SUN GRID UTILITY COMPUTING

BROOKHAVEN NATIONAL LABORATORY:

THE STAR EXPERIMENT

Key Features

- Access to the power of hundreds of CPUs to meet large-scale requirements for high-performance computing
- A true pay-as-you-go model for high-performance computing, based on a flat rate of $1 per CPU hour with no additional fees, contracts, or minimum hours
- The flexibility to choose from a library of widely used applications or port a specific application to Solaris™ 10 for use with Sun™ Grid
- Built-in, multiple-layered defenses at every level to secure data
- Powerful, industry-standard architecture with Sun Fire™ dual-processor servers (4 GB RAM per CPU), Solaris 10 OS, and Sun N1™ Grid Engine 6 software

Sun introduced Sun Grid Utility Computing to deliver large-scale computing power and related resources on a utility basis to projects that require them. And there are few better examples of a project that could be described as large-scale than the STAR experiment at Brookhaven National Laboratory.

An acronym for Solenoidal Tracker At RHIC — the laboratory's Relativistic Heavy Ion Collider — STAR tracks the thousands of particles produced by ion collisions at RHIC, searching for signs of something called the quark-gluon plasma (QGP), a form of matter that is believed to have last existed just after the Big Bang, at the dawn of the universe. The goal of STAR is to bring about a better understanding of the universe in its earliest stages, by making it possible for scientists to better understand the nature of the QGP.

The STAR experiment is a massive collaboration of 570 scientists and engineers representing 60 institutions in 12 countries. As the size of the collaboration and the scope of its work continue to grow, so does the challenge of having the computing power and data processing resources to carry out that work efficiently. Sun Grid Utility Computing has become a part of the distributed computing strategy that STAR is constructing to meet this challenge.

>

A two-fold IT challenge

The quest to understand the properties of matter and reaction involves accumulating and analyzing millions of images of ion collisions, a process that is both computing-intensive and data-intensive. The STAR detector captures images at a rate of about 100 per second and has accumulated several hundred million images so far in the course of the experiment. To provide an idea of the computing challenge associated with analyzing these images, consider that it can take as long as several minutes to reconstruct just one image using a typical Linux box. At that rate, some processing passes take as long as eight to ten months. If they could be done faster, it would leave more time for physicists' analysis of large datasets.

It's an extremely simple way to gain access to the power of hundreds of CPUs in one utility, anytime it's needed.

That brings us to the second part of the challenge: managing data. "A scientist will look at the initial analysis and then go on to look at the details, which requires even larger data samples," explains Jerome Lauret, RHIC/STAR Software and Computing Project Leader, "so the more scientists that are involved, the greater the scope of the data and dataset challenge."

Sun Grid Utility Computing has proven useful on the computing side of the equation, as a resource for the simulations of design, collisions, and other models that are essential to the research conducted by the experiment's physics working groups. Sun™ Grid has also supported simulations associated with ongoing research related to upgrades of the STAR detector — upgrades that will allow further advances in the experimental physics of heavy ion collisions.

Sun Grid Utility Computing

How it works

The STAR experiment relies on the Open Science Grid (OSG), a distributing computing infrastructure project of the National Science Foundation and the U.S. Department of Energy, to extend its computing capabilities. Recently, there has also been an effort to explore additional computing technologies and seek synergies among computing resources. Sun Grid is part of this effort.

"Sun Grid aims to deliver computing power and resources over the Internet, enabling us to optimize performance, speed our time to results, and accelerate innovation," as Maxim Potekhin, simulation leader for the STAR experiment, describes it.

The delivery of grid computing power over the Internet is something that sets Sun Grid Utility Computing apart. Sun Grid provides Web portal-based access to large-scale computing resources on an as-needed basis, at a flat cost of $1 per CPU hour.[1] It's an extremely simple way to gain access to the power of hundreds of CPUs in one utility, anytime it's needed. There are no contracts or minimum hours; customers pay only for the CPU time they consume. In some cases, especially for organizations that, unlike the STAR experiment, only require extended computing power at occasional peak times, this may be more cost-effective than building their own infrastructures for high-performance computing.

Organizations that use Sun Grid Utility Computing have access not just to the computing resources, but also to a library of on-demand high-performance computing applications. A list of available applications is published on Network.com, the Sun Grid Web portal, and it is searchable by industry and market. The customer simply browses the catalog, chooses an application, and pays Sun for the CPU hours used. License management is provided by Sun.

Many customers, of course, do not need any of the popular applications that are available for use with Sun Grid. Rather, they have their own applications to run. This is the case with the STAR experiment.

The STAR experiment's experience with Sun Grid

The STAR experiment has focused on using Sun Grid to run simulations using its own StarSim application. This required porting the application to Solaris 10. Regression testing was used to ensure that what was generated could be comparable to what the STAR simulation team generates in the platform that it normally uses. The only other requirement was for an interactive node, which Sun provided.

In STAR's case, the Web-based approach proved to be inappropriate, because of the constant need to bring tremendous amounts of data back to the team in the form of ongoing progress reports—a need better served by more of a scripting approach, rather than a Web-based approach. The STAR experiment team collaborated with Sun Grid developers to come up with a command line interface to use instead of the Web interface. The command line was built on the API library for Sun Grid.

> "Throughout the project, whether porting the application or running it, we found the Sun Grid team to be extremely responsive. They understood what we were trying to do, and they were always quick to provide the answers and resources we needed."
>
> **Dr. Jerome Lauret**
> RHIC/STAR Software and Computing Project Leader
> –Brookhaven National Laboratory

"Throughout the project, whether porting the application or running it, we found the Sun Grid team to be extremely responsive," notes Dr. Lauret. "They understood what we were trying to do, and they were always quick to provide the answers and resources we needed."

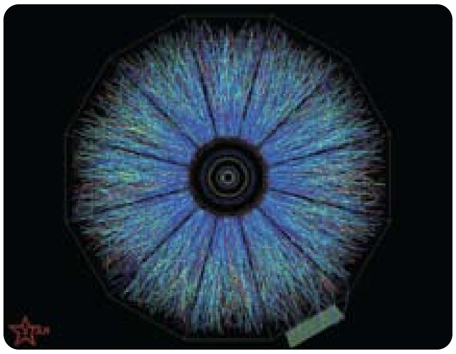

Courtesy of Brookhaven National Laboratory–RHIC Collaboration

The grand vision of the grid

Sun Grid is part of a much larger vision of grid computing that Dr. Lauret proposes for the STAR experiment. "This is the idea of not just a grid, but *the* grid," explains Dr. Lauret, "which would be the aggregation of a multitude of grids, both commercial and open, to create an extraordinary computing resource."

Realizing the vision hinges on defining what the grid should be, and on determining how Sun Grid and other grids can interoperate in the service of that definition. How, for example, can STAR interface not just with Sun Grid and various North American grids, but also with Asian and other grids, which may represent entirely different visions of what a grid should be?

"My main concern is to create bridges and opportunities, to consider the approach to grid computing from different sides, and to arrive at some ultimate paradigm of the grid," says Dr. Lauret. "As part of that process, we envision exchanging feedback with Sun about this paradigm and how it will be defined, and about what we would like to see in Sun Grid in support of the larger grid."

In this sense, the vision of grid computing in the STAR experiment is shaping the ongoing development of Sun Grid—and vice-versa.

1. **CPU hour is defined as follows**: For each job a customer submits and runs, CPU usage is aggregated and rounded up to the nearest whole hour. For example, if a job uses 1000 CPUs for one minute, the usage is aggregated as 1000 CPU minutes, or 16.67 hours. This is rounded up to 17 hours and the job is billed as $17.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

SunWIN # 512106 Lit # SWSS13221-0 8/07