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

Evolution of ENEAGRID/CRESCO HPC Infrastructure and Big Data analysis opportunities

Workshop "Big Open Data Analysis" – ENEA Roma Sede 2/2/2018

G. Bracco
giovanni.bracco@enea.it

HPC a Lab in ENEA DTE-ICT Division

Information and Communication Technology Division
Services & HPC, (~75 staff) HPC (~20 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**

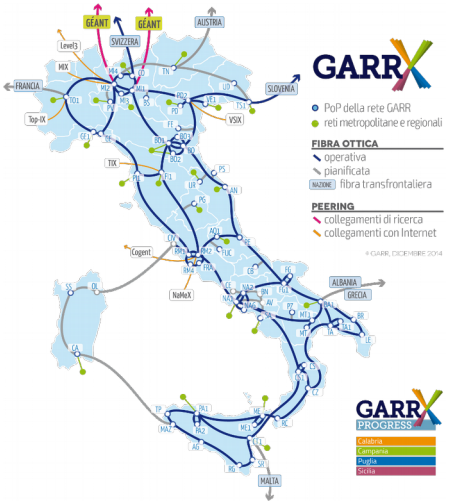

12 main sites connected by GARR, the Italian NREN, 6 Research Centres with HPC/storage resources

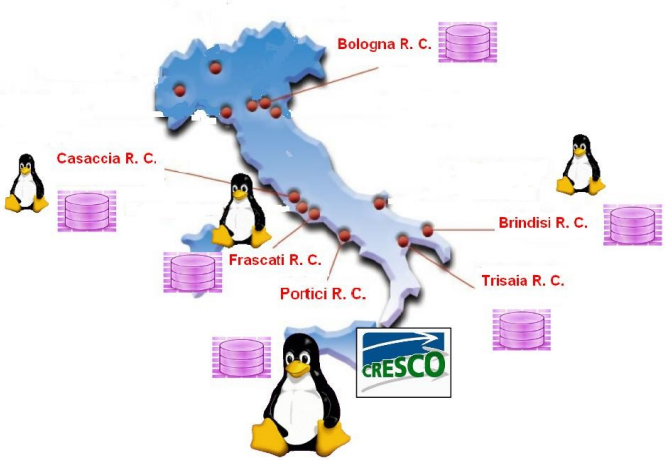

ENEA HPC Infrastructure: ENEAGRID/CRESCO Clusters

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

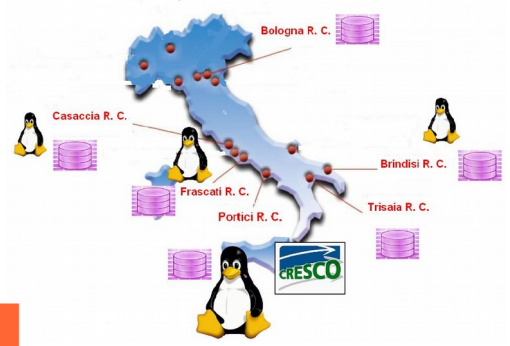

Cluster name (Linux x86_64)	Network	Cores/Tflops
CRESCO3 AMD	IB QDR	1152/~10
CRESCO4 INTEL	IB QDR	4864/~100
CRESCO5 INTELv3	IB QDR	672/~25
CRESCOF AMD	IB QDR	456/~4
CRESCOC AMD	IB DDR	192/~1.5
K40/PHI/Large RAM	IB QDR	300/~17
Total in production		**~7500/~160**
Q1/18 CRESCO6	**SKL OPA**	**~10000/~700**

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, 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, Agostino Funel, Guido Guarnieri, Filippo Palombi, Giovanni Ponti.

ENEA Main HPC Site: Portici

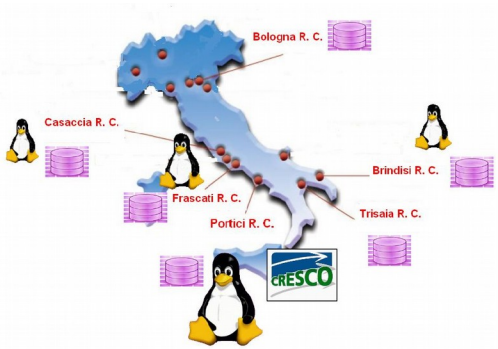

ENEA HPC: current developments (1)

CINECA (Bologna) is the main HPC site in Italy, being 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 includes the installation of a TIER1 HPC cluster (~0.7 Pflops, HPL ~0.5 Pflops) in ENEA Portici site in 2018/Q1. The cluster is now in Portici in a upgraded computer room and installation, configuration & qualification are currentely in progress.

- **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

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

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

A3+A4
SKL 2x24 cores 2.1 GHz
~2300 nodes >7 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 in 5 racks
Lenovo Stark
fat-twin 2U, 4 nodes
SKL 8160 2.1 GHz
2x24 cores
192 GB RAM
4 GB/core
6 NSD servers

Space for additional nodes, 2 empty racks

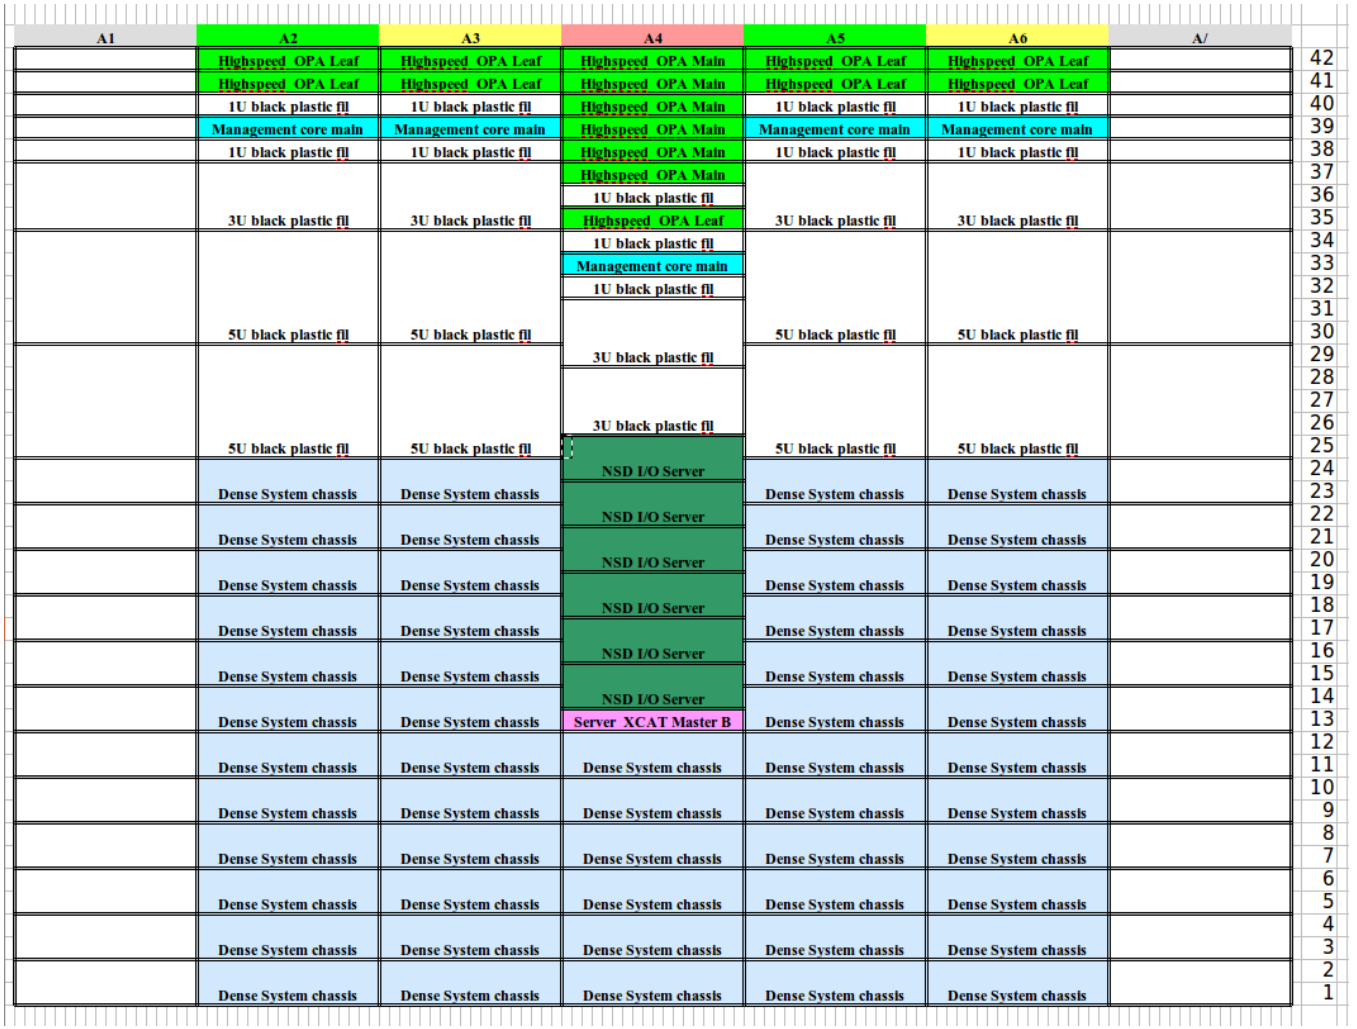

2016/2017 complete reorganization of Portici HPC site, new conditioning system, 2 computer rooms, separation beween clusters and data area.

Current developments (4): CRESCO6 OmniPath Fabric

OPA Fabric

48 ports switch 100Gb/s bandwidth, port latency 100ns

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

216 node ports

6 I/O nodes

1 subnet manager

Expansion:

72 node

1:2 fat-tree

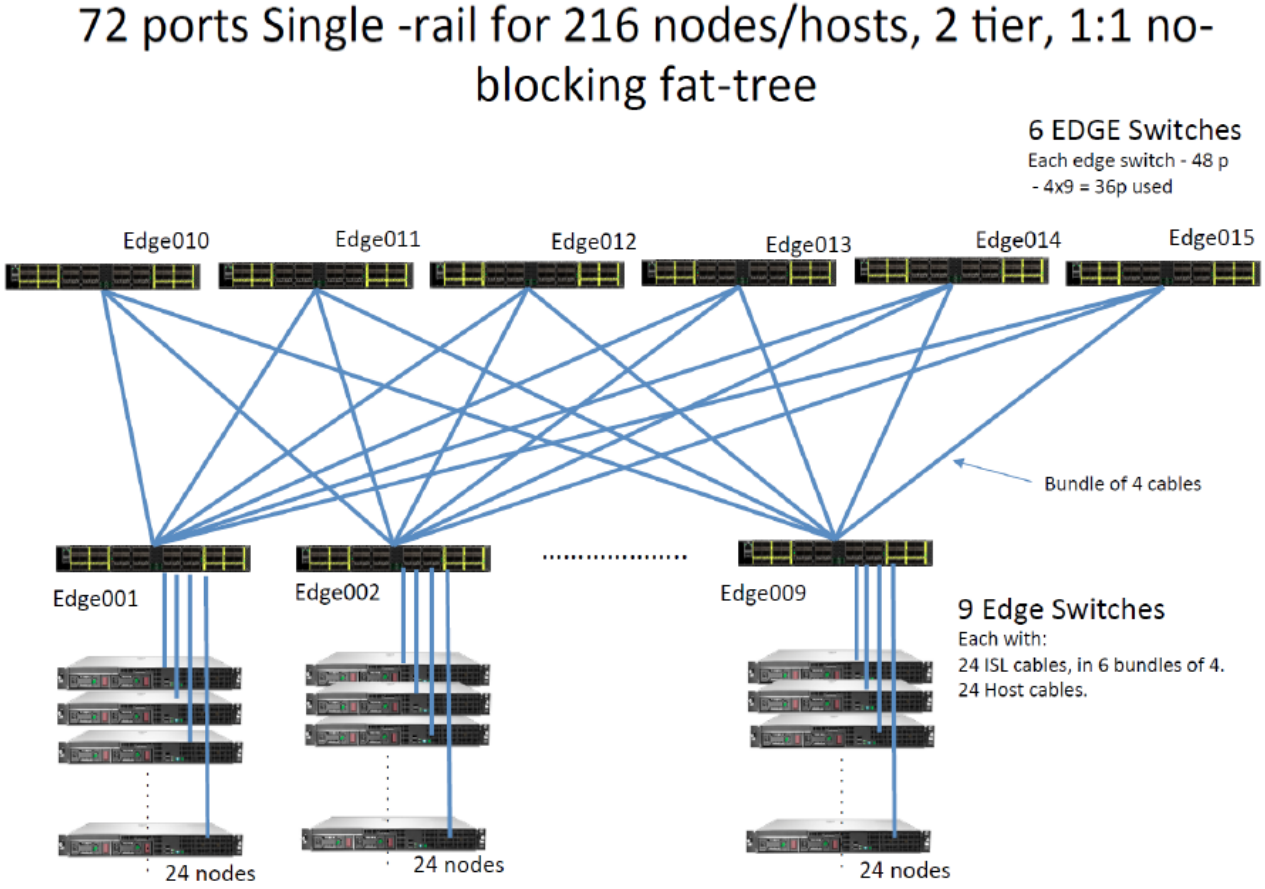

Current developments (5): CRESCO6 2018/01/28

Installation of CRESCO6 cluster

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

OPA Fabric

6 “core” switches

ENEAGRID DATA & BIGDATA: storage resources

While ENEA has a long HPC experience and the performances of ENEA clusters in the last 10 years are relevant at national level, the data and storage management resources have not attained up to now the same level of relevance..

The situation is now changing due to the arrival of the Big Data paradigm and also to the request of the applications based on HPC technology.

Data & Storage Resources available in ENEAGRID/CRESCO

- **AFS distributed filesystem for user homes & projects (~30 TB)**
- **GPFS(Spectrum Scale) parallel file system (multisite, mostly in Portici 1.8 PB)**
- **Big RAM nodes at Portici (5 nodes with 768 GB RAM, 2 nodes with K40 GPU)**
- **DDN storage in Portici (~1.7 PB raw)**
- Lower performance storage for HPC (~500 TB raw)
- Cloud dedicated storage, distributed in 7 centres (~1 PB raw)
- Tape Library (Portici 2.5 PB)

ENEAGRID DATA & BIGDATA Sources

The availability of significant amount of computed results can be the basis for the application of the Big Data, Deep learning techniques

Examples of the data available in ENEAGRID/CRESCO storage systems:

- Climate/energy related forecasts (~100TB)
- **Air pollution models and air quality evaluation (~200 TB, 7x10^6 files)** collab. with ENEA SSPT Dept
- **Web crawling (~15 TB, 340x10^6 web pages):** see next D'Agostino presentation in the Workshop
- Nuclear fusion experiments
- HPC Cluster monitoring system (MySQL, 9 GB 110x10^6 tuples, loop recording) ZABBIX
- Bioinformatics
- ....

BD&Deep Learning: an example from environmental science

“A Machine Learning Framework to Forecast Wave Conditions”
Scott C. James, Yushan Zhang, Fearghal O'Donncha, 2017 https://arxiv.org/abs/1709.08725

- Monterey Bay, California, a deep-learning framework trained to forecast wave conditions A physics-based **Simulating WAves Nearshore (SWAN)** model generated training data for the deep learning network. The model driven by measured wave conditions, ocean currents from an operational forecasting system, and wind data was run between April 1st, 2013 and July 31st, 2017 generating forecasts at three hour intervals to provide a total of **12,400 distinct model outputs**.
- A deep learning network was trained and results were compared to observed buoy wave data within the model domain with **comparable levels of accuracy** to the traditional physics-based, SWAN model and a computing time reduction of a factor 12000.

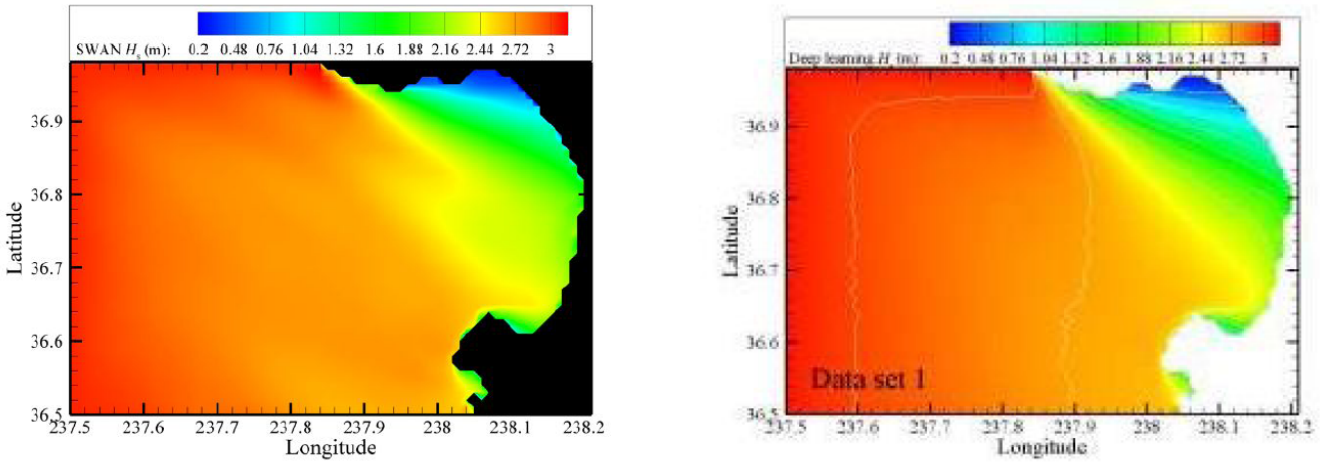

12,000% faster

Big Data/Deep Learning: air quality context

Air quality models and deep learning

- The AMS (Atmospheric Modelling System) - MINNI is the reference modelling system of Italian Ministry of Environment, developed in ENEA (SSPT Dept.) with the aim to support air quality policy at national and regional level. Large amount of experimental and simulation data in periods starting from 1990 are available.
- the system is based on a complex modelling, including atmospheric dynamics and chemical species evolution, and a large base of computed results have been produced and stored on ENEAGRID/CRESCO data storage systems.
- **A collaboration activity to apply deep learning techniques to the available modeling results has just started**, targeting the estimate of the simulation precision and the identification of anomalous situations (collab. with M.Adani,M.D'Isidoro et al.)

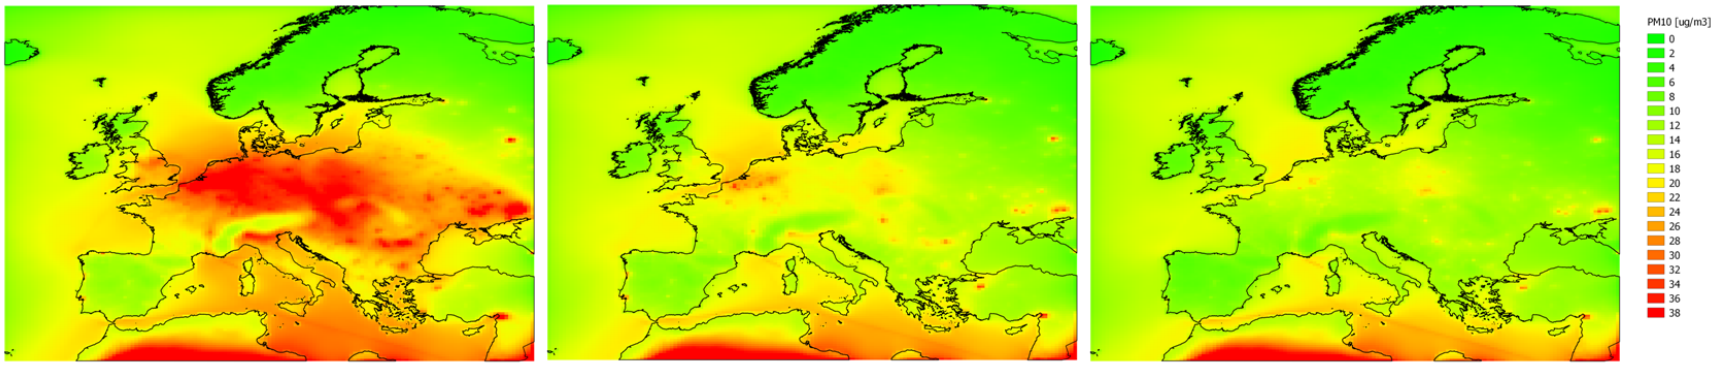

PM10, 1990, 2000, 2010 Briganti et al.

BigData: Web Crawling and Web Data Analysis in ENEAGRID

Setting up in ENEAGRID an environment for web crawling activity

Web Data Retrieving: tools for systematically and automatically exploring a network content in order to create web data archives (***BUbiNG***);

Web Data Analysis: tools to storing, indexing, displaying and clustering the web data (***Solr***, ***Wayback***, ***Carrot²***).

Access to GUI from FARO2 menu for authorized users

The operation must be notified and discussed with network people!

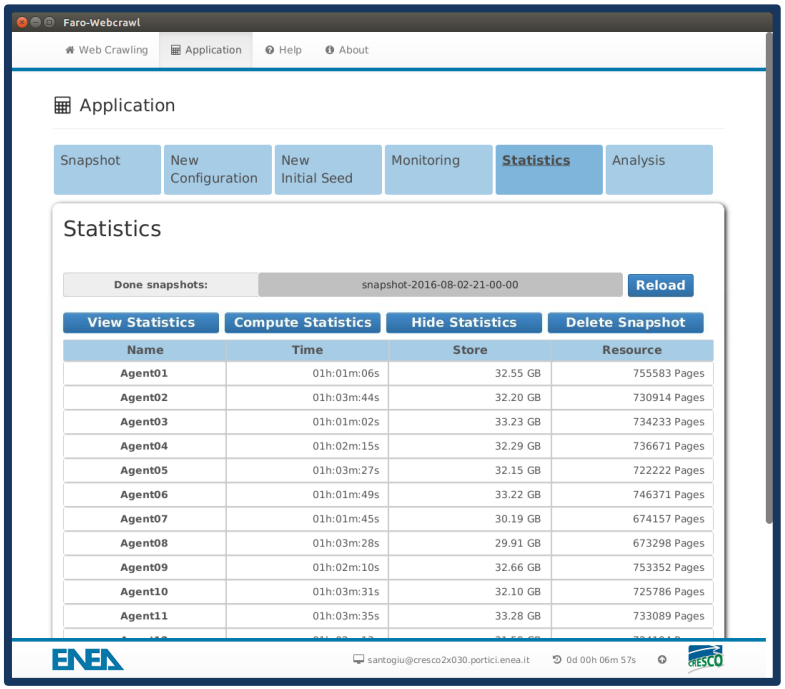

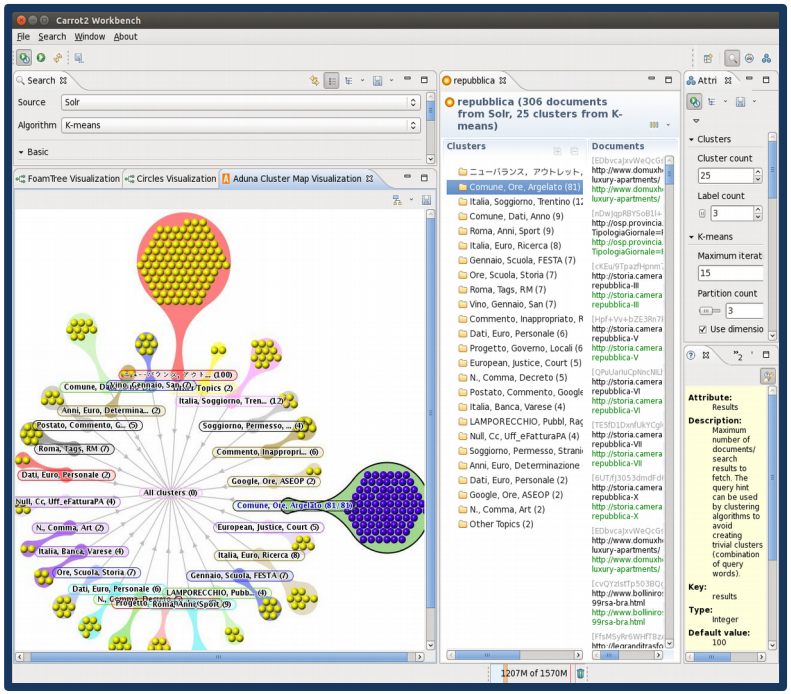

Application contexts: Web Searching; Intelligence & security; Blog analysis; User behaviors; Marketing; Critical Infrastructure.

Big Data: Web Crawling test

Efficiency & Robustness

Number of agents: 16;

Running time: 8h;

Amount data: 2.94TB (66M Pages);

Downloading speed: 850Mbps (2305 Pps)

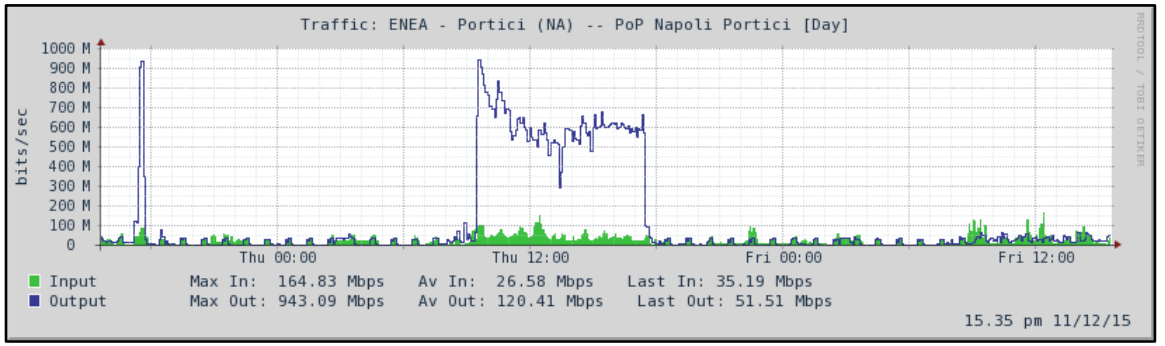

Single web crawling session, kept alive for **8 one hour**, by considering only web pages from *.it* domain.

Reliability

Downloaded data:
15 TB (484 GB/snapshot);

Saved on the storage:
3,3 TB (111 GB/snapshot);

Downloading Speeds:
μ: 1,00 Gbps, σ: 0,0005 Gbps.

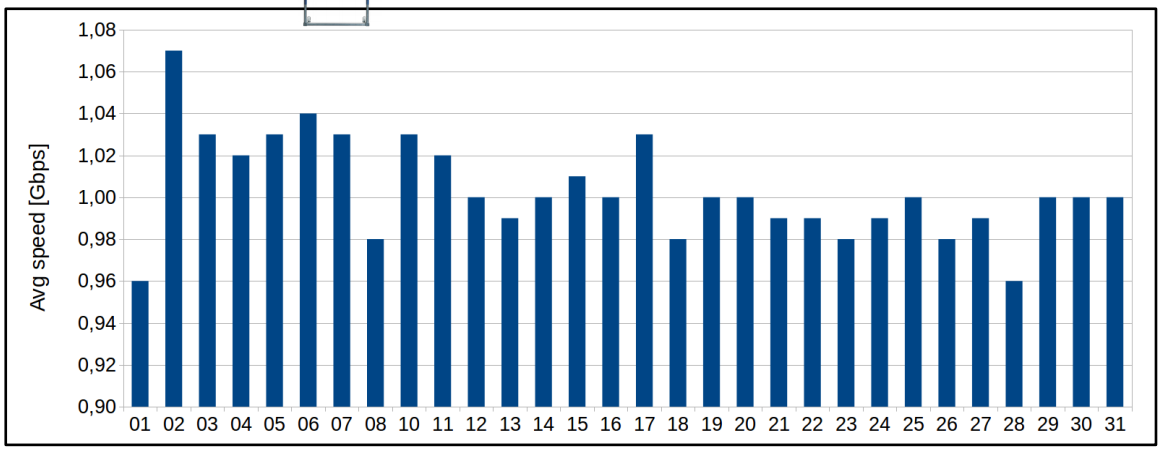

Daily web crawling sessions, during **one month**, each of them kept alive for **one hour** (from 21:00 to 22:00), by considering only web pages from *.it* domain. (by courtesy of G.Santomauro)

ENEA HPC and Big Data Analysis Opportunities

Conclusions

The outcomes of ENEAGRID/CRESCO computional/storage resources and related development and research activities:

- ENEA & collaborators have access to a state-of-art HPC resources of a size relevant at national level. The new CRESCO6 cluster will be soon in production on Portici site
- in Big Data/Deep Learning context new tools are available and the amount of existing modelling data can provide an opportunity for new research approaches, by applying state-of-art analysis techniques.

Thank you for your attention!