

# **HPC at CNAF** **LASER PLASMA ACCELERATION**

**WHEN**  
**24<sup>th</sup> June, 2014**  
**14.00 - 20.00**

**WHERE**  
**CNAF Meeting room**  
 Viale Berti Pichat 6/2, Bologna (first floor)

**SPEAKERS** • G. Turchetti (UniBO/INFN) • G. Maron (CNAF/INFN) •  
 G. Arduini (CERN) • A. Bazzani (UniBO/INFN) • M. Giovannozzi  
 (CERN) • G. Franchetti (GSI) • A. Franchi (ESRF) •  
 L. Gizzi (CNR/INO) • G. Sarri (QUB) • F. Boscherini (UniBO) •  
 A. Sgattoni (CNR/INO) • S. Sinigardi (UniBO/INFN) •  
 G. Benedetti (LBNL) • F. Rossi (UniBO/INFN) •  
 D. Cesini (CNAF/INFN) • F. Giacomini (CNAF/INFN)

**#HPCATCNAF**

[tiny.cc/hpcatcnaf](http://tiny.cc/hpcatcnaf)







- Beams and plasma dynamics
- Computing activities at CNAF
- HPC infrastructure at CNAF: status and perspectives
- Support infrastructures to software development projects
- Beam dynamics studies for the LHC and Injector upgrade; some examples
- Multipolar non-linearities and correction strategies
- Transport and space charge studies with MICROMAPS
- Transport and space charge studies with HALODYN
- Laser acceleration of electrons and Thomson X-rays at CNR/INO
- Generation of ultra-bright, multi-MeV gamma-ray beams via non-linear Thomson scattering
- Time resolved X-ray spectroscopies
- Rayleigh-Taylor Instability in high energy gain radiation pressure ion acceleration
- Prepulse and preplasma studies for a proton acceleration record experiment: 2D analysis
- Electrons acceleration at Berkeley: simulations and experiments
- Laser acceleration of protons and electrons: simulations on GPUs at the HPC@CNAF facility

- CNAF hosts the Italian Tier1 computing centre for the LHC experiments ATLAS, CMS, ALICE and LHCb....
- ... but also one of the main Italian processing facilities for several other experiments:



a Payload for Antimatter Matter Exploration and Light-nuclei Astrophysics




- CDF, SuperB, KLOE, LHCf, NA62
- AMS, ARGO, AUGER, FERMI, MAGIC, PAMELA, Icarus, Xenon100, Borexino, Gerda, CTA, Opera, Darkside, Cuore, Virgo
- Involved in a number of **Grid** and **Cloud** national and international projects
  - WLCG
  - EGI-Inspire
  - MCLOUD
  - Open City Platform
  - !CHAOS
  - many other in preparation towards H2020

→ collaboration with other disciplines


# INFN-T1 in Numbers



	CPU (HS06)	Disk (TB-N)	Tape (TB)
<b>2014</b>	150000	15000	25000
<b>2015</b>	205000	16000	35000



For the LHC VOs is the second WLCG computing centre, after CERN, in terms of CPU time



- In 2014 CNAF started to operate an HPC cluster but remains mainly an High Throughput Center
  - In the short term is not going to become a supecomputing center
- The HPC cluster was created as a collaboration between
  - Physics Dept. Bologna University
  - INFN-Bologna
  - INFN-CNAF
  - INFN-COKA project
- To support local HPC users and their close collaborators
- To offer an HPC testbed for developers before moving to other facilities that operates at a greater scale
- To acquire expertise at CNAF in a going-parallel world

# **HPC Cluster at CNAF**

## **Status and Perspectives**

*Daniele Cesini - INFN-CNAF*

# People and Collaboration




- **The HPC cluster at CNAF was created thanks to a collaboration between:**
  - Physics Dept. Bologna University
  - INFN-Bologna
  - INFN-CNAF
  - INFN-COKA project
- **Many operative contributions and support from:**
  - INFN-BO: Vincenzo Vagnoni (ib, storage, gpfs and main screwdriver interventions 😊)
  - UNIBO: Stefano Sinigardi and Francesco Rossi (gpu, ib, cabling, main testers)
  - CNAF Farming: Stefano Dal Pra (lsf), Massimo Donatelli (ldap and user account)
  - CNAF Storage: Vladimir Sapunenko (gpfs)
  - CNAF Network: Stefano Zani, Lorenzo Chiarelli (switch and uplink)
  - CNAF Infrastructure: Michele Onofri, Andrea Ferraro (power and cooling)
  - CNAF R&D: Matteo Manzali (mic)

# The Cluster Status

- 13 Worker Nodes
  - CPU: 392 + 24 HT cores
    - 320 HT cores E5-2640
    - 48 HT cores X5650
    - 24 HT cores E5-2620
    - 24 HT cores E5-2620
  - 15 GPUs:
    - 8 Tesla K40
    - 5 + 2 Tesla K20
    - 2x(4GRID K1)
  - 2 MICs:
    - 2 x Xeon Phi 5100

	CPU	GPU	MIC	TOT
TFLOPS (DP - PEAK)	3.2	19.2	2.0	24.4

- 2 disks server
- 60 TB shared disk space
- 4 TB shared home
- 1 Infiniband QLOGIC switch 18 ports
- 1 Eth switch BROCADE 48x1Gb/s + 8x10Gb/s

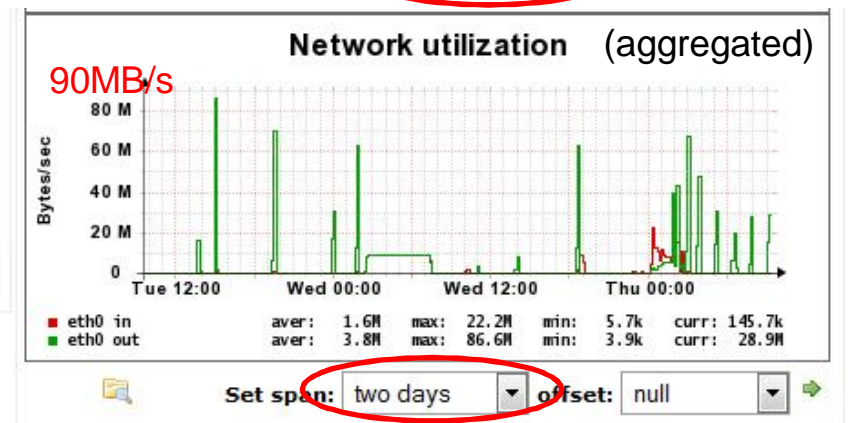

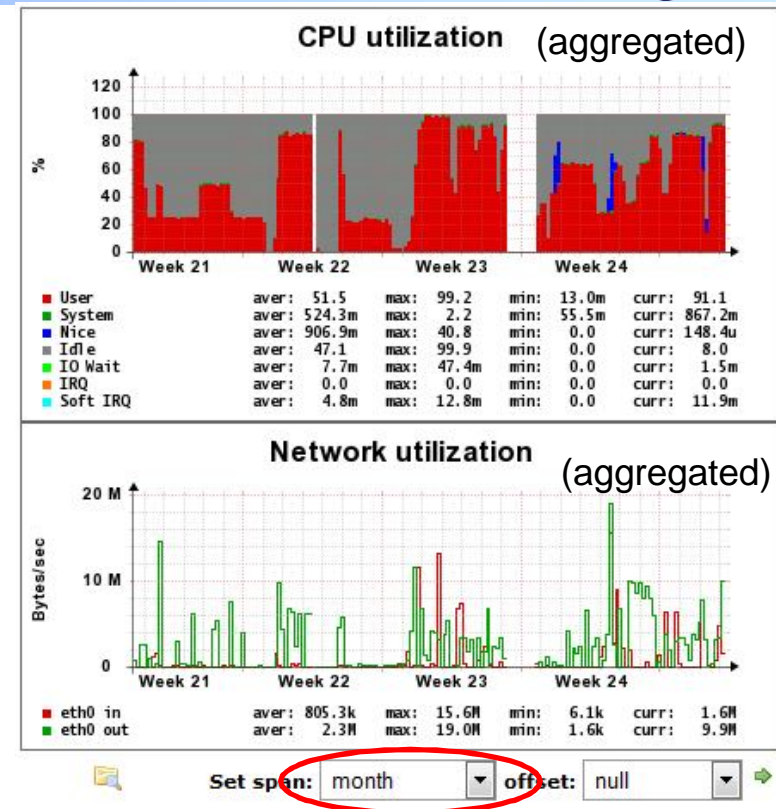


- LSF9.1 to access the cluster
- GPFS for the shared file systems
- Compilers
  - CUDA5.5, CUDA6.0
  - GCC4.6 (system)
    - Compiled by us 4.7/4.8 – working on 4.9
  - icc, ifort
  - Python 2.7/3.3, ipython
  - Openmpi-1.5.4
  - cmake, cmake28, gdb
- Others libs
  - HDF5, FFTW, BLAS, CUBLAS, GSL
- Tools
  - Virtualenv, matplotlib, git, hg, paraview, numpy, scipy, pylab, gnuplot



# Usage Statistics

- LSF
  - 2 batch queues (short, infinite)
  - 1 interactive
  - Flat configuration
    - no priorities, no shares
  - About 1k jobs (last 2 months)
  - 26 registered users
    - 13 active (last month)



# Open Issues



- Improve **availability** and storage performance
  - Fix the redundant disk controller problems (this week, no down)
  - Upgrade the storage-switch link to 2x10Gb/s
    - Study the possibility to use infiniband also for storage data
  - Add **monitoring** and **alarm** systems
- **Accounting** for CPU and GPU/MIC usage
- Dedicated **GPUs queue** with reserved cores
- Fix the Infiniband problems on 3 nodes
  - Add a new 24 core 2xTeslaK20 machine
    - Need the IB adapter
  - Change the IB adapters of the two X5650 servers (24x2 core)

# Cluster Upgrade



- The physical infrastructure is not the limiting factor....
- ...it depends on budget
- For further upgrades in the short term:
  - 4 working ports are still available on the Infiniband switch
  - 6x3 new ports can be enabled with just a license upgrade:  
3x480 euro+VAT

# Links and Contacts

- <http://wiki.infn.it/strutture/cnaf/clusterhpc/home>
- <https://lemon.cr.cnaf.infn.it/lemon-web/info.php?entity=hpc>
- `hpc-support<_at>lists.cnaf.infn.it`
- `hpc-users<_at>lists.cnaf.infn.it`

