



**SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures**

FP7-ICT-2011-7



DEEP

Dynamical Exascale Entry Platform

Grant Agreement Number: 287530

D1.2

Midterm management report at month 6

Approved

Version: 2.0
Author(s): E. Suarez, JUELICH
Contributions: F.Affinito (CINECA), D.Alvarez(JUELICH), A.Auweter(BADW-LRZ)
V.Beltran (BSC), F.Delalandre (EPFL), N.Eicker (JUELICH), J.Gimenez
(BSC), W.Guerich (JUELICH), D.Henaux (CGGVS), H.Ch.Hoppe (Intel),
M.E.Innocenti (KULeuven), J.Kreutz (JUELICH), H.Merx (CYI), L.Palm
(BADW-LRZ), G.Staffelbach (CERFACS)
Date: 12.06.2012

Project and Deliverable Information Sheet

DEEP Project	Project Ref. №: 287530	
	Project Title: Dynamical Exascale Entry Platform	
	Project Web Site: http://www.deep-project.eu	
	Deliverable ID: D1.2	
	Deliverable Nature: Report	
	Deliverable Level: CO * (the present document contains only the public part of the deliverable)	Contractual Date of Delivery: 31 / May / 2012
		Actual Date of Delivery: 29 / May / 2012
EC Project Officer: Leonardo Flores		

* - 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

Document	Title: Midterm management report at month 6	
	ID: D1.2	
	Version: 2.0	Status: Approved
	Available at: http://www.deep-project.eu	
	Software Tool: Microsoft Word	
	File(s): DEEP_D1.2_Midterm_Management_Report_at_M6_v2.0_ECapproved-PublishablePart	
Authorship	Written by:	E. Suarez, JUELICH
	Contributors:	F.Affinito (CINECA), D.Alvarez(JUELICH), A.Auweter(BADW-LRZ) V.Beltran (BSC), F.Delalondre (EPFL), N.Eicker (JUELICH), J.Gimenez (BSC), W.Guerich (JUELICH), D.Henaux (CGGVS), H.Ch.Hoppe (Intel), M.E.Innocenti (KULeuven), J.Kreutz (JUELICH), H.Merx (CYI), L.Palm (BADW-LRZ), G.Staffelbach (CERFACS)
	Reviewed by:	F.Geier (ParTec), W.Guerich (JUELICH), N.Eicker (JUELICH)
	Approved by:	BoP/PMT

Document Status Sheet

Version	Date	Status	Comments
1.0	29/May/2012	Final	Submitted to EC
2.0	12/June/2012	Approved	Approved by EC

Document Keywords

Keywords:	DEEP, HPC, Exascale
-----------	---------------------

Copyright notices

© 2011-2012 DEEP Consortium Partners. All rights reserved. This document is a project document of the DEEP project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the DEEP Partners, except as mandated by the European Commission contract 287530 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

Project and Deliverable Information Sheet	2
Document Control Sheet.....	2
Document Status Sheet	3
Document Keywords	4
Table of Contents	5
List of Figures	5
Executive Summary	6
1 Publishable summary	7
1.1 Project objectives.....	7
1.2 Work performed and main results.....	8
1.3 Expected final results	12
Annex A	13
A.1 Listing of dissemination activities	13
List of Acronyms and Abbreviations	16

List of Figures

Figure 1: Sketch of DEEP hardware architecture (CN: Cluster Node; BN: Booster Node; BI: Booster Interface)	11
Figure 2: Sketch of DEEP software architecture.....	11

Executive Summary

The Dynamical Exascale Entry Platform (DEEP) project started on 1st December 2011 and will last three years. The main goal of the project is to develop a prototype hardware and software supercomputing system with the potential to reach a peak performance of 100 PFlop/s in 2014, paving the way towards Exascale systems by the end of the decade. DEEP will optimise a set of grand-challenge applications with high societal impact and generic algorithmic structure for this platform. The key innovation of the DEEP project is its holistic Exascale-enabling concept integrating the architectural, system software and application level. The strategic goals of DEEP are *(i)* to contribute to an independent provision of general purpose Exascale performance supercomputers for the European HPC research infrastructure PRACE, *(ii)* to advance the growth of ICT and HPC hardware and software technology developed and produced in Europe, and *(iii)* to expand worldwide leadership and competitiveness of Europe's computational scientists and engineers.

This report describes the objectives, work performed, resources used, and achievements during the first six months of the project. The main achievements in the reporting period are enumerated below:

- Establish the project's organisation structure setting up all its bodies and selecting their members.
- Define the procedure to control the quality of reports and deliverables.
- Define the procedure for dissemination of the DEEP results.
- Organise a training workshop for application and middleware software developers.
- Define the concept details of the DEEP System hardware.
- Build hardware prototypes to study the most critical parts of the hardware architecture in detail.
- Define the environmental requirements for the installation of the DEEP System with special focus on the energy efficiency and the usage of hot water cooling.
- Sketch the software environment with all its layers taking into account the interactions between them.
- Analyse different programming models and the needs for their implementation.
- Analyse the structure and main characteristics of the pilot applications and measure their performance in a standard cluster for later reference.

1 Publishable summary

Numerous challenges must be overcome to reach Exascale by the end of the decade. When starting from current PFlop systems a factor of 1000 in performance increase is required. Investigating the roadmaps of standard CPUs reveals that it will not be sufficient to update and scale the concept of current cluster systems. In order to meet the requirements of energy-efficiency the use of accelerators becomes inevitable. However, today's solution of accelerated clusters –i.e. cluster with accelerators attached to each node– will not carry us to Exascale. On the one hand this is due to the inflexibility originating from the static assignment of standard CPUs and accelerators; on the other hand the competing use of the system bus by both accelerator and interconnect, combined with the lack of ability of the accelerator to use the interconnect autonomously, limits the scalability of this solution seriously.

Therefore, the DEEP Architecture proposes to detach the accelerators from the standard CPUs and to gather them into a separate cluster of accelerators that is called Booster. It is foreseen to run the highly scalable parts of the applications on this part of the DEEP System. The Booster is connected to a standard Cluster qualified to handle those parts of an application not suited for the Booster. The benefits are manifold: there is more flexibility on the ratio of standard CPUs and accelerators to be used by an application, the ability of the accelerators in the Booster to act autonomously allows for off-loading more complex and more parallel kernels, and an extended programming-model will support the application-developers in the identification of these offload-kernels and in porting their workload to the proposed architecture.

1.1 Project objectives

The specific objectives of the DEEP project are:

- Development of a prototype hardware platform consisting of a Cluster element based on multi-core-chips, a Booster element based on many-core technology and a commensurate connectivity, following the components-off-the-shelf philosophy. This prototype of Cluster Booster Architecture will serve as proof-of-concept for a next-generation PRACE production system reaching up to 100 PFlop/s in the time-frame 2014/2015, with the potential to achieve Exascale between 2018 and 2020.
- Combination of innovative technologies for the Booster element: novel Intel many-core processors, EXTOLL high speed interconnect, hot water cooling.
- Development of a reliable, open source cluster operating system, interconnect and runtime software stack with high resilience while exploiting millions of cores.
- Development of programming models, scientific libraries and performance tools for standard x86-based many-core processors, in order to achieve high productivity and enabling unprecedented scalability.
- Improvement of current cluster energy efficiency by an order of magnitude exploiting novel many-core chip technologies and advanced software-aided cooling technologies with a power usage effectiveness approaching a value of 1.
- Optimisation of a set of application codes on the DEEP System that are representative for future Exascale computing and data handling requirements, chosen from the fields

of Health and Biology, Climatology, Seismic Imaging, Computational Engineering, Space Weather, and Superconductivity and proving safe extrapolation to millions of cores as will be required with future Exascale systems.

- Demonstration of scalability of the new hardware-software concept with respect to the generic multi-scale, adaptive grid and long-range force parallelisation models underlying the application codes.
- Dissemination of the innovations and results of the project to the public.

With four PRACE supercomputing centres involved, DEEP will foster the alignment and synergies with the PRACE activities on the deployment of second-generation PRACE leadership-class supercomputers under the INFOSO Capacities Programme with the prospects to deploy >20 PFlop/s DEEP Supercomputers as future PRACE production systems. The involvement of the centres in the PRACE technical advisory group STRATOS will guarantee the dissemination of knowledge amongst the entire PRACE Tier-0/Tier-1 consortium. Through the participation of DEEP Partners in the PROSPECT association, the DEEP project will help pushing Europe's HPC industry and scientific community forward on its path towards the creation of a European HPC technology platform.

1.2 Work performed and main results

In the first six months the DEEP project has reached all the five milestones that were set for this reporting period, including the establishment of all the structures and procedures that will be needed during the further progress of the project.

Legal and administrative tasks

The DEEP consortium is constituted by 16 Partners (one of them is a Third Party Partner), including 4 PRACE supercomputing plus 3 further Research Centres, 5 Industry Partners, and 4 Universities. The management of this relatively large consortium requires legal and administrative regulations.

The Consortium Agreement, signed by all DEEP Partners, is the legal frame that describes the rights and responsibilities of all Partners in the consortium. With all these internal regulations in place, the Grant Agreement with the European Commission was signed on 18th November 2011. About three weeks later, on the 1st December 2011, the DEEP project officially started.

DEEP's effective kick-off was the face-to-face meeting of the whole consortium that took place in Juelich (Germany) on the 5th and 6th December 2011. This meeting gave the DEEP members the opportunity to meet each other in person and to identify the people with whom they will most closely work in the next three years. In addition the motivation and goals of the project were presented; the project management structure with all its bodies was established; the procedure for quality control of the deliverables was presented; and the tasks to be done in all Work Packages (WPs) during the first six months of the project were discussed, assigning responsibilities and contact people for each subject. Such face-to-face meetings are going to be organised on a regular basis twice a year. The second consortium meeting took place on the 8th – 9th May 2012 in Garching (Germany). Here the status of the project and the preparation of the first review meeting were discussed.

To guarantee the quality of Deliverables and Reports an internal review process has been set up. One or two people are selected from each Partner as internal reviewers. Before its submission to the European Commission, each Deliverable is reviewed by one internal reviewer plus one member of the PMT. Reasonable internal deadlines for this reviewing

process have been set, to be on time with the submission deadline given by the European Commission.

Dissemination, training and outreach

The implementation of the innovative DEEP concept constitutes a challenge that will require the development of new techniques and tools never tested before. Access to the know-how achieved in this process shall not remain limited to the group of people directly involved in the project, but must be made available for a wider community. For this reason, WP2 in DEEP is entirely devoted to the dissemination of the knowledge accumulated along the project's duration, and in training the users on its application.

The centre of the dissemination activities of DEEP is its web site: www.deep-project.eu. The DEEP web page will be updated regularly during the lifetime of the DEEP project and referred to in all other materials (articles, press releases, brochures, presentations, etc.). It is used to publish general information about the project, current activities, training opportunities, job vacancies, publications, tutorials, success stories, and achievements of the project.

Although the official starting point of the project was in December 2011, several dissemination activities took place even before that time. Partners from the DEEP consortium presented the project's concept in several conferences and workshops, including the two most important events in the HPC community: the International Supercomputing Conference (ISC) and the Supercomputing Conference (SC) that took place in Europe and the USA, respectively. During ISC'11 two members of the consortium gave presentations where DEEP was mentioned (at this time it was still only a proposal). On SC'11 the DEEP project co-organised, together with the two European Exascale projects CRESTA and MontBlanc, a BoF (Birds of a Feather) session where the three projects were presented to the interested public. A flyer explaining the goals and most important aspects of the DEEP project was distributed in the BoF and during the whole duration of the SC'11 conference from the booths of the DEEP Partners present in the event.

To foster cooperation activities with European industry and European R&D organisations, a liaison programme between the DEEP project and relevant industrial and business partners will be established. For this purpose DEEP collaborates and relies heavily on the existing PRACE advisory group for Strategic Technologies (STRATOS) as a vehicle to promote the use of multi-Petascale to Exascale systems to industrial and academic users.

Training the community on how to use the software and hardware developed in DEEP is also an important part of the project. The main goal of the training events in DEEP is to teach the application developers participating in the project on how to use the software tools and programming environment running on the DEEP System and other intermediate prototypes. The first DEEP Training Workshop took place from 30th January to 3rd February 2012 in Barcelona (Spain). An intensive course including lectures and numerous hands-on sessions on performance analysis tools (Scalasca (JSC) and Extrae/Paraver (BSC)), the OmpSs programming environment (BSC), and the Intel's MIC architecture was offered to all members of the DEEP and MontBlanc projects¹.

¹ The session on Intel's MIC, given the confidential character of the topic, was restricted to DEEP members only.

Technical Work

The technical work in DEEP is grouped into the three main parts: system hardware, system software, and applications.

In DEEP's hardware concept (see figure 1) a standard Cluster, composed of multi-core processors and the high flexibility network InfiniBand, is attached to a Booster of KNC many-core accelerators connected via a highly scalable EXTOLL torus network. New HPC middleware will be developed to efficiently communicate between Cluster and Booster and to allow for I/O directly from the latter. DEEP's software stack (see figure 2) focuses on scalability and an efficient use of resources. The parts of the application with a complex communication pattern (low to medium scalability) will run on the DEEP Cluster, while the parts with regular communication patterns (highly scalable) will run on the Booster. The DEEP's programming environment includes OmpSs (BSC) to help the developers in decomposing their applications into tasks in order to efficiently overlap the use of the Cluster and the Booster sides.

On the hardware side the work in the first six months of the project has been focused on defining the details of the DEEP hardware concept and the requirements for its environment infrastructure. Four hardware prototypes have been built to study in detail the most critical parts of the Booster:

- MIC evaluator: 3 standard server machines each containing a KNF card installed at JUELICH to start testing software tools and applications on the Intel's MIC architecture.
- EXTOLL evaluator: based on commodity hardware connecting multiple KNF through EXTOLL (used for software development)
- Interconnect evaluator part1: a board assembled at UniHD with one EXTOLL FPGA and one KNF to make tests of communication between EXTOLL and KNF.
- Interconnect evaluator part2 (BNC evaluator): board containing one KNF and one EXTOLL FPGA assembled at Eurotech to test signal integrity, mechanical and thermal issues.
- BIC evaluator: system installed at UniHD and constituted by three PCs, one with Infiniband, one with EXTOLL, and one with both used to develop Cluster-Booster protocol.

The lessons learned using these prototypes are a fundamental ingredient for the definition of the DEEP hardware design in all its details. The system hardware, including the Booster system architecture, the Booster Interface (BI) solution, the resiliency concept, and the cooling and power monitoring concept have been described in D3.1, the first Deliverable from WP3.

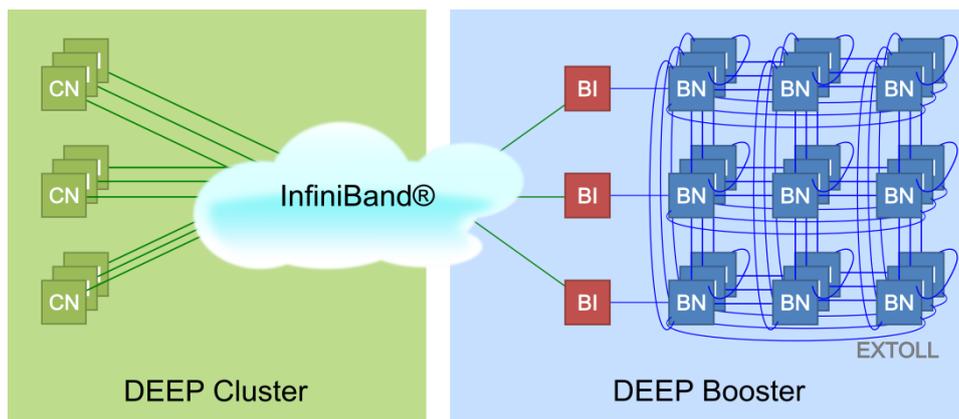


Figure 1: Sketch of DEEP hardware architecture (CN: Cluster Node; BN: Booster Node; BI: Booster Interface)

The environmental requirements of the DEEP System are strongly influenced by the importance of building it as energy efficient as possible. To reduce its energy consumption the cooling system of both Cluster and Booster is based on the “hot-water”-cooling concept, which allows “free cooling”. The distinctive features of the DEEP cooling system in comparison to a standard water- or air-cooled supercomputer imply special conditions for the installation of the DEEP System, conditions which are described in Deliverable D7.1.

On the software side, the most important activities of the first six months in the DEEP project were focused on the discussions needed to define the details of the software environment and the programming model. The overview of the DEEP software environment has been described in the DoW and is shown in the figure below. The different Work Packages involved in the development or directly affected by the characteristics of the system software (DDG, WP4, WP5 and WP8) have engaged in numerous discussions to define the details of all the software layers depicted in figure 2. The requirements from the application developers from WP8 are a fundamental input for the developers in WP4 and WP5, indicating them which features are needed in the final middleware.

The lowest layer of the software environment, the Cluster-Booster protocol, has been already defined and the implementation on the BIC evaluator has started. Its characteristics are detailed in Deliverable D4.1. In addition, first components of the ParaStationMPI and some parts of the OmpSs programming environment have been ported to the MIC architecture using the MIC evaluator.

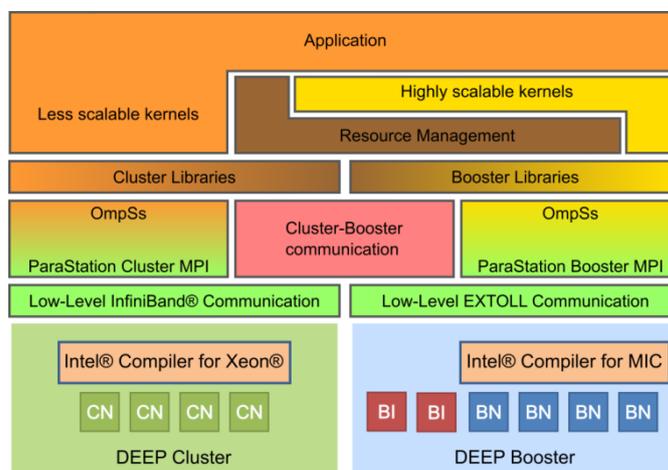


Figure 2: Sketch of DEEP software architecture

The application developers, with the help of the support team from Task 8.7, have worked on the characterisation and description of their applications to provide this information as an input for WP4, WP5, and the DDG (Design and Development Group). After the Training Workshop in Barcelona the developers were able to make also performance measurements on JUDGE, a cluster located at JUELICH, which is available as in-kind contribution. This exercise serves not only for the developers to gather experience in using the performance analysis tools, but also to deliver results that will be used as a reference for comparison with the measurements done later on using the DEEP System.

1.3 Expected final results

The DEEP project, at the end of its three years duration, will have installed the DEEP System in Juelich (Germany), composed of two parts, the Cluster and the Booster, running with a software stack that allows applications to distribute their code on both parts of the machine, dynamically assigning Cluster Nodes to Booster Nodes and vice versa. The experience gained running six scientific applications on the DEEP System, with different code structures, requirements, and scientific goals, will demonstrate if the proposed DEEP concept is suited for the next generation Exascale supercomputers.

Annex A

A.1 Listing of dissemination activities

Conferences, workshops, and meetings:

- **PARS Workshop**, Rüslikon, Switzerland, May 27, 2011:
 - N.Eicker and Th.Lippert (JUELICH): "The DEEP Project: an enhanced Cluster-Architecture for the Exascale" (presentation)
- **ISC 2011**, Hamburg, Germany, June, 2011
 - N.Eicker (JUELICH) (presentation) (22.06.2011)
 - Th.Lippert (JUELICH), "Trans-Petaflop/s initiatives in Europe" (presentation) (20.06.2011)
- **IESP 7th Workshop**, Cologne, Germany, October 6, 2011:
 - N.Eicker (JUELICH) (presentation) (22.06.2011)
- **EESI Final International Conference**, Barcelona, Spain, October 10-11, 2011:
 - W.Guerich (JUELICH): "Dynamical Exascale Entry Platform: DEEP" (presentation)
- **European Exascale Labs DEEP Dive**, Hillsboro, USA, November 10, 2011:
 - N.Eicker (JUELICH): (presentation)
- **Supercomputing Conference 2011**, Seattle, USA, November 17, 2011:
 - N.Eicker, W.Guerich (JUELICH): BoF European Exascale Projects
 - DEEP flyer distributed on the booths from DEEP Partners
- **Internal Exascale Cross-Lab Workshop**, Juelich, Germany, February 13-14, 2012
 - Th.Lippert (JUELICH): "The Exascale race: ready, steady, go!" (presentation)
 - N.Eicker (JUELICH): "Boosting Clusters: a view on the DEEP Architecture" (presentation)
 - E.Suarez (JUELICH): "ECL and DEEP applications" (presentation)
 - G.Lapenta (KULeuven): "Space weather" (presentation)
- **9th Intel EMEA HPC Roundtable**, Paris, France, March 27-28, 2012:
 - N.Eicker (JUELICH) "Boosting Clusters - Exploring heterogeneous architectures in DEEP" (presentation)
- **CEA-FZJ Internal Workshop**, Juelich, Germany, April 4-5, 2012:
 - E.Suarez (JUELICH) "Update on DEEP" (presentation)
- **Computational Methods in High Energy Density Plasmas, Workshop II: Computational Challenges in Magnetized Plasma**, Los Angeles, USA, April 16-20, 2012
 - G.Lapenta (KULeuven): "The challenge of multiphysics: federation or unification?" (presentation)
- **Meeting of ZKI "supercomputing" working group**, Kaiserslautern, Germany, April 19-29, 2012:

- J.Kreutz (JUELICH): “GPGPUs am Jülich Supercomputing Centre” (presentation, with some slides about DEEP)
- **“Tag der Informatik” at the University of Erlangen, Germany, April 20, 2012**
 - H.Ch.Hoppe (Intel): “The way to Exascale” (presentation)
- **ACM International Conference on Computing Frontiers 2012** Conference, Cagliari, Italy, May 15-17, 2012:
 - A.Auweter (BADW-LRZ): “DEEP: An Exascale prototype architecture based on a flexible configuration” (Invited Talk)

Publications, proceedings, press-releases, and newsletters:

- **PARS-Mitteilungen**, Mitteilungen - Gesellschaft für Informatik e.V., Parallel-Algorithmen und Rechnerstrukturen, ISSN 0177-0454, Nr. 28, October 2011 (Workshop 2011), 110 – 119
 - N.Eicker and Th.Lippert (JUELICH): “An accelerated Cluster-Architecture for the Exascale”
- **inSiDE; Innovative Supercomputing in Deutschland**, Vol.9, No.2, Autumn 2011.. Published twice a year by the The German National Supercomputing Centers HLRS, LRZ, JSC. Available online in: http://inside.hlrs.de/htm/Edition_02_11/article_12.html. Open Access: Yes
 - E.Suarez, N.Eicker, and W.Guerich (JUELICH): “Dynamical Exascale Entry Platform: the DEEP Project”.
- **Exascale - Newsletter des Forschungszentrums zum Supercomputing**, Nr. 03/2011, S. 3. Available online: http://www.fz-juelich.de/SharedDocs/Downloads/UK/DE/publikationen/themenhefte/exascale_nl_03.pdf?__blob=publicationFile. Open Access: Yes
 - Ch.Hohlfeld (JUELICH): “Tausendmal schneller mit neuem Booster”.
- **Gauß-Allianz Infobrief**, Nr. 5, February 2012. Available online: http://www.gauss-allianz.de/files/downloads/ga-infobrief_nr-5.pdf, Open Access: Yes
 - E.Suarez (JUELICH): “ Start des europäischen Exascale-Projekts DEEP”.
- **Press release from the Forschungszentrum Juelich**, 15.11.2011. Published online: <http://www.fz-juelich.de/SharedDocs/Pressemitteilungen/UK/EN/2011/11-11-15deep.html> Open Access: Yes
 - W.Guerich (JUELICH): “Booster for Next-Generation Supercomputers. Kick-off for the European exascale project DEEP”.
- **Press release from Eurotech**, 05.04.2012. Published online: <http://www.eurotech.com/en/press+room/news/?527&Eurotech%3a+%801.2M+order+for+Aurora+HPC+from+German+Research+Centre+J%26%23252%3blich>. Open Access: Yes
 - A.Barbaro (Eurotech): “Eurotech: €1.2M order for Aurora HPC from German Research Centre Jülich”.

- **HPC Wire**, 05.04.2012. Available online: http://www.hpcwire.com/hpcwire/2012-04-05/eurotech_to_supply_hpc_system_for_dynamical_exascale_entry_platform_project.html. Open Access: Yes
 - A.Barbaro (Eurotech): “Eurotech to Supply HPC System for Dynamical Exascale Entry Platform Project”

List of Acronyms and Abbreviations

A

- API:** Application Programming Interface
ASIC: Application Specific Integrated Circuit, Integrated circuit customised for a particular use
Aurora: The name of Eurotech's cluster systems

B

- BADW-LRZ:** Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften. Computing Centre, Garching, Germany
BAR: Base Address Register
BBP: Blue Brain Project
BI: Booster Interface (functional entity)
BIC: Booster Interface Card, Interface card to connect the Booster to the Cluster InfiniBand network
BLAS: Basic Linear Algebra Subprograms, Standard application programming interface to publish basic linear algebra libraries
BMC: Baseboard Management Controller
BN: Booster Node (functional entity)
BNC: Booster Node Card is a physical instantiation of the BN
BNC evaluator: Same as EXTOLL evaluator
BoF: Birds of a Feather Session: informal meeting during a Conference where people can discuss about the topic of their common interest
BoP: Board of Partners for the DEEP project
Booster System: Hardware subsystem of DEEP comprising of BNC, BIC and Intra-Booster network
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

- CANBus:** Controller Area Network Bus: vehicle bus standard designed to allow microcontrollers and devices to communicate with each other within a vehicle without a host computer.
CERFACS: Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, Toulouse, France
CGGVS: CGGVeritas Services SA, Paris, France
CINECA: Consorzio Interuniversitario, Bologna, Italy
CN: Cluster Node (functional entity)
Coordinator: The contractual Partner of the European Commission (EC) in the project
COI: Coprocessor Offload Infrastructure
CPU: Central Processing Unit
CRESTA: Collaborative Research into Exascale Systemware, Tools & Applications: EU-FP7 Exascale Project led by the University of Edinburgh.
CYI: Cyprus Institute, Nicosia, Cyprus

D

- DC:** Direct Current (electricity)
- DDG:** Design and Developer Group of the DEEP project
- DEEP:** Dynamical Exascale Entry Platform: EU-FP7 Exascale Project led by Forschungszentrum Juelich.
- DEEP Architecture:** Functional architecture of DEEP (e.g. concept of an integrated Cluster Booster Architecture)
- DEEP Booster:** Booster part of the DEEP System
- DEEP Supercomputer:** A future Exascale supercomputer based on the DEEP Architecture
- DEEP System:** The production machine based on the DEEP Architecture developed and installed by the DEEP project
- DFP:** Dense Form Factor
- DoW:** Description of Work: Annex I of the Grant Agreement

E

- EC:** European Commission
- ECHAM:** ECHAM/MESSy Atmospheric Chemistry code
- ECL:** ExaCluster Laboratory, A collaboration of Intel, ParTec and JUELICH to develop cluster management software for Exascale computing
- EESI:** European Exascale Software Initiative (FP7)
- EMAC:** ECHAM/MESSy atmospheric chemistry simulation code, developed by the Cyprus Institute
- EMEA:** Europe, the Middle East and Africa, Regional designation used for government, marketing and business purposes
- EPFL:** École Polytechnique Fédérale de Lausanne, Switzerland
- 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 between 10^{15} and 10^{18} Floating point operations per second
- EXTOLL:** High speed interconnect technology for cluster computers developed by University of Heidelberg
- EXTOLL evaluator:** Platform for evaluation of EXTOLL technology, Developed and used in the DEEP project
- Extrae/Paraver:** Performance analysis tool developed by BSC

F

- FFTW:** Fastest Fourier Transform in the West: a particular implementation of the FFT algorithm
- FLOP:** Floating point Operation
- FPGA:** Field-Programmable Gate Array, Integrated circuit to be configured by the customer or designer after manufacturing

G

GIT: Distributed revision control and source code management system
GPU: Graphics Processing Unit
GRS: German Research School for Simulation Sciences GmbH, Aachen and Juelich, Germany

H

HCA: Host Channel Adapter
HDF: Hierarchical Data Format: a set of file formats and libraries designed to store and organise large amounts of numerical data
HPC: High Performance Computing
HW: Hardware

I

IA: Intel Architecture
IB: InfiniBand
ICT: Information and Communication Technologies
IEEE: Institute of Electrical and Electronics Engineers
IMM: Implicit Moment Method
INFSO: Information Society
Intel: Intel GmbH Braunschweig, Germany
Interconnect evaluator: Hardware for interconnect studies on physical and mechanical layer, Developed and used in the DEEP project
I/O: Input/Output
IP: Intellectual Property or Internet Protocol (depending on the context)
iPIC3D: Programming code developed by the University of Leuven to simulate space weather
IPC: Instructions Per Cycle
ISC: International Supercomputing Conference, Yearly conference on supercomputing which has been held in Europe since 1986

J

JSC: Juelich Supercomputing Center
JUDGE: Juelich Dedicated GPU Environment: a cluster at the Juelich Supercomputing Centre
JUELICH: Forschungszentrum Jülich GmbH, Jülich, Germany

K

KB: Kilo Byte
KNC: Knights Corner, Code name of a processor based on the MIC architecture
KNF: Knights Ferry, Intel first prototype processor based on the MIC
KULeuven: Katholieke Universiteit Leuven, Belgium

L

LAPACK: Linear Algebra PACKage

M

- MB:** Mega Byte
- MC:** Monte Carlo
- MESSy:** Modular Earth Submodel System simulation code
- MLNX:** Mellanox Technologies, Ltd., Sunnyvale, California and Yokneam, Israel
- MIC:** Intel Many Integrated Core architecture
- MIC evaluator:** Platform for evaluation of the MIC architectural concept, Used only in the DEEP project
- MIC-OS:** Operating System of the MIC architecture
- Mini Booster prototype:** Minimal instantiation of a DEEP Booster used for analysis of the energy-aware functionality, Developed and used in the DEEP project
- Mini DEEP System:** A fully featured DEEP System of minimal size comprising the Mini Booster
- MontBlanc:** European scalable and power efficient HPC platform based on low-power embedded technology: EU-FP7 Exascale Project led by the Barcelona Supercomputing Centre.
- MPI:** Message Passing Interface, API specification typically used in parallel programs that allows processes to communicate with one another by sending and receiving messages
- MPICH:** Freely available, portable implementation of MPI

N

- NIC:** Network Interface Card, Hardware component that connects a computer to a computer network

O

- OmpSs:** BSC's Superscalar (Ss) for OpenMP
- OpenMP:** Open Multi-Processing, Application programming interface that support multiplatform shared memory multiprocessing
- OS:** Operating System

P

- ParaStationMPI:** Software for cluster management and control developed by ParTec
- ParTec:** ParTec Cluster Competence Center GmbH, Munich, Germany
- PC:** normally Personal Computer, but in the context of the proposal also Project Coordinator
- PCB:** Printed Circuit Board, board used in electronic to mechanically support and electrically connect electronic components
- 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
- PIC:** Particle In Cell:

- PM:** Project Manager of the DEEP project or Person Month (depending on the context)
- PMT:** Project Management Team of the DEEP project
- PR:** Normally Public Relations, but in the context of this report Particle Related
- PRACE:** Partnership for Advanced Computing in Europe (EU project, European HPC infrastructure)
- Project Coordinator:** Leading scientist coordinating and representing the DEEP project
- PROSPECT:** Promotion of Supercomputing Partnerships for European Competitiveness and Technology (registered association, Germany)
- PUE:** Power Usage Effectiveness

Q

- QDR:** Quad Data Rate, Communication signalling technique of InfiniBand
- QPACE:** specialised supercomputer for QCD Parallel Computing on CELL processors

R

- RAS:** Reliability Availability and Serviceability
- RDMA:** Remote Direct Memory Access
- R&D:** Research and Development
- RTD:** Research and Technological Development

S

- SATA:** Serial Advanced Technology Attachment: computer bus interface for connecting host bus adapters to mass storage devices such as hard disk drives and optical drives.
- 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
- SISSA:** International School of Advanced Studies, Trieste, Italy
- SMFU:** Shared Memory Functional Unit
- SS:** Superscalar, programming environment developed by BSC
- SSD:** Solid State Disk.
- STRATOS:** PRACE advisory group to foster development of HPC technologies in Europe
- SVN:** Apache Subversion: a software versioning and revision control system
- SW:** Software

T

- TCO:** Total Cost of Ownership
- TFlop/s:** Teraflop, 10^{12} Floating point operations per second
- Tier-0, Tier-1, ...:** Different classes of supercomputers ordered by their performance
- ToW:** Team of Work Package leaders within the DEEP project
- TurboRVB:** Quantum Monte Carlo Software for electronic structure calculations developed by SISSA

U

UniHD: University of Heidelberg, Germany

UniReg: University of Regensburg, Germany

V

VELO: Virtualised Engine for Low Overhead: an EXTOLL communications channel.

W

WP: Work Package

WG: Work Group

X

x86: Family of instruction set architectures based on the Intel 8086 CPU

Y***Z***