

Gestalt Computing: Hybrid Traditional HPC and Cloud Hardware and Software Support

Jay Lofstead, Andrew Younge

Sandia National Laboratories

Exceptional service in the national interest

U.S. DEPARTMENT OF ENERGY

National Nuclear Security Administration

SAND2022-5177 C

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.

Supercomputer or Cloud?

Sandia National Laboratories

- Traditional HPC workloads
  - Very large single task per job running for a long period of time.
  - Write to storage dominates IO behavior
  - Regular peer-to-peer communication requires fast, low latency interconnect
- Traditional Cloud (and general HPC) workloads
  - Lots of independent tasks, many short running
  - Frequently read intensive for data analysis or machine learning-style tasks
  - Interconnect doesn't really matter
- We still deploy different platforms for each workload, but that is expensive!

On Prem or remote?

Sandia National Laboratories

- IF the machine can be kept busy most, if not all, of the time, on prem, self managed can be cheaper
- IF applications require specialized hardware, on prem may be the only option

- Few places are exclusively one or the other
  - Supercomputer workloads have intense analysis phases that match Cloud pretty well
  - Cloud (general HPC) applications may write a “large” output at the end (e.g., the generated ML model) or use checkpoints to guard against failures.

Cloud Bursting and Sky Computing

Sandia National Laboratories

- Cloud bursting
  - Offload workloads from specialized platforms to general use platforms to ensure availability.
  - Reduce size of on prem resources to most common needs controlling costs
- Sky computing
  - Multi-cloud and hybrid cloud setup
  - Choose where to run based on spot prices or workload requirements
  - Infrastructure to migrate data and apps to alternative platforms
    - Reminiscent of the grid work from the 1990s-2000s

Some workloads are sensitive

Sandia National Laboratories

- National security workloads require tight data controls
  - Corporate/trade secrets are similar
- Budgets are not infinite
- New cloud tools quite useful for modeling and simulation data analysis demanding both infrastructures
- Need a way to offer a hybrid machine that adapts to the workload
  - Underlying hardware must adapt with a software layer to manage dynamically

Persistent Memory (SCM)

Sandia National Laboratories

- Devices on memory bus that can persist rather than just store data
  - Intel Optane 3-D Crosspoint one example
- Offered as higher capacity for DRAM at a lower cost, lower performance point.

- RAM disks have existed for eons, why not adapt PMEM to be a RAM disk for workloads that need storage and DRAM for those that need on node memory capacity?

Competitors

Sandia National Laboratories

- SATA-based devices
  - Too slow an interface to be interesting as Infiniband can be a faster way to store data remotely
- NVMe devices, expensive storage, but fast for storage
  - Cheaper for capacity than DRAM or PMEM, faster than SATA devices
- CXL for remote storage access
  - Micron/PNNL exploring PMEM + CXL models

Challenges

Sandia National Laboratories

- FAST devices hampered by slow software stacks
  - NVMe demands carefully tuned IO routines
  - PMEM demands proper API calls to avoid IO stack overheads

- In spite of this, assume the following:
  - Cost is acceptable
  - Infrastructure exists to change machine view
  - Software infrastructure to use hardware effectively incorporated

How would we use PMEM to deploy a single machine for different workloads?

Use Case 1: ModSim + Analytics

Sandia National Laboratories

- LAMMPS molecular dynamics + analysis
  - Consider a melt or fracture simulation (moderate physics)

- LAMMPS spreads data across several nodes
- Analysis runs separately at the same time
- Feedback loop to control how LAMMPS runs when phenomena first appear (e.g., backtrack slightly and then run with finer time resolution)

- LAMMPS nodes: PMEM for extra memory
- Analysis: PMEM for storage

Use Case 2: Cloud Analytics

Sandia National Laboratories

- Iterative learning algorithms
  - Read the same dataset MANY times, randomly

- PMEM devices storage
- Data pre-staged into or cached as read in PMEM devices
- Main memory used for model creation

Use Case 3: Traditional ModSim

Sandia National Laboratories

- LAMMPS by itself
  - Just need to push data for analysis and checkpoints periodically
- PMEM as storage (burst buffer)
- Asynchronously drain data to off-node storage during compute intensive application phases

Software Infrastructure

Sandia National Laboratories

- SLURM vs. Kubernetes, Spark, Mesos, Yarn, etc.?
- What about on-node operating system?
- Hybrid is possible
  - CIQ Fuzzball is nascent
  - Flux from LLNL is nascent (more hybrid on-prem/cloud)

Discussion

Sandia National Laboratories

- Different applications have different needs
- Hardware can adapt offering a gestalt—with the right software

- PMEM offers adaptable solution that costs slightly more for a single machine, but enables efficient running of a wider variety of workloads on a single platform

Thank you

Sandia National Laboratories

- Questions gflofst@sandia.gov or online