

SUMMIT: *Scale new heights. Discover new solutions.*

How does the universe work? What is life? How can we make this world a better place? These are questions we have been asking ourselves for centuries—and now, with high-performance computing (HPC), they are questions that we are becoming increasingly more capable of answering.

The Oak Ridge Leadership Computing Facility's (OLCF's) 27-petaflop supercomputer, Titan, is America's largest leadership-class system. Located at Oak Ridge National Laboratory, it has enabled scientific breakthroughs in renewable energy, climate prediction, and materials science. Titan uses a heterogeneous architecture, or one that couples central processing units (CPUs) with general-purpose graphics processing units (GPUs) which allowed it to achieve a 1,000 percent increase in peak computing performance over its predecessor Jaguar, with only a 20 percent increase in peak power consumption.

OLCF's next machine will also have a heterogeneous architecture. The machine, named Summit, will arrive in 2017 and be ready for users in 2018. Summit will deliver more than five times the computational performance of Titan's 18,688 nodes, using only approximately 4,600 nodes. Each Summit node will contain multiple IBM POWER9 CPUs and NVIDIA Volta GPUs all connected together with NVIDIA's high-speed NVLink and a huge amount of memory. Each node will have over half a terabyte of coherent memory (HBM "high bandwidth memory" + DDR4) addressable by all CPUs and GPUs, plus an additional 800 gigabytes of NVRAM.

Core communication

Moving data is one of the most challenging tasks for any supercomputer: moving data from CPUs to GPUs, node to node, or into permanent storage. Summit will make huge advances over Titan in moving data. To improve latency between the processors, Summit will utilize NVIDIA NVLink high-speed interconnect which will allow direct communication between CPUs and GPUs 5x-12x faster than Titan. For moving data across the system, Summit will use dual-rail Mellanox EDR InfiniBand interconnects, providing 23 gigabytes per second of data sharing between the nodes. The interconnect will be configured as a non-blocking fat-tree, an arrangement that reduces data congestion (traffic jams) in the system.

In addition, each node will come with an extra 800GB of NVRAM, which can be used as a burst buffer or as additional node memory. Burst buffers act as intermediaries between the application and the file system. Since they are built from solid-state drive devices and have a high I/O throughput with low access latency, they can absorb the bulk data at a rate much higher than the file system. Data can be held there in the

extended memory, or it can be held temporarily while the burst buffers drain the data to the file system for scratch storage.

And, because supercomputers in general produce so much data, storage is a must. Summit will be connected to an IBM general parallel file system (GPFS) with 1 terabyte per second I/O bandwidth and 120 petabytes of disk capacity. The GPFS is highly scalable with a low latency rate, and provides parallel access to data and shared disks.

Summit will also connect to the OLCF High Performance Storage System (HPSS) for archival storage. The mass storage facility consists of tape and disk storage components, servers, and the HPSS software.

Summit as a science solver

Summit will break new ground and bring new understanding to many areas of science and engineering such as combustion science, nuclear power, biofuels, fusion energy, climate change, solar energy, energy storage, and catalysis, to name a few. These achievements will have an impact on America's economy, and provide new insights for improved quality of life worldwide.

By scaling up and parallelizing Jaguar's support systems and codes, Titan considerably advanced users' time to solution. With at least five times the performance of Titan, Summit will be able to perform several simulations in the time it takes Titan to perform just one.

With Summit, scientists who have been using HPC machines like Titan will be able to add much more complexity to their codes, which in turn will lend greater resolution to phenomena they are studying. Generally speaking, Summit users will be able to conduct much more realistic work than they were able to do on Titan—and as the HPC community knows, researchers are already using Titan to perform breakthrough science.

Transition to Summit: Preparing Titan Users for Fivefold Computing Power

Summit will provide revolutionary performance by way of evolutionary changes to the current Titan heterogeneous architecture, making Summit an ideal follow on system to Titan. By developing and refactoring applications to improve performance portability on accelerated architectures, Titan users will be better positioned to take advantage of Summit, as well as other next-generation leadership computing resources, and beyond. Users can create applications that explore performance portability and exploit untapped parallelism by:

1. Using accelerated [programming libraries](#) whenever possible
2. Preferring high-level compiler directives such as Open MP/Open ACC over low-level frameworks such as CUDA or OpenCL
3. Exposing as much node-level parallelism as possible
4. Relying on a suite of [development tools](#) to maximize parallelism.

The OLCF has been dedicated to providing best practices expertise for leadership-class computing for two decades. The OLCF created the Center for Accelerated Application Readiness, or CAAR, to help prepare codes for future generation systems. CAAR will establish eight partnership teams to prepare scientific applications for highly effective use on Summit. The partnership teams, consisting of the core developers of the application and OLCF staff, will receive support from the IBM/NVIDIA Center of Excellence at ORNL and have access to multiple computational resources. For more information about CAAR, please visit <https://www.olcf.ornl.gov/summit/caar-call-for-proposals>.

In addition to maintaining a robust programming environment on Titan, the OLCF is working with other high-performance computing centers to identify common strategies and best practices for developing performance portable applications. As part of the OLCF Center for Accelerated Application Readiness, early Summit users and OLCF user assistance staff will collaborate with the Argonne Leadership Computing Facility and the National Energy Research Supercomputing Center to ensure user applications will perform on multiple HPC architectures in the future.

Coinciding with the Summit launch, the OLCF will offer various tutorials, webinars, and workshops to help make sure users are ready on day one. During this three-year period of planning and development, please be aware that the OLCF's goal is to smoothly transition from Titan to Summit by providing an abundance of guidance to users during the process.