7 things you should know about...

Grid Computing

Scenario

For years, Dr. Rayburn has been looking for tools to help his architecture students move beyond paper sketches and scaled-down models. He knows that as working architects, they will be using computer simulations that require not just design skill but proficiency with increasingly complex software and hardware. Unfortunately, his department cannot afford to purchase and support a computing system with the necessary processing capacity to run such advanced applications.

Over the summer, the university's IT staff, working with the computer science department, set up a computer grid running on the campus network. The grid connects nearly all university-owned computers, including those in labs, the library, as well as faculty and staff offices. The software that runs the grid gives local users priority for those machines, but when they are idle, their processors can be used over the grid. Using the power of the campus grid, Dr. Rayburn's students can now use sophisticated architectural design software that previously was unavailable because of its processing requirements. With the software, students can design buildings and other structures as well as the areas surrounding them, and create three-dimensional, interactive animations of their designs. As presentations, the animations allow viewers to "fly" over and around the scenes the students generate, zooming in and out and moving in any direction they want to go. The university's grid supplies enough unused computing power to process the animations fast enough for it all to function smoothly.

After several weeks of using the software, two of Dr. Rayburn's students persuade faculty in the meteorology department to connect a very large climatic database to the grid. The database includes data about the exact positioning of the sun and moon at any latitude on the globe during daily, monthly, and yearly cycles, as well as historical data on weather conditions for most parts of the world. With the database available on the grid, the students can incorporate seasonal changes into their animations. They can render a building at a particular latitude, at a specific time of the year or spanning weeks or months. Dr. Rayburn sees that with the new capabilities, his students are able to create better designs, ones that make more creative use of natural light—even as seasons change—and that demonstrate students' deliberation about how their structures interact with the environment.

1 What is it?

Computing grids are conceptually not unlike electrical grids. In an electrical grid, wall outlets allows us to link to an infrastructure of resources that generate, distribute, and bill for electricity. When you connect to the electrical grid, you don't need to know where the power plant is or how the current gets to you. Grid computing uses middleware to coordinate disparate IT resources across a network, allowing them to function as a virtual whole. The goal of a computing grid, like that of the electrical grid, is to provide users with access to the resources they need, when they need them.

Grids address two distinct but related goals: providing remote access to IT assets, and aggregating processing power. The most obvious resource included in a grid is a processor, but grids also encompass sensors, data-storage systems, applications, and other resources. One of the first commonly known grid initiatives was the SETI@home project, which solicited several million volunteers to download a screensaver that used idle processor capacity to analyze data in the search for extraterrestrial life. In a more recent example, the Telescience Project provides remote access to an extremely powerful electron microscope at the National Center for Microscopy and Imaging Research in San Diego. Users of the grid can remotely operate the microscope, allowing new levels of access to the instrument and its capabilities.

2 Who's doing it?

Many grids are appearing in the sciences, in fields such as chemistry, physics, and genetics, and cryptologists and mathematicians have also begun working with grid computing. Grid technology has the potential to significantly impact other areas of study with heavy computational requirements, such as urban planning. Another important area for the technology is animation, which requires massive amounts of computational power and is a common tool in a growing number of disciplines. By making resources available to students, these communities are able to effectively model authentic disciplinary practices.

How does it work?

Grids use a layer of middleware to communicate with and manipulate heterogeneous hardware and data sets. In some fields—astronomy, for example—hardware cannot reasonably be moved and is prohibitively expensive to replicate on other sites. In other

more ➪

EDUCAUSE Learning Initiative
advancing learning through IT innovation
Formerly NLII
www.educause.edu/eli

instances, databases vital to research projects cannot be duplicated and transferred to other sites. Grids overcome these logistical obstacles and open the tools of research to distant faculty and students. A grid might coordinate scientific instruments in one country with a database in another and processors in a third. From a user's perspective, these resources function as a single system—differences in platform and location become invisible.

On a typical college or university campus, many computers sit idle much of the time. A grid can provide significant processing power for users with extraordinary needs. Animation software, for instance, which is used by students in the arts, architecture, and other departments, eats up vast amounts of processor capacity. An industrial design class might use resource-intensive software to render highly detailed three-dimensional images. In both cases, a campus grid slashes the amount of time it takes students to work with these applications. All of this happens not from additional capacity but through the efficient use of existing power.

4 Why is it significant?

Grids make research projects possible that formerly were impractical or unfeasible due to the physical location of vital resources. Using a grid, researchers in Great Britain, for example, can conduct research that relies on databases across Europe, instrumentation in Japan, and computational power in the United States. Making resources available in this way exposes students to the tools of the profession, facilitating new possibilities for research and instruction, particularly at the undergraduate level.

Although speeds and capacities of processors continue to increase, resource-intensive applications are proliferating as well. At many institutions, certain campus users face ongoing shortages of computational power, even as large numbers of computers are underused. With grids, programs previously hindered by constraints on computing power become possible.

5 What are the downsides?

Being able to access distant IT assets—and have them function seamlessly with tools on different platforms—can be a boon to researchers, but it presents real security concerns to organizations responsible for those resources. An institution that makes its IT assets available to researchers or students on other campuses and in other countries must be confident that its involvement does not expose those assets to unnecessary risks. Similarly, directors of research projects will be reluctant to take advantage of the opportunities of a grid without assurances that the integrity of the project, its data, and its participants will be protected.

Another challenge facing grids is the complexity in building middleware structures that can knit together collections of resources to work as a unit across network connections that often span oceans and continents. Scheduling the availability of IT resources connected to a grid can also present new challenges to organizations that manage those resources. Increasing standardization of protocols addresses some of the difficulty in creating smoothly functioning grids, but, by their nature, grids that can provide unprecedented access to facilities and tools involve a high level of complexity.

6 Where is it going?

Because the number of functioning grids is relatively small, it may take time for the higher education community to capitalize on the opportunities that grids can provide and the feasibility of such projects. As the number and capacity of high-speed networks increase, however, particularly those catering to the research community and higher education, new opportunities will arise to combine IT assets in ways that expose students to the tools and applications relevant to their studies and to dramatically reduce the amount of time required to process data-intensive jobs. Further, as grids become more widespread and easier to use, increasing numbers and kinds of IT resources will be included on grids. We may also start to see more grid tie-ins for desktop applications. While there are obvious advantages to solving a complex genetic problem using grid computing, being able to harness spare computing cycles to manipulate an image in Photoshop or create a virtual world in a simulation may be some of the first implementations of grids.

7 What are the implications for teaching and learning?

Higher education stands to reap significant benefits from grid computing by creating environments that expose students to the "tools of the trade" in a wide range of disciplines. Rather than using mock or historical data from an observatory in South America, for example, a grid could let students on other continents actually use those facilities and collect their own data. Learning experiences become far richer, providing opportunities that otherwise would be impossible or would require travel. The access that grid computing offers to particular resources can allow institutions to deepen, and in some cases broaden, the scope of their educational programs.

Grid computing encourages partnerships among higher education institutions and research centers. Because they bring together unique tools in novel groupings, grids have the potential to incorporate technology into disciplines with traditionally lower involvement with IT, including the humanities, social sciences, and the arts. Grids can leverage previous investments in hardware and infrastructure to provide processing power and other technology capabilities to campus constituents who need them. This reallocation of institutional resources is especially beneficial for applications with high demands for processing and storage, such as modeling, animations, digital video production, or biomedical studies.

EDUCAUSE Learning Initiative

advancing learning through IT innovation

www.educause.edu/eli

January 2006