Italian National Agency for New Technologies,
Energy and Sustainable Economic Development

The new CRESCO6 HPC cluster, the latest upgrade of the ENEAGRID computational infrastructure

ENEA CR Portici 30/5/2018

G. Bracco
giovanni.bracco@enea.it

HPC a Lab in ENEA DTE-ICT Division

Information and Communication Technology Division
Services & HPC, (~75 staff) HPC (17 staff)

Services & Activities

- **HPC CRESCO clusters**
- **ENEAGRID Infrastructure**
- **Storage, Databases, Big Data**
- **Scientific Software**
- **3D/Graphic tools & resources**
- **Remote operation of experiments & lab data preservation**
- **Virtual Labs**
- **Application area: material science for energy, fission & fusion nuclear technology, climate and environmental science,complex systems,...**

- **http://www.eneagrid.enea.it**
- **http://www.cresco.enea.it**

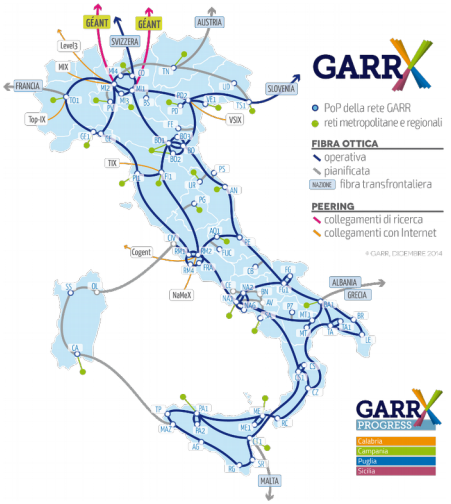

12 main sites connected by GARR, the Italian NREN, 6 Research Centres with HPC/storage resources

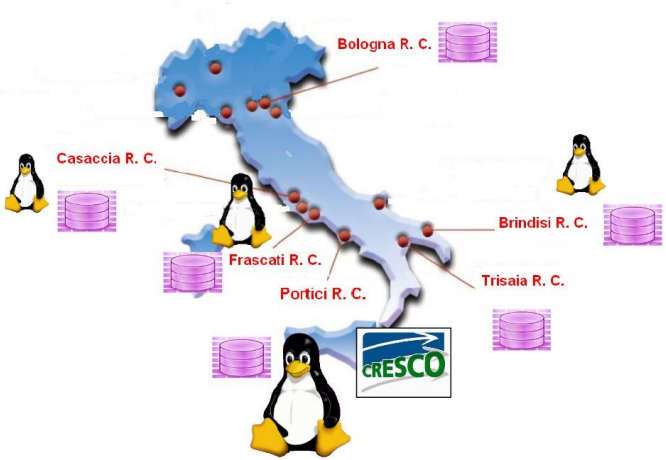

ENEA HPC Infrastructure: ENEAGRID/CRESCO Clusters

ENEA HPC resources are integrated into **ENEAGRID**
The main computing facilities are **CRESCO clusters**
(now in production ~160 Tflops very soon >800 Tflops)
Scientific/Technical Software
Annual reports: **http://www.cresco.enea.it/CRESCO_reports**

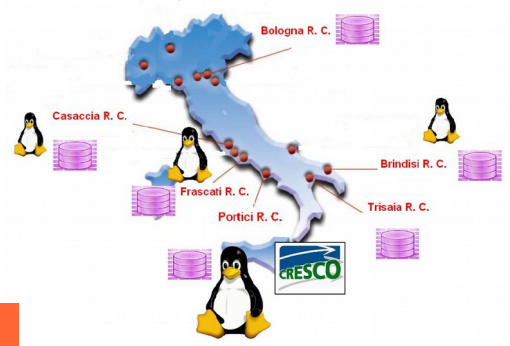

Cluster name (Linux x86_64)	Network	Cores/Tflops
CRESCO3 AMD	IB QDR	1152/~10
CRESCO4 INTEL	IB QDR	4864/~100
CRESCO5 INTELv3	IB QDR	672/~25
CRESCOF AMD	IB QDR	456/~4
CRESCOC AMD	IB DDR	192/~1.5
K40/PHI/Large RAM	IB QDR	300/~17
Total in production		**~7500/~160**
2018 CRESCO6	**SKL OPA**	**~10000/~700**

Research area 2017	Usage (%)
Nuclear energy	33.6
Climate/Env.	25.3
Mat. Sciences	22.7
CFD/Combust.	10.8
Bio Tech.	3.1
Complx Systems	1.7
Aerospace	0.6
Other domains	2.33

Type of users 2017	% wall clock time
ENEA	75%
Italian Univ./Inst.	17%
Foreign Univ./Inst.	7 %
Enterprises	1%
Tot core*hour	34 Millions
Total # users	134





ENEA HPC participation in H2020 projects

Topics of interest and opportunities: **(2016 -> ETP4HPC member)**

- HPC infrastructures & technologies,applications (material science), distributed data management, I/O performances, collaboration and remote access tools...
- Many different systems for testing and benchmarking

Running projects:

- **H2020 Centre of Excellence: EoCoE -** *http://www.eocoe.eu/*

  Energy oriented Centre of Excellence for computing application (Transversal basis, Meteorology, Materials,Water, Nuclear Fusion)
  - WP1/Task3 Efficient I/O (Juelich,PSNC,ENEA,CEA)
  - WP3 Materials (WP leader M.Celino, ENEA DTE-ICT)
- **Support tools** for several H2020 projects: SEADATACLOUD, NEXTOWER, M4F, INSPYRE, GEMMA... and to EERA-JPNM (Joint Programme on Nuclear Materials): e.g. Report Management

ENEAGRID/CRESCO 2018 People

http://www.eneagrid.enea.it/people/2018EneaGridPeople.html

Fiorenzo Ambrosino, Giuseppe Aprea, Tiziano Bastianelli, Riccardo Bertini, Irene Bellagamba, Giovanni Bracco, Luigi Bucci, Francesco Buonocore, Marco Caporicci, Michele Caiazzo, Beatrice Calosso, Massimo Celino, Marta Chinnici, Antonio Colavincenzo, Aniello Cucurullo, Pietro D'Angelo, Davide De Chiara, Matteo De Rosa, Daniele Di Mattia, Stefano Ferriani, Gianclaudio Ferro, Claudio Ferrelli, Agostino Funel, Dante Giammattei, Marcello Galli, Simone Giusepponi, Giuseppe Glorioso, Roberto Guadagni, Guido Guarnieri, Michele Gusso, Francesco Iannone, Massimo Marano, Angelo Mariano, Giorgio Mencuccini, Silvio Migliori, Marialuisa Mongelli, Patrizia Ornelli, Simonetta Pagnutti, Filippo Palombi, Salvatore Pecoraro, Antonio Perozziello, Samuele Pierattini, Salvatore Podda, Giovanni Ponti, Andrea Quintiliani, Giuseppe Santomauro, Alberto Scalise, Fabio Simoni, Daniele Visparelli.

Between them a special acknowledgement to the people mostly involved in the daily user support and system management:
Fiorenzo Ambrosino, Antonio Colavincenzo->Giuseppe Glorioso, Agostino Funel, Guido Guarnieri, Filippo Palombi, Giovanni Ponti.

ENEA HPC: current developments (1)

CINECA (Bologna) is the main HPC site in Italy, the national TIER0 centre for PRACE. CINECA is now #14 in Top500 2017/11 with **MARCONI** system, HPL 7.4 Pflops, Peak 15 Pflops. ~2 Pflops added 2018/Q1.

In the context of the Italian HPC ecosystem, an agreement was signed between ENEA and CINECA in 2015 to promote joint activities.

In this framework ENEA/CINECA answered and won an EUROFUSION call in 2015 for what has now become the **MARCONI-FUSION** section (5 Pflops) of the **MARCONI** system, in operation since 2016/Q3, helping CINECA in obtaining a system relevant at world scale. (new tender now in progress!)

The same ENEA/CINECA agreement triggered the procurement of **CRESCO6**, a TIER1 HPC cluster (~0.7 Pflops, HPL ~0.5 Pflops) in ENEA Portici. Cluster configuration & qualification are currently in progress.

Evolution of HPC resources in Portici:

- **2008 - CRESCO PON 2001-2006 : CRESCO2 -> 25 Tflops (#126 Top500)**
- **2013 - CRESCO4 PON 2007-2013: CRESCO4 -> 100 Tflops**
- **2018 - CRESCO6 - CINECA Collaboration: CRESCO6 -> 700 Tflops**

ENEA HPC: current developments (2) CINECA state now

ENEA
Italian National Agency for New Technologies, Energy and Sustainable Economic Development

CINECA SCAI
SuperComputing Applications and Innovation

MARCONI:New Tier-0 system

Technical Features:

- Fabric: Intel OmniPath
- Architecture: Lenovo NeXtScale

A2
KNL 68cores, 1.4 GHz;
3600 nodes, 11 PFs

A3+A4
SKL 2x24 cores 2.1 GHz
~2300 nodes >7 PFs

A1
BRD 2x18 cores 2.3 GHz
~720 nodes ~1 PFs

EUROfusion

2018/Q1 at Portici
a Tier1 system
~216 dual SKL,
~10000 core, ~0.7
PFlops

CRESCO6
~
CRESCO4
x 7

[F. Iannone]

ENEA HPC: current developments (3) CRESCO6 Layout

216 nodes - 5 racks

Lenovo Stark

chassis D2, fat-twin 2U

4 nodes SD530

SKL 8160 2.1 GHz

2x24 cores

192 GB RAM 4 GB/core

6 NSD servers

Space for additional nodes, 2 more racks

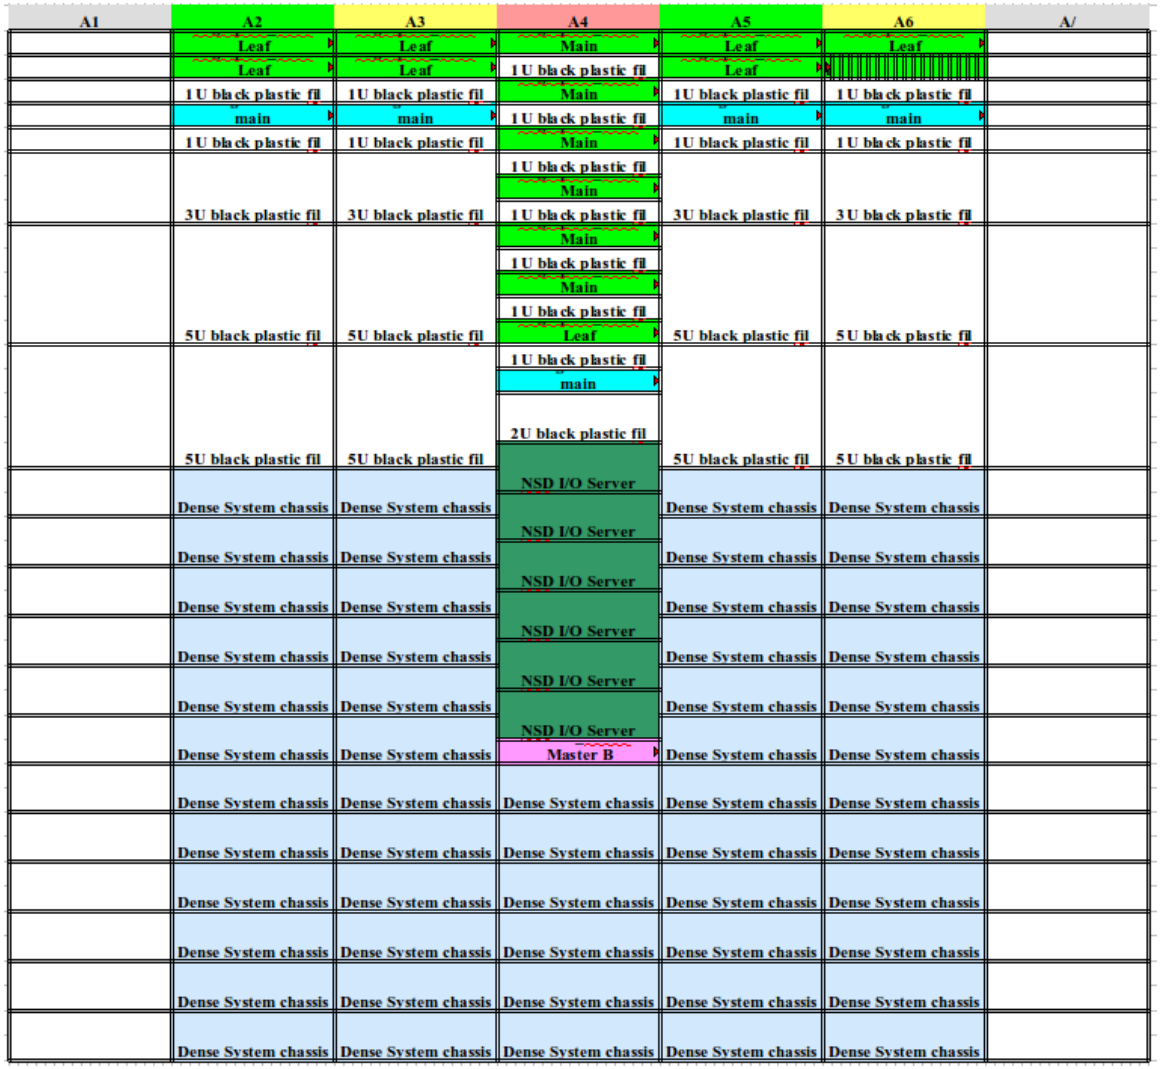

2016/2017 complete re-organization of Portici HPC site, new conditioning system, 2 computer rooms, separation beween clusters and data area.

Current developments (4): CRESCO6 OmniPath Fabric

INTEL OPA Fabric

48 ports switch 100Gbps bandwidth, port latency 100ns

[old QDR: 32 Gb/s, port latency 140ns]

216 node ports

6 I/O nodes

1 fabric manager

Possible expansion:

+72 node 1:2 fat-tree

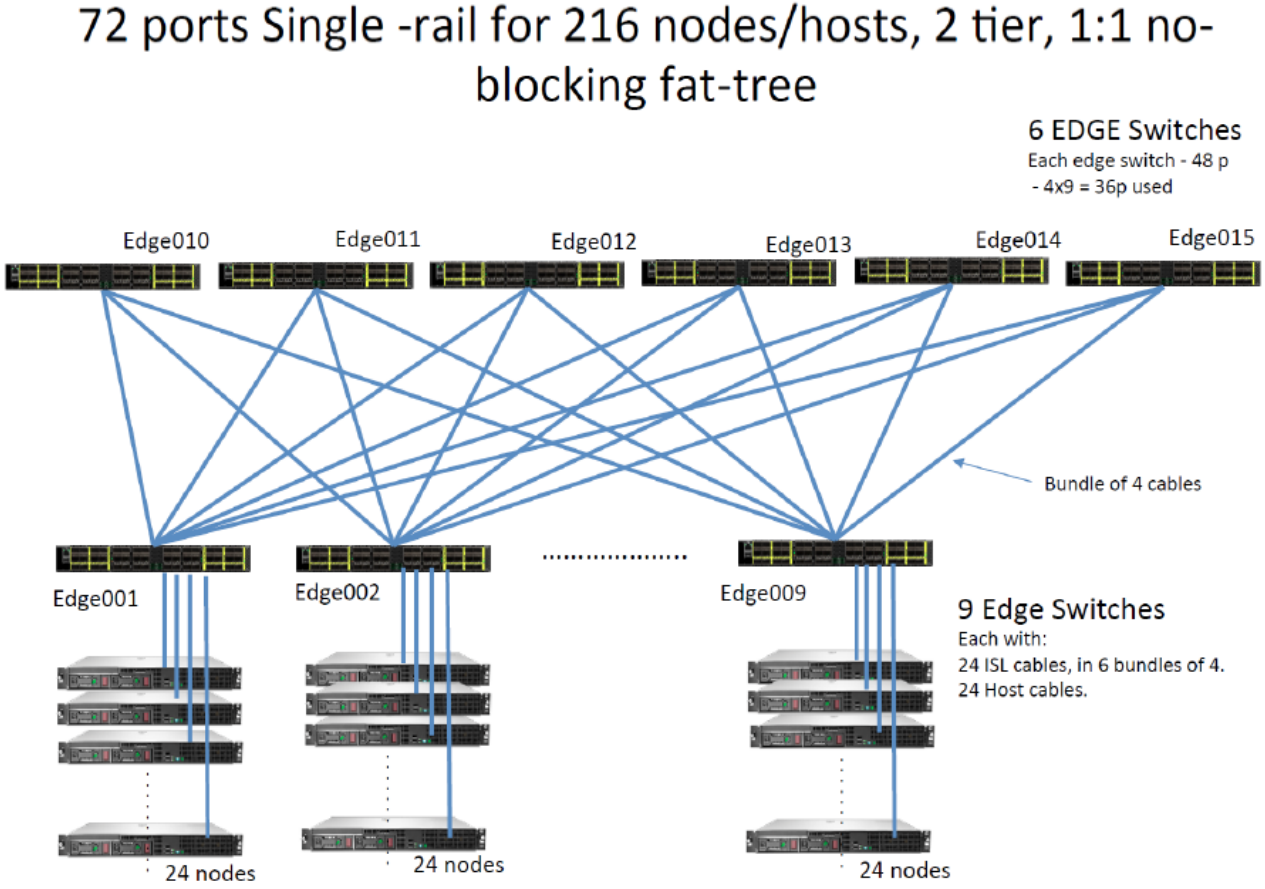

Current developments (5): CRESCO6 Storage Systems

GPFS/Spectrum Scale shared with the other clusters

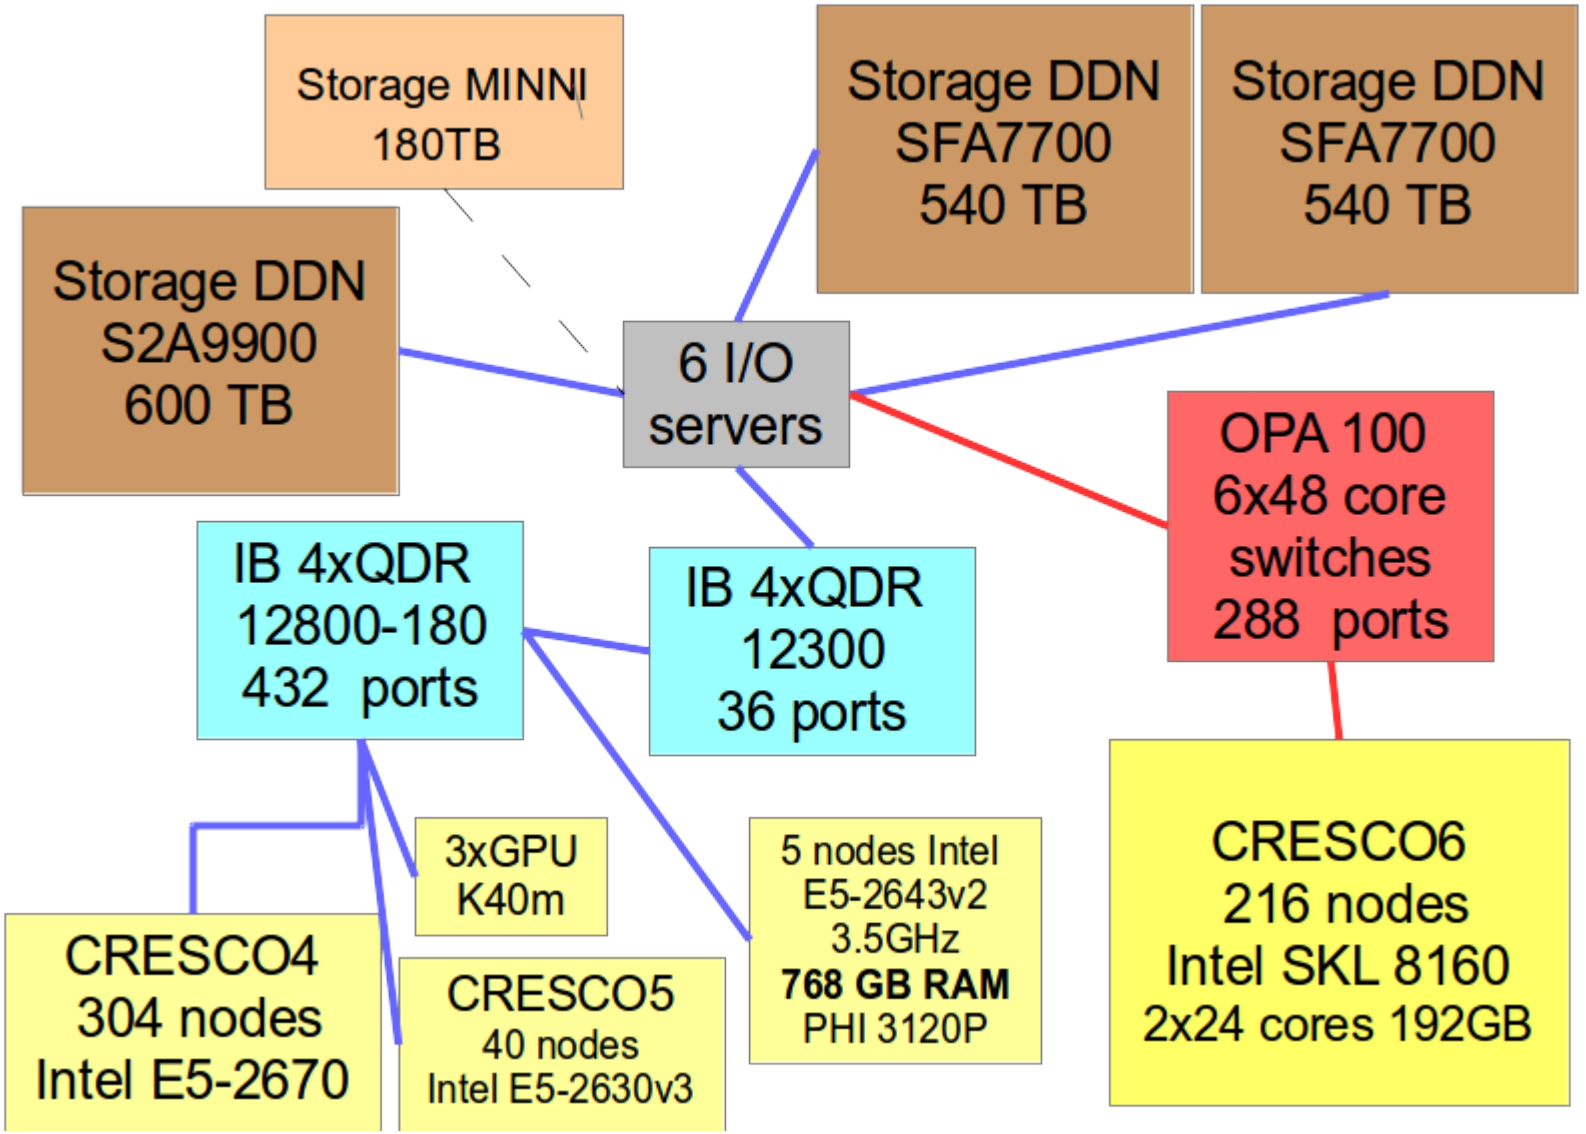

Current developments (6): CRESCO6 installation

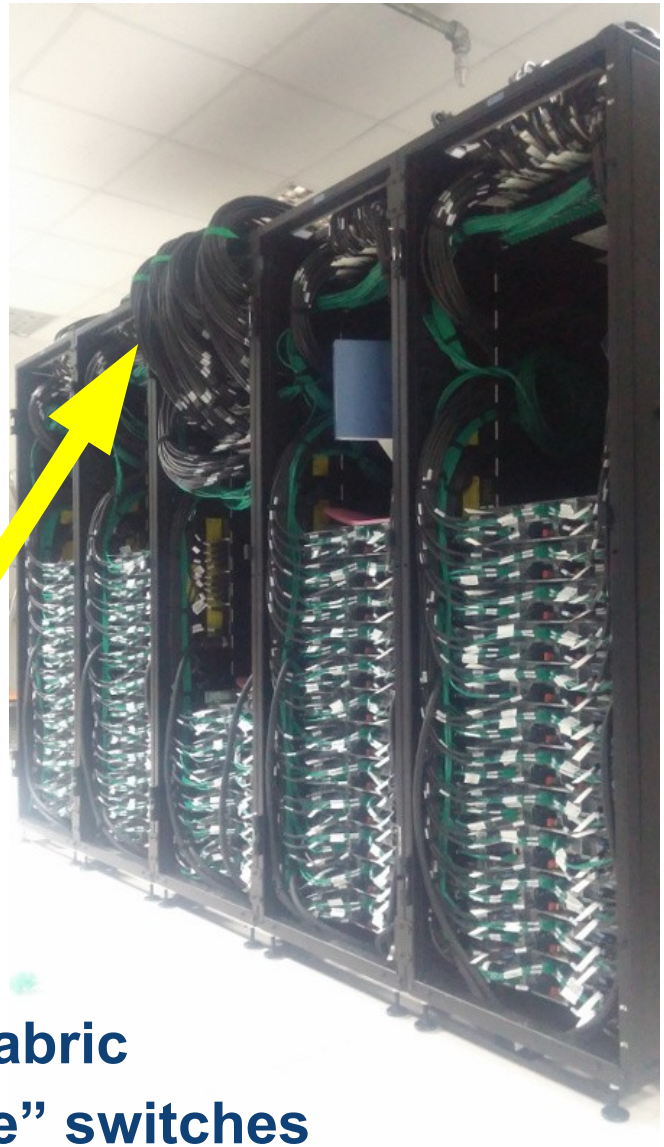

Installation of CRESCO6 cluster

smart-aisle cold air containment & grids still to be installed

OPA Fabric

6 “core” switches

Current developments (7): CRESCO6 now

New computer room cooling units

ready for free cooling using cold water from Portici site (to be implemented)

Installation of CRESCO6 cluster

smart-aisle cold air containment & grids

Current developments (8): cooling system

CRESCO clusters dry cooler system on the roof of the building, 27/2/2018

Current status and first results on CRESCO6

CRESCO6 cluster has been installed using confluent/xCAT by a Lenovo expert (David Rajendra, thanks for his help!) as a self-standing cluster with minimal network connection and without the integration into the ENEAGRID environment.

The integration in ENEAGRID is currently in progress. Some results have been obtained on HPL (max eff. 63%) and also HPCG (efficiency 0.73 %) [MARCONI reference data have been obtained in the normal production setup]

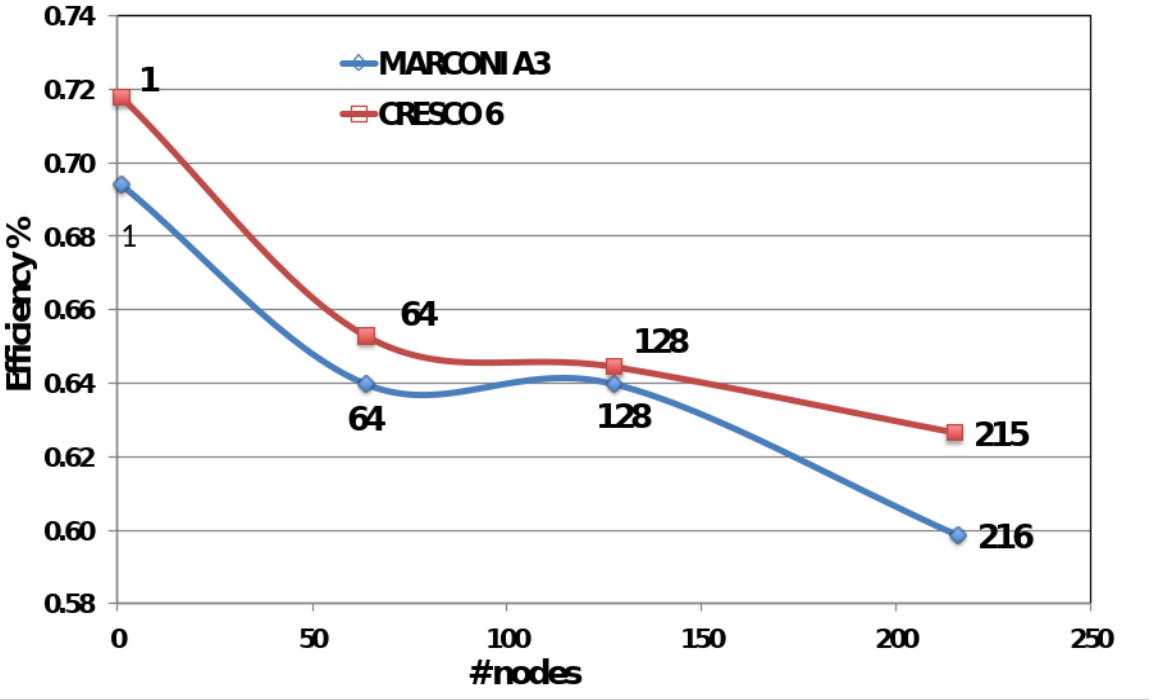

Status and first results on CRESCO6 (2)

Bandwidth benchmark OSU - OmniPath versus TrueScale (QDR)

- max bandwidth 100 Gbps (CRESCO4 QDR 32 Gbps)
- min latency 1μs => 1.4μs (CRESCO4 QDR 1.4μs =>1.8μs)

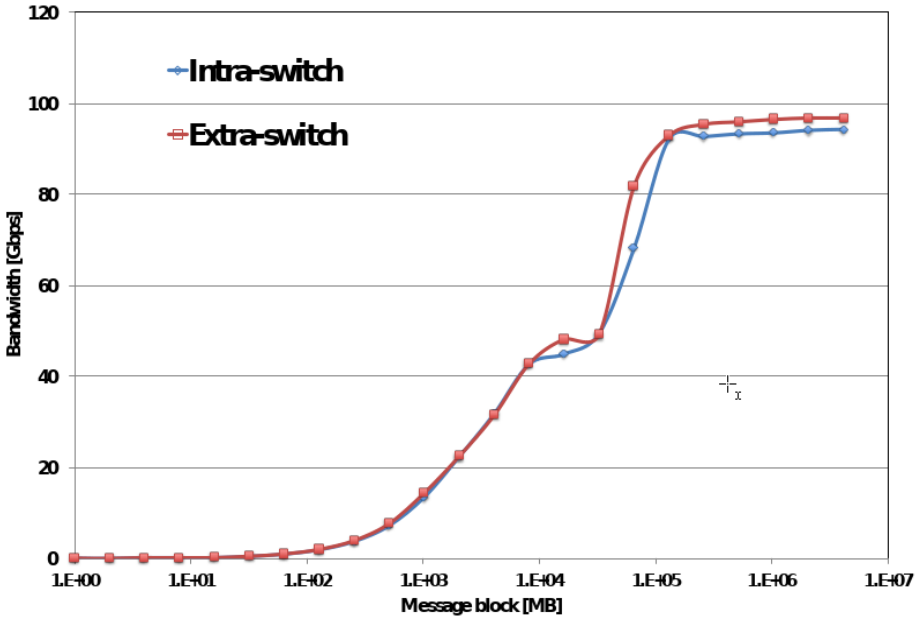

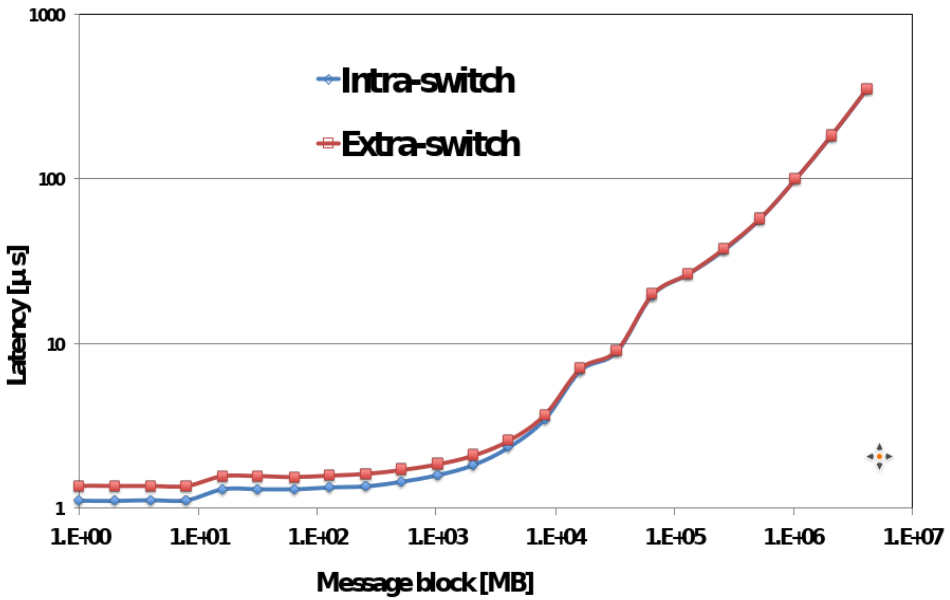

8th CRESCO Annual Report, 2016 - 45 papers

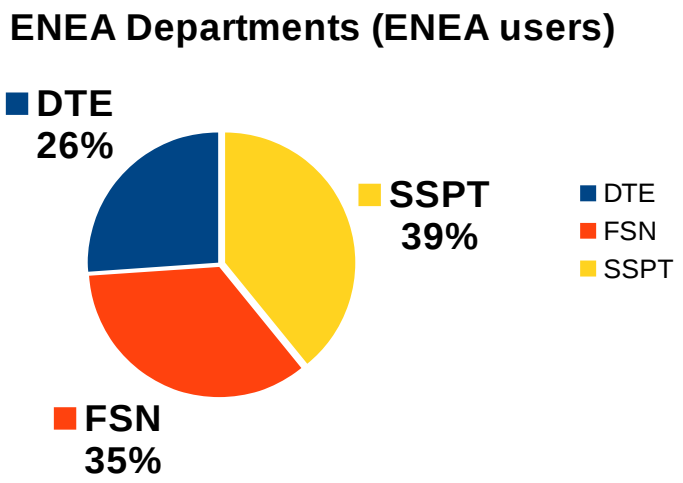

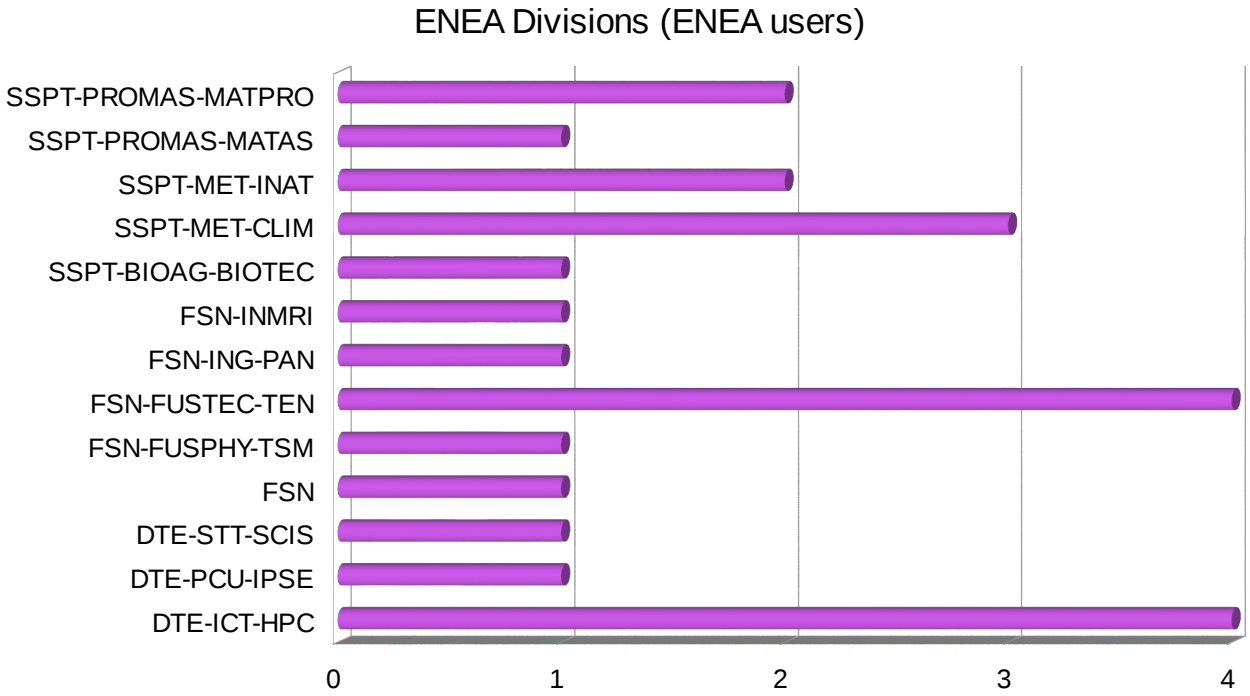

42 contact persons
corresponding to 72 HPC users;

23 ENEA , 19 external
contributors

External user affiliation

UniSA	2
Uni Madrid	1
UniLaSapienza	5
ISPRA	1
Consorzio RFX	1
UniCamerino	1
UniTorVergata	1
Vinča Institute Belgrade	1
CREATE Consortium	1
UniRC	1
UniFI	1
INM-Germany(Julich)	1
UniNA	1

Annual Report 2016: Research & Development Domain

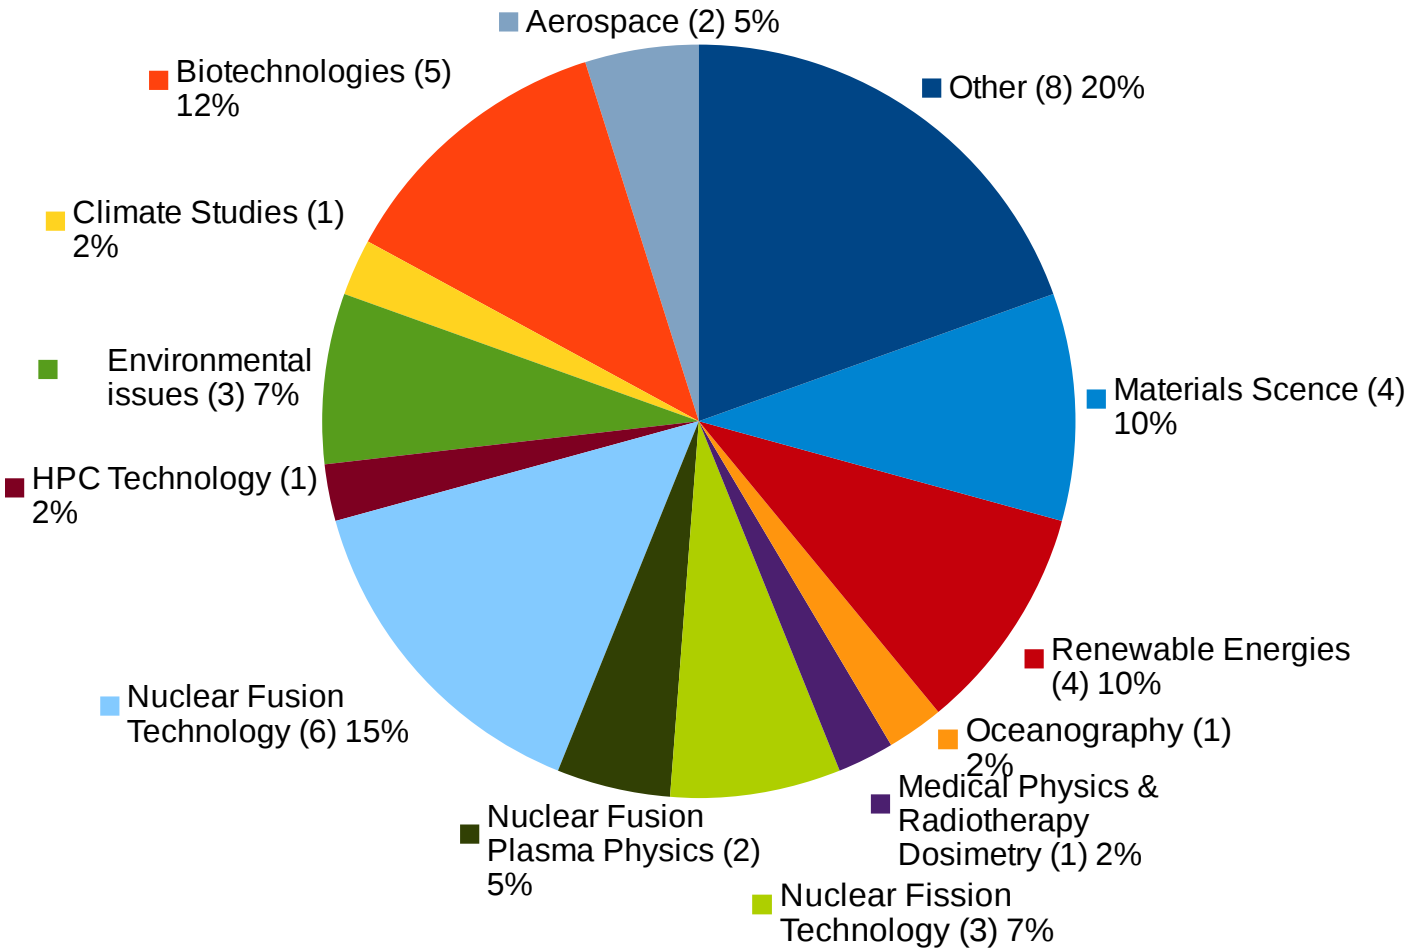

Other(8)
Interaction waves structures (1)
Atmospheric sciences (1)
Network Science for Complex Syst (1)
Radiologically characterization (1)
In silico drug design (1)
Industrial flows with heat and mass transfer (1)
Complex Systems (1)
Solid State Physics(1)

Accounting 2017

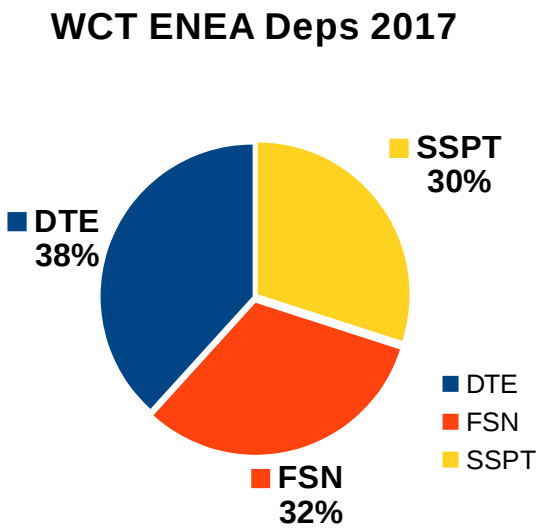

ENEA Deps	wct_y
DTE	1501,23
FSN	1231,22
SSPT	1162,45
tot	**3894,9**

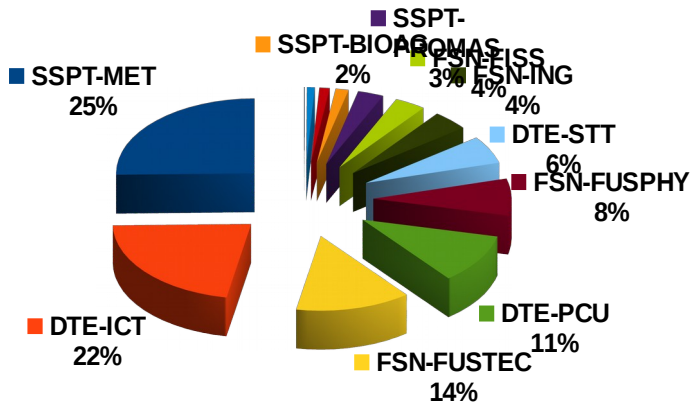

ENEA Divisions	wct_y
SSPT-MET	985,409
DTE-ICT	850,927
FSN-FUSTEC	525,896
DTE-PCU	417,496
FSN-FUSPHY	311,467
DTE-STT	232,803
FSN-ING	166,619
FSN-FISS	142,129
SSPT-PROMAS	118,007
SSPT-BIOAG	59,0303
FSN	46,8156
FSN-SICNUC	35,5117
FSN-INMRI	2,77281
Divisions with	< 2
tot	**3894,9**

Inside CRESCO6 during the installation phase