Q&A from June 2022 OLCF User Conference Call
Frontier Announcement & Overview

Can you comment on the different aspects of the acceptance test?
There are several phases. One is a general set of correctness tests of simple benchmarks, another is specific application performance benchmarks (in contract), then is performance testing on other applications (agreement between us+DOE), and finally a stability test where the system is run under constant workload for an extended time.

What is the state of AMD compiler/tool maturity on Frontier?
We're working with AMD to ensure available tools meet users' needs. (Some of this falls under acceptance testing.) AMD has received feedback from the CAAR and ECP teams that they are using to prioritize tool development efforts.

Will a cluster of similar architecture be available for training?
Yes, we will have a training cluster with an architecture identical to Frontier.

Will some multi-process GPU sharing (similar to NVIDIA's MPS) be available? Will the solution be time sharing or proper "process joining"?
Yes. Multi-process GPU sharing is natively supported on AMD GPUs, so no MPS options are needed. It is proper sharing of the GPU by the processes. The easiest use case is one or more MPI ranks to a single device. Use of multiple devices from a single MPI rank may be a different (and more difficult) use case.

Can the GPUs talk via MPI?
Yes, GPU-aware MPI is available on Frontier.

It was mentioned each MI250 contained 2 Graphics Compute Dies (GCDs) and thus would appear as 2 GPUs, and they have an infinity fabric shared between them. Are there any benchmarks on usage of that link or information on how applications might use it?
We're not aware of examples of codes using the "numa" nature of that packaging so far. AMD's suggestion is to treat each device as a single GPU.

Has the networking been tested as scales larger than Crusher, and what are the expectations for network at scale relative to what's been seen on Crusher?
We have some test runs at full scale, tests have not been strenuous but are not trivial. We anticipate scaling will be at least as good as Crusher within limits of Dragonfly topology.

Will it be difficult to adapt AI pipelines from Summit to Frontier?
Many basic AI tools have been ported to the system, and our expectation is for things to be supported well. You may need to seek replacements for NVIDIA-specific tools/libraries/intrinsics. (i.e. CUDA Graphs, etc.)

What scheduler will Frontier use?
Slurm

If I am at ORNL, can I visit/see/take pictures of Frontier?
Absolutely! Frontier is visible from the observation deck in Building 5600 and you are welcome to take pictures. (You can also go down the hall to the Summit viewing area and take pictures as well!)

Can the compute nodes talk directly to the Internet via SlingShot?
No, although if someone had an appropriate use case it may be possible to set up a proxy to allow limited connectivity. Even with that, there may be a bottleneck in the connection so the best way to move data from external sources is likely our Data Transfer Nodes.

When is the multi-tiered storage solution expected to be in place?
It should be in place before the end of the year, but won't be in production until we complete acceptance on it (which will likely be after acceptance of the rest of the system).

Will current Director's Discretion (DD) projects "roll over" to Frontier?
No, you will need to submit a separate application for a DD project on Frontier.

I have a current DD project on Summit, what is the best way to see if that application will work on Frontier or to see what is needed to extend functionality to Frontier? Should I apply for a new DD project?
Yes, you should apply for a new DD project on Frontier once we begin accepting applications.