University of Florida's highly secure Grid bolsters scientific research.

Overview

- ***Challenge***
  University of Florida recognized a barrier to research progress caused by scientists' lack of trust in the security of collaborative high-performance computing systems
- ***Why Become an On Demand Business?***
  Resilient, highly secure high-performance computing resources could help improve research productivity and advance biological, environmental, computer science and engineering disciplines
- ***Solution***
  A Grid computing environment that harnesses the power of the users' existing systems to create virtual machines with resources on demand, providing excellent security to encourage information sharing
- ***Key Benefits***
  Grid is designed to help keep data and applications safe from tampering and to enable collaboration; meets variable demands for computing power, with very low maintenance costs

By harnessing the power of virtual machines, the University of Florida is taking advantage of computing resources that are far greater than any one laboratory can afford.

Gauging by the activity at the Advanced Computing and Information Systems Laboratory (ACIS) at the University of Florida, Gainesville, some of the greatest advances in medicine, computer science, manufacturing and other realms of human endeavor may soon come from virtual sources. The laboratory is performing research on computing and information processing systems to support vast expanses of scientific study.

Funded in part by the National Science Foundation and top corporations, ACIS (*www.acis.ufl.edu*) needs to continually push the envelope of

> *"Researchers — concerned with the security of their data and applications in a collaborative environment with shared resources — have not interacted much, which has hampered productivity."*
>
> *–Dr. José Fortes and Renato Figueiredo, Professor and Assistant Professor of Electrical and Computer Engineering, University of Florida*

On Demand Business Benefits

- Researchers can utilize compute-intensive research applications to make discoveries in the fields of medicine, engineering and education, among others
- Grid is designed to help keep data and applications safe from tampering, offering resilient, highly secure computing resources that research scientists can use
- The Grid environment can help meet variable demand for computing power by generating virtual servers
- Grid management software is designed to help reduce maintenance costs by automatically provisioning new servers and quickly updating new application releases

knowledge to keep the facility at the top echelon of research institutions in the United States. That's why it recently decided to tackle one of the most challenging barriers in science: enabling researchers to collaborate.

Searching for a resilient, secure system

"Collaboration among biologists, physicists, chemists, electrical engineers and mechanical engineers could speed scientific research and discovery," explain ACIS lab leaders Dr. José Fortes and Renato Figueiredo, professor and assistant professor of electrical and computer engineering at the University of Florida. "But researchers in these areas—concerned with the security of their data and applications in a collaborative environment with shared resources—have not interacted much, which has hampered productivity."

It's not that researchers have lacked access to computing power to support their work. For a number of years, they have had at their disposal highly powerful supercomputing clusters capable of performing the data-intensive computation required to conduct exploration in many areas of research. But these scientists have feared that when existing supercomputing resources are made available over the Internet—the typical vehicle for sharing data and scientific tools—unauthorized individuals might gain access to the supercomputing system and steal or destroy their research. These security concerns have made data sharing meager at best. "We wanted to advance the field of computer science by bringing about a quantum increase in collaboration," says Fortes.

To dissolve the collaboration barrier, Fortes and Figueiredo needed a security-rich computing resource that could better facilitate the sharing of computational tools. Such a resilient computing solution could provide scientists with the confidence they needed to share their code and collaborate with other scientists across disciplinary lines. Scientists would be able to tap unlimited computing resources in geographically dispersed environments that would help to improve research productivity and bring about unprecedented results. "We felt we had a wonderful opportunity to help science evolve as well as keep our institution at the forefront of knowledge," Fortes adds.

To help make its supercomputing resources more resilient—impervious to security breaches—ACIS decided to tap into the capabilities of virtual computing (networks of computers that appear to users as one large computer). Also known as Grids, these virtual computers have processing power that is aggregated from a large number of distributed machines.

Like other virtual computing solutions, the Grid that ACIS constructed is a network of physical machines that appears to researchers as one system. It provides the same tools that individual physical machines have for communicating

with networks and managing workloads. As such, the ACIS Grid can facilitate collaboration across departments, institutions and even countries. In addition, the Grid generates virtual servers in response to users' requests, so it can grow on demand.

However, what makes the ACIS Grid highly secure is the fact that its virtual servers are isolated—from each other, the network and its physical servers—by virtualization software. Although the virtual servers can be accessed from the Internet, only authorized users can use volumes of data and applications from within the virtual machines assigned to them by the middleware. When the researcher is finished, the management software simply destroys the virtual server connection and nothing remains.

Virtual machines working together

To advance the science of Grid computing, the ACIS group decided to concentrate its attention on a technology invented by IBM and associated with the IBM mainframe: IBM VM, or Virtual Machine. VM has been used to create guest images, or virtual servers, to cost-effectively distribute the mainframe's computing power. Of late, the combination of the Linux operating system and virtual servers running on the IBM **@server** zSeries® has become one of IBM's most powerful—and successful—server offerings. The zSeries 800 server is a dedicated Linux server, capable of generating hundreds of Linux virtual servers, each providing mainframe scalability and stability along with the reliability and cost savings associated with the Linux operating system.

"We chose the IBM zSeries 800 because it has the most advanced virtualization technology of any server to date," says Fortes. In addition to the z800, ACIS is using an IBM **@server** xSeries® 1300 cluster with 32 compute nodes running VMware software from IBM Business Partner VMware. VMware enables IBM Intel processor-based servers such as the xSeries to create virtual servers in similar fashion to the zSeries. The Grid runs both on Red Hat Linux and on SUSE LINUX Enterprise Server, Version 8. Backing up data on the Grid is an IBM TotalStorage® Enterprise Storage Server®. IBM Business Partners Cornerstone Systems and e-TechServices provided the hardware and initial system integration.

Open standards-based system for greatest efficiency

The multi-tiered infrastructure using the z800 server and the xSeries cluster comprise the physical nodes of the Grid. Using open standards in part created by the Globus Alliance, ACIS is using the Globus 2.0 Toolkit, the Java™ CoG kit 0.9.3 and a virtual file system. ACIS has created a system called In-VIGO to apportion workloads and storage resources on demand. The Grid is designed to apportion data and applications to its virtual resources automatically, making this a virtual system in all aspects: network, machine, data and applications.

Key Components

Software

- Red Hat Linux
- SUSE LINUX Enterprise Server, Version 8

Servers

- IBM **@server** zSeries 800
- IBM **@server** xSeries
- IBM TotalStorage Enterprise Storage Server

Business Partners

- VMware
- Cornerstone Systems, Inc.
- e-TechServices

"IBM virtual machine technologies help reduce administrative complexity and maximize utilization capabilities, increasing the reliability, security and manageability of its servers."

–Renato Figueiredo, Assistant Professor of Electrical and Computer Engineering, University of Florida

IBM tools provide outstanding reliability

According to Figueiredo, this breakthrough in highly secure Grid computing gives researchers much more confidence in sharing information than they have had before. "Virtual servers are isolated from physical machines (and other virtual servers) by VM monitors—a software layer that is leaner and easier to secure than a full-blown operating system," he says. "This isolation provides a degree of security that is difficult to match."

IBM's role in this innovative study was practically a given since the ACIS research focuses on applying IBM technology to Grid architecture. The use of hardware provided by IBM—a leader in grid computing technology—was a necessity. However, the group asserts that IBM equipment stands out in the marketplace for its exceptional price and performance. "IBM virtual machine technologies also help reduce administrative complexity and maximize utilization capabilities, increasing the reliability, security and manageability of its servers." says Figueiredo.

Now in the initial stage of testing, it is evident that the ACIS Grid offers enormous computing capabilities to scientists at the University who are hungry for computing power. ACIS will use the Grid to help advance research on such widely disparate fields as transnational digital government, implantable intelligent devices and memory hierarchies. Additional prospects for the future include the ability to provision and update virtual servers with the latest software tools automatically, saving large amounts of maintenance time, while enabling collaboration among large groups of scientists.

In addition to running simulations of nanodevices, the Grid itself is the subject of its own research. "We hope to find ways to provision virtual servers automatically with some of the complex tools that they will need," says Figueiredo. "A large part of our research is to endow the machine with the properties we want it to have so that scientists, engineers and other researchers can take advantage of this model in the future." The group also plans to implement the Open Grid Services Architecture (OGSA)—the open standards-based architecture developed by the Globus Project—in the near future. To enable other groups to work with them, the ACIS lab is joining its Grid to Grids at Purdue University and other universities.

"We hope that this will be the beginning of the interdisciplinary research that we envisioned at the start of the project. We are proving and will continue to prove that we can create highly secure Grids using IBM and VMware technology," says Fortes.

For more information

Please contact your IBM sales representative or IBM Business Partner.

Visit us at:
ibm.com/ondemand

For more information about Cornerstone Systems and e-TechServices, please visit:
www.csihome.com
www.e-techservices.com

IBM Corporation
Corporate Marketing
New Orchard Road
Armonk, NY 10504
U.S.A.

Produced in the United States of America
12-03

GM13-0446-00