Projects

Going DEEP-ER to Exascale

October 1, 2013 marked the start of the EU-funded project "DEEP-ER" (DEEP–Extended Reach). As indicated by its name, DEEP-ER aims for exztending the Cluster-Booster Architecture proposed by the currently running EU-project "DEEP" [1] with additional functionality. The DEEP-ER project will have a duration of three years and a total budget of more than 10 million Euro. Its coordinator, the Jülich Supercomputing Centre, leads a consortium of 14 partners from 7 different European countries.

The heterogeneous Cluster-Booster Architecture implemented by DEEP consists of two parts: a Cluster based on multi-core CPUs with InfiniBand interconnect, and a Booster of many-core processors connected by the EXTOLL network. DEEP-ER (see Fig. 1) will simplify this concept by unifying the interconnect merging Cluster and Booster Nodes into the same network. The DEEP-ER prototype will explore new memory technologies and concepts like non-volatile memory (NVM) and network attached memory (NAM). Additionally, with respect to DEEP the processor technology will be updated: next generation Xeon processors will be used in the Cluster and the KNL generation of Xeon Phi will populate the Booster.

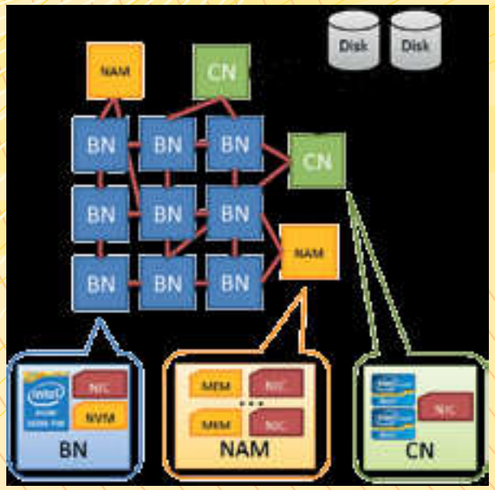

Figure 1: Cluster-Booster Architecture as implemented in DEEP-ER. (CN: Cluster Node; BN: Booster Node; NAM: Network Attached Memory; NVM: Non-Volatile Memory; MEM: Memory). In contrast to DEEP, here Cluster and Booster Nodes are attached to the same interconnect. Additional memory and storage is added at the node, the network, and system levels.

The DEEP-ER multi-level I/O infrastructure has been designed to support data-intensive applications and multi-level checkpointing/restart techniques. The project will develop an I/O software platform based on the Fraunhofer parallel file system (BeeGFS), the parallel I/O library SIONlib, and the I/O software package Exascale10. It aims to enable an efficient and transparent use of the underlying hardware and to provide all functionality required by applications for standard I/O and checkpointing.

DEEP-ER proposes an efficient and user-friendly resiliency concept combining user-level checkpoints with transparent task-based application restart. OmpSs is used to identify the application's individual tasks and their interdependencies. The OmpSs runtime will be extended in order to automatically restart tasks in the case of transient hardware failures. In combination with a multi-level user-based checkpoint infrastructure to recover from non-transient hardware-errors, applications will be able to cope with the higher failure rates expected in Exascale systems.

DEEP-ER's I/O and resiliency concepts will be evaluated using seven HPC applications from fields that have proven the need for Exascale resources.

Figure 2: Participants of the DEEP-ER kickoff meeting.

Acknowledgements

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 610476 (Project "DEEP-ER").

References

[1] Suarez, E., Eicker, N., Gürich, W.: "Dynamical Exascale Entry Platform: the DEEP Project", inSiDE Vol. 9 No.2, Autumn 2011, http://inside.hlrs.de/htm/Edition_02_11/article_12.html

contact: Estela Suarez, e.suarez@fz-juelich.de

- Estela Suarez
- Norbert Eicker

Jülich Supercomputing Centre (JSC)

Spring 2014 • Vol. 12 No. 1 • inSiDE