

Computing grid is racing the clock

by Katie Yurkewicz

To deal with the computing demands of the LHC experiments, scientists have created the world's largest, most international distributed-computing system. Developers of this intricate and innovative system have been pushing the boundaries of networking and grid computing for years, but time is finally running out. With start-up of the accelerator only a year away, will LHC computing stand up to the demands of thousands of physicists and a flood of data measured in petabytes?

The LHC experiments together will generate more than ten petabytes (or 10 billion megabytes) of data every year for more than a decade—a veritable deluge of information that will be in demand by at least 7000 scientists scattered around the globe. The convergence of so much data and so many physicists eager to use it meant that scientists couldn't rely on the old approach to particle-physics computing.

"The old model was that everyone had a computer account at the accelerator laboratory," explains Dario Barberis, ATLAS computing coordinator from the University of Genoa, Italy. "People would log in remotely to the central computer cluster, select events, take those selected samples home, and finish their work at their institution."

That model, in which all the data are stored at CERN and all 7000 physicists use CERN computing power for processing and analysis, was not practical for the LHC. The funds, electrical power, and human resources necessary for a single, all-purpose computing site would be too great for one laboratory to provide. So physicists and computer scientists from around the world came together to create a grid-computing system for the experiments, in which more than 100 small and large computing centers share the responsibility for storing, generating, and processing the data.

A need for speed

The new computing approach for the LHC was made possible by advances in two information technologies—high-speed networking and grid computing—over the past decade.

A global computing system wouldn't work without a mechanism to send huge amounts of data quickly and easily between computing centers. At the time the LHC computing system was conceived, high-speed, high-capacity networks were just beginning to circle the globe. It was up to the leaders of networking for the LHC to negotiate connections among participating computer centers all over the world. Today, after years of technical and organizational work, data can be sent between CERN and most of the 11 large, "Tier-1" LHC computing centers at an astonishing rate of 10 gigabits per second.

"This is really amazing," says Fermilab's Lothar Bauerdick, computing coordinator for the CMS experiment. "The speed of the connection between the disk and CPU in your laptop at home is less than 10 gigabits per second." So if all 10 gigabits between CERN and the Tier-1 center in Vancouver, for example, were to be used at once, a gigabyte of data would travel more than 5000 miles in less time than it takes to move millimeters within your computer.

These high-speed connections won't just benefit particle physicists. "We're pushing a lot to get good networks because we need to distribute the data," says CERN's Federico Carminati, ALICE computing coordinator, "but then the network stays there and can be used for other things. This is especially important for developing countries."

Five-year mission

While the networking challenges for LHC computing were mainly organizational, the same could not be said for the grid computing technologies. The tools and technologies that would turn the

Illustration: Sandbox Studio



grid computing idea into reality were in their infancy when work on the computing system for the LHC began. With the success of the experiments hinging on the success of grid computing, particle physicists were thrust to the forefront of international grid development. It has taken more than five years of collaboration by hundreds of physicists and computer scientists to build a working grid system, and there is still much more to be done.

"We're in a particular position because we are the only discipline—and I think this is true—that will live or die with the Grid," says Carminati.

The Worldwide LHC Computing Grid is made up of a system of interconnected national and international grid computing projects, including Enabling Grids for E-Science, NorduGrid and the Open Science Grid. Some of these projects have built on the LHC computing infrastructure to create something that benefits researchers from many other sciences. All the projects have banded together and devoted money, manpower, and resources to creating a grid system that physicists can live with.

About 100 worldwide computing centers have joined together under the umbrella of the Worldwide LHC Computing Grid Collaboration. The computing centers in the WLCG are organized into tiers, with each assigned a specific role in generating, storing, processing, or providing data for physicists to use. Raw and processed data are shipped from the Tier-0 center at CERN and shared among the eleven Tier-1 centers in Asia, Europe, and North America. Dozens of Tier-2 computing centers host smaller amounts of data, provide some storage and computing space for physicists, and generate the simulated data that is vital to particle physics discoveries.

Preparing for chaos

With thousands of computers, hundreds of institutions, dozens of international networks, and several large grid projects involved, a reliable system in which all parts work smoothly together is not easy to achieve. The international system that is now in place may be the first of its kind, but it is fragile and can't yet support thousands of users.

"People will be using the system in an unpredictable way, and how the system reacts to this chaos is a major question," says LHCb computing coordinator Nick Brook from the University of Bristol.

The full chain of data movement, from the detectors to the CERN computing center, to the Tier-1 and Tier-2 centers all over the world, and then into the hands of physicists, hasn't yet been fully tested. Automatically moving and processing petabytes of data will be a challenge, but one that will be surpassed by the demands placed on the system by thousands of physicists. About 100 researchers are currently testing the system, allowing developers time to increase capability as quickly as possible.

"There still needs to be a lot of work on the grid system to support thousands of scientists using it regularly," adds Bauerdick.

Each experiment and grid project has a plan to ramp up its computing capabilities over the next two years, so that no data will go to waste for lack of computing power or capability, and physicists will be able to start their search for the Higgs boson, supersymmetry, or new forms of matter, using all the latest tools and techniques as soon as the first collisions take place in the LHC.

"If you look at what we've achieved, it's a lot," adds Carminati. "But there are still many pieces that are missing, and the clock is running."

CERN, as the collection point for data, is the Tier-0 center of the LHC computing grid. It then distributes the data to 11 Tier-1 centers around the world. The amount of data is on a much larger scale than most people ever encounter. More than 10 petabytes (PB) of data will be collected each year by the experiments at CERN. For comparison, one gigabyte (GB) of data, a little more than can be stored on a CD, is one thousandth of a terabyte (TB), which in turn is one thousandth of a petabyte.

