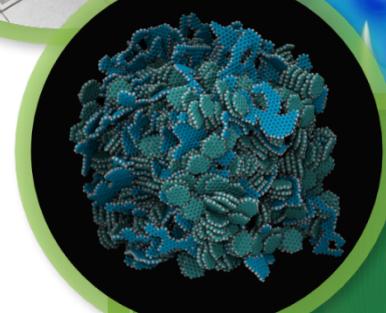


ARM At ORNL: A One Year Review



ORNL is managed by UT-Battelle
for the US Department of Energy

Introduction

- Have been officially looking at ARM architecture for about a year
 - Probably some sporadic individual efforts before that
- Have one cluster from Cray, buying another from HPE
- Technology evaluation
 - No plans to deploy ARM supercomputers in production, but we want to know how feasible something like that would be
- Not just the hardware, want to understand the current state of the software ecosystem
 - Current state is actually pretty good
 - We'll try to push it forward in some areas where we have some experience
- Users from not just ORNL, but other labs and universities

Cray Envoy System

- Operational since Nov '16
- Hardware:
 - 16 Compute nodes
 - Cavium ThunderX CPU's
 - 96 cores per node (2 sockets) and 128 GB RAM per node
 - 10G Ethernet between nodes
 - NFS-mounted disk array in the rack for storage
- Not really intended for performance tests
 - Users still built and ran Linpack and HPCG
- Cray PE is still NDA, so I can't talk about performance numbers

HPE Collaboration

- HPE started a collaboration between itself, several national labs and various vendors to push development of ARM servers for HPC
- ORNL joined the collaboration in February/March 2017
- We're deploying a test cluster using this new hardware
 - 16 nodes (15 compute, 1 login)
 - 2 socket ThunderX2, 256GB ram, 500GB SSD, EDR IB
 - 4 nodes will have AMD GPU's (2 per node)
- First two nodes arrived last week, remainder delivered by mid-December

What Are We Doing With This Hardware

- Lustre client
- Researchers' favorite apps
 - <https://gitlab.com/arm-hpc/packages/wikis/home>
- OpenSHMEM (though since our hardware doesn't support RDMA, it's not very efficient)
- “Regular” Linux software
 - Might want to use ARM architecture on something smaller than a leadership-class supercomputer
 - Need standard post-processing, analysis & vis software
 - Fortunately, most of that seems to work
 - Python (including NumPy & Matplotlib), ParaView, NetCDF, HDF5

Lustre Client

- ORNL has committed to supporting a Lustre client for ARM
- James Simmons is working on it
- Client is currently partially functional
 - Can mount a filesystem and do metadata operations
 - Getting error messages from the MDS when we try to write
- Performance testing will be done with the HPE system
- Want the ARM client to be a first class citizen
 - Want it integrated into the Lustre C.I. process
 - Not a “one-off” build
 - Would like to see it integrated into Linux distros (but even `x86_64` isn’t, so AArch64 is going to take some work...)

Summary

- Overall, I'm optimistic
- Experience building & running software on the Cray has been positive
 - Spack made things a lot easier once we got it set up
- We've got fancy new hardware showing up soon
 - Hopefully, I can come back next year and talk about all the work I've done with the GPU's
- Software is available and working; I'd like to see more official support
 - Bit of a bootstrap problem: software vendors won't support ARM unless their customers ask for it; customers won't buy ARM if they can't run their software
 - This is one of the points we'll be discussing at the ARM BoF session

Questions?

ACKNOWLEDGMENTS

This work was supported by, and used the resources of, the Oak Ridge Leadership Computing Facility, located in the National Center for Computational Sciences at ORNL, which is managed by UT Battelle, LLC for the U.S. DOE, under the contract No. DE-AC05- 00OR22725.