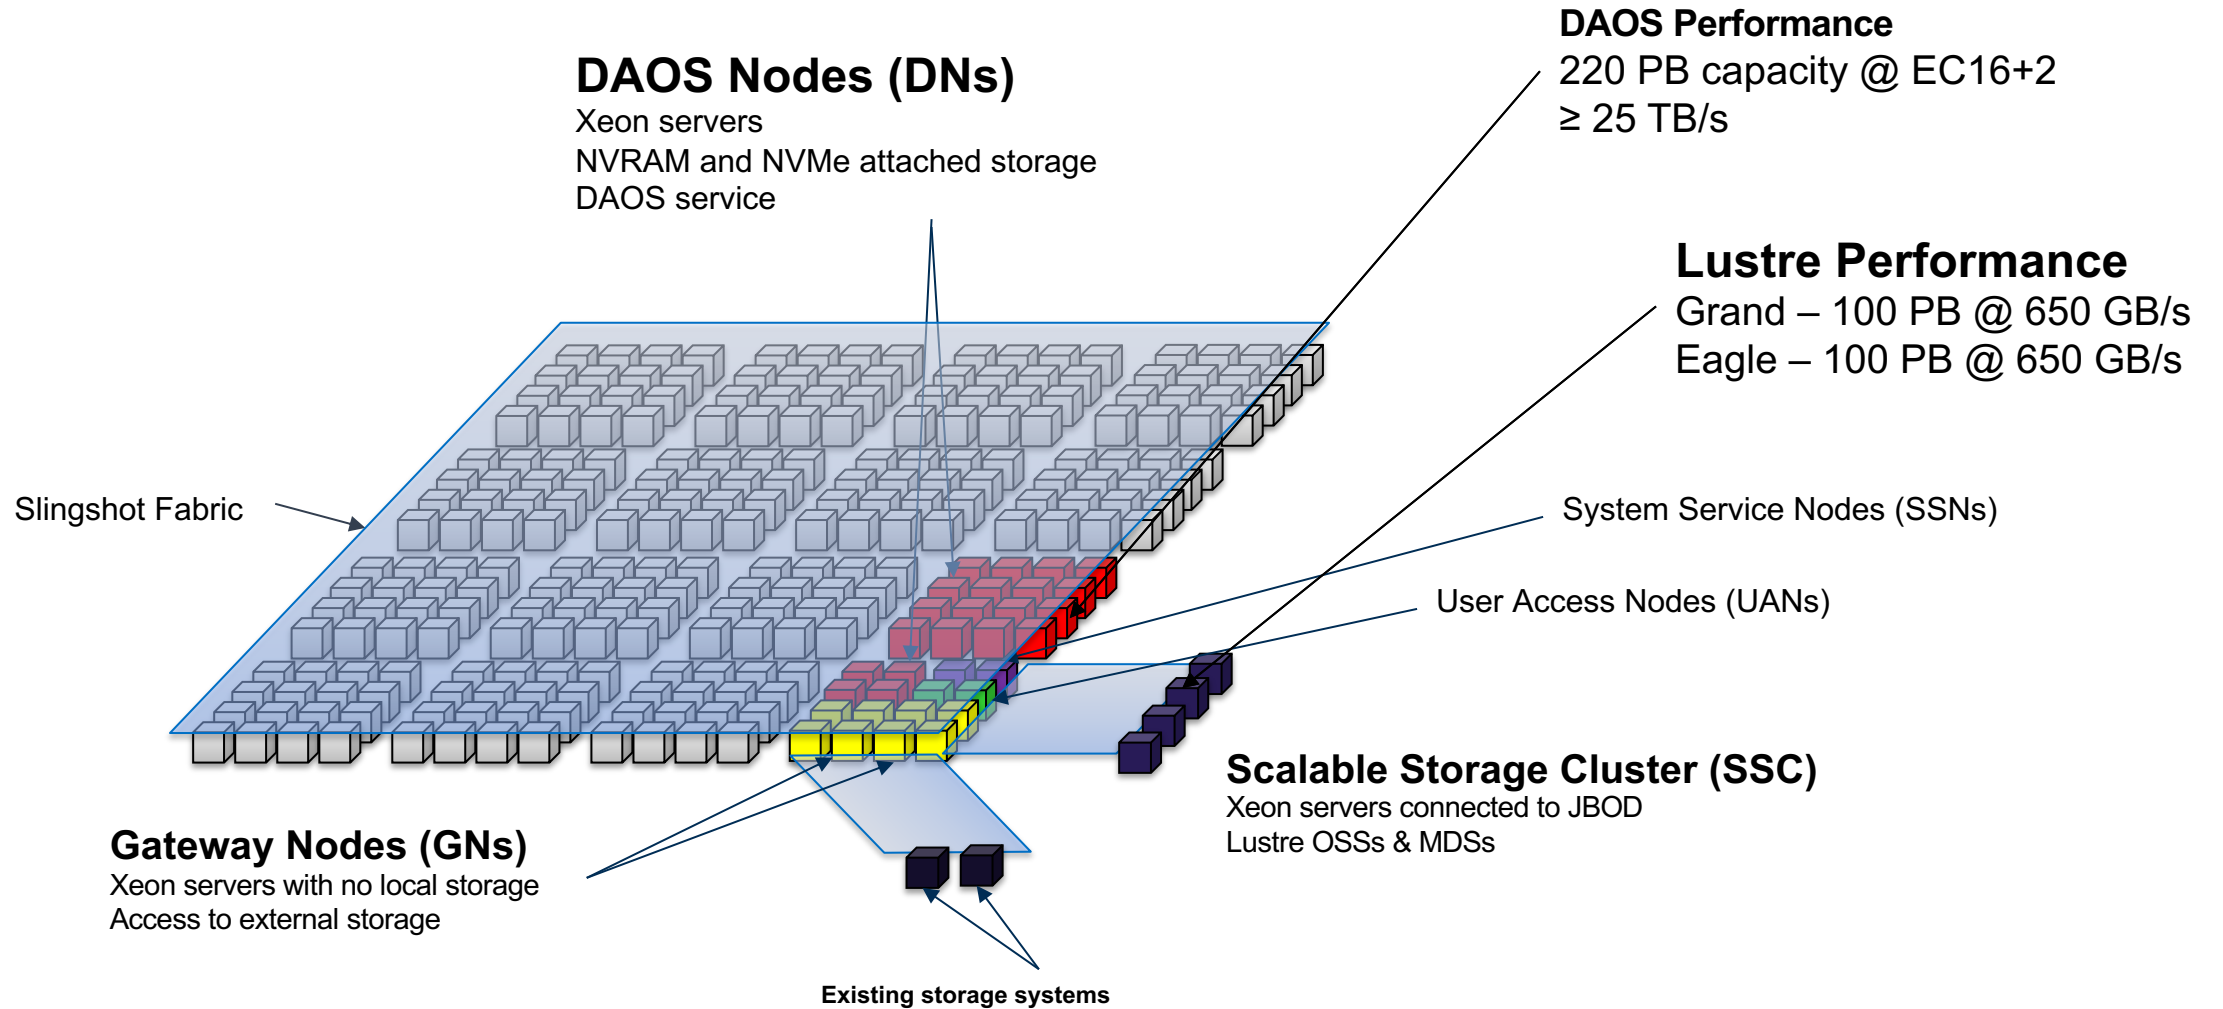




Kevin Harms
Performance Engineering Team Lead – Argonne Leadership Computing Facility



Aurora Overview



Goals

- Users can easily access storage with limited initial knowledge
 - Data access starts with Lustre
 - User can create DAOS POSIX containers with a single command
 - DAOS POSIX containers can be access seamlessly through Lustre path
- Users can adopt DAOS at their own pace
 - Our DAOS storage has 20x performance of Lustre storage
 - Intentional design to “force” users with high I/O demands to use DAOS
 - Can move workloads to be more DAOS centric to improve performance overtime
- Lustre (Grand/Eagle) still serve as our center-wide storage systems
 - Provided data exchange with other systems, both internal and external
 - Supported by a wide variety of new and legacy systems

User Setup

- Aurora Storage Allocations
 - Projects will receive a Lustre allocation on Grand or Eagle
 - Projects will receive a DAOS allocation
 - Size of allocations for both Lustre and DAOS are large enough to for data to be maintained for duration of the project
 - 6 months for Discretionary
 - 1 year for ALCC
 - 1-3 years for INCITE
 - Expectation that DAOS is used for high performance I/O
 - Lustre will be used for data staging, code and binaries, backup or cold storage

Functionality

- Login Node

- When users login they can 'cd' to their project space under Lustre to work on files/data which are shared with their project team
- Using 'daos' commands, users can create DAOS POSIX containers and they will be linked within the Lustre filesystem
 - When a user 'cd' into a DAOS POSIX container, using the Lustre foreign file feature to symlink to DAOS mount point
 - dfuse mounts the pool/container
 - User can seamlessly open files within the container

- Compute Node

- Same process as login node
- User gives the Lustre path of the file to open and then application will access DAOS seamlessly if file data is within DAOS container
 - Using either dfuse or the interception library

Acknowledgements

This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

Thank you!