

DUG 2019



DAOS STATUS - ALCF

GORDON MCPHEETERS
ALCF Operations
File System Admin

KEVIN HARMS
ALCF
Performance Engineer

KAL ALFIZAH
Intel
Onsite DAOS Support

TEST BED OVERVIEW

Current:

- 2 DAOS Nodes
 - SuperMicro SYS-2029UZ Dual Xeon Gold 6248 (10c/20t @ 2.5GHz)
 - Hyperthreading - Yes or No?
 - 12 * 128 GiB DCPMM / 384 GiB RAM
 - 12 * Intel SSD DC P4510 (2 TB) NVMe
- EDR IB
- DAOS configuration of various combinations have been exercised
 - single server primarily
 - use of DCPMM and without DCPMM (RAM only)
- 6 Haswell client nodes

FUTURE TEST BED OVERVIEW

- New 18 node test bed purchased from Cray
 - installed in a River Rack
 - same CPU/Motherboard as current test bed
 - 384 GiB Ram, 12 * 128 GiB DCPMM (Apache Pass) and 12 * 2 TB NVMe
 - 2 * 100 GbE NICs
- Primary test bed running into the A21 production time frame
- Enough nodes to support:
 - full erasure code testing
 - online server additional/removal
 - host several smaller multi-node file systems for test flexibility

RESULTS

- Initial standup of one I/O server running on each node, daos00 and daos01
 - (daos_server, daos_io_server, daos_agent)
 - OFI verbs
- Pool and container creation
- Test of POSIX container with dfuse and libioil
 - Pool/container on 1 server, no replication
 - ~450 MB/s with 'dd' of 1 MiB transfer size
 - ~1400 MB/s for the same but LD_PRELOAD of libioil