DEEP-ER IN A NUTSHELL

OBJECTIVES:

- Extend the DEEP concept addressing two important challenges for future Exascale systems:
  - I/O
  - Resiliency
- Enhance the existing DEEP hardware architecture by:
  - Next generation processors
  - Unified interconnect
  - Exploration of new memory technologies

BUDGET:

€ 10 Mio. (EU-funding: € 6.4 Mio.)

TIME PERIOD:

Oct 2013 – Sept 2016

CONSORTIUM:

- Co-ordinator: Jülich Supercomputing Centre
- 14 Partners
- 7 European countries

MEET DEEP-ER AT SC14

• European Exascale Projects	Booth #1039
• Emerging Tech Exhibits	DEEP-ER I/O

MEET DEEP-ER PARTNERS AT SC14

• BSC	Booth # 3427
• CINECA	Booth # 733
• Eurotech	Booth # 3339
• Fraunhofer ITWM	Booth # 3147
• Intel	Booth # 1215 / # 1315
• Inria	Booth # 3533
• JSC	Booth # 639
• LRZ	Booth # 806
• Seagate	Booth # 3239

DEEP-ER

GET IN TOUCH WITH THE DEEP-ER PROJECT

Email: pmt@deep-er.eu
Web: www.deep-er.eu
Twitter: @DEEPprojects

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement n° 610476.

DEEP-ER

Resiliency

Highly scalable I/O

True Co-Design

EUROPE TOWARDS EXASCALE

CLUSTER-BOOSTER ARCHITECTURE

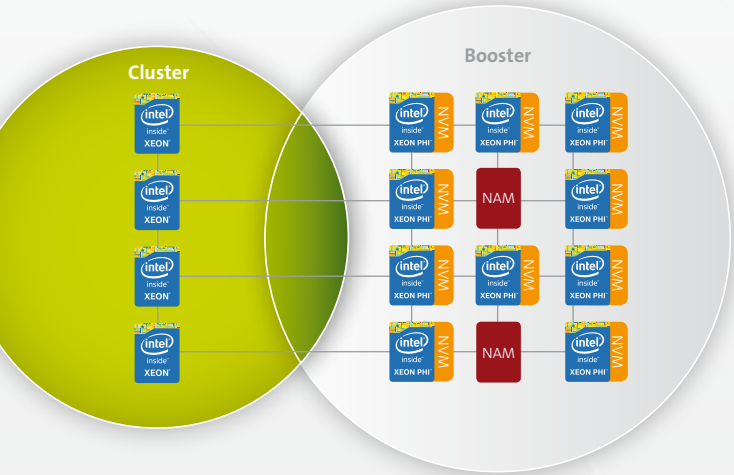

DEEP-ER extends and improves the **Cluster-Booster architecture** of the predecessor project DEEP. Building on the core DEEP design principle of a **standard HPC Cluster** and a tightly connected **Booster**, DEEP-ER **simplifies the architecture** and introduces support for highly **scalable I/O** and **checkpoint/restart**:

- DEEP-ER uses future Intel® Xeon Phi™ CPUs (code named Knights Landing) for the Booster and the **same interconnect** network to span both Cluster and Booster.
- **Non-volatile memory** (NVM) devices attached to the Booster nodes **provide fast, persistent local storage**.
- Novel **Network Attached Memory** (NAM) devices connected to the DEEP-ER interconnect offer shared, **persistent memory resources** to both Cluster and Booster nodes.
- PCI Express generation 3 serves as the **integration layer for the Booster nodes**, enabling the project to easily upgrade to new NVM or interconnect technology.

MASTERING I/O AND RESILIENCY

The DEEP-ER project will feature a **highly scalable, efficient, and user-friendly parallel I/O system** based on the Fraunhofer parallel file system BeeGFS. Extensions to the Posix I/O API will enable applications to **efficiently use the different levels of the memory/storage** subsystem. These extensions will originate from **BeeGFS** itself, the parallel I/O library **SIONlib**, and **Exascale10** – a novel I/O concept developed by the Exascale10 Workgroup.

Exploiting the multiple levels of non-volatile memory and storage added to the DEEP architecture, DEEP-ER will provide a **low-overhead, unified user-level checkpointing system**. Building on the OmpSs programming model and facilitated by the ParaStation Global MPI, a **dual-approach resiliency concept** will combine a coarse-grained application-based multi-level **checkpoint/restart** mechanism with a less intrusive and more fine-grained scheme for **task-controlled recovery from component failures.**

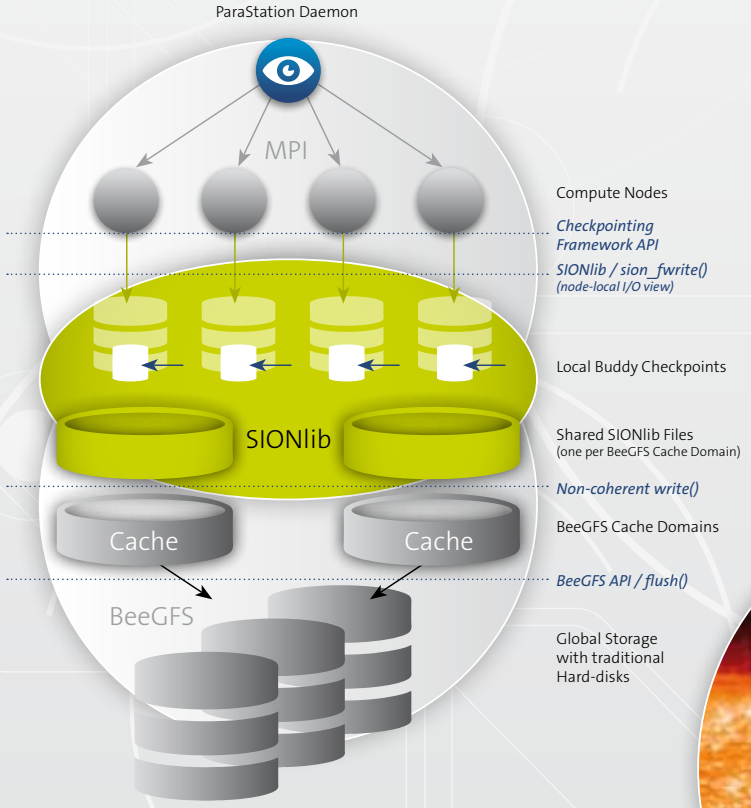

DRIVER TO EXASCALE: APPLICATIONS

In DEEP-ER seven scientific codes have been carefully chosen to explore the path to Exascale proposed by the Cluster-Booster architecture. Reaching a **new level of scalability** requires a tight **collaboration**. Thus, the application **requirements** and structure have been **feed into the design process** of the check-pointing infrastructure, parallel file system, and the scalable I/O libraries developed as part of DEEP-ER.