



Sandia
National
Laboratories

Exceptional service in the national interest

BEEGFS ON DEMAND EXPERIENCES, INFRASTRUCTURE, AND GOALS

Matthew L. Curry, mlcurry@sandia.gov

Michael Aguilar, mjaguil@sandia.gov

Shyamali Mukherjee, smukher@sandia.gov

25 May 2023

I/O in the Datacenter Workshop

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.





BEEGFS ON DEMAND

- BeeGFS On Demand (BeeOND) is a parallel file system that can be easily deployed on-the-fly with a selection of temporary resources
- Many potential use cases, not necessarily tied to HPC
- Use of non-shared resources to construct a storage pool (typically a shared resource) is extremely powerful
 - Network
 - Space
 - IOPs
 - Bandwidth



INITIAL PLATFORM: STRIA

- Stria is the development platform for Astra, the first petascale Arm platform
 - Quasi-testbed
- Tiered storage
 - 400TB NVMe Lustre
 - Off-machine lustre systems (14PB/19PB disk-based Lustre)
- 288TB SATA SSDs (one per node)
- Prototyping/evaluating multiple capabilities and flows
 - SSD-based deployment
 - Configurable geometries
 - Stage-in and stage-out facilities
 - RAM-based deployment





BEEOND SITE INTEGERATION

- Custom Slurm integration
 - Requires some site effort, but flexible
- mkfs (and TRIM)
 - No re-use between jobs
- Flexibility
 - Multiple metadata servers
 - Mapping to different portions of allocations
 - Media (SSD or RAM)
- Parameter passing for Slurm prolog/epilog is very limited
 - Binary parameters (--constraint=booleanArgument)
 - Integer parameters (--licenses=numericalArgument:value)
 - Impossible: enums, general strings, file paths



USE CASES

- Interesting use cases exist for any combination of staging pattern
 - Upside – Can't fail!
 - (Except... what needs a PFS but no staging?)
- Stage-in and stage-out is the primary barrier of use
 - MPIFileUtils work very well
 - Users don't want to do it
 - Requires MPI to operate
- Better BeeOND staging tool promised
- For some stage-in-only use cases, making BeeOND act as a read-through cache would be ideal (see: bcachefs)

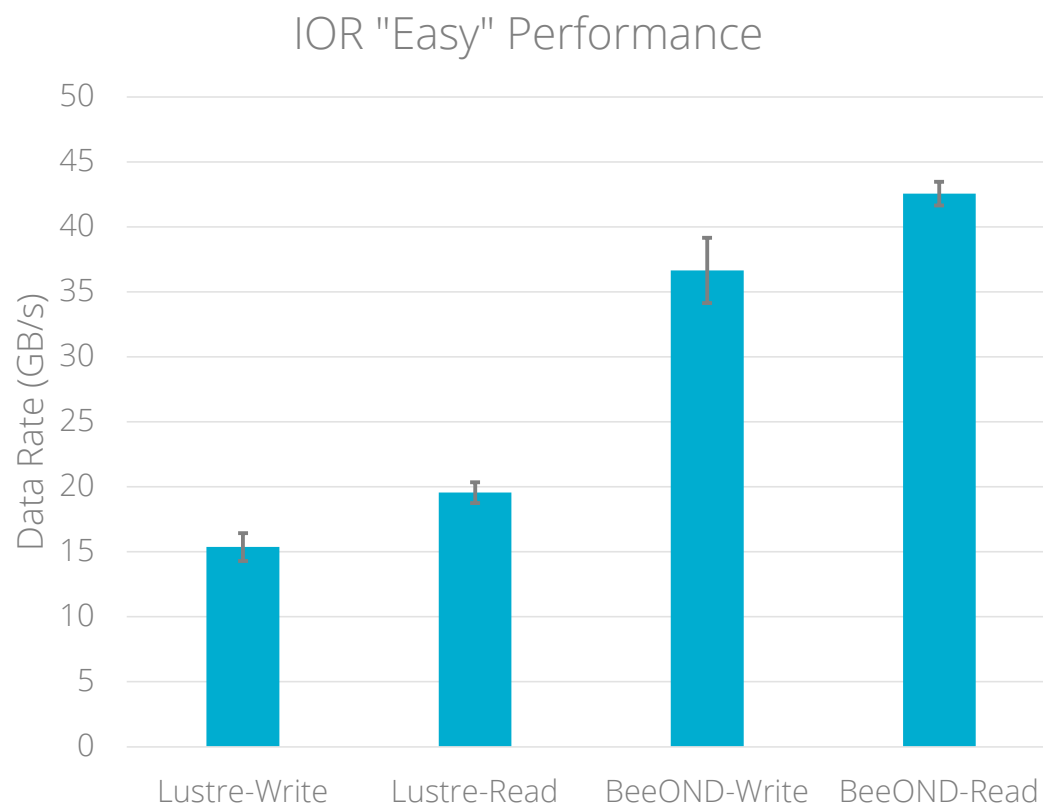
	No Stage In	Stage In
No Stage Out	Compilation CI/CD Out-of-Core	AI/ML Training Read Caching
Stage Out	Checkpointing Visualization Write-back Caching	Transactional Workloads

Categories of storage system interactions and example workloads



CHECKPOINTING/BURST BUFFER


- 100 nodes compared to NVMe-based Lustre
- Significant performance boost with BeeOND, but not holistic
 - Non-automatic stage-out
 - No resilience available on this tier
- Read variance is much better





STRIA BEEOND PERFORMANCE – AI/ML

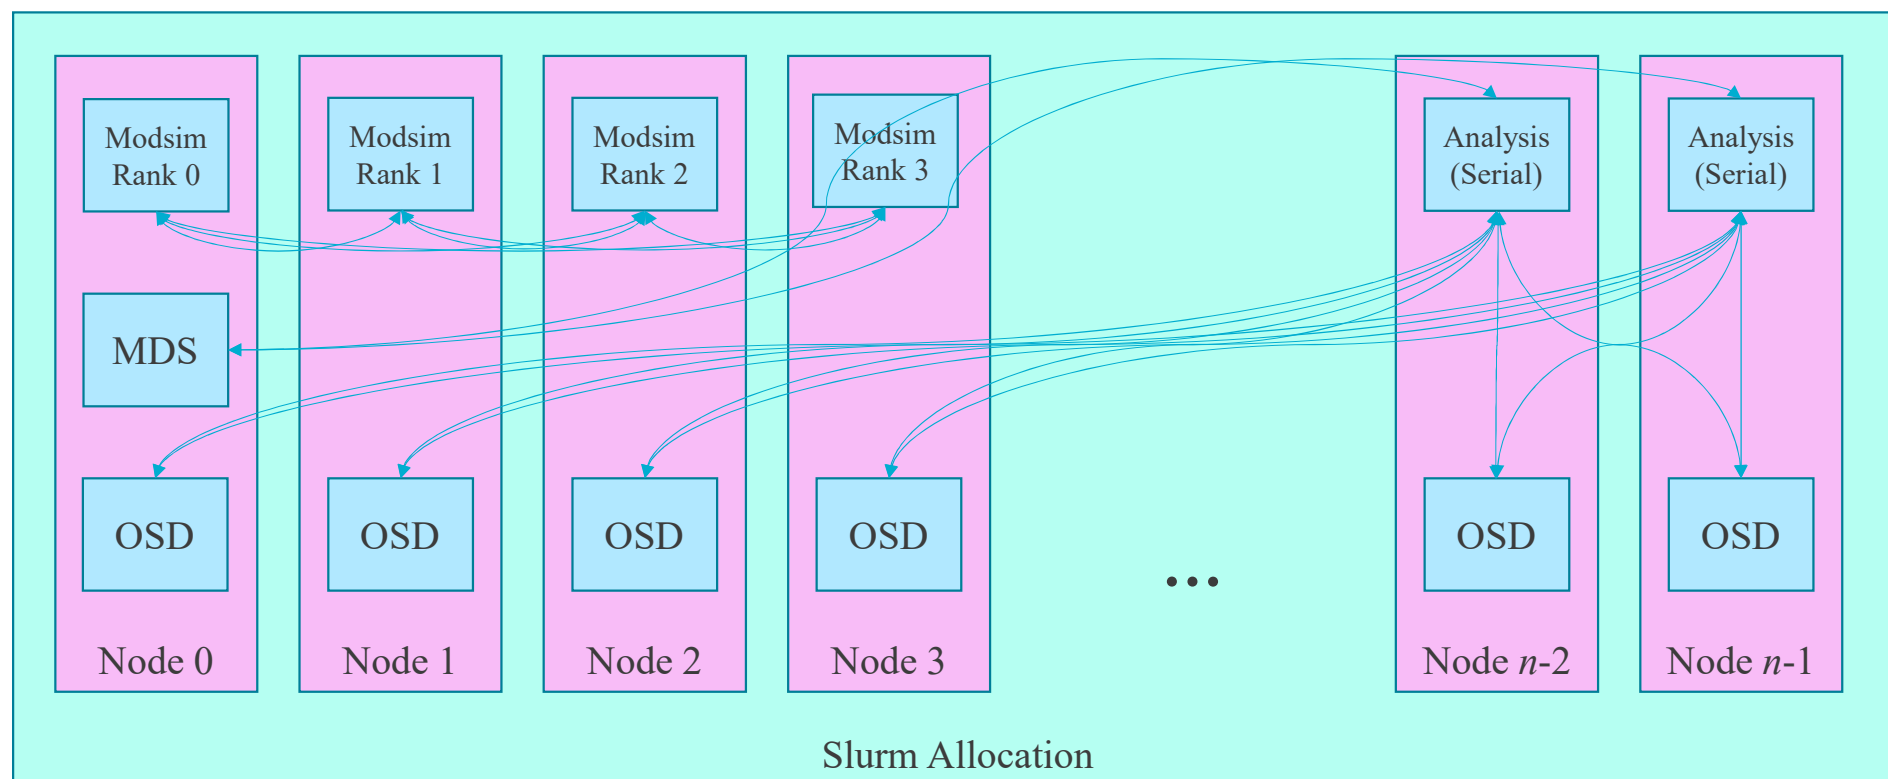
- ImageNet classification benchmark, I/O-only
- Three configurations:
 - RAM-based
 - SSD-based
 - Lustre
- BeeOND performance is *comparable*, and that's good enough
- Note significant performance variation in low ranks – Possible metadata server contention





BEEOND-INDUCED PERFORMANCE DEGRADATION

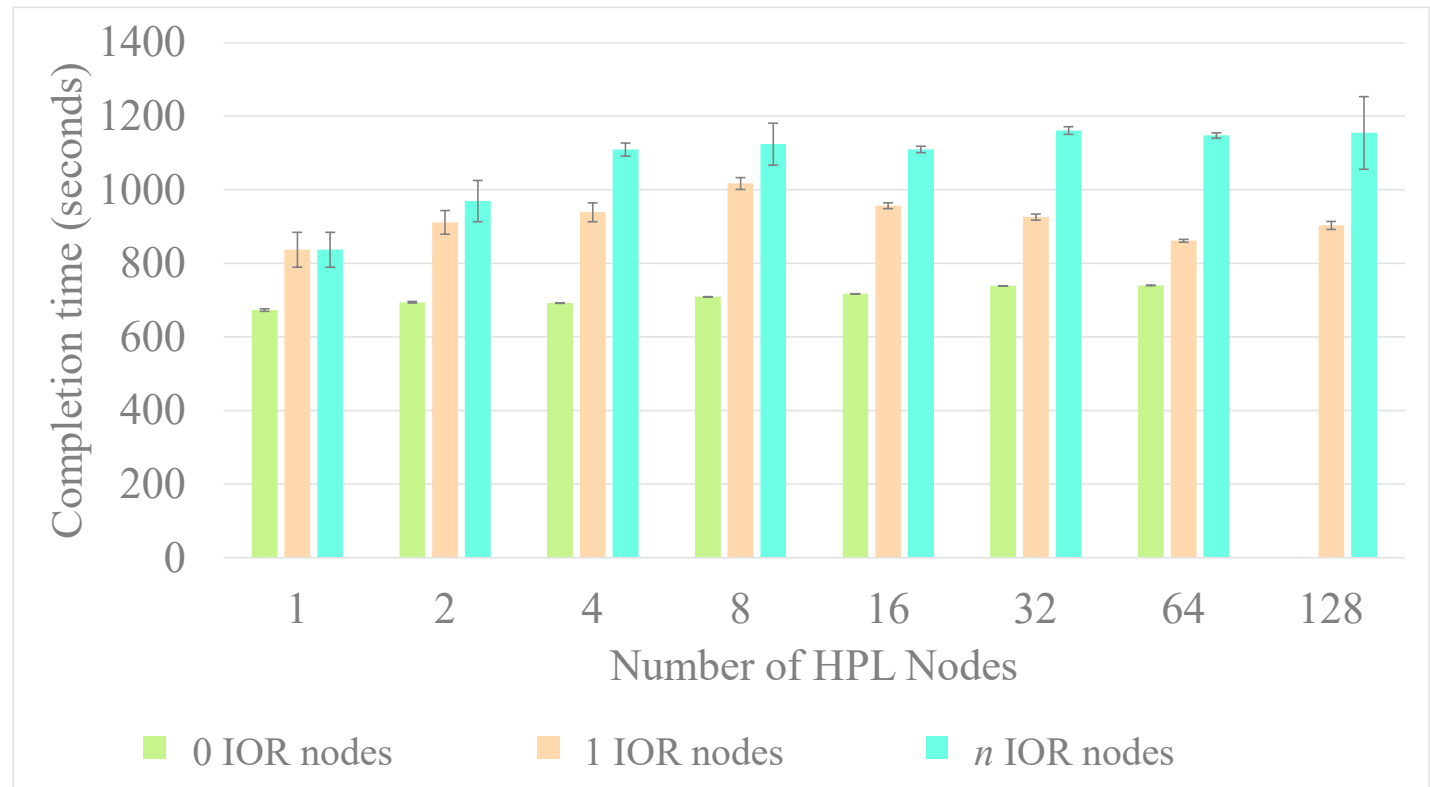
Can I/O to BeeOND processes cause spooky performance degradation at a distance?





PERFORMANCE OF CO-LOCATED COMPUTE AND BEEOND


- Clearly, the answer is yes
- Even small numbers of processes can cause perturbation
- If tightly coupled progress is desired, core- or host-specialization may be needed





10 DEPLOYMENT CHALLENGES


- SSD allocation is static per node
 - Not “right-sized” for the job
 - Reliant on system designer to get the 90% answer
- Jobs that need *more* SSD capacity have to allocate extra compute nodes
- No per-node resilience included
- Recovering from job failure is tricky, as placement is important
- Data service geometry is bound to job geometry





11 DEPLOYMENT CHALLENGES

- Disaggregated/composable infrastructure may be the answer
- NVMeoF can offer some of this flexibility today, at the expense of noise on the network
- CXL may offer high-speed storage devices over a separate network type, easing congestion on main HPC fabric





CONCLUSIONS

- BeeOND is seeing wide deployment at Sandia in the coming months
 - Even to systems that do not include on-node SSDs
- Applicable to a large number of workloads
 - But users resist manual data movement/staging
- Promises to decrease storage system variance and increase performance
 - But if used without care, can *increase* variance and *decrease* performance
- Current techniques (like user configurability) are key to realizing benefits
 - But configurability can be challenging
- This (and other solutions like it) are key to taking full advantage of commodity clusters