CINECA SCAI
SuperComputing Applications and Innovation

MARCONI

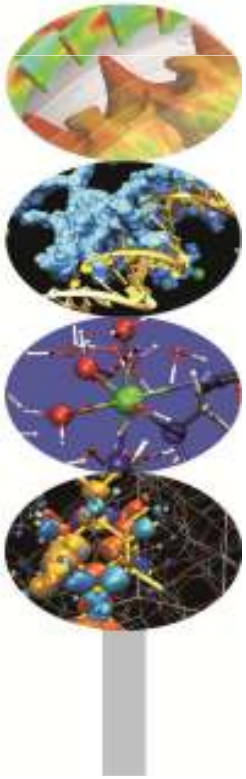

Installation roadmap and resources allocation procedure

Elda Rossi

CINECA

CINECA SCAI
SuperComputing Applications and Innovation

A new Supercomputer (codename: MARCONI) is being installed at CINECA, available for Italian and European research community.

It is a Lenovo NeXtScale system based on Intel technology, with a final peak performance around 20PFlop/s.

Deployment of Marconi is starting mid-April, the complete delivery expected as July 2017.

The first part (2 PFlop/s conventional "Broadwell" based) is expected to be in full production as July 2016.

CINECA

hpc.cineca.it → ForUsers → Documentation

A short story

- In 2015 the computing resources in Cineca were:
  - Tier-0: **FERMI** (acquired in summer 2012)
  - Tier-1: **GALILEO** (acquired in Jan 2015)
  - Front-end, Viz, BigData: **PICO** (acquired in Nov 2014)
- FERMI arrived at the expected end of its activity.
- The Cineca governing bodies, aimed at supporting scientific research, approved a development plan, with an investment of Euro 50 million in two phases, from 2016 to 2020:
  - 2 x 5 → 10 Pflops in 2016-2017
  - 10 x 5 → 50 Pflops in 2019-2020

Marconi

CINECA SCAI
SuperComputing Applications and Innovation

MARCONI: the new Tier-0 system

- A tender was issued in 2015 and assigned Jan 2016 to lenovo
- The system will be delivered in three phases:
  - A1: April 2016 (BRD 2 PFs)
  - A2: Sept 2016 (KNL 11 PFs)
  - A3: July 2017 (SKL 5 PFs)
- In total:
  - 18 PFs peak performance
  - 17PB row storage
  - 3MWatt power consumption

CINECA

CINECA SCAI
SuperComputing Applications and Innovation

MARCONI: New Tier-0 system

Technical Features:

- Intel based
- Architecture: Lenovo NeXtScale
- Fabric: Intel OmniPath

A2

KNL 68cores, 1.4 GHz;
3600 nodes, 11 PFs

A3

SKL 2x20 cores, 2.3 GHz;
1500 nodes, 5 (+2) PFs

CINECA

Timings

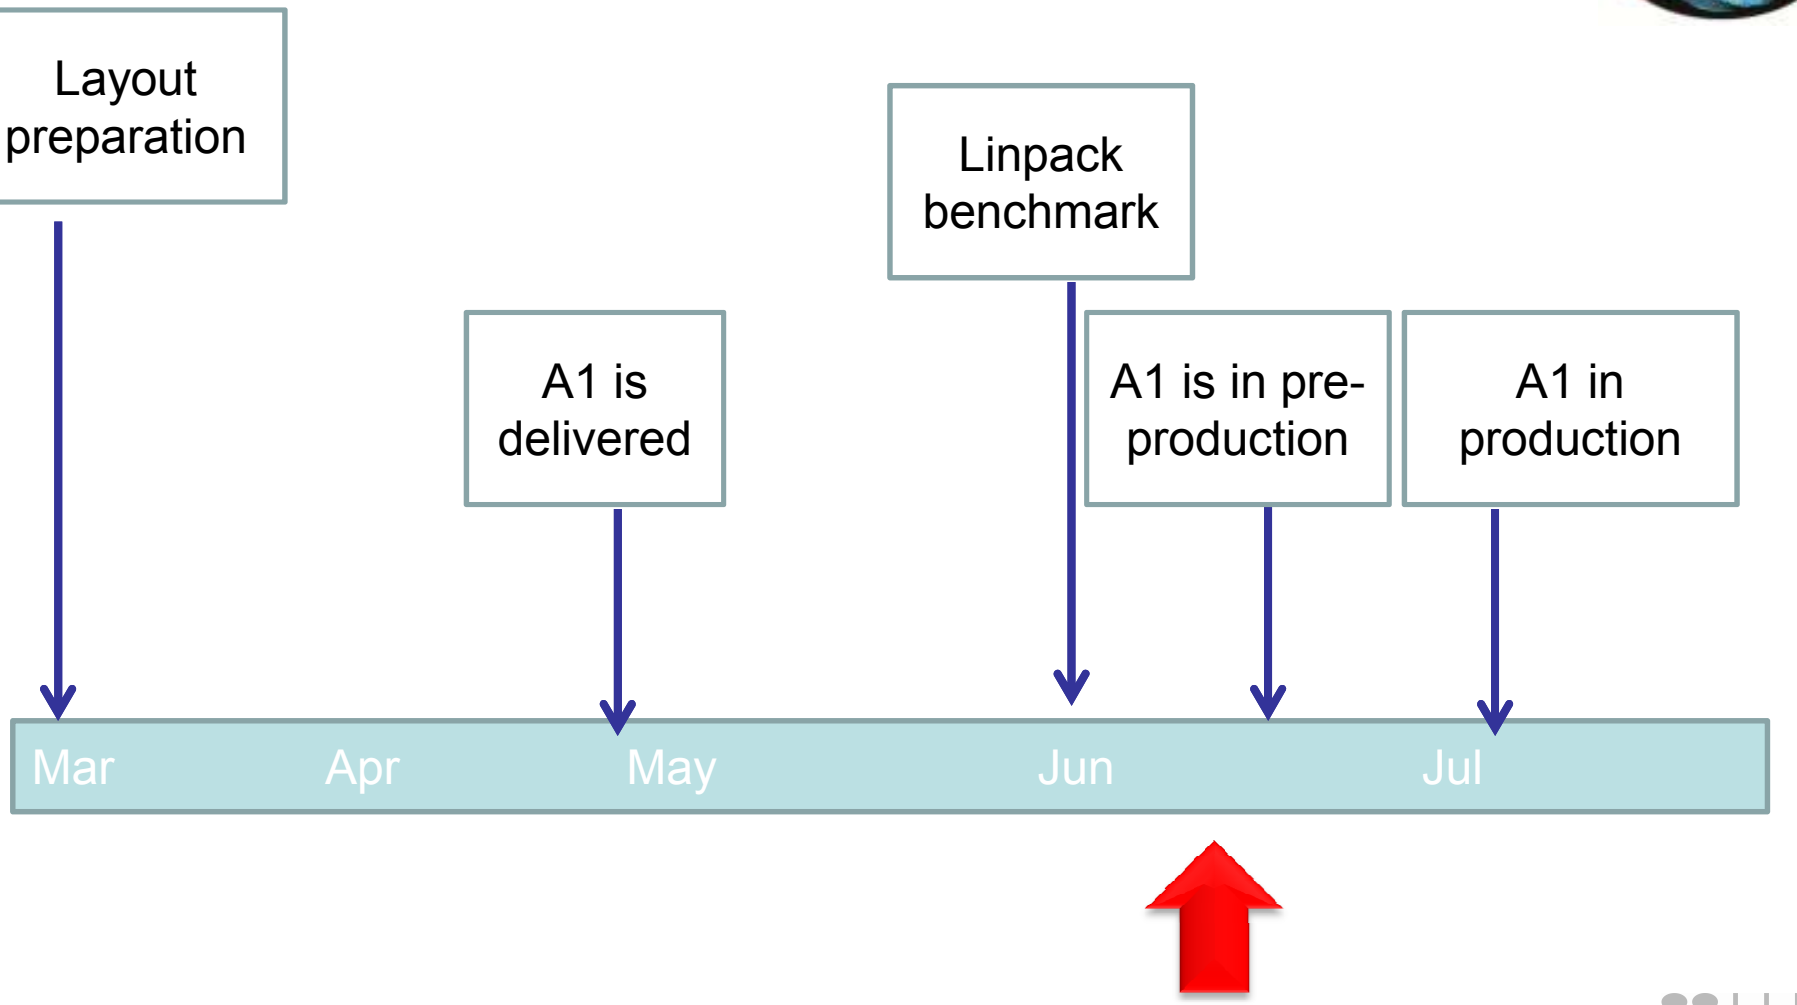
Layout preparation
A1 is delivered
Linpack benchmark
A1 is in pre-production
A1 in production
Mar
Apr
May
Jun
Jul
(on time)

CINECA SCAI
SuperComputing Applications and Innovation

MARCONI installation news

lenovo

- All hardware equipments already in Cineca (last shipping on Friday, Apr 29)
- Still missing 8 login nodes (using 2 management nodes in the meantime)
- Set up of the machine started on (Apr 26)
- Linpack benchmark for classification in TOP500 and acceptance: (June 1-3)
- System administrators starting configuration (June 6)
- User environment configuaration by UserSupport (June13)
- Pre-production (June 20)
- Full production (July 4)

CINECA

CINECA SCAI
SuperComputing Applications and Innovation

Technical details

A1 (half reserved to EUROfusion)

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
2PFs	1512	2x Intel Broadwell 18cores @2.3GHz	128 GB	Intel OmniPath 2:1 100Gb/s	21	10 Front End Nodes (2xBDW 18c +128GB RAM)+ 2 MGMT nodes

Core tot: 54.432
Core-h/anno=476.824.320

A2

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
11 PFs	3600	Intel KnightsLanding 68cores @1.4 GHz	96 GB	Intel OmniPath 2:1 100Gb/s	50	4 Front End Nodes (2xBDW 18c +128GB RAM)+ 48 I/O nodes

Core tot: 244.800
Core-h/anno=2.144.448.000

CINECA

CINECA SCAI
SuperComputing Applications and Innovation

Technical details

A3 (great part reserved to EUROfusion)

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
5PFs	1512	2x Intel SkyLake 20cores @2.3GHz	192 GB	Intel OmniPath 2:1 100Gb/s	21	2 Front End Nodes + 2 MGMT nodes (2xSKL 20c +192GB RAM) (2xSKL 20cc +

Core tot: 60.480

Core-h/anno=529.804.800 (500 M core-h/y)

CINECA

Preliminary schema of A1

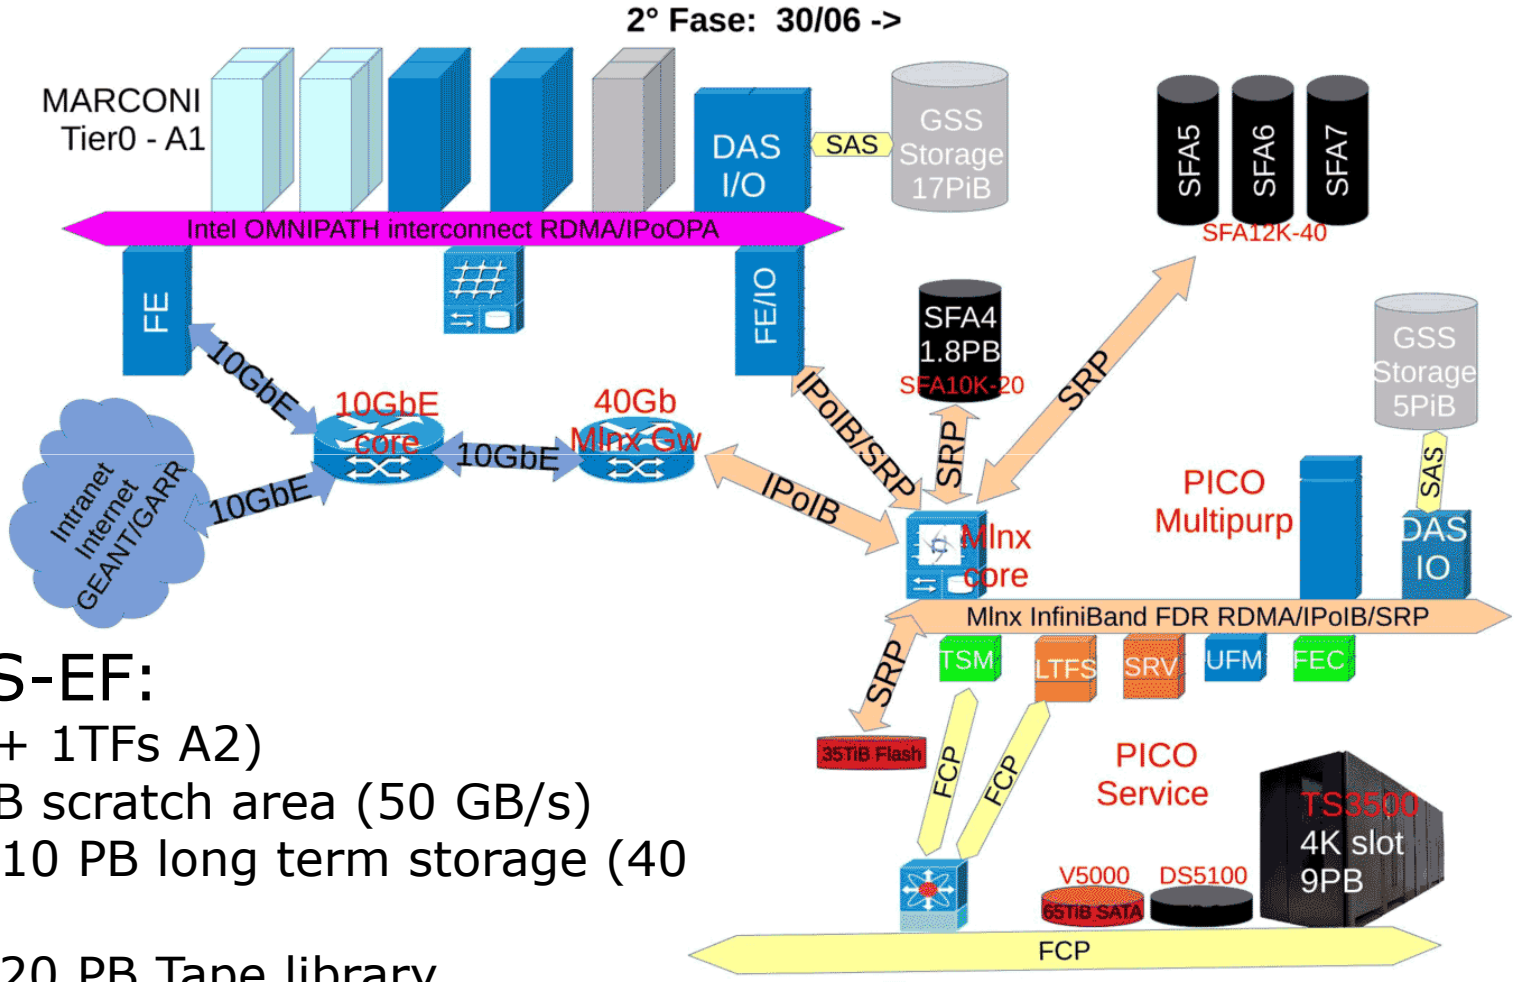

POLARIS-EF:

1 TFs A1 (+ 1TFs A2)
GSS*: 5 PB scratch area (50 GB/s)
by Jan17: 10 PB long term storage (40 GB/s)
20 PB Tape library

*GPFS: NSD servers layout under design

CINECA

Further details

- NeXtScale architecture
- nx360M5 nodes:
  - Supporting Intel HSW & BDW
  - Able to host both IB network Mellanox EDR & Intel Omni-Path
  - Twelve nodes are grouped into a Chassis (6 chassis per rack)
- The compute node is made of:
  - 2 x Intel Broadwell (Xeon processor E5-2697 v4)
  - 18 cores, 2,3 HGz
  - 8 x 16GB DIMM memory (RAM DDR4 2400 MHz), 128 GB total
  - 1 x 129 GB SATA MLC S3500 Enterprise Value SSD
  - 1 x link OPA 100GBs
  - 2*18*2,3*16 = 1.325 GFs peak
- The login node (10 servers) – x3550M5 1U
  - 2 x Intel Broadwell (Xeon processor E5-2697 v4)
  - 18 cores, 2,3 HGz
  - 8 x 16GB DIMM memory (RAM DDR4 2400 MHz), 128 GB total
  - 2 x 1 TB SAS
  - 1 x link OPA 100GBs + 1 link 1GbE + 2 link 10GbE
- 24 rack in total:
  - 21 rack → compute
  - 1 rack → service nodes
  - 2 racks → core switch

How to get access

- Since the pre-production phase, all active projects on FERMI will be migrated to MARCONI (1:5)
- What about if you do not have projects on FERMI?
  - Agreements: your coordinator can ask to move some standard-h to MARCONI (ratio to be defined)
  - ISCRA, LISA: next calls will open access to MARCONI
  - ??? Ask to UserSupport
- Accounting not active in the first production month (working from Sept, 1st)

CINECA SCAI
SuperComputing Applications and Innovation

CINECA

CINECA SCAI
SuperComputing Applications and Innovation

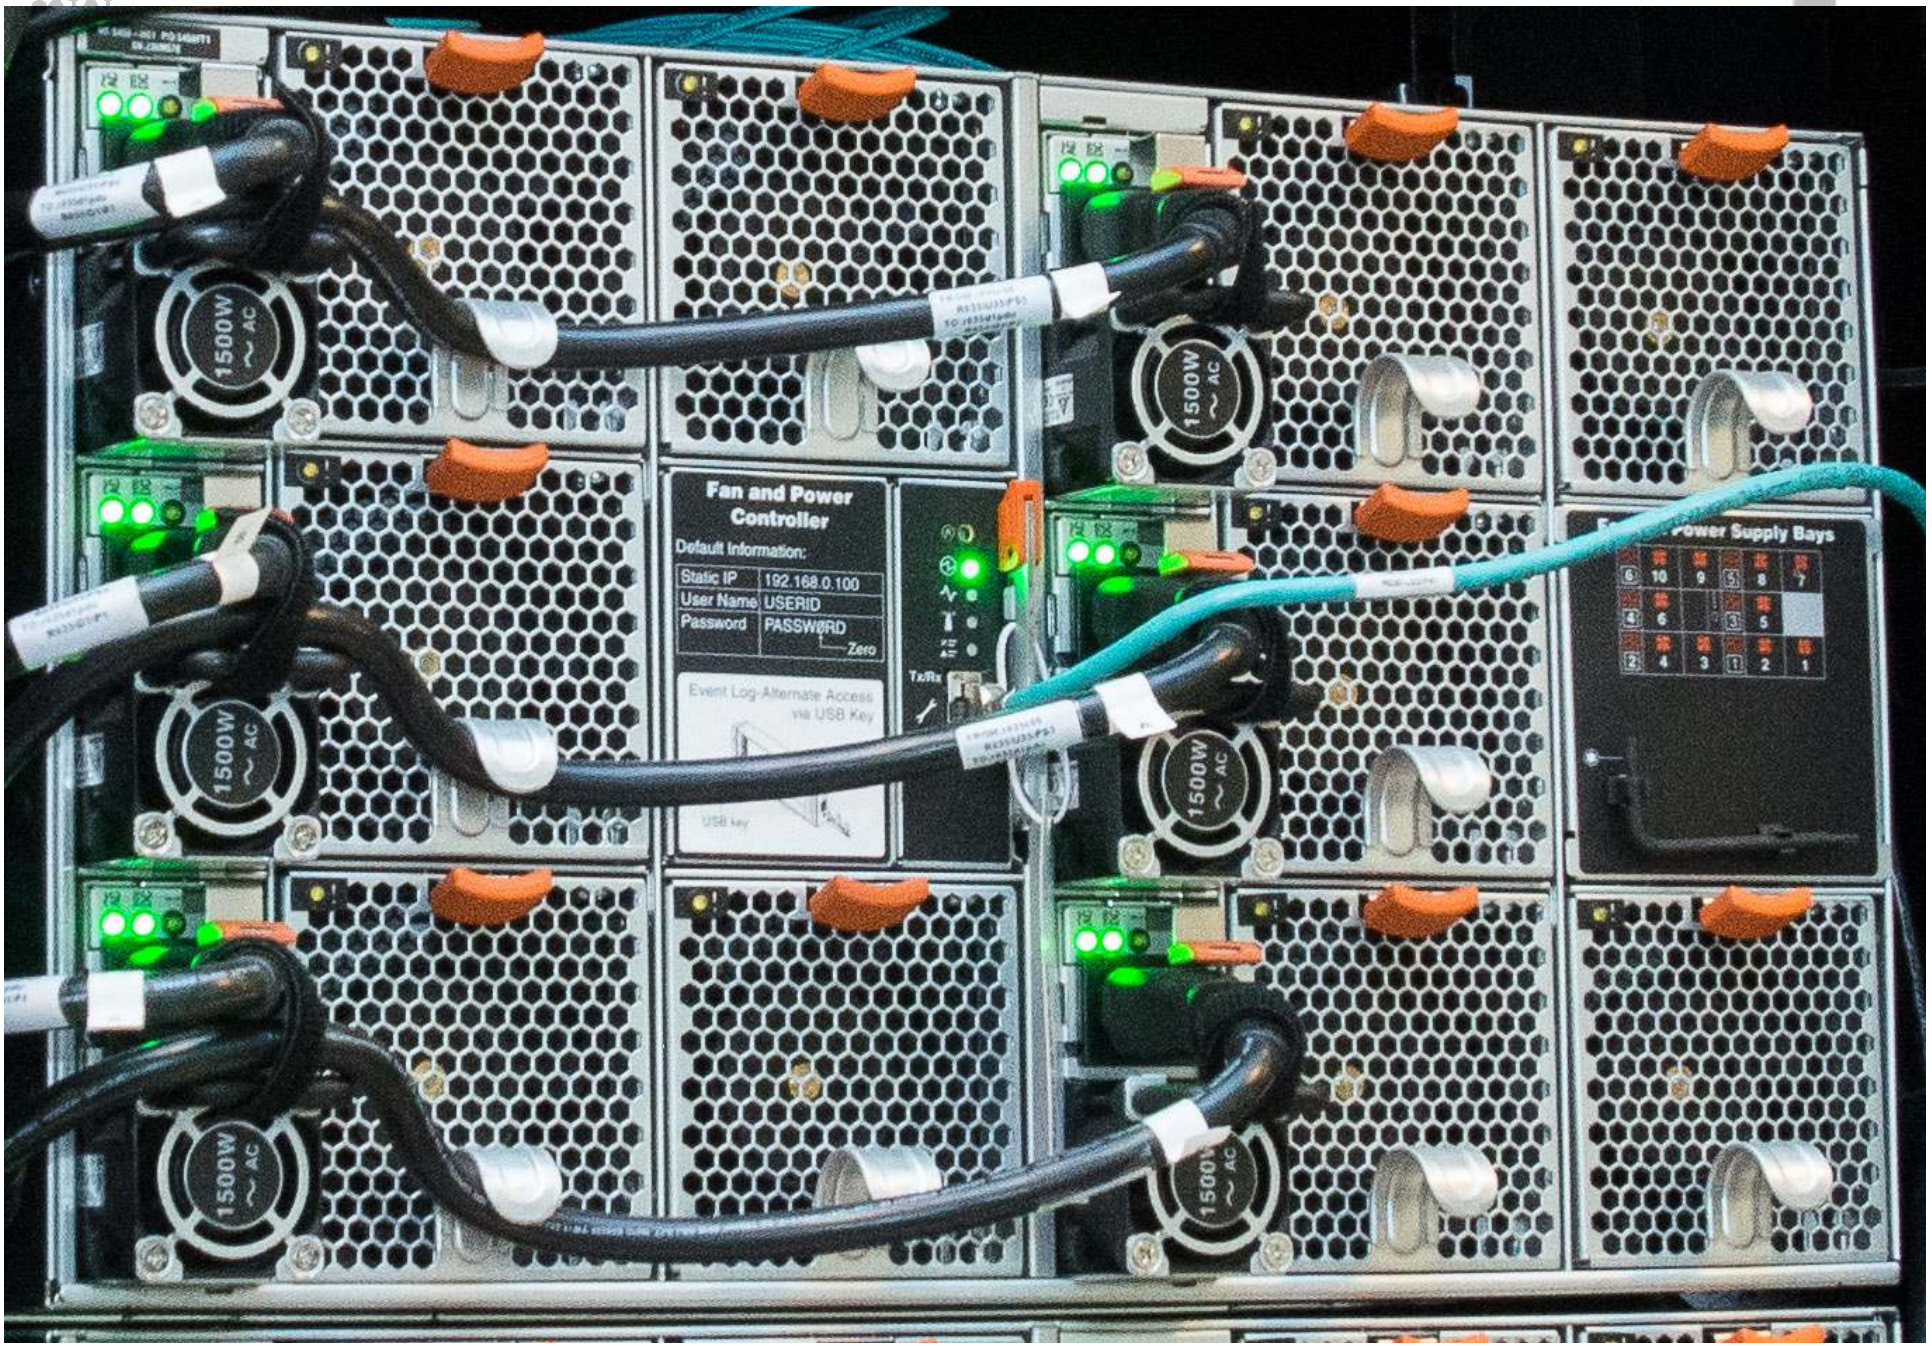
1500W ~ AC
Fan and Power Controller
Default Information:
Static IP 192.168.0.100
User Name USERID
Password PASSWORD
Zero
Event Log-Alternate Access via USB Key
Power Supply Bays

CINECA SCAI

SuperComputing Applications and Innovation

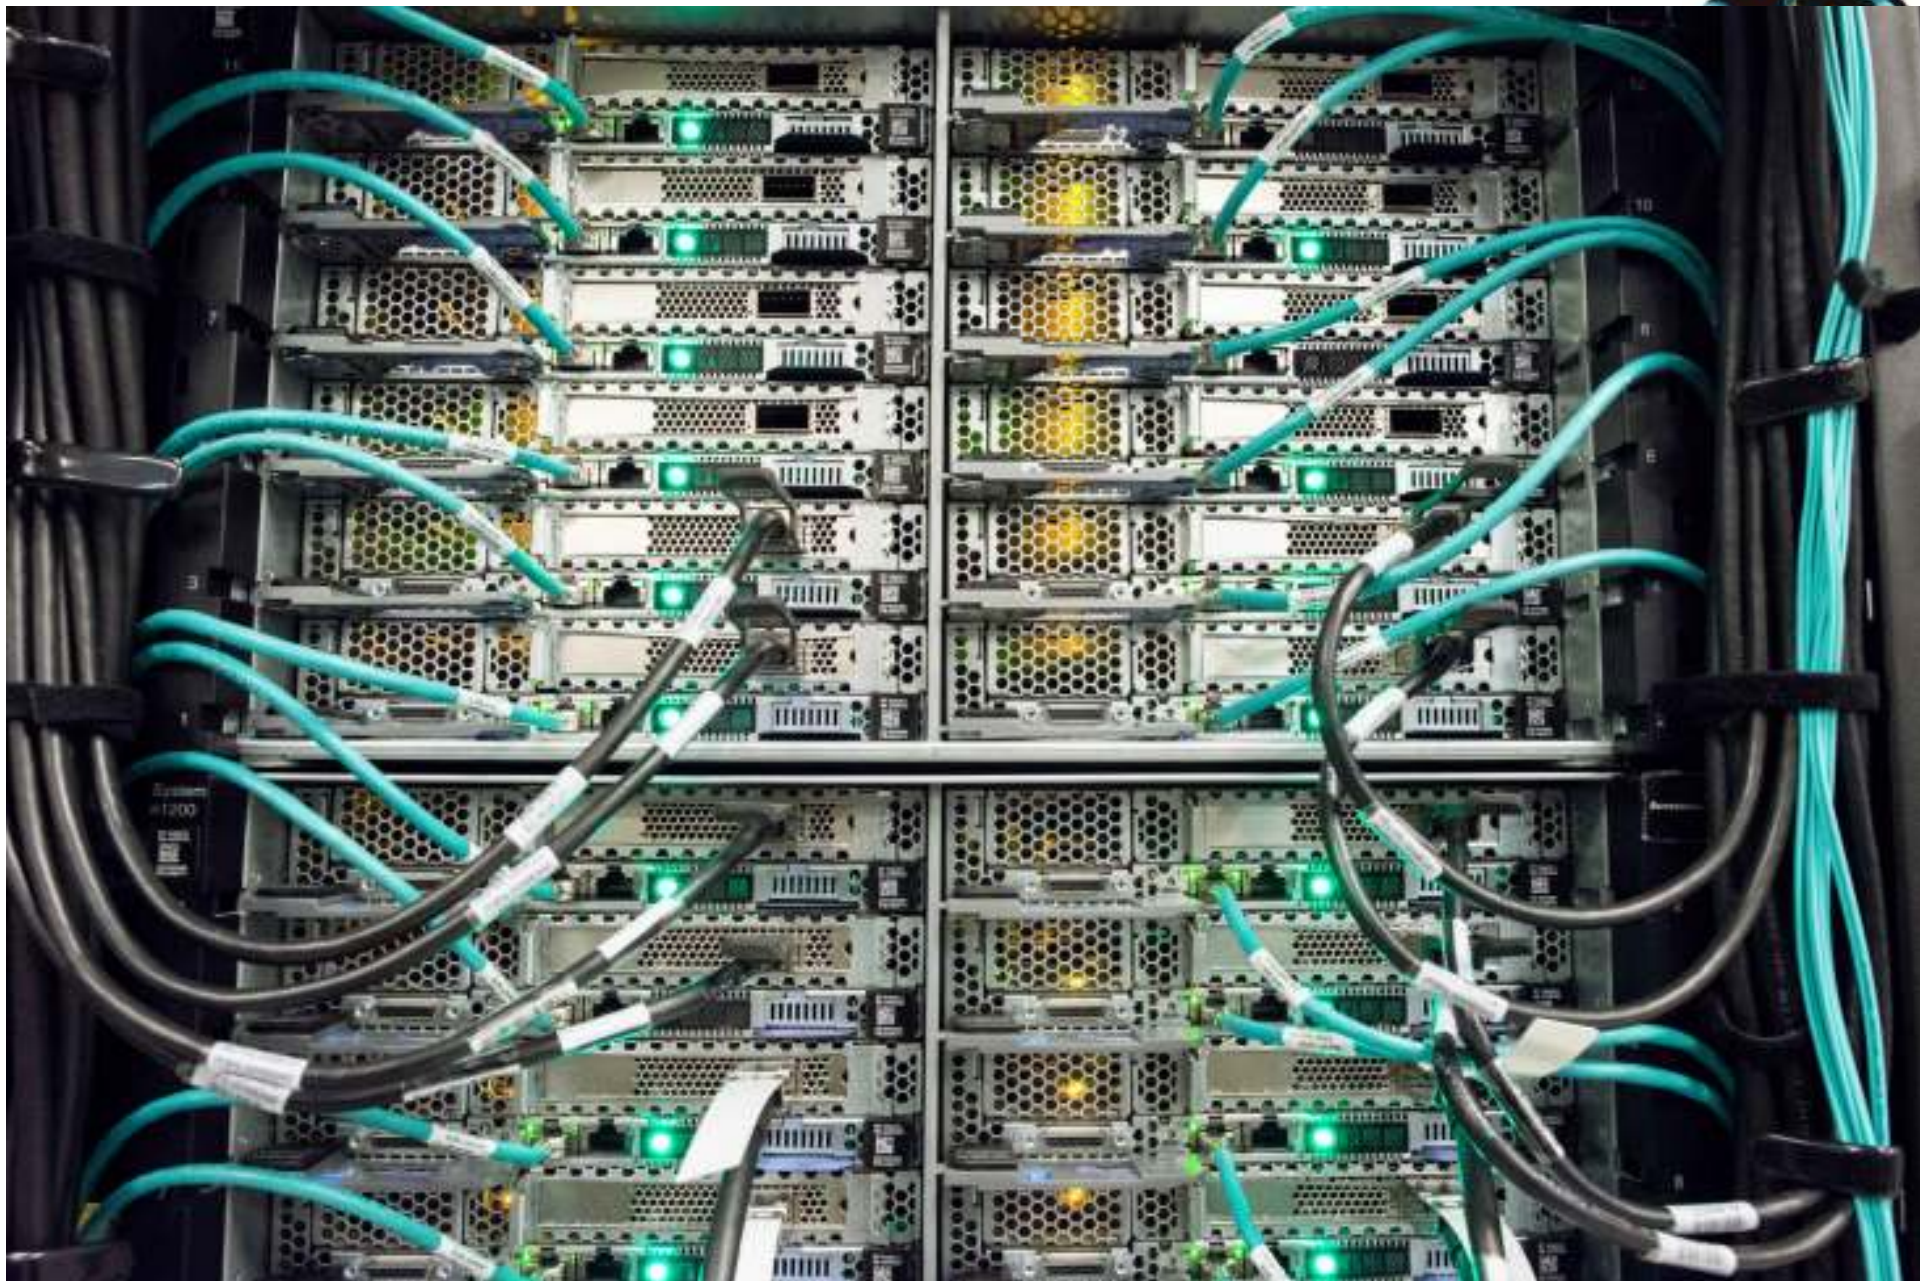

CINECA SCAI

SuperComputing Applications and Innovation