COMPETENCE CENTER HIGH PERFORMANCE COMPUTING

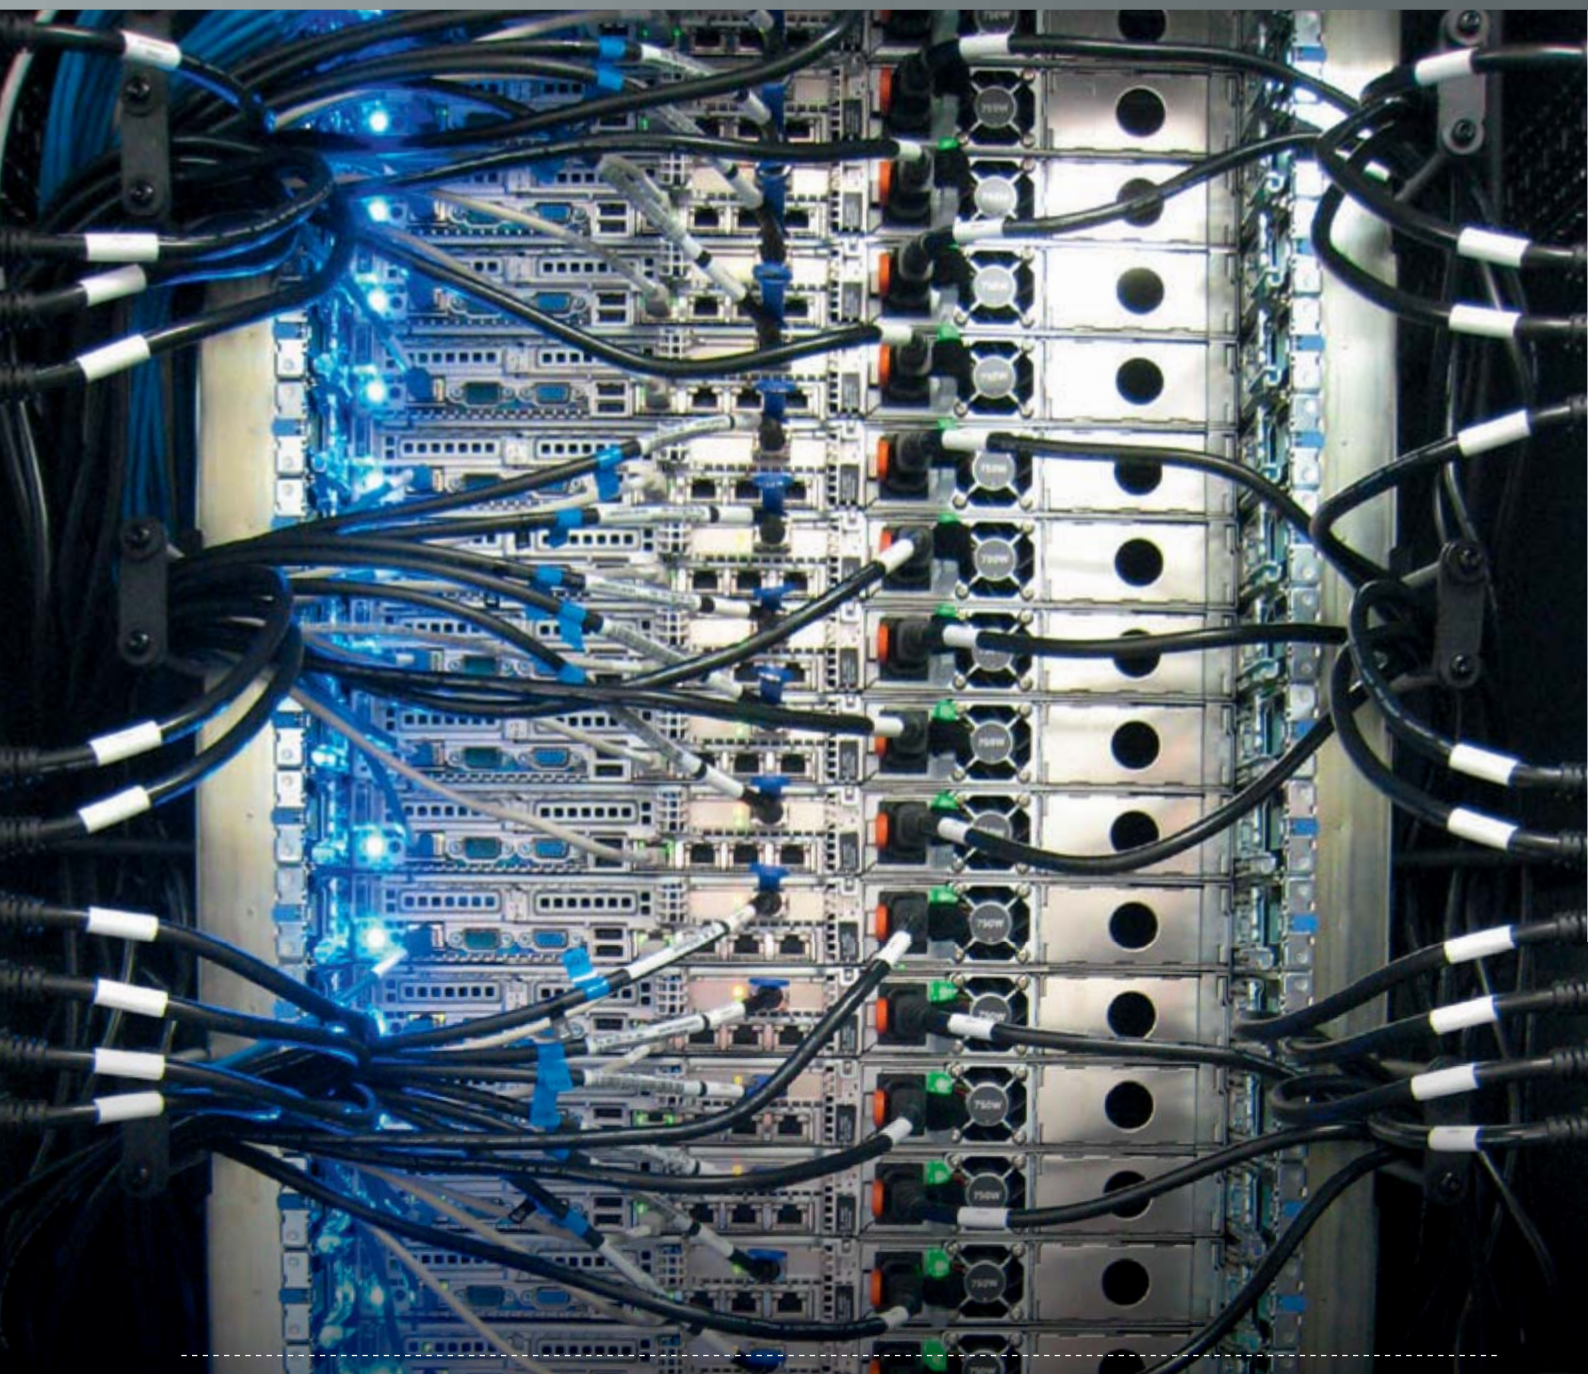

Back side of the computer racks: Separate data arrays are stacked in such “cabinets” and connected to each other in fast networks, for example, Infiniband. The close packing and linking of several such racks creates a mainframe computer that provides modern data centers with the computing power required for future-oriented HPC applications – like Deep Learning or other types of simulation.

DR. FRANZ-JOSEF PFREUNDT
HEAD OF DEPARTMENT

High Performance Computing (HPC) is indispensable for research and economic competitiveness. Basic research in the fields of energy, the material and life sciences and climate research is unthinkable today without sophisticated simulations. The same applies to key sectors of the German economy. Deep learning and machine learning methods are poised to significantly change our society.

The European Union has declared HPC a key strategic issue for Europe and is making a substantial investment in the field. One major element in all of this is the convergence of HPC and Big Data. The department contributes in this environment by participating in the key EU projects and was nominated for the EU Innovation Radar Prize for creating the GPI programming model. In 2016, the ecosystem of myPowergrid technology was significantly expanded to support the management of distributed power storage systems and our priority research topic Deep Learning established international status with algorithmic concepts and deep analyses.

MAIN TOPICS

- HPC System Software and Applications
- Deep Learning and Big Data
- ICT and Renewable Energies
- Parallel File Systems – BeeGFS and BeeOND
- Simulation and Data Analysis in Seismology

Contact
franz-josef.pfreundt@itwm.fraunhofer.de
www.itwm.fraunhofer.de/en/hpc

TWO NEW PROJECTS BASED ON THE MYPOWERGRID TECHNOLOGY

1 *Schoonschip Community in Amsterdam Noord*

2 *Realised peak shaver for load reduction of industrial consumers with load profile metering; the battery supports a maximum power of 60 kW and has a capacity of 60 kWh.*

The integration of alternative energy from sources like photovoltaic plants and wind turbines poses a great challenge for the power grids. The group "Green by IT" is dedicated to creating a reliable, cost efficient, safe and future-oriented energy supply by developing intelligent IT systems and innovative algorithms to enable the transition and deal with the fluctuating production from renewable sources. The year 2016 marked the launch of two new research projects: GreenPowerGrid and Grid-Friends, each with a project term of three years. Both of these projects use the technological foundation provided by the successfully completed myPowerGrid project.

GreenPowerGrid is sponsored by the state of Rhineland-Palatinate and the European Fund for Regional Development (EFRE) to pursue the goal, in cooperation with Stadtwerken Speyer, of enabling a regional green energy supply. The myPowerGrid technology serves as the basis for the development of decentralized PV storage power plant. The goal is to build more than 100 PV storage systems within the city limits of Speyer and provide a local power supply directly to customers using reliable alternative energies. The project has already produced a metering concept and an innovative billing/accounting system. The power supply for 250 households for an entire year has been simulated to a resolution computed every second and shows a positive local self-supply rate by the PV power storage station and, after the initial assessment, a positive new business model for the public utility. A German-Dutch business consortium is developing and evaluating a coordination mechanism between individual power consumers and producers in the Grid-Friends project, which is sponsored by ERA-Net Smart Grids Plus and the German Federal Ministry for Economic Affairs and Energy (BMWi). ITWM's part of the Grid-Friends project is to develop a comprehensive energy management platform for micro-grids and energy communities. In addition to the controls for the storage batteries, the institute is orchestrating a merger of two sectors: that is, the flexible control of heat accumulators in combination with heat pumps and charging stations for electric drive vehicles with respect to the availability of local photovoltaic power supplies. The resulting energy management platform is to be deployed in a demonstrator with 35 floating apartment buildings on a canal in Amsterdam and a residential neighborhood in Cologne-Widdersdorf, Germany. The two demonstrators pursue different objectives: maximum self-supply in Amsterdam and maximum cost efficiency in Widdersdorf.

In addition to these research projects, an agreement with a leading grid operator was signed for consulting and demonstration project: ITWM is planning the installation and follow-on operations for two high-performance battery storage systems for load reduction in industrial operations with load profile metering. The optimized control of the batteries is based on the local energy management system adopted from the myPowerGrid technology.

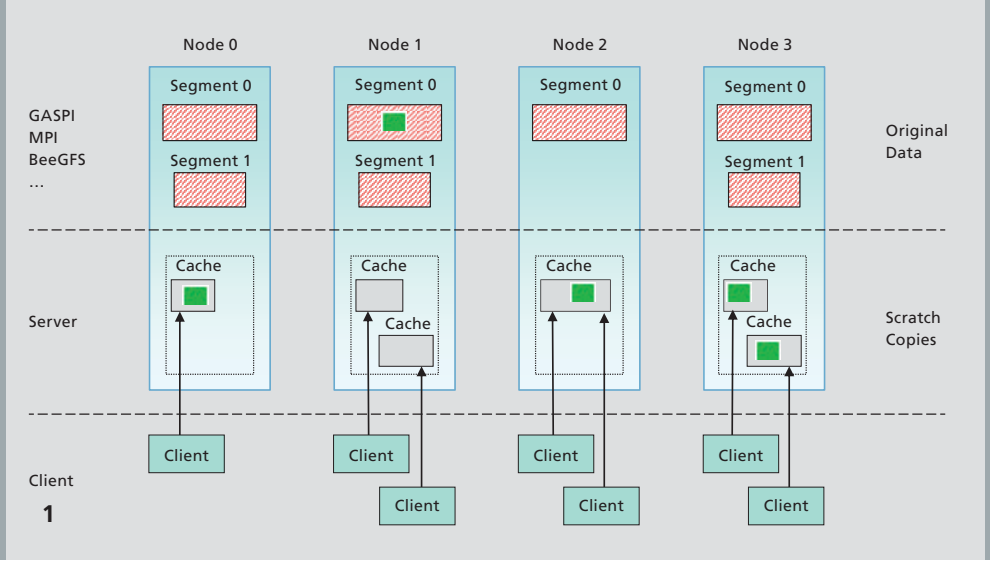

HIGH PERFORMANCE COMPUTING FOR EUROPE

1 *INTERWinE's Directory Cache Client-Server-Architektur*

The CC HPC has been involved in the European funding programs for several years. We are currently involved in six ongoing projects in HPC and transport. We present two highlights of the ongoing projects.

Fraunhofer ITWM's Communication Model GPI (Global Address Space Programming Interface), a key technology in the HPC, is fostered within the scope of the EU-funded INTERWinE project. INTERTWinE promotes the interoperability of communication models available on the HPC market. Communication models are used to communicate data between compute nodes and compute cores. This is important for parallel programs that efficiently utilize the computing power of large computer clusters. GPI allows for asynchronous, multi-threaded communication, avoiding intermediate copies of data. As a result, the computation time can be optimally overlapped with the time required for the communication of the required data. Interoperability between GPI and other communication models allows users to mix already developed code with GPI-based code and to take advantage of the benefits of GPI. In the scope of the INTERWinE project, Fraunhofer ITWM is developing a tool that allows applications that use task-based programming models, which often run on one node, to be scaled on many nodes.

Our distributed runtime environment GPI-Space forms the basis of the IT infrastructure in the EU-funded SafeClouds project. SafeClouds will revolutionize air traffic management. Air traffic will continue to grow in the coming years. In order to improve the security standards cost-neutrally, it is necessary that data is exchanged and intelligent algorithms are applied. This is the only way to quickly identify risk factors from the massive data volumes and initiate counter-measures. GPI-Space forms the basis for the data exchange. The components of GPI-Space enable a division between the algorithmic expertise of the domain experts and the parallelization of the programs for an efficient execution on computer clusters. The runtime environment of GPI-Space is responsible for distributing and executing the program on the available resources (also via geographically distributed systems). The workflow is defined using the runtime environment and executed dynamically and concurrently using GPI. For us the SafeClouds project is exciting as we use our software GPI-Space in a completely new application area. In the next two years, we will accompany the SafeCloud project and analyze its results.

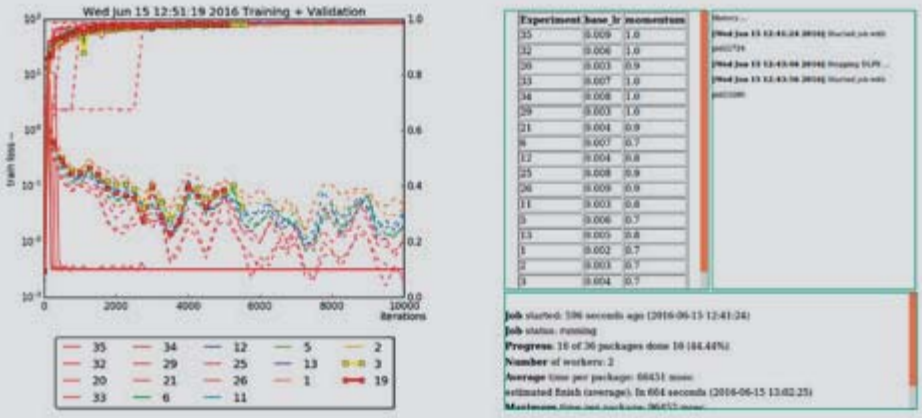

1

SCALABLE DEEP LEARNING BASED ON HPC TECHNOLOGY

1 *DLPS system automates meta-parameter optimization of deep neuronal networks in the cloud.*

In recent years, huge advances have been made in the field of machine learning and its many application areas – from voice recognition to automatic image analysis to autonomous cars to computers that play Go at the World Championship level. Behind all of these achievements are the so-called Deep Learning algorithms. This learning method models the learning problem using very large, complex, artificial neuronal networks. Training such networks requires not only very large data volumes, but also enormous computing power. In terms of the mathematics, the training process corresponds to a non-convex and non-linear optimization problem in the high-dimensional space. Such challenging problems claim days to weeks of calculation time before producing a practical model using the usual optimization methods.

According to the literature, the previously attained acceleration of the training process through parallelization has been rather modest because of the sequential nature of established optimization processes. CC HPC, based on experience in high-performance computing and HPC tools like GPI, GPI space, and BeeGFS, has been able to advance the development of new algorithms in this area to the point of being able to introduce the first scalable solutions.

CaffeGPI was created on the basis of the widely used Caffe open source software and enables – via multiple computing nodes – the distributed training of Deep Learning models on HPC clusters. Thanks to the use of a new optimization algorithm and the GPI-library, CaffeGPI achieves a significantly better scalability than other distributed approaches.

Deep Learning in the Cloud (DLPS): The DLPS system, based on GPI space, enables the automatic and redundant optimization of the meta-parameters of a deep learning model. Adapting these parameters to a specific problem is usually a time consuming and data computing intensive process. DLPS, however, can automate and cost efficiently "outsource" it to the cloud.

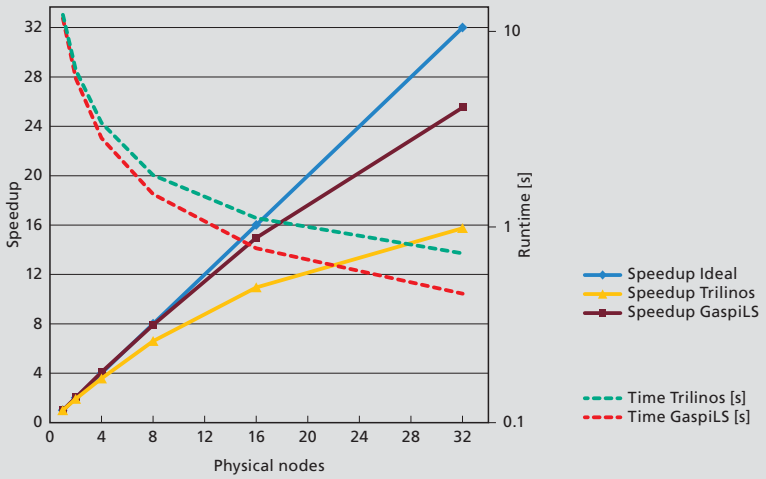

1

GaspiLS – SCALABLE LINEAR SOLVER LIBRARY

Nowadays, in a lot of research fields, simulations form the third pillar of science beside the traditional ones represented by theory and experiment. Computational models are supposed to create a virtual image of nature. In most of the cases, these simulations are based on the discretization of a differential equation. Ultimately, this yields a big system of algebraic equations which are solved e. g. by an iterative solver. The solution of these systems of equations is extremely compute intensive and the effort usually increases with a higher level of detail in the underlying model. The demand for more and more compute power is obvious.

1 *Scalability and absolute runtime of GaspiLS in comparison to Trilinos for a CG based reservoir simulation (Intel Haswell, 256 GB RAM, two sockets per node, six threads per socket)*

However, the CPU clock rate has reached an upper limit on the hardware side. In order to gain more compute power, one is increasing the parallelism on the compute units instead. For example, the amount of floating point numbers which can be simultaneously processed by a single arithmetic operation is raised. On the other hand, the number of compute cores is increased on a single CPU. On top of this, several CPUs are arranged in parallel in a network to form a single big unit. On modern systems, there are up to 20 cores per CPU. This implies that the dual socket systems, which are generally used in the HPC sector, provide up to 40 cores per compute node. One has a huge amount of hardware provided parallelism on a single device which also needs to be reflected in the software layer. The so called accelerator architectures which host up to 72 cores are even more extreme within that respect.

If one wants to exploit the increasing parallelism provided by modern hardware for the simulations, software with good scalability features is required. Here, scalability is a measure which quantifies the additional benefit which is generated by adding additional hardware resources. That means, software which optimally scales exploits the hardware provided resources to 100 %. This has to be the ultimate goal in order to use the hardware efficiently.

GaspiLS is a numerical solver library for iterative methods, which has been designed right from the beginning to provide a good scalability. It is based on the GPI-2 programming model and implements an asynchronous data dependency driven execution model with perfect overlap of communication and computation. It avoids global synchronization points as much as possible. These are the basic principles and necessary conditions to achieve a good scalability. The object oriented design of GaspiLS defines abstract interfaces for matrices, vectors, iterative solvers and preconditioners. As such, it is easy to extend. GaspiLS provides several solvers like (P)CG, BiPCGStab and GMRES. It is also providing preconditioners like e. g. Jacobi, ILU(0) and ILUM(0). GaspiLS is going to be released OpenSource under GPLv3 license. At the moment, GaspiLS is used in two industry projects.

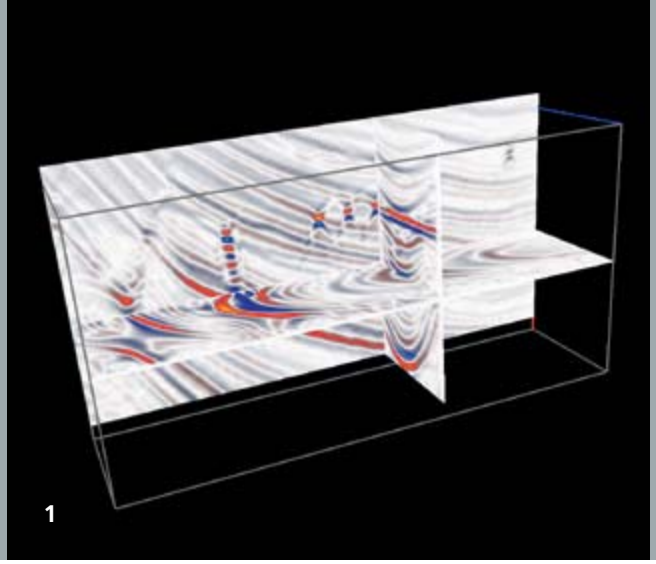

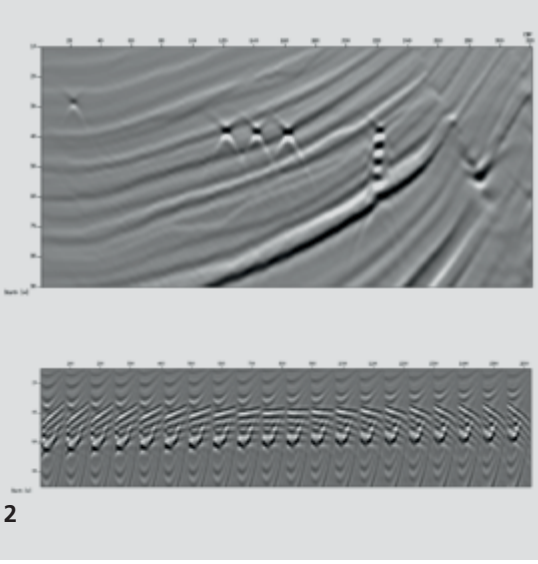

USING SEISMIC DATA TO DETECT BOULDERS

1 *Position of underground diffractors as result of GRT migration. GRT interim results shown as fold-outs highlight the linear appearance of the diffractions.*

2 *Subsurface image points of diffractions as obtained by GRT migration (upper figure); the linear events with varying dip are an indicator for existing diffraction (lower figure).*

Reflection seismology ranges over a large area of transmitted acoustic wavelengths and is not limited to the frequency bands around 50 Hz so relevant for oil and gas exploration at reservoir depths of several thousand meters. Besides providing structure, stratigraphy, and rock parameter data at such depths, it can also provide detailed, high resolution exploration of the shallow underground.

In this BMWi sponsored project, CC HPC brings many years of experience and competence in developing analysis methods for seismic data and, in cooperation with Fraunhofer IWES, is now working on a process chain to find subsurface boulders having a diameter larger than one meter at subsurface depths up to 100 meters. Energy companies have a great interest in discovering such rock obstacles when planning offshore wind parks and making adjustments for the exact positioning of the turbine towers.

Boulders resting just below the surface or embedded at a shallow depth in the sea bed are easily detected by bathymetric surveys and very high resolution sonar systems. At depths of 10–80 meters below the sea bed, the surveys require the use of deeper penetrating seismologic methods. The idea for the planning of the offshore wind parks is to rely on the seismic acquisition measurements already collected regarding the issues of structural stability. At frequencies around 300 Hz and the resulting wavelengths of 5 meters, boulders of the dimension specified above can be detected using the diffraction response patterns in the seismic data sets.

To separate the diffractions from the comparatively much higher amplitude reflections, CC HPC developed a two-step process, consisting of a multi-focusing technique and our GRT pre-stack depth migration. The reflections are weakened in favor of the diffractions in both steps to produce a limited background image of the diffracting object, which is amplified and inserted in the classic structural migration result to expose the spatial relationships. An interpretation of the results enables conclusions to be drawn about the frequency of occurrence as well as the location of individual objects.

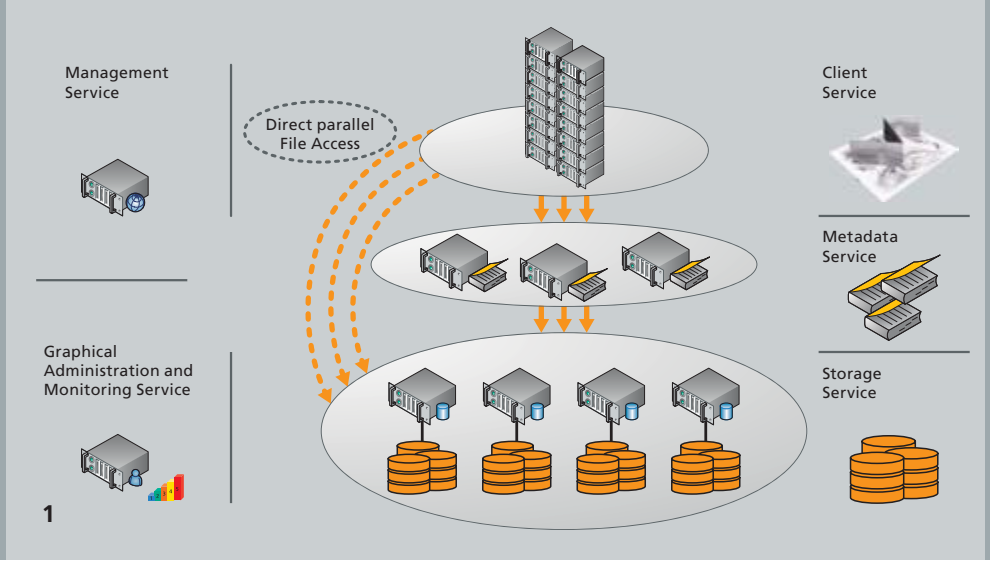

1

BeeGFS – PARALLEL FILE SYSTEM WITH MAXIMUM PERFORMANCE AND HIGH AVAILABILITY

The performance of state of the art processors and network technologies is constantly improving and new possibilities are arising for handling ever more complex problems and for obtaining more realistic and detailed simulation results. However, this also requires the ability to work with very large data sets, often in the terabyte range. To manage these huge amounts of data and ensure optimum computing performance, CC HPC has been working for several years on the BeeGFS parallel file system where parallel read and write function is possible. In this system, individual files are distributed to multiple servers in so called chunks. This method makes it possible to process data sets at speeds many times faster than usual. From the outset, scalability, maximum performance, and flexibility as well as ease of use have been the cornerstones of the development effort and, today, these features are the reasons why the system is increasingly attractive to a wide and global user base.

Last year, BeeGFS was extended by high availability mechanisms, which enables customers to use the software to create a highly reliable global file system. This system requires no other third party software or special high availability hardware. In addition, BeeGFS supplies the BeeOND tool that creates a parallel file system on a defined set of hardware at the push of a button. If required, BeeOND can be used as a dynamic parallel file system directly on the compute nodes. Temporary calculation data can be processed directly, without the need to access a global storage system over the network.

BeeGFS is supplied free of charge at www.beegfs.com. The software is distributed under an open-source license and the source files can be obtained from the website. A Fraunhofer ITWM spin-off company, ThinkparQ, has provided worldwide commercial support for BeeGFS since 2014. Meanwhile, CC HPC continues its development of the software.

The CC HPC development team is also contributing its extensive knowledge to a number of successful EU-sponsored projects. New methods and computer architectures are being developed in cooperation with other project partners, with special attention on exascale computing. When working with such high performance computing units, itneeds to be ensured, that the data to be processed can be provided fast enough. This fact makes BeeGFS one of the key elements of the project.

1 *Overview of BeeGFS architecture: Separate services for metadata and user data enable independent scalability of both functions.*

BeeGFS
developed by Fraunhofer