SEVENTH FRAMEWORK PROGRAMME

FP7-ICT-2013-10

DEEP-ER

DEEP Extended Reach

Grant Agreement Number: 610476

D1.9

Final report

Approved

Version:	2.0
Author(s):	E.Suarez (JUELICH)
Contributor(s):	S.Eisenreich (BADW-LRZ), H.Ch.Hoppe (Intel), K.Thust (JUELICH), V.Beltran (BSC), A.Zitz (JUELICH), J.Schmidt (UHEI), I.Zacharov (Eurotech)
Date:	04.05.2017

Project and Deliverable Information Sheet

<table>
<tr><td rowspan="7">DEEP-ER Project</td><td colspan="2">Project Ref. №: 610476</td></tr>
<tr><td colspan="2">Project Title: DEEP Extended Reach</td></tr>
<tr><td colspan="2">Project Web Site: http://www.deep-er.eu</td></tr>
<tr><td colspan="2">Deliverable ID: D1.9</td></tr>
<tr><td colspan="2">Deliverable Nature: Report</td></tr>
<tr><td rowspan="2">Deliverable Level:
PU*</td><td>Contractual Date of Delivery:
31 / March / 2017</td></tr>
<tr><td>Actual Date of Delivery:
31 / March / 2017</td></tr>
<tr><td colspan="2">EC Project Officer: Juan Pelegrín</td></tr>
</table>

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

<table>
<tr><td rowspan="6">Document</td><td colspan="2">Title: Final report</td></tr>
<tr><td colspan="2">ID: D1.9</td></tr>
<tr><td>Version: 2.0</td><td>Status: Approved</td></tr>
<tr><td colspan="2">Available at: Publishable part at: http://www.deep-er.eu</td></tr>
<tr><td colspan="2">Software Tool: Microsoft Word</td></tr>
<tr><td colspan="2">File(s):DEEP-ER_D1.9_Final_report_v2.0-ECapproved</td></tr>
<tr><td rowspan="4">Authorship</td><td>Written by:</td><td>E.Suarez (JUELICH)</td></tr>
<tr><td>Contributors:</td><td>S.Eisenreich (BADW-LRZ), H.Ch.Hoppe (Intel), K.Thust (JUELICH), V.Beltran (BSC), A.Zitz (JUELICH), J.Schmidt (UHEI), I.Zacharov (Eurotech)</td></tr>
<tr><td>Reviewed by:</td><td>S.Höffler-Thierfeldt (JUELICH), I.Schmitz (ParTec)</td></tr>
<tr><td>Approved by:</td><td>BoP/PMT</td></tr>
</table>

Document Status Sheet

Version	Date	Status	Comments
1.0	31/March/2017	Final	EC submission
2.0	04/March/2017	Approved	EC approved

Document Keywords

Keywords:	DEEP-ER, HPC, Exascale, progress report, month 42, results, impact, dissemination, exploitation

Table of Contents

List of Figures

List of Tables

1 Publishable summary

1.1 Executive Summary

The DEEP – Extended Reach (“DEEP-ER”) project was active from October 2013 through March 2017 for a term of 42 months. The project extended the **Cluster-Booster Architecture** – first realised in the DEEP project[1] – by local, fast storage-class memory that enables a highly scalable I/O system and an efficient mechanism to recover applications from hardware failures. In this way, DEEP-ER has addressed two significant Exascale challenges: the growing gap between **I/O** bandwidth and compute speed, and the need to significantly improve system **resiliency**.

The hardware prototype built in DEEP-ER belongs to the second generation Cluster-Booster systems and includes several improvements with respect to its DEEP predecessor: new processor generations provide substantial performance and efficiency enhancements in particular for the Booster side, a fast, new fabric technology delivers unmatched bandwidth per node, and most importantly an **innovative memory hierarchy** enables unrivalled I/O performance. The Booster designed and built within the DEEP-ER project uses Eurotech’s Aurora direct liquid cooling technology and consists of 72 Intel® Xeon Phi™ (KNL) nodes, delivering a peak performance of around 200 TFlop/s. Each compute node has 400 GByte of fast local storage-class memory attached via PCI Express.

The DEEP-ER multi-level memory/storage hierarchy and I/O infrastructure has been designed to support data-intensive applications and multi-level checkpointing/restart techniques. The project has developed a scalable and efficient I/O software platform based on the BeeGFS parallel file system, the parallel I/O library SIONlib, and the Exascale10 (E10) optimisations to MPI-IO. Together, they enable efficient and transparent use of the DEEP-ER system innovations and provide all the required I/O functionality.

On top of this I/O infrastructure, DEEP-ER has developed an efficient and user-friendly resiliency concept combining transparent task-based application restart techniques with user-level checkpoints. The OmpSs runtime, which allows identifying the application's individual tasks and their interdependencies, has been extended to automatically re-start tasks in case of transient hardware failures. In combination with a multi-level user-based checkpoint infrastructure that uses the local storage to recover from non-transient hardware-errors, applications are able to better cope with the failure rates expected in Exascale systems.

All the hardware and software components in the DEEP-ER project have been developed in intense, iterative co-design collaboration. The seven high impact applications participating in DEEP-ER have provided requirements, benchmarking results and user feedback to the hardware and software developers who, themselves, have supported the application scientists in porting and adapting theirs codes. The DEEP-ER applications have demonstrated that, as a result of DEEP-ER, I/O parts of HPC codes can run significantly faster and scale up much better, and that writing checkpoints and restarting tasks or full applications can be done with minimal overhead.

[1] www.deep-project.eu

1.2 Project context and objectives

The DEEP-ER project started in a time in which the HPC community was still struggling to find solutions to the multitude of challenges along the way from Peta- to Exascale computing. The most important ones had been identified[2] as: *i)* reducing the power consumption of HPC systems, *ii)* achieving the high levels of parallelism needed to extract performance from a machine with millions of cores, *iii)* dealing with the frequent hardware failures expected in such huge systems, and *iv)* addressing the increasing gap between the growth rate of compute power and the amount and performance of memory and storage available in HPC systems.

Preceding DEEP-ER, the DEEP project had addressed the first two topics by introducing the Cluster-Booster Architecture and first demonstrating it within a highly efficient hardware packaging an adapted software stack. The DEEP Prototype wisely combined general purpose processors with energy efficient many-core co-processors, integrating both into a dense chassis and rack design equipped with direct warm-water cooling and a full-fledged power monitoring system. Furthermore, DEEP's approach to heterogeneous computing was pioneer in making accelerator/co-processor technology becoming autonomous – able to operate and communicate with each other without a host CPU –, and in supporting arbitrary combinations of "regular" and accelerator/co-processor nodes. With this concept the inherent mix of different scalability levels present in complex HPC codes could be directly mapped onto the HW, allowing applications to scale further and deal better with higher concurrency levels. The DEEP software stack made sure that code developers could profit from this specific architecture without having to rewrite their codes and, therefore, keeping them portable.

Building upon the DEEP achievements, DEEP-ER set its goals at addressing the two remaining Exascale challenges (*iii* and *iv* from above) by extending the Cluster-Booster Architecture with additional I/O and resiliency functionalities. New, non-volatile memory technologies have been integrated to create a multi-level memory/storage hierarchy and I/O infrastructure, which, in combination with a highly scalable I/O software stack provides support for data-intensive applications. Additionally, and based on this advanced I/O functionality, an efficient and user-friendly resiliency concept combining user-level checkpoints with transparent task-based application restart has been developed.

The development of the DEEP-ER Prototype and its I/O and resiliency software packages has been driven by seven significant European HPC applications. This exemplary co-design effort has produced a hardware & software solution that perfectly fits the application's needs. Deep technical discussions took place throughout the whole project duration between SW and HW experts, and between the I/O and resiliency teams, making DEEP-ER a unique example of an all-round co-design project. The ported and optimised application codes showed the usability, performance and resiliency of the DEEP-ER concepts, demonstrating the success of this unique co-design project.

The specific objectives of the DEEP-ER project and the key results achieved in the project towards them are (text in italics copied verbatim from the DoW):

[2] See DARPA report by P. Kogge et al. on "Exascale Computing Study: Technology Challenges in Achieving Exascale Systems", 2008: http://www.cse.nd.edu/Reports/2008/TR-2008-13.pdf

1. *Address two main Exascale challenges: I/O and resiliency. DEEP-ER will **extend the DEEP Architecture by**: i) a highly scalable, efficient and easy-to-use parallel **I/O system**; ii) providing a combination of **low-overhead user-level checkpoint/restart and automatic task recovery**.*

   The DEEP-ER system prototype (see objective 2 and 3) enhances the Cluster-Booster Architecture from DEEP by including fast, local storage-class memory in the Cluster and Booster nodes, and by introducing the concept of a shared, network attached memory resource. The DEEP-ER optimised I/O and resiliency software stack (see objectives 4 and 5) has been developed to leverage these new capabilities.

   Intensive discussions took place between the hardware, software, and application developers during the whole project duration, making co-design the core of the DEEP-ER project. This approach made sure that the project developments fulfil the user requirements with respect to the focus topics of I/O and resiliency, paying attention also to other important aspects such as programmability, portability and scalability.

2. ***Develop a prototype system of the extended DEEP Architecture** that leverages advances in hardware components (Intel's second generation Intel® Xeon Phi™ processors, high-speed interconnects and non-volatile memory devices) to further improve the performance and efficiency of the DEEP-ER Prototype and realise the novel I/O system and resiliency improvements. This prototype will allow proving the viability of the concept for 500 PFlop/s-class of supercomputers*.

   The installation of the DEEP-ER Prototype – composed of a Cluster and a Booster – will be completed after end of the project, during April 2017. The DEEP Cluster is an off-the-shelf system with 16 Intel® Xeon® (Haswell generation) processors. The Booster, developed following the project requirements, is based on Eurotech's direct liquid-cooled Aurora product line and contains 72 Intel® Xeon Phi™ (Knights Landing/KNL) boards in a blade-center design. Cluster and Booster are connected internally and with each other by the European EXTOLL TOURMALET interconnect fabric, which provides high-speed connectivity over the whole DEEP-ER system. Each node includes a 400 GByte storage-class memory device, which is attached via PCI Express and supports the NVMe protocols. Together with the KNL on-package memory, the on-node DRAM, the external global storage, and two Network Attached Memory (NAM) devices, these constitute the innovative DEEP-ER memory hierarchy (see objective 3).

   While the Cluster and NAM nodes are air-cooled, all components of the DEEP-ER Booster are direct liquid cooled, using Eurotech's coldplate technology. Inlet temperatures of up to 50 degrees Celsius are possible, allowing year round free cooling and attaining high energy efficiency. All together the DEEP-ER Prototype provides a peak performance of about 200 TFlop/s, offers a DRAM capacity of about 8 TByte, and has a total energy consumption of approx. 30 KW at full load.

   The DEEP-ER Cluster was installed early in the project, as it was used as a software development and application porting platform. A small-sized

DEEP-ER Booster – with 8 Booster nodes – was attached to it already in Q2/2016, and has been widely used by application and system software developers to complete their work. The combination of the DEEP Cluster and this 8-node Booster has received the name of “Software Development Vehicle”, or SDV.

As mentioned above, the installation of the full-fledged 72-node Aurora Booster is scheduled for April 2017. The late installation is the consequence of delays in obtaining critical components, designing, qualifying and manufacturing the required new boards, and finally in putting together and bringing up the final system. These difficulties demonstrate the challenges involved in building a leading-edge HPC system. With the high impact of the hardware-development risks in mind, progress of key tasks was tracked closely throughout the project, and delays were mitigated by deploying the DEEP-ER Software Development Vehicle (SDV), mentioned above.

The applications ported to the SDV have clearly demonstrated the potential of the DEEP-ER approach and its advantages with respect to existing systems. The project has developed modelling tools for extrapolation to 500 PFlop/s-class supercomputer configurations – but due to the very limited scale of the SDV, a meaningful extrapolation was not feasible. However, the clear performance improvements by the local storage and the adapted I/O stack will carry over to large configurations and materially help scalability. The same is true about the local, “buddy” checkpointing scheme. The EXTOLL interconnect by its nature scales up linearly without the need for additional switches, and applications with local communications will directly profit from this property.

3. ***Explore*** *the potential of* ***new storage technologies*** *(non-volatile and network attached memory) for use in HPC systems, with a focus on parallel I/O and system resiliency by integrating them with the DEEP-ER Prototype*.

   Each of the 88 nodes of the DEEP-ER Prototype includes a 400 GByte non-volatile storage-class memory device, which is attached via PCI Express and supports the highly efficient NVMe protocol. Extensive experiments and a wide range of measurements with I/O benchmarks, application mock-ups and finally the full applications clearly show substantial performance increases over best-of-breed conventional SSDs, in particular for scenarios with many parallel I/O requests, and over a state-of-the-art Intel Xeon based storage system. The DEEP-ER I/O and resiliency software has been extended to exploit the node-local storage (see Objectives 4 and 5).

   DEEP-ER has also introduced an innovative memory concept: the Network Attached Memory (NAM), a high-speed memory device that is directly attached to the network and enables remote memory accesses from all nodes in the system. This memory resource is made available to system and application software via the libNAM library, and the local compute capability of the NAM has been demonstrated in a checkpoint/restart use case, where the NAMs pulls the data for checkpoints from the nodes and computes all parity information itself.

4. ***Develop*** *a highly scalable, efficient and user-friendly* ***parallel I/O system*** *tailored to HPC applications. The system will exploit innovative hardware features, optimise I/O routes to maximise data reuse, and expose a user friendly interface to applications. Its design will meet the requirements of traditional, simulation-based as well as emerging data-intensive HPC applications*.

   The I/O software combines three main components: the global file system *BeeGFS*, the scalable parallel I/O library *SIONlib*, and the *E10* optimized MPI-IO implementation. Together, they provide the I/O functionality required by applications, covering different kinds of I/O approaches and access patterns.

   In BeeGFS cache domain handling, user-level stripe-size definition, and native support for the EXTOLL network have been implemented. The cache domain – based on *BeeGFS on demand* (*BeeOND)* – runs on the node-level fast storage devices and can be used synchronously or asynchronously, providing a most efficient alternative to the traditional POSIX I/O access and reducing the frequency of accesses to the global storage, and thus ultimately increasing the file system's scalability.

   SIONlib has been refactored and improved to eliminate code replications and to increase its overall modularity and manageability. Furthermore, it has been extended to support *buddy-checkpointing* and the NAM (see Objective 5).

   The newly developed E10 optimisations to the MPI-IO parallel I/O system are based on the ROMIO implementation and improve the performance of applications using collective I/O operations. E10 automatically intercepts MPI I/O calls, so that its use is fully transparent to applications. Furthermore, E10 supports ParaStation MPI and the caching functionality of BeeGFS.

5. ***Develop*** *a unified user-level system that significantly reduces the checkpointing overhead by exploiting multiple levels of storage and new memory technologies. Extend the DEEP programming model to combine* ***automatic re-execution of failed tasks and*** *recovery of long-running tasks from* ***multi-level checkpoints****, and introduce easy-to-use annotations to control checkpointing*.

   DEEP-ER's checkpointing strategy is orchestrated by the Scalable Checkpointing Library SCR, which has been extended to support the DEEP-ER memory/storage hierarchy, I/O system and software architecture. SCR determines the frequency and location of checkpoints within the system's multi-level memory hierarchy based on a specially developed failure model. The fast local system storage is used for *buddy-checkpointing*, – implemented combining SCR with SIONlib –, which saves checkpoint data in the NVM devices of both the local and a companion node. In case of hardware failure, the data lost on the damaged node can be then recovered from its "buddy" and the application can continue with minimum time-lost.

   OmpSs and ParaStation have been extended to jointly enable the recovery of failed tasks. *Task-recovery* drastically reduces the compute time lost due to hardware failure, especially for applications that offload tasks, i.e. with large tasks offloaded from Cluster to Booster or vice versa. When a failure occurs, the application recovers by simply making the parent task re-spawn the failed task, without any of the other tasks being affected.

Creating functionalities such as buddy-checkpointing or task-recovery is only possible when software packages with complementary functionality are combined smoothly. The unique co-design spirit of the DEEP-ER project and all its members has made such joint development possible.

6. ***Analyse the requirements of HPC codes*** *carefully selected to represent the needs of future Exascale applications with regards to I/O and resiliency,* ***guide the design of the DEEP-ER hardware and software components****, optimise these applications for the extended DEEP Architecture and use them to evaluate the DEEP-ER Prototype. Selected applications cover the fields of Health, Earthquake Physics, Radio Astronomy, Oil Exploration, Space Weather, Quantum Physics, and Superconductivity.*

   Continuous co-design discussions took place to identify the application requirements in terms of hardware capabilities, I/O and resiliency functionalities. Starting with a simple questionnaire to the application developers, the co-design discussions evolved during project execution gradually increasing the level of detail and including additional aspects of hardware, software and application development and testing. DDG teleconferences and face-to-face meetings have been used as the co-design discussion platform.

   The applications were analysed and performance and scaling tests took place. In several cases mock-ups of the applications have been developed, to easily implement code changes and analyse their impact on the overall performance. These analysis activities have helped to improve the applications in important ways (e.g. with regards to vectorisation, communication strategies and/or numbering schemes, I/O, checkpointing, etc.). Furthermore, the I/O and resiliency strategies implemented in DEEP-ER (objectives 4 and 5) have been integrated with the applications, selecting the most appropriate I/O software components for each one.

   This collaboration between applications and the other project teams was complemented by co-design discussions between hardware and software developers, as well as between developers for different software packages. This global approach to co-design makes the DEEP-ER project and its developments unique in its class.

7. ***Demonstrate and validate the benefits of the extended DEEP Architecture*** *and its first implementation (the DEEP-ER Prototype) with the DEEP-ER pilot applications and for applications that exploit generic multi-scale, adaptive grid and long-range force parallelisation models.*

   Benchmarking campaigns – notably on the Software Development Vehicle (SDV) – demonstrated the improvements created by the DEEP-ER hardware and software innovations on the overall application performance, scalability and resiliency. It is important to highlight that code modifications have been performed assuring application portability, and that performance improvements have been seen also on other systems outside of the project.

   A combination of analysis and modelling tools has enabled us to predict application performance and characteristics to larger system configurations,

given a sequence of application traces created on smaller system scales. Initially, this was targeting pre-Exascale systems with 100,000s of nodes, and prediction of the compute and communication characteristics was possible by taking tracefiles from large systems outside of the project.

To also cover I/O characteristics, the Extrae event tracing library was extended to record I/O traces, which contain events for relevant POSIX and MPI-IO calls. These traces allow for a close inspection of I/O behaviour, and they are the base for computing efficiency measures for the I/O operations. These in turn can then be extrapolated to arrive at predictions for larger system configurations. The delay in making the full DEEP-ER Prototype available has severely limited the basis of scale of traces that could be created, and consequently, no meaningful extrapolations could be made before the end of the project.

Finally, a MoU has been put in place between DEEP-ER and PRACE-4IP to make the full DEEP-ER Prototype available to PRACE users after the end of the DEEP-ER project. This collaboration shall provide further validation results, increase the impact of the DEEP-ER project, and facilitate the adoption of its developments by the European HPC user community.

1.3 Main scientific and technical results/foregrounds

The DEEP-ER project has proven that applying co-design systematically throughout the full HPC development chain fosters innovation in all the individual areas involved and enables new global solutions to the Exascale challenges. The results and foregrounds generated in the project involve and combine very diverse aspects of HPC. For the purpose of this report, the main project results have been classified into three categories: hardware, software, and applications.

1.3.1 Hardware

The DEEP-ER project has defined the second generation of the Cluster-Booster Architecture and designed and built an actual prototype implementation (see Figure 1). In this DEEP-ER Prototype the second generation of Intel® Xeon Phi™ processors (code name “Knights Landing” or KNL) provide compute power for the Booster Nodes, while Intel® Xeon® processors (“Haswell” generation) populate the Cluster Nodes. All nodes in the system (in both Cluster and Booster) feature additional non-volatile memory (NVM) devices for efficiently buffering I/O and storing checkpoints. A uniform high-speed EXTOLL interconnect runs across Cluster and Booster, and network-attached memory (NAM) devices connected to it provide high-speed shared memory access. They add to the node-local memory and external pool of hard-disk storage, building an intermediate level in the system’s memory hierarchy. A final, most innovative ingredient is the new Network Attached Memory (NAM), which provides high speed, non-volatile, memory globally accessible to all nodes in the system.

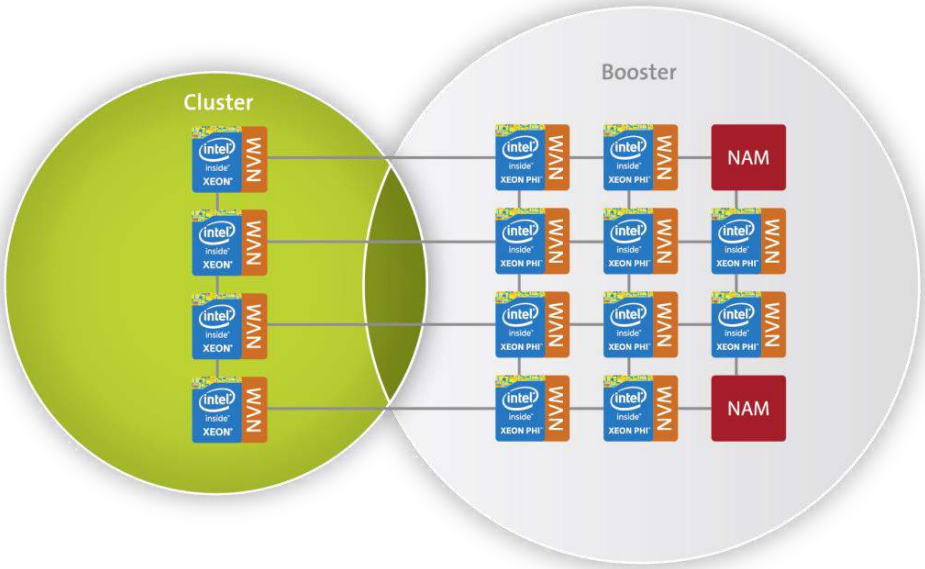

Figure 1: The DEEP-ER system architecture

The DEEP-ER Prototype is based on the Eurotech Aurora Blade architecture. With this product line Eurotech has specialised in high-density, energy efficient direct liquid cooled cluster systems for HPC. The KNL blade that Eurotech developed in DEEP-ER integrates the half-width, 1U Intel Server board S7200AP and creates a new addition to the Aurora product line. A highlight of its design is its space-efficient cold plate technology, which has been enhanced to support memory DIMMs (of the ultra-low profile or ULP-DIMM variety), as shown in Figure 2. Each KNL node is populated with 6x ULPDIMMS, with a total capacity of 96 GByte DDR4 memory, in addition to KNL's 16 GByte of fast on-package MCDRAM. With this innovative approach Eurotech achieves highest density, fitting 18 KNL nodes into a single 19" wide chassis. The direct liquid cooling system also provides high energy efficiency, as its support inlet temperatures of up to 50^{o} Celsius.

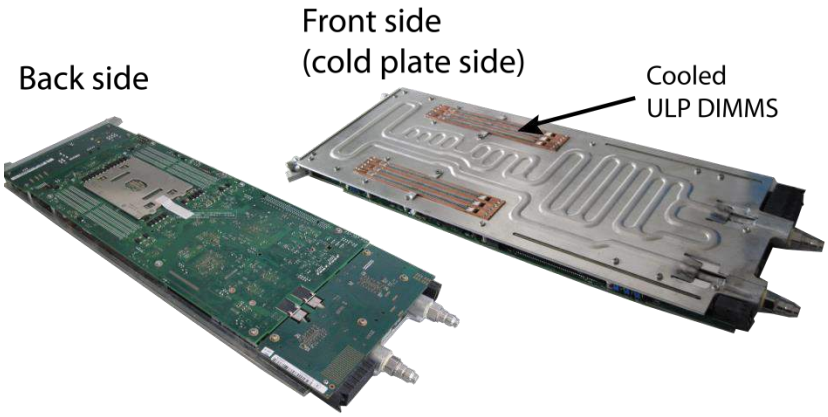

Figure 2: KNL Blade of the Eurotech Aurora Booster.

The 18 KNL blades are plugged into a PCIe backplane to drive the signals to the Root Card located in the upper part of the chassis (see Figure 3), which itself hosts the system's peripheral devices. In the DEEP-ER Booster the Root Card is populated with 18 EXTOLL Tourmalet NICs and 18 fast NVM devices. Each of the KNL slots in the lower part of the chassis is connected via PCIe to one NIC and one NVM device in the Root Card, giving each of the KNL nodes access to the EXTOLL high speed network and additional 400 GByte

storage-class memory. The backplane and Root Card were designed to support PCIe gen3 speeds.

The chosen NVM technology is Intel's DC P3700, an SSD replacement device with 400 GByte capacity that provides high speed, non-volatile local memory, attaches to the node with 4 lanes of PCIe, and supports the highly efficient NVMe protocols. Extensive experiments and a wide range of measurements with I/O benchmarks and application mock-ups have been performed, which clearly show substantial performance increases over conventional best-of-breed SSDs, in particular for scenarios with many parallel I/O requests, and over state-of-the art I/O servers.

Figure 3: Front view of an Aurora Booster chassis.

The DEEP-ER Booster has been designed and built in the project and it consists of 4 Aurora chassis, with a total of 72 KNL nodes and the same amount of NVM devices and EXTOLL TOURMALET NICs. Its companion – the DEEP-ER Cluster, see Figure 4 –, is an off-the-shelf cluster with 16 Intel Xeon (Haswell generation) servers, each of them having a EXTOLL Tourmalet NIC and a NVM device attached. The DEEP-ER Cluster was available early in the project and used as Software Development Vehicle (SDV). All together the DEEP-ER Prototype achieves a total peak performance of about 200 TFlop/s, and a total memory capacity of 8 TByte.

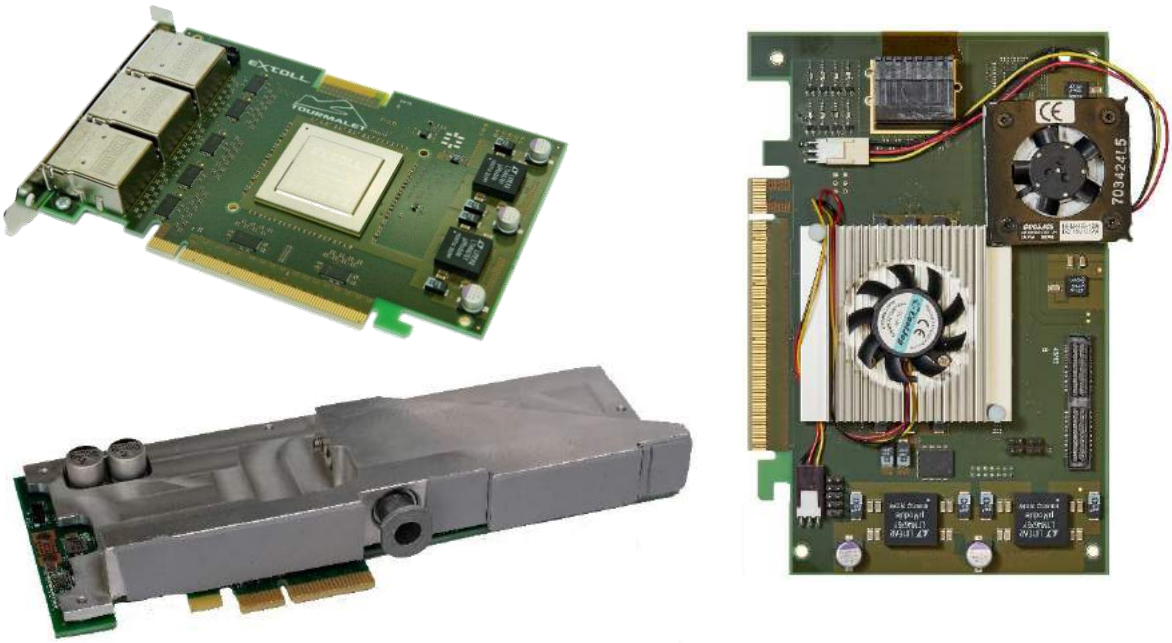

Figure 4: EXTOLL TOURMALET NIC (upper left), Intel NVM device with custom-made Eurotech coldplate (lower left) and NAM device (right).

DEEP-ER has also introduced an innovative memory concept: the Network Attached Memory (NAM). Packed into a PCIe add-on card form factor, the NAM combines Hybrid Memory Cube (HMC) devices with a state-of-the-art Xilinx Virtex 7 FPGA to create a high-speed memory device that is directly attached to the EXTOLL fabric and therefore can be accessed by all nodes in the system. The FPGA implements the HMC controller, the EXTOLL network interface providing two links with the full TOURMALET speed of 100 Gbit/s per link, and finally the NAM logic for remote memory access and the parity calculations for the checkpoint/restart use case. The HMC controller has been developed by UHEI and its design has been released as Open Source (http://www.uni-heidelberg.de/openhmc). The first implementation of the NAM provides just 2 GByte capacity, due to limitations of current HMC technology. Future implementations can, however, increase capacities and may trigger a rethinking of memory architectures for HPC and data analytics. The libNAM library has been implemented to give system and application software on the nodes access to the NAM memory pool and enable the execution of any pre-defined functions in the NAM logic. In DEEP-ER, a checkpointing/restart use case has been implemented, which uses the NAM FPGA to pull the required data from the compute nodes and locally calculate the parity information.

The NAM exploits the remote DMA capabilities of the EXTOLL fabric, which enable remotely accessing memory resources without the intervention of an active component (e.g. a CPU). The EXTOLL TOURMALET fabric was developed outside the DEEP-ER project. It relies on an ASIC-based NIC that provides six links of 100 Gbit/s bandwidth each, and contains logic for highly efficient message transmission and DMA. EXTOLL is a direct network – each NIC contains a switch that routes packets between its six links, and topologies of arity 6 or less can be built without the need for additional switches. With its performance and flexibility, EXTOLL establishes itself as the leading European interconnect technology for HPC.

The specific design and implementation of the DEEP-ER Prototype has been chosen to match the requirements from the software and application users. An intensive co-design effort driven by the Design and Development Group (DDG), a body from technical experts in the various parts of the project, has made this possible.

At the time of writing, the integration of the DEEP-ER Prototype system is still work in progress. All components are available and have been tested and validated, yet a number of unexpected technical difficulties have slowed down the process of bringing up the full 72-node Booster system. It is still planned to have the system up and running with PCIe generation 2 connectivity between Booster nodes and NVM devices and EXTOLL TOURMALET NICs in time for the final DEEP-ER review.

1.3.2 Software stack, programming environment and tools

1.3.2.1 I/O software

The DEEP-ER memory hierarchy and multi-level I/O infrastructure has been designed to support data-intensive applications and multi-level checkpointing/restart techniques. Employing this infrastructure, the project developed a scalable and efficient I/O software platform based on the BeeGFS parallel file system, the parallel I/O library SIONlib, and the Exascale10 (E10) optimisations for MPI-IO, as schematically represented in Figure 5. All together they enable the efficient and transparent use of the underlying hardware and provide all functionality required by applications for I/O and checkpointing.

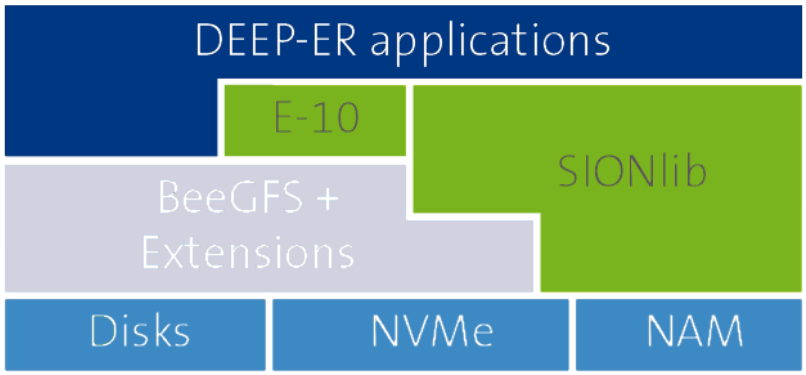

Figure 5: DEEP-ER I/O software layers.

1.3.2.1.1 BeeGFS

As a global parallel file system BeeGFS provides a solid, common basis for high-performance, parallel I/O operations. Beyond the advantages and new functionalities achieved when combining it with the rest of the DEEP-ER software, BeeGFS has itself been significantly improved during the DEEP-ER project. Five new functions have been developed: a local cache layer in the file system, the capability to define the striping size at user level, a feature for following symbolic links, cyclic redundancy check (CRC) checksum calculation, and native support for the EXTOLL network. The cache domain – based on *BeeGFS on demand* (BeeOND) – stores data in the fast node-local NVM devices; it can be used in a synchronous or asynchronous mode, and is selected via a particular directory path. Thanks to BeeOND, the frequency of accesses to the global storage is reduced, increasing the overall scalability of the file system. Letting users define the striping size, based on their knowledge of the application data layout, allows to efficiently determine the best physical distribution of the data in the storage media. With the new follow-symlink feature the user can

decide whether a symlink should be created in the cache/global file system or a copy of the destination file is required. The integration of the CRC checksum calculation avoids additional reads from the cache/global file system. And finally, the full performance is obtained when transferring data through the EXTOLL interconnect thanks to the new native support for the EXTOLL protocol.

1.3.2.1.2 SIONlib

On top of BeeGFS, SIONlib acts as a concentration-layer for applications to most efficiently use the file system. SIONlib compacts all the data that applications performing task-local I/O need to store into one or very few large files. Without SIONlib each task of such an application would save its own file, together generating a huge amount of small files and easily overloading the file system. Measurements done by the DEEP-ER applications have demonstrated the performance advantage of this approach. SIONlib plays a key role in DEEP-ER, building a bridge between the I/O and resiliency components of the software stack. It supports multi-level checkpointing including the efficient storing of data in the global file system and the copying of local data into the fast NVM of a companion (“buddy”) node for redundancy. Both functions work in combination with SCR. SIONlib has been also adapted to use libNAM to trigger parity computation on the NAM, allowing applications to transparently use this new feature for checkpointing without modifying the application code. Furthermore, the communication layers of SIONlib have been refactored and improved to eliminate code replications and increase its overall modularity and manageability.

1.3.2.1.3 Exascale 10 (E10)

Last but not least, E10 addresses the so-called small I/O problem. Here the ever-increasing level of parallelism results in a huge number of small I/O operations in order to obtain global access to file data in Exascale parallel applications. E10 includes an optimised version of ROMIO, which improves the performance of applications relying on the MPI-IO API. Furthermore, E10 supports the caching functionality of BeeGFS, as well as user-level data-striping. The integration of E10 into the applications is seamless thanks to its support library, which catches MPI-I/O calls without the intervention of the user or modifications of the application code.

1.3.2.2 Resiliency Software

The DEEP-ER resilience architecture is based on user-level application checkpoint/restart – which is a proven technique providing a high level of resiliency and avoiding the I/O overheads caused by OS-driven system checkpointing. The project complements this approach with novel task-based recovery techniques (see Figure 6). With this combination, DEEP-ER has developed new resiliency features that isolate partial failures of the system without requiring a full application restart, resulting in a more resilient, fine-grained and flexible architecture. The DEEP-ER resiliency software stack integrates the capabilities of the I/O software packages described in Section 1.3.2.1 with the functionality of SCR and ParaStation MPI and OmpSs, which are the two main components of the DEEP/-ER programming environment. As a result, the scalability of application checkpointing is substantially improved.

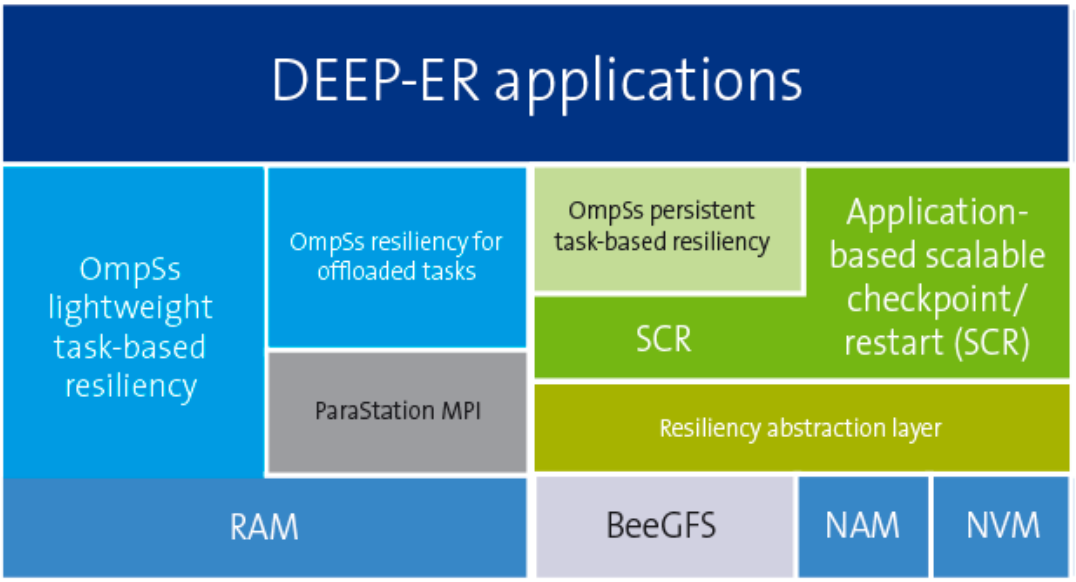

Figure 6: DEEP-ER resiliency layers

1.3.2.2.1 Multi-level checkpointing/restart: SCR

The Scalable Checkpoint-Restart library SCR offers a flexible interface for applications to perform checkpoints and restart from them in case of failure. The user only needs to call the library and indicate the data that needs to be stored for the application to be able to restart. SCR decides itself, based on a failure model (Section 1.3.2.2.3), where and how often checkpoints are performed, and keeps a database of checkpoints and their locations in preparation for eventual restarts.

In DEEP-ER multi-level memory hierarchy, data can be stored – ordered from fastest but smallest capacity to slowest but largest capacity: locally into the node's attached NVM device, into the NVM device of another node, and into the global file system's storage. SCR employs the capabilities of BeeGFS and SIONlib to most efficiently save data at each of the levels.

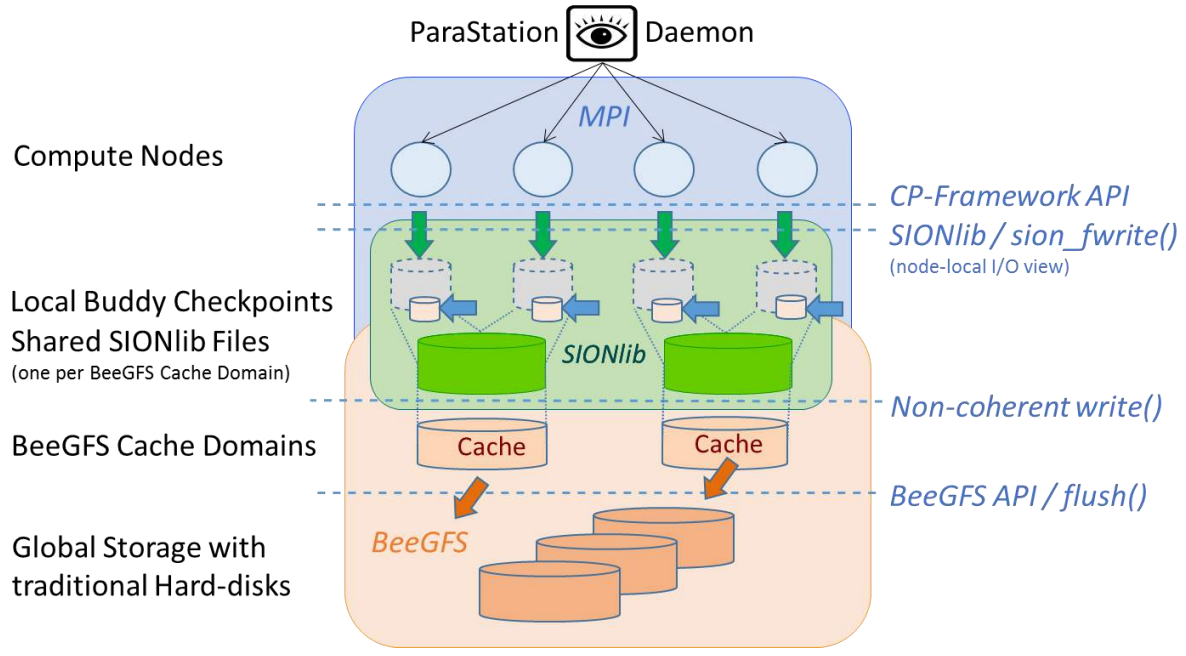

Figure 7: Buddy-checkpointing software stack.

For example, the DEEP-ER buddy-checkpointing functionality (Figure 7) combines SCR, ParaStation MPI, SIONlib, and BeeGFS to save data not only on the node-local storage but also on the NVM device of a companion node. In case of hardware failure, the data lost on the damaged node can be recovered from its “buddy” and the application can continue with minimum time-loss. SCR keeps track of the association between host nodes and buddies; SIONlib takes care that all MPI processes running on a single node jointly write their checkpoint-date into a single file in the buddy-node; and BeeOND saves the data itself in the cache file system residing on the NVM device, transferring it later to permanent global storage.

Furthermore, SCR and SIONlib together can use libNAM to access the NAM memory pool and perform checkpointing/restart to/from the NAM.

DEEP-ER’s work on SCR has been performed in a branch of the original SCR code, spending significant effort in re-organising and cleaning it to allow a better integration of the DEEP-ER modifications. The ideas developed in DEEP-ER have been presented to and discussed with the main developer of the SCR library, Adam Moody, who expressed his interest in future development of DEEP-ER's SCR modifications.

1.3.2.2.2 Task-based resiliency: ParaStation MPI + OmpSs

The runtime of the task-based programming environment OmpSs has been extended in DEEP-ER in order to automatically re-start tasks in the case of transient hardware failures. Failed OmpSs tasks can now be restarted at their last checkpoint, without losing the work that had been performed in parallel by other OmpSs tasks of the same code. Several use cases and applications have demonstrated the functionality and very low overhead of the approach.

To support this task-based resiliency functionality, the ParaStation management daemon has been extended to provide an interface for querying resiliency-related status information from the MPI layer and thus also from the OmpSs runtime environment. ParaStation MPI itself has been enhanced and is now able to detect, isolate and clean up failures of MPI-offloaded tasks, which can be then independently restarted, avoiding the need of a full application recovery.

1.3.2.2.3 Failure model

Finally, a failure model has been developed in the form of an event based Monte Carlo simulation to optimise policies that for each application determine the frequency, redundancy level and storage-location of each checkpoint. This model takes into account the probability and type of failures, the performance of the user-based and task-based checkpoint/restart implementations on the DEEP-ER hardware architecture, as well as application specific characteristics. A closed formula, developed to validate the failure model, has been integrated with the SCR library enabling it to decide on the required checkpointing frequency for each application.

1.3.2.3 Benchmarking software: JUBE

The Jülich Benchmarking Environment (JUBE) was used to perform benchmarking activities that complemented the I/O and resiliency software developments and served to evaluate their performance and functionality. Benchmarks, proxy-applications, and full-fledged applications have been integrated into JUBE. With this activity the parameters required in

BeeGFS for an optimal performance of metadata handling have been determined, and issues on the network performance on the SDV were identified and finally solved.

1.3.2.4 Performance analysis and modelling: Extrae/Paraver/Dimemas

The BSC performance analysis and modelling tools were extended to handle I/O operations, including time, duration, file handle and data volume into the Extrae event trace-files. This allows for detailed analysis of I/O behaviour and performance. On the performance modelling side, a method to calculate I/O efficiency has been devised, which can be combined with the multiplicative performance model used for compute and network characterisation.

These tools have been used in DEEP-ER to predict the behaviour of applications running on a DEEP-ER-like system of much larger scale to assess the potential of the DEEP-ER architecture for Exascale class systems.

1.3.3 Applications

The application developers play a crucial role in the DEEP-ER project. Their work is two-folded: on the one hand their co-design input drives the development and future of both hardware and software architectures; on the other hand, they validate the work done by other technical work packages by testing it with their codes.

During three and half years the application developers have been working side-by-side with a highly skilled and interdisciplinary support team. Experts on the Intel Xeon and Xeon Phi architecture, application optimisation, I/O and resiliency software, programming models, and performance analysis tools have guided them on how to improve their codes through training workshops and frequent personal communication. This intensive support has enabled the applications to achieve remarkable performance and scalability improvements, which are not only valid on the DEEP-ER System but also carry over to most modern HPC platforms.

During the initial phase of the project, the applications were thoroughly analysed, studying their structure to identify the best code distribution between Cluster and Booster, the best opportunities for I/O optimisation, a strategy for resiliency support as well as for their general optimisation. The defined strategies where application-specific, as all applications are different in their nature, algorithms and purpose. All applications initially used MPI, but only few of them had been threaded. In DEEP-ER the developers became aware of the advantages of an optimal hybridisation of their codes, either based on OpenMP or on OmpSs, and they have identified and integrated the most suited threading strategy into their software. Additionally, the codes have been ported to Intel Xeon Phi platforms and various improvements have been made to better exploit this many-core architecture. Two examples are aligning the data for better automatic vectorisation by the compiler, and improving data locality for a more efficient use of the cache memory.

Furthermore the application developers have improved the I/O performance and resiliency of their codes by applying the different techniques developed in the DEEP-ER project. Which software packages each of them specifically employed was determined taking into consideration the intrinsic nature and characteristics of the application codes. The varied selection of applications made it possible to test all the I/O and resiliency strategies of the project demonstrating the advantages that they bring to real HPC developers, and to identify for which kind of codes is each of them best suited.

Each application has made a big step forward from various perspectives:

- KULeuven started with a complex implementation of the particle-in-cell application iPic3D, in which it was hard to introduce software changes and to foresee their overall impact in terms for performance. With the development of a mock-up a much better analysis of performance bottlenecks was possible and different software implementations could be tested, which served to identify the most performant approaches and opened the door for a wide range of code improvements. For instance, a new field solver has been developed: it takes advantage of the PETSc library and scales much better than the original one. Also the particle mover was significantly optimised. A new application (xPic) came out of this effort, which takes over all scientific capabilities of the original iPic3D but executes them much more efficiently on modern HPC platforms. The xPic code will now become the next generation particle-in-cell code for space weather simulations. xPic runs the particles simulation on the Booster and offloads the field-solver to the Cluster part of the DEEP-ER system architecture. The code profits from using SIONlib and exploiting the local NVMe devices for near-node storage. Performance gains running on Intel Xeon Phi, when compared to other architectures such as Xeon and GPGPU, have also been demonstrated. The resiliency strategy of application based checkpoint-restart was extended by the integration of SCR.
- The quantum Monte-Carlo TurboRVB application from CINECA on High temperature superconductivity has seen impressive improvements in the I/O performance when using SIONlib to execute task-local I/O. The impact of SIONlib improvement is highest when applying checkpointing, what by itself improves the application's resiliency to hardware failures. This resiliency strategy was also improved by using SCR. TurboRVB's code complexity was also challenging with regards to software changes and in particular when trying to introduce multi-threading. Learning from KULeuven's experience, a mock-up code called 2degas was also introduced, which runs and achieves good performance with OpenMP.
- GERShWIN – an application from Inria to study the impact of electromagnetic fields in human tissues – experienced remarkable code improvements. Here, only some of them are highlighted: lower cache-miss rate and better thread balancing thanks to better data locality; minimal communication latency by using non-blocking MPI calls; speed-up thanks to better data alignment and exposed multithreading in packing/unpacking of send/receive buffers, resilience to failures thanks to multi-level checkpointing, etc. Using the OmpSs offload model GERShWIN offloads the I/O intensive parts and non-threadable parts of the application to the Cluster, while the rest of the code exploits the high performance and scalability of the Booster. The integration of SIONlib into the application phases that perform intensive process-local outputs provided an up to 3.6x speed-up. Performance gains thanks to the use of the local NVMe for checkpointing and out operations have been also observed. Additionally the application was made resilient by integrating the OmpSs persistent task-based checkpointing and SCR.
- The SeisSol simulation of rapid crustal deformation & earthquake source dynamics from BADW-LRZ was already achieving good performance in standard HPC systems. However adaptations were needed to fully exploit many-core architectures such as Intel Xeon Phi. The main bottleneck in the code stemmed from several small element-local matrix-matrix multiplications with diverse sparsity patterns, which depend on the

order of the basis-functions and the physical use case considered. The desired performance improvements were achieved by applying a three-way optimisation strategy: (1) exploit the prior knowledge of limited number of sparsity patterns involved in the sparse-dense matrix multiplication kernels, (2) introduce OpenMP threading, and (3) improve the I/O behaviour. Furthermore, SIONlib-based I/O for checkpointing and restart was implemented, and the code uses the MPIWRAP library of E10 to transparently use the NVMe as fast storage. Also the checkpoint-restart strategy was extended by using SCR. SeisSol is now an Open Source development and close collaboration with the main development team will encourage the integration of the improvements done within DEEP-ER into the main branch of the code.

- ASTRON's application for analysis of radio astronomy data combines two components: a correlator and an imager. The computational core of the correlator was first heavily optimised for Intel Xeon Phi, achieving almost 80% of the processors' peak performance. Most of the effort in the project was therefore focused on the imager part of the pipeline, which was first ported to Xeon Phi. The gridding and degridding parts are usually the most time consuming parts in an imaging run, so much effort was put into optimising these operations for the Xeon and Xeon Phi. A "mock-up imager" based on a new gridding/degridding library was introduced for easier testing of DEEP-ER specific features such as the local NVMe and the NAM. Additionally, a better communication strategy was defined by analysing five different solutions based on point-to-point communications vs. collective communications, blocking vs. non-blocking principles, and combinations of both. Finally, SCR and MPI fault-tolerance have been integrated to improve the resilience of both the imager and the correlator applications.
- The Full Wave Inversion (FWI) seismic imaging application from BSC was studied and a mock-up code that very well replicates the real execution was developed for better manageability. FWI processes the seismic shots on the Booster and uses the Cluster for managing the Booster-workers and aggregate their results. OmpSs is employed for intra-node task handling and data distribution (task offloading), while MPI is used for data processing. Low-level optimisations of the Booster part of the code were done, including manual vectorisation and employing Mercurium's features like prefetching, which provide a portable way to generate highly optimised code. In total, a speedup of about 11.4 could be reached compared to the baseline version of the application. Finally, using OmpSs task-based resiliency for offloaded tasks achieved application resiliency with very low overhead.
- The Chroma simulation framework for lattice QCD developed by UREG runs fully on the Booster. Performance and scaling improvements on Intel Xeon Phi were achieved by using OpenMP, and increasing vectorisation, both manually (intrinsics) and automatically with compiler options. The latter was possible thanks to the development of a new "LibHadronAnalysis" library. The code's limiting factor for large production runs was parallel I/O. Modifications in Chroma's original MPI rank structure enabled the use of HDF5 output files and showed speed-ups compared to usual binary output. The implementation allows easy adjustment to all hardware architectures, it can fully utilize parallel file systems and it can be further improved by e.g. asynchronous writing while continuing the computation. On top of HDF5 version

the MPIWRAP library is used to exploit the SSD cache features provided by E10 using the NVMes.

Altogether, the application results do show the benefits of DEEP-ER's I/O and resiliency strategies embedded into a highly flexible architecture. Furthermore, the applied modifications brought improvements in performance and code-manageability that maintain the application's portability and outperform the original codes not only on DEEP-ER hardware systems, but also on standard HPC platforms.

1.3.4 Co-design

The innovative solutions developed in DEEP-ER would not have been possible without a stringent co-design approach taken throughout the whole project's timeframe and across all its components: hardware, system software, and applications. The demonstrated importance and impact of such intensive collaboration is a result of the project in itself.

The body driving the co-design discussions in DEEP-ER is the Design and Development Group (DDG), in which the project experts jointly determine the high-level design of the project developments through discussions taking place in the DDG bi-weekly teleconferences. The DDG decides on all proposals that have impact across work packages, and reaches agreements on all open questions of such nature. In this way, the DDG allows for a coherent development of all software and hardware components of the project, and their proper integration with each other. For topics that require a more deep analysis and discussion, specific cross-WP and co-design activities have taken place.

Face-to-face meetings of the whole consortium took place twice per year and were the scenario of further detailed co-design discussions. One third of the available time was reserved for the application developers to described their codes and present the results recently obtained, as well as the planned next steps. Lively discussions took place during and after the applications' presentations as the rest of the team (hardware and software experts) used this opportunity to learn about the codes and ask questions on their particular requirements, which constituted very important co-design input.

The focus of co-design activities varied along the project duration. At the very beginning of the project, when the hardware and software designs were being shaped, the most important aspect was to gather the main application requirements, for which purpose a detailed questionnaire was formulated. The evaluation of their answers and further detailed discussions on each of the subtopics followed up. Hardware and system software developers adapted their original design sketches accordingly, to make sure that they addressed all the collected application requirements. At this stage, some aspects that influenced the hardware design were the amount of memory needed both in total and per single node, the minimum required bandwidth or the communication patterns and preferred network topology.

In the mid-phase of the project hardware and system software were developed, while application codes were improved and adapted to the project's needs. Trainings were offered to the latter to facilitate the code-porting. At this stage hardware-software and software-software co-design were dominating. A particularly close collaboration was established between the I/O and resiliency software developers, which led to a tight interleaving of software packages in an optimal combination of each other's capabilities. Without this very intense collaboration between the teams functionalities such as buddy-checkpointing that combines a hardware component (NVMe) with four individual software packages (SCR,

ParaStation, SIONlib and BeeGFS) would not have been possible. During this implementation phase also questions arose and were addressed regarding the API or interface preferred by the applications for the use of the various software components.

In the final stage of the project software and hardware platforms became fully available to the users. These were instructed by the respective developers on the possible solutions for each use case, on how to implement available functionality, and how to adapt the application codes to best exploit them. During these test, validation and benchmarking activities bugs were identified and reported to the developers. For instance, issues with the OmpSs offload functionality were identified when testing FWI and the found bugs were solved. From the hardware perspective the application results brought the final answers with regard to the advantages and caveats of each of the system components, and on potential improvement for future evolution of the system architecture.

1.4 Potential impact and main dissemination and exploitation

1.4.1 Potential impact

DEEP-ER has created far-reaching impact beyond its focus areas of high-performance I/O and system resiliency. Its prototype solutions address a wide range of HPC topics such as scalability and end-to-end performance, energy efficiency and ease of programming. The results cut across the complete HPC stack and amount to a fully integrated system prototype combining hardware with system software, programming environments and highly tuned applications. Fully embracing the co-design approach has made this possible. It led to widespread innovation and substantially reinforced the position of European industry and academia in HPC and demonstrated the strength of both players working together and the mutual gains achievable in such close collaboration. This approach will continue in DEEP-EST and serves as a reference for upcoming HPC projects.

One must not forget that the success of a project is significantly influenced by the quality of its members. The DEEP-ER team is composed of a group of highly engaged and motivated, technically and scientifically excellent people, most of them in the early stage of their careers. Adding to that, training sessions and workshops have been organised in the project to further promote them. This, together with the very tight contact between experts of different disciplines that DEEP-ER has fostered, has provided the team with valuable experience and networking contacts, turning them into a group of world-class experts that are in high demand. Depending on their personal interests and individual paths, these individuals will play their part in shaping the future HPC landscape.

1.4.1.1 Strengthen the European industry supplying and operating HPC systems

The three leading European HPC technology enterprises Eurotech, EXTOLL and ParTec together with Intel were responsible for designing, manufacturing, and installing the integrated DEEP-ER HW/SW prototype. The success of DEEP-ER is a testimonial on the expertise of these companies to create highly innovative HPC solutions. DEEP-ER's success will boost the market position of the three companies and open up new opportunities.

The DEEP-ER Prototype was designed and engineered to maximize system density and energy efficiency, provide highly efficient direct liquid cooling and high system reliability. The result is a new generation in Eurotech's line of Aurora systems, integrating the latest generation of server (Intel Xeon) and manycore processors (Intel Xeon Phi) into a blade-

based system with a Backplane that Eurotech designed to run at full PCI Express generation 3 speeds and to support up to two PCI Express extension cards per blade. All components, including the EXTOLL TOURMALET network cards and the NVM storage cards, are either directly liquid cooled by cold plates or attached to the liquid cooling circuit. The DEEP-ER prototype system achieves its density by a very innovative and space-efficient cold plate design and by innovative integration of ultra-low profile (ULP) DIMMs with direct liquid cooling. These unique innovations will create a sizeable competitive advantage for Eurotech, who are amongst the world-wide pioneers of direct liquid cooling for HPC.

The University of Heidelberg realized in DEEP-ER a first proof of concept for the innovative “network attached memory” (NAM) idea. While this prototype was necessarily limited in scale, it will lead to successive implementations (judging by interest of many of today’s leading memory players) and has the potential to lead to a rethinking of memory and storage architectures in HPC and data analytics.

DEEP-ER for the first time uses the new, ASIC based network controller designed and produced by EXTOLL in Europe. In addition to the best-in-class messaging and remote access functionality already demonstrated in DEEP, TOURMALET delivers 100 Gbit/s bandwidth for each of its six network links. In effect, TOURMALET is by far the leading European interconnect technology, and DEEP-ER constitutes a key proof point for TOURMALET that will greatly improve its market position.

1.4.1.2 European software innovation for scalable I/O and resiliency

Key elements of the DEEP-ER software stack adapted and extended from DEEP are the ParaStation MPI library and the OmpSs task-based programming model. The former now supports a clean and automatic restart of applications in case of HW failures, and the later was enhanced by support for restarting failed tasks. Both capabilities shall improve the resiliency of applications with minimal effort on the user side, making the software packages ParaStation and OmpSs even more attractive to users. This, added to their support for hardware heterogeneity, prepares it for the high complexity of future Exascale systems.

The integration of non-volatile, fast local storage at the compute nodes has resulted in substantial increases in application performance and system throughput. Key elements here were BeeGFS as the transparent file-system interface used by SIONlib and E10 to provide truly scalable I/O performance. This success will accelerate the take-up of local storage in HPC and Big Data systems. In combination with advances in storage class and non-volatile memory as announced in the meantime by vendors, this has the potential to create a breakthrough in supporting data-intensive science and engineering.

The BeeGFS filesystem is at the core of DEEP-ER’s I/O software stack. It provides a Posix-compliant interface and has been extended in DEEP-ER to make best use of the node-local fast memory to minimize global storage accesses and further increase I/O scalability. BeeGFS is distributed as Open Source by the Fraunhofer Institute and commercially supported by its spin-off company ThinkParQ. The results of DEEP-ER will enable BeeGFS to quickly take advantage of future node-local memory technologies, as well as strengthen its global market position.

E10 by Seagate, SIONlib and SCR by JUELICH (all available as Open Source) have gained valuable functionality and were optimised for the I/O patterns exhibited by the co-design applications. The integration with BeeGFS into a coherent I/O and resiliency system SW

stack provides significant value above and beyond the sum of its parts to future HPC and data analytics systems and end users.

JUELICH's Open Source benchmarking environment JUBE has proven its worth in managing the real-world DEEP-ER codes and enabling and managing an ambitious benchmarking and measurement campaign. Also the BSC tools Extrae/Paraver/Dimemas gained new functionalities and users within the DEEP-ER project, strengthening their position as one of the best established toolkits for performance analysis and modelling of applications in HPC.

1.4.1.3 Foster European excellence in HPC simulation codes

Seven relevant, real-world European applications in important scientific and engineering fields did drive the co-design in DEEP-ER. Their unique requirements were thoroughly analysed, the pre-existing code was modernized to significantly improve performance and scalability, and adapted for efficient use of the DEEP-ER I/O and resiliency interfaces as well as for the Cluster-Booster partitioning.

The resulting applications will now enable faster and more energy- and cost-effective scientific discovery and better engineering solutions, benefitting European research and industry alike. Since the DEEP-ER software interfaces are based on standards and anchored in long-lasting development teams, they will be adapted to future heterogeneous platforms, enabling the seven applications to take advantage of such new systems. In addition, the codes continue to run on conventional architectures, sometimes showing surprising performance and efficiency improvements compared to their former formulation. This makes evident that adapting a code for DEEP-ER is equivalent to preparing it for new generation HPC platforms.

Even more importantly, the expertise generated in DEEP-ER was distilled into a "best known methods" playbook for modernizing a wide range of applications, ultimately leading to a profound beneficial effect on the entire HPC application ecosystem.

The DEEP-ER Prototype is planned to remain operational at Jülich for at least two years. It will remain available for the DEEP-ER users as long as it is operational so that the potential for achieving further scientific results can be fully exploited in the future. Furthermore, the use of the platform will be opened for external users. A Memorandum of Understanding (MoU) has been put in place between the DEEP-ER and the PRACE-4IP projects for that purpose.

1.4.1.4 Reinforce cooperation in international endeavours on Exascale

The contributions from academic partners in managing the complex DEEP-ER project and its co-design cycle, and in bringing in cutting-edge hardware and software innovations were absolutely critical for the ultimate success. The excellence shown here sets these institutions (JUELICH for the overall project, BSC for the resiliency software, and BADW-LRZ for the dissemination of the project results) up as prime partners for the next round of HPC research and demonstration projects in Europe. But the project has made an impact not only within the continent, but well beyond.

JUELICH, ParTec and Intel established already in 2010 the ExaCluster Laboratory (ECL) a long-term collaboration that aims at jointly investigating approaches to improve the scalability, energy efficiency and performance of cluster computing. The DEEP projects family (including DEEP, DEEP-ER and soon DEEP-EST) are at the core of the ECL and

enable it to spread and extend the know-how to other European partners. The long-term nature of the ECL collaboration ensures that the results and experiences gathered within the projects will be taken into consideration, exploited and disseminated in the future, not only within the European borders but worldwide. For this purpose, the collaborations of the three partners with international institutions will be exploited.

Through the partners' participation in international research and innovation collaborations and initiatives, the results obtained in DEEP-ER have reached a wider visibility. Examples are the Joint Laboratory for Extreme Scale Computing (JLESC), in which JUELICH, BSC and Inria are members, and the Big Data and Extreme Scale Computing (BDEC), in which a large amount of the DEEP-ER consortium is represented. Worth mentioning is also the strong involvement of many of the DEEP-ER partners in ETP4HPC.

Finally, DEEP-ER members have participated in several international conferences and workshops in Europe, America and Asia to present their ideas and results. The interest raised in the audience has been remarkable and collaboration opportunities arouse. New research projects and industrial collaborations will come up in the next years, which have their roots in the DEEP-ER project.

1.4.1.5 The next step: DEEP-EST

The DEEP-EST Project already stands on the shoulders of DEEP-ER and generalizes the Cluster-Booster approach to create a truly Modular Supercomputing system. It combines three compute modules with distinct architectures adapted to highly scalable and general-purpose HPC codes, as well as data analytics and data-intensive applications. This will make the Modular Supercomputing architecture fit to best support workloads that emerge from the confluence of HPC and Big Data, bringing tangible benefits in system throughput and energy efficiency to large computer centres.

1.4.2 Main dissemination activities

The DEEP-ER project and its results have been presented and discussed in many meetings, conferences, and workshops, in the form of presentations, invited talks, panel sessions, posters, papers, and publications. Particular effort has been invested in being present and active on the two main events of the worldwide HPC community: the Supercomputing Conference (SC) and the International Supercomputing Conference (ISC), which take place every year in the US and Germany, respectively. The DEEP-ER activities on ISC and SC ranged from presentation of results in talks and posters to the presence at the exhibition floor in a joint research booth passing by the organisation of Birds of a Feather (BoF) sessions and workshops as well as participation in panel discussions, and presentations at other vendors' booths. This strong presence has only been possible joining efforts with the European projects DEEP, Mont-Blanc, CRESTA, EPiGRAM, EXA2CT, NUMEXAS, and more recently some of the H2020 FET projects such as NEXTGenIO. The DEEP and DEEP-ER projects drove the European Exascale Projects (EEP) initiative for four and half years, until the coordination of joint dissemination activities was taken over by the EXDCI project. On top of the SC and ISC participations, PRACEdays and the more recently introduced European HPC Summit Week have been important outreach events. The DEEP projects have been present with talks at these events and benefitted greatly from the interaction among the European HPC and Exascale scene. Particularly worth mentioning is the satellite event we hosted together with the EEP at PRACEdays15 in Dublin. We developed and realized a

whole day of sessions with a specifically industry-related focus. The idea was to showcase how (industrial) application developers can benefit from the developments in the European Exascale projects. Also with a clear industry focus the DEEP-ER project seized the opportunity in 2016 to be present at various other conferences and trade fairs like CeBIT or the Forum Teratec. Another highlight was attending and giving a keynote at SAI computing conference 2016, where exchange was possible with other fields of IT.

Figure 8: Joint EEP booth at SC15.

Furthermore, the goals and results from DEEP-ER have been actively disseminated online via the project's website and social media channels. Especially, Twitter worked well for outreach activities. Regular updates with new content have resulted in a steady increase of follower numbers, proof of the large interest that the DEEP-ER ideas have caused in the HPC community.

A detailed list of publications and dissemination activities performed during the DEEP-ER project is given in Section 2 of the present report. Efforts for disseminating the project's results and impact will continue beyond the end of the project. The DEEP projects family continues now with DEEP-EST, which guarantees a seamless communication effort also after the official end of DEEP-ER. Further scientific results are expected to be published after the end of the project. In particular, opening the use of the DEEP-ER Prototype to external users shall generate further results, the publication of which will be encouraged.

The centre of the dissemination activities of DEEP-ER is its Web site: www.deep-er.eu, which will continue being accessible after the end of the project. With the upcoming of DEEP-EST, the existing website will be transformed into a general www.deep-projects.eu website. This adapted website will jointly present the results of all three projects, making the information easily accessible for the wider HPC community.

Training the community on how to use the software and hardware developed in DEEP-ER has also been an important part of the project. The main goal of the training events within the project was to teach the application developers participating in the project on how to use the software tools and programming environment running on the DEEP-ER Prototype and other intermediate software development vehicles.

Now that the project arrived to its end, is the time to offer training on the use of the DEEP-ER concepts to the rest of the HPC community, especially those users coming through the collaboration with PRACE. DEEP-ER users will continue having access to the project hardware, to enable them obtaining the maximum amount of results possible. Presentations,

articles and publications should be the outcome of this effort. Without interfering with these activities, JUELICH has already started evaluating the DEEP-ER hardware also with in-house applications, to use it for training workshops, and to make it available to interested collaboration partners. The goal is here to maximise the amount of results and experience gathered with the hardware systems and software tools developed in the project, increasing in this way the impact of the project.

1.4.3 Exploitation of results

The DEEP-ER project has created remarkable scientific and technical results and know-how that will be exploited by the partners after the end of the project. With this, DEEP-ER significantly contributes to the evolution and implementation of the Strategic Research Agenda (SRA) defined by the European Technology Platform for High Performance Computing (ETP4HPC), reinforcing the worldwide competitiveness of both European industry and research.

The hardware and software results in DEEP-ER were largely achieved by European industrial and academic partners – this goes a long way to ensure that Europe will in the future be able to develop key parts of a HPC infrastructure independently, should this be necessary. The system work performed by Eurotech, University of Heidelberg (and indirectly, their industrial spinoff Extoll GmbH) in close collaboration with Intel will benefit these European players and their ecosystem of design and manufacturing partners in Europe both directly and indirectly. The former impact is constituted by market-leading products in the shape of the DEEP-ER System and components, by having reference installations of these at JUELICH, one of the PRACE hosting members, and by having been able to push key technologies (dense packaging, direct liquid cooling, non-volatile memory, network attached memory…) forward that will be instrumental to stay at the leading edge of the market with future products. Additionally, having gained immense experience in the tight integration of high-end processor, memory, and interconnect technology within a high-speed signalling environment will greatly assist the DEEP-ER technology players in their future product development.

The DEEP-ER software stack developed by the industrial partner ParTec and the leading research centres JUELICH, BSC, and FHG-ITWM is the key for the generalisation of the Cluster-Booster concept to become the Modular Supercomputer architecture. The I/O and resiliency functionality that they now provide constitute an important asset both for the upcoming Exascale computing generation and for the convergence of HPC with HPDA (high performance data analytics) application fields.

Requirements and applications that are important to PRACE have shaped the design of the DEEP-ER software and hardware environments. This, together with the recently established collaboration that will open the use of the project developments to a wider PRACE community, ensures that the results of DEEP-ER will match the requirements of the European Research Infrastructure. The seven highly optimised DEEP-ER co-design applications and the best known methods distilled from the application work in DEEP-ER will give the European Research Area a head-start for making their applications Exascale-ready.

The DEEP project is fully aligned with the ETP4HPC SRA as released in 2013 and its update SRA2 released in 2015. In fact, through the partners' participation in the various working groups the project has significantly been influencing the evolution of this document. The

DEEP-ER results directly address the challenges of extreme scalability, programmability energy efficiency, resiliency, and I/O as identified in the SRA. They contribute innovative European solutions for the R&D priorities of “HPC System Architecture and Components”, “System Software and Management”, “Programming Environment”, “Energy and resiliency”, and “Balance Compute, I/O and Storage Performance”.

The European Research Area (ERA) and industry will directly profit from having access to the installed DEEP-ER hardware systems and to the software stack, with the PRACE centres JUELICH and BSC, together with ParTec, FHG-ITWM, and Seagate being able to provide excellent support. In addition, the methods used and lessons learned in the adaptation/optimisation of the seven DEEP-ER co-design applications will serve to speed up the transition of further European applications to more efficient and scalable formulations. Having such applications at their disposal will, in turn, greatly support domain scientists and engineers in making ground-breaking discoveries faster and developing better products.

In addition, the European Research Area (ERA) and industry will profit by being able to materially influence the further development of the Cluster-Booster Architecture and the DEEP-ER software stack to best match their specific requirements. This joint effort of the DEEP-ER technology partners will continue in DEEP-EST, which aims at bringing the Cluster-Booster Architecture to Exascale era in the form of a Modular Supercomputer, building upon the results achieved in the DEEP and DEEP-ER projects.

1.5 Project website

The DEEP-ER webpage can be reached in the following address: www.deep-er.eu

1.6 Contact

Project Coordinator: coordinator@deep-er.eu

Project Management Team: pmt@deep-er.eu

Twitter channel: @DEEPprojects

LinkedIn group: DEEPprojects

1.7 Project logo

1.8 List of beneficiaries

List of Beneficiaries			
Participant number	**Partner**	**Acronym**	**Country**
1	Forschungszentrum Juelich GmbH	JUELICH	Germany
Third Party under Clause 10 (TP10) from JUELICH	ParTec Cluster Competence Center GmbH	ParTec	Germany
2	Intel Deutschland GmbH	Intel	Germany
TP10 from Intel	Intel Iberia S.A.	Intel Iberia	Spain
3	Bayerische Akademie der Wissenschaften	BADW-LRZ	Germany
4	Ruprecht-Karls-Universitaet Heidelberg	UHEI	Germany
5	Fraunhofer-Gesellschaft zur Foerderung der Angewandten Forschung e.V.	FhG-ITWM	Germany
6	Eurotech S.P.A.	Eurotech	Italy
TP10 from Eurotech	ETH Lab SRL	ETH Lab	Italy
7	Barcelona Supercomputing Center – Centro Nacional de Supercomputación	BSC	Spain
8	Seagate Systems UK Ltd.	Seagate	United Kingdom
9	Consorzio Interuniversitario CINECA	CINECA	Italy
10	Katholieke Universiteit Leuven	KULeuven	Belgium
11	Institut National de Recherche en Informatique et en Automatique	Inria	France
12	Stichting Astronomisch Onderzoek in Nederland	ASTRON	Netherlands
13	Universitaet Regensburg	UREG	Germany

2 Use and dissemination of foreground

2.1 A. Dissemination measures

As explained in Section 1.4.2, dissemination and training has been very intensive in the DEEP-ER project. A notable effort has been made to increase the visibility of the results and achievements of the project, in the HPC community and beyond. The following tables display lists of the publications and general dissemination activities performed by the DEEP-ER project within its lifetime. It is expected that the lists will increase with articles about results obtained in the last period of the project or even after its end.

2.1.1 Scientific, peer reviewed publications

Table A1: List of Scientific (peer reviewed) publications, starting with the most important ones										
No.	**Title**	**Main author**	**Title of the periodical or the series**	**Number, date or frequency**	**Publisher**	**Place of publication**	**Year of publication**	**Relevant pages**	**Permanent identifiers[3] (if available)**	**Is/Will open access[4] provided to this publication?**
1	Collective Offload for Heterogeneous Clusters	Sainz, F.	Proc. Of 22nd IEEE International Conference on High Performance Computing (HiPC)		IEEE	HiPC	2015		10.1109/HiPC.2015.20	no
2	Adapting a Finite-Element Type Solver for Bioelectromagnetics to the DEEP-ER Platform	Leger, R.	Proc of ParCo 2015			ParCo 2015	2016		ISBN: 978-1-4673-8488-9/15 DOI: 10.1109/HiPC.2015.20	yes
3	Progress towards Physics-Based Space Weather Forecasting with Exascale Computing	Innocenti, M. E.	Advances in Engineering Software		Elsevier	Netherlands	2016		DOI: 10.3233/978-1-61499-621-7-349	no

[3] A permanent identifier should be a persistent link to the published version full text (if open access or abstract if article is pay per view) or to the final manuscript accepted for publication (link to article in repository).

[4] Open Access is defined as free of charge access for anyone via Internet. Please answer "yes" if the open access to the publication is already established and also if the embargo period for open access is not yet over but you intend to establish open access afterwards

Table A1: List of Scientific (peer reviewed) publications, starting with the most important ones										
No.	**Title**	**Main author**	**Title of the periodical or the series**	**Number, date or frequency**	**Publisher**	**Place of publication**	**Year of publication**	**Relevant pages**	**Permanent identifiers[3] (if available)**	**Is/Will open access[4] provided to this publication?**
4	Improving Collective I/O Performance Using Non-Volatile Memory Devices	Congiu, G.	Proc of IEEE Cluster 2016		IEEE	IEEE Cluster 2016	2016	pp. 120-129	http://www.sciencedirect.com/science/article/pii/S0965997816301363 http://dx.doi.org/10.1016/j.advengsoft.2016.06.011	no
5	Exploring Time and Energy for Complex Accesses to a Hybrid Memory Cube	Schmidt, J.	Proc of Memsys 2016		ACM	Memsys 2016	2016	pp. 142-150	**DOI:** 10.1109/CLUSTER.2016.37 **Electronic ISBN:** 978-1-5090-3653-0 **Print on Demand(PoD)** ISBN: 978-1-5090-3654-7 **Electronic ISSN:** 2168-9253	yes

Table A1: List of Scientific (peer reviewed) publications, starting with the most important ones										
No.	**Title**	**Main author**	**Title of the periodical or the series**	**Number, date or frequency**	**Publisher**	**Place of publication**	**Year of publication**	**Relevant pages**	**Permanent identifiers[3] (if available)**	**Is/Will open access[4] provided to this publication?**
6	Image-Domain Gridding on Graphics Processors	Veenboer, B.	Proc of IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)			IPDPS '17	2017		not yet published	No
7	Exactly energy conserving semi-implicit particle in cell formulation	Lapenta, G.	Journal of Computational Physics	Volume 334	Elsevier		2017		not yet published	No
8	Supporting Automatic Recovery in Offloaded Distributed	Peña, A.	Prof of ISC 2017			ISC 2017	2017		http://dx.doi.org/10.1016/j.jcp.2017.01.002	tbd

Table 1: A1. List of Scientific (peer reviewed) publications.

2.1.2 Dissemination activities

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
1	Website	BADW-LRZ			DEEP-ER Website Up and Running	25/09/2013	Munich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	http://www.deep-er.eu/press-corner/news/2-deep-er-website-up-and-running.html
2	Conference	N. Eicker (JUELICH)		Human Brain Project Summit (HBP) 2013	DEEP and DEEP-ER: Boosters for HPC	09/10/2013	Lausanne, Switzerland	HPC Community Scientific Community	n/a	Europe	https://www.humanbrainproject.eu/hbp-summit-2013-overview
3	Press Release	JUELICH			Mit DEEP-ER noch schneller zum Exascale Rechner	09/10/2013	Jülich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	Germany	http://www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/DE/2013/13-10-09-DEEPER.html English version: http://www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/EN/2013/13-10-09-DEEPER.html;jsessionid=E4A1055123C4026C5158A08BF56BB320

[5] A drop down list allows choosing the dissemination activity: publications, conferences, workshops, web, press releases, flyers, articles published in the popular press, videos, media briefings, presentations, exhibitions, thesis, interviews, films, TV clips, posters, Other.

[6] A drop down list allows choosing the type of public: Scientific Community (higher education, Research), Industry, Civil Society, Policy makers, Medias, Other ('multiple choices' is possible).

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
4	Conference	N. Eicker (JUELICH)		European Research and Innovation Conference (ERIC 2013)	The DEEP-ER Project - Extending the Reach of the Cluster-Booster Architecture	23/10/2013	Nice, France	HPC Community Scientific Community	n/a	Europe	
5	Flyer	BADW-LRZ		SC13	DEEP-ER Flyer: first version	01/11/2013	Munich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	http://www.deep-er.eu/files/DEEP-ER_Flyer_SC13.pdf
6	Flyer	BSC		SC13	European Exascale Projects Flyer - first developed for SC13 (talking about CRESTA, DEEP, Mont-Blanc)	01/11/2013	Barcelona, Spain	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	
7	Publication	JUELICH	Exascale Newsletter JSC		Safer and Faster with DEEP-ER	01/11/2013	Jülich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	Europe	http://www.fz-juelich.de/SharedDocs/Downloads/PORTAL/EN/publications/exascale-newsletter/exascale_nl_03_2013.pdf?__blob=publicationFile
8	Publication	JUELICH	Newsletter JSC		Start of the Exascale Projects DEEP-ER and Mont-Blanc 2	01/11/2013	Jülich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	Europe	http://www.fz-juelich.de/SharedDocs/Meldungen/IAS/JSC/EN/2013/2013-11-exascale-projects.html?nn=897918

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
9	Exhibition	European Exascale Projects		SC13	Joint European Exascale Booth at SC13 (organised for 1st time); booth #3741 DEEP displayed first hardware	17/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	http://iebms.heiexpo.com/sc/SC13Floorplan.pdf
10	Presentation	H.-CH. Hoppe (INTEL)		SC13	DEEP and DEEP-ER Architecture (demonstration at the Intel booth)	18/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Media	n/a	International	
11	BoF	European Exascale Projects		SC13	Building on the European Exascale Approach	19/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	
12	Presentation	N. Eicker (JUELICH)		SC13	DEEP and DEEP-ER: Innovative Cluster Architecture for Intel Xeon Phi (held at the BoF)	19/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	
13	Presentation	N. Eicker (JUELICH)		SC13	DEEP and DEEP-ER: Innovative Cluster Architecture for Intel Xeon Phi (held at Intel Theatre / Intel Booth)	19/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
14	Conference	E. Suarez (JUELICH)		SC13	Emerging Technologies and Big Data (Euro-Centric): Participation in Panel Discussion	21/11/2013	Denver, Colorado	Scientific Community HPC Community Industry Policy Maker	n/a	International	
15	Media Briefing / Interview	E. Suarez + N. Eicker (JUELICH)	Computer World	SC13	Media Briefing: Computer World	21/11/2013	Denver, Colorado	Media	1	International	
16	Workshop	JUELICH			DEEP-ER Co-Design Meeting	24/01/2014	Jülich, Germany	Project Internal	n/a	Europe	
17	Workshop	BSC		13th VI-HPS Tuning Workshop	13th VI-HPS Tuning Workshop	14/02/2014	Barcelona, Spain	HPC Community Scientific Community	n/a	Europe	http://www.vi-hps.org/training/archive/tws/tw13.html
18	Presentation	E. Suarez (JUELICH)		Joint European Exascale Projects Workshop	DEEP and DEEP-ER	18/03/2014	Edinburgh, Scotland	Project Internal	n/a	Europe	
19	Presentation	N. Eicker (JUELICH)		Joint European Exascale Projects Workshop	Global MPI and the DEEP Programming Model	18/03/2014	Edinburgh, Scotland	Project Internal	n/a	Europe	
20	Presentation	J. Labarta (BSC)		Joint European Exascale Projects Workshop	OmpSs	18/03/2014	Edinburgh, Scotland	Project Internal	n/a	Europe	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
21	Presentation	B. Mohr (JUELICH)		Joint European Exascale Projects Workshop	Scalasca	18/03/2014	Edinburgh, Scotland	Project Internal	n/a	Europe	
22	Workshop	J. Labarta (BSC)		G8 ECS internal workshop, K-computer facility	Behind DEEP and Mont-Blanc	26/03/2014	Kobe, Japan	HPC Community Scientific Community	n/a	International	
23	Publication	JUELICH	inside		Going DEEP-ER into Exascale	01/04/2014	Stuttgart, Germany	Scientific Community HPC Community	n/a	Europe	http://inside.hlrs.de/editions/14spring.html#projects3
24	Presentation	S. Narasimhamurthy (Seagate)		LUG 2014	Collective I/O for Exascale I/O Intensive Applications	09/04/2014	Miami, US	Scientific Community HPC Community Industry	n/a	International	http://opensfs.org/lug14/ Slides are available here: http://www.opensfs.org/wp-content/uploads/2014/04/D2_S21_CollectiveIOforExascaleIOIntensiveApplications.pdf
25	Videos	S. Narasimhamurthy (Seagate)		LUG 2014	Collective I/O for Exascale I/O Intensive Applications	09/04/2014	Miami, US	Scientific Community HPC Community Industry	n/a	International	YouTube video: https://www.youtube.com/watch?v=uSo0K82E3e0
26	Publication	JUELICH			Neues EU-Projekt DEEP-ER	15/04/2014	Stuttgart, Germany	Scientific Community HPC Community	n/a	Germany	https://gauss-allianz.de/de/infobrief/382-infobrief-nr-26#beitrag5

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
27	Flyer	BSC		ISC'14	Joint European Exascale Projects Flyer - Updated for ISC'14	01/05/2014	Barcelona, Spain	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	
28	Website	BADW-LRZ		ISC'14	Joint Website for European Exascale Projects set up	01/05/2014	Munich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	http://exascale-projects.eu/
29	Workshop	Fraunhofer ITWM		Transtec Petabyte Workshop	BeeGFS used in DEEP-ER	14/05/2014	Tübingen, Germany	HPC Community Industry Scientific Community	n/a	Europe	http://www.transtec.de/petabyte/
30	Social Media	BADW-LRZ			DEEPprojects: Official Launch of DEEP Twitter Account	15/05/2014	Munich, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	https://twitter.com/DEEPprojects
31	Workshop	E. Suarez (JUELICH)		PRACEdays14 Satellite Event: 'Workshop on Exascale and PRACE prototypes'	DEEP and DEEP-ER	19/05/2014	Barcelona, Spain	HPC Community Scientific Community	n/a	Europe	Agenda: http://www.prace-ri.eu/IMG/pdf/workshop_on_exascale_and_prace_prototypes_-_agenda.pdf

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
32	Flyer	BADW-LRZ		ISC'14	DEEP-ER Flyer: Update (intro features hardware)	01/06/2014	München, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	http://www.deep-er.eu/files/DEEP-ER_Flyer_ISC14.pdf
33	Media Coverage	Th. Lippert (JUELICH)	ISC'14 Blog	ISC'14	Smart Acceleration for Clusters	18/06/2014	Leipzig, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	http://www.isc-events.com/isc14/isc_blog/items/smart-acceleration-for-clusters.html
34	Exhibition	European Exascale Projects		ISC'14	Joint European Exascale Projects Booth	23/06/2014	Leipzig, Germany	Scientific Community HPC Community Industry Policy Maker Media	n/a	International	
35	Media Briefing / Interview	JUELICH / UHEI / Eurotech	insideHPC	ISC'14	Media Briefing with R. Brückner	23/06/2014	Leipzig, Germany	Media	1	International	
36	Presentation	H.-CH. Hoppe (INTEL)		ISC'14	Demonstration at the Intel booth, showing DEEP and DEEP-ER Results	23/06/2014	Leipzig, Germany	HPC Community Industry Scientific Community	n/a	International	
37	Workshop	E. Suarez (JUELICH)		Intel Cross-Lab Workshop	DEEP and DEEP-ER	23/06/2014	Leipzig, Germany	HPC Community Industry Scientific Community	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
38	BoF	European Exascale Projects		ISC'14	Joint European Exascale Projects BoF, titled: The European Approach to Exascale	24/06/2014	Leipzig, Germany	Scientific Community HPC Community Industry Policy Maker Media	75	International	http://www.isc-events.com/isc14_ap/sessiondetails.htm?t=session&o=108&a=select
39	Presentation	N. Eicker (JUELICH)		ISC'14	The DEEP and DEEP-ER Projects - presentation held at European Exascale Projects BoF	24/06/2014	Leipzig, Germany	Scientific Community HPC Community Industry Policy Maker Media	75	International	http://www.isc-events.com/isc14_ap/personendetails.htm?t=speaker&o=215&a=select&ra=sessiondetails
40	Videos	BADW-LRZ / JUELICH	Primeur Magazine	ISC'14	Six European Exascale Projects are Dealing with the Hardware and Software Challenges in Exascale	26/06/2014	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	http://primeurmagazine.com/weekly/AE-PR-08-14-37.html
41	Videos	H.-CH. Hoppe (INTEL)	InsideHPC / Rich Report	ISC'14	Applications for the DEEP and DEEP-ER Project	30/06/2014	Portland, USA	HPC Community Industry Scientific Community	n/a	International	https://www.youtube.com/watch?v=oKwKuuIwrwA
42	Videos	JUELICH / UHEI / Eurotech / Intel	InsideHPC / Rich Report	ISC'14	DEEP and DEEP-ER Project Updates	30/06/2014	Portland, USA	HPC Community Industry Scientific Community	n/a	International	https://www.youtube.com/watch?v=flO-KOn3qKE
43	Media Coverage	JUELICH / UHEI / Eurotech	insideHPC	ISC'14	DEEP and DEEP-ER Project Updates: Video	01/07/2014	Portland, USA	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2014/07/video-deep-deep-er-project-updates-isc14/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
44	Website	BADW-LRZ / JUELICH		ISC'14	Image Gallery: DEEP Impressions from ISC'14	01/07/2014	Jülich, Germany	HPC Community Industry Scientific Community Media	n/a	International	http://www.deep-er.eu/press-corner/gallery.html
45	Workshop	E. Suarez (JUELICH)		JSC-internal seminar with visit of Dr. Tjerk P. Straatsma, Oak Ridge Leadership Computing Facility, National Center for Computational Sciences, Oak Ridge National Laboratory	DEEP and DEEP-ER	08/07/2014	Jülich, Germany	HPC Community	n/a	International	
46	Media Coverage	BADW-LRZ / JUELICH	International Innovation		Extreme Computing	22/07/2014	UK	HPC Community Industry Scientific Community Policy Makers	n/a	International	https://www.deep-er.eu/files/IntInnovation_DEEP_DEEP-ER.pdf
47	Website	BADW-LRZ			DEEP-ER Website: Reworked (new categories, new content)	15/08/2014	Munich, Germany	HPC Community Industry Scientific Community Policy Makers Media	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
48	Website	BADW-LRZ / JUELICH			DEEP-ER Into Exascale Computing	25/08/2014	Munich, Germany	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://www.deep-er.eu/press-corner/news/46-deep-er-into-exascale-computing.html
49	Publication	H.-CH. Hoppe (INTEL)	Intel European Exascale Labs Report 2013		DEEP-ER: Bringing Europe Closer To Exascale	01/09/2014	Jülich, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.exascale-labs.eu/Intel%20European%20Exascale%20Labs%20Annual%20Report%202013.pdf
50	Presentation	R. Leger (INRIA)		4th Brazil – France Workshop on High Performance Computing and Scientific Data Management	A parallel Discontinuous Galerkin Time-Domain solver of Maxwell's equations	15/09/2014	Gramado, Brasil	HPC Community Industry Scientific Community	n/a	International	
51	Website	BADW-LRZ / UHEI / INTEL			Experimenting With Innovative Memory Technology	02/10/2014	Munich, Germany	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://www.deep-er.eu/press-corner/news/50-hmc.html
52	Training	JUELICH		Training Session - organised by Intel	Intel Parallel Computing Workshop	10/11/2014	Jülich, Germany	Project Internal	n/a	Europe	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
53	Videos	BADW-LRZ		SC14	Launch of the DEEP Image Video	14/11/2014	Munich, Germany	General Public HPC Community Industry Scientific Community Media Policy Makers	n/a	International	https://www.youtube.com/watch?v=eWaFRYYJxK0
54	Exhibition	INRIA		SC14	DEEP-ER work presented at INRIA booth	17/11/2014	New Orleans, USA	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	
55	Exhibition	European Exascale Projects		SC14	Joint European Exascale Booth at SC14 booth #1039 DEEP displayed hardware	17/11/2014	New Orleans, USA	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
56	Poster	S. Breuner (ITWM) / W. Frings, K. Thust, N. Eicker (JUELICH), G. Congiu, S. Narasimhamurthy (Seagate)		SC14	DEEP-ER I/O	17/11/2014	New Orleans, USA	HPC Community Industry Scientific Community	n/a	International	http://www.deep-er.eu/files/DEEP-ER_IO_Poster.pdf
57	Website	BADW-LRZ / JUELICH		SC14	Event Info + Gallery with SC impressions	25/11/2014	Munich, Germany	Scientific Community Media HPC Community	n/a	International	http://www.deep-er.eu/press-corner/gallery.html
58	Presentation	E. Suarez (JUELICH)		JUELICH-JSC meeting (Visit C.Aubley)	DEEP and DEEP-ER	19/01/2015	Jülich, Germany	HPC Community	n/a	Germany	
59	Workshop	E. Suarez (JUELICH)		BDEC (Big Data and Extreme Scaling) Workshop	The DEEP and DEEP-ER Projects	28/01/2015	Barcelona, Spain	HPC Community Scientific Community	n/a	Europe	http://www.eesi-project.eu/pages/menu/bdec.php

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
60	Workshop	N. Eicker (JUELICH)		SUMA (Supermassive Computations in Theoretical Physics) Workshop	The European Supercomputer Projects DEEP & DEEP-ER (presentation)	11/02/2015	Trento, Italy	HPC Community Scientific Community	n/a	Europe	
61	Media Coverage	BADW-LRZ	iSGTW		Working to Make Exascale Supercomputing A Reality (Image Video)	25/02/2015	Switzerland	HPC Community Industry Scientific Community	n/a	Europe	http://www.isgtw.org/visualization/working-make-exascale-supercomputing-reality
62	Press Release	European Exascale Projects		PRACEdays15	Email Invite to European Exascale Satellite Event at PRACEdays15	30/03/2015	Munich, Germany	HPC Community Industry Scientific Community Policy Makers Media	n/a	Europe	
63	Media Coverage	European Exascale Projects	Scientific Computing World	PRACEdays15	Europe's Exascale on Display	01/04/2015	UK	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://www.scientific-computing.com/news/news_story.php?news_id=2650
64	Media Coverage	European Exascale Projects	insideHPC	PRACEdays15	Europe's Exascale on Display	05/04/2015	Portland, USA	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://insidehpc.com/2015/04/european-exascale-display-may/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
65	Flyer	European Exascale Projects			European Exascale Projects Flyer - re-worked for events in 2015 (PRACEdays15, ISC'15, SC15)	01/05/2015	Barcelona, Spain	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	current version available at: http://exascale-projects.eu/Joint_EuropeanExascaleProjects_flyer.pdf
66	Flyer	BADW-LRZ		PRACEdays15	Enabling HPC Applications For Exascale - Flyer devolped for Potential Users of a DEEP system	01/05/2015	Munich, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.deep-er.eu/images/materials/DEEP_ER_User_web_062015.pdf
67	Flyer	European Exascale Projects		ISC'15	Is Europe Ready for Exascale? - Flyer promoting the ISC'15 Workshop	01/05/2015	Barcelona, Spain	HPC Community Industry Scientific Community Media	n/a	International	http://www.exascale-projects.eu/ISC15_ExascaleWorkshop_final.pdf
68	Training	JUELICH / BSC			DEEP-ER Application Developer Training	11/05/2015	Barcelona, Spain	Project Internal	15	Europe	http://www.deep-er.eu/press-corner/news/83-in-medias-res-hands-on-training-workshop-with-deep-er-application-developers.html
69	Presentation	E. Suarez (JUELICH)		ECL Meeting Jülich	DEEP and DEEP-ER: Status Update	13/05/2015	Jülich, Germany	HPC Community	n/a	Germany	
70	Presentation	B. Mohr (JUELICH)		2nd International HPC Forum	Jülich On The Way To Exascale	20/05/2015	Tianjin, China	HPC Community	n/a	International	http://ihpcf.org/Program.html

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
71	Presentation	N. Eicker (JUELICH)		JSC-LBL meeting (Visit S.Dosanjh)	DEEP and DEEP-ER	22/05/2015	Jülich, Germany	HPC Community	n/a	Europe	
72	Conference	European Exascale Projects		PRACEdays15	Satellite Event titled: Enabling Exascale in Europe for Industry	26/05/2015	Dublin, Ireland	HPC Community Industry Scientific Community	45	Europe	http://www.prace-ri.eu/IMG/pdf/PRACEDAYS15_SatelliteEvent_EuroExa_Agenda_FINAL.pdf
73	Presentation	E. Suarez (JUELICH); M. Tchiboukdkjian (CGG); G. Staffelbach (CERFACS)		PRACEdays15	DEEP and DEEP-ER: Innovative Exascale architectures in the light of user requirements	26/05/2015	Dublin, Ireland	HPC Community Industry Scientific Community	45	Europe	http://www.prace-ri.eu/IMG/pdf/DEEP_EstelaSuarez_GabrielStaffelbach_MarcTchiboukdjian.pdf
74	Media Coverage	European Exascale Projects	Scientific Computing World	ISC'15	ISC'15 High Performance Show Preview	01/06/2015	UK	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://www.scientific-computing.com/show/show.php?show_id=28

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
75	Presentation	E. Suarez (JUELICH)		NorduGrid Conference	The DEEP-ER way towards exascale I/O and resilience	04/06/2015	Bern, Switzerland	HPC Community Scientific Community	n/a	Europe	http://indico.hep.lu.se//contributionDisplay.py?contribId=8&sessionId=7&confId=1578 Slides: http://indico.hep.lu.se//getFile.py/access?contribId=8&sessionId=7&resId=0&materialId=slides&confId=1578
76	Presentation	G. Lapenta (KU Leuven)		Astronum 2015	Using HPC Kinetic Simulations to Help the MMS Mission Find its Target: Reconnection Diffusion Regions	08/06/2015	Avignon, France	HPC Community Scientific Community	n/a	Europe	http://irfu.cea.fr/ASTRONUM2015/abstract.html
77	Media Coverage	BADW-LRZ	Primeur Magazine	ISC'15	DEEP/-ER at ISC'15	11/06/2015	Munich, Germany	HPC Community Industry Scientific Community Policy Makers Media	n/a	International	http://primeurmagazine.com/flash/AE-PF-05-15-30.html
78	Website	BADW-LRZ / JUELICH		ISC'15	Announcement: DEEP/-ER at ISC'15	11/06/2015	Munich, Germany	HPC Community Industry Scientific Community Policy Makers Media	n/a	International	http://www.deep-er.eu/press-corner/events/past-events/3-isc-15.html
79	Flyer	BADW-LRZ		ISC'15	DEEP-ER Flyer: Completely updated / re-worked version	15/06/2015	Munich, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	http://www.deep-er.eu/images/materials/Flyer_DEEP_ER_web_062015.pdf

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
80	Media Coverage	J. Labarta (BSC)	ISC'15 Blog	ISC'15	Survival Machine and Surviving the Machine	23/06/2015	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.isc-hpc.com/blog/the-survival-machine-and-surviving-the-machines.html
81	Website	BADW-LRZ		PRACEdays15	Enabling Exascale in Europe For Industry - Event Review	25/06/2015	Munich, Germany	HPC Community Industry Scientific Community Policy Makers	n/a	International	http://www.deep-er.eu/press-corner/news/82-enabling-exascale-in-europe-for-industry.html
82	Media Briefing / Interview	E. Suarez + N. Eicker (JUELICH)	The Platform	ISC'15	Background Interview on DEEP + DEEP-ER	26/06/2015	Jülich, Germany	Media	1		
83	Presentation	W. Frings (JUELICH)		3rd JLESC Workshop	Description of buddy checkpointing concepts	29/06/2015	Barcelona, Spain	HPC Community	n/a	Europe	
84	Presentation	D. Alvarez Mallon (JUELICH)		ISC'15	DEEP and DEEP-ER: From a Concept to Application Reality	01/07/2015	Jülich, Germany	HPC Community Scientific Community	n/a	International	https://prezi.com/mvtzcpeoodgh/deep-and-deep-er-isc15-final/?utm_campaign=share&utm_medium=copy
85	Presentation	D. Alvarez Mallon (JUELICH)		Third NESUS Working Groups Meeting	DEEP & DEEP-ER: European HW and SW Innovations on the Way to Exascale	02/07/2015	Leuven, Belgium	HPC Community Scientific Community	50	Europe	http://www.nesus.eu/event/third-nesus-working-groups-meeting?instance_id=136

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
86	Presentation	J. Amaya (KU Leuven)		12th International School/Symposium for Space Simulations (ISSS-12)	The Numerical Magnetosphere: Fully Kinetic Simulations of the Solar Wind–Magnetosphere Interaction	03/07/2015	Prague,	Czech Republic	n/a	Europe	
87	Presentation	M. E. Innocenti (KU Leuven)		12th International School/Symposium for Space Simulations (ISSS-12)	Performance Analysis of the Multi-level Multi-domain Approach on Emerging Computing Architectures	03/07/2015	Prague,	Czech Republic	n/a	Europe	
88	Presentation	F. Bacchini (KU Leuven)		12th International School/Symposium for Space Simulations (ISSS-12)	New Approaches to Large Scale Particle in Cell Simulations based on Fluid Methods Applied to Emerging Hybrid Architectures	03/07/2015	Prague,	Czech Republic	n/a	Europe	
89	Website	BADW-LRZ		ISC'15	Relaunch DEEP-ER Website	10/07/2015	Munich, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	www.deep-er.eu
90	Exhibition	European Exascale Projects		ISC'15	Joint European Exascale Booth at ISC'15 DEEP - displayed examples of final hardware	12/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community Media Policy Makers	150	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
91	Presentation	H.-CH. Hoppe (INTEL) + D. Alvarez Mallon (JUELICH)		ISC'15	Presentation + Demo at the Intel booth	13/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	
92	Presentation	E. Suarez (JUELICH)		ISC'15	Architecture Innovation with Intel Xeon Phi - presentation held at the Intel Booth	13/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community Media Policy Makers	n/a	International	
93	Media Coverage	Eurotech	idw Online	ISC'15	Eurotech delivers the "Booster" system to Jülich to complete the DEEP supercomputer	14/07/2015	Amaro, Italy Frankfurt, Germany	HPC Community Industry Scientific Community Media	n/a	International	https://idw-online.de/en/news634748
94	Press Release	Eurotech		ISC'15	Eurotech delivers the "Booster" system to Jülich to complete the DEEP supercomputer	14/07/2015	Amaro, Italy Frankfurt, Germany	HPC Community Industry Scientific Community Media	n/a	International	http://www.eurotech.com/en/press+room/news/?723&Eurotech+delivers+the+%26quot%3bBooster%26quot%3b+system+to+J%26%23252%3blich+to+complete+the+DEEP+supercomputer
95	Videos	R. Gimenez (BSC)		ISC'15	European Exascale Projects at ISC'15: Day 1	14/07/2015	Frankfurt, Germany	HPC Community Scientific Community Media	100	International	https://www.youtube.com/watch?v=ZP4uUzG7GJc&feature=youtu.be

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
96	Presentation	N. Eicker (JUELICH)		ISC'15	Presentation at EEP Workshop: Taming Heterogeneity by Segregation	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community	40	International	http://www.isc-events.com/isc15_ap/sessiondetails.htm?t=session&o=249&a=select
97	Workshop	European Exascale Projects		ISC'15	Is Europe Ready for Exascale?	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community	40	International	http://www.isc-events.com/isc15_ap/sessiondetails.htm?t=session&o=249&a=select
98	Media Coverage	Eurotech	insideHPC	ISC'15	Eurotech Delivers Booster to DEEP Project	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2015/07/eurotech-delivers-booster-to-deep-project/
99	Media Coverage	Eurotech	HPCwire	ISC'15	Eurotech Delivers “Booster” System to Jülich to Complete DEEP Supercomputer	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.hpcwire.com/off-the-wire/eurotech-delivers-booster-system-to-julich-to-complete-deep-supercomputer/
100	Presentation	N. Eicker (JUELICH)		ISC'15	Taming Heterogeneity by Segregation – An Innovative Approach to Heterogeneous Exascale Architectures	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community Media	25	International	http://www.isc-events.com/isc15_ap/presentationdetails.htm?t=presentation&o=388&a=select&ra=personendetails
101	Workshop	European Exascale Projects		ISC'15	Is Europe Ready for Exascale?	16/07/2015	Frankfurt, Germany	HPC Community Industry Scientific Community Media	25	International	http://www.isc-events.com/isc15_ap/sessiondetails.htm?t=session&o=249&a=select

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
102	Media Coverage	Eurotech	Scientific Computing	ISC'15	Booster System Installed at Jülich, Completes DEEP Supercomputer	17/07/2015		HPC Community Industry Scientific Community	n/a	International	http://www.scientificcomputing.com/news/2015/07/booster-system-installed-j%C3%BClich-completes-deep-supercomputer
103	Website	BADW-LRZ / JUELICH		ISC'15	Image Gallery: DEEP Impressions from ISC'15	17/07/2015	Jülich, Germany	HPC Community Industry Scientific Community Media	n/a	International	http://www.deep-project.eu/deep-project/EN/News/Multimedia/Galleries/_node.html
104	Workshop	G. Lapenta (KU Leuven)		PIC Methods for Emerging Architectures: a shootout between explicit and implicit approaches	Workshop organised at Santa Monica Center for Emerging Technologies	20/07/2015	Santa Monica, USA	HPC Community Scientific Community	n/a	International	
105	Media Coverage	UHEI / Extoll	insideHPC	ISC'15	Extoll rolls out Tourmalet Network Chip at ISC'15	22/07/2015	Mannheim, Germany	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2015/07/extoll-rolls-out-tourmalett-network-chip-at-isc-2015/
106	Media Coverage	UHEI / Extoll	Innovations Report	ISC'15	Extoll Introduces HPC Network Chip Tourmalet	22/07/2015	Mannheim, Germany	HPC Community Industry Scientific Community	n/a	Germany	http://www.innovations-report.com/html/reports/information-technology/extoll-introduces-the-hpc-network-chip-tourmalet.html

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
107	Press Release	UHEI / Extoll			Extoll Introduces HPC Network Chip Tourmalet	22/07/2015	Mannheim, Germany	HPC Community Industry Scientific Community Media	n/a	International	http://www.deep-project.eu/SharedDocs/Downloads/DEEP-PROJECT/EN/Pressreleases/extoll-july-2015.pdf?__blob=publicationFile
108	Workshop	JUELICH		3rd JLESC Workshop	Description of buddy checkpointing concepts	29/07/2015	Barcelona, Spain	HPC Community	n/a	International	
109	Workshop	W. Frings (JUELICH)		CLUSTOR Workshop - Workshop on Cluster Computing Technology	Description of buddy checkpointing concepts	30/07/2015	Hamburg, Germany	HPC Community	n/a	Germany	
110	Presentation	V. Beltran (BSC)		13th US National Conference on Computational Mechanics	Enabling Complex Applications on Heterogeneous Clusters with OmpSs MPI Offloading	30/07/2015	San Diego, USA	HPC Community	n/a	International	http://13.usnccm.org/
111	Media Coverage	E. Suarez (JUELICH)	insideHPC	ISC'15	DEEP Moves Towards Exascale	03/08/2015	Portland, USA	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2015/08/deep-project-moves-towards-exascale-at-isc-2015/
112	Media Coverage	E. Suarez (JUELICH)	Primeur Magazine	ISC'15	European Exascale Projects DEEP-ER and Mont-Blanc to Investigate New Exascale Technologies (VIDEO)	10/08/2015	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	https://www.youtube.com/watch?t=16&v=tr_co6vu-4s
113	Media Coverage	E. Suarez (JUELICH)	Primeur Magazine	ISC'15	European Exascale Projects DEEP-ER and Mont-Blanc to Investigate New Programming and Network-attached Memory technologies	10/08/2015	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	http://primeurmagazine.com/weekly/AE-PR-09-15-48.html

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
114	Media Coverage	J. Schmidt (UHEI)	Primeur Magazine	ISC'15	Demonstration: technology developments of the exascale project DEEP-ER	10/08/2015	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	https://www.youtube.com/watch?v=aA42Me4s-nI
115	Presentation	G. Lapenta (KU Leuven)		International Conference on Numerical Simulation of Plasmas	Using HPC Kinetic Simulations to help the MMS mission find its target: reconnection diffusion regions	11/08/2015	Golden (Colorado), US	HPC Community Scientific Community	n/a	Europe	http://www.icnsp.org/
116	Presentation	R. Leger (INRIA)		ParCo 2015	Assessing the DEEP-ER Cluster/Booster Architecture with a Finite-Element Type Solver for Bioelectromagnetics	01/09/2015	Edinburgh, Scotland	HPC Community Scientific Community	n/a	Europe	https://www.conftool.net/parco2015/index.php?page=browseSessions&presentations=show
117	Presentation	N. Attig (JSC)		CSP 2015	Impacts of Current Hardware and Software Developments on Simulation Sciences	08/09/2015	Moscow, Russia	HPC Community	n/a	International	http://csp2015.ac.ru/Presentations/Attig.pdf
118	Presentation	N. Eicker (JUELICH)		596. WE-Heraeus Seminar “Science Applications for Exascale Computing - Exploring New Avenues towards Scalability and Fault-Tolerance”	Taming Heterogeneity by Segregation – The DEEP View on Exascale	09/09/2015	Bad Honnef, Germany	HPC Community	n/a	Germany	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
119	Media Coverage	A. Auweter (BADW-LRZ)	Primeur Magazine	ISC'15	System monitoring for energy efficiency in the Mont-Blanc and DEEP-ER project	29/09/2015	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	http://primeurmagazine.com/weekly/AE-PR-10-15-92.html
120	Media Coverage	N. Eicker (JUELICH)	Primeur Magazine	ISC'15	Developing Hardware for the Exascale Era	06/10/2015	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	http://primeurmagazine.com/weekly/AE-PR-11-15-24.html
121	Press Release	JUELICH		Release Final Project Brochure DEEP	Next Generation Supercomputing - Boosting Science in Europe	15/10/2015	Jülich, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.deep-project.eu/SharedDocs/Downloads/DEEP-PROJECT/EN/deep-final-brochure.pdf?__blob=publicationFile
122	Workshop	N. Eicker (JUELICH)		LENS2015 International Workshop	Taming Heterogeneity by Segregation – Taming Heterogeneity by Segregation -- The DEEP and DEEP-ER take on Heterogeneous Cluster Architectures	29/10/2015	Akihabara, Japan	HPC Community	n/a	International	http://wallaby.aics.riken.jp/lens/invited.html
123	Exhibition	European Exascale Projects		SC15	Joint European Exascale Booth	16/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	150	International	
124	Poster	J. Schmidt (UHEI)		SC15	openHMC - an Open-Source Hybrid Memory Cube Controller --> presented at the Emerging Technology Track of the Conference	17/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
125	Media Briefing / Interview	E. Suarez + N. Eicker (JUELICH)	Intersect 360 Research	SC15	Interview with Analysts from 360 Research	17/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	2	International	
126	Media Briefing / Interview	E. Suarez + N. Eicker (JUELICH)	Scientific Computing World	SC15	Interview on DEEP results and DEEP-ER	18/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	2	International	
127	Presentation	E. Suarez (JUELICH)		SC15	The DEEP-ER project presentation held at the Intel Theatre "An Update on European HPC Initiatives"	19/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	15	International	
128	Presentation	S. Breuner (ITWM)		SC15	Presentation on DEEP-ER project during SC (at partner booth of BeeGFS)	19/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	15	International	
129	BoF	F. Mantovani (BSC)		SC15	BoF Session: Taking on Exascale Challenges: Key Lessons and International Collaboration Opportunities	19/11/2015	Austin, Texas, USA	HPC Community Industry Scientific Community	100	International	http://www.deep-er.eu/press-corner/news/132-bof-sc15.html
130	Exhibition	K. Thust (JUELICH)		SC15	SIONlib promted at JSC booth	19/11/2015	Austin, Texas, USA	HPC Community Scientific Community	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
131	Exhibition	S. Narasimhamurthy (Seagate)		SC15	E10 presented at Seagate booth	19/11/2015	Austin, Texas, USA	HPC Community Scientific Community	n/a	International	
132	Workshop	V. Beltran (BSC)		JLESC Workshop	Task-based resiliency in OmpSs	03/12/2015	Bonn, Germany	HPC Community Scientific Community	n/a	International	http://www.fz-juelich.de/ias/jsc/EN/Expertise/Workshops/Conferences/JLESC-4/Programme/Abstracts/beltran-resilience-s.html?nn=1893370
133	Poster	J. Schmidt (UHEI)		International Conference on RECONFIGurable Computing and FPGAs	openHMC - A configurable open-source Hybrid Memory Cube controller	07/12/2015	Mayan, Mexico	HPC Community Scientific Community	200	International	http://www.reconfig.org/
134	Poster	J. Amaya (KU Leuven)		AGU Fall Meeting	First-principle modeling of planetary magnetospheres: Mecury and the Earth	14/12/2015	San Francisco, USA	HPC Community Scientific Community	n/a	International	http://fallmeeting.agu.org/2015/
135	Poster	J. Amaya (KU Leuven)		AGU Fall Meeting	Fully Kinetic 3D Simulations of the Interaction of the Solar Wind with Mercury	14/12/2015	San Francisco, USA	HPC Community Scientific Community	n/a	International	http://fallmeeting.agu.org/2015/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
136	Publication	H.-CH. Hoppe (INTEL)	Intel European Exascale Labs Report 2014/2015		DEEP and DEEP-ER: Brining Europe Closer to Exascale	31/12/2015		HPC Community Industry Scientific Community	n/a	International	http://www.exascale-labs.eu/Intel%20European%20Exascale%20Labs%20Annual%20Report%202014.pdf
137	Workshop	E. Suarez (JUELICH)		HPC-Leap: School on HPC architectures and large-scale numerical computation	Implementing a new computer architecture paradigm	15/01/2016	Jülich, Germany	HPC Community Scientific Community	n/a	International	http://indico-jsc.fz-juelich.de/event/16/other-view?view=standard
138	Website	S. Eisenreich (BAdW-LRZ)	DEEP-ER Website	CeBIT 2016	DEEP-ER @ CeBIT16	15/02/2016	Hannover, Germany	HPC Community Industry	n/a	International	http://www.deep-er.eu/press-corner/events/18-deep-er-cebit.html
139	Media Briefing / Interview	Th. Lippert (JUELICH)	Science Node		Bylined article by Th. Lippert on "Boosting Science with the next generation of Supercomputers"	15/02/2016	Jülich, Germany	HPC Community Scientific Community	n/a	International	
140	Media Coverage	S. Breuner (ITWM)	insideHPC		BeeGFS Parallel File System goes OpenSource	23/02/2016	USA	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2016/02/beegfs-parallel-file-system-now-open-source/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
141	Press Release	S. Breuner (ITWM)			BeeGFS Parallel File System now OpenSource	23/02/2016	Kaiserslautern, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.beegfs.com/docs/press_releases/2016-02-23-BeeGFS_Open_Source_and_User_Meeting.pdf
142	Media Coverage	S. Breuner (ITWM)	HPCwire		BeeGFS Parallel File System now OpenSource	23/02/2016	Kaiserslautern, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.hpcwire.com/off-the-wire/beegfs-parallel-file-system-now-open-source/
143	Media Coverage	S. Breuner (ITWM)	Scientific Computing World		BeeGFS Parallel File System now OpenSource	23/02/2016	UK	HPC Community Industry Scientific Community	n/a	International	http://www.scientific-computing.com/news/news_story.php?news_id=2782
144	Media Coverage	S. Eisenreich (BAdW-LRZ)	Primeur Magazine	CeBIT 2016	Exascale project DEEP-ER to present at CEBIT	01/03/2016	Almere, The Netherlands	HPC Community Industry Scientific Community	n/a	International	http://primeurmagazine.com/flash/AE-PF-03-16-5.html
145	Exhibition	E. Suarez (JUELICH)		CeBIT 2016	JSC presents DEEP + DEEP-ER at booth 'Innovationsland NRW"	14/03/2016	Hannover, Germany	HPC Community Industry	n/a	International	http://www.deep-er.eu/press-corner/events/18-deep-er-cebit.html
146	Media Briefing / Interview	E. Suarez (JUELICH)	DataCenter Insider	CeBIT 2016	Interview with editor-in-chief Ulrike Ostler	14/03/2016	Hannover, Germany	HPC Community Industry	n/a	Germany	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
147	Media Briefing / Interview	E. Suarez (JUELICH)	Süddeutsche Zeitung	CeBIT 2016	Interview with Tech Editor Helmut Martin-Jung	16/03/2016	Hannover, Germany	Industry General Public	n/a	Germany	
148	Media Briefing / Interview	E. Suarez (JUELICH)	Informatik Spektrum	CeBIT 2016	Interview with editor-in-chief Hartmut Engesser	16/03/2016	Hannover, Germany	Scientific Community HPC Community	n/a	Germany	
149	Media Coverage	Th. Lippert (JUELICH)	Science Node		Boosting Science with the next generation of supercomputers	16/03/2016	Switzerland	HPC Community Scientific Community	n/a	International	https://sciencenode.org/feature/boosting-science-with-the-next-generation-of-supercomputers.php
150	Videos	E. Suarez (JUELICH)	Twenty4 Picutres	CeBIT 2016	Video production on promoting DEEP/ER prototype system	17/03/2016	Hannover, Germany	HPC Community Industry Scientific Community	n/a	Germany	
151	Website	S. Eisenreich (BAdW-LRZ)	DEEP-ER Website		Update Events List	31/03/2016	Munich, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.deep-er.eu/press-corner/events.html
152	Website	S. Eisenreich (BAdW-LRZ)	DEEP-ER Website		Update Application Pages	31/03/2016	Munich, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.deep-er.eu/applications.html

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
153	Presentation	C. Clauss & T. Moschny (ParTec)		parallel 2016	Verhalten von MPI Programmen im Fehlerfall (= Behaviors for MPI programmes when errors occur)	07/04/2016	Heidelberg, Germany	HPC Community	n/a	German	http://www.parallelcon.de/veranstaltung-5099-dem-tausendf%C3%BC%C3%9Fler-ein-bein-gestellt%3A-verhalten-von-mpi-programmen-im-fehlerfall.html?id=5099
154	Presentation	J. Amaya (KU Leuven)		EGU Conference 2016	Innovative HPC Architectures for the Study of Planetary Plasma Environments	17/04/2016	Vienna, Austria	Scientific Community	n/a	International	
155	Presentation	J. Amaya (KU Leuven)		EASC 2016	Towards exascale simulations of space plasmas using the DEEP-ER architecture	28/04/2016	Stockholm, Sweden	Scientific Community HPC Community	n/a	International	https://docs.google.com/viewer?a=v&pid=sites&srcid=ZWFzYzIwMTYuY29tfGVhc2MtMjAxNi1zdG9ja2hvbG18Z3g6MjFlMzE2MzI4NzljMzIzMA
156	Presentation	E. Suarez (JUELICH)		EXDCI Workshop at European HPC Summit Week	The DEEP and DEEP-ER Projects	10/05/2016	Prague, Czech Republic	HPC Community	n/a	International	http://www.deep-er.eu/images/materials/DEEP_DEEPER_20160510_EXDCI_Prague_Final.pdf
157	Presentation	C. Manzano (JUELICH)		BeeGFS User Meeting	BeeGFS in the DEEP-ER Project	17/05/2016	Kaiserslautern, Germany	HPC Community	n/a	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
158	Presentation	F. Kautz (Fraunhofer ITWM)		BeeGFS User Meeting	BeeGFS User APIs	17/05/2016	Kaiserslautern, Germany	HPC Community	n/a	International	
159	Media Coverage	S. Eisenreich (BAdW-LRZ)	Scientific Computing World	ISC'16	ISC'16 Show Preview	01/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	print
160	Presentation	E. Suarez (JUELICH)		EMiT Conference	Technology Emerging from the DEEP and DEEP-ER projects	03/06/2016	Barcelona, Spain	HPC Community Industry Scientific Community	90	International	http://emit.tech/EMiT2016/Suarez-EMiT2016-Barcelona.pdf
161	Media Coverage	S. Eisenreich (BAdW-LRZ)	ISC'16 Blog	ISC'16	Programming Models: Slow transition or complete disruption?	10/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	now availalbe at: http://www.deep-er.eu/press-corner/news/174-programming-models-for-exascale-transition-or-disruption.html
162	Publication	S. Eisenreich (BAdW-LRZ)		ISC'16	Europe Towards Exascale - A Lookback on 5 Years of European Exascale Research Collaboration	14/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://exascale-projects.eu/EuroExaFinalBrochure_v1.0.pdf

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
163	Media Coverage	E. Suarez (JUELICH)	INSIDE (Innovatives Supercomputing in Deutschland)	CeBIT 2016	DEEPprojects at CeBIT 2016	15/06/2016	Stuttgart, Germany	HPC Community Scientific Community	n/a	German	http://inside.hlrs.de/#deep-project-at-cebit16
164	Press Release	U. Krackhardt (EXTOLL / UHEI)	EXTOLL Website	ISC'16	EXTOLL's Network chip enables network attached accelerators of any kind	16/06/2016	Mannheim, Germany	HPC Community Industry Scientific Community	n/a	international	http://www.deep-er.eu/images/EXTOLL_Tourmalet_ISC_v1.0.pdf
165	Media Coverage	U. Krackhardt (EXTOLL / UHEI)	insideHPC	ISC'16	Extoll Network Chip enables network attached accelerators	17/06/2016	Portland, USA	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2016/06/extolls-network-chip-enables-network-attached-accelerators-of-any-kind/
166	Presentation	J. Schmidt (UHEI)		ISC'16	Network Attached Memory - Presentation at the PhD Forum	20/06/2016	Frankfurt, Germany	HPC Community	n/a	International	
167	Presentation	I. Zacharov (Eurotech)		ISC'16	Aurora Tigon v4 with KNL, a system from research for research (presentation at the Intel Collaboration Hub)	20/06/2016	Frankfurt, Germany	HPC Community Scientific Community	15	International	
168	Press Release	A. Somma (Eurotech)	Eurotech Website	ISC'16	Eurotech introduces the Aurora Tigon v4	21/06/2016	Amaro, Italy Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.eurotech.com/en/press+room/news/?775

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
169	Presentation	E. Suarez (JUELICH)		ISC'16	The DEEP-ER Project (presentation at the Intel Collaboration Hub)	21/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	15	International	
170	Presentation	J. Labarta (BSC)		ISC'16	The OmpSs Programming Model Vision (BoF presentation)	21/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	40	International	
171	Press Release	A. Somma (Eurotech)		ISC'16	Eurotech introduces the Aurora Tigon v4	21/06/2016	Amaro, Italy Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://www.eurotech.com/en/press+room/news/?775
172	Presentation	I. Zacharov (Eurotech)		ISC'16	BoF Presentation 'Monitoring large-scale HPC Systems'	22/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	45	International	
173	Presentation	N. Eicker (JUELICH)		ISC'16	BoF Presentation 'Exascale I/O Challenges'	22/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	45	International	
174	Presentation	N. Eicker (JUELICH)		ISC'16	Hardware Prototyping in DEEP-ER presentation at ISC'16 Workshop	23/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	30	International	
175	Workshop	E. Suarez (JUELICH)		ISC'16	Workshop organised by DEEP-ER and MontBlanc on 'Hardware prototyping for next-gen HPC architectures'	23/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	30	International	

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
176	Presentation	R. Leger (INRIA)		ISC'16	A feedback on approaching the DEEP-ER platform with a DGTD-based simulation software for Bioelectromagnetics applications	23/06/2016	Frankfurt, Germany	HPC Community Industry Scientific Community		International	Workshop report availalbe via: http://arxiv.org/pdf/1607.02835v1.pdf
177	Exhibition	Intel, Seagate, Inria		Teratec Forum 2016	DEEP-ER represented at partner booth via poster, flyer, leaflets etc	28/06/2016	BRUYERES LE CHATEL, France	HPC Community Industry Scientific Community	n/a	International	http://www.teratec.eu/gb/forum/
178	Flyer	A. Somma (Eurotech)		SAI Computing Conference 2016	Roll-up display for DEEP-ER booth at SAI: To Exascale with Co-Design	01/07/2016	Amaro, Italy	HPC Community Industry Scientific Community			
179	Presentation	E. Suarez (JUELICH)		KIT-JSC Meeting	DEEP/-ER Cooling Concept	01/07/2016	Karlsruhe, Germany	HPC Community Industry Scientific Community		Germany	
180	Media Coverage	E. Suarez (JUELICH)	insideHPC	ISC'16	DEEP-ER Project Moves Europe Closer to Exascale	05/07/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	http://insidehpc.com/2016/07/deep-er-project/
181	Media Coverage	E. Suarez (JUELICH)	DataCenter Insider		Supercomputing neu gedacht	05/07/2016	Augsburg, Germany	HPC Community Industry	n/a	Germany (DACH region, including Austria + Switzerland)	http://www.datacenter-insider.de/100mal-effizienter-bei-1018-operationen-pro-sekunde-a-541242/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
182	Exhibition	N. Eicker (JUELICH) + T. Moschny (ParTec)		SAI Computing Conference 2016	DEEP-ER as official knowledge partner to SAI computing conference, including booth on show floor	13/07/2016	London, GB	Industry Scientific Community	200	International	http://saiconference.com/Conferences/Computing2016
183	Presentation	N. Eicker (JUELICH)		SAI Computing Conference 2016	Taming Heterogeneity in HPC - Keynote	15/07/2016	London, GB	Industry Scientific Community		International	
184	Media Coverage	E. Suarez (JUELICH)	Top500 blog		A Dive into DEEP-ER, Exascale Research with a Distinctly Europea Flair	18/07/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	https://www.top500.org/news/a-dive-into-deep-er-exascale-research-with-a-distinctly-european-flair/
185	Media Coverage	U. Krackhardt (EXTOLL / UHEI)	Top500 blog		EXTOLL's Network Marches to the Beat of a different Drummer	25/07/2016	Frankfurt, Germany	HPC Community Industry Scientific Community	n/a	International	https://www.top500.org/news/extolls-network-marches-to-the-beat-of-a-different-drummer/
186	Videos	E. Suarez (JUELICH)	YouTube	EMiT Conference	Interview to Keynote Speaker E. Suarez	28/07/2016	Barcelona, Spain	HPC Community Industry Scientific Community	n/a	International	https://www.youtube.com/watch?v=5KL0RMYW4A4
187	Media Coverage	N. Eicker (JUELICH)	YouTube	SAI Computing Conference 2016	Taming Heterogeneity in HPC - Keynote	08/08/2016	London, GB	HPC Community Industry Scientific Community		International	https://www.youtube.com/watch?v=aM9AkgG5ud4

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
188	Media Coverage	S. Eisenreich (BAdW-LRZ)	insideHPC	ISC'16	New Report Looks at European Exascale Projects	12/08/2016	Portland, USA	HPC Community Industry Scientific Community		International	http://insidehpc.com/2016/08/european-exascale-projects/
189	Media Coverage	S. Eisenreich (BAdW-LRZ)	insideHPC	SAI Computing Conference 2016	Taming Heterogeneity in HPC - The DEEP-ER take	19/08/2016	Portland, USA	HPC Community Industry Scientific Community		International	http://insidehpc.com/2016/08/taming-heterogeneity-in-hpc-the-deep-er-take/
190	Presentation	E. Suarez (JUELICH)		Kleine Nacht der Wissenschaft, Kulturbahnhof Jülich	The future of Supercomputing	02/09/2016	Jülich, Germany	General Public		Germany	http://www.fz-juelich.de/portal/DE/UeberUns/Veranstaltungen/docs/2016/16-09-02-kleine-nacht-der-wissenschaft.html
191	Media Coverage	S. Eisenreich (BAdW-LRZ)	Europa.eu	n/a	Europe towards Exascale	02/09/2016	Brussels, Belgium	HPC Community Industry Scientific Community		International	https://ec.europa.eu/digital-single-market/en/news/europe-towards-exascale
192	Conference	Congiu, G. (Seagate)		IEEE Cluster 2016	Improving Collective I/O Performance Using Non-Volatile Memory Devices	13/09/2016	Teipei, Teiwan	HPC Community Industry Scientific Community		International	https://ssl.linklings.net/conferences/ieeecluster/ieeecluster2016_program/views/at_a_glance.html
193	Media Coverage	S. Eisenreich (BAdW-LRZ)	HPCwire	SC16	HPCwire Readers' Choice Award 2016 (category: academia-industry collaboration)	30/09/2016	USA	HPC Community Industry Scientific Community		International	https://www.hpcwire.com/2016-hpcwire-readers-choice-awards/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
194	Media Coverage	E. Suarez + N. Eicker (JUELICH)	Mitarbeitermagazine des Forschungszentrums Jülich (Employee Magazine by FZJ)		Internationale Forschungskooperationen: Gemainsam Stark (Interview + article in German on international collaboration)	01/10/2016	Jülich, Germany	HPC Community Scientific Community		Germany	
195	Media Coverage	S. Eisenreich (BAdW-LRZ) / E. Suarez (JUELICH)	INSIDE (Innovatives Supercomputing in Deutschland)	ISC'16	Prototyping next-generation supercomputing architectures - report on ISC'16 workshop	01/10/2016	Stuttgart, Germany	HPC Community Industry Scientific Community		International	http://inside.hlrs.de/#prototyping-next-generation-supercomputing-architectures-isc16-workshop
196	Conference	J. Schmidt (UHEI)		Memsys 2016	Exploring Time and Energy for Complex Accesses to a Hybrid Memory Cube	04/10/2016	Washington, USA	HPC Community Industry Scientific Community	n/a	International	http://memsys.io/2016/program/
197	Website	S. Eisenreich (BAdW-LRZ)			Latest SIONlib version includes DEEP-ER developments	14/10/2016	Jülich, Germany	HPC Community Industry Scientific Community			http://www.deep-er.eu/press-corner/news/216-latest-sionlib-version-includes-deep-er-developments.html

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
198	Exhibition	S. Breuner (ITWM)		SEG - Society of Exploration Geophysicists; 86th annual meeting	BeeGFS booth onsite	16/10/2016	Dallas, Texas	Scientific Community, Industry	n/a	International	http://seg.org/Annual-Meeting-2016
199	Website	S. Eisenreich (BAdW-LRZ)			Earth System Modelling: Benefitting from Cluster-Booster division - A DEEPprojects Use Case	09/11/2016	Munich, Germany	HPC Community Scientific Community		International	http://www.deep-er.eu/applications/deep-er-use-cases/211-earth-system-modelling-benefitting-from-cluster-booster-division.html
200	Website	S. Eisenreich (BAdW-LRZ)		SC16	DEEP-ER on the finishing streight - Status Update report for SC16	11/11/2016	Munich, Germany	HPC Community Industry Scientific Community		International	http://www.deep-er.eu/press-corner/news/214-deep-er-on-the-finishing-straight.html
201	Exhibition	S. Eisenreich (BAdW-LRZ) + E. Suarez (JSC)		SC16	Booth space at official JSC booth	14/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	
202	Conference	J. Schmidt (UHEI)		SC16	Doctoral Showcase NAM	15/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	http://sc16.supercomputing.org/presentation/?id=drs106&sess=sess269
203	BoF	N. Eicker (JUELICH)		SC16	MPICH: A High-Performance Open-Source MPI Implementation BoF	15/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	http://sc16.supercomputing.org/presentation/?id=bof112&sess=sess329

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
204	Conference	M. Ostaz (BSC)		SC16	ETP4HPC: BoF Session	16/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	http://www.etp4hpc.eu/en/sc16-bof.html
205	Presentation	N. Eicker (JUELICH)		SC16	Modular Supercomputing - Talk at Intel Community Hub	17/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	
206	Conference	W. Frings (JUELICH)		SC16	The DEEP-ER take on exascale I/O - Talk at NEXTGenIO Workshop	18/11/2016	Salt Lake City, USA	HPC Community Industry Scientific Community		International	http://www.nextgenio.eu/events/workshop-sc-16-exascale-io-challenges-innovations-and-solutions
207	Website	S. Eisenreich (BAdW-LRZ) / F. Kautz (ITWM)			BeeGFS DEEP-ER Cache API availalbe	30/11/2016	Kaiserslautern, Germany	HPC Community Industry Scientific Community		International	http://www.deep-er.eu/press-corner/news/215-beegfs-deep-er-cache-api-available.html
208	Presentation	W. Frings (JUELICH)		JLESC Workshop	“HPC-Tools JUBE, LLview and SIONlib at JSC: Recent developments” (presentation containing DEEP-ER’s buddy checkpointing and NAM-XOR-checkpointing using SIONlib)	01/12/2016	Kobe, Japan	HPC Community Industry Scientific Community		International	
209	Poster	J. Amaya (KU Leuven)		AGU Fall Meeting	Global fully kinetic models of planetary magnetospheres with iPic3D	15/12/2016	San Francisco, USA	HPC Community Scientific Community		International	https://agu.confex.com/agu/fm16/meetingapp.cgi/Paper/164592
210	Media Coverage	Intel	HPCwire		Advancing Modular Supercomputing with DEEP and DEEP-ER	24/02/2017	San Diego, USA	HPC Community Industry Scientific Community		International	https://www.hpcwire.com/2017/02/24/modular-supercomputing-deep-deep-er-architectures/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
211	Videos	Intel		SC16	DEEP-ER Modular Supercomputing - Recording of a talk given at SC16 at the Intel Community Hub	24/02/2017		HPC Community Industry Scientific Community		International	http://www.intel.com/content/www/us/en/high-performance-computing/julich-deeper-projects-video.html
212	Conference	E. Suarez (JUELICH)		ISUM 2017 (8th International Supercomputing Conference In Mexico)	Keynote: Modular Supercomputing: the DEEP approach to hardware heterogeneity	02/03/2017	Guadalajara, Mexico	HPC Community Scientific Community		International	http://www.isum.mx/en/contenido/dr-estela-suarez
213	Conference	E. Suarez (JUELICH)		ISUM 2017 (8th International Supercomputing Conference In Mexico)	Round-table: Women in TICS: Women at Technology World	03/03/2017	Guadalajara, Mexico	HPC Community Scientific Community		International	http://isum.mx/en/evento/women-tics-women-technology-world
214	Presentation	E. Suarez (JUELICH)		ParFlow developers workshop	The DEEP project(s) and the Cluster-Booster Architecture	28/03/2017	Jülich, Germany	HPC Community Scientific Community Industry	n/a	German	
215	Publication	G. Lapenta (KU Leuven)	Journal of Computational Physics		Exactly energy conserving semi-implicit particle in cell formulation	01/04/2017		HPC Community Scientific Community	n/a	International	http://dx.doi.org/10.1016/j.jcp.2017.01.002

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
216	Publication	J. Romein (Astron)	Proc. Of IEEE International Parallel and Distributed Processing Symposium (IPDPS' 17), Orlando, FL, May 2017	IEEE International Parallel and Distributed Processing Symposium (IPDPS'17), Orlando, FL, May 2017	Image-Domain Gridding on Graphics Processors	29/05/2017	Orlando, Florida	HPC Community Scientific Community		International	http://www.ipdps.org/
217	Presentation	J. Romein (Astron)	Proc. Of IEEE International Parallel and Distributed Processing Symposium (IPDPS' 17), Orlando, FL, May 2017	IEEE International Parallel and Distributed Processing Symposium (IPDPS'17), Orlando, FL, May 2017	Image-Domain Gridding on Graphics Processors	29/05/2017	Orlando, Florida	HPC Community Scientific Community		International	http://www.ipdps.org/

Table A2: List of dissemination activities											
No.	**Type of activities[5]**	**Main leader**	**Name of publication**	**Event**	**Title**	**Date/Period**	**Place**	**Type of audience[6]**	**Size of audience**	**Countries addressed**	**Link**
218	Publication	A. Pena (BSC)	Proc of ISC'17	ISC'17	Supporting Automatic Recovery in Offloaded Distributed Programming Models Through MPI-3 Techniques	01/07/2017		HPC Community Scientific Community Industry	n/a	International	

Table 2: A2. List of dissemination activities.

2.2 B. Exploitable foreground (confidential[7] information marked clearly)

2.2.1 Applications for patents, trademarks, registered designs, etc.

A patent on the direct liquid cooled solution for ULP DIMMS has been filed by Eurotech (see table below).

Table B1: List of applications for patents, trademarks, registered designs, etc.					
Type of IP Rights[8]	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Application reference (s) (e.g. EP123456)**	**Subject of title of application**	**Applicant(s) (as on the application)**
Patent	Yes	24/04/2018	102016000106718	SCHEDA ELETTRONICA REFRIGERATA	EUROTECH SpA
Patent	Yes	24/04/2018	102016000107037	Dispositivo di dissipazione termica per una scheda elettronica	EUROTECH SpA

[7] Note to be confused with the "EU CONFIDENTIAL" classification for some security research projects.

[8] A drop down list allows choosing the type of IP rights: Patents, Trademarks, Registered designs, Utility models, Others.

2.2.2 Exploitable foreground

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge	Heterogeneous architecture	No	n/a	Cluster-Booster architecture improved with a multi-level memory hierarchy	C26 - Manufacture of computer, electronic and optical products J62.0.3 - Computer facilities management activities M71 - Architectural and engineering activities; technical testing and analysis M72 - Scientific research and development	Immediate, research usage	None	JUELICH
General advancement of knowledge	Library for managing Network Attached Memory	No	n/a	libNAM	J58.2 - Software publishing J62.0.1 - Computer programming activities M72 - Scientific research and development	Immediate, research usage	To be published as Open Source (licence decision pending)	JUELICH and UHEI

[9] A drop down list allows choosing the type of foreground: General advancement of knowledge, Commercial exploitation of R&D results, Exploitation of R&D results via standards, exploitation of results through EU policies, exploitation of results through (social) innovation.

[10] A drop down list allows choosing the type sector (NACE nomenclature) : http://ec.europa.eu/competition/mergers/cases/index/nace_all.html

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge	New features in parallel I/O libraries	No	n/a	SIONlib	J58.2 - Software publishing J62.0.1 - Computer programming activities	Immediate, research usage	Open Source (BSD licence)	JUELICH
General advancement of knowledge	New features in benchmarking environment	No	n/a	JUBE	J58.2 - Software publishing J62.0.1 - Computer programming activities	Immediate, research usage	Open Source (GPLv3)	JUELICH
General advancement of knowledge	New features in checkpointing software	No	n/a	SCR	J58.2 - Software publishing J62.0.1 - Computer programming activities	Immediate, research usage	Open Source (BSD licence)	LLNL & JUELICH
General advancement of knowledge	Utilizing optional MPI features as defined by the MPI standard to implement application-level resiliency.	No	n/a	ParaStation MPI	- J58.2.9 - J62.0.1 - M72.1.9	Immediate	Yes	ParTec
General advancement of knowledge	Enhanced ParaStation MPI to support resiliency features provided to upper layers	No	n/a	ParaStation MPI	- J58.2.9 - J62.0.1 - M72.1.9	Immediate	Yes	ParTec

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
Commercial exploitation of R&D results	Commercial offering	No	n/a	Computer server for HPC and research market	C26 - Manufacture of computer, electronic and optical products J62.0.2 - Computer consultancy activities J62.0.9 - Other information technology and computer service activities	Immediate	Patents	Eurotech
General advancement of knowledge	Heterogeneous architecture & application optimisation	No	n/a	Experience enhanced in Xeon Phi architectures as well as in performance evaluation and code optimization	J62.0.1 - Computer programming activities M72 - Scientific research and development	Immediate, research usage	None	BADW-LRZ
General advancement of knowledge	Open-source memory interface controller for Hybrid Memory Cubes	No	N/A	openHMC	J63.9.9 - Other information service activities n.e.c. M72 - Scientific research and development	Immediate research usage	openHMC LGPLv3	UHEI

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge (Commercial Exploitation of R&D Results)	Network Attached Memory A compute engine attached to the EXTOLL interconnection network	Partly	N/A	NAM	M71 - Architectural and engineering activities; technical testing and analysis M72 - Scientific research and development	Immediate and near to mid term research usage	None	UHEI & EXTOLL
Commercial exploitation of R&D results	Scalable storage architecture with two tiers, cache domains and common storage	No	n/a	BeeGFS and BeeOND	J62.0.3 - Computer facilities management activities J63.1.1 Data processing, hosting and related activities M71 - Architectural and engineering activities; technical testing and analysis M72 - Scientific research and development	Immediate	BeeGFS EULA	Fraunhofer ITWM
Commercial exploitation of R&D results	I/O Programming libraries, which extends the POSIX interface to handle the cache domains	No	n/a	BeeGFS	J62.0.1 - Computer programming activities M72 - Scientific research and development	Immediate	BeeGFS EULA	Fraunhofer ITWM

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
Commercial exploitation of R&D results	data mirroring system with a built-in high availability functionality for the storage servers	No	n/a	BeeGFS	J62.0.3 - Computer facilities management activities J63.1.1 Data processing, hosting and related activities??? M71 - Architectural and engineering activities; technical testing and analysis	Immediate	BeeGFS EULA	Fraunhofer ITWM
General advancement of knowledge	Lightweight task-based C/R	No	n/a	In-memory C/R feature to transparently improve task-based programming models		Immediate, research usage	None	BSC
General advancement of knowledge	Application-based C/R based on pragmas	No	n/a	Portable interface to abstract low-level C/R APIs		Immediate, research usage	None	BSC

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge	Resiliency features for OmpSs offloading extensions	No	n/a	OmpSs offloading extensions has been enhanced with automatic resiliency features		Immediate, research usage	None	BSC & ParTec
Commercial Exploitation of R&D Results	Openly available Software to improve storage and I/O performance	No	n/a	Exascale10 I/O middleware software for boosting I/O performance exploiting compute node local storage resources	J62.0.3 - Computer facilities management activities M71 - Architectural and engineering activities; technical testing and analysis M72 - Scientific research and development J58.2 - Software publishing	Immediate usage for HPC I/O community in the near term (3-5 months) and inclusion in Seagate's exploitation plans in the medium term(6 months to 1 year)	None	Seagate

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge	The code xPic	No	n/a	Improved software for the simulation of Space Weather	J.62.01 - Computer programming activities M.72.19 - Other research and experimental development on natural sciences and engineering P.85.42 - Tertiary education	Immediate, research usage	GPLv3	KULeuven
General advancement of knowledge	Simulation software	No	n/a	Bioelectromagnetics solver enhanced with hybrid MPI/OpenMP parallelism	M72 - Scientific research and development	Immediate, research usage	None	Inria
General advancement of knowledge	New imaging application based on a novel algorithm	no	n/a	software; now being integrated in the LOFAR pipeline	J58.2.9 - Other software publishing J62.0.1 - Computer programming activities M72.1.9 - Other research and experimental development on natural sciences and engineering	immediate research usage		ASTRON

Table B2: Exploitable foreground								
Type of Exploitable Foreground[9]	**Description of exploitable foreground**	**Confidential (Yes/No)**	**Foreseen embargo date dd/mm/yy**	**Exploitable product(s) or measure(s)**	**Sector(s) of application[10]**	**Timetable, commercial or any other use**	**Patents or other IPR exploitation (licences)**	**Owner & Other Beneficiary(s) involved**
General advancement of knowledge	Parallel file system usage	No	N/A	Exploitation of I/O libraries to accelerate Lattice QCD I/O on parallel file systems	M72 - Scientific research and development	Immediate, research usage	None	UREG
General advancement of knowledge	Architecture-specific application optimization	No	N/A	Full Lattice QCD application optimized for KNL	M72 - Scientific research and development	Immediate, research usage	None	UREG

Table 3: B2. Exploitable foreground

The DEEP-ER project has performed a diversity of research and development activities. The foreground generated with them ranges from hardware to applications, passing by software, programming environment, and tools. The use and exploitation of these foregrounds in the future is specific to the scope of the owner partner/s and their roadmaps.

The industrial partners of the project will exploit their foreground to bring new products to the market, based on the designs and experience gathered in DEEP-ER, contributing to their own and Europe's economic growth. On the hardware side, the Aurora-based Booster constitutes a unique asset for Eurotech. The design of each of its components, the integration strategy, firmware, and also the experience gathered in their bring-up and operation are important results of DEEP-ER perfectly exploitable and even marketable after the project's end. Similarly the network attached memory device (NAM) developed by University of Heidelberg is a technology with a high potential for further advancements and a bright future in the HPC market. The same applies for the developments and improvements done in ParaStation by the software company ParTec and in BeeGFS by FHG-ITWM. Although both are Open Source, ParTec and the start-up ThinkParQ (spin-off of FHG-ITWM) will benefit by providing its specialised commercial support.

For the application developers the goal is not to obtain an economical benefit from the project, but to advance in knowledge, producing scientific results, and finally transmit it to the next generation through education.

The computing centres JSC and BSC will exploit the software that they have developed and/or improved in the project: SIONlib, libNAM, SCR, JUBE, OmpSs, Extrae/Paraver, and Dimemas). As public institutions both follow and are committed to Open Source initiatives and provide support free-of-charge to their user communities. The mentioned components are also used for training and educational purposes.

The universities and centres that have participated in DEEP-ER as application developers, have been able to modernise their codes, making them more scalable and performant, and preparing them for running on the next generation HPC systems. In consequence they will be potentially able to obtain more scientific results and publications.

The following subsections describe in more detail the nature and exploitation potential of the most important foreground achieved in the DEEP-ER project, grouped according to the partners which are largest responsible for them.

2.2.2.1 JUELICH

The Jülich Supercomputing Centre (JSC) of Forschungszentrum Jülich GmbH is the largest Supercomputing Centre in Germany. As a member of the German Gauss Centre for Supercomputing – the German national HPC institution – JSC provides access to Tier-0 computing resources through PRACE and national mechanisms. Nevertheless, JSC is not simply running and operating large-scale HPC systems. At the same time it has also a renowned reputation for co-developing – together with provider industry partners – production-ready systems based on the latest and most promising technologies. Remarkable examples are the JUROPA and QPACE1 clusters. Amongst others, JSC contributed in their construction with the definition of the systems architecture and their network topology. Beyond that, the design and development of the system management middleware and the system's software stacks are executed under JSC's supervision. JSC's goal with this kind of activities is to keep up with technological evolution while making sure that the production

systems procured by the computing centre will perfectly match the needs of its user community.

The DEEP projects family constitute an important step forward on this strategic path. JSC proposed the Cluster-Booster concept due to its potential for enabling a wider range of applications reaching higher scalability, while keeping the system cost and energy consumption at bay. Furthermore, from the system administrator's point of view, its full flexibility in the assignment of standard multi-core and accelerator resources to each application allows for maximising the overall usage of the hardware, avoiding any idling or underused resources. Results obtained in DEEP and DEEP-ER encouraged the JSC team to extend the concept to the more general Modular Supercomputing architecture. The multi-level hierarchy of (partially) non-volatile memory and the I/O and resiliency functionalities developed by DEEP-EST position enlarge the portfolio of applications able to profit from the modular architecture, making it even more interesting for a data centre like the JSC.

The DEEP and DEEP-ER projects and the experience gathered therein serve JSC to demonstrate the benefits of its innovative Cluster-Booster concept and its potential for future HPC systems. The deployment of the JURECA cluster in 2015 is the first step of the realisation of the concept at production level. During 2017, JUELICH plans attaching to the existing JURECA cluster a production Booster built with Intel Xeon Phi processors. The knowledge acquired by building and running the DEEP and DEEP-ER platforms will be exploited in the setup and operation of the future JURECA Cluster-Booster system.

Applications will be able to run across both Cluster and Booster parts of JURECA, exploiting the offload process developed in DEEP. The Open Source software stack developed in the DEEP and DEEP-ER projects already provides the environment required to run the JURECA production system. ParaStation, a shared development from the ParaStation consortium in which JSC is actively involved, is the cluster management software of the current production cluster systems at JSC. Thanks to the developments done in DEEP and DEEP-ER, ParaStation is already today a solid foundation for the JURECA Cluster-Booster production system.

The applied co-design methodologies – based on direct contact between hardware, software and application developers – and even the management strategies developed and successfully employed in DEEP-ER constitute an invaluable add-on to JSC's already demonstrated expertise as coordinator of large R&D projects. These will be directly applied into the DEEP-EST project and beyond. JSC and its partners will profit in current and future projects from the knowledge and experience acquired in DEEP-ER.

But the experience gathered by JSC is not limited to how best manage and coordinate a large consortium. DEEP-ER is a large, multi-disciplinary R&D project in which all its participants had the rare opportunity of closely interacting with experts from across the whole spectrum of HPC fields. The young, very talented JSC members participating in DEEP-ER got in this way a global view on HPC technology that, together with the networking opportunities and visibility provided by the project, opens them new perspectives for their future development in the HPC world. JSC itself profits from their enlarged skills, their gained maturity and experience, and their strong motivation to continue putting their talent in the institute's service.

Finally, the DEEP projects family has notably contributed to increase JSC's visibility and reputation as an internationally leading HPC centre. The dissemination activities organised in

the DEEP and DEEP-ER contexts have enabled JSC to present its in-house developments, ideas, strategies and roadmap to the HPC community and key players therein. In numerous events in which the institute has presented the concept and results of the projects, new opportunities of collaboration have emerged with key-players on the international community. These will open new opportunities for JSC in its dedicated mandate to advance science, industry and society through High Performance computing, increasing the European expertise in the field.

2.2.2.2 ParTec

Building on the achievements made in the DEEP project, ParTec has generated exploitable foreground in DEEP-ER especially in the advancements of ParaStation MPI, the communication and cluster operation library for heterogeneous systems. Within DEEP-ER, the focus was on adding resiliency features, to make the most of the DEEP-ER hardware architecture.

ParaStation MPI comprises three basic components:

1. A low-level communication library (pscom), supporting various high-speed interconnects like Extoll, and providing integral communication facilities for upper layers.
2. The ParaStation MPI library and runtime (psmpi), providing the well-known and established Message Passing Interface to parallel applications.
3. The ParaStation Management (psmgmt) is responsible for starting and controlling parallel applications including process pinning, signal handling, I/O forwarding, accounting, proper cleanup and more. Pluggable modules are used for the integration with the batch system.

Following the spirit of the DEEP project, we strived, for adding new features, this time to support resiliency features in higher-level libraries and runtimes like SCR and OmpSs, to stay within or atleast as close as possible to the MPI standard.

Following requirements from SCR and OmpSs, the resiliency features in ParaStation MPI are focused around the MPI_Comm_spawn mechanism. In a first step, we added the ability to disconnect spawned process groups, powered by the new asynchronous connect anddisconnect capability of pscom.

In a second step, ParTec utilized the rarely implemented, yet very useful option provided by the MPI standard to handle communication problems in a non-fatal way. Normally, any communication problem affecting a parallel application leads to termination of the entire application. However, the MPI standard foresees the option to propagate communication problems to the application by means of error codes from MPI calls and error handlers.

The required extensions within ParaStation MPI for implementing this option were two-fold:

- On the one hand, the process manager had to be made aware of the notion of process groups, as they are formed in terms of spawned MPI kernels. Building upon this, we added user-controllable, fine-grained policies for handling problems in parent and child groups independently. The option to let parent groups survive the termination of one of their children groups is heavily used by OmpSs to restart offloaded, highly-scalable code parts.

- On the other hand, the pscom lower-level communication layer had to be made capable of detecting and reporting broken network connections, as they may indicate dead peer processes within the offloaded tasks. Consequently, such broken network connections should not influence in any way network connections between unrelated peers.

By now, both of these extensions - the process group awareness as well as the connection guard feature - have finally left experimental stage and already are now part of the common ParaStation release branch.

2.2.2.3 Intel

For partner Intel, the exploitable foreground includes the following:

- *Intel Xeon Phi compiler and library bug fixes and optimisations*: As a result of Intel's involvement in WP6 and the application support team, a number of bugs were identified in the compiler, threading and math libraries for the 2nd product generation of Intel Xeon Phi, and optimisations were proposed in light of the DEEP-ER application requirements and porting/tuning experience. The bugs have been fixed outside of the project, and optimisations were added to the Intel Xeon Phi compiler and libraries. This has contributed to improving the quality and performance achieved by said software components marketed by Intel.
- *Proof points for the 2nd generation of Intel Xeon Phi*: Some of the applications ported and optimised in work package 6 will serve as proof points for the substantial performance improvements compared to the first Intel Xeon Phi generation as well as to state-of-the-art Intel Xeon systems.
- *Proof points for fast local storage-class memory*: One of the main innovations of DEEP-ER is the integration of fast storage-class memory with the Cluster and Booster nodes, based on Intel SSD solutions attached via PCI Express and supporting the NVMe protocols. The seamless integration of these devices into the BeeGFS parallel file system, the SIONLib I/O concentrator library, the SCR checkpointing system and the E10 implementation of MPI-IO enables straightforward use of this key innovation, and results from work packages 3 and 6 clearly show the substantial performance gains compared to a state-of-the-art storage system. These proof points will support the creation of a market for HPC nodes with integrated storage-class memory, which Intel is keen to service with the next generation of PCIe-attached SSD systems (Intel® Optane™ technology).
- *High redundancy local checkpointing*: The "buddy" checkpointing system prototyped in the DEEP-ER project substantially reduces the overhead of creating regular checkpoints (since it uses local fast storage), while providing redundancy that covers the case of single node failures. The approach carries over to systems that have I/O "burst buffers" at a subset of nodes, and it will profit from advances in Intel storage technology (Intel 3D XPoint).

2.2.2.4 BADW-LRZ

The Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (BADW-LRZ) is the general IT service provider for the Munich universities and a growing number of research institutions across Bavaria. On top, it is one of the three German Tier-0 HPC providers organised in the Gauss Centre for Supercomputing. SuperMUC, LRZ's prime

HPC resource, reaches 6 PFlop peak compute performance. Yearly, close to 3000 academic users are granted access to the machine, one fifth of which are European researchers admitted through the Partnership for Advanced Computing in Europe (PRACE).

As a Tier-0 HPC provider, BADW-LRZ has a strong interest in enabling researchers to best exploit the massive computational resources offered. To achieve this, BADW-LRZ has a long-standing tradition in training application developers and in application support. For the latter, efforts have been increased by substantially expanding BADW-LRZ's high-level technical assistance. One of the initiatives here is to focus on domain specific and community-oriented support and to provide the matching research infrastructure. These efforts are organised in so called 'application-labs', which have been set up for five domains: (1) astro- and plasma-physics, (2) big-data sciences, (3) biology & life sciences, (4) computational fluid dynamics and last but not least (5) geo- and environmental sciences.

Against this background, it was a natural fit for BADW-LRZ to join the DEEP-ER project in the field of application experience. The focus has been on the application SeisSol, a state of the art research code for earthquake faulting and scenario simulations. The code became part of the BADW-LRZ HPC benchmarking suite and as such a tool for the evaluation of new platforms in procurement processes of next generation installations. BADW-LRZ also supports the key developers of the SeisSol-team from Technical University of Munich and Ludwig-Maximilians-Universität Munich in the framework of an Intel Parallel Computing Center (IPCC) and the SuperMUC extreme scaling workshops.

Within WP 6 of the DEEP-ER project, BADW-LRZ took the opportunity to strengthen its knowledge base in code performance evaluation, code optimisation, as well as evaluating emerging technologies and concepts through the co-design process of the project. In more concrete terms BADW-LRZ worked amongst others on improving I/O efficiency taking into account the DEEP-ER I/O software stack and therefore making the code more resilient and scalable.

The experience gained will directly stimulate developments and support provided by BADW-LRZ's application lab to other research codes, workflows, and infrastructures in the field of geology and environmental sciences. It helps BADW-LRZ to enable application developers in bringing their codes to the next level and be prepared to use the next generation of large-scale compute-cluster installations. Last but not least, the acquired knowledge can be passed on directly to application developers and the next generation of HPC experts through the extensive course programme BADW-LRZ offers e.g. with the extreme scaling workshops or the PRACE PATC courses.

2.2.2.1 University of Heidelberg – EXTOLL

The Computer Architecture Group (CAG) of UniHD was involved in WP3 and has carried out leadership of WP7. In WP3, CAG developed the Network Attached Memory (NAM) as an additional component that can be seamlessly integrated with the EXTOLL network. The NAM can be used as a remote, shared memory device or to perform computation. In DEEP-ER, it acts as an active component to autonomously fetch checkpoint data from nodes that participate in a checkpoint/restart process in order to improve resiliency. The NAM then calculates parity over all checkpoints which can be used to recover the system from any single node failure. Since NAMs can be attached to any unused link of an EXTOLL NIC the maximum number of NAMs in a system scales with the number of nodes.

Development of the NAM required a strong cooperation between WP4 and WP3 since features had to be defined upon the user's requirements and technical feasibility. As a result, a library called libNAM was developed. libNAM provides an easy to use interface for programmers to explore the capabilities of the NAM. For checkpoint / restart, libNAM was integrated with the scalable IO library SIONlib.

The NAM concept will be adopted in the DEEP-EST project. While the actual application may differ, basic building blocks such as the memory controller and the FPGA EXTOLL link implementation can be used directly and will lower the risk associated with developing a new prototype.

The NAM prototype was prominently displayed at several conferences and poster and presentation sessions which led to several, academia based research collaborations for UHEI. It also increased public awareness of the relatively new EXTOLL network technology and extended its portfolio by a state of the art application.

EXTOLL, as the provider of the Tourmalet ASIC NIC has continued to develop and improve corresponding software and network management tools. Multiple, non-uniform network topologies within the DEEP-ER SDV and the final prototype required very complex routing algorithm to correctly setup the network. Therefore EXTOLL investigated and implemented a software component for multi-topology capable, deadlock-free networks. This component will ease the setup of future heterogeneous systems with non-uniform topologies.

2.2.2.1 FHG-ITWM

The Fraunhofer-Gesellschaft is the leading organisation for applied research in Europe. Its research activities are conducted by 69 institutes and research units at locations throughout Germany, each focusing on different fields of applied science. The Fraunhofer Institute for Industrial Mathematics ITWM supports companies in the development and optimisation of products, services, communication and working processes.

In the context of the DEEP-ER project the scalable storage architecture with a common cluster storage and cache domains were implemented. Hard-drives provide the high capacity for globally shared data. To avoid as much communication as possible with this centralised layer of globally shared storage servers, a cache layer of subdomain storage servers (cache domains) based on new storage technologies like NVM is used. The cache layer allows applications to write data that is not shared with other subdomains (i.e. non-coherent) and thus provides the ability to scale I/O performance linearly with the number of compute nodes.

In the DEEP-ER project two I/O programming libraries were developed to improve the I/O performance of the applications and benefit from the new storage architecture. The BeeGFS striping API allows creating files in the filesystem with a data stripe configuration that matches the I/O requirements of the application. The cache API implements flush and prefetch operations to copy the data between the cache domains and the global storage. The flush and prefetch operations can be done synchronously or asynchronously. The asynchronous implementation allows improving the performance compared to the synchronous version in several ways. First the flush or prefetch functions are offloaded to a cache daemon, which is running on the same node. A second performance improvement is implemented by splitting big files into byte range flushes or prefetches to use multiple threads. The striping API is still part of the beegfs-client-devel package. The package of the cache API will be integrated into our release repository. New features will be added to the

APIs on customer request and also in other research projects. BeeGFS can offer customers new solutions to solve their I/O problems with the two API and the new I/O architecture.

A data mirroring system has been implemented, featuring a built-in High Availability (HA) functionality in order to achieve better resiliency and reliability. Using this system, the data is still accessible to the cluster even if a node that contains a mirrored file should fail. Data mirroring on storage targets is based on so-called buddy groups, in general consisting of two targets, which should be located on different storage servers. In the meantime metadata mirroring with high availability functionality is implemented. In the next step the management daemon gets the build-in high availability feature to complete the build-in high availability functionality of BeeGFS. The storage and metadata HA is used by several BeeGFS customers world-wide. The build-in HA solution allows the usage of BeeGFS in business critical 24/7 environments. It opens a new market for BeeGFS.

2.2.2.2 Eurotech

Eurotech designed the DEEP-ER Booster Prototype and has announced the commercial product “Aurora Tigon v4” based on this development. While Aurora is the primary exploitable foreground from the DEEP-ER project, there is a number of advantages Eurotech gained from participating in the project:

2.2.2.2.1 Development of 3D integration for 3rd party boards within the Aurora chassis

Eurotech’s previous products were based on boards of its own design that used direct liquid cooling. Within the DEEP-ER project Eurotech successfully adopted Intel’s S7200AP (“Adams Pass”) board for direct cooling within the Aurora chassis. The challenge is the heat extraction from vertically oriented sources, such as the DIMMs on the industry standard board in a very dense packaging. The application of the direct cooling principle to the LPDIMMs (Low Profile DIMMs) is an industry first. To complete this task Eurotech developed a 3D model of the board with the heat extraction elements.

This successful implementation of the 3D integration opens up possibilities for faster adoption of new technological capabilities in future Eurotech products.

2.2.2.2.2 Experience with high speed PCB design

Eurotech has always designed products with high speed backplanes and interconnects. Case in point is the development of the DEEP backplane. Regarding the Eurotech design for the DEEP-ER Booster the challenge was to implement the 4.6 Tbit/s backplane to connect the peripheral devices to the corresponding nodes. To master tis challenge, Eurotech built up internal expertise for this kind of design, which requires also the inside manufacturing. With this knowledge Eurotech is well positioned for the next generation of designs. In particular, in the follow-up DEEP-EST project it will be important to implement a state of the art interconnect.

2nd generation direct water cooling

Eurotech did pioneer the direct water cooling approach for all its products about 10 years ago. In the DEEP-ER Booster prototype we have used newer manufacturing techniques that allow implementing this principle at lower cost. While the principle of the new technology was verified in previous projects (i.e. QPACE2), we have pushed the technology in DEEP-ER to include much larger board area for the cooling. Looking forward, this technology will be used in the follow up products using the direct water cooling principle.

2.2.2.2.3 Chassis controller with embedded network

Eurotech implemented the chassis controller scheme at two levels: low level control that collects sensor data and supervises the operational safety, and high-level control with a Linux mini-computer. In addition, the administrative Gigabit Ethernet network is also implemented within the same framework. This combination makes extensive monitoring of the machine possible. The low level functions perform independent power and reset on the NICs, which facilitates the maintenance of a fully working HPC network even when nodes are not present. This operational principle has proven itself in this project and Eurotech will maintain this principle in future generation products.

2.2.2.3 BSC

BSC has a key role as a scientific and technological player in the research field in Europe.

Due to its research centre vocation, the first interest is to exploit the DEEP-ER project results for producing innovative outstanding scientific products, in the form of Open Source software and publications, contributions to conferences and workshops, with a focus on programming models, resiliency and analysis tools for HPC systems. The possibility of performing experiments on the innovative DEEP-ER systems can lead to important evidences for the development of future programming models with improved resiliency support for highly heterogeneous HPC platforms.

Within the system software, a very promising role is played by the OmpSs programming model that, thanks to the Mont-Blanc 1/2/3 and DEEP/-ER projects and leveraging its open source model, can now count on a wider community of users and developers considering it as de-facto forerunner of the OpenMP standard. The innovative offloading extensions developed in the context of DEEP, and novel resiliency features (such as lightweight task-based checkpoint-restart (C/R), application-based C/R based on pragmas and the enhanced offloading extensions with resiliency features) developed in DEEP-ER will further increase the impact and outreach of OmpSs, especially on heterogeneous clusters.

Performance tools are another key topic for BSC and the DEEP-ER project offered a framework to demonstrate the benefits of using BSC tools and models. In DEEP-ER the BSC efficiency model has been extended to obtain a diagnosis of the impact of I/O on the application performance and efficiency.

The last vocation of BSC is about liaising with very high profile industrial and technological companies at national and international level. The knowledge in developing programming models, analysis tools and projections for the DEEP System, gained by collaborating with application developers, will enable BSC to be a leader in the development of programming models and tools, not only in the HPC environment, but also in the industrial and academic sectors. This will help BSC obtaining more valuable collaborations and regional, national and international funding in the future, e.g. EU H2020.

2.2.2.4 Seagate

Seagate has worked on the I/O part of the DEEP-ER project, which was one of the proposed key extensions to the DEEP Architecture. Before the project started, compute node local NVRAM technologies needed to be handled explicitly by applications (there was no automatic mechanism for users to easily employ these memories in their codes). In DEEP-ER Seagate has provided automatic mechanisms (through MPI-IO hints) to select the

preferred storage technology to be used (local NVMs vs global HDDs) by the application. Seagate demonstrated that particularly for the collective I/O use case, performance could be greatly improved when using local NVRAM devices.

The DEEP-ER enabled improvements (Exascale10 extensions for DEEP-ER) have been implemented in the ROMIO middleware (a popular implementation of the MPI-IO specifications) and have been made openly available to the community.

In the near term discussions are in place for including this code into the upstream ROMIO release, as part of MPICH, which implements the ubiquitous and most prevalent MPI programming model for HPC. Furthermore, Seagate will look to exploit the outcomes of DEEP-ER to improve the performance of its future products post DEEP-ER, especially in the area of utilising node-local storage resources to drastically boost I/O performance.

Seagate has also gathered very valuable inputs from DEEP-ER use cases – as it continues to work on architecting storage solutions for Exascale. Participation in DEEP-ER has indeed strengthened relations with key European players (e.g.: JUELICH) which has been carried over to some of its H2020 initiatives.

2.2.2.5 CINECA

CINECA is the largest computer centre in Italy and plays a key role in providing computer resources and technological innovation to both national users and to the international community research through infrastructures such as PRACE. As an application developer of DEEP-ER, the principal exploitable foreground of the project has been in the improvements and optimisations of the two quantum Monte Carlo simulation codes, TurboRVB and 2degas. These can be summarised as follows:

- TurboRVB: Thanks to the project this program has been extensively re-factored and now has a checkpoint facility either via rank-local files or SIONlib. In addition, it is also possible to use SCR for checkpointing.
- 2degas: This program has been re-factored with modern Fortran90 constructs and converted from an MPI application to an OpenMP code; a hybrid MPI/OpenMP version will be available soon. A hybrid version of the program will be extremely useful for researchers running on exascale architectures.

Beyond these application codes, the exploitable foreground of DEEP-ER has allowed CINECA to improve its services in many ways, the most important being the use and programming of the Intel KNL processor. This experience has been exploited recently with the introduction of the cluster Marconi at CINECA, currently the largest KNL cluster in Europe, to guide the configuration of the KNL nodes. For example, a comparison of some performance benchmarks between the DEEP-ER and CINECA KNL installations helped to identify a misconfiguration of Marconi which could then be corrected. Since a large partition of Marconi has been allocated to PRACE Tier-0 calls, as well as to other large projects such as EuroFusion, this sort of improvement has benefited many European HPC users directly. Thanks to DEEP-ER it has also been possible to enhance our understanding of programming models such as MPI and OpenMP and the importance of task and thread affinity in application performance. Many of the lessons learnt in KNL programming have been integrated in courses we give at CINECA, for example those forming part of the PATC (PRACE Advanced Training Centre) programme.

Other foreground being considered or exploited at CINECA, or in collaborations in which CINECA is involved, has derived from the innovative hardware and software technologies developed or investigated in DEEP-ER. For example,

- SIONlib and SCR for checkpointing and resiliency.
- The use of NVM. Similar technology is already in use on CINECA's PICO system and comparison between the DEEP-ER system and PICO was helpful in using the NVM technology on PICO.
- Performance and profiling tools such as Extrae, Vtune and Scalasca. Although we were aware of such tools before the project, DEEP-ER allowed us to gain more competence in their use and suggest them as tools for our users.

The availability of experts in all fields of HPC from engineers to middleware developers to application experts, in industry and academia, has resulted in a transfer of knowledge and expertise, which would have been impossible without being part of this project.

2.2.2.6 KULeuven

The Centre for mathematical Plasma-Astrophysics (CmPA) is one of the research units of the Mathematics Department of KU Leuven University in Belgium. The centre produces basic scientific research on astrophysical plasmas, from the environment around the planets, to the Sun and the stars. The unit is also very actively involved in the education of students from all the university, and offers specialized courses in advanced plasma physics, solar physics, astrophysics and space weather.

The team lead by professor Giovanni Lapenta works on subjects related to Space Weather. All plasma effects observed in the interaction between the Sun and the Earth (or other planets) are studied in this group, including magnetic reconnection, solar wind-magnetosphere interaction, plasma turbulence, particle-fields interactions, among others. The team's main research tools are the computer simulations of different plasma flows of interest. The codes developed by the team perform self-consistent numerical simulations of different astrophysical plasmas, modelling the small-scale movement of individual electrons and the flow of plasma at the scale of planets.

To achieve the goal of performing realistic simulations of plasma flows for Space Weather research, the team relies on the Particle-in-Cell (PIC) method. Simulations performed with this method are computationally intensive and require the use of the biggest supercomputers available.

The DEEP-ER project has allowed refining the concept of Cluster-Booster Architecture for the PIC code of KULeuven. Today the code resolves the equations of electromagnetism in the Cluster while transporting billions of particles in the Booster. This allows performing simulations much faster than in traditional systems. The inclusion of multiple levels of parallelism, from vectorisation and multi-threading inside a compute node, to message passing among nodes, has allowed to better make use of the computational resources available in DEEP-ER and outside the project in other computing centres.

The code can also produce a large amount of data and simulations frequently require multiple restarts. While information about the electromagnetic fields is small, particle data can generate hundreds of Gigabytes of information in one writing iteration. Such files need to be administered efficiently by the code and the file system. Thanks to the I/O and checkpoint new developments in DEEP-ER, the PIC code is much more efficient for the researchers.

KULeuven is currently using the DEEP-ER Prototype to perform simulations of Space Weather physics that could not be performed before. We expect that new scientific insights will be available soon thanks to the accomplishments in the DEEP-ER project and thanks to the use of the DEEP-ER computer system.

2.2.2.7 Inria

The project team from Inria that has participated in the DEEP-ER project is developing the GERShWIN (discontinuous GalERkin Solver for microWave INteraction with biological tissues) simulation software for the numerical modelling of microwave interaction with biological tissues. The work undertaken in the DEEP-ER project has allowed to substantially increase the high performance capabilities of this software, in particular through its hybrid MPI/OpenMP parallelisation and algorithmic optimisation to the hardware characteristics of the DEEP-ER Cluster-Booster architecture. Application-level checkpointing have also been implemented in view of introducing some resiliency features in GERShWIN. Finally, input/output performance has also been improved through the integration of SIONlib and the exploitation of NVMe. Overall, the new version of GERShWIN offers new functionalities and capabilities that will be very useful for performing electromagnetic wave propagation simulations at Exascale.

In the future, the new version of GERShWIN will be exploited for research activities in relation with the study of potential adverse effects of human exposure to microwave form wireless systems. Thanks to the improved scalability and overall computational efficiency of the DGTD (Discontinuous Galerkin Time-Domain) solver implemented in this software, parametric studies will be possible in shorter wall clock times. The result of these studies will be made available to the community of scientists actively involved in the research field at the international level, through publications in dedicated journals. More importantly, the know-how acquired on the DGTD method at the heart of the GERShWIN software will be exploited for a new generation simulation software developed by the same Inria group for nanoscale light/matter interactions in relation with nanophotonics.

2.2.2.8 ASTRON

Within ASTRON, the main exploitable result is the new Image-Domain-Gridding application, which was developed with a significant contribution from DEEP-ER. The developed code was integrated and is being tested in the WSClean imager, which is used to image data from LOFAR, world's largest low-frequency radio telescope. This will allow LOFAR to create images not only more efficiently, but also with a wider field of view than was previously possible.

The work on the imager resulted in a paper that is accepted for IPDPS'17 – one of the top-ranking conferences in HPC – and to which, again, DEEP-ER made a significant contribution. We prepare an extended journal version of this paper, which includes a thorough performance and energy efficiency analysis of this application running on the Xeon Phi.

In a follow-up project, we will port the Image-Domain Gridder to FPGAs using Intel's new OpenCL/FPGA framework, to achieve even higher energy efficiency. The imaging code of the future SKA telescope poses Exascale requirements, making it an important HPC application.

2.2.2.9 UREG

Explanation to the two items included by UREG in Table 3 (B2).

- We evaluated the use of BeeGFS in general for I/O for Lattice QCD, and explored the use of libraries that enable higher performance for parallel I/O. The MPIWrap library can be used directly with our existing HDF5 routines, which allows easy integration in existing code just by re-linking. This can be generally applied to systems that provide suitable parallel file systems. We have therefore chosen BeeGFS as the main file system for the storage of your production system QPACE-3. This system will serve 360 TeraBytes for users of QPACE-3 via BeeGFS. For redundancy, this system consists of 4 OSSes, 2 MDSs and a direct attached SAS storage, using Linux' multipath facilities and BeeGFSs processes for failover.
- The Chroma Lattice QCD application has been ported to and optimised for the KNL architecture. Chroma is a well-established software package with ample features for most aspects of Lattice QCD. This package can now take full advantage of KNL-based machines. The full Lattice QCD software stack can now run efficiently on KNL machines, providing much needed computing power to the community. Architectures with the same instruction set will also benefit directly from this work, and it is expected that the work done here will facilitate porting to future CPUs with similar vector instructions. Software ported for the KNL will run on our production systems QPACE-2 (KNC based) and QPACE-3 (KNL based). As we will stick with this codebase in the foreseeable future, the experiences gained hereby will help us to enable Chroma for emergent architectures.

3 Report on societal implications

A General Information *(completed automatically when **Grant Agreement number** is entered.*

Grant Agreement Number:	**610476**
Title of Project:	**DEEP Extended Reach**
Name and Title of Coordinator:	**Prof. Dr. Dr. Thomas Lippert**

B Ethics

1. Did your project undergo an Ethics Review (and/or Screening)? • If Yes: have you described the progress of compliance with the relevant Ethics Review/Screening Requirements in the frame of the periodic/final project reports? Special Reminder: the progress of compliance with the Ethics Review/Screening Requirements should be described in the Period/Final Project Reports under the Section 3.2.2 'Work Progress and Achievements'	○ ⊗	**Yes** **No**
2. Please indicate whether your project involved any of the following issues (tick box) :		***NO***
RESEARCH ON HUMANS		
• Did the project involve children?		
• Did the project involve patients?		
• Did the project involve persons not able to give consent?		
• Did the project involve adult healthy volunteers?		
• Did the project involve Human Genetic Material?		
• Did the project involve Human biological samples?		
• Did the project involve Human data collection?		
RESEARCH ON HUMAN EMBRYO/FOETUS		
• Did the project involve Human Embryos?		
• Did the project involve Human Foetal Tissue / Cells?		
• Did the project involve Human Embryonic Stem Cells (hESCs)?		
• Did the project on human Embryonic Stem Cells involve cells in culture?		
• Did the project on human Embryonic Stem Cells involve the derivation of cells from Embryos?		
PRIVACY		
• Did the project involve processing of genetic information or personal data (eg. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)		
• Did the project involve tracking the location or observation of people?		
RESEARCH ON ANIMALS		

• Did the project involve research on animals?	
• Were those animals transgenic small laboratory animals?	
• Were those animals transgenic farm animals?	
• Were those animals cloning farm animals?	
• Were those animals non-human primates?	
RESEARCH INVOLVING DEVELOPING COUNTRIES	
• Did the project involve the use of local resources (genetic, animal, plant etc)?	
• Was the project of benefit to local community (capacity building, access to healthcare, education etc)?	
DUAL USE	
• Research having direct military use	
• Research having the potential for terrorist abuse	

C Workforce Statistics

3 Workforce statistics for the project: Please indicate in the table below the number of people who worked on the project (on a headcount basis).

Type of Position	Number of Women	Number of Men
Scientific Coordinator	1	1
Work package leaders	4	6
Experienced researchers (i.e. PhD holders)	4	33
PhD Students	0	3
Other	10	25

4 How many additional researchers (in companies and universities) were recruited specifically for this project?	**12**
Of which, indicate the number of men:	10

D Gender Aspects

5 Did you carry out specific Gender Equality Actions under the project ? Note: Most of the partners have defined and implement their own gender action plans. For instance, job openings encouraging applications by female staff, special work-life-balance, and equal chances programs.	○ ⊗	Yes No

6 Which of the following actions did you carry out and how effective were they?

	Not at all effective				Very effective
❑ Design and implement an equal opportunity policy	○	○	○	○	○
❑ Set targets to achieve a gender balance in the workforce	○	○	○	○	○
❑ Organise conferences and workshops on gender	○	○	○	○	○
❑ Actions to improve work-life balance	○	○	○	○	○

○ Other:

7 Was there a gender dimension associated with the research content – i.e. wherever people were the focus of the research as, for example, consumers, users, patients or in trials, was the issue of gender considered and addressed?

○ Yes- please specify:
⊗ No

E Synergies with Science Education

8 Did your project involve working with students and/or school pupils (e.g. open days, participation in science festivals and events, prizes/competitions or joint projects)?

○ Yes- please specify
⊗ No

9 Did the project generate any science education material (e.g. kits, websites, explanatory booklets, DVDs)?

⊗ Yes- please specify: poster explaining the Cluster-Booster concept for children at the “Tag der Neugier” (open doors day) at JUELICH.
○ No

F Interdisciplinarity

10 Which disciplines (see list below) are involved in your project?

⊗ Main discipline[11]: 1.1 (Computer Science)

⊗	Associated discipline[11]: 2.2 (Com. Engin.)	⊗	Associated discipline[11]: 1.2, 1.3, and 1.4, as pilot applications. 5.3 (HPC training)

G Engaging with Civil society and policy makers

11a Did your project engage with societal actors beyond the research community? *(if 'No', go to Question 14)*	⊗ ○	Yes No

[11] Insert number from list below (Frascati Manual)

11b If yes, did you engage with citizens (citizens' panels / juries) or organised civil society (NGOs, patients' groups etc.)?

- ○ No
- ○ Yes- in determining what research should be performed
- ○ Yes - in implementing the research
- ⊗ Yes, in communicating /disseminating / using the results of the project

11c In doing so, did your project involve actors whose role is mainly to organise the dialogue with citizens and organised civil society (e.g. professional mediator; communication company, science museums)?	○	Yes
	⊗	No

12 Did you engage with government / public bodies or policy makers (including international organisations)

- ○ No
- ⊗ Yes- in framing the research agenda
- ⊗ Yes - in implementing the research agenda
- ⊗ Yes, in communicating /disseminating / using the results of the project

13a Will the project generate outputs (expertise or scientific advice) which could be used by policy makers?

- ○ Yes – as a **primary** objective (please indicate areas below- multiple answers possible)
- ⊗ Yes – as a **secondary** objective (please indicate areas below - multiple answer possible)
- ○ No

13b If Yes, in which fields?

Agriculture		Energy		Human rights	
Audiovisual and Media		Enlargement		Information Society	X
Budget		Enterprise		Institutional affairs	
Competition		Environment		Internal Market	
Consumers		External Relations		Justice, freedom and security	
Culture		External Trade		Public Health	
Customs		Fisheries and Maritime Affairs		Regional Policy	
Development Economic and Monetary Affairs		Food Safety		Research and Innovation	X
Education, Training, Youth		Foreign and Security Policy		Space	
Employment and Social Affairs		Fraud		Taxation	
		Humanitarian aid		Transport	

13c If Yes, at which level?

- ○ Local / regional levels
- ⊗ National level
- ⊗ European level
- ○ International level

H Use and dissemination

14 How many Articles were published/accepted for publication in peer-reviewed journals?	**8**
To how many of these is open access[12] provided?	**2**
How many of these are published in open access journals?	**0**
How many of these are published in open repositories?	**2**
To how many of these is open access not provided?	**6**
Please check all applicable reasons for not providing open access:	
☒ publisher's licensing agreement would not permit publishing in a repository ❑ no suitable repository available ☒ no suitable open access journal available ❑ no funds available to publish in an open access journal ❑ lack of time and resources ❑ lack of information on open access ❑ other: ……	

15 How many new patent applications ('priority filings') have been made? *("Technologically unique": multiple applications for the same invention in different jurisdictions should be counted as just one application of grant).*		**2**
16 Indicate how many of the following Intellectual Property Rights were applied for (give number in each box).	Trademark	**0**
	Registered design	**0**
	Other	**0**
17 How many spin-off companies were created / are planned as a direct result of the project?		**0**
Indicate the approximate number of additional jobs in these companies:		

18 Please indicate whether your project has a potential impact on employment, in comparison with the situation before your project:

☒	Increase in employment, or	☒	In small & medium-sized enterprises
❑	Safeguard employment, or	☒	In large companies
❑	Decrease in employment,	❑	None of the above / not relevant to the project
❑	Difficult to estimate / not possible to quantify		

19 For your project partnership please estimate the employment effect resulting directly from your participation in Full Time Equivalent (*FTE = one person working fulltime for a year*) jobs:	*Indicate figure:*
Difficult to estimate / not possible to quantify	☒

[12] Open Access is defined as free of charge access for anyone via the internet.

I Media and Communication to the general public

20 As part of the project, were any of the beneficiaries professionals in communication or media relations?

⊗ Yes ○ No

21 As part of the project, have any beneficiaries received professional media / communication training / advice to improve communication with the general public?

○ Yes ⊗ No

22 Which of the following have been used to communicate information about your project to the general public, or have resulted from your project?

☒	Press Release	☒	Coverage in specialist press
☒	Media briefing	☒	Coverage in general (non-specialist) press
❑	TV coverage / report	❑	Coverage in national press
❑	Radio coverage / report	❑	Coverage in international press
☒	Brochures /posters / flyers	☒	Website for the general public / internet
☒	DVD /Film /Multimedia	☒	Event targeting general public (festival, conference, exhibition, science café)

23 In which languages are the information products for the general public produced?

☒	Language of the coordinator	☒	English
☒	Other language(s)		

***Question F-10*:** Classification of Scientific Disciplines according to the Frascati Manual 2002 (Proposed Standard Practice for Surveys on Research and Experimental Development, OECD 2002):

FIELDS OF SCIENCE AND TECHNOLOGY

1. NATURAL SCIENCES

1.1 Mathematics and computer sciences [mathematics and other allied fields: computer sciences and other allied subjects (software development only; hardware development should be classified in the engineering fields)]

1.2 Physical sciences (astronomy and space sciences, physics and other allied subjects)

1.3 Chemical sciences (chemistry, other allied subjects)

1.4 Earth and related environmental sciences (geology, geophysics, mineralogy, physical geography and other geosciences, meteorology and other atmospheric sciences including climatic research, oceanography, vulcanology, palaeoecology, other allied sciences)

1.5 Biological sciences (biology, botany, bacteriology, microbiology, zoology, entomology, genetics, biochemistry, biophysics, other allied sciences, excluding clinical and veterinary sciences)

2 ENGINEERING AND TECHNOLOGY

2.1 Civil engineering (architecture engineering, building science and engineering, construction engineering, municipal and structural engineering and other allied subjects)

2.2 Electrical engineering, electronics [electrical engineering, electronics, communication engineering and systems, computer engineering (hardware only) and other allied subjects]

2.3. Other engineering sciences (such as chemical, aeronautical and space, mechanical, metallurgical and materials engineering, and their specialised subdivisions; forest products; applied sciences such as geodesy, industrial chemistry, etc.; the science and technology of food production; specialised technologies of interdisciplinary fields, e.g. systems analysis, metallurgy, mining, textile technology and other applied subjects)

3. MEDICAL SCIENCES

3.1 Basic medicine (anatomy, cytology, physiology, genetics, pharmacy, pharmacology, toxicology, immunology and immunohaematology, clinical chemistry, clinical microbiology, pathology)

3.2 Clinical medicine (anaesthesiology, paediatrics, obstetrics and gynaecology, internal medicine, surgery, dentistry, neurology, psychiatry, radiology, therapeutics, otorhinolaryngology, ophthalmology)

3.3 Health sciences (public health services, social medicine, hygiene, nursing, epidemiology)

4. AGRICULTURAL SCIENCES

4.1 Agriculture, forestry, fisheries and allied sciences (agronomy, animal husbandry, fisheries, forestry, horticulture, other allied subjects)

4.2 Veterinary medicine

5. SOCIAL SCIENCES

5.1 Psychology

5.2 Economics

5.3 Educational sciences (education and training and other allied subjects)

5.4 Other social sciences [anthropology (social and cultural) and ethnology, demography, geography (human, economic and social), town and country planning, management, law, linguistics, political sciences, sociology, organisation and methods, miscellaneous social sciences and interdisciplinary , methodological and historical S1T activities relating to subjects in this group. Physical anthropology, physical geography and psychophysiology should normally be classified with the natural sciences].

6. HUMANITIES

6.1 History (history, prehistory and history, together with auxiliary historical disciplines such as archaeology, numismatics, palaeography, genealogy, etc.)

6.2 Languages and literature (ancient and modern)

6.3 Other humanities [philosophy (including the history of science and technology) arts, history of art, art criticism, painting, sculpture, musicology, dramatic art excluding artistic "research" of any kind, religion, theology, other fields and subjects pertaining to the humanities, methodological, historical and other S1T activities relating to the subjects in this group] .

List of Acronyms and Abbreviations

A

API: Application Programming Interface.

Aurora Booster: Name given occasionally to the DEEP-ER Booster, as its design belongs to the Eurotech Aurora line of HPC clusters.

B

BADW-LRZ: Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften. Computing Centre, Garching, Germany

BeeGFS: The Fraunhofer Parallel Cluster File System (previously acronym FhGFS). A high-performance parallel file system to be adapted to the extended DEEP Architecture and optimised for the DEEP-ER Prototype.

BN: Booster Node (functional entity)

BNC: Booster Node Card is a physical instantiation of the BN

BoP: Board of Partners for the DEEP-ER project

BSC: Barcelona Supercomputing Centre, Spain

BSCW: Basic Support for Cooperative Work, Software package developed by the Fraunhofer Society used to create a collaborative workspace for collaboration over the web

C

CINECA: Consorzio Interuniversitario, Bologna, Italy

CN: Cluster Node (functional entity)

Coordinator: The contractual partner of the European Commission (EC) in the project

CP/RS: Checkpoint / Restart

CPU: Central Processing Unit

CRB: Customer Reference Board. An early version of a KNL board developed by Intel.

CRESTA: Collaborative Research into Exascale Systemware Tools & Applications: EU-funded Exascale project.

D

DDG: Design and Developer Group of the DEEP-ER project

DEEP: Dynamical Exascale Entry Platform

DEEP-ER: DEEP Extended Reach: this project

DEEP-ER Network: high performance network connecting the DEEP-ER BN, CN and NAM; to be selected off the shelf at the start of DEEP-ER

DEEP-ER Prototype: Demonstrator system for the extended DEEP Architecture, based on second generation Intel® Xeon Phi™ CPUs, connecting BN and CN via a single, uniform network and introducing NVM and NAM resources for parallel I/O and multi-level checkpointing

DEEP Architecture: Functional architecture of DEEP (e.g. concept of an integrated Cluster Booster Architecture), to be extended in the DEEP-ER project

DEEP System: The prototype machine based on the DEEP Architecture developed and installed by the DEEP project

E

E10: Exascale 10. Parallel I/O software developed by a consortium of partners around the EOFS community. Partner Seagate is responsible for the development needed for the DEEP-ER project.

EC: European Commission

EC-GA: EC-Grant Agreement

EEP: European Exascale Projects

EESI: European Exascale Software Initiative (FP7)

EOFS: European Open File System.

EU: European Union

Eurotech: Eurotech S.p.A., Amaro, Italy

Exaflop: 10^{18} Floating point operations per second

Exascale: Computer systems or Applications, which are able to run with a performance above 1018 Floating point operations per second

EXDCI: European Extreme Data & Computing Initiative (EXDCI), a H2020 coordination initiative.

EXTOLL: High speed interconnect technology for cluster computers developed by University of Heidelberg

ETP4HPC: European Technology Platform for High Performance Computing.

F

FhGFS: Acronym previously used to refer to BeeGFS.

FLOP: Floating point Operation

FP7: European Commission 7th Framework Programme.

FPGA: Field-Programmable Gate Array, Integrated circuit to be configured by the customer or designer after manufacturing

G

GRS: German Research School for Simulation Sciences GmbH, Aachen and Juelich, Germany

H

H5hut: Library implementing several data models for particle-based simulations that encapsulates the complexity of parallel HDF5.

HDF5: Hierarchical Data Format: A set of file formats and libraries designed to store and organise large amounts of numerical data

HMC: Hybrid Memory Cube

HPC: High Performance Computing

HW: Hardware

I

ICT: Information and Communication Technologies
IEEE: Institute of Electrical and Electronics Engineers
Intel: Intel Germany GmbH Feldkirchen,
IP: Intellectual Property
iPic3D: Programming code developed by the University of Leuven to simulate space weather
ISC: International Supercomputing Conference, Yearly conference on supercomputing which has been held in Europe since 1986

J

JUBE: Jülich Benchmarking Environment
JUDGE: Juelich Dedicated GPU Environment: A cluster at the Juelich Supercomputing Centre
JUELICH: Forschungszentrum Jülich GmbH, Jülich, Germany

K

KNC: Knights Corner, Code name of a processor based on the MIC architecture. Its commercial name is Intel® Xeon Phi™.
KNL: Knights Landing, second generation of Intel® Xeon Phi™
KULeuven: Katholieke Universiteit Leuven, Belgium

L

M

MIC: Intel Many Integrated Core architecture
Mont-Blanc: European scalable and power efficient HPC platform based on low-power embedded technology
Mont-Blanc 2: Follow-up project of Mont-Blanc
MPI: Message Passing Interface, API specification typically used in parallel programs that allows processes to communicate with one another by sending and receiving messages
MTBF: Mean Time Between Failures.

N

NAM: Network Attached Memory, nodes connected by the DEEP-ER network to the DEEP-ER BN and CN providing shared memory buffers/caches, one of the extensions to the DEEP Architecture proposed by DEEP-ER
NASA: National Aeronautics and Space Administration, Washington, USA
NEF: Network of European Foundations: name of server where financial data is uploaded to provide it to the EC.

NetCDF: Network Common Data Form. A set of software libraries and data formats that support the creation, access, and sharing of array-oriented scientific data
NVM: Non-Volatile Memory
NVMe: NVM Express. Specification for accessing solid-state drives attached through the PCIe bus.

O

OEM: Original Equipment Manufacturer. Term used for a company that commercialises products out of components delivered by other companies.
OmpSs: BSC's Superscalar (Ss) for OpenMP
OpenMP: Open Multi-Processing, Application programming interface that support multiplatform shared memory multiprocessing
OS: Operating System

P

ParaStation Consortium: Involved in research and development of solutions for high performance computing, especially for cluster computing
ParaStationMPI: Software for cluster management and control developed by ParTec
Paraver: Performance analysis tool developed by BSC
Paraview: Open Source multiple-platform application for interactive, scientific visualisation
ParTec: ParTec Cluster Competence Center GmbH, Munich, Germany
PCI: Peripheral Component Interconnect, Computer bus for attaching hardware devices in a computer
PCIe: PCI Express, Standard for peripheral interconnect developed to replace the old standards PCI, improving their performance
PFlop/s: Petaflop, 10^{15} Floating point operations per second
PM: Person Month or Project Manager of the DEEP project (depending on the context)
PMT: Project Management Team of the DEEP-ER project
PRACE: Partnership for Advanced Computing in Europe (EU project, European HPC infrastructure)
PROSPECT: Promotion of Supercomputing Partnerships for European Competitiveness and Technology (registered association, Germany)
PTC: Persistent Task-based Checkpoint

Q

QCD: Quantum Chromodynamics
QPACE: QCD Parallel Computing Engine. Specialised supercomputer for QCD Parallel Computing

R

R&D: Research and Development

S

SC: International Conference for High Performance Computing, Networking, Storage, and Analysis, organised in the USA by the Association for Computing Machinery (ACM) and the IEEE Computer Society

Scalasca: Performance analysis tool developed by JUELICH and GRS

SCR: Scalable Checkpoint/Restart library

SDV: Software Development Vehicle: a HW system to develop software in the time frame where the DEEP-ER Prototype is not yet available.

SEO: Search Engine Optimisation: the process of improving the visibility of a website or a web page in a search engine's results.

SSD: Solid State Disk

SW: Software

T

TFlop/s: Teraflop, 10^{12} Floating point operations per second

ToW: Team of Work Package leaders within the DEEP-ER project

TP10: Third Party under special clause 10.

U

UHEI: University of Heidelberg, Germany

UREG: University of Regensburg, Germany

V

VI-HPS: Virtual Institute for High Productivity Supercomputing

VTune: Commercial application for software performance analysis

W

WP: Work Package

X

Y

Z