

Max Planck Computing & Data Facility (MPCDF)*
Gießenbachstraße 2, D-85748 Garching bei München

MPCDF, the new name of the RZG

Following an organizational restructuring of the computing facilities within the Max Planck Society the name of the RZG (Rechenzentrum Garching) was changed. Thus

in short, MPCDF is just the new name for the RZG, and stands for 'Max Planck Computing and Data Facility'.

High-throughput data transfers using Globus Online

John Alan Kennedy

During the past 12 months the MPCDF has gained increasing experience in the usage of Globus Online through our engagement in projects such as Virgo and EUDAT. Globus Online provides a user-friendly way to enable high-performance, reliable and secure file transfers.

Taken at its simplest the Globus Online system is comprised of three components:

1. A web-based user layer
2. A file transfer service
3. Several data servers – known as endpoints

The **web-based user portal** allows users to create and discover data endpoints; as well as to schedule and manage data transfers. This is all achieved in a few mouse clicks.

The **file transfer service** manages all the aspects of data transfer including automatic tuning to maximize bandwidth usage, automatic fault recovery (retries), and notifications to the users regarding transfer completion or problems. This removes a great deal of burden from the end users and allows them to concentrate on their research.

The **data servers (endpoints)** can be easily deployed at computing centers as well as on personal laptops through Globus Connect Servers and/or existing gridFTP services can also be used that way providing access to the source data.

The following figure shows the workflow for a simple data transfer using Globus Online:

1. The user submits a transfer request.

2. Globus Online contacts the service endpoints and moves the data.
3. The user is notified on completion.

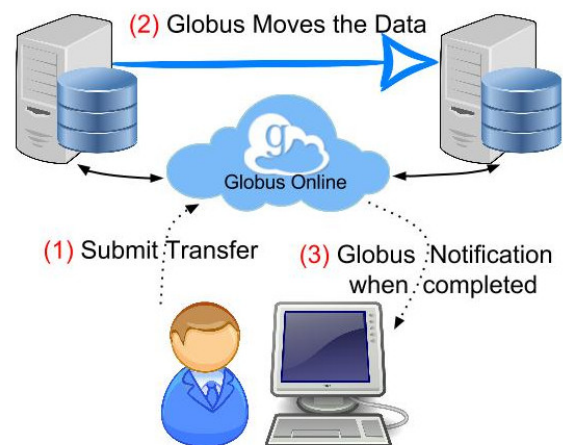


Fig. 1: Globus Online Workflow

Transfer tests have been undertaken between numerous centres both in Europe and the US with rates of 1 Gb/s observed even for large transfers. For example a 10-TB Virgo dataset, containing 23 directories and 11,776 files, was transferred from the MPCDF to the Johns Hopkins University in Baltimore USA with an average transfer rate of 1.2 Gb/s. For the user this was a single logical data transfer (highlight 23 directories and click go).

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

Interactive graphics with x3d(om)

Elena Erastova, Markus Rapp

Scientific journals, like for example the *Astrophysical Journal* enable authors to supplement manuscripts about simulation results with interactive, three-dimensional graphics. In addition to classical visualizations typically shown as images in the article or some pre-rendered movie, simulation data can be represented as 3D objects which can be interactively explored (e. g. by zooming, rotating, flipping objects) using a web browser which implements the X3D standard and object model (no browser plugin required). The approach is implemented as a Java script file embedded in an html page which links to an X3D data file. A basic example of a 3D image taken from the x3dom page together with the html source code is available by clicking on [this link](#). In this example a basic, predefined x3dom 'box' shape was used to define the object `<scene>...<box>/<box>/</scene>`.

Meanwhile many popular visualization tools like Paraview, VMD, Blender, MeshLab and others support the generation of X3D files which allows representing complex visualization results as interactive 3D images. An application with [Myoglobin molecular structure](#) is available under the highlighted hyperlink.

In the following the basic workflow used for preparing interactive graphics for the simulation data of Melson et al. (APJ 801, L24, 2015, see also [H.-Th. Janka: 'Three-dimensional computer simulations support neutrinos as cause of supernova explosions', MPA, August 2015](#)) is illustrated. Once the data have been visualized in ParaView, the File -> Export Scene menu option needs to be used and 'X3D File' to be selected (cf. Figs. 2 and 3).

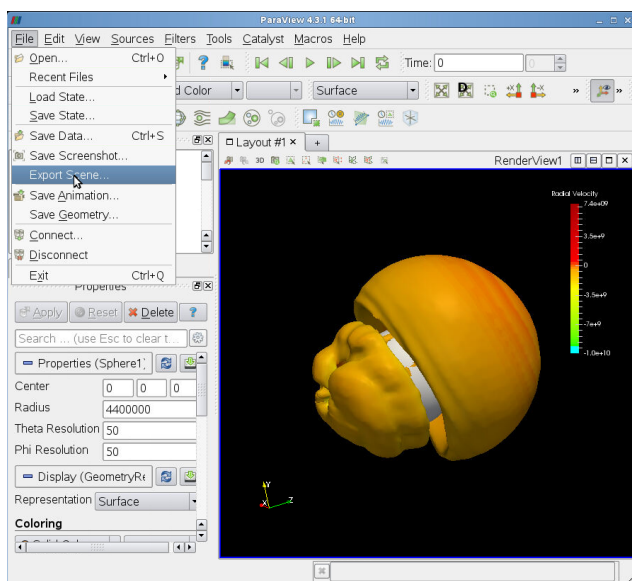


Fig. 2: Exporting a model as X3D file in ParaView

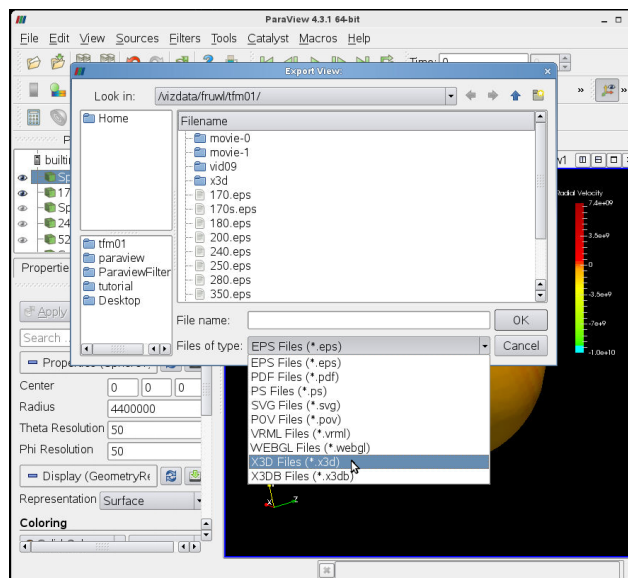


Fig. 3: ParaView Export Scene dialog

An [example webpage](#) shows the object which can be rotated and zoomed with the standard mouse controls, together with corresponding html source code. The html source code of this page is shown below the 3D image. Here, `<scene><inline url='170ms.x3d'>/<inline>/</scene>` is used to import the object(s) from the .x3d file.

Note that we have originally employed the VisIt visualization tool which does not provide the X3D export feature yet. Until this feature becomes available in VisIt (planned for the next release scheduled for autumn 2015) the following workaround can be used: The visualization with VisIt is exported to a VTK file format (using the VisIt dialogue File -> Set Save Options). The VTK file can be imported into ParaView and finally be re-exported as X3D file.

Quick links

- [x3dom project](#)
- [X3D files on wikipedia](#)
- [web3d consortium](#)
- [ParaView](#)
- [VisIt](#)

Projects

John Alan Kennedy, Hermann Lederer

Virgo Data Centre

Virgo is an international consortium of computational cosmologists who perform Cosmological Supercomputer Simulations on leading HPC resources (see <http://www.virgo.dur.ac.uk/>). The Virgo Consortium has a core membership of about a dozen science groups in the UK, Germany, the Netherlands, Canada, the USA and Japan. The group of Simon White from the MPA plays a key role both in performing simulations and through the realization of the Virgo Data Centre at the MPCDF.

As the cosmological supercomputer simulations increase in both size and number the need for a dedicated data centre has arisen. This data centre will not only house the data, but provide access and value-added services to open the data to the cosmological community as a whole. This is especially relevant after the ground-breaking Millennium run, which was published in 2005. It is now clear that this data is of great value to a wide community of cosmologists undertaking galaxy formation studies and also for observational cosmology.

In order to manage the increasing data volumes as well as to enable data evaluation and knowledge extraction from existing and future simulation data, the Virgo Data Centre has been established at the MPCDF. This initiative is driven by the collaboration of the Virgo members with domain specific knowledge, data and computing experts at the MPCDF and other Virgo partners.

The data centre consists of a 2 PB high-performance storage system, a database containing numerous post-processing results from cosmological models, and compute clusters attached for post-processing. Moreover, Globus Online services have been deployed to enable reliable, high-speed data transfer to and from the Virgo Data Centre.

EU Horizon 2020 Projects

First projects from the Horizon 2020 programme have started or are starting in 2015. The MPCDF is engaged in some of them together with respective Max Planck Institutes:

EUDAT2020 – a continuation of the FP7 EUDAT projects – provides research data services in a Pan-European infrastructure, through an engagement of 35 partners from 15 European countries (see <http://www.eudat.eu>).

RDA Europe is the European plug-in to the Research Data Alliance (RDA), a world-wide activity building the social and technical bridges to enable global and open sharing of data. Currently Australia, Europe and the US are represented (see <http://www.rd-alliance.org>).

PRACE-4IP is the continuation of the FP7 PRACE-IP projects to co-finance the PRACE activities for supercomputing in Europe. Germany is represented by the Gauss Centre for Supercomputing (GCS) with the FZJ (Jülich), the HLRS (Stuttgart) and the LRZ (Garching). As in the past the MPCDF is engaged as a Third Party to GCS (see <http://www.prace-ri.eu>).

BioExcel is a new Centre of Excellence for Computing Applications in the field of Life Sciences to start in November. It is coordinated by the KTH (Sweden). From the MPG, the MPI for Biophysical Chemistry and the MPCDF are engaged.

NoMaD is a new Centre of Excellence for Computing Applications in the field of Materials Science to start in November. It is coordinated by the FHI (see <http://nomad-lab.eu>).

100-Gbps wide area network connection

Hermann Lederer

A wide area network connection with a bandwidth of 100 Gigabits per second has been realized between the Max Planck Institutes in Martinsried and the MPCDF in Garching. The need arises especially from the generation of very large data sets through new high-resolution technologies in the areas of electron microscopy and mass spectrometry used in the MPIs for Biochemistry and for Neurobiology. The implementation has been realized by T Systems and DACOSO, the winners of a European procurement, in close collaboration with the Leibniz Rechen-

zentrum (LRZ). The LRZ is operating the Munich Science Network (MWN), where existing fibres have been used as backbone. This backbone has been equipped with 100-Gbps-enabled DWDM (Dense Wavelength Division Multiplexing) systems, and the MPIs in Martinsried have been connected to the south-west end of the MWN with a 100-Gbps transponder, while the MPCDF has been connected to the northern end of the MWN at the LRZ in Garching also with a 100-Gbps transponder.

News and Announcements

Andreas Schmidt, Andreas Schott, Ingeborg Weidl

New gate machines

For many years, actually since July 1994, the RZG provided under the names *ibmr*, *sp*, and later *gate*, IBM-AIX systems for general-purpose usage and gateway login functionality. After 20 years it is time for a change.

Nowadays almost every user has access to either his personal computer or to other computing cluster resources. So the need for a general-purpose computer has been superseded by the requirement of a secure gateway computer to access the computing facilities at the MPCDF. Since the maintenance of the AIX systems is no longer possible, it was quite obvious to switch to today's standard of Linux systems.

In order to provide better stability and also redundancy, the old gate system has been substituted by [two new computers, gateafs and gatezero](#). The first one, *gateafs*, will take over the old name *gate* and provides the known home directories in the afs including e-mail access. The other one, *gatezero*, is very minimalistic aiming to reduce the number of reboots by offering only a small local home directory and no access to the afs. This machine will therefore be acting just as a gateway or tunnel.

On both machines there are no resources available for general-purpose computing, like compiling and running programs. For those needs please use the respective [dedicated computing facilities](#). Furthermore, for security reasons the possible allowed key exchange and encryption methods when connecting to these Linux systems have been reduced to cover only those methods currently considered safe enough.

Switching from AIX to Linux actually also implies that several old legacy programs known from the AIX systems may no longer be available or be substituted by new ones sometimes with different names. A prominent candidate is the former 'pine' for reading e-mails, which is now available as 'alpine' on *gateafs* only.

Interactive usage of hydra login nodes

The *hydra* login nodes [hydra.rzg.mpg.de](#) (*hydra01*, *hydra02*) are only for login, editing and managing your data. Any application should be run only on the pool of so-called 'interactive' nodes [hydra-i.rzg.mpg.de](#) (*hydra03* – *hydra08*). Graphical applications like IDL should be run on the visualization cluster [viz.rzg.mpg.de](#).

Test runs of your codes should not run for more than 2

or 3 hours and should not use more than 8 cores. Longer runs (even on a single core) can be submitted to the LoadLeveler batch system.

And please, always keep an eye on the memory consumption of the application that you start interactively on [hydra-i.rzg.mpg.de](#) (e. g. with the command 'top').

What to do with expired and forgotten passwords

All owners of an MPCDF account have to renew the password of their account once a year as it will expire otherwise. In case you miss the reminding e-mails and your password expires, you are not able to log onto any MPCDF machine.

However, as long as you know your old, expired password you are able to reactivate your account using the [MyRZG Recovery Portal](#). To use this special secure section of the MPCDF webpages you don't need to log in, but only have to tell the recovery portal your account name. You are then able to reactivate your account using the 'Change Expired Password' link from the menu. The form is similar to the [MyRZG Change Password form](#).

In the future we will also enable a self service to reset a forgotten password on your own. In order to be able to use this functionality, we need additional information you have to enter in the [MyRZG Change Security/Recovery Options form](#).

New user regulations

As of May 1st, the [user regulations](#) or in German the so-called [Nutzungsordnung](#) have taken effect. There are no changes concerning the rights and duties of the users. The main change concerns the new name – now MPCDF instead of RZG – and the adjustments formally connected to the changes in the organizational structure, MPCDF now being an autonomous part of the MPG. New users have to accept the new regulations in order to get an account, while there is no need for an action for existing users.

There is one alteration, which was in fact practically implemented already some time ago, concerning the deletion of old e-mails: now there is no longer a cleanup for e-mails lying in the inbox on the mail server of the MPCDF for more than half a year.

Tips and Tricks

Klaus Reuter, Andreas Schott

Python setuptools

In recent years, the Python programming language has become increasingly popular in the field of scientific computing. On MPCDF systems, a selection of important Python modules is provided, in particular an optimized build of the SciPy stack. The language versions 2.7 and 3.3 are supported.

For Python, a huge number of packages serving virtually any purpose is available on the Internet (see <https://pypi.python.org>). MPCDF users request many of them regularly, however it is often not useful to provide the particular package site-wide. To enable users to easily download, build, and install Python packages, the Python setuptools package was made available recently (modules `python27/setuptools` and `python33/setuptools`). Among other functionalities, setuptools provides the command `'easy_install'` which can be used to automatically install a certain package together with its dependencies locally as follows:

```
easy_install --user PACKAGE_NAME
```

The Python package is then located at `~/local` and is automatically found by the interpreter.

How to efficiently make use of only one gateway session

Often one has the need for multiple screens with only one putty/ssh session running. In addition, the connection may be unstable and one would lose all the work when getting disconnected. Furthermore, if connected to a machine behind the gateway, this connection would also be lost. The solution to overcome such situations is the command `screen`. With `screen` one can open multiple screens in just one session, switch between them and resume them. They survive even if the connection to the gateway gets lost. One just reconnects and resumes the respective sessions.

Here comes, how this can be achieved. Assume one wants a connection to *hydra* from anywhere in the world, thus one first needs to connect to one of the gateway machines with putty or any other ssh command. After being connected to e.g. *gatezoro* one does not simply invoke `ssh hydra` but:

```
screen -S hydra ssh hydra
```

The option `'-S'` is not required, but names the screen session for later identification. The whole will look as if one would be simply connected to *hydra* via *gatezoro*, but in fact this connection to *hydra* will survive disconnects from *gatezoro*.

Now assume the session to *gatezoro* is being dropped. After reconnecting to *gatezoro* the command `screen -ls` will display existing screen sessions, e.g. here something like:

```
1962.hydra (Attached)
```

To resume this connection to *hydra*, returning to where one has been just before the session to *gatezoro* got lost, one has to give the command:

```
screen -d -r 1962
```

The option `'-d'` is required to drop the existing line from the old session, which seems still to be attached, while the option `'-r'` will actually resume then this session. May be one needs another session to *hydra* for editing, then pressing `'Ctrl-a'` followed by `'c'` will create a new session, leaving the old one attached. Then

```
screen -S hydra-editor ssh hydra
```

will open a new session to *hydra*. With `'Ctrl-a d'` one can detach a session from the display, while the connection behind that session will stay, so one can resume it later.

One can switch between all attached sessions in forward direction with `'Ctrl-a n'` (next) and in backward direction with `'Ctrl-a p'` (previous). Having multiple screen sessions one can get an overview with the command `'Ctrl-a w'`. One can directly switch to one of these sessions with `'Ctrl-a'` followed by the number of the session. One can name the session with `'Ctrl-a A'`.

There is also the possibility to copy data from one screen to another, or to save the contents of the current screen (`'Ctrl-a h'` for hardcopy). For more information about screen consult the manpage, which explains also how to configure defaults in `$HOME/.screenrc`.

Events

Hermann Lederer, Markus Rampp

Software training

This year's software training for Nvidia GPU programming (related to the Max Planck supercomputer which is equipped with 676 Nvidia K20X GPUs) took place at the MPCDF in May 2015. The one-day event attracted more than 30 Max Planck scientists from 12 different institutes, not only from the Munich area. The program comprised an overview of the relevant GPU computing and visualization resources and activities at the MPCDF and lectures by Nvidia on their GPU technology roadmap, on OpenACC programming, and on in-situ visualization. The training material is available at [the workshop's website](#).

International HPC Summer School

This year's International HPC Summer School on Challenges in Computational Sciences took place in Toronto, Canada from June 21st to 26th, hosted by Compute-Canada. 80 participants (PhD students and post-docs) from institutions in Canada, Europe, Japan and the US had been selected from more than 550 applications. Besides an extraordinary programme covering many disciplines and programming techniques, intensive mentoring and networking made this again a world-wide unique and highly successful event. The 2016 HPC Summer School is expected to take place in June/July in Ljubljana, Slovenia (see <http://www.ihpcss.org>).