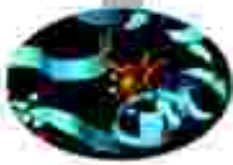

A2 - Knights Landing configuration on Marconi

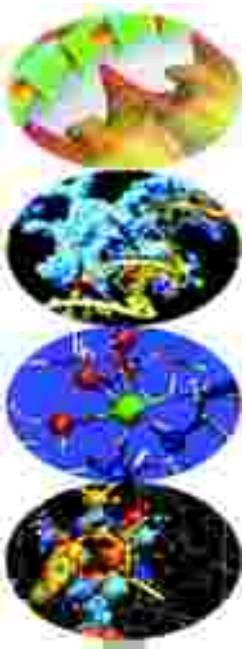

Silvia Giuliani @ SCAI Cineca
Alessandro Grottesi @ SCAI Cineca

A2 Location

50 racks

Sala F - 46 racks (pink)

Sala M - 4 racks (pink)

A2 Cooling system

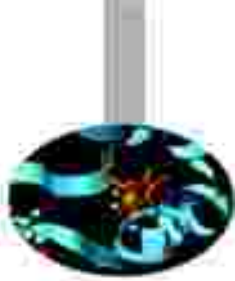

Water based free cooling system

- Free cooling is an approach to lowering the air temperature in a building or data center by using naturally cool air or water instead of mechanical refrigeration

- The water cooling system is installed on racks doors of A2 system

Marconi A2 Configuration

- 50 racks
- 72 nodes/rack

TOT nodes 3,600

- 68 cores/node

TOT cores 244,800

- 45 racks for accademic users
- 5 racks for eurofusion users

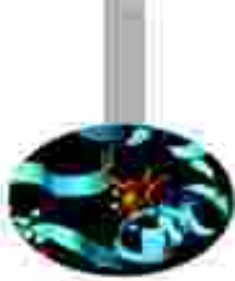

A2 Compute node

- 1x Intel Xeon Phi 7250
- **Xeon Phi**
  - Many Integrated Core (Mic)
    - ➔ Knights Corner (Galileo): compute node accelerator
    - ➔ Knights Landing (Marconi A2): standalone processor

A2 - Compute node

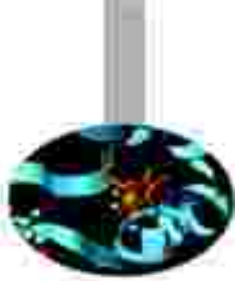

- **Processor number: 7250**
  - best performance/watt
  - 68 core
  - Hyperthreading: 272 threads
  - 1.4 GHz of clock
  - 1.60 GHz of clock with Turbo Boost → off
  - Intel(R) AVX-512 (512-bit extensions to the 256-bit Advanced Vector Extensions)

A2 Login node

- 2 x Intel Xeon Processor E5-2697 v4
- 2.30GHz and 128 GB of memory.
- 3 nodes available for regular users
- 6 service nodes for cluster management

- Users login nodes are shared between A1, A2 and A3.
- Service login nodes are separated (6 for each partition), although they are identically configured.

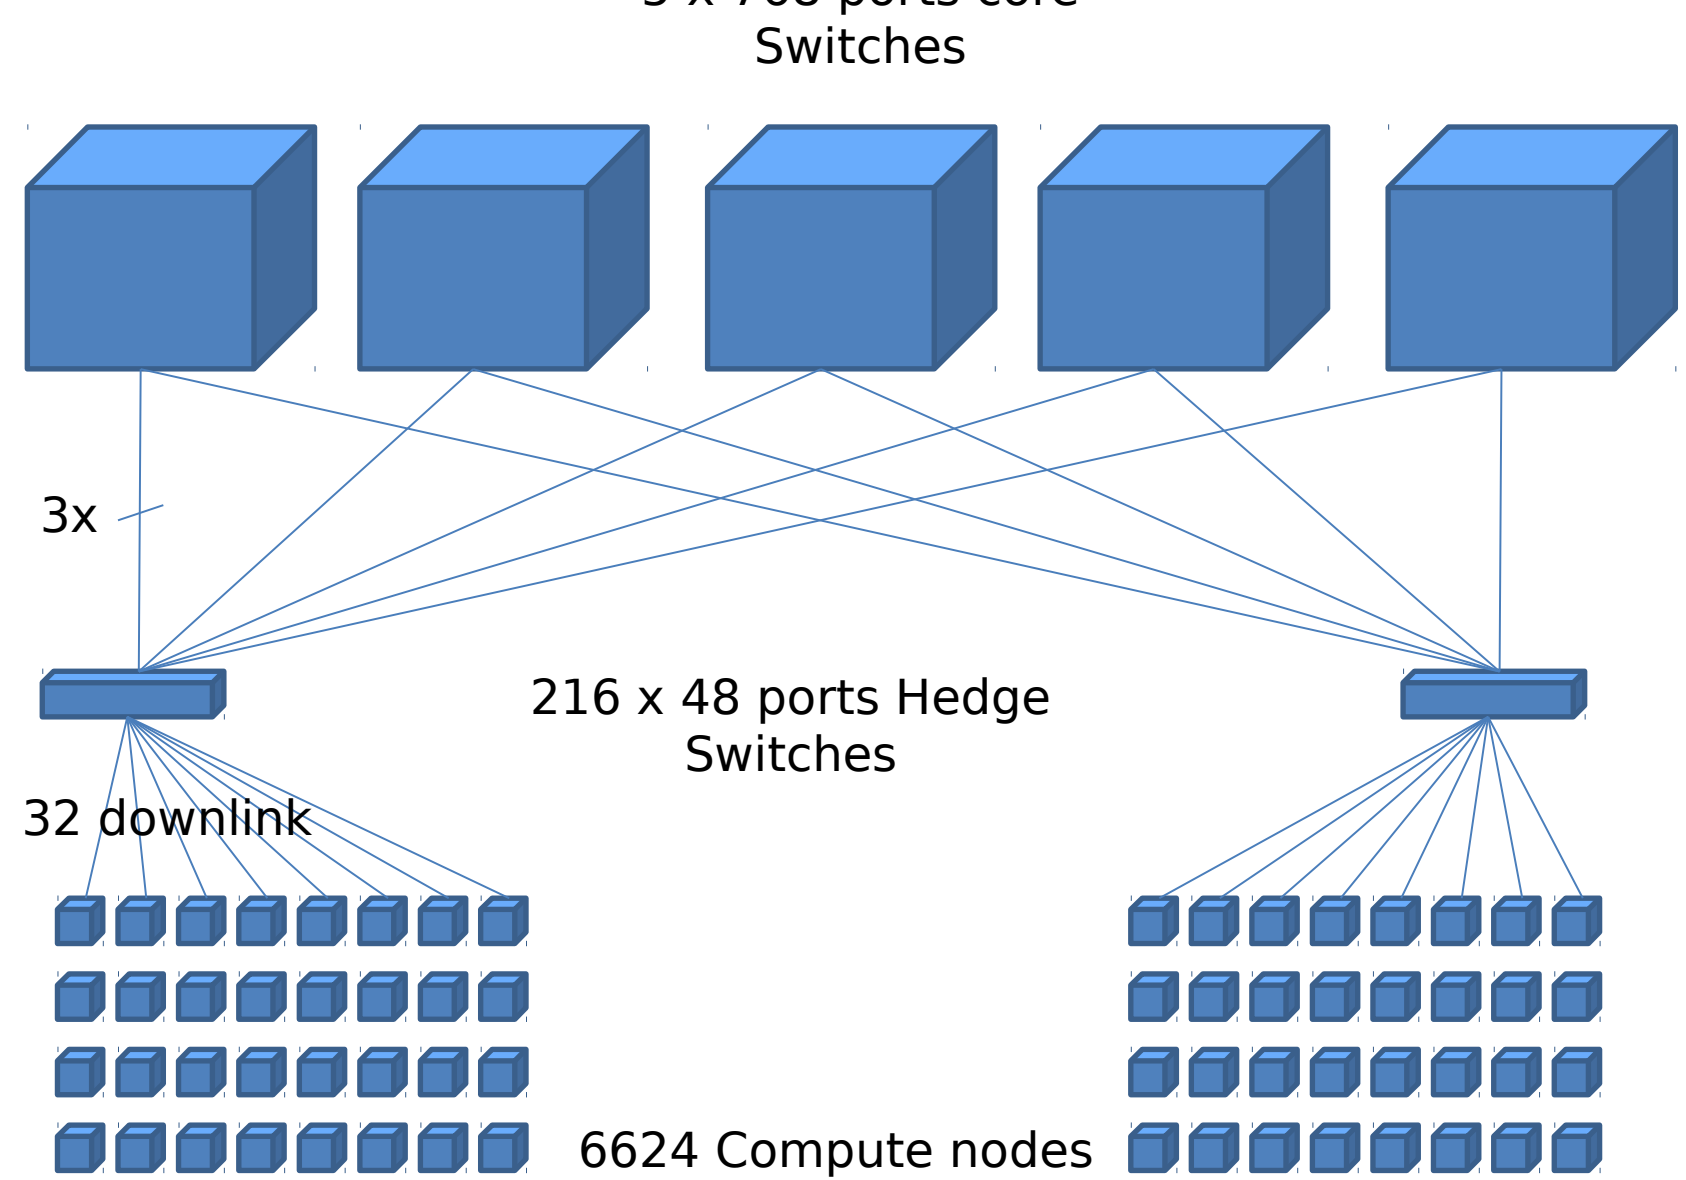
5 x 768 ports core
Switches
3x
216 x 48 ports Hedge
Switches
32 downlink
6624 Compute nodes
32 nodes
fully interconnected island

Marconi Network

- Network type: new Intel Omnipath, 100 Gb/s. MARCONI is the largest Omnipath cluster of the world.
- Network topology: Fat-tree 2:1 oversubscription tapering at the level of the core switches only.
- Core Switches: 5 x OPA Core Switch "Sawtooth Forest", 768 ports each.
- Edge Switch: 216 OPA Edge Switch "Eldorado Forest", 48 ports each.
- Maximum system configuration: 5(opa) x 768 (ports) x 2 (tapering) → 7680 servers.

How to login to Marconi A2

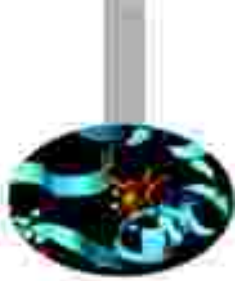

- **SSH – to access the cluster from your pc**
  ssh <username>@login.marconi.cineca.it
  - by giving the password
  - by a public key (without using the password)
  - create the key on your pc:
    $ ssh-keygen
    $ ls –l .ssh
    -rw------- 1 ... 668 Apr 26 14:56 id_rsa
    -rw-r--r-- 1 ... 601 Apr 26 14:56 id_rsa.pub
  - copy the key to the destination cluster
    scp id_rsa.pub <username>@login.marconi.cineca.it:/$HOME/
    $ cat id_rsa.pub >> $HOME/.ssh/authorized_keys

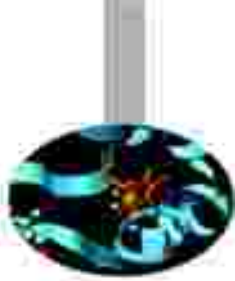

Local Disk Spaces

User Spaces

- **$HOME** →/marconi_home/userinternal/<username>
  - Permanent (dependent of the life of the user) and backed-up
  - Quota = 50 GB by default
  - For storing source code, executables, configuration files or important input files
  - For compiling your program
  - Mounted on the login and compute nodes

- **$CINECA_SCRATCH** →/marconi_scratch/userinternal/<username>
  - Temporary (files older than 50 days automatically deleted)
  - Not backed-up and parallel filesystem (GPFS)
  - No quota
  - For production testing, for temporary output files
  - Mounted on the login and compute nodes

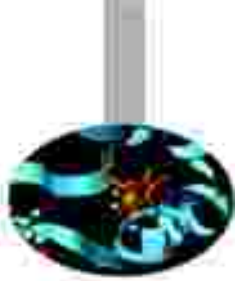

Local Disk Spaces

Project Spaces

- **$WORK** →/marconi_work/<project_name>
  - Permanent (deleted six months after the end of the corresponding project)
  - NOT backed-up and parallel filesystem (GPFS)
  - 1 TB default quota
  - For production activity
  - Mounted on the login and compute nodes

Local Disk Spaces

Project Spaces

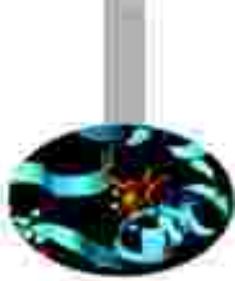

$WORK

- "chprj" command
  - -l list your PROJECTS
  - -d <project_name> set your default project for $WORK

- Unix file permissions:

  $WORK →/marconi_work/<default project name>
  Owner: PI
  UNIX group: project_name
  drwxrwx--- 29 PI project_name 4096 May 17 15:11 .

  All collaborators of the project and the PI can write into $WORK

  $WORK/subdir or file
  Owner: subdir creator (PI or collaborator)
  UNIX group: interactive
  drwxr-xr-x 29 Collaborator Interactive 4096 May 17 15:11 .

  In order to sharing personal data between all collaborators of the project:

```
chgrp –R subdir project_name      # change unix group
chmod –R 770 subdir               # add rwx permissions to group
```

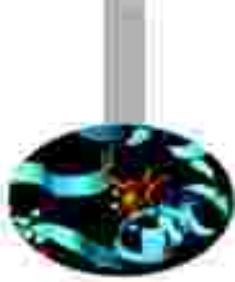

Shared Disk Spaces

You need to ask for this kind of resource explicitly, it does not come as part of a project
(mailto: superc@cineca.it)

- USER SPACE
- $TAPE → /gss/gss_work/tape/userexternal/<username>
  - conceived for saving “personal” data on magnetic media
  - shared among platforms
  - Quota=500 GB
  - mounted on the login nodes
  - mounted on the compute nodes only of PICO cluster

PROJECT SPACE

- $DRES
  - conceived for saving “project” data on magnetic media
  - shared among platforms and projects
  - mounted on the login nodes
  - mounted on the compute nodes only of PICO cluster
  - FS: normal filesystem access on hight throughput disks
  - ARCH: magnetic tape archiving with a disk-like interface via LTFS
  - REPO: smart repository based on iRODS