

----- Original Message -----

Subject: Grid Computing Security and Policy Management

Date: Wed, 05 Jun 2002 13:03:55 -0400

From: Deni Connor <DConnor@nww.com>

To: [rajkumar@csse.monash.edu.au](mailto:rajkumar@csse.monash.edu.au)

I am writing a story for Network World on the security of grid computing. I saw your name on a usenet group and was interested in your opinion. I may use your comments in my story.

1. Have you implemented a grid project? To do what? If so, is it within the organization? Or between organizations?
2. Is there reason to be concerned about the security of data flowing across the Internet between different companies? Why?
3. Once a company has been authenticated to another companies resources, how do you control its ability to spawn tasks and run applications?
4. How do you keep one company out of anothers data?
5. In intra-company grids, how do you make sure policies set for one department, do not hinder or alter tasks in another department?
6. Why aren't the security measures like PKI and Kerberos which we use in local and wide-area networks sufficient in grids?
7. What is your title? Who do you work for? Where are they located?

Thanking you in advance for your answers.

Deni Connor  
Senior Editor

Network World, The Leader in Network Knowledge  
Tele: 512-345-3850  
Fax: 512-345-3860  
8815 Mountain Path Circle  
Austin, Texas 78759  
Covering Storage, Novell, Linux and open source  
For Networking News & Info: [www.nwfusion.com](http://www.nwfusion.com)  
Find out more about Network World, Inc.: [www.networkworld.com](http://www.networkworld.com)

## **Rajkumar's answers to Network World's Press Inquiry:**

*1. Have you implemented a grid project? To do what? If so, is it within the organization? Or between organizations?*

I have been involved in the research and development of Grid technologies, particularly economic and users' quality-of-service requirements driven resource management and scheduling system for World-Wide Grid since 1999. I have contributed towards the development of a grid resource broker called Nimrod-G that has the ability to schedule large-scale computational and data intensive applications such as drug design on world-wide distributed resources. The Nimrod-G software technology research and development is carried out by Jon Giddy (Distributed Systems Technology Centre) and Rajkumar Buyya (currently at the University of Melbourne) under the leadership of Professor David Abramson (Monash University). I have developed a grid modelling and simulation environment called the GridSim Toolkit in collaboration with Dr. Manzur Mushed (Gippsland School of Computing and IT at Monash). I have also developed the Virtual Laboratory environment for distributed molecular modeling for drug design on the Grid in collaboration with Kim Branson (WEHI). In fact, all these collaborative efforts have truly contributed towards my Ph.D. thesis, Economic-based Distributed Resource Management and Scheduling for Grid Computing, submitted to Monash University.

I have been continuing with these efforts and developing Grid technologies and tools that enable distributed processing of data-intensive science applications such as High Energy Physics (in collaboration with the Experimental Particle Physics group at the University of Melbourne) and Brain Activity analysis (in collaboration with Osaka University, Japan) on the World-Wide Grid (WWG). The WWG serves as a testbed for experimentation and evaluation our research ideas, tools, and applications. It has resources distributed around the world and they are owned by different organizations with whom we collaborate. Recently, we have initiated a new project that aims to design and develop advanced economic schedulers for enterprise Clusters and Grids in collaboration with Wolfgang Gentsch (Sun Microsystems, USA).

For further information on the Nimrod-G project, please see:

<http://www.csse.monash.edu.au/~david/nimrod/> For pointers to our most recent and emerging Grid projects, please visit the Grid Computing and Distributed Systems (GRIDS) Laboratory webpage: <http://www.cs.mu.oz.au/~raj/grids/>. The GRIDS Lab's major thrust is in the design and development of next-generation computing systems and applications that aggregate or lease services of distributed resources depending on their availability, capability, performance, **cost**, and users' quality-of-service requirements.

*2. Is there reason to be concerned about the security of data flowing across the Internet between different companies? Why?*

Yes, there is enormous concern about data and application security both during its flow across the Internet and also when staged on the grid resource. The first concern is mainly

because it is possible for someone to tap your data (passive listeners) and possibly modify it (active listeners who may change its meaning) on its path. The second concern is that when you use others computers in the grid, it is possible that the owners of those computers may read your data. These can be addressed by sophisticated encryption techniques both during transmission and also during their representation/storage on external resources.

*3. Once a company has been authenticated to another companies resources, how do you control its ability to spawn tasks and run applications?*

If one is able to authenticate successfully, means that the user has authorized to spawn tasks and run applications by resource owners. At the moment, the Grid as a technology does not provide any special control mechanisms. By default, existing operating systems do provide the means to enforce certain policies such as who is allowed to read, write, and update which data and execute which programs.

*4. How do you keep one company out of anothers data?*

Operating systems already provide means to control who is allowed to access data. For example, on Unix systems there is a support to set permissions such as only the owner of resource is permitted to access data.

*5. In intra-company grids, how do you make sure policies set for one department, do not hinder or alter tasks in another department?*

The cluster resource management systems such as Sun Grid Engine (SGE) provide mechanisms for defining resource share and allocation policies to various departments across the enterprise. They also allow the users to define their quality-of-service requirements such as the deadline, so that the system can allocate resources to meet the users' requirements irrespective of their resource share. However, if everyone states high-demand and tight deadline, then it will strictly enforce allocation policies to ensure that each user gets their minimum share.

*6. Why aren't the security measures like PKI and Kerberos which we use in local and wide-area networks sufficient in grids?*

The PKI and Kerberos technologies essentially provide means for handling authentication issues (who is allowed to access a resource, but they don't provide technologies to protect your data during transmission and also from other users or even resource owners.

*7. What is your title? Who do you work for? Where are they located?*

I am a faculty member (the US equivalent of Assistant Professor) and I am leading/directing research activities of the **Grid Computing and Distributed Systems (GRIDS) Laboratory** within the Department of Computer Science and Software Engineering at the University of Melbourne, Australia.