

# Bits & Bytes

No.212, April 2023

Computer Bulletin of the Max Planck Computing and Data Facility (MPCDF)\*  
<https://docs.mpcdf.mpg.de/bnb>

## High-performance Computing

### New Supercomputer of the MPG - *Cobra* successor

In the course of the last year MPCDF, in collaboration with MPG-GV and in close consultation with committees of the MPG-BAR and the MPCDF board, conducted a procurement for the next-generation HPC system of the Max Planck Society which will replace the current *Cobra* machine by the end of this year. The corresponding proposal to the president of the MPG was supported by 58 departments and 19 research groups of 37 Max Planck Institutes from all three sections of the MPG. On January 16th, 2023 a contract was signed with the winner of the European tender, Atos<sup>1</sup>.

The new machine consists of a CPU-only and a GPU-accelerated system, both based on AMD processors of the latest generations, an Nvidia/Mellanox InfiniBand (NDR) network, as well as disk (ca. 20 PiB) and NVMe (ca. 500 TiB) storage with IBM SpectrumScale file system technology. The CPU-only system consists of 768 compute nodes, each with two AMD EPYC “Genoa” processors, providing 128 Zen4 cores per node and 512 GiB (609 nodes), 768 GiB (90 nodes), 1024 GiB (66 nodes), or 2048 GiB (3 nodes) of DDR5 memory per node, respectively. This system is scheduled for installation in the second half of this year. The GPU-accelerated system comprises 192 compute nodes, each with two of the new AMD Instinct MI300A “APU” processors with CPU cores and GPU compute units integrated on the same chip and coherently sharing the same high-bandwidth memory (128 GiB HBM3 per APU). This system is scheduled for installation during the first half of 2024.

In order to help preparing applications for the new system, MPCDF will provide further technical details to its users as they become available and will organize dedi-

cated trainings and workshops in due course, in particular covering the new APU technology. For more details and support please contact Markus Rampp<sup>2</sup>.

The procurement of the *Raven* successor system is scheduled for 2025/2026.

*Erwin Laure & Markus Rampp*

### Documentation of HPC hardware characteristics

Recently, the *Raven* user guide<sup>3</sup> was extended with more detailed information on the *Raven* hardware characteristics<sup>4</sup>, including a schematic of a GPU node showing its building blocks and the theoretical bandwidths between them. Moreover, roofline plots for the CPUs and GPUs and other performance numbers from measurements based on microbenchmarks are presented which may be helpful to the advanced user when analyzing and optimizing HPC code.

*Klaus Reuter*

### CMake Recipes Repository

Have you ever asked yourself questions like “How do I link against third-party libraries correctly with CMake?”, “How do I create custom CMake targets to build my documentation?” or “How do I get the current Git hash in my source code to print it to the log files?”. We are here to help! The MPCDF maintains a repository with a growing number of CMake recipes, i.e. ready-to-use CMake code snippets for various tasks. You can check them out in our GitLab<sup>5</sup>. And, if there is something missing or you need further clarifications, feel free to open an issue in our helpdesk.

*Sebastian Eibl*

\*Editors: Dr. Renate Dohmen & Dr. Markus Rampp, MPCDF

<sup>1</sup>[https://atos.net/en/2023/press-release\\_2023\\_02\\_09/atos-to-build-max-planck-societys-new-bullsequana-xh3000-based-supercomputer](https://atos.net/en/2023/press-release_2023_02_09/atos-to-build-max-planck-societys-new-bullsequana-xh3000-based-supercomputer)

<sup>2</sup><mailto:markus.rampp@mpcdf.mpg.de>

<sup>3</sup><https://docs.mpcdf.mpg.de/doc/computing/raven-user-guide.html>

<sup>4</sup><https://docs.mpcdf.mpg.de/doc/computing/raven-details.html>

<sup>5</sup><https://gitlab.mpcdf.mpg.de/mpcdf/training/cmake-recipes>

## MPCDF HPC Cloud

### Introduction

The MPCDF HPC Cloud provides on-demand computing and storage resources to research projects of Max Planck Institutes. Implementing the infrastructure-as-a-service model, the HPC cloud offers servers, networking features, and storage resources through a high-level API, CLI, or browser-based GUI. The moniker *HPC Cloud* highlights the co-location of the cloud and HPC systems at MPCDF. While the HPC systems provide massive computing power, they are limited to applications that can be run within a batch system. The cloud complements the HPC workflows by enabling projects to define flexible computing solutions, including the deployment of workflow engines, databases and long-running jobs (exceeding the 24hr limit of the HPC batch system). Interaction with the HPC systems can be achieved, most notably, using the Nexus storage systems: the Nexus-POSIX file system which is mounted on *Raven* and can be mounted on the HPC Cloud on demand, as well as Nexus-S3, the globally accessible object store. The HPC Cloud and Nexus storage systems are implemented with OpenStack, Ceph, and IBM Spectrum Scale.

### Hardware Resources



Figure 1: HPC Cloud Hardware at MPCDF

The initial deployment of the HPC Cloud took place in 2021 in collaboration with the Fritz Haber Institute (FHI), and the Max Planck Institutes (MPI) for Human Cognitive and Brain Sciences, and for Iron Research, comprising 60 Intel Icelake-based compute nodes for a total 4320 cores, 44 TB of main memory, 12 Nvidia A30 GPUs, and 80 TB of local SSDs. An extension in collaboration with FHI, MPI of Animal Behavior, and the MPDL, is currently being deployed and will provide additional resources including 2688 cores, 96 TB of main memory, 60 Nvidia A100 GPUs, and 130 TB of NVMe-based local storage. All compute nodes have redundant 25 Gb ether-

net links to a 100 Gb backbone, which itself has 100 Gb uplinks to the MPCDF core network as well as a direct connection to *Raven*. This second installment provides significant new resources and also allows for better coverage of use cases such as machine learning and memory or I/O intensive applications.

### Project Support

The HPC Cloud is open to Max Planck Institutes and already hosts projects from numerous institutes including Biblioteca Hertziana, MPI for Biology of Ageing, MPI of Neurobiology and MPI for Chemistry. In collaboration with the MPCDF Cloud team, research teams can design, deploy and manage solutions within the HPC Cloud, taking advantage of the proximity to the *Raven* HPC system when needed. In addition MPCDF has developed a set of recipes covering common requests, for example deployment of Kubernetes clusters “on top” of cloud-based resources.

Starting in spring 2023 Max Planck Institutes can rent resources within the HPC Cloud based on a flexible and cost-effective renting model. Projects may start with an evaluation phase in which free tier resources are used to test project-specific use cases for a period of three to six months. Upon the completion of an evaluation phase, projects may transition into production where resources are rented on a rolling basis. More information about the MPCDF HPC Cloud and rental model can be found in our documentation<sup>6</sup>.

### Summary

The HPC Cloud has been supporting cloud and hybrid Cloud-plus-HPC projects since 2021 and the opportunity now exists for new projects to evaluate use cases and rent cloud resources. The core aspects of the MPCDF HPC Cloud offerings are:

- Standard cloud services (compute, storage, networking)
- Solutions for hybrid Cloud-HPC projects (best of both worlds)
- Enabling/design support for MPG research projects (and evaluation projects)
- A flexible and cost-effective billing model

*John Alan Kennedy, Frank Berghaus, Brian Standley*

<sup>6</sup><https://docs.mpcdf.mpg.de/doc/cloud/index.html>

## News

### MPCDF SelfService

Over the past months we have received a number of user requests concerning the registration pages for new MPCDF accounts. These pages and the entire registration workflow have now been modernized to provide a more comfortable experience to both applicants and approvers and to fit in with the overall aesthetics of the platform. Both applicants and approvers are informed about the registration process in more detail. Applicants can now also edit their application data up until its acceptance or rejection.

Furthermore, phone numbers and E-mail addresses now need to be verified before they can be used for 2FA token creation to avoid typos and accidental registration of unauthorized addresses. E-mail addresses used for 2FA must differ from the main account E-mail address to avoid possible circumvention of the 2FA mechanism.

*Amazigh Zerzour*

### Pushing Fusion-Plasma Simulations Towards Exascale

Together with the Max Planck Institute for Plasma Physics (IPP), the MPCDF engages in two projects to improve the performance and scalability of fusion-plasma simulations (particularly the GENE code<sup>7</sup>) towards exascale, that is  $10^{18}$  floating-point operations per second. Precise simulations of fusion plasmas are essential for the development of fusion reactors, like the European ITER experiment or IPP's ASDEX-Upgrade and Wendelstein-7X.

In the Darexa-F (data reduction for exascale applications in fusion research) project, funded by the BMBF as part of the Scalexa program, MPCDF, who is also leading the projects, works together with IPP, the Technical University Munich, the Friedrich-Alexander-University Erlangen-Nürnberg, and ParTec on improving data handling in fusion simulations. We are particularly looking at compression techniques as well as mixed-precision data formats to reduce the overhead introduced by I/O, communication, and memory access. The exploitation of novel hardware, e.g. Data Processing Units (DPUs) is also an aim of the project. Darexa-F started on December 1st, 2022 and will run for three years.

Under the leadership of the Royal Institute of Technology (KTH) in Stockholm, Sweden, a consortium of ten Euro-

pean partners is improving simulations of different plasma applications (space, laser, fusion) towards exascale in a EuroHPC Centre of Excellence called Plasma-PEPSC<sup>8</sup>. The project applies different techniques to improve performance and scalability, including optimizations for GPUs, advanced memory management, novel communication mechanisms, and the exploitation of novel hardware, including upcoming European processors. These techniques are applied to four codes: BIT (Czech Academy of Sciences), GENE (IPP), PIConGPU (Helmholtz-Zentrum Dresden-Rossendorf), and Vlasiator (Univ. of Helsinki). MPCDF is collaborating with IPP on the GENE code and Dr. Tilman Dannert from MPCDF is responsible for the overall technical developments as the project's Technical Director. Plasma-PEPSC started on January 1st, 2023 and will run for four years.

*Erwin Laure*

### Base4NFDI: Creating NFDI-wide basic services in a world of specific domains

NFDI is a German initiative to set up research data infrastructures within all disciplines, covering humanities and social sciences, life sciences, natural sciences and engineering sciences. To ensure sustainability, it will integrate national with international activities.

In addition to domain-specific NFDI consortia, Base4NFDI<sup>9</sup> has been formed. Base4NFDI is a unique joint effort of all NFDI consortia to develop and deploy NFDI-wide basic services. These services will be integrated into the emerging infrastructures at the European level, especially the EOSC. The target group for basic services is the wider NFDI community and, in particular, operators of community-specific services. The resulting NFDI-wide basic service portfolio will be beneficial for all disciplines.

MPCDF has a co-leading role in Task Area 2: service integration and ramping-up for operation. MPCDF is also represented on the *Technical Expert committee* which evaluates proposals for basic service development and gives recommendations on funding those.

*Raphael Ritz*

<sup>7</sup><https://genecode.org/>

<sup>8</sup><https://plasma-pepsc.eu>

<sup>9</sup><https://base4nfdi.de/>

## Events

### Meet MPCDF

Our monthly online-seminar series “Meet MPCDF” has developed to a well-visited and valued training event. On every first Thursday of the month at 15:30 you can participate in an online seminar consisting of a talk usually given by a member of the MPCDF and subsequent discussion. All material can later be found on our training webpage<sup>10</sup>. The schedule of upcoming talks is:

- April, 6th, *The HPC Cloud at the MPCDF*
- May, 4th, *Containers in HPC*
- June, 1st, To be decided
- July, 6th, *The MPCDF Metadata Tools*

We encourage our users to propose further topics of their interest, e.g. in the fields of high-performance computing, data management, artificial intelligence or high-performance data analytics. Please send an E-mail to [training@mpcdf.mpg.de](mailto:training@mpcdf.mpg.de)<sup>11</sup>.

*Tilman Dannert*

### Introduction to MPCDF services

The next issue of our semi-annual workshop “Introduction to MPCDF services” will be held on April 20th, 14:00-16:30 via Zoom. Topics comprise login, file systems, HPC systems, the Slurm batch system, and the MPCDF services remote visualization, Jupyter notebooks and datashare, together with a concluding question & answer session. Basic knowledge of Linux is required. Registration is open<sup>12</sup>.

*Tilman Dannert*

### AI Training Course

In May the MPCDF, in collaboration with Nvidia, will host an “AI for Science Bootcamp” training event. At May 12th-13th, you will learn in an online event how to apply AI tools, techniques, and algorithms to real-life problems. Among others, you will study the key concepts of deep neural networks, how to build deep-learning mod-

els, and how to measure and improve the accuracy of your models. This online bootcamp is a hands-on learning experience where you’ll be guided by step-by-step instructions with mentors on hand to help throughout the process. Registration is open<sup>13</sup>. Since the capacity for the hands-on session is limited, a “first-come-first-serve” policy has to be applied for the registrations.

*Andreas Marek*

### Course on “Python for HPC”

The next iteration of our well-established course on “Python for High Performance Computing” is scheduled for July 25th to 27th, 2023. The event takes place online via Zoom and teaches how to use the Python ecosystem efficiently on HPC systems. We detail on topics such as using NumPy, SciPy, Cython, Numba, JAX, writing compiled extensions in C, C++, Fortran, making use of multithreading, GPU programming, and leveraging distributed memory parallelization using mpi4py and Dask. The lectures in the morning are complemented by exercises and Q&A sessions in the late afternoon. Registration is open<sup>14</sup>.

*Klaus Reuter*

### RDA-Deutschland Conference

During this year’s Love Data Week<sup>15</sup> (February 13th-14th, 2023) the annual conference of the Research Data Alliance Germany<sup>16</sup> took place. The online event featured numerous contributions organized in 14 sessions ranging from RDA for Newbies<sup>17</sup> to Data-Driven Decision Making<sup>18</sup>. Former MPCDF colleague Peter Wittenburg presented the opening keynote reviewing 10 years of RDA and 5 years of RDA Germany<sup>19</sup>. The conference program as well as summaries of the sessions and most of the slides presented are available from the conference website<sup>20</sup>. As in previous years MPCDF contributed to the organization of the event.

*Raphael Ritz*

<sup>10</sup><https://www.mpcdf.mpg.de/services/training>

<sup>11</sup><mailto:training@mpcdf.mpg.de>

<sup>12</sup><https://plan.events.mpg.de/event/56/>

<sup>13</sup><https://www.mpcdf.mpg.de/events/33495/14192>

<sup>14</sup><https://plan.events.mpg.de/e/mpcdf-python-for-hpc-2023>

<sup>15</sup><https://forschungsdaten.info/fr/fdm-im-deutschsprachigen-raum/deutschland/love-data-week/>

<sup>16</sup><https://www.rda-deutschland.de>

<sup>17</sup><https://indico.desy.de/event/37011/contributions/132884/>

<sup>18</sup><https://indico.desy.de/event/37011/contributions/132898/>

<sup>19</sup><https://indico.desy.de/event/37011/contributions/132897/>

<sup>20</sup><https://indico.desy.de/event/37011/>