

# The DEEP (and DEEP-ER) projects

Estela Suarez - Jülich Supercomputing Centre

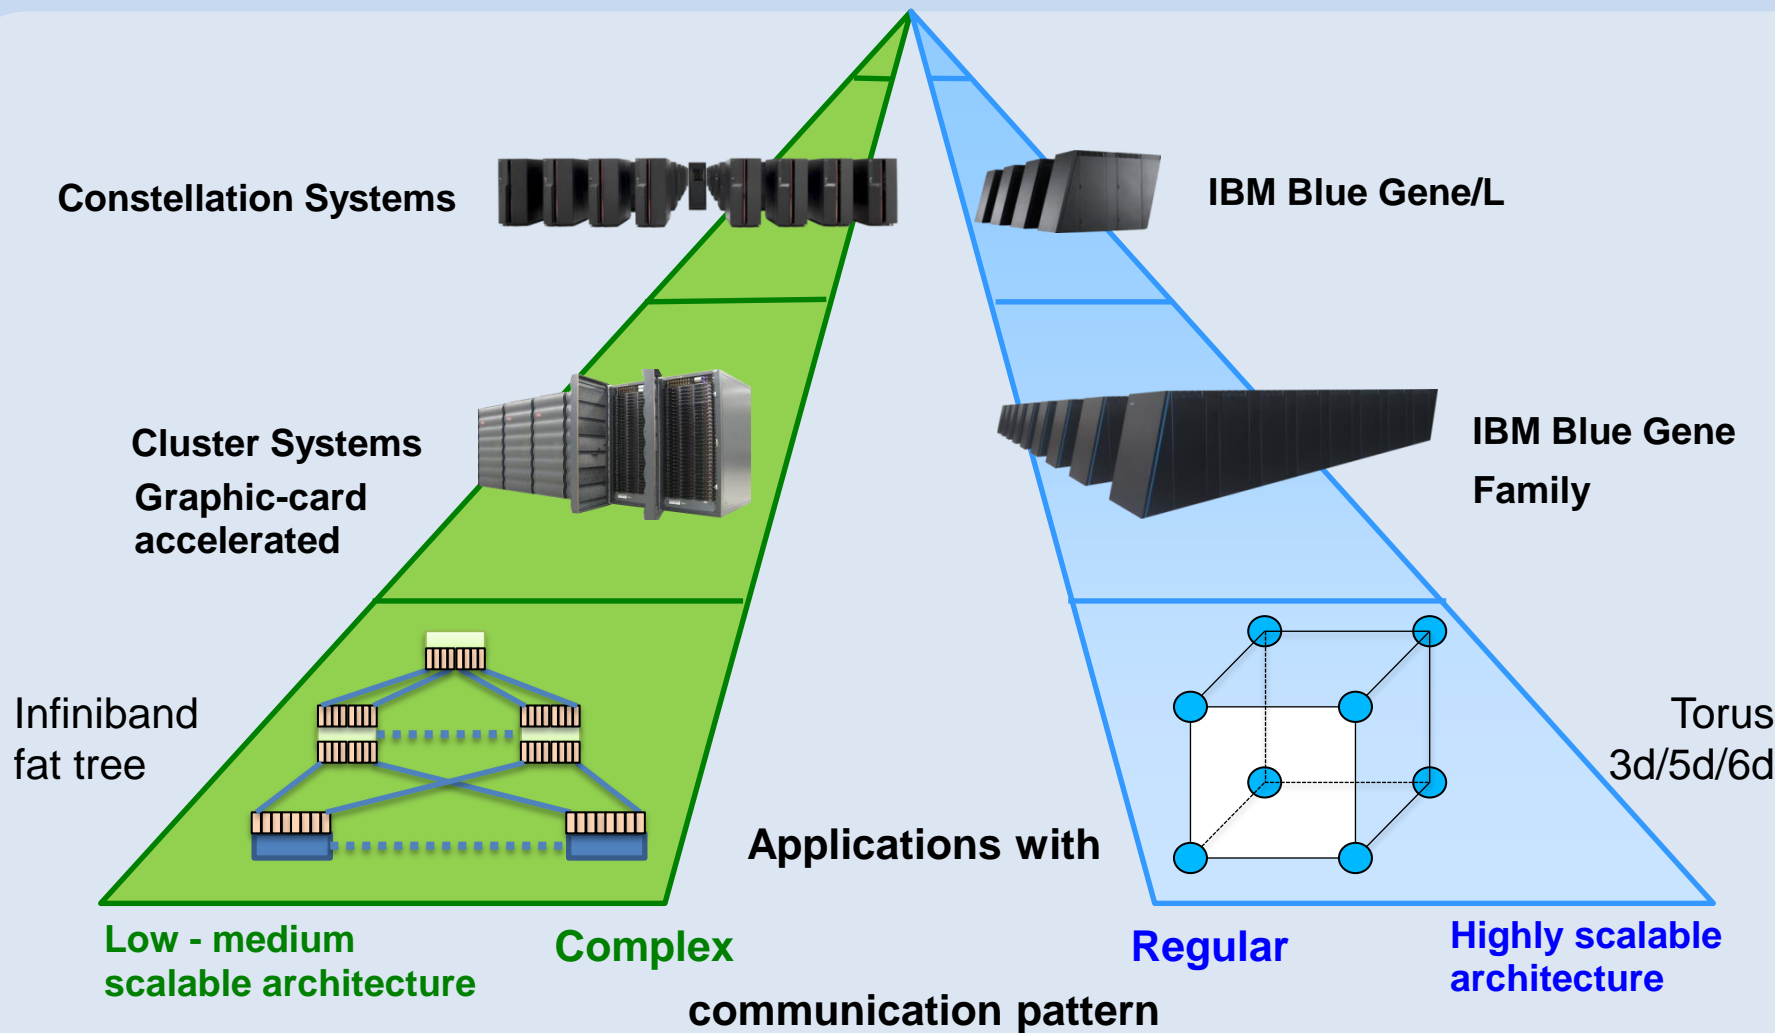
BDEC for Europe Workshop  
Barcelona, 28.01.2015

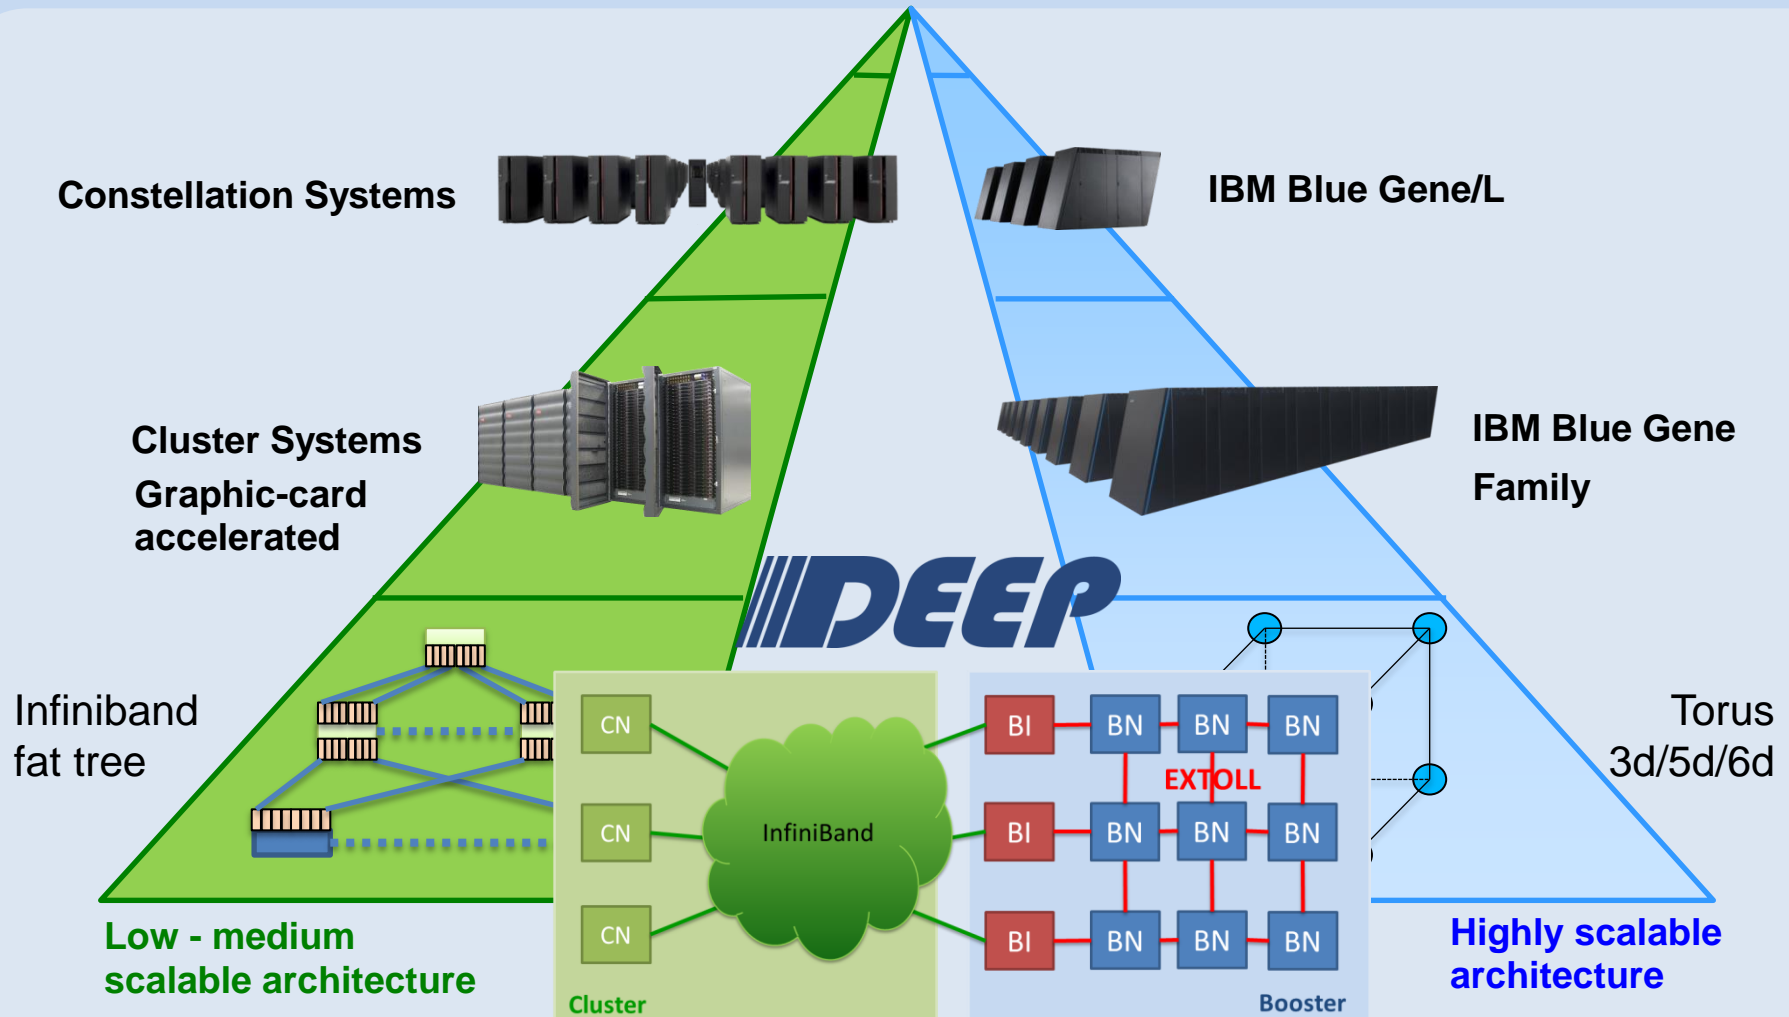
## DEEP

- Cluster-Booster **concept**
- Software stack
- Programming environ.
- Performance tools
- Energy efficiency
- Applications:
  - Co-design
  - Evaluation/demonstration

## DEEP-ER

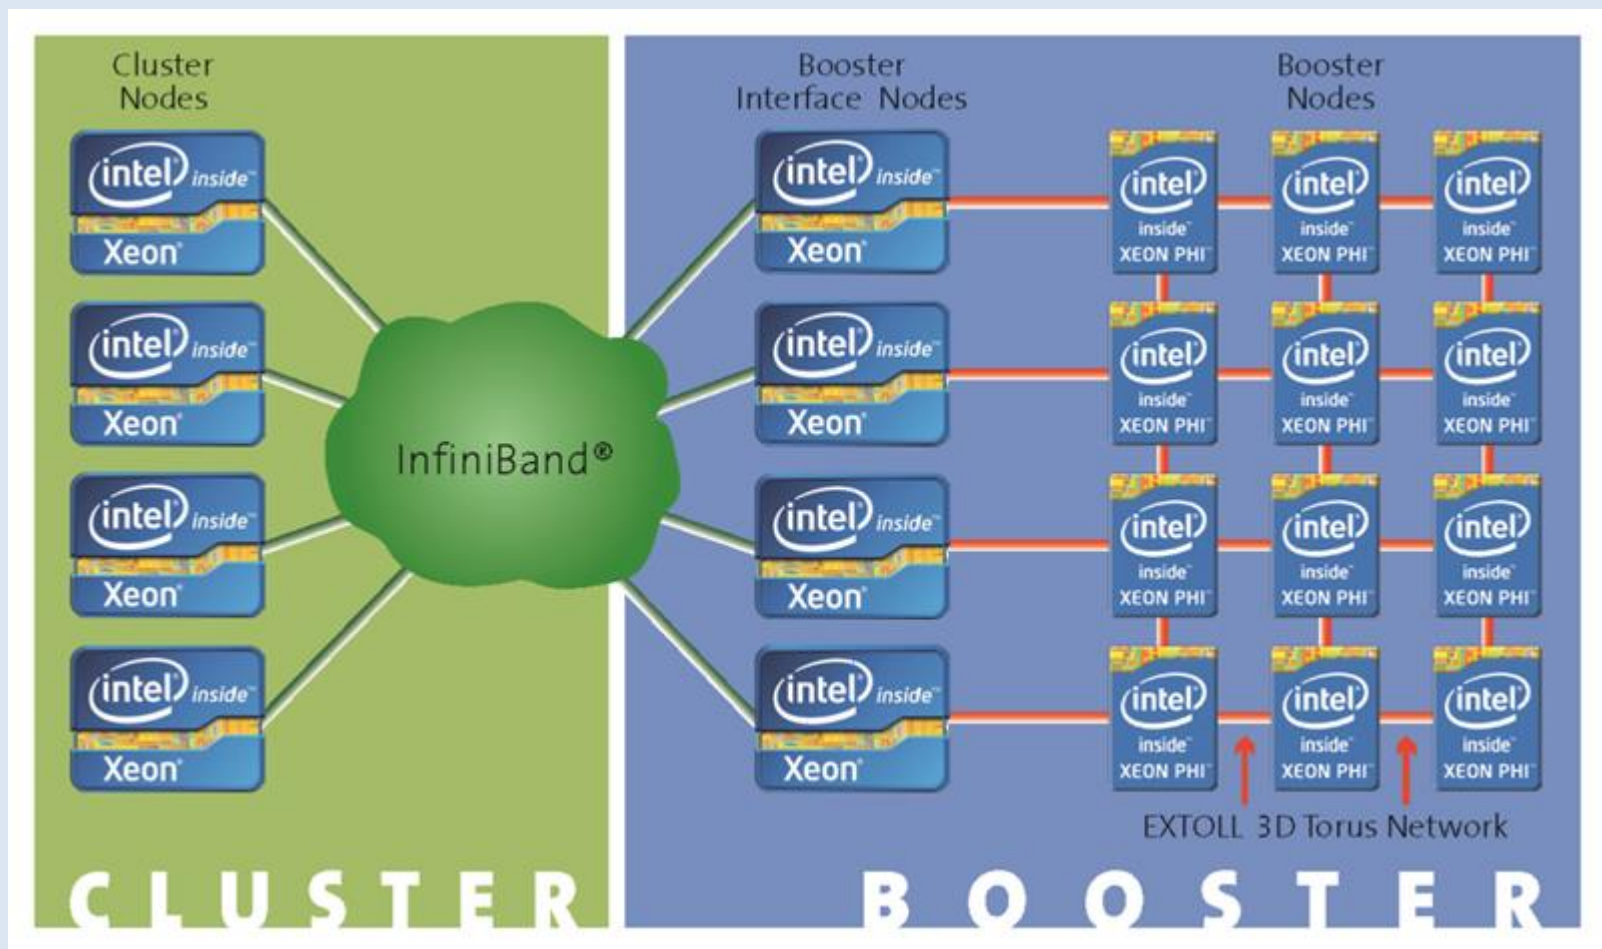
- Improve **I/O**
- Improve **resiliency**
- New memory technology
- Applications:
  - Co-design
  - Evaluation/demonstration





128 Xeon (Sandy Bridge)

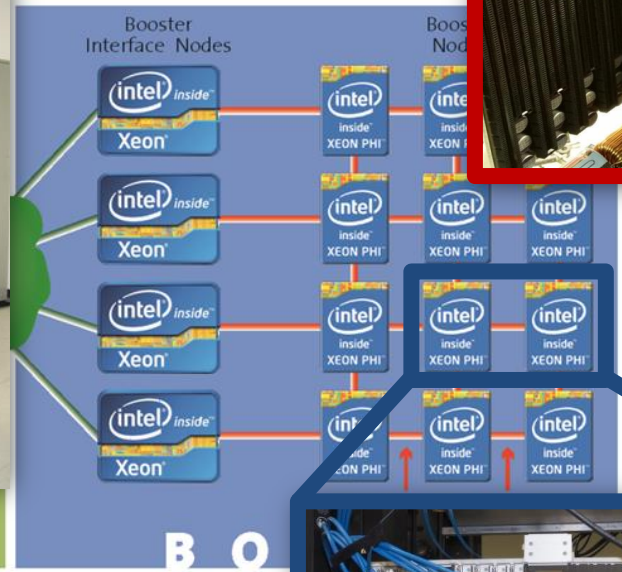
384 Xeon Phi (KNC)




**DEEP  
Cluster  
at JSC**



**CLUSTER**



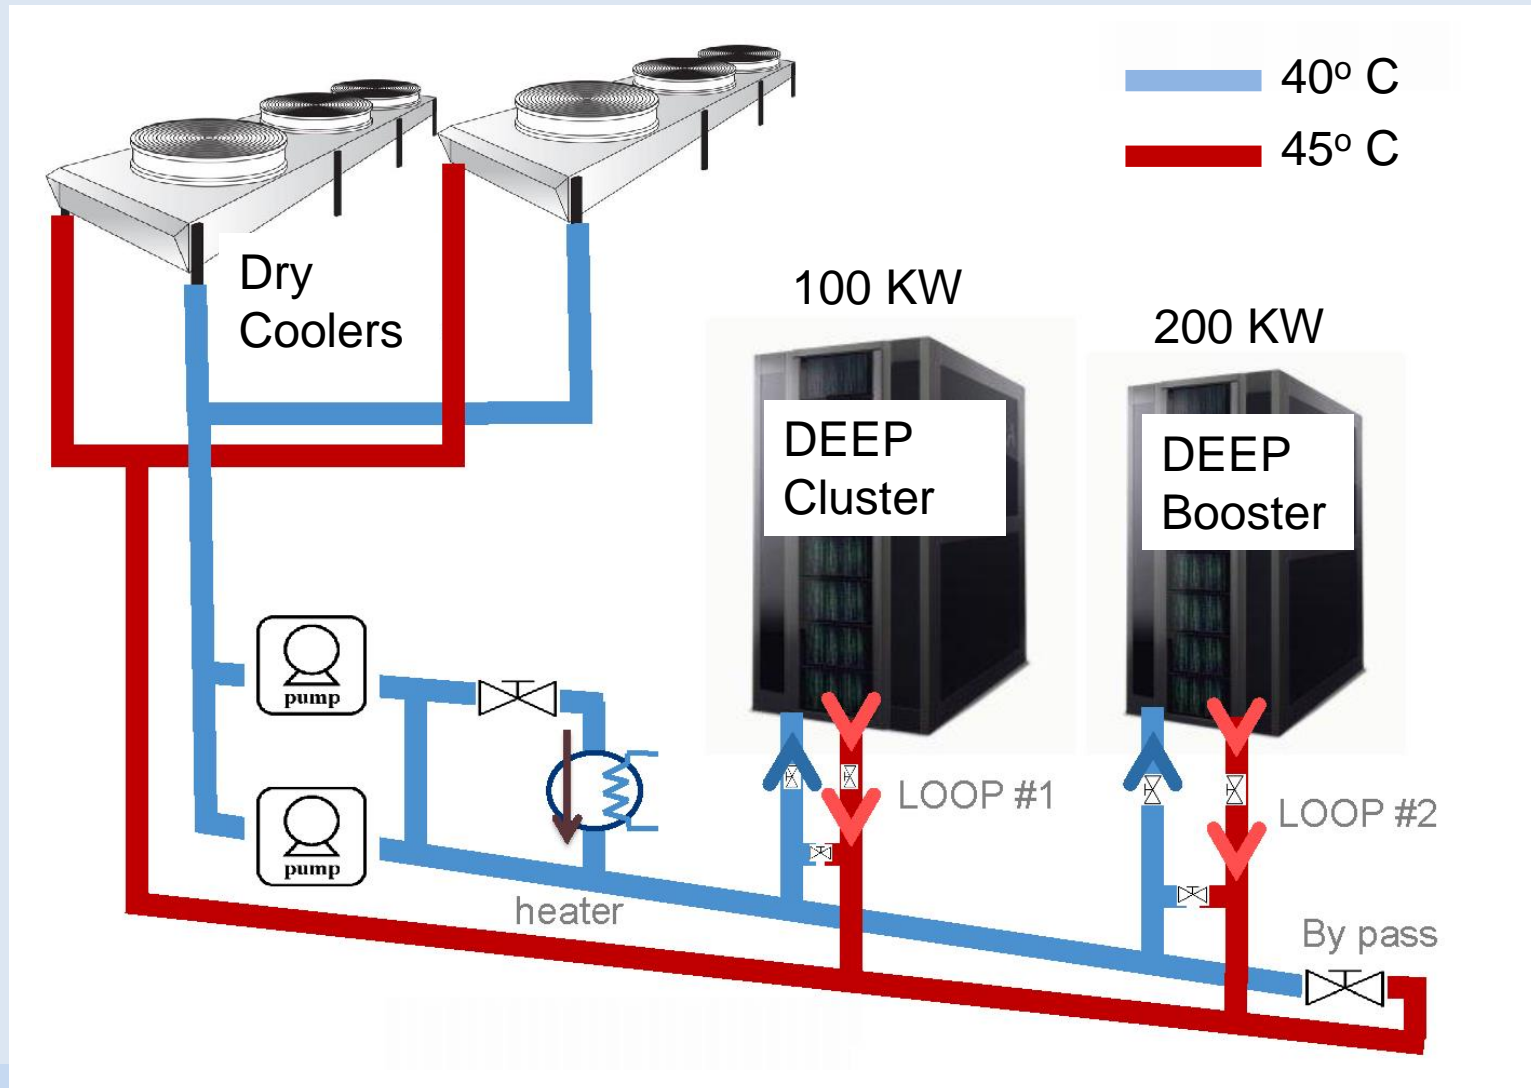
**Booster  
Chassis  
(32 KNCs)**

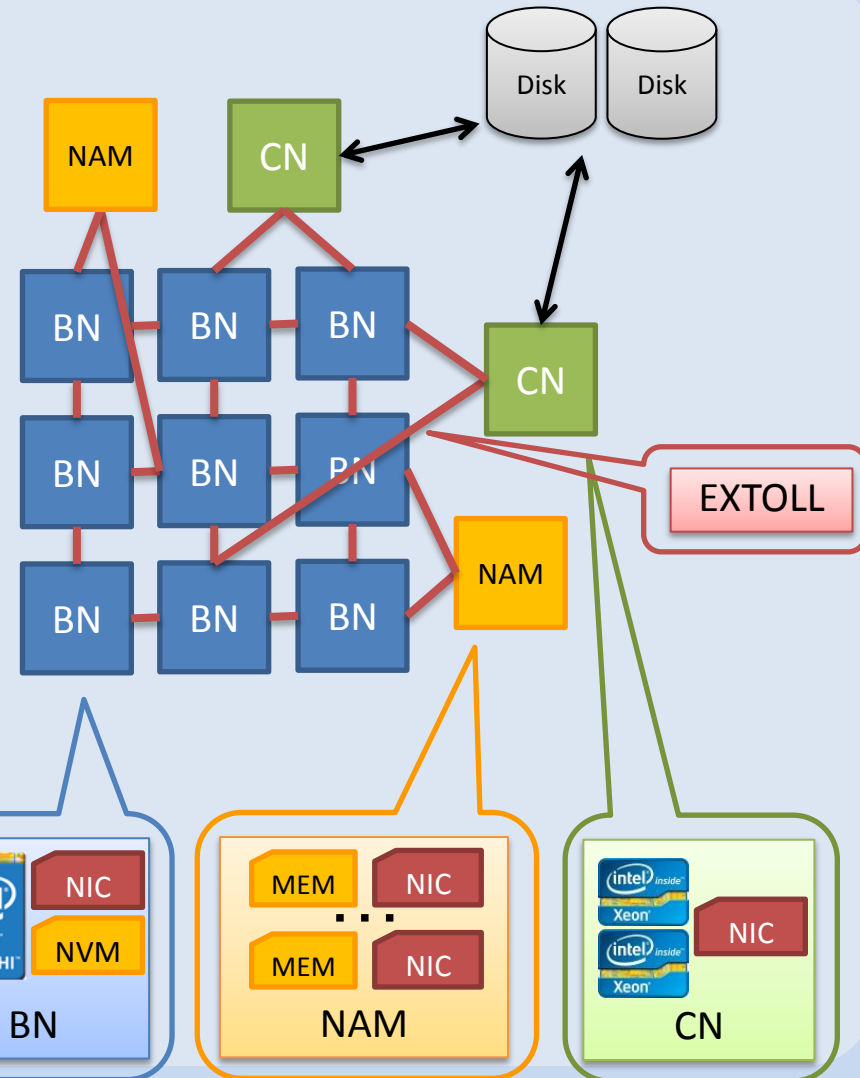



**ASIC  
Evaluator  
(32 KNCs)**









## Legend:

CN: Cluster Node

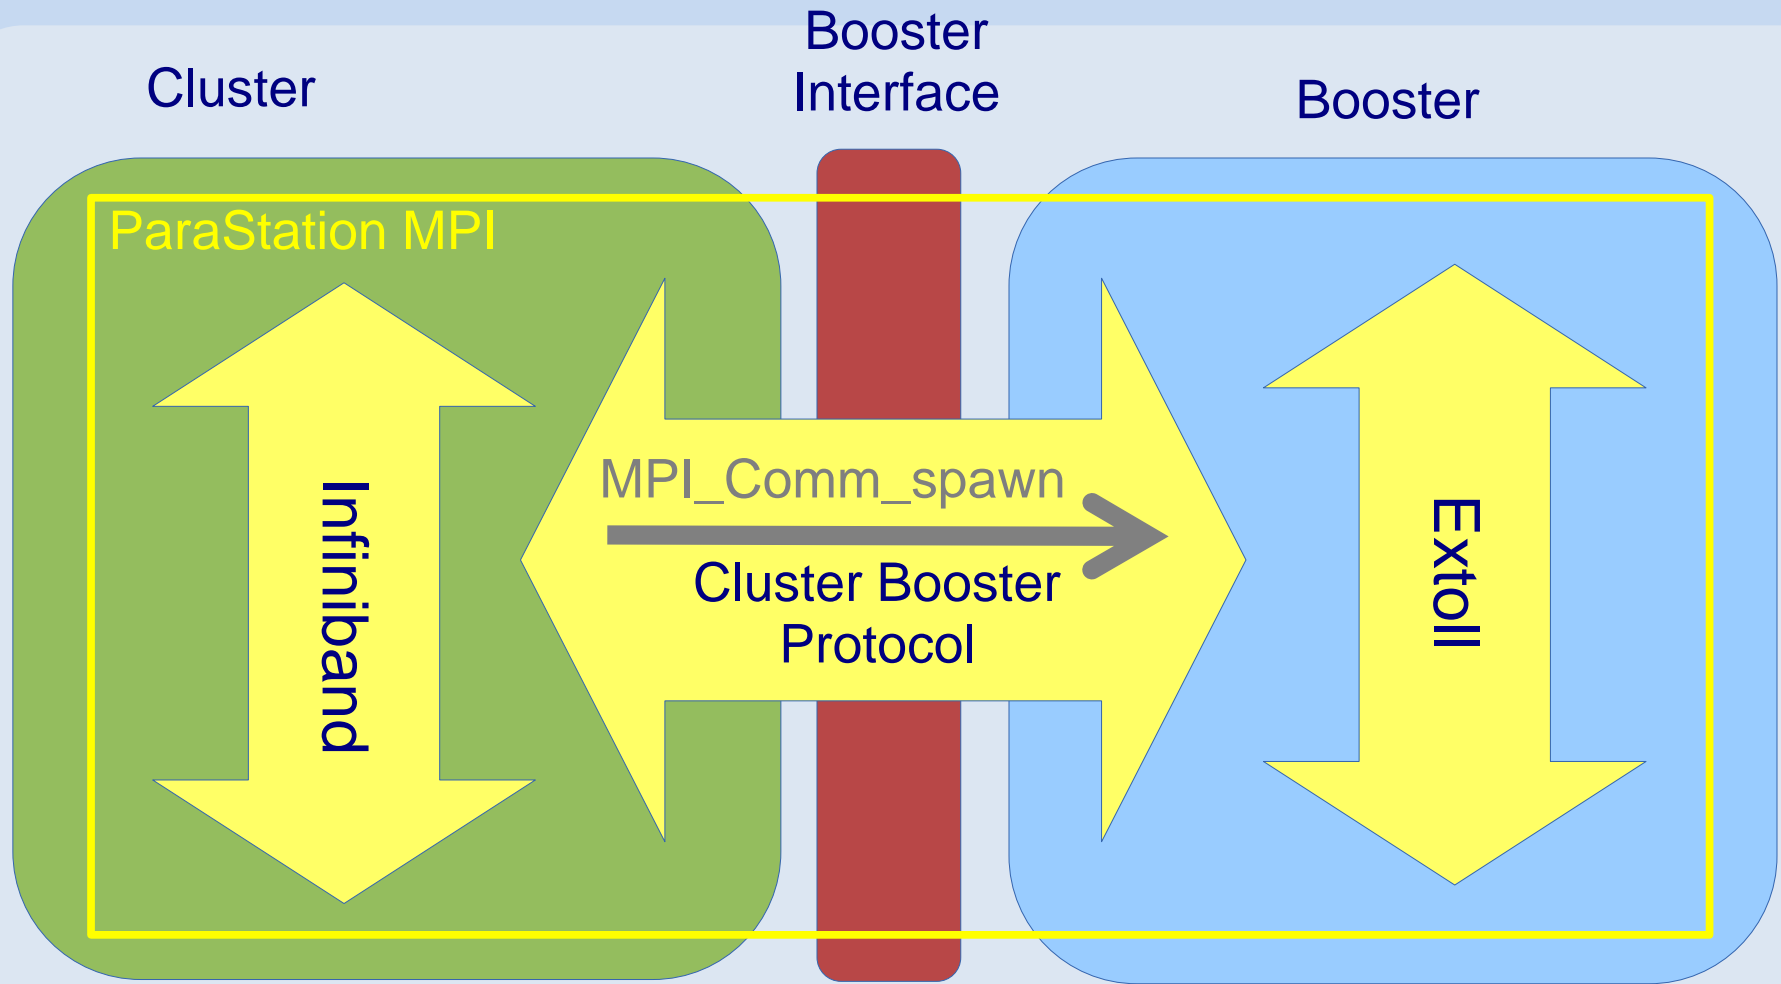
BN: Booster Node

BI: Booster Interface

NAM: Network Attached Memory

NVM: Non Volatile Memory





OmpSs on top of MPI provides pragmas to ease the offload process

Source code

```
int main(int argc, char *argv[]){
    /*...*/
    for(int i=0; i<3; i++){
        #pragma target device (comm:size*rank+i) copy_deps
        #pragma omp task input(...) output(...)
        foo_mpi(i, ...);}}

```

Compiler

OmpSs Compiler

Application binaries

Cluster Executable

Booster Executable

DEEP Runtime

Cluster MPI

ParaStation Global MPI

DEEP Runtime

Booster MPI

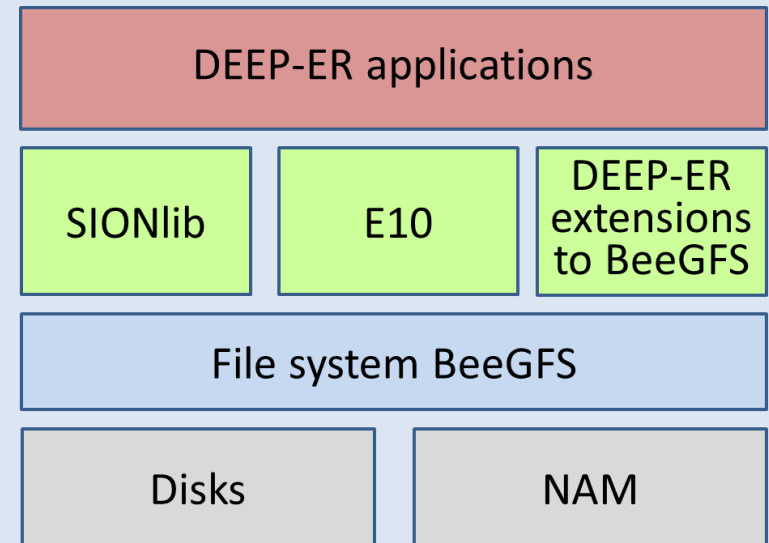
OmpSs Runtime

CLUSTER

BOOSTER

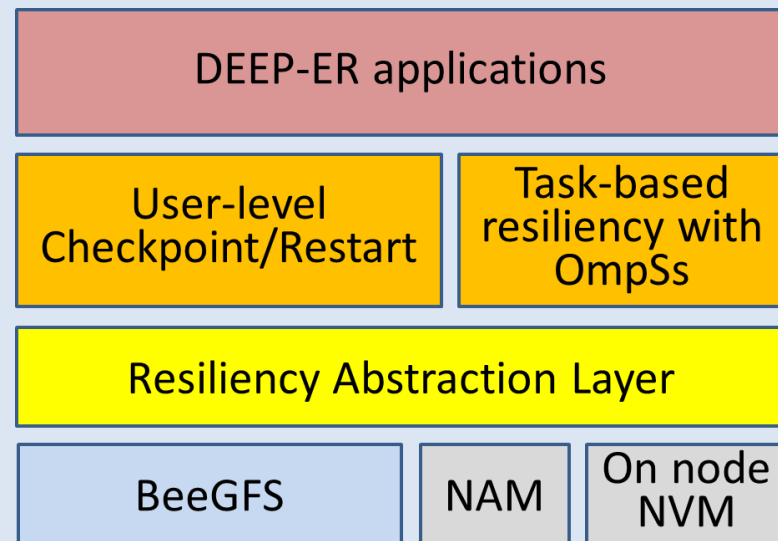
- **Improve I/O scalability on all usage-levels**

- BeeGFS leverages DEEP architecture and novel memory technology
- Extended I/O APIs combine performance with ease of use
  - SIONlib
  - E10



**guide the development by synthetic I/O benchmarks, resiliency scheme and real-world applications**

- **Develop a hierarchical, distributed checkpoint/restart scheme leveraging DEEP-ER architecture**
  - Stage checkpoints in NVM and NAM close to the Booster Nodes
  - Provide checkpoint/restart APIs for task-based “MPI offload” model
  - Develop OmpSs extensions for automatic task resiliency




**with applications guiding the development and validating the results**

- **DEEP+DEEP-ER applications:**

- Brain simulation (EPFL)
- Space weather simulation (KULeuven)
- Climate simulation (CYI)
- Computational fluid engineering (CERFACS)
- High temperature superconductivity (CINECA)
- Seismic imaging (CGGVS)
- Human exposure to electromagnetic fields (INRIA)
- Geoscience (BADW-LRZ)
- Radio astronomy (Astron)
- Oil exploration (BSC)
- Lattice QCD (UREG)

- **Goals:**

- Co-design and evaluation of DEEP architecture and its programmability
- Analysis of the I/O and resiliency requirements of HPC codes



- **Hardware status:**

- DEEP Cluster (@JSC)
- Booster Chassis (32 KNCs) (@JSC)
- ASIC Evaluator (32 KNCs) (@UniHD)
- Energy Efficiency Evaluator (16 KNCs) (@LRZ)

- **Software status:**

- System software implemented
- Validation on DEEP Hardware ongoing
- Programming model completed:
  - Global MPI + OmpSs: offload of highly parallel tasks


- **Scientific Applications:**

- Optimised (vectorisation, threading)
- Application division implemented

**Booster Chassis**



**ASIC Evaluator**





- **Hardware status:**

- Overall architecture design finished
- NVM under evaluation with applications
- NAM in development

- **Software status:**

- Same environment as in DEEP
- Extensions for I/O and resiliency:
  - I/O: BeeGFS, SIONlib, Exascale10
  - Resiliency: application-based + task-based checkpoint

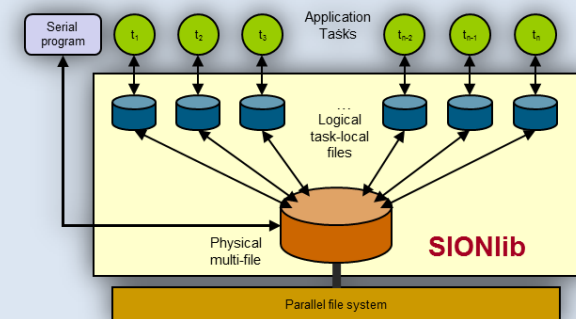
- **Scientific Applications:**

- Applications analysed, optimisations ongoing

## On-node NVM (Intel DC P3700)



## SIONlib



- Exascale poses challenges
  - Energy, Resiliency, Scalability, Programmability
  - Have to face more and huger levels of parallelism
  - Computing will become (even more) heterogeneous
- Some new ideas are around → DEEP
  - allows to map application's levels of scalability onto hardware
  - follows new approaches for the programming paradigm
  - **handles heterogeneity in an innovate way**
- Address also I/O and resiliency → DEEP-ER
- More info: <http://www.deep-project.eu>  
<http://www.deep-er.eu>

EU-Exascale projects  
20 partners  
Total budget: 28,3 M€  
EU-funding: 14,5 M€  
Nov 2011 – Sept 2016



Visit us at  
PRACE Days'15  
Dublin,  
26.05.2015