OAK RIDGE National Laboratory

Summit Training Workshop – Knoxville, TN

Tom Papatheodore

Oak Ridge Leadership Computing Facility

December 3-6, 2018

U.S. DEPARTMENT OF ENERGY

Summit Training Workshop Agenda

	Monday, December 3, 2018	Tuesday, December 4, 2018	Wednesday, December 5, 2018	Thursday, December 6, 2018
8:00	Registration	Registration	Registration	Registration
8:00	Working Breakfast	Working Breakfast	Working Breakfast	Working Breakfast
8:15	Welcome Address + Notes	Welcome Address + Notes	Welcome Address + Notes	Welcome Address + Notes
8:30	Summit Architecture (overview) *Scott Atchley*	CUDA Unified Memory *Jeff Larkin*	GPFS *George Markomanolis*	Debugging (Arm DDT) *Nick Forrington*
9:00	Summit Architecture (IBM P9, NVIDIA V100) *Brian Thompto & Jeff Larkin*	GPU Direct, RDMA, CUDA-Aware MPI *Steve Abbott*	SSDs / Burst Buffers *George Markomanolis & Chris Zimmer*	Arm MAP / Performance Reports *Nick Forrington*
9:30	Programming Environment *Matt Belhorn*	Targeting Summit's Multi-GPU Nodes *Steve Abbott*	Network Features and MPI Tuning *Chris Zimmer*	Mixing OpenMP + OpenACC *Lixiang (Eric) Luo*
10:00	Programming Environment (cont.) *Matt Belhorn*	Targeting Summit's Multi-GPU Nodes (cont.) *Steve Abbott*	IBM P9 SMT Deep Dive *Brian Thompto*	Python Environments *Matt Belhorn*
10:30	Break	Break	Break	Break
10:45	Storage Areas / Data Transfers *George Markomanolis*	GPU Accelerated Libraries *Jeff Larkin*	Experiences using the Volta Tensor Cores on Summit *Wayne Joubert*	CAAR Porting Experience: QMCPACK *Andreas Tillack*
11:15	LSF Batch Scheduler & jsrun Job Launcher *Chris Fuson*	NVIDIA Profilers *Jeff Larkin*	Experiences Porting/Optimizing Codes for Acceptance Testing *Bob Walkup*	E3SM Application Readiness Experiences on Summit *Matt Norman*
11:45	LSF Batch Scheduler & jsrun Job Launcher (cont.) *Chris Fuson*	Using V100 Tensor Cores *Jeff Larkin*	Targeting GPUs Using OpenMP Directives on Summit w/GenASiS: A Simple and Effective Fortran Experience *Reuben Budiardja*	Experiences in Porting XGC to Summit *Ed D'Azevedo*
12:15	Working Lunch	Working Lunch	Working Lunch	Working Lunch
12:30	Working Lunch - Practical Tips for Running on SUMMIT *David Appelhans*	Working Lunch - Directive-Based GPU Programming *Oscar Hernandez*	Working Lunch - Node Performance *Wayne Joubert*	Working Lunch - ML/DL Frameworks on Summit *Junqi Yin*
13:15	Hands-On with OLCF Staff and Vendors	Hands-On with OLCF Staff and Vendors	Hands-On with OLCF Staff and Vendors	Hands-On with OLCF Staff and Vendors
17:00	Adjourn	Adjourn	Adjourn	Adjourn

Ascent

- 18-Node cabinet with the same architecture/nodes as Summit (only 16 nodes available for scheduling).
- **NOTE:** All in-person participants will be sharing access to Ascent during the workshop, so the following resource policies will be in place:
  - 1 Node per job, 20 minute walltime limit, 1 job running & 1 job eligible
  - As the week progresses, if you need more nodes, we will handle requests individually

OLCF Summit User Guide

- OLCF Documentation on using Summit
- **NOTE:** Ascent mounts different file systems than Summit, so please refer to info in these slides instead.

Available File Systems / Storage Areas on Ascent

Upon logging in to Ascent, you will be placed in your personal home (NFS) directory, **`/ccsopen/home/userid`**. The **`/ccsopen/proj/gen113`** directory (also NFS) can be accessed by all members of the GEN113 project (so all in-person workshop participants); If you need a collaborative workspace, I would suggest that you create a directory here (with some appropriate name). Both of these directories are within NFS, *and are places you might want to keep source code and build your application.*

You also have access to a (GPFS) parallel file system, called wolf. *This is where you should write data when running on Ascent's compute nodes.* Under **`/gpfs/wolf/gen113`**, there are 3 directories:

```
[ascent ~]$ ls /gpfs/wolf/gen113/
proj-shared  scratch  world-shared
```

- **`proj-shared`** can be accessed by all members of GEN113, so this is where you should create a directory (with some appropriate name) to collaborate with your team. So you would possibly have 2 shared directories for collaboration – 1 on NFS (for source code, scripts, compiling, etc.) and 1 on GPFS (for writing data from compute nodes that needs to be shared among your team members).

- **`scratch`** contains directories for each user, and only that user can access his/her own scratch directory. So you could write here from the compute nodes if you don't need to share the resulting files with the rest of your team.

- **`world-shared`** can accessed by any user on the system in any project (e.g. STF007 or something). This is meant for collaborating across groups, but you should likely not need this ability during the workshop since we all have access to GEN113 anyway.

Additional Notes

- Please make sure to wait for a mic when asking questions
  - If not, (presenters) please repeat the question
- If you are attending in-person, and you do not know how to obtain access to Ascent, please
  - Find attachment to email I sent titled "Summit Training Workshop – Ascent Access & Logistics"
  - Or, email Tom Papatheodore at papatheodore@ornl.gov

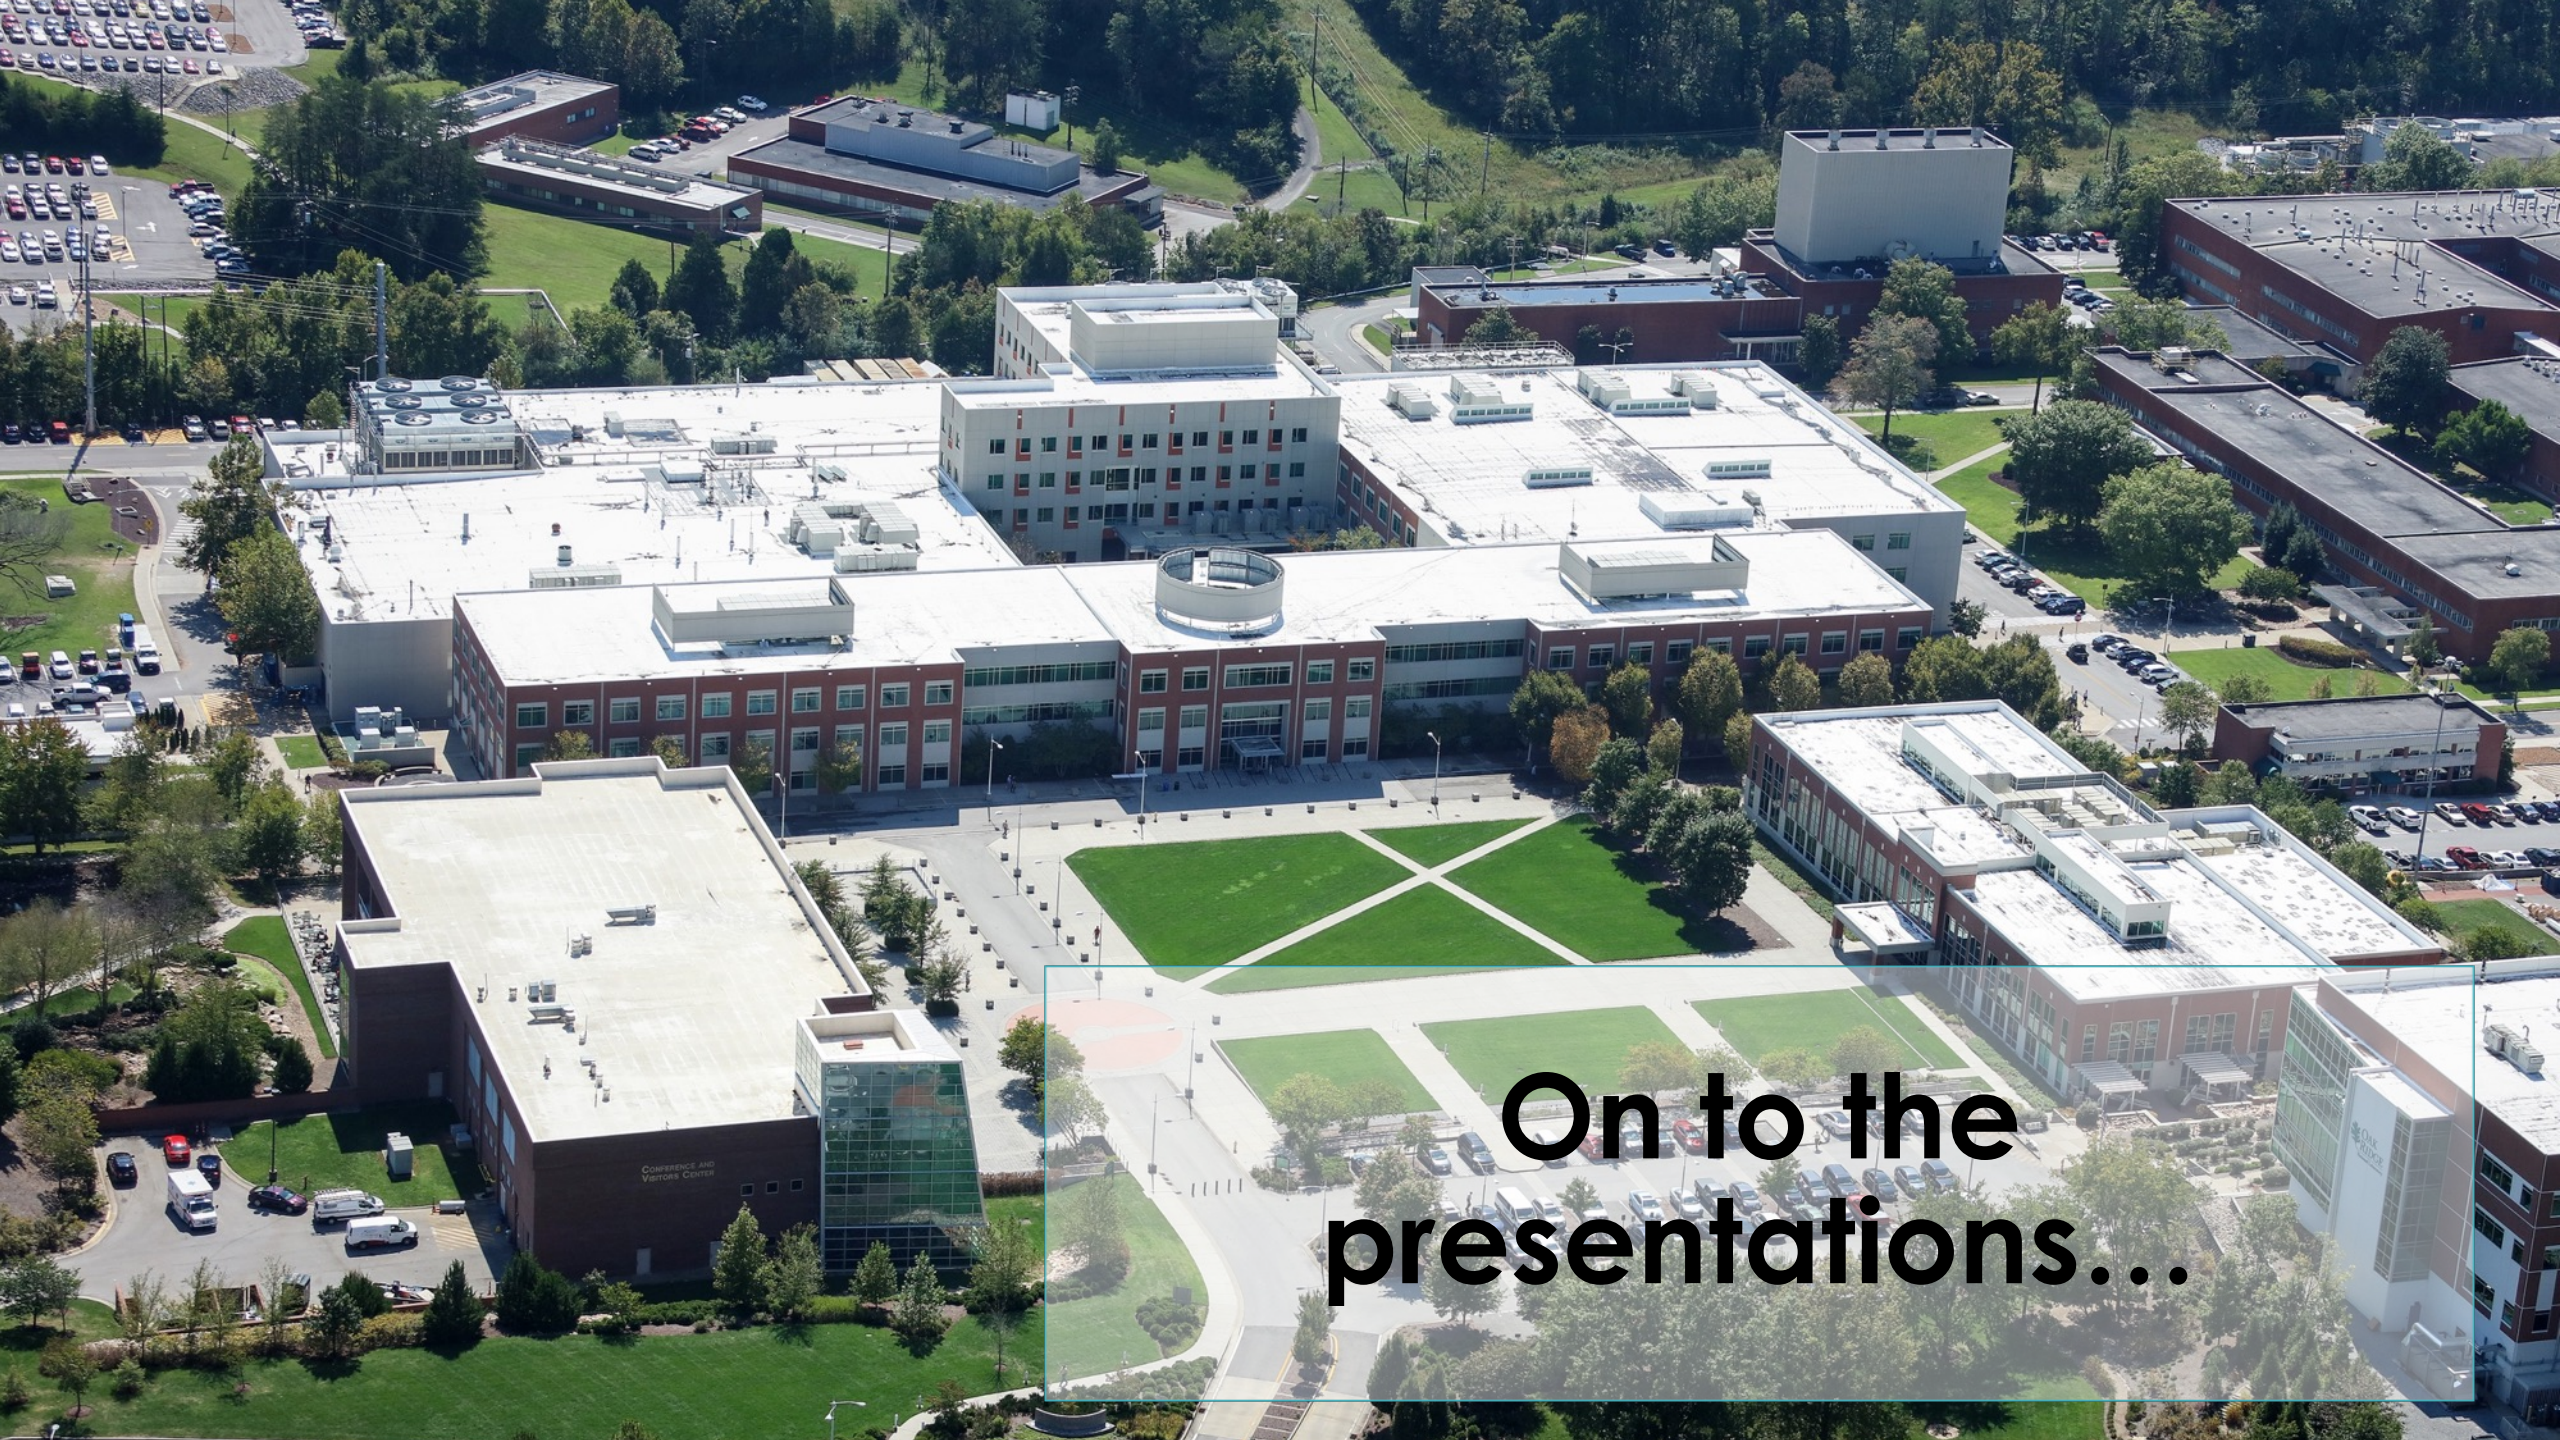

On to the presentations…