

Garching Computing Center of the Max Planck Society and the Institute for Plasma Physics*
Boltzmannstraße 2, D-85748 Garching bei München

Next Generation Max Planck Supercomputer at RZG

Hermann Lederer

The current supercomputers of the Max Planck Society at RZG, an IBM Power6 and a BlueGene/P system, are to be replaced in 2012/2013 with a PetaFlop/s class system. A Europe-wide procurement was done in 2011, supervised by the corresponding advisory boards, and finalized last December. The decision was taken for an IBM system. The contract includes a two-phase delivery: an initial smaller system in summer 2012 and the main delivery in Q2 2013.

The initial system will be based on the Intel SandyBridge processor and an FDR InfiniBand Interconnect, together with the I/O subsystem. It will consist of 610 compute nodes (590 with 64 GB RAM, 20 with 128 GB RAM),

with 16 cores per node, a total of 9760 cores with a main memory of 40 TB and a peak performance of about 200 TeraFlop/s. The software stack will be very similar to the one which is currently being provided on the Linux clusters hosted by RZG, i.e. it will be based on Intel and GNU compilers, Intel MKL and MPI, together with the common extensions like FFTW, WSMP and other tools and libraries.

The new system will be operated in parallel with the current Power6 HPC system (with 120 TF peak) until the end of 2012. At the end of 2012 the Power6 system will be taken out of operation, while the BlueGene/P system will be operated until the installation of the new main system has been completed.

HPSS - High Performance Storage System at RZG

Manuel Panea, Andreas Schott, Ingeborg Weidl

Since decades, RZG is providing migrating file systems, such as the r-tree of the HPC cluster and the m-tree of AFS. The migration facilities were based on IBM's TSM/HSM storage management software for the hierarchical storage management. To cope with the growing demand for data storage, RZG switched to a new and powerful software package, which is used in many large computing centers around the world where huge amounts of data have to be handled. The new HPSS (High Performance Storage System) software was activated in November 2011. HPSS is characterized by extreme scalability and throughput. The new system will provide much faster access to migrated data on tape for both, the migrating r-tree of the HPC cluster, and the m-tree of AFS.

From the user's point of view, beside the improved ac-

cess, there is no change in functionality. As in the past, the underlying software is transparent to the end user. The data currently located in either of the file system trees was transferred to the new HPSS system without any user action required.

To remind you, data written to these file systems is automatically moved from disk to tape. You can manually force the recall of a migrated file by using any command which opens the file (e.g. the command 'file'). On the Power6 HPC cluster 'vip', you can see which files are resident on disk in /r and which ones have been migrated to tape with the command 'ghi_ls' (optionally with the option -l), while in AFS the command 'fs ls' provides information about the status of a file.

*Tel.: +49(89) 3299-01, e-mail: benutzerberatung@rzg.mpg.de, URL: <http://www.rzg.mpg.de/>

VisIt, a parallel visualization and data analysis tool

Klaus Reuter

VisIt, a parallel visualization and data analysis tool

The RZG is offering remote visualization services based on a central visualization infrastructure, enabling scientists to analyze and visualize complex datasets produced by numerical simulations on the supercomputers at RZG. This article introduces the [VisIt software](#) which is a free visualization and analysis tool for scientific data. It has proven in practice to be a stable and scalable solution for a wide range of visualization requirements.




Figure 1: Core collapse supernova (Type-II) explosion dynamics in 3D. Simulation by F. Hanke et al. (MPA), Visualization by E. Erastova & M. Rampp (RZG) using VisIt. A single timeslice comprises $400 \times 60 \times 120$ zones on a non-uniform, time-dependent polar grid. In total 1000 HDF5 output files a 1 GB were rendered for generating an animation

VisIt can be used to visualize scalar and vector fields defined on two- and three-dimensional structured and unstructured meshes. Users can quickly generate static or animated visualizations from their data. VisIt offers a large number of tools such as contour plots on planes or arbitrary surfaces, iso-surface plots, volume rendering, and arrow or streamline plots to visualize vector fields. It is able to read most of the commonly used file formats and also can be easily extended with new reader plugins. The VisIt rendering engine is MPI parallel and is able to handle very large data set sizes. Using the RZG visualization cluster it is possible to process data volumes of up to 1 terabyte per time step. Up-to-date versions of VisIt are available on the cluster. The MPI-parallel rendering mode is preconfigured and can be started conveniently from the VisIt GUI.

VisIt can be used interactively, or it can be controlled using Python scripts. In practice, it is useful to find suitable parameters and views by first using the interactive mode. This information can be used to write a Python script describing the complete (animated) visualization. As a second step, the rendering process can then be run as an MPI job on the VIZ cluster or on other Linux clusters the user has access to. The resulting images can finally be saved to disk and converted into movies suitable for presentations.

Detailed information on [how to use the visualization cluster and software](#), as well as a [selection of completed visualization projects](#) is available on the RZG web pages. Users may contact the HPC application group for getting support on data conversion and for jointly performing complex visualization projects.

Software News

Renate Dohmen

NAG numerical libraries

Several releases of the NAG libraries comprising numerical and statistical routines are provided for all architectures and compilers supported by the RZG. Users are encouraged to switch to the most recent versions Mark22 or Mark23, respectively, which are FORTRAN 90 based (re-

placing the FORTRAN 77 legacy of the older releases).

Starting 1st March a two-month trial license will be available for the SMP version of the NAG library on the Linux clusters, on Solaris systems and on the Power6. Simply by linking the SMP library instead of the sequential one you can gain performance from the shared-memory parallelism of multi-core systems. Please try this library. Any feed-

back is welcome. During the trial period NAG technical support is open to all users for any questions or problems. All NAG libraries are made available via environment modules. By using the generic environment variables provided by the modules in your Makefiles one can easily switch between different releases and also between the sequential and SMP versions of the NAG libraries.

New webpage 'Software news'

Earlier this year, a [new webpage](http://www.rzg.mpg.de) (see www.rzg.mpg.de, 'News and Events') was established for the documentation of relevant changes to the RZG software environment. Important updates like, e.g., changes to the default version of specific environment modules are documented there chronologically for all relevant platforms.

EU Project EUDAT

Johannes Reetz

The European FP-7 project EUDAT (www.eudat.eu) aims at the establishment of a sustainable collaborative data infrastructure for Europe. As an important stakeholder, the Max Planck Society is involved in this project with the RZG, the MPI for Psycholinguistics and the MPI for Meteorology. The RZG is member of the EUDAT executive board and leads the work package corresponding to operations of the data infrastructure.

Europe's science and research communities from a wide range of scientific fields are faced with increasingly large amounts of valuable data that stem from new sources such as powerful new sensors and scientific instruments used in analyses, experiments and observations as well as growing volumes of data from simulations and from the digitization of library resources. The Max Planck Society with its researchers is engaged in many large national and international data-oriented scientific projects. They

are facing the challenges of a massive expansion in the volume of data to be handled and preserved as well as in the degree of complexity of that data. EUDAT is a consortium of 25 European partners from 13 countries. It includes national data centers, technology providers, research communities, and funding agencies. The project is co-funded by the European Commission. It started in October 2011 with a runtime of three years.

Apart from the discipline-oriented data infrastructures which are existing or being built, EUDAT will establish a cross-disciplinary infrastructure upon which common data services will be provided by diverse communities. EUDAT follows the vision formulated by the [European High Level Expert Group on Scientific Data](#) and the [Blue Ribbon Task Force on Sustainable Digital Preservation and Access](#).