Cineca and Marconi100

Fabio Affinito

Cineca in a nutshell

- Cineca is a consortium of italian Universities and institutions of research
- Cineca has headquarters in Bologna, with offices in Milan, Rome, Naples and Chieti
- Cineca is the main italian public HPC center, it is affiliated to PRACE and it is one of next sites of the EuroHPC pre-exascale systems

Marconi briefly

The main Supercomputer @ Cineca is MARCONI.
It was installed at CINECA in 2016,
available for Italian and European research community.

It was a Lenovo NeXtScale system based on Intel technology, expected to be incrementally update so to reach a final performance of 50 Pflops/peak in few years.

The last (and final) upgrade occurred early this year (May 2020) by substituting the KNL partition with an accelerated cluster engineered by IBM.

Marconi briefly

- A1: mid 2016 (Intel BDW)
- A2: Jan 2017 (Intel KNL)
- A3: mid 2017 (Intel SKL)
- A4: Jan2018 (Intel SKL +) substitute BDW
- A5: Jan2019 (Intel SKL ++)
- **M100: May 2020** (IBM Power9+GPUs) substitute KNL

Marconi

Next phase, from 2021, foresee the installation of a pre-exaflops system, thanks to the contribution of European Community

Leonardo

Cineca HPC infrastructure

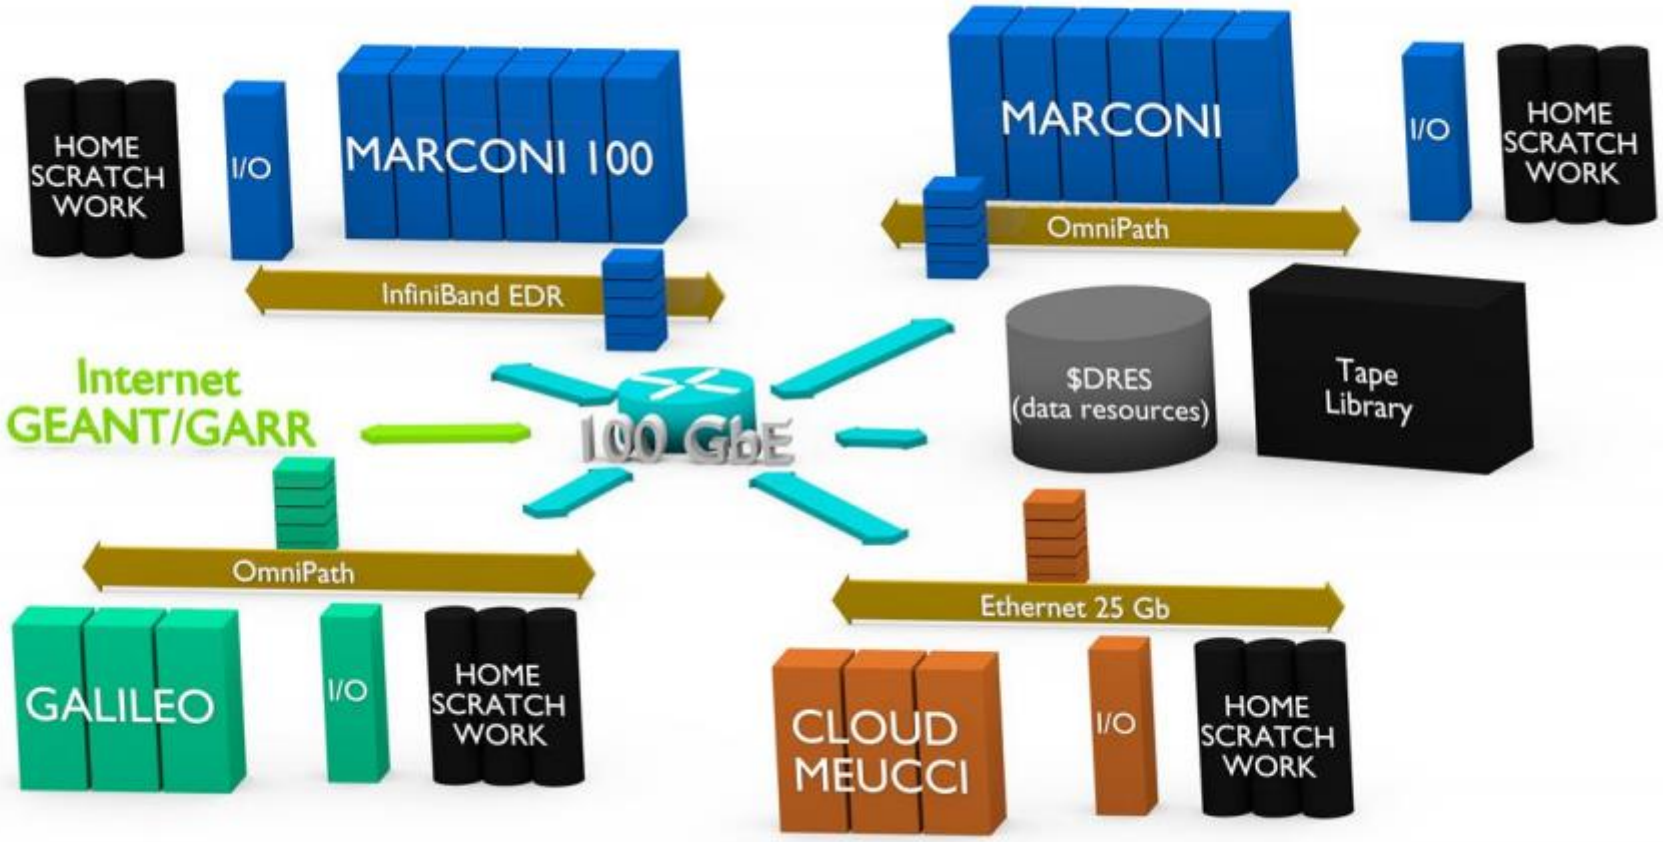

Marconi100 system overview

Marconi 100 is an IBM AC922 (Whiterspoon) cluster

55 racks

980 nodes

Mellanox IB EDR DragonFly++

2 x Power9 CPU

16 cores each
4 HW threads each

4 x NVIDIA Volta V100 GPU
Nvlink 2.0, 16GB

256 GB/node

Marconi100 node architecture

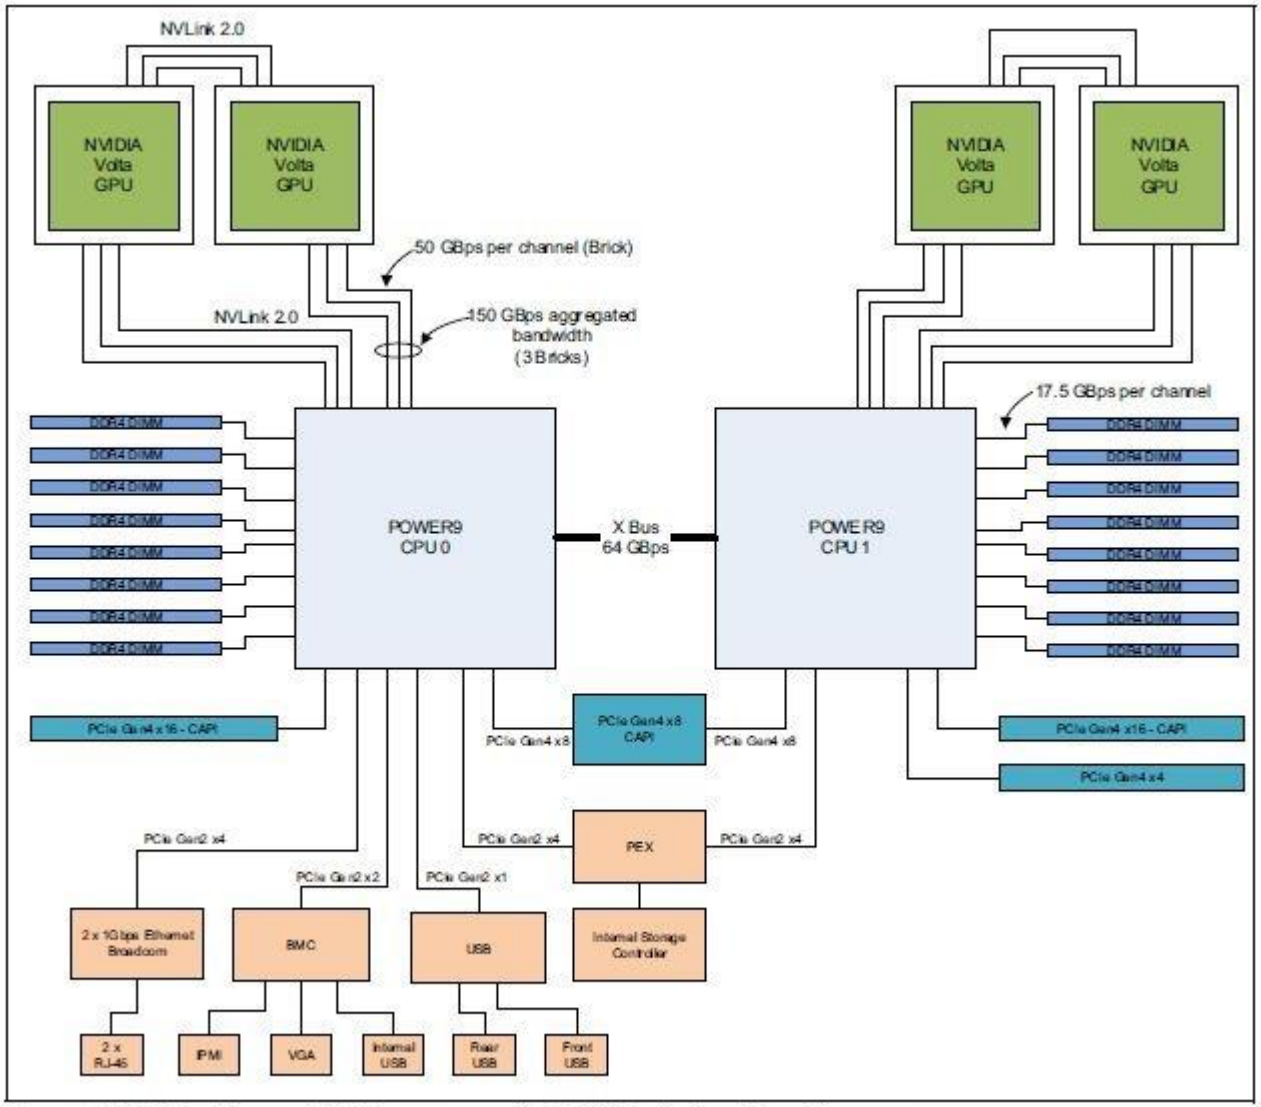

Figure 2-5 The Power AC922 server model GTH logical system diagram

IBM Power9

Each node has two IBM Power9 sockets

2 (socket) x 16 (cores) x 4 HW threads

Total: 128 threads on the node

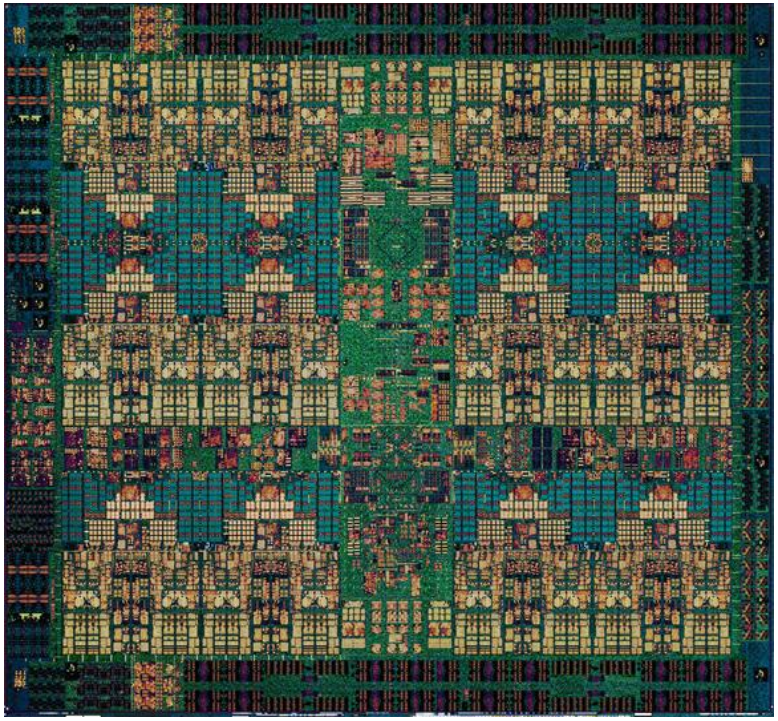

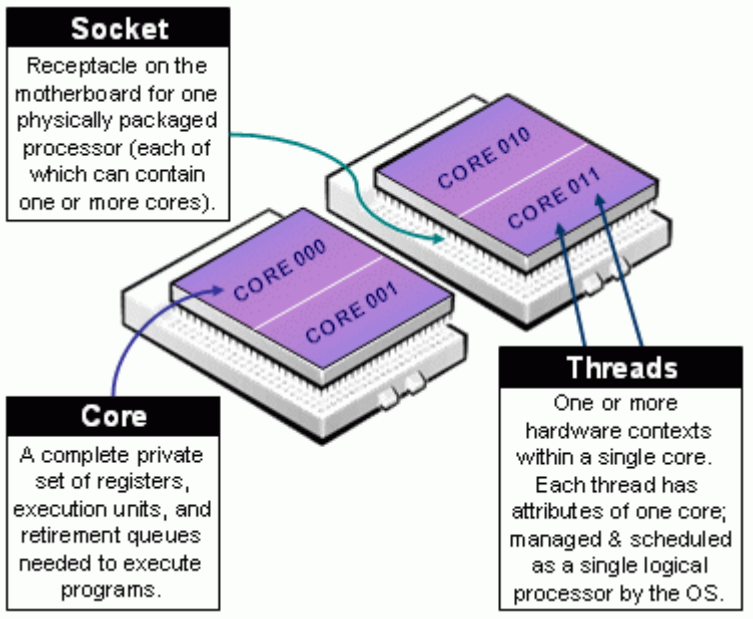

NVIDIA V100 GPU

Each Marconi100 node has four V100 GPUs

Each Tesla V100 GPU has:

- 150+150 GB/s total BW (**NVLink** v2.0)
- 5,120 CUDA cores (64 on each of 80 SMs)
- 640 Tensor cores (8 on each of 80 SMs)
- 20MB Registers | 16MB Cache | 16GB HBM2 @ 900 GB/s
- 7.5 DP TFLOPS | 15 SP TFLOPS | 120 FP16 TOPS

nVIDIA **GPUDirect** technology is fully supported (shared memory, peer-to-peer, RDMA, async), enabling the use of CUDA-aware MPI

Marconi100 Software Stack

- **Compilers**
  - **XL (IBM compilers: xlf90, xlC, etc.)**
  - GNU (gcc, gfortran)
  - **PGI**
  - CUDA
- **Communication Libraries**
  - **Spectrum_MPI**
  - OpenMPI
- **Libraries**
  - ESSL, BLAS, LAPACK, FFTW
  - HDF5, …

Support for OpenACC and CUDA Fortran

Fully optimized for M100 architecture

Module environment – Compilers and libraries

```
[faffinit@login03 ~]$ ml
Currently Loaded Modulefiles:
 1) profile/base
[faffinit@login03 ~]$ ml av
------------------------------------------- /cineca/prod/opt/modulefiles/profiles --------------------------------------------
profile/advanced  profile/base  profile/candidate  profile/chem-phys  profile/deeplrn  profile/global  profile/lifesc

---------------------------------------- /cineca/prod/opt/modulefiles/base/environment ----------------------------------------
autoload

----------------------------------------- /cineca/prod/opt/modulefiles/base/libraries -----------------------------------------
blas/3.8.0--gnu--8.4.0                      lapack/3.9.0--pgi--19.10--binary              zlib/1.2.11--gnu--8.4.0
boost/1.72.0--spectrum_mpi--10.3.1--binary  nccl/2.6.4--cuda--10.1
essl/6.2.1--binary                          netcdf/4.7.3--gnu--8.4.0
fftw/3.3.8--gnu--8.4.0                      netcdf/4.7.3--spectrum_mpi--10.3.1--binary
fftw/3.3.8--spectrum_mpi--10.3.1--binary    netcdff/4.5.2--gnu--8.4.0
gsl/2.6--gnu--8.4.0                         netcdff/4.5.2--spectrum_mpi--10.3.1--binary
hdf5/1.12.0--gnu--8.4.0                     petsc/3.12.4--spectrum_mpi--10.3.1--binary
hdf5/1.12.0--spectrum_mpi--10.3.1--binary   scalapack/2.1.0--spectrum_mpi--10.3.1--binary
lapack/3.9.0--gnu--8.4.0                    szip/2.1.1--gnu--8.4.0

----------------------------------------- /cineca/prod/opt/modulefiles/base/compilers -----------------------------------------
cuda/10.1  gnu/8.4.0  pgi/19.10--binary  python/3.8.2  spectrum_mpi/10.3.1--binary  xl/16.1.1--binary

------------------------------------------ /cineca/prod/opt/modulefiles/base/tools ------------------------------------------
anaconda/2020.02  cmake/3.17.1  singularity/3.5.3  spack/0.14.2-prod  superc/2.0
[faffinit@login03 ~]$
```

For more detail, see the Marconi100 User Guide:
https://wiki.u-gov.it/confluence/display/SCAIUS/UG3.2%3A+MARCONI100+UserGuide

How to access HPC resources

http://iscra.cineca.it

https://prace-ri.eu/hpc-access/