



MAX-PLANCK-GESELLSCHAFT

MPCDF in a Nutshell

Raphael Ritz

Max Planck Computing and Data Facility (MPCDF)

Garching bei München

Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V.
- an independent research organisation

High Performance Computing

Hosted Compute and Data Systems

Mass Data Archive Systems

Application Support

Data Projects

Computational Support

International Projects

MPG contracted an Intel SkyLake processor based system

- Delivery in Nov + Dec 2017
- Installation in Jan 2018
- More than 2500 compute nodes with 40 cores each => 100.000 cores
- Half of the nodes are equipped with 96 GB RAM, the other half with 192 GB RAM and eight nodes contain 768 GB RAM
- The batch system SLURM as on DRACO
- Peak performance 8 PF
(The HYDRA system without GPUs had roughly 1.6 Peak)



Compute and Data Systems hosted at MPCDF (2017-2018)



Fritz-Haber-Institut, Berlin

Administrative Headquarters of the Max Planck Society, Munich

Max Planck Institute for Astrophysics, Garching

Max Planck Institute for Astronomy, Heidelberg

Max Planck Institute for Biochemistry, Martinsried

Max Planck Institute of Biophysics

Max Planck Institute for Brain Research, Frankfurt

Max Planck Institute for Chemistry, Mainz

Max Planck Institute for Chemical Physics of Solids, Dresden

Max Planck Institute for Evolutionary Anthropology, Leipzig

Max Planck Institute for Extraterrestrial Physics, Garching

Max Planck Institute for Iron Research, Düsseldorf

Max Planck Institute of Neurobiology, Martinsried

Max Planck Institute for Plant Breeding Research, Cologne

Max Planck Institute for Plasma Physics, Garching

Max Planck Institute for Physics, Munich

Max Planck Institute for the Physics of Complex Systems, Dresden

Max Planck Institute for Polymer Research, Mainz

Max Planck Institute of Psychiatry

Max Planck Institute of Quantum Optics, Garching

Max Planck Institute for Ornithology, Radolfzell

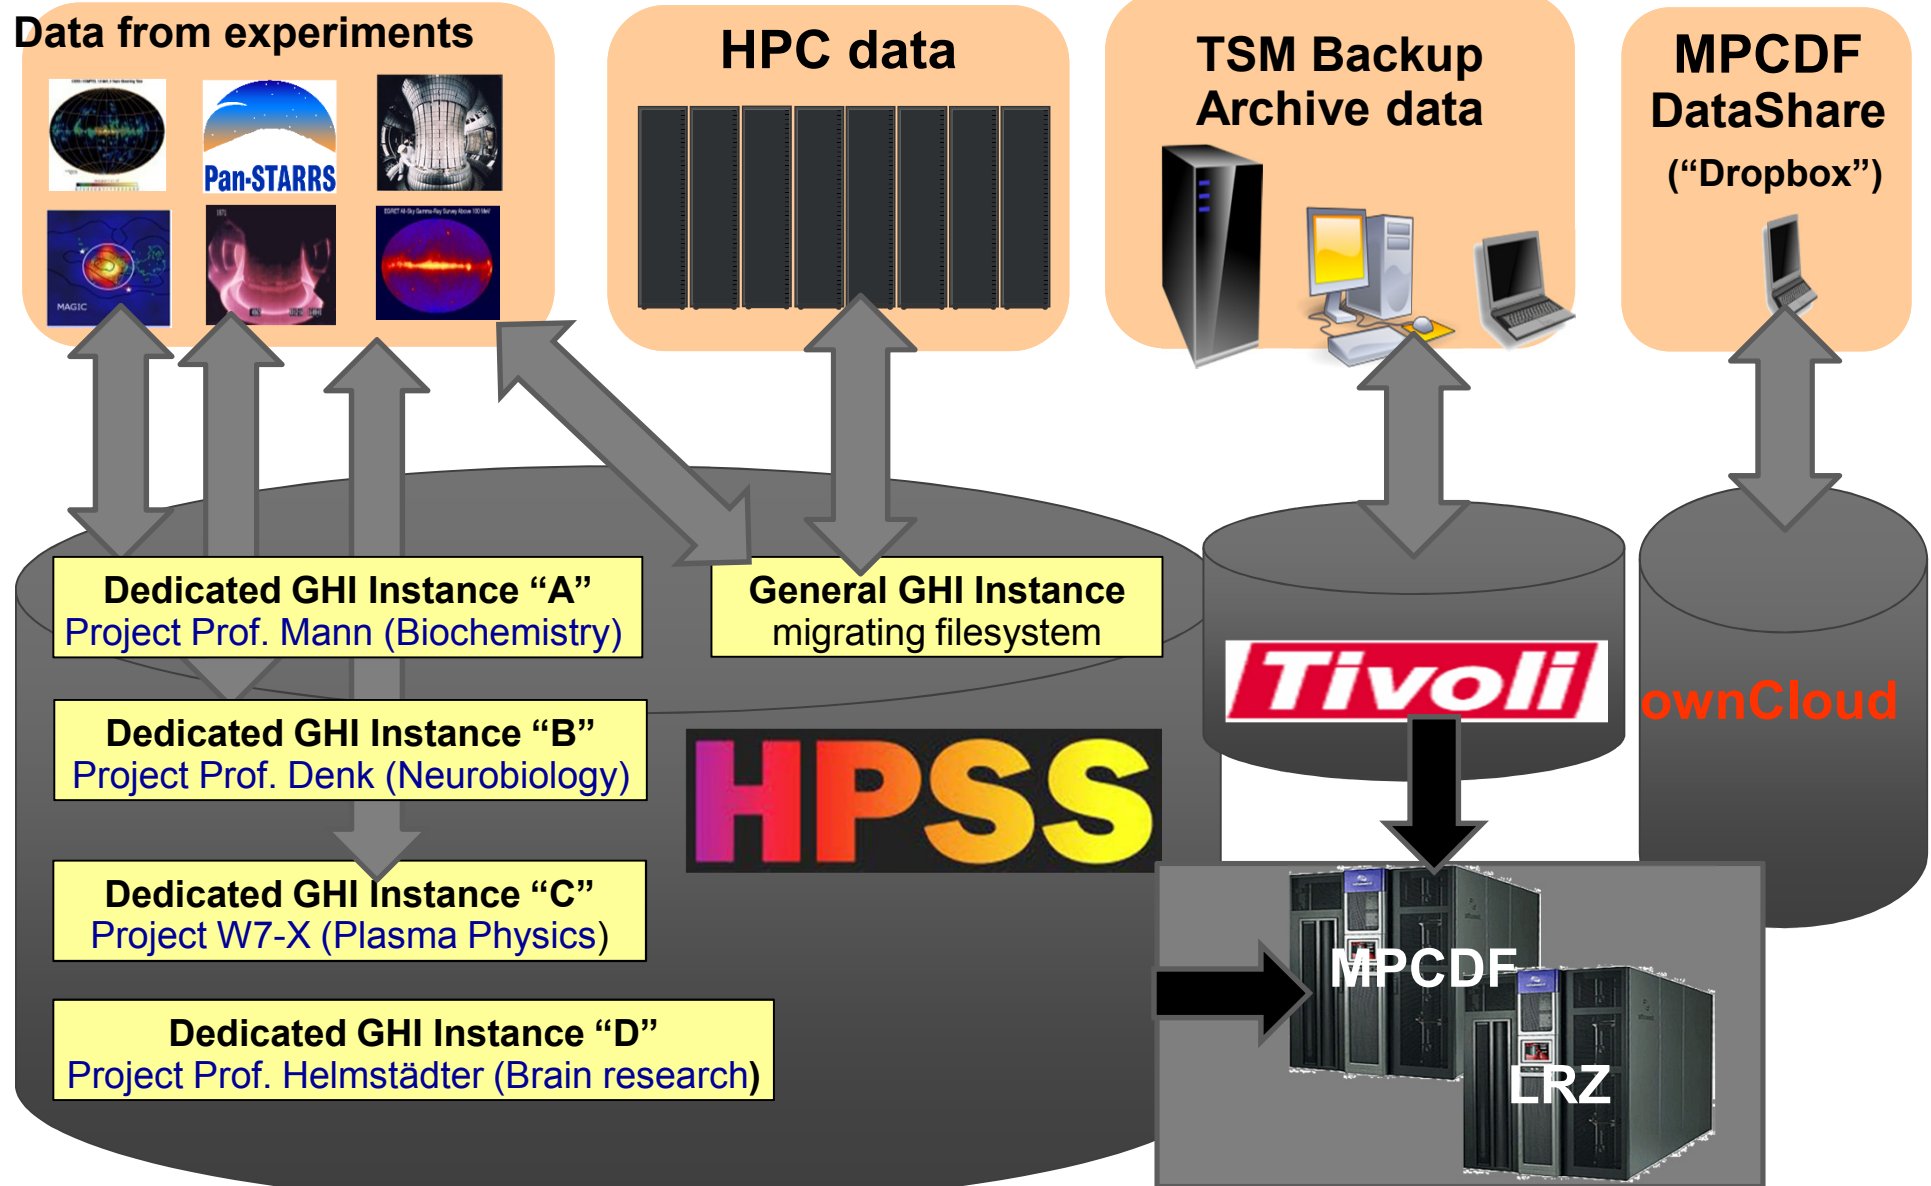
Max Planck Institute for Science of Light, Erlangen

Max Planck Institute for Solid State Physics, Stuttgart

Max Planck Institute for Terrestrial Microbiology, Marburg

Fritz-Haber-Institut (Berlin), Max Planck Institute for the Physics of Complex Systems (Dresden),

Max Planck Institute for the Structure and Dynamics of Matter (Hamburg)





The top Scientific Data Center (concerning size)



ECMWF European Center of Medium Weather Forecast

Site	Storage (Disk + Tape) in Bytes	PBs	Disk/Tape Percentage ¹	Files	Millions of Files	Non-file Namespace Percentage ²	Statistics Updated	HPSS in production since
ECMWF	335,878,755,008,621,000	335.88	0.35%	317,219,062	317.2	9.62%	10/03/2017	Nov 2002
UKMO	199,559,198,222,001,000						09/19/2017	Jun 2009
NOAA-RD	120,895,045,847,033,000						01/04/2017	Nov 2002
BNL							09/29/2017	Sep 1998
LBNL-User	113,013,671,922,991,000						09/25/2017	Mar 1998
Meteo-France	92,846,654,883,109,800	92.85	9.78%	355,734,321	355.7	7.74%	09/21/2017	Feb 2015
CEA TERA						14.72%	09/25/2017	Jul 1999
NCAR						2.48%	10/01/2017	Apr 2011
MPCDF						0.50%	10/04/2017	Nov 2011
ORNL	74,662,859,118,843,800	74.66	14.76%	77,965,233	78.0	3.07%	09/18/2017	Nov 1997
LANL-Secure	74,662,859,118,843,800						09/18/2017	Nov 1997
LLNL-Secure	69,404,916,079,982,000						09/18/2017	Jul 1998
DKRZ	63,688,513,389,795,200						09/18/2017	Jul 2009
IN2P3	52,557,509,363,302,200						09/18/2017	Jul 1999
ANL	45,802,836,088,011,500	47.82	1.18%	379,175,632	379.2	3.67%	09/26/2017	Sep 2008
NCSA	45,802,836,088,011,500	45.60	2.74%	272,872,629	272.9	1.85%	09/30/2017	Nov 2012
DWD	45,168,744,871,000,000		3.76%	49,711,637	49.7	0.40%	09/18/2017	Jul 2011
LLNL-Open	40,768,116,527,723,900		1.46%	1,026,247,376	1,026.3	4.42%	09/28/2017	Jan 1997
HLRS	38,685,500,000,000,000						09/27/2017	Sep 2008
SLAC	38,190,300,000,000,000						09/15/2017	Jul 1999
CEA TGCC	34,450,000,000,000,000						09/25/2017	Nov 2010
IU	35,200,000,000,000,000	35.21	2.79%	247,275,615	247.3	13.06%	09/19/2017	Jun 1999
JAXA	24,981,791,000,000,000		0.92%	62,384,671	62.4	4.10%	10/03/2017	Jan 2009
LBNL-Backup	23,863,250,066,508,500		1.14%	19,782,567	19.8	4.24%	09/25/2017	Jun 1997
Purdue	20,028,143,404,906,400			83,405,165	83.4	4.44%	09/18/2017	Oct 2011
SSC	16,379,089,900,000,000						09/28/2017	Jun 2017
KEK	12,539,220,000,000,000						09/18/2017	Jan 2001
LaRC	11,744,513,700,000,000						09/21/2017	Dec 1998
PNNL	11,175,393,810,966,200	11.18	10.18%	108,991,538	109.0	12.85%	09/18/2017	Apr 2010
SciNet	10,346,797,825,796,000	10.35	0.39%	36,544,109	36.5	5.77%	09/18/2017	Jun 2011
NOAA-C Boulder	9,875,856,100,817,410	9.88	5.23%	246,117,486	246.1	0.00%	09/18/2017	Oct 2014
NOAA-C Asheville	9,872,930,478,566,420	9.87	5.19%	246,161,880	246.2	0.00%	09/18/2017	Oct 2014
CNES	9,740,582,118,968,410	9.74	19.88%	3,771,860	3.8	1.73%	09/28/2017	Apr 2015
LANL-Open	6,246,438,965,365,600	6.25	0.00%	93,442,800	93.4	4.32%	10/03/2017	Oct 1997
KIT	5,338,102,287,178,880	5.34	2.49%	29,996,962	30.0	6.71%	10/04/2017	Jul 2015
NCEI	3,720,165,000,000,000	3.72	2.00%	90,773,427	90.8	1.15%	09/26/2017	Jan 1998
SNL-Open	2,854,937,058,000,000						09/07/2017	Jan 1970
SNL-Secure	2,146,211,784,000,000						09/07/2017	Jan 1970
Total	2,067,933,330,690,000,000						NA	"-"

BNL Brookhaven National Laboratory

MPCDF Max Planck Computing and Data Facility

ORNL Oak Ridge National Lab

DKRZ- German Climate Center

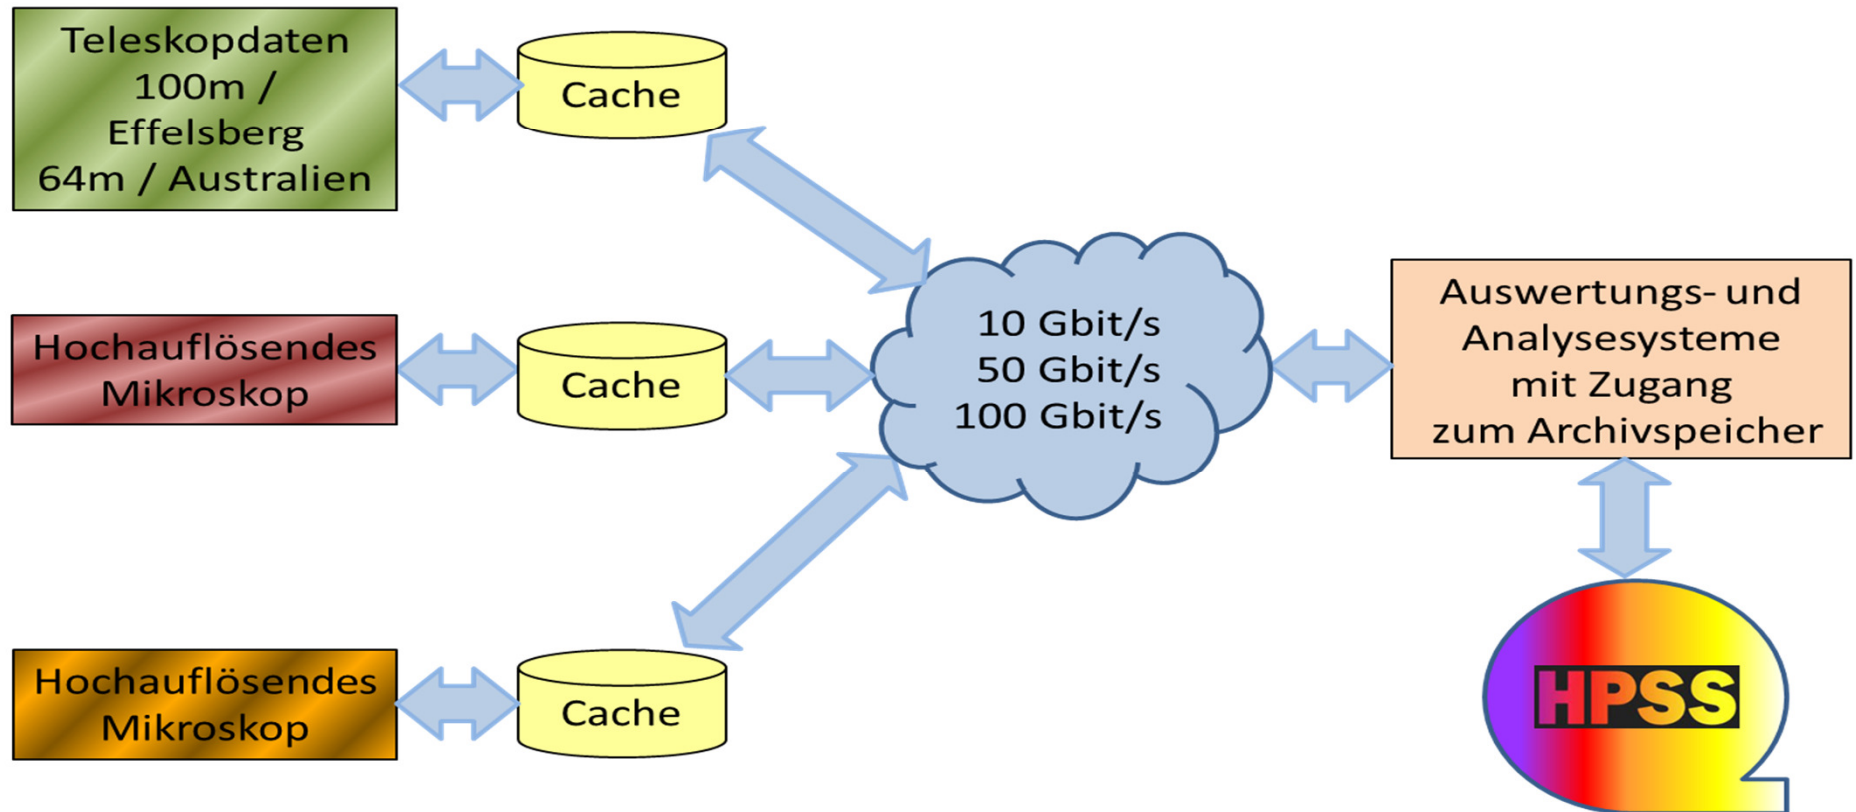
SLAC National Accelerator Laboratory

KIT Karlsruhe Institute of Technology

MPI für Radioastronomie Bonn
MPI für Hirnforschung Frankfurt
MPI für Neurobiologie München

VPN via DFN oder
Münchner Wissenschaftsnetz

MPCDF Garching





MAX PLANCK COMPUTING & DATA FACILITY

RECHENZENTRUM GARCHING DER MAX-PLANCK-GESELLSCHAFT

Contact | Login | External Links

Search Site



HOME

SERVICES

ABOUT MPCDF

USERSPACE

HELPDESK

Home » Userspace » Forms, etc » Registration

NEW USERS

FORMS, ETC

■ REGISTRATION

REGULATIONS

NUTZUNGSORDNUNG

TERMS OF USE FOR
GUEST USERS

PC FORMS

COURSES

MYMPCDF

Registration at the Max Planck Computing and Data Facility

If you are working with one of the Max-Planck-Institutes below, you can register here for use of the machines at the Max Planck Computing and Data Facility.

Participants of projects (e.g. EU-H2020, NOMAD, IFERC, ITM), please apply [here](#).

Your application must be approved by one of your institute's representatives (directors) whose name has been added to our database. Your account will be created only after that approval. If you have any concerns, please feel free to call our office at (089) 3299-2175.

<http://www.mpcdf.mpg.de/userspace/forms/onlineregistrationform>




https://subscriptions.rzg.mpg.de/manage?service_name=collaboration_management&instance_na

MPCDF MAX PLANCK COMPUTING & DATA FACILITY
RECHENZENTRUM GARCHING DER MAX-PLANCK-GESELLSCHAFT

Hello **Raphael Ritz**,

Home
Services
Invite
Help
Login Details
Logout

Welcome to the MPCDF subscription management service

-  **1. Subscribe yourself to the MPCDF services**
-  **2. Invite internal or external guest users for collaboration**
-  **3. Grant access to MPCDF services to invited guest users**

On this platform you can subscribe to MPCDF services (currently, these are [GitLab](#) and [DataShare](#)). In addition, you can invite internal and external colleagues to participate in the services. If an invited colleague doesn't have an MPCDF account, a **guest user account** is created for her or him. Guest users have per default limited permissions. Additional permissions can be granted by the inviter.

Quick link: [Click here for registering at the MPCDF services.](#)

<https://subscriptions.rzg.mpg.de>



https://subscriptions.rzg.mpg.de/manage/invite_new_collaborator

MPCDF MAX PLANCK COMPUTING & DATA FACILITY
RECHENZENTRUM GARCHING DER MAX-PLANCK-GESELLSCHAFT

Hello **Raphael Ritz,**

Invite an External Guest User

You can invite guest users who do not have an MPCDF account to use the MPCDF services. For these colleagues, *guest accounts* are generated. In contrast to a full qualified MPCDF user account, a guest account has access only restricted access to the MPCDF services. For example, a guest user of the DataShare service has no own storage: if the inviter grants access to the guest, he can upload or edit existing data in the inviter's account.

First name:

Last name:

Email-address:

<https://subscriptions.rzg.mpg.de>



- DataShare:
 - ownCloud-based sync-and-share solution (think Dropbox)
 - Command line client: <https://pypi.python.org/pypi/pocli>
 - <https://datashare.mpcdf.mpg.de>
- GitLab:
 - Collaborative development
- DataHub
 - Staging area for massive data transfers in and out of MPCDF
 - GlobusOnline endpoint available
 - <https://www.mpcdf.mpg.de/services/data/data-transfer/datahub>




- Hardware provisioning at MPCDF
- Implementation and porting support – also available at your site!
- Experience with, e.g., CuPy <https://cupy.chainer.org/>
 - CuPy is an open-source matrix library accelerated with NVIDIA CUDA. It also uses CUDA-related libraries including cuBLAS, cuDNN, cuRand, cuSolver and NCCL to make full use of the GPU architecture.
- GeNN
 - GeNN is a GPU-enhanced Neuronal Network simulation environment based on code generation for Nvidia CUDA.
 - <http://genn-team.github.io/genn/>



- No major project so far but ...
- ... we are building up in-house expertise at the moment
- Some software in place, e.g. Tensorflow
- Very interested in pursuing a concrete project
- Contact at MPCDF: Andi Marek (andreas.marek@mpcdf.mpg.de)



- Cloud computing via OpenStack (think AWS EC2)



- Tutorial: <https://developer.openstack.org/firstapp-shade/>

<https://en.wikipedia.org/wiki/OpenStack>

- **Keeper** – *Self-service archiving* mit MPDL
- **Castellum** – Probandendatenbank, mit MPI Bildungsforschung, Psychiatrie, Kognitionsforschung, GV, ...
- **DIMAG** – digitales Magazin, mit MPG Archiv

- **NOMAD** Center of Excellence
- **EUDAT**
- European Open Science Cloud (**EOSC**)
 - EOSC Pilot
 - EOSC HUB
- Research Data Alliance (**RDA**)
 - RDA Europe (1–4)
 - RDA Deutschland

THE RESEARCH DATA ALLIANCE

www.rd-alliance.org

*building the social and technical bridges
that enable open sharing of data*

18 FLAGSHIP OUTPUTS

of which 4 ICT
Technical
Specifications

75 ADOPTION CASES

across multiple
disciplines,
organisations &
countries

91 GROUPS WORKING ON GLOBAL DATA INTEROPERABILITY CHALLENGES

of which 33 WORKING GROUPS
& 58 INTEREST GROUPS

6,889 INDIVIDUAL MEMBERS FROM 137 COUNTRIES

67% Academia & Research
15% Public Administration
11% Enterprise & Industry

43 ORGANISATIONAL MEMBERS & 8 AFFILIATE MEMBERS



Vision

Researchers and innovators openly share data across technologies, disciplines, and countries to address the grand challenges of society.

Mission

RDA builds the **social and technical bridges** that **enable open sharing** of data.

WWW.RD-ALLIANCE.ORG

@RESDATALL



CC BY-SA 4.0

Datenapplikationen (Raphael.Ritz@mpcdf.mpg.de)

- Gemeinsame Entwicklung von Datenprojekten
- Circa 15 Projekte in den letzten Jahren: Unterstützung von Großexperimenten (W7X, Asdex Upgrade), internationalen Projekten (CERN, MAGIC, Belle) und institutsspezifischen Projekten
- ownCloud: Unterstützung dediziert für Institute mit Anbindung an das Institut User Management
- GitLab: aktuell fast 1500 „repositories“

Compute- und Dateninfrastrukturprojekte (Christian.Guggenberger@mpcdf.mpg.de)

- Compute-, Plattensysteme, Filesysteme (GPFS, BeeGFS), Archivsysteme
Beratung und Unterstützung zu Filesystemen
(z.B. MPI für Biophysik, MPI für Pflanzenzüchtungsforschung)
- Ausschreibungen
- BAR-Anträge
- Cloud Services
- Virtuelle Maschinenumgebung

Optimierung und Entwicklung von HPC Applikationen (Markus.Rampp@mpcdf.mpg.de)

- Höhere Skalierung
- GPU-Unterstützung
- Entwicklung von neuen Algorithmen
- Visualisierungs-Projekte
- Applikations-Stack für dedizierte Mid-range Systeme und den MPG-Hochleistungsrechner

Questions?

Comments?

Feedback?

Ideas?