

---

# Bits & Bytes

No. 179

RZG Computer Bulletin

Dec 2006

Computing Center of the Max Planck Society and the Institute for Plasma Physics\*

---

## The Next Generation Supercomputer of the Max Planck Society at RZG

At the end of 2005, the scientific council of the RZG had triggered the process to start a procurement for a follow-up system of the current MPG supercomputer at RZG, an IBM Power4 based system. This system had started operation in 2001 and consists of about 1000 processors with a peak performance of 5 TF/s and a main memory of about 2 TB.

In order to provide adequate supercomputing resources for those groups in the MPG which heavily depend on such a system for their research, a follow-up system at least 10-fold more powerful should be procured.

30 groups from 20 institutes applied for procuring such a new powerful supercomputer. After obtaining “green light” from both the BAR of the MPG in mid March 2006, and from the “Wissenschaftsrat” of the German government at the end of March 2006, a European procurement was started at the beginning of April 2006.

The procedure selected was the so-called “Wettbewerblicher Dialog”, newly introduced into German law in September 2005. This procedure includes a dialogue phase to be able to optimize the technical configurations and solutions for the specific and complex requirements. By mid May, the six companies Bull, Cray, HP, IBM, NEC with partner SUN, and SGI had qualified as bidders for the procurement. Final offers were received by end of August. The most economic offer was obtained from IBM. A contract was signed before the end of November 2006.

The new system will consist of Power6 processor based 32-way nodes with an Infiniband Interconnect. The peak performance of the system will exceed 100 TeraFlop/s. In addition, a small BlueGene system (two racks) will be installed. Highly scalable applications of the Max Planck Society will be able to use the system on several thousands of processors already for production. Even more important, it will become possible to prepare, develop and optimize applications for future petascale computing.

Hermann Lederer

## Refurbishment of the Old RZG Computer Machine Hall

In 2001/2002, a new RZG computer hall was built as an attachment of the old room dating from 1968. Developments of demands for hosting systems at RZG, and the increasing demands of systems especially with respect of energy consumption and cooling capacity, made it inevitable to expand the RZG facilities by additionally refurbishing the old computer room and installing further electrical and cooling systems.

These works have been terminated in early December 2006 and lead to a doubling of available floor space and cooling capacity.

First Linux clusters will be installed in the new room still this year. Installation of the new supercomputer will start next year.

Hermann Lederer

## SPAM Reduction by Greylisting

### Customizing Your Server-Side Email Setup

The possibilities of customizing your server-side email setup on the page [www.rzg.mpg.de/networking/mailsetup.html](http://www.rzg.mpg.de/networking/mailsetup.html) have been extended: To reduce the amount of spam you can switch on a technique called “greylisting” for your email addresses. In that case incoming emails from an unknown sender will be rejected by us with a temporary failure code. For wanted mail, the sender’s mailserver will retry after a few minutes and we then accept the message. Future mails from the same sender will be accepted without delay. On the other hand, the spam senders’ software usually doesn’t bother to retry a temporarily failed delivery attempt, so the amount of spam is greatly reduced. In rare cases, the greylisting mechanism may fail with wanted mail; to work around this we keep a “whitelist” of mail server addresses that are known not to work with greylisting. This list will be extended as needed.

We already described another technique to sort out spam by setting up a filter in your email client, based on the tagging we provide: [www.rzg.mpg.de/networking/spam/index-en.html](http://www.rzg.mpg.de/networking/spam/index-en.html)

By common request you can now delete incoming mails that exceed a configurable “spaminess” on the server. But be advised that there is no way to recover these. We recommend to sort mail on the client side.

---

\*Max-Planck-Institut für Plasmaphysik, Boltzmannstraße 2, D-85748 Garching bei München, tel.: +49(89) 3299-01, e-mail: [benutzerberatung@rzg.mpg.de](mailto:benutzerberatung@rzg.mpg.de), URL: <http://www.rzg.mpg.de/>  
Editorial: Dr. Roman Hatzky, Tel. -1707

In addition to sending vacation notes and redirecting emails during an absence, it is now possible to redirect emails depending on specified patterns in the sender's address or subject.

Klaus Desinger

### **Why Was This Message Delivered to Me? – A Frequently Asked Question**

Sometimes you will receive email messages that are apparently not addressed to you. As with paper mail, an email has a recipient address on the “envelope” and in the “To: ” field of the header, but only the envelope address is relevant for delivery. Both addresses are not necessarily identical; they are usually different in the case of mailing lists, blind carbon copies or when you resend (not forward) an email to another person. Spammers just misuse this feature in an attempt to make their messages more interesting. If you receive a message, then your address was certainly on the envelope, even if it does not appear in the header.

Klaus Desinger

### **Virtual Private Network – Accessing Local Resources from the Internet**

Virtual Private Network (VPN) technology allows remote users to access resources that are otherwise only available on campus. A secure connection is established by tunneling encrypted traffic through the public Internet. The remote computer is assigned an IP address from the Garching MPG campus address range – the campus network is virtually extended to the remote user.

For the necessary client software and the configuration parameters please see:

[www.rzg.mpg.de/networking/vpn/](http://www.rzg.mpg.de/networking/vpn/).

Klaus Desinger

### **New Tape Library at RZG**

A new automated tape library for use by the TSM backup and archive servers has recently been installed at RZG to replace our old ADIC AML/E library, which has been in operation for the past 12 years.

A tape library (also known as tape silo) is a large and enclosed cabinet with many slots to store computer tapes and with several tape drives. One or more robot arms move the tapes from the slots to the tape drives and vice versa. A tape library can store huge amounts of data.

The new library is a Sun SL8500. In the configuration purchased by RZG, it has 6500 slots for tapes of type LTO3 and 24 LTO3 tape drives linked through Fibre-Channel connections to two different multiprocessor computers hosting a total of 16 TSM server databases. LTO3

tapes have a capacity of 400 GB uncompressed data (or about 600 GB after compression, which is done in hardware by the drives themselves), which gives us a total capacity of about 2.5 Petabytes (~ 2500 TB) uncompressed or 3.8 PB compressed data. The LTO3 tape drives can read and write native data at 80 MB/s. They can also read and write the older generation LTO2 tapes, in use by the old tape library, allowing for a smooth transition to the new system.

The library can be expanded to a total of about 9000 tape slots and 64 tape drives. Further expansion is possible by connecting together up to 32 single libraries, giving a grand total of about 300 000 tapes and over 2000 tape drives.

Manuel Panea

### **Linux-Clusters: Consolidation of the Used Distribution**

When the first Linux Clusters in blade-center technology were installed at RZG in January 2003, the desktop operating system “RedHat 9” was used. This distribution became outdated soon and a distribution with a longer life-cycle was required. Another drawback were missing software-drivers for the “state-of-the-art” hardware usually installed at RZG. For this reason, the commercial distribution SuSE Linux Enterprise Server 9 (SLES9) has been chosen. It is guaranteed to be supported for many years and most hardware-vendors provide drivers for it, when there is none for this hardware in the official kernel. The 64 Bit-Opteron clusters have always been run with SLES9, the 32 Bit clusters are now upgraded to this distribution. This eases the management and the synchronization of the software stack available on both the 32 Bit and 64 Bit linux cluster-nodes dramatically. On the server side, however, some applications may require different distributions, such as “Scientific Linux”, so that also in the future more than one Linux-distribution will be in use at RZG.

Christof Hanke

### **Data Visualization Support**

Data visualization support is available through a new RZG member, Ralph Bruckschen ([Ralph.Bruckschen@rzg.mpg.de](mailto:Ralph.Bruckschen@rzg.mpg.de)). Focus is on the visualization of complex and very large data sets as they may arise from supercomputing simulation runs. Current visualization projects include supernovae and cosmological simulation data. The main software used is Amira 4.1, augmented by self-developed out-of-core algorithms, on a graphics workstation with a high I/O rate (800 MB/s).

Ralph Bruckschen