High Performance Computing

Marco Moro
CRS4 HPC Technologist
marco.moro@crs4.it

XIII Seminar on Software for Nuclear Subnuclear and Applied Physics – Alghero 5-10 June 2016

Agenda

- CRS4 Center for Advanced Studies, Research and Development in Sardinia
- HPC systems: an introduction
- HPC infrastructure at CRS4: some numbers
- *Clustering* resources (scheduler,system software,application software, libraries, compilers)
- Management monitoring maintenance
- HPC research at CRS4

About CRS4

CRS4, Center for Advanced Studies, Research and Development in Sardinia, is an interdisciplinary research center that **promotes the study, development and application of innovative solutions** to issues in natural, industrial and social environments.
These developments and solutions are based on **Computational Science, Information Technology and High-Performance Computing**

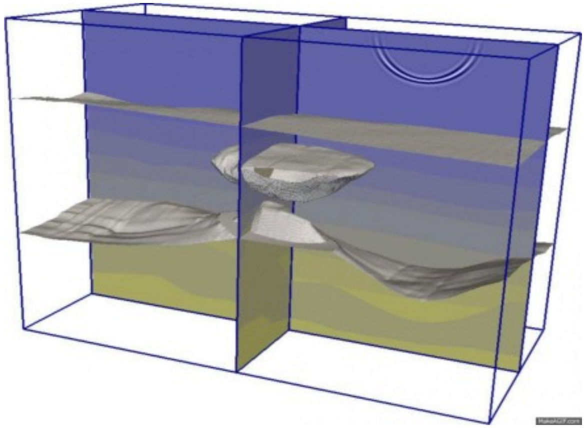

Since 1991, year of foundation, CRS4 cooperates with academic, scientific and industrial entities, **participating in relevant national and international projects**.

The Center is located at the Science and Technology Park of Sardinia in Pula Cagliari - Italy

HPC systems: an introduction

High Performance Computing (HPC) is the method by which scientists and engineers **solve** complex problems using apps that require high **bandwidth**, low **latency** networking and high **computing capabilities**.

Thanks to efficient numerical methods, HPC can solve extremely detailed mathematical and physical models within **reasonable timeframes**.

Global Climate

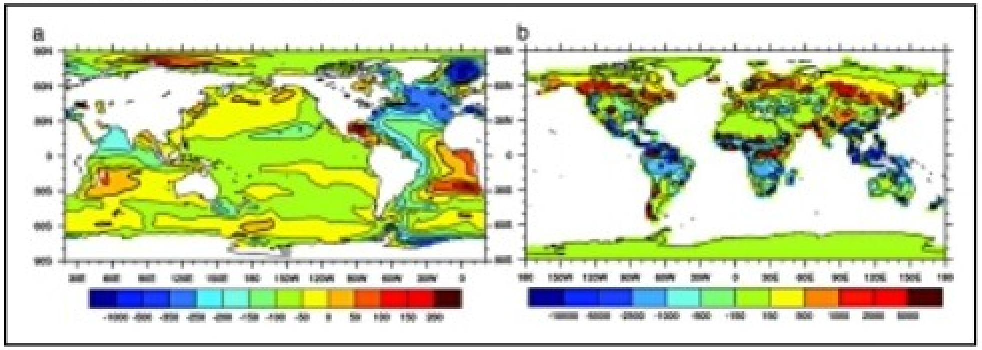

Biology

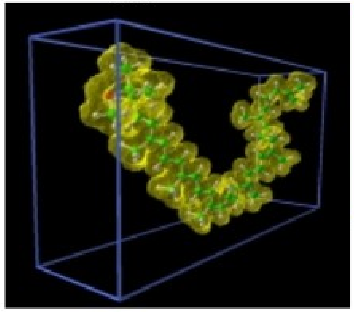

Nanoscience

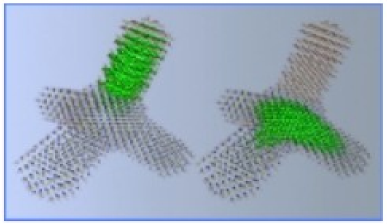

Combustion

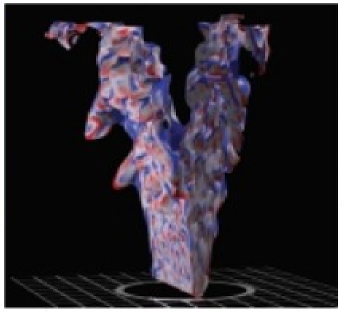

Astrophysics

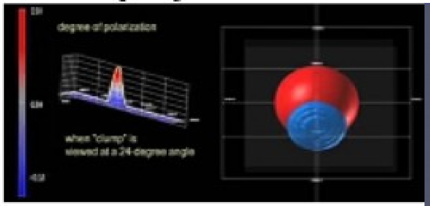

CFD

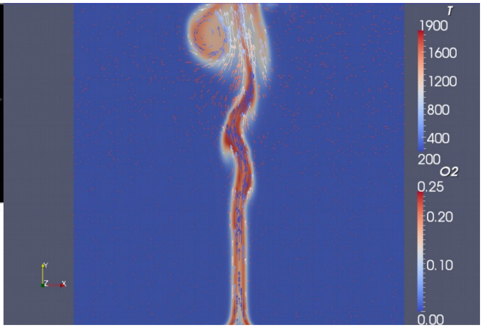

All CRS4 High Performance Computing resources are **Linux Computer Clusters**

A computer cluster is composed by single (quite simple) computers networked into a local fast area network with libraries and programs installed which allow processing to be shared among them

The result is an **high-performance parallel computing cluster** from “standard” computer hardware

Dictionary

- **Bandwith (data transfer rate)**

  the amount of data that can be carried from a point to another in a given time. Usually expressed in bytes per second (B/sec) for memory systems (e.g. hard disk, RAM), or in bit per second (bps) for network communication

- **Network Latency**

  time from the source sending a packet to the destination receiving it plus the on-way time from the destination back to the source

- **Compute Performance**

  Expressed in FLOPS – FLoating point Opearations Per Second

Examples @ CRS4:

10 Gigabit Ethernet network

- Bandwith: 10 Gigabps
- latency: 40 us

Infiniband FDR network

- Bandwith: 56 Gigabps
- latency: 5-10 us

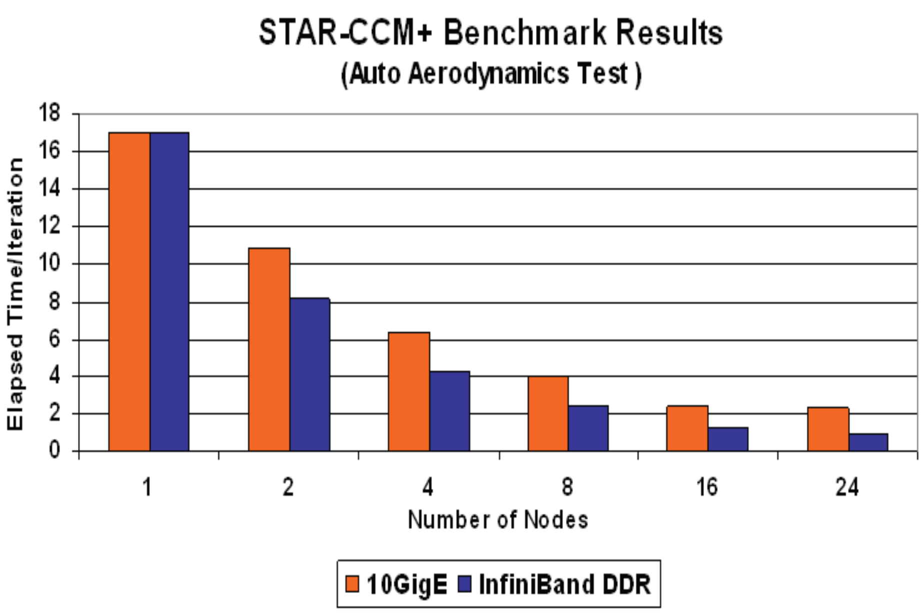

STAR-CCM+ is an entire engineering process for solving problems involving flow (of fluids or solids), heat transfer and stress.

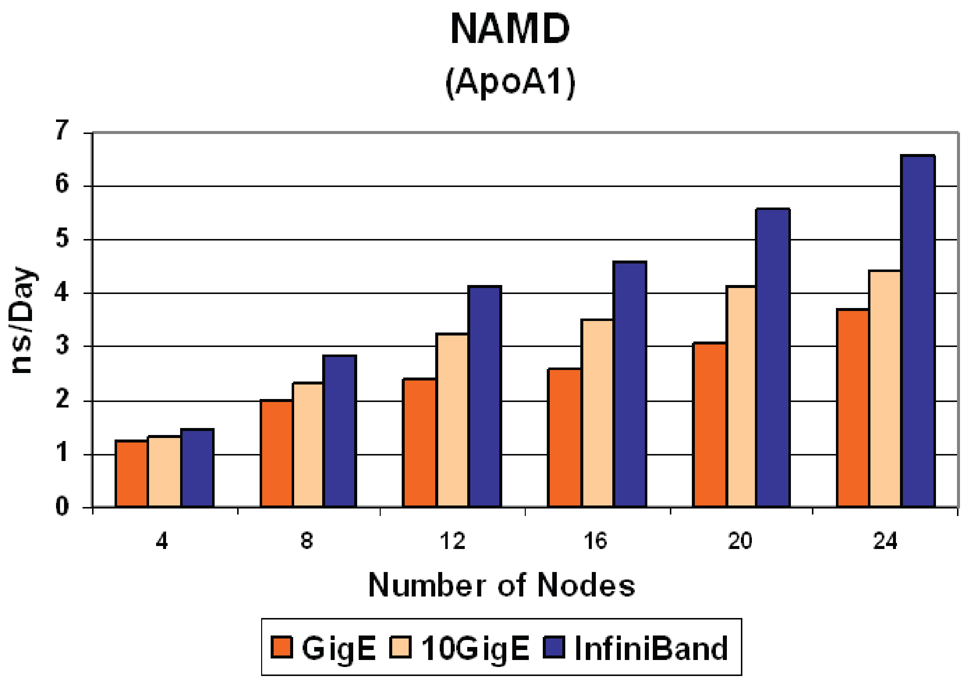

NAMD is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems

From: HPC Advisory Council

High availability

A system or a service must be UP: running and in health

- Redundancy of all network links and switches. If primary link fails, the second becomes active
- Servers in configuration active/standby
- Backup device uses heartbeat mechanism to send a signal to the primary device
- If the primary device stops responding, then a failover occurs
- The system/service is fault tolerant

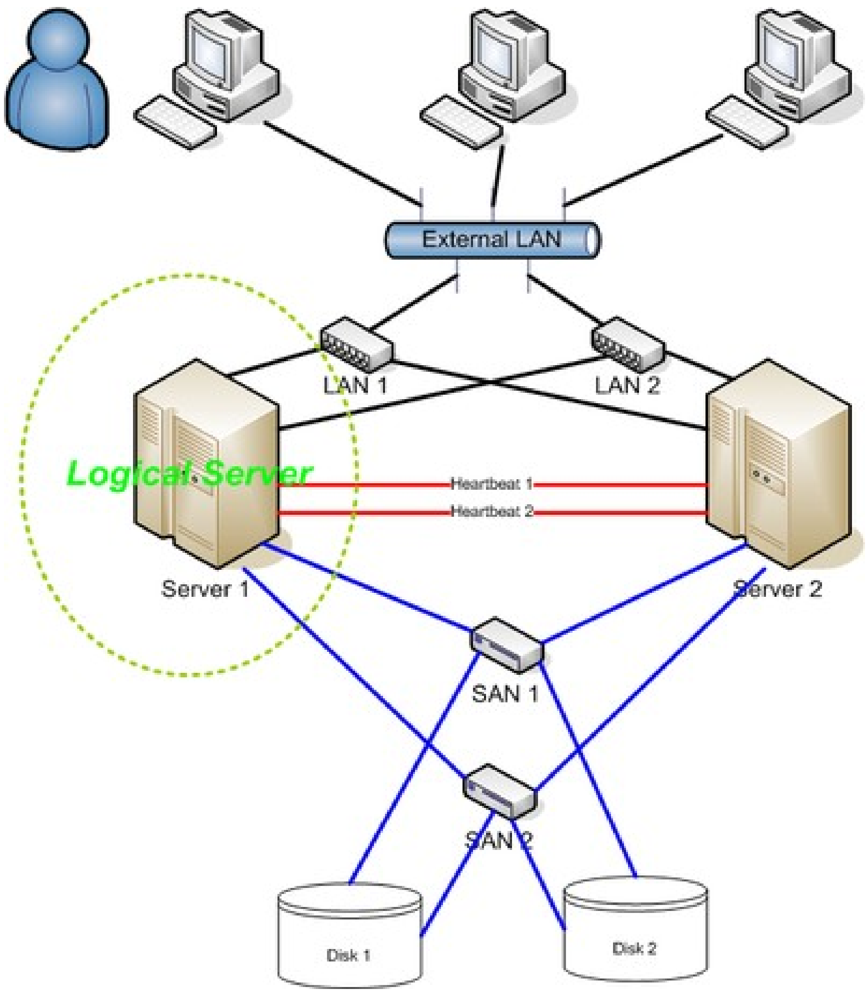

High availability

Keep my data safe!

RAID (Redundant Array of Indipendent Disk)

Is a fault tolerance solution for hard drives implemented in servers and storage system

RAID 1

consists of an exact copy (or mirror) of a set of data on two or more disks

Simplest RAID 1 system contains at least two disks

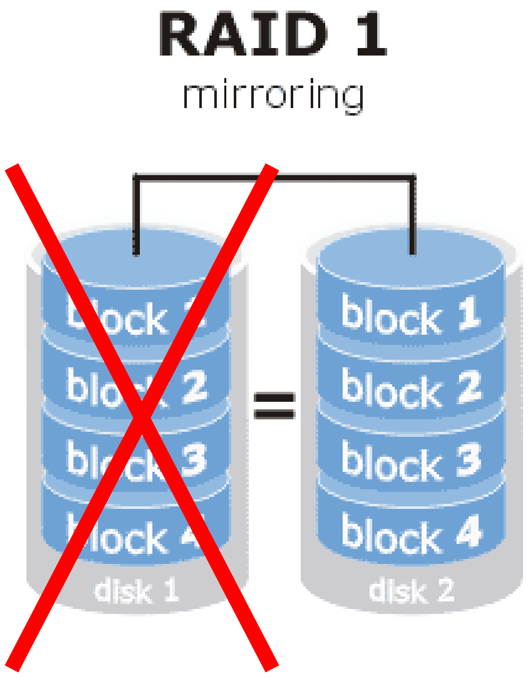

HPC infrastructure at CRS4: some numbers

HPC infrastructure at CRS4: some numbers

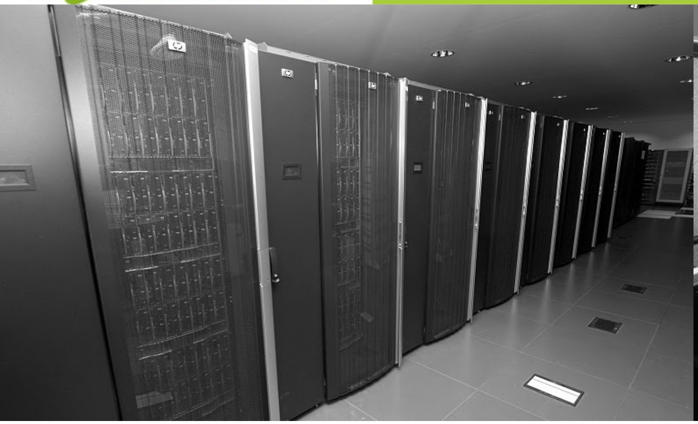

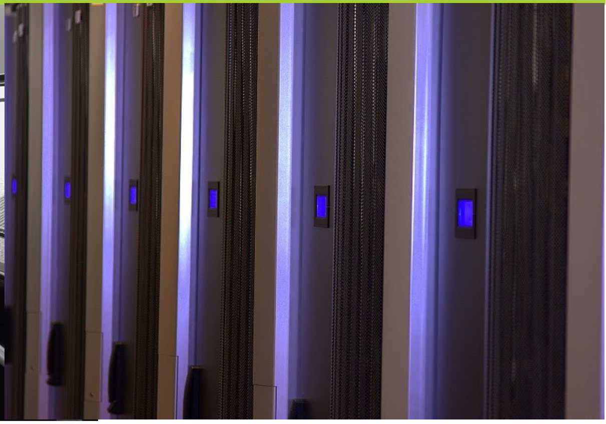

- ~500 compute nodes
- Hybrid architectures (Nvidia GPUs, Intel Phi, AMD GPUs, FPGA)
- >200 10GBps Ethernet ports, >1200 1Gps Ethernet ports
- ~280 Infiniband port
- 1 Gbps primary internet connection, 10Gbps next year
- ~280 TeraFlops peack processing speed (CPU+GPU)
- 4.5 Petabyte storage system
- New acquisitions of node, storage and network devices in progress

HPC infrastructure at CRS4: some numbers

	Total memory	Core #	networking	Notes
HP std Cluster				
CPU	6.7 TB	3200	IB-DDR20Gbps+GbEth	34,5Tflops
Huawei std Cluster				
CPU	4.4TB	656	IB-FDR56Gbps+GbEth	14.5Tflops
HP in house std Cluster BOTIN PARTNERS NAVAL ARCHITECTURE				
CPU	4.4TB	656	IB-FDR56Gbps+GbEth	14.5Tflops
Hybrid GPU Cluster				
CPU Intel	640GB	160	IB-QDR40Gbps+GbEth	2.8Tflops
GPU nVidia K10	160GB	61440	IB-QDR40Gbps+GbEth	91Tflops
Intel Phi coprocessor	80GB	600	IB-QDR40Gbps+GbEth	20Tflops
Hybrid CPU Cluster				
CPU Intel	1TB	96	Gb eth	2Tflops
GPU nVidia K40	96GB	23040	Gb eth	34Tflops
Hybrid ATI Cluster				
CPU Intel	1.3TB	200	IB-FDR56Gbps+GbEth	3.7Tflops
GPU AMD FirePRO	160GB	28160	IB-FDR56Gbps+GbEth	50Tflops

HPC infrastructure at CRS4: some numbers

Huawei eolo Cluster

33 nodes
14,5Tflops
4,4TB RAM
20TB disks space
656 Intel Xeon E5-2680 v2 @ 2.80GHz cores
66 1Gbps ports
33 Infiniband FDR 56Gbps low latency ports
Redundant power supplies
Redundant network links to the data center resources

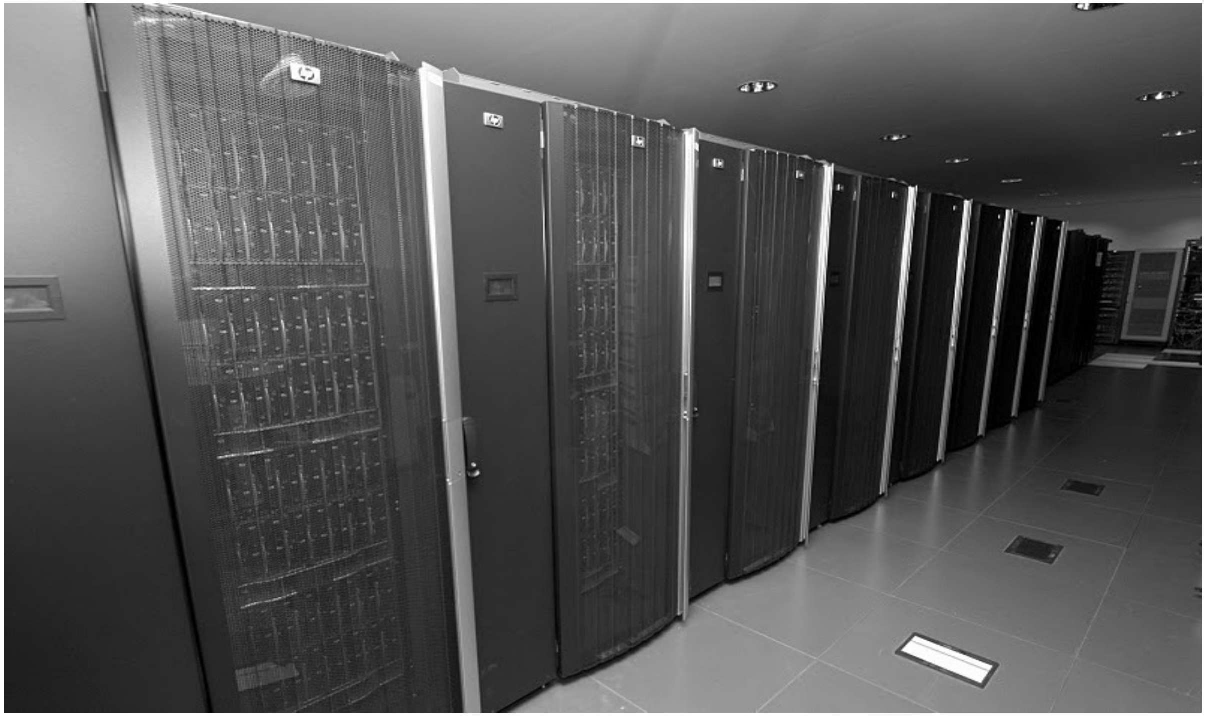

HPC infrastructure at CRS4: some numbers

Hybride architectures

- Mix of CPU and GPU
- An efficient way to run parallel and serial code
- CPU: powerful cores designed to run serial processes.Traditional compilers
- GPU: thousand of small cores optimised for parallel tasks. CUDA, NVIDIA parallel computing platform
- MIC (Many Integrated Core): x86-compatible multiprocessor architecture that utilize existing parallelization software tools (OpenMP)and specialized version of Intel Compilers
- FPGA: Integrated Circuit designed to be configured by a customer or a designer after manufacturing – hence "field-programmable". Usually requires proprietary software and complex programming tecniques

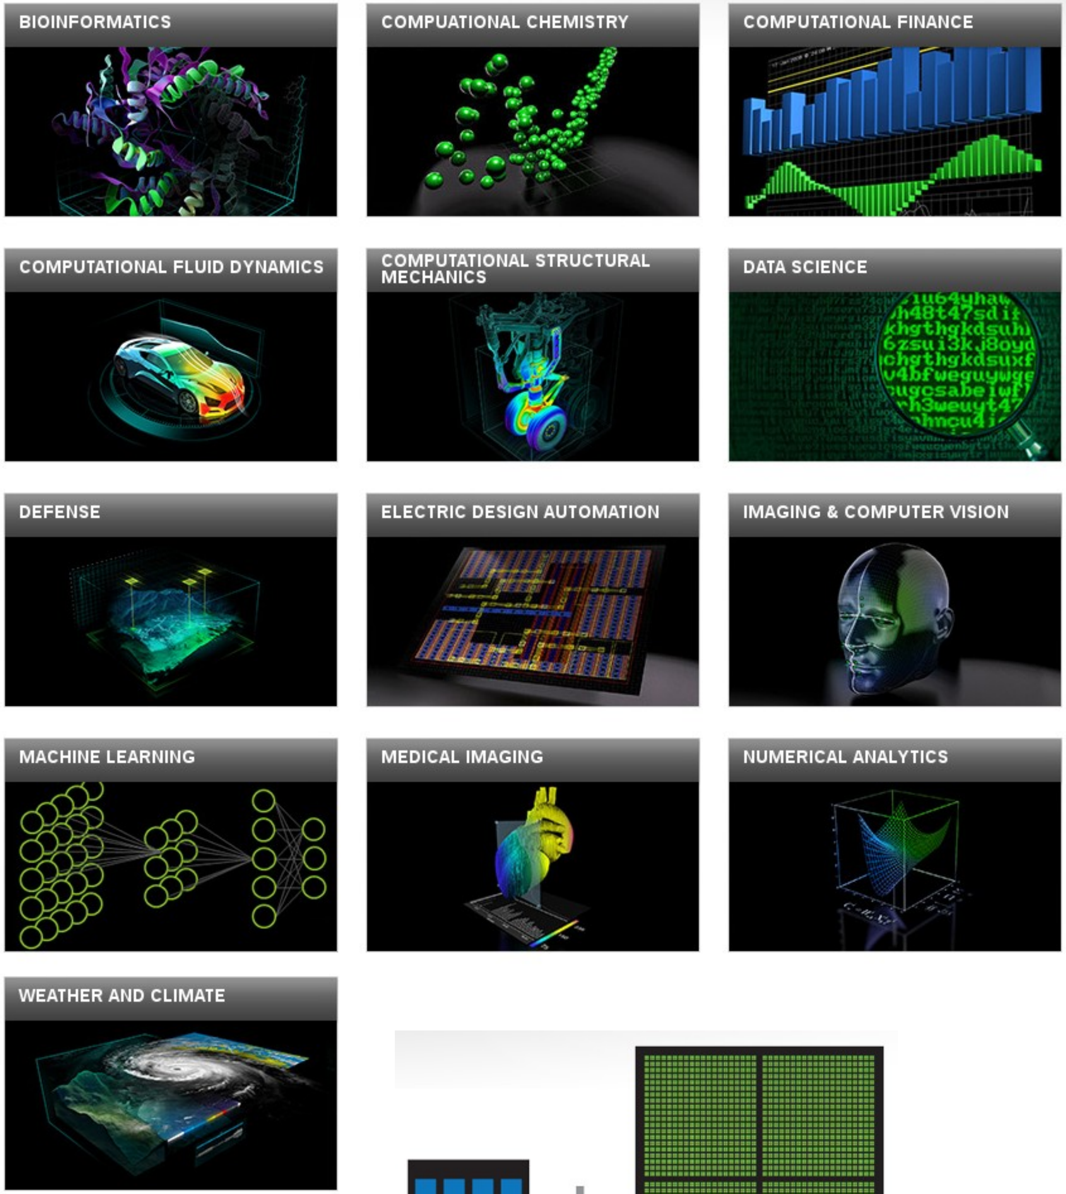

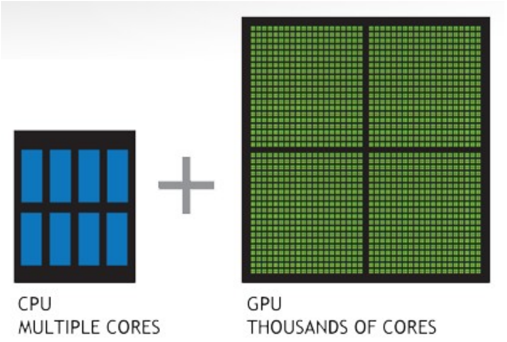

Clustering resources

Software @ CRS4

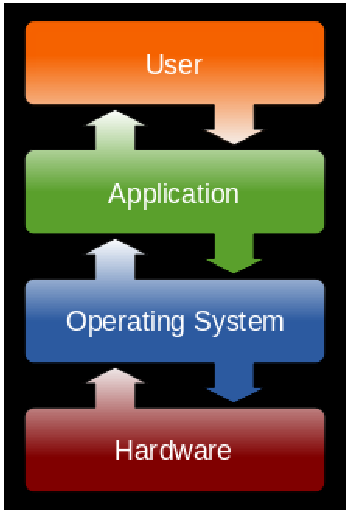

Compilers and programming tools:

- **Intel Compiler Suite (up to version 2016)**
- **PortlandGroup (PGI) compiler and debugger (up to version 14.9)**
- GNU compilers (up to 6.1)
- CUDA (up to 7.5)
- Java
- Valgring
- ACMLib
- MKLib

Libraries:

atlas3.10.0	Fftw-3.3.2	Igraph-0.6.5	Picard-tools-1.123	Blacs
Scalapack	Boot	GotoBlas	Gatk-2.1-8	Libevent-2.0.22
Glibc-2.14	Libgtextutils-0.6.1	sparsehash-2.0.2	Boost	Glibc-2.17
Libint-2.0.3	srma-0.1.15	Gmp-5.0.5	Lzo-2.06	SuiteSparse
magma	Yaml-0.1.5	Mpc-0.9	zlib-1.2.8	Bzip2-1.0.6
mpfr-3.1.1	Clapack-3.1.1.1	Mymedialite-3.10	LessTiff	Libxc
Libxp	Metis	Parmetis	cgal	curl-7.31.0
grib_api-1.8.0	gsl-1.11	ice-3.4.2	....	

Manage your environment on HPC cluster

- Availability of **hundreds of different softwares**, libraries and tools (also different versions of the same package)

- The Environment Modules Package (Modules) provides an help to the dynamic modification of a user's environment via modulefiles:
  - Allows the system admins to install differents software versions
  - Users are able to customize own environment by creating and loading their modulefiles
  - Dynamically modifies system variables (PATH, LD_LIBRARY_PATH ....) and/or application variables

Examples:

```
-bash-4.2$ module load GCC-5.2.0
```

```
-bash-4.2$ module list
Currently Loaded Modulefiles:
  1) GCC-5.2.0_CentOS7      3) perl-5.20.2_CentOS7    5) profilo_v3             7) galaxy_v3
  2) python-2.7.9_CentOS7   4) jdk1.8.0_40            6) samtools-1.3
```

```
-bash-4.2$ module display starccm11.02.009
-------------------------------------------------------------------
/Archive/Software/Modules/Modules/3.2.10/modulefiles/OS_TYPE/CentOS7/SOFTWARE/starccm11.02.009:

module-whatis    puts starccm11.02.009 in your PATH
prepend-path     PATH /Archive/Software/CentOS7/StarCD/STAR-CCM+11.02.009-R8/star/bin:/Archive/Software/CentOS7/StarCD/STAR-CCM+11.02.009-R8
:/Archive/Software/CentOS7/StarCD/STAR-View+11.02.009
prepend-path     CDLMD_LICENSE_FILE /Archive/Software/StarCD/license.dat
-------------------------------------------------------------------
```

Job Scheduler: GridEngine

Sharing and organizing a large number of computational resources for a large number of jobs requires a DRM (Distributed Resource Manager) software

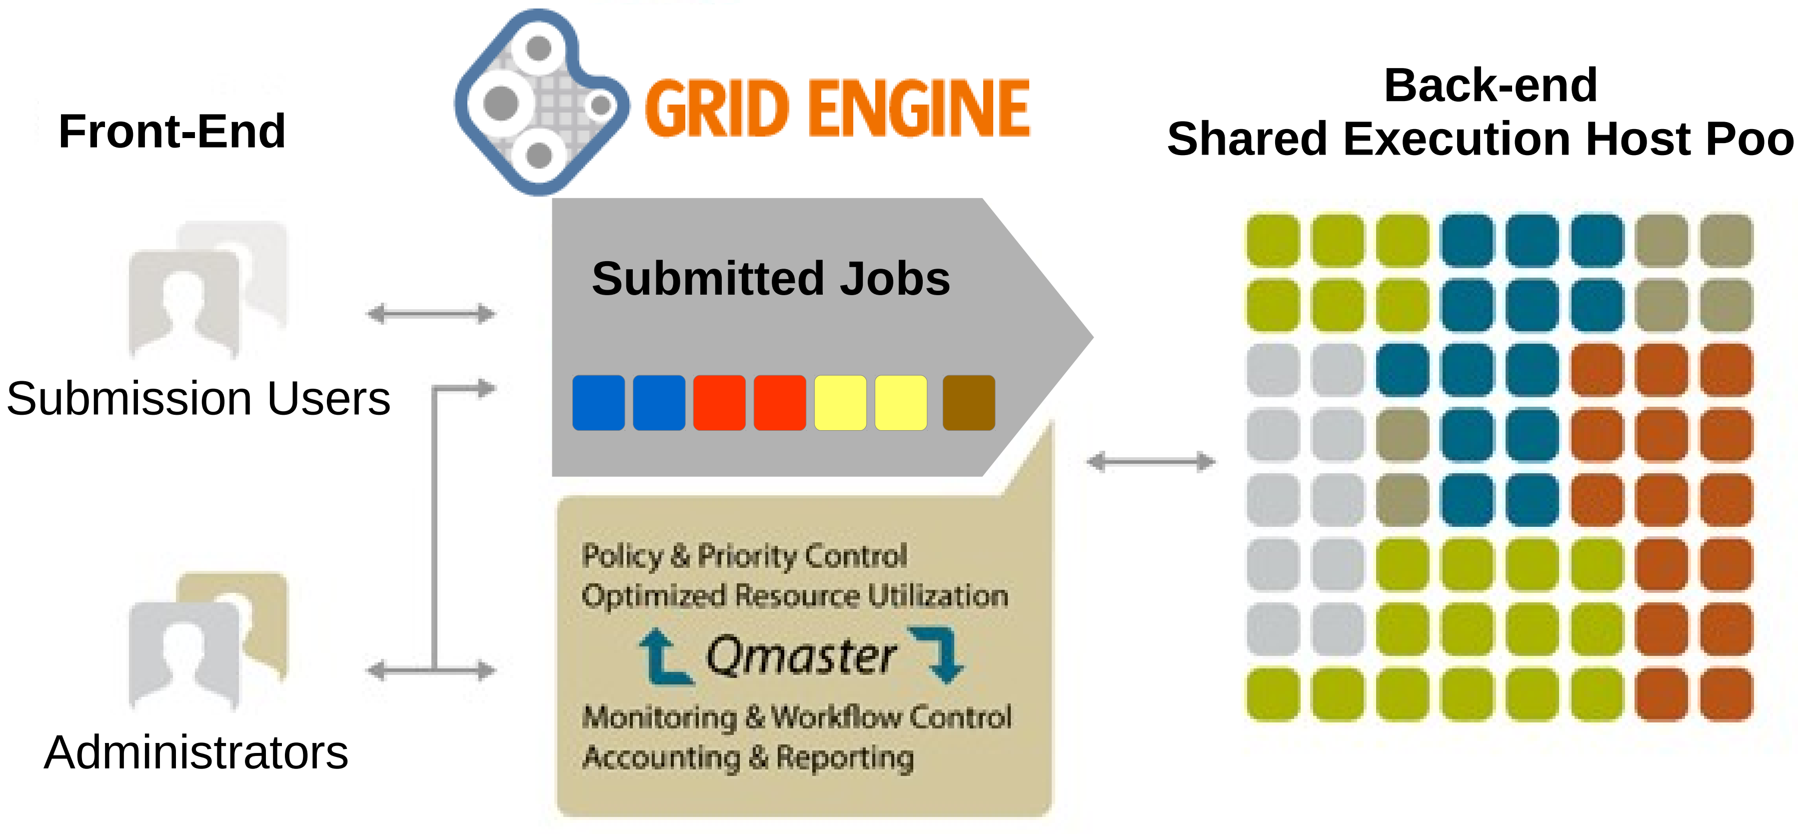

Job Scheduler: GridEngine

```
[root@mommoti01 ~]# qstat
job-ID  prior   name       user         state submit/start at     queue                          slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------
1750496 0.60500 hbasebig   surfer       r     10/10/2012 10:26:36 hdfs@entu065.crs4.int            512
1762426 0.50500 flexcryst_ hofmann      r     10/11/2012 15:42:53 interactive@entu003.crs4.int       1
1763213 0.50500 QRLOGIN    julie        r     10/12/2012 12:40:39 interactive@entu003.crs4.int       1
1767998 0.50559 TD16.0_pha zara         r     10/14/2012 10:49:47 genoma.gwa@oghe017.crs4.int        4
1768342 0.50637 c20_1mMSAL zara         r     10/15/2012 00:49:32 genoma.gwa@oghe042.crs4.int        8
1768625 0.50500 QRLOGIN    julie        r     10/15/2012 13:02:38 interactive@entu002.crs4.int       1
1768628 0.50520 GA-codesa3 mdentoni     r     10/15/2012 13:34:17 doublew@entu029.crs4.int           1
1768735 0.50520 case4      karalit      r     10/15/2012 14:14:21 doublew@entu032.crs4.int           1
1770750 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:12:35 doublew@entu019.crs4.int           1
1770751 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:13:20 doublew@entu023.crs4.int           1
1770757 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:16:43 doublew@entu034.crs4.int           1
1770758 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:17:13 doublew@entu052.crs4.int           1
1770759 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:17:58 doublew@entu053.crs4.int           1
1770762 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:18:58 doublew@entu054.crs4.int           1
1770764 0.50500 GA-codesa3 mdentoni     r     10/15/2012 16:19:43 doublew@entu055.crs4.int           1
1771287 0.50500 chr1_XHG   maria        r     10/15/2012 17:41:21 genoma.gwa@oghe015.crs4.int        1
1771288 0.50500 chr2_XHG   maria        r     10/15/2012 17:41:21 genoma.gwa@oghe015.crs4.int        1
```

GridEngine is the **main way** to access CRS4's HPC clusters resources

hadoop framework

- Hadoop provides an effective and scalable way to process large quantities of data (big data)
- Hadoop has two goals:
  - scalable storage (HDFS)
  - scalable computation (MapReduce)
- Today Hadoop scale up to thousands of nodes

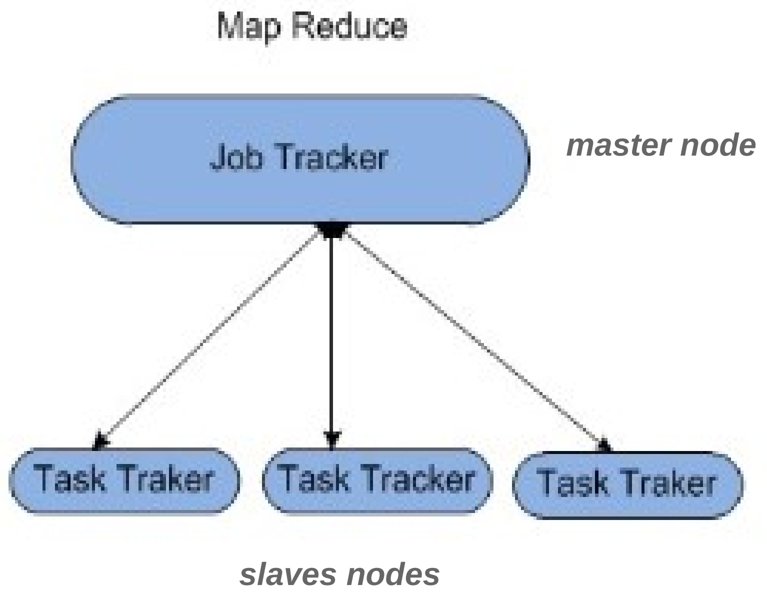

HADOOCCA

Developed by

- Hadoop assumes it has **exclusive and long-term** use of its nodes
- It has its own job submission, queueing, and scheduling system
- But CRS4 HPC resources are accessed exclusively via GridEngine queue system

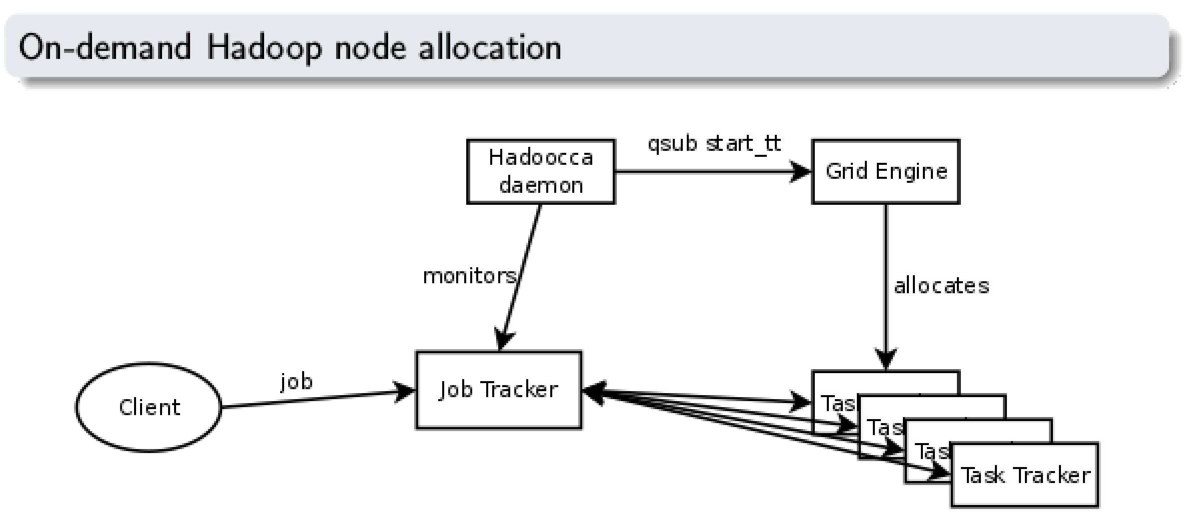

- HADOOCCA is a tool that controls and implements dynamic Hadoop MapReduce clusters creation by the use of GridEngine queue system

Management monitoring maintenance

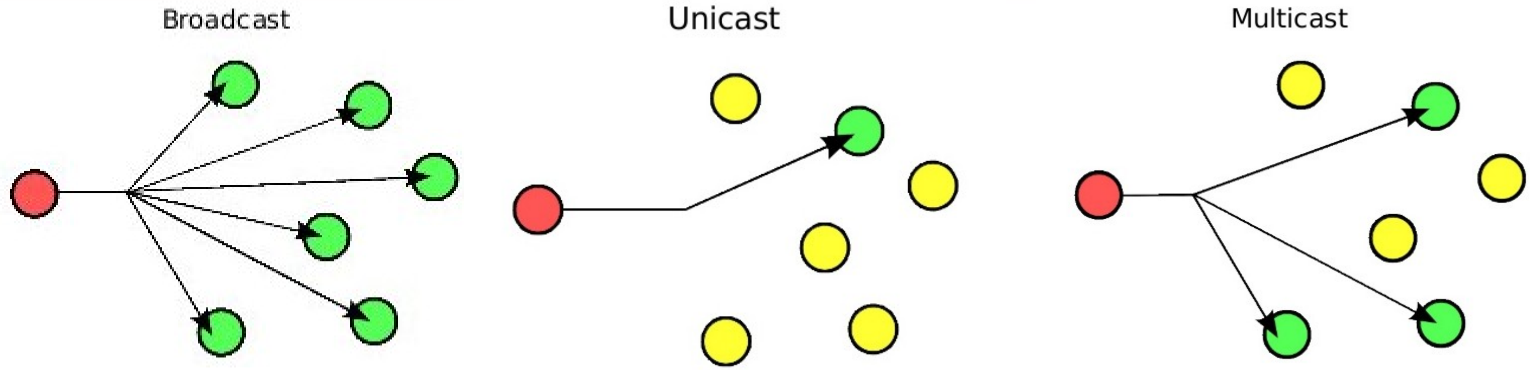

- Scalable (>2000 nodes) distributed monitoring system
- Each node of the group runs gmond daemon and sends its state to a multicast address
- Every node receives and stores the status of each other node of the group
- Ganglia gmetad server acquires the status of the group by querying a random node.

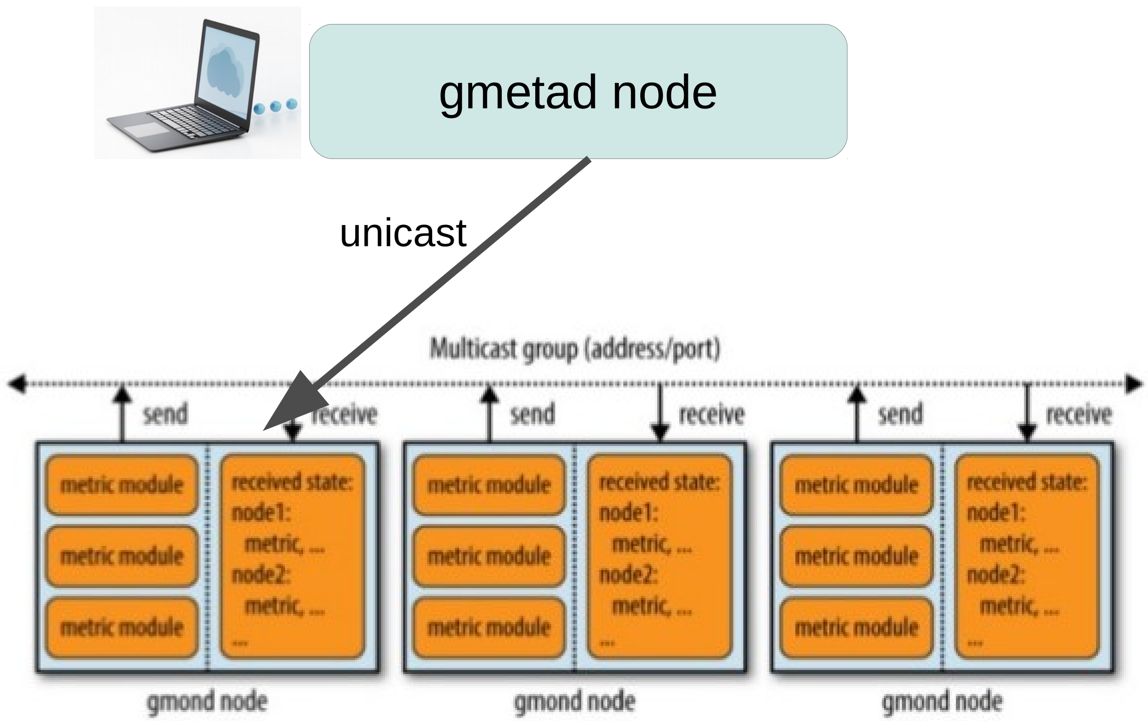

Default multicast topology

management monitoring maintenance

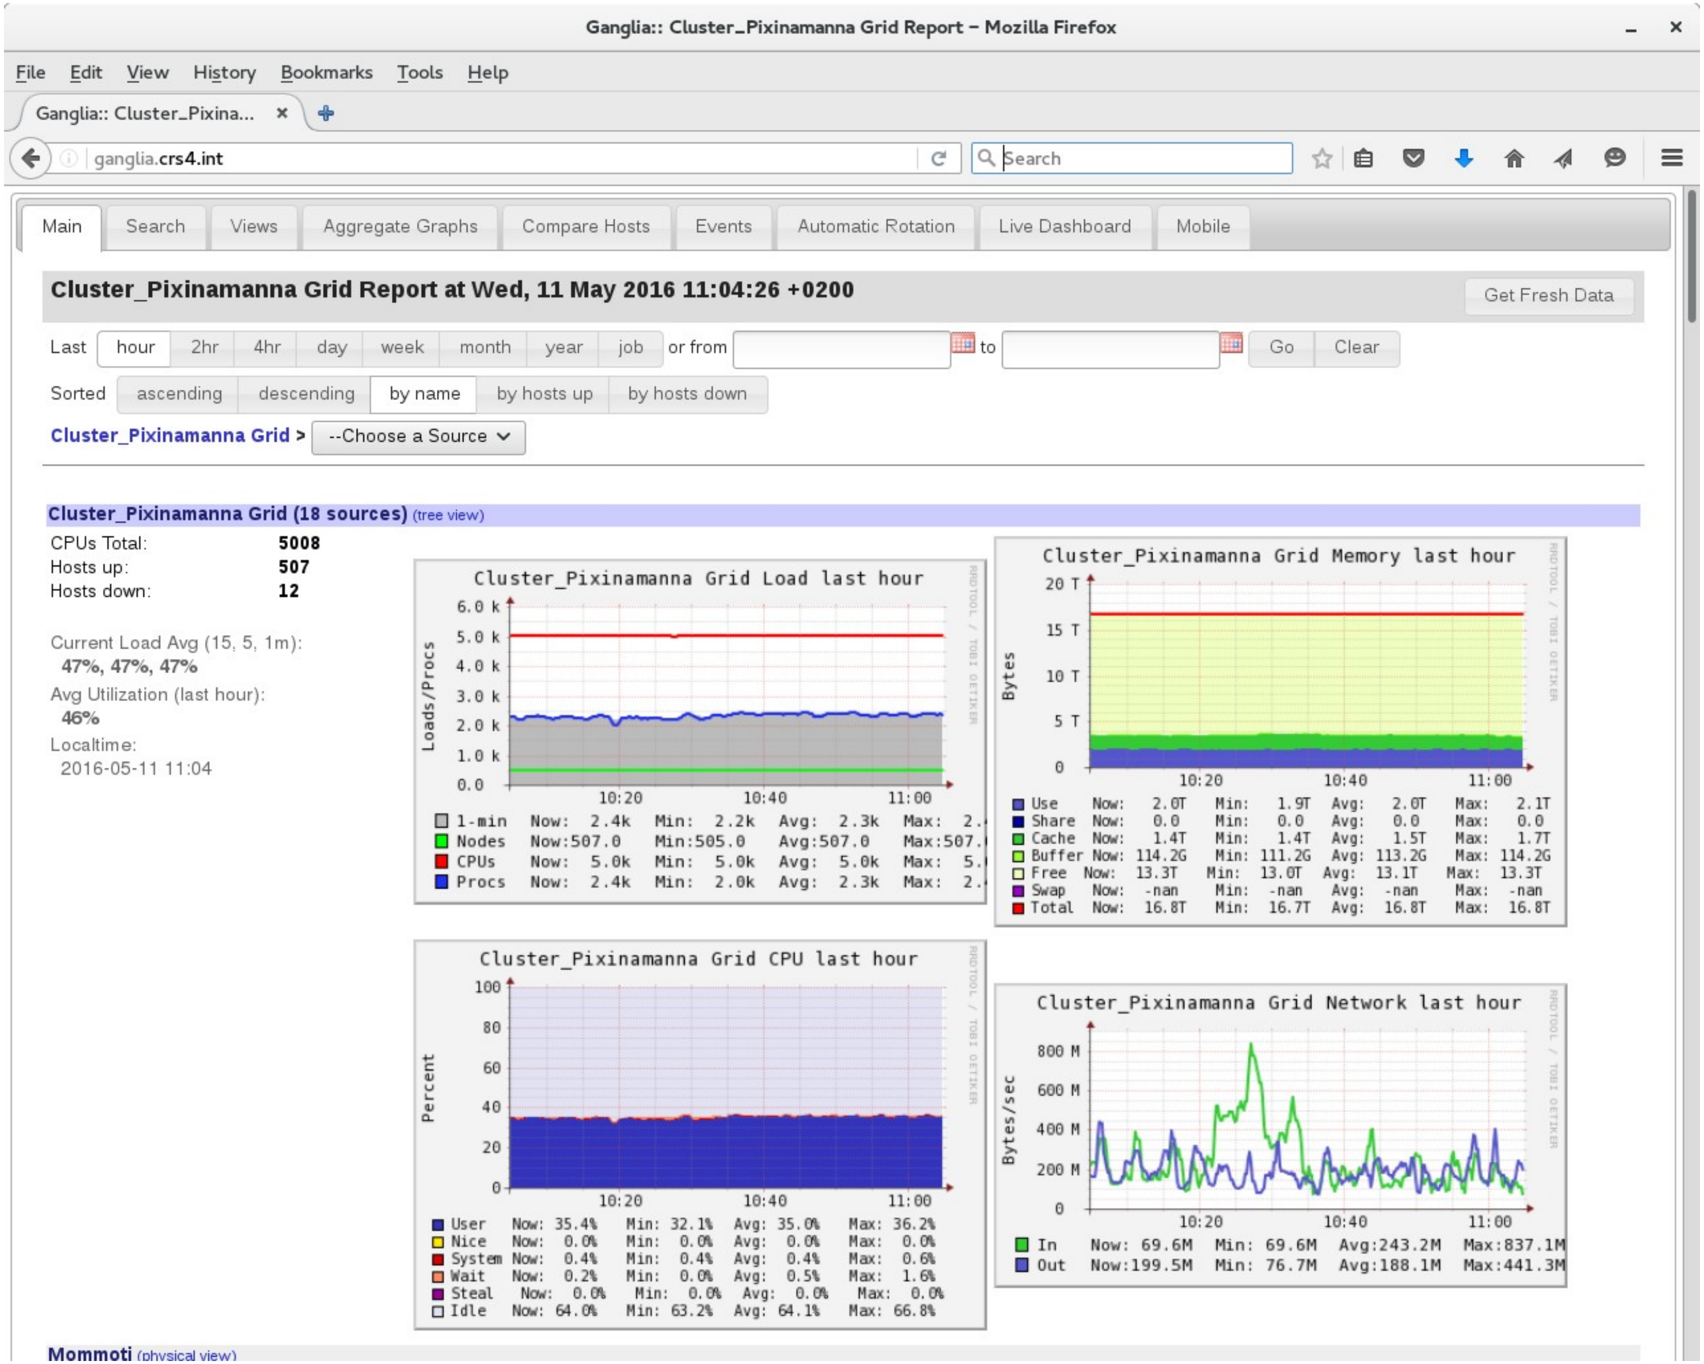

Xymon - MonitorX

- monitoring tool for servers, custom applications and networks e.g. :
  - network services: http, ftp, smtp and so on
  - local resources like disk utilisation,memory,cpu load, logfiles, processes
- The information is collected by a central server and presented in a set of simple, intuitive webpages frequently updated to reflect changes in the status of systems
- It also records the history of monitored item and generates reports and graphs
- Alerts for issues may be sent in form of e-mails or SMS-messages

management monitoring maintenance

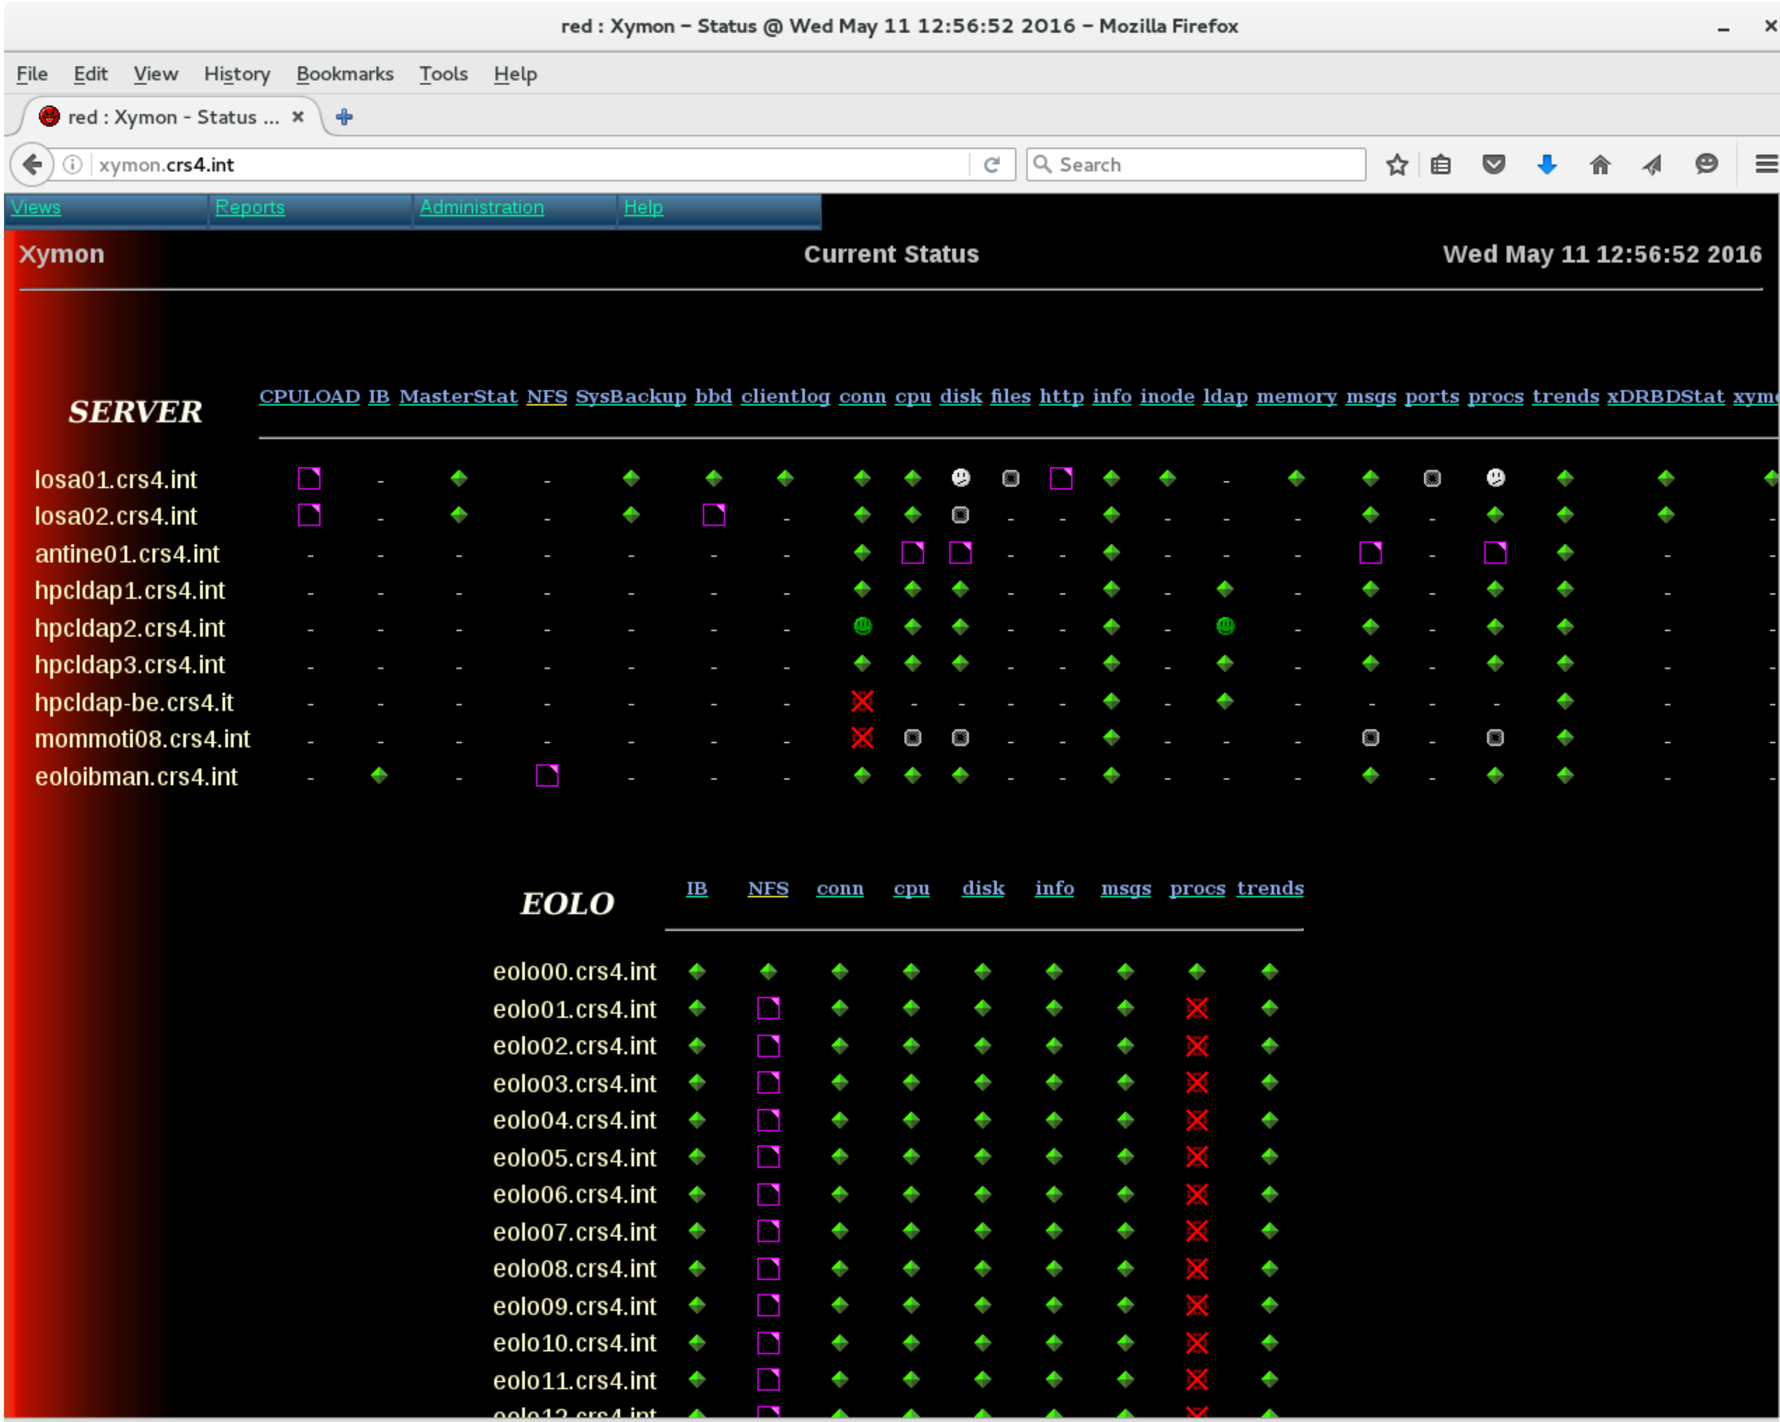

management monitoring maintenance

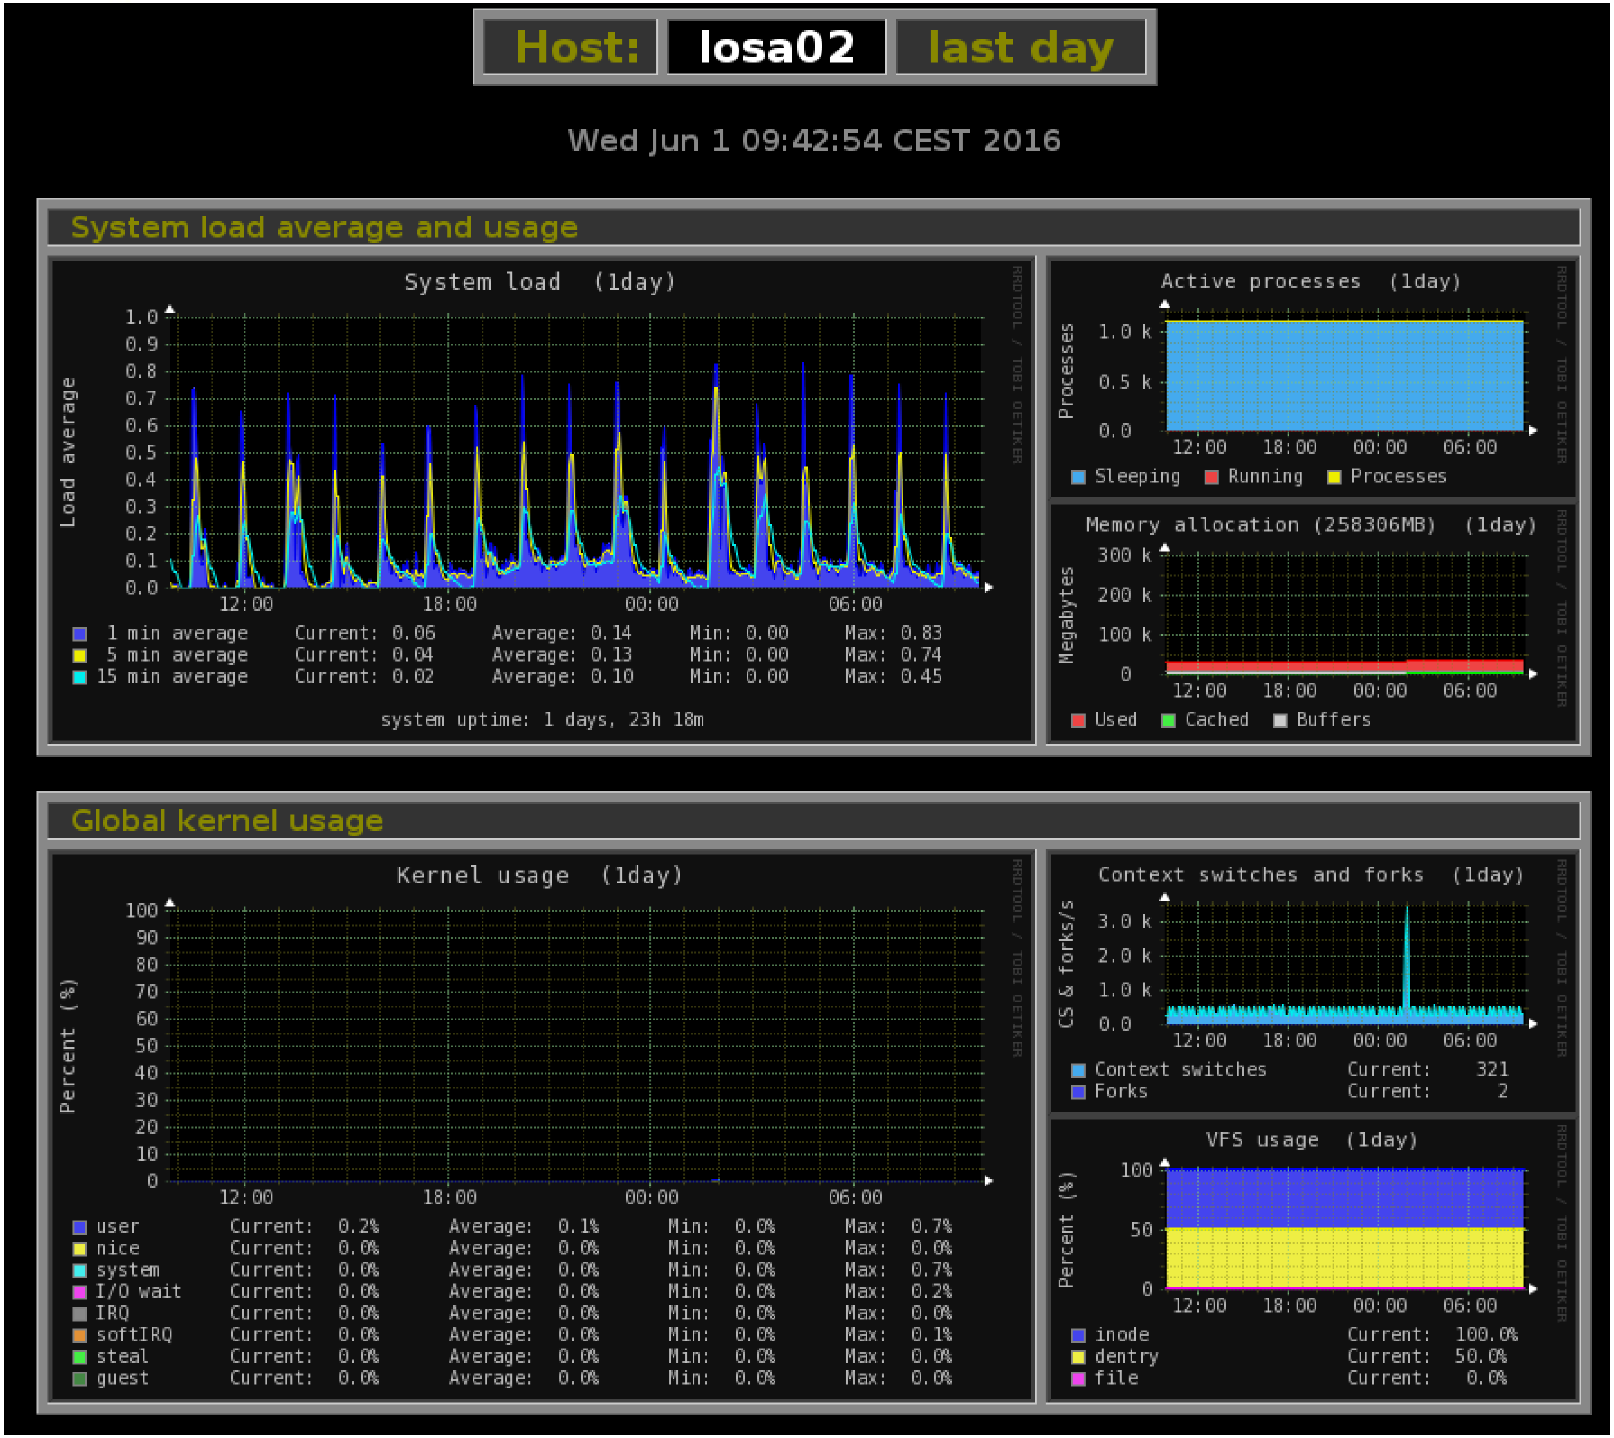

management monitoring maintenance

MUCCA

MANAGER UTILITY FOR CLUSTER AND COMPUTER ADMINISTRATION

Mucca is a collection of useful system administrator tools and scripts developed by CRS4 HPC group.

All collected information is stored into a SQL DB.

Mucca provides an useful web interface showing the status of Clusters.

With a click of mouse We can select hosts, groups of hosts and entire racks.

Topics:

- Power off/on (soft or cold power on/off)
- Reinstall hosts or restart hosts with an alternative OS
- Check the status in graphical or textual mode
- Plan a reboot disabling queues on active hosts
- Display hosts where user jobs are running
- Show the load values (RAM, CPU or both) for a single host
- Enable/disable scheduler queues
- Enable/disable “CRS4 Green” utility
- And much more!

management monitoring maintenance

MANAGER UTILITY FOR CLUSTER AND COMPUTER ADMINISTRATION

CRS4 MUCCA Green utility

- The goal is to reduce significantly the electricity consumption by turning off unused compute node
- The access to cluster resources is allowed exclusively through a centralised queue system
- Thus by monitoring the queue system, MUCCA GREEN interacts with Grid Engine deciding if turn on/off compute hosts

power on/off

IPMI
ILO protocols
ILOM

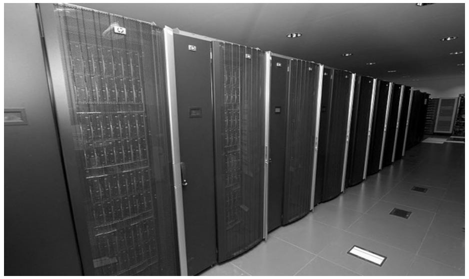

monitors

```
326577 4.02258 Artery_Aor ostefano     r     05/20/2016 11:16:02 gpu@digo08.crs4.int            1
326578 4.02258 Artery_Cor ostefano     r     05/20/2016 11:16:02 gpu@digo09.crs4.int            1
326579 4.02258 Artery_Tib ostefano     r     05/20/2016 11:16:02 gpu@digo03.crs4.int            1
326621 5.41781 m2s3-npt8  akumar       r     05/20/2016 11:46:17 eolo@eolo17.crs4.int         160
994037 2.68623 sleep      carlo        qw    10/29/2015 11:51:35                                1
994038 2.13968 test       carlo        qw    10/29/2015 11:52:37                                1
326580 2.13882 Brain_Ante ostefano     qw    05/20/2016 11:15:49                                1
326581 2.04091 Brain_Caud ostefano     qw    05/20/2016 11:15:49                                1
326582 1.95988 Brain_Cere ostefano     qw    05/20/2016 11:15:49                                1
326583 1.89100 Brain_Cere ostefano     qw    05/20/2016 11:15:49                                1
```

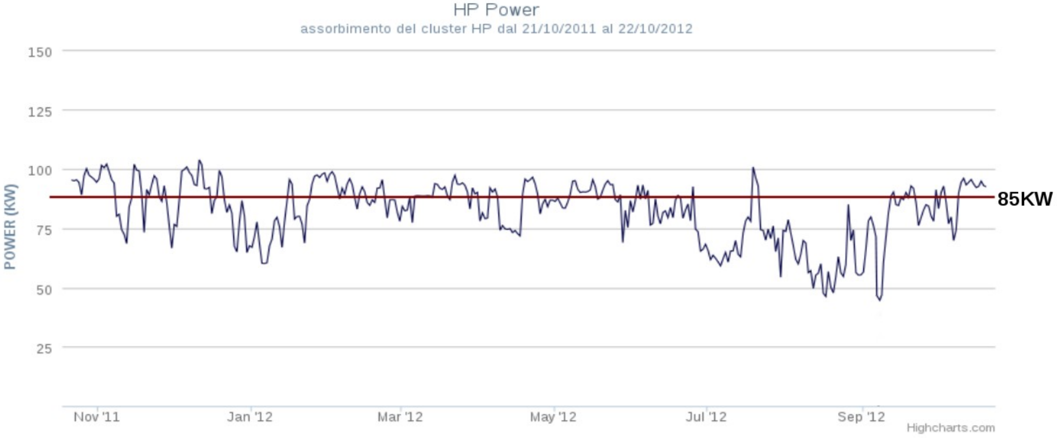

management monitoring maintenance

management monitoring maintenance

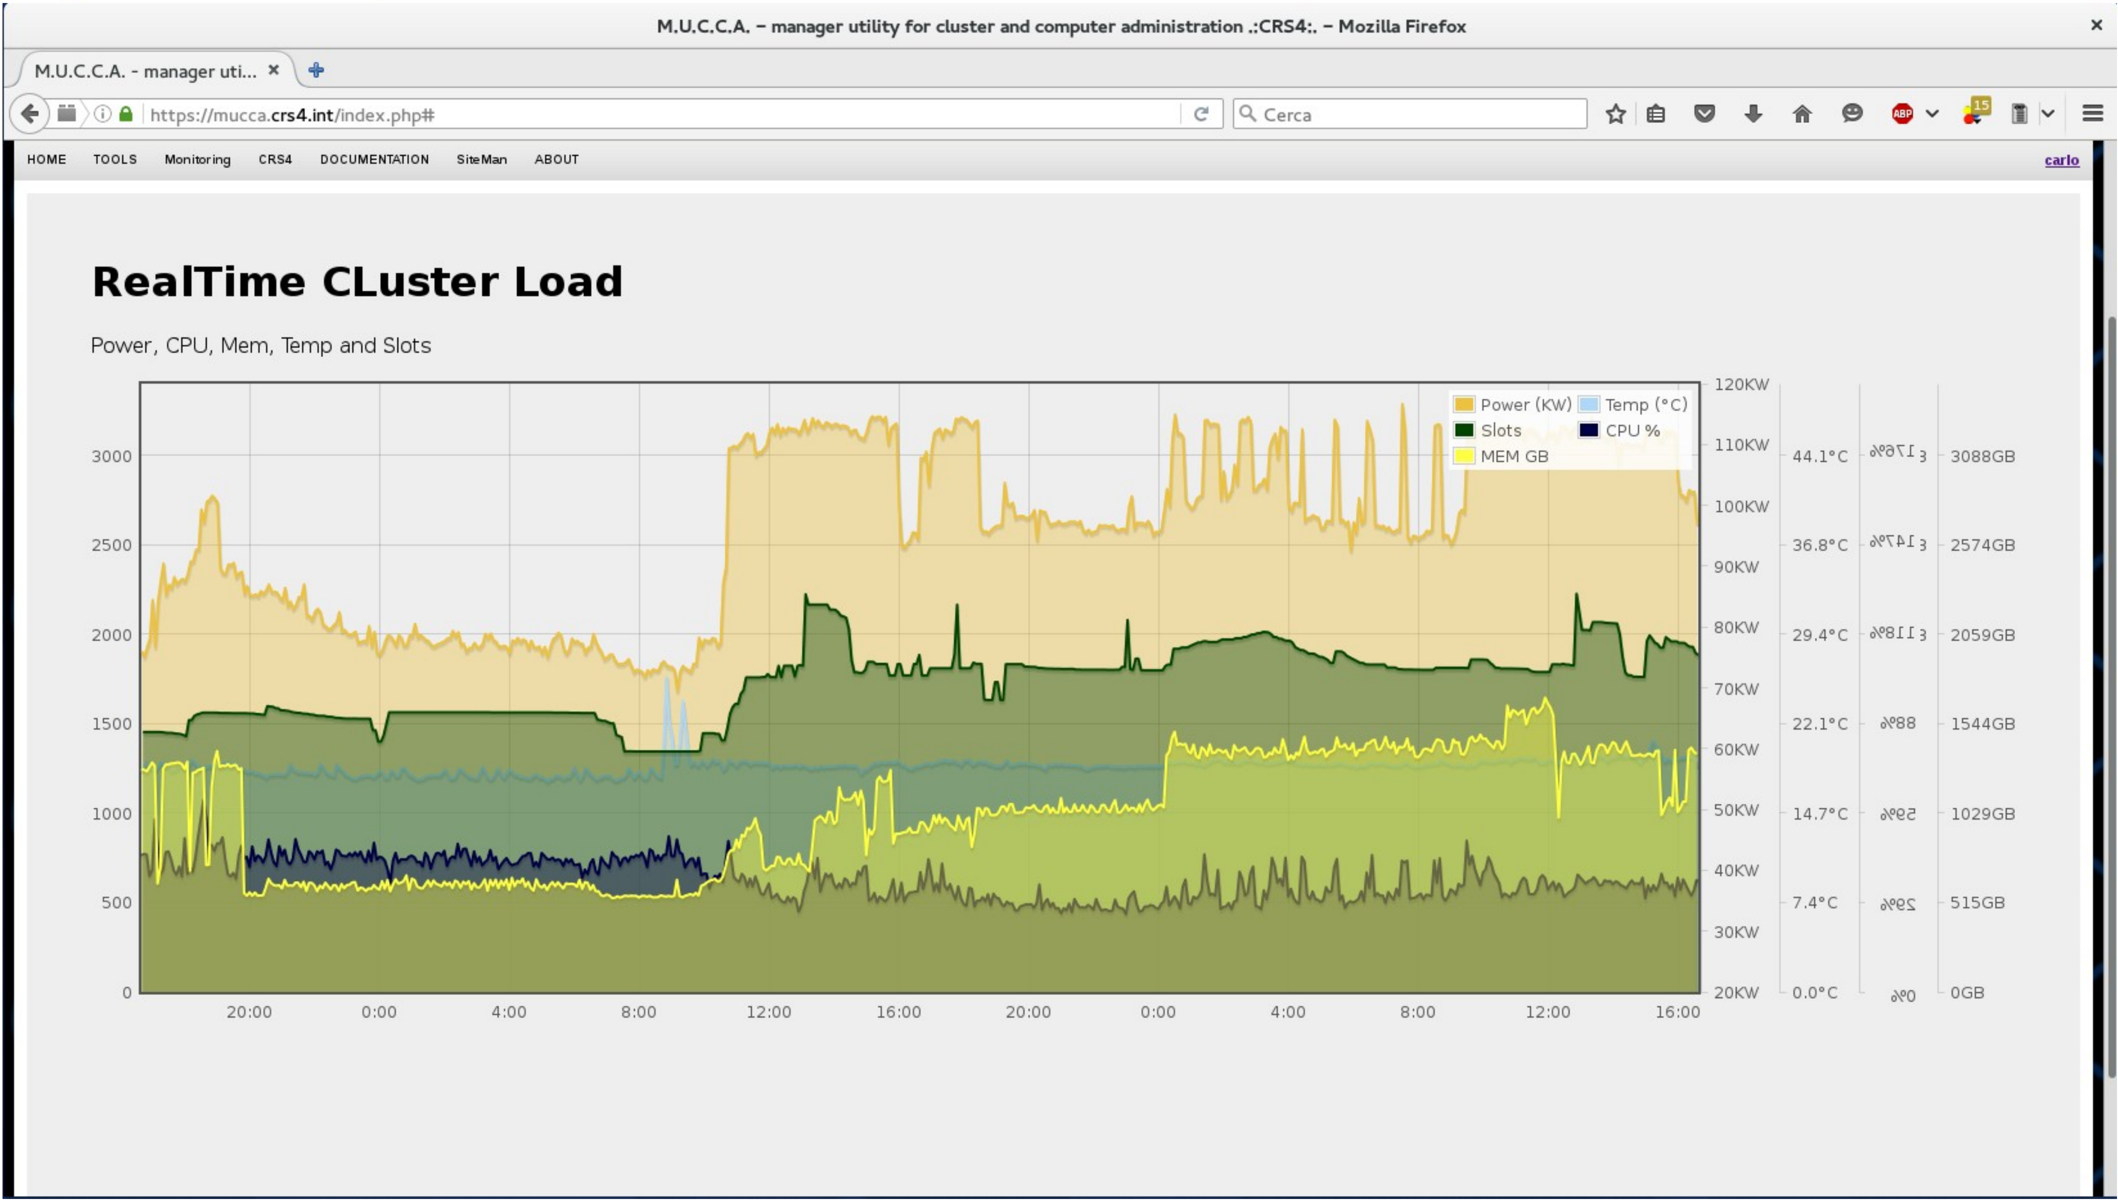

management monitoring maintenance

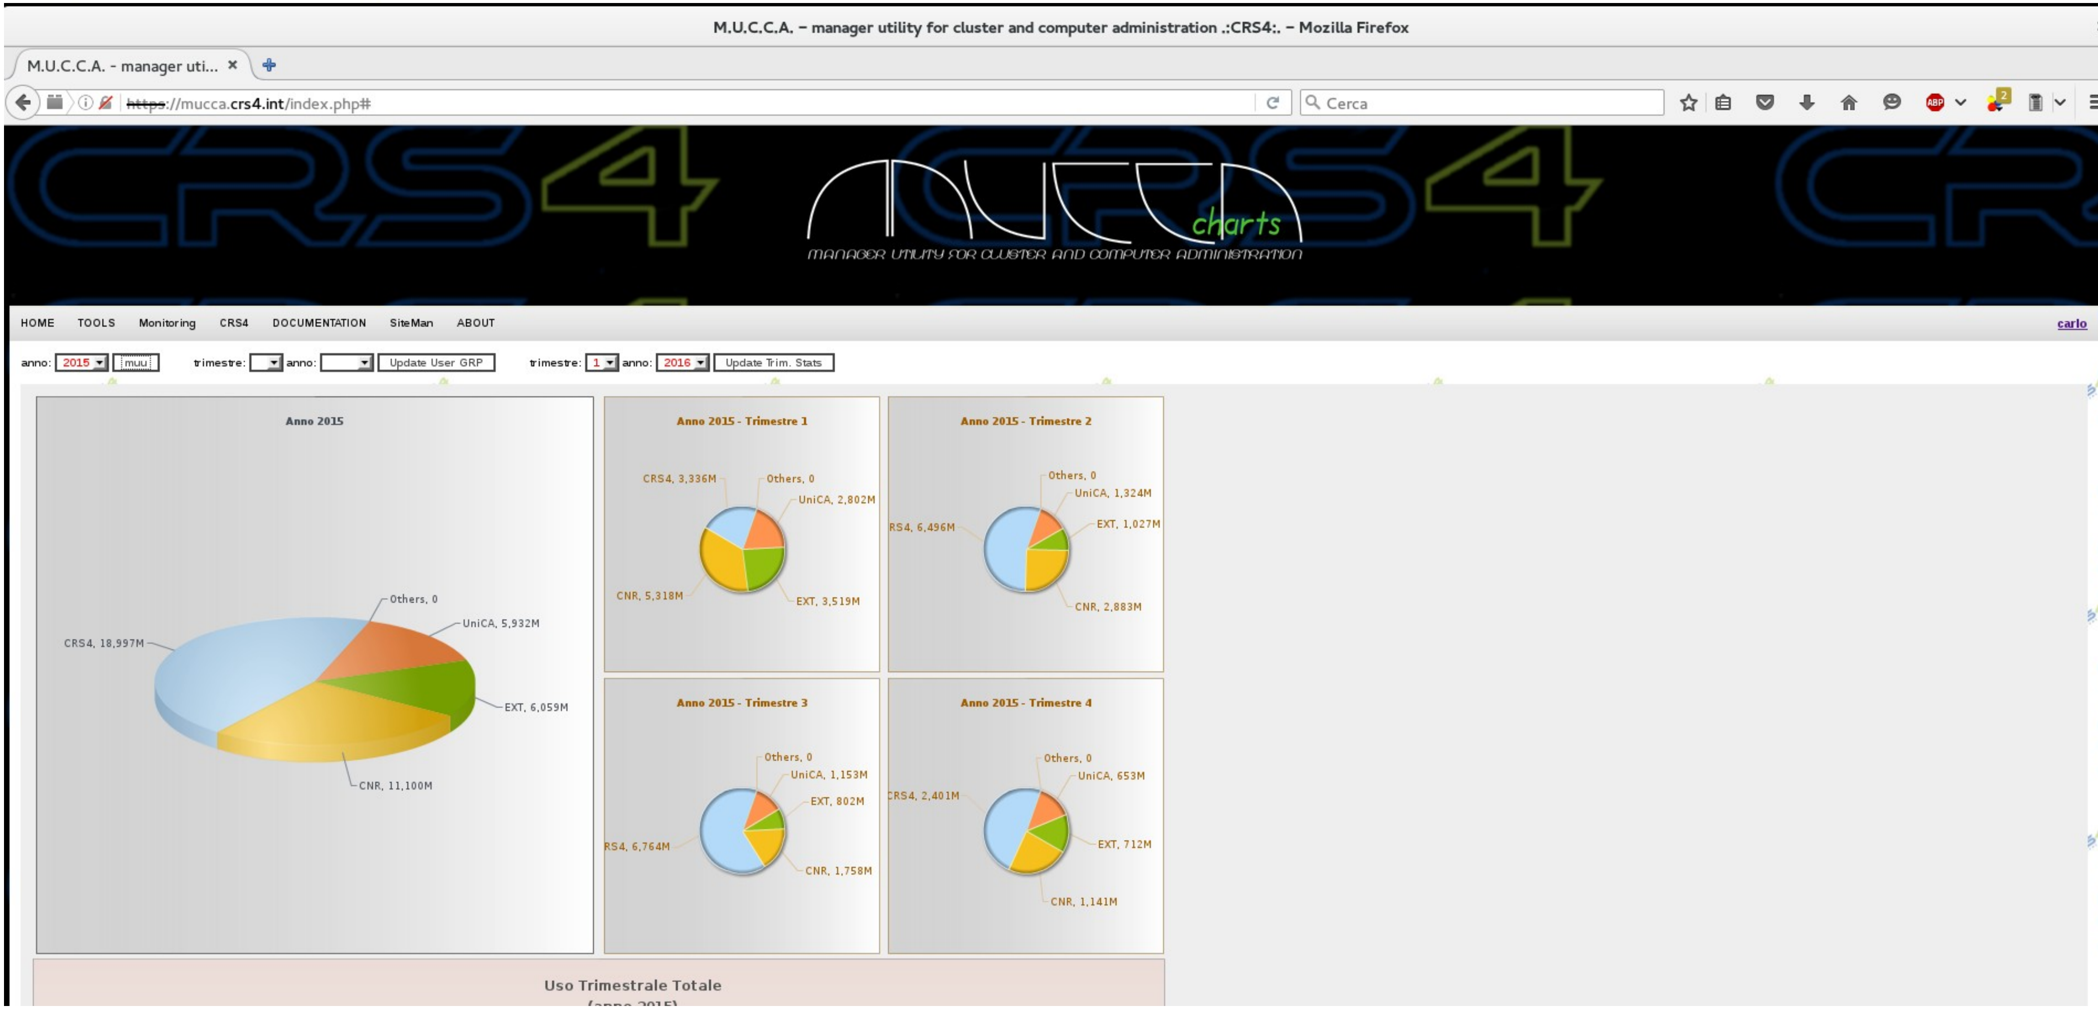

management monitoring maintenance

M.U.C.C.A. – manager utility for cluster and computer administration .:CRS4:. – Mozilla Firefox

M.U.C.C.A. - manager uti...

https://mucca.crs4.int/index.php#

MANAGER UTILITY FOR CLUSTER AND COMPUTER ADMINISTRATION

CLUSTER

HOME TOOLS Monitoring CRS4 DOCUMENTATION SiteMan ABOUT carlo

Filtra Risultati:

mese: [all, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12] anno: 2015 cluster: gruppo: mu

utilizzo clusters mensile per gruppo

CLUSTER	TOT S totale richi	% TOT SLOT	TOT ALLOC totale slot effettivamente allocati	% TOT ALLOC	TOT WC sec tempo di utilizzo sui core richiesti	TOT WC h	% TOT WC	TOT WC ALLOC sec tempo di utilizzo sui core effettivamente allocati	TOT WC ALLOC h	% TOT WC ALLOC
October 2015 (2682000 sec.)										
biohosts (112 core)	462		4624		492955	136.93		492955	136.93	0.16%
BIOINFO (CRS4)	4419	95.57%	4419	95.57%	383282	106.47	77.75%	383282	106.47	77.75%
GWAS (CNR)	193	4.17%	193	4.17%	109427	30.4	22.2%	109427	30.4	22.2%
HPCN (CRS4)	12	0.26%	12	0.26%	246	0.07	0.05%	246	0.07	0.05%
TOT biohosts October 2015					492955	136.93		492955	136.93	
digo (160 core)	7212		12897		276676214	76854.5		277160279	76988.97	64.59%
BIOINFO (CRS4)	4589	63.63%	10169	78.85%	5708819	1585.78	2.06%	6166034	1712.79	2.22%
ENERGY (CRS4)	336	4.66%	336	2.61%	71200	19.78	0.03%	71200	19.78	0.03%
ENV (CRS4)	6	0.08%	51	0.4%	458177	127.27	0.17%	472907	131.36	0.17%
GWAS (CNR)	4	0.06%	64	0.5%	808	0.22	0%	12928	3.59	0%
HPCN (CRS4)	5	0.07%	5	0.04%	154138	42.82	0.06%	154138	42.82	0.06%
IOM (CNR)	2272	31.5%	2272	17.62%	270283072	75078.63	97.69%	270283072	75078.63	97.52%
TOT digo October 2015					276676214	76854.5		277160279	76988.97	
entu (2152 core)	60219		60248		1173396906	325943.59		1173401029	325944.73	20.33%
AGCT (CRS4)	1	0%	1	0%	16	0	0%	16	0	0%
BIOENG (CRS4)	1719	2.85%	1719	2.85%	63816863	17726.91	5.44%	63816863	17726.91	5.44%
BIOINFO (CRS4)	20240	33.61%	20240	33.59%	50371566	13992.1	4.29%	50371566	13992.1	4.29%
BIOMED (CRS4)	6080	10.1%	6080	10.09%	13167968	3657.77	1.12%	13167968	3657.77	1.12%
DC (CRS4)	18305	30.4%	18305	30.38%	60724999	16868.06	5.18%	60724999	16868.06	5.18%
ENERGY (CRS4)	1280	2.13%	1280	2.12%	2762240	767.29	0.24%	2762240	767.29	0.24%
ENV (CRS4)	539	0.9%	539	0.89%	136130850	37814.13	11.6%	136130850	37814.13	11.6%
FUEL (CRS4)	4	0.01%	4	0.01%	323401	89.83	0.03%	323401	89.83	0.03%

HPC research at CRS4

HPC research at CRS4

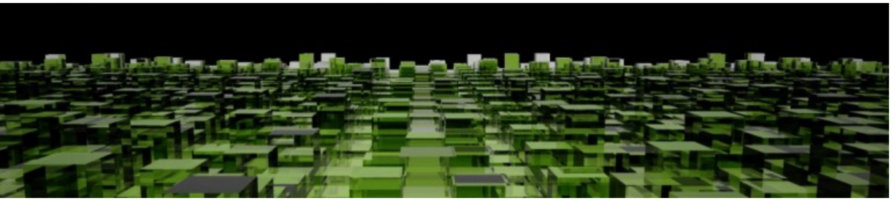

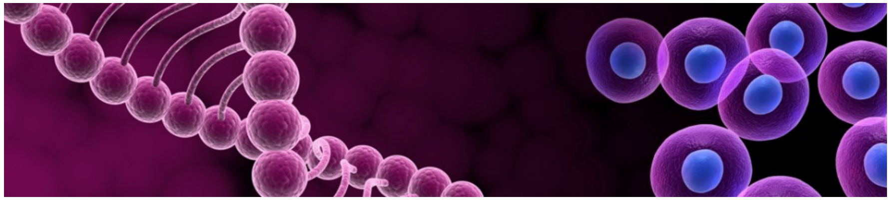

For each research:

- Computations were done using the CRS4 HPC facility
- System was tuned by HPCN group according to the needs of:
  - Number of cores
  - Amount of memory RAM
  - Storage space
  - Low/medium latency connection
  - OS and software to install
  - Custom software environment
  - License management
  - Useful tools to monitor software/hardware performance

OPERATIVE CHAIN FOR NUMERICAL METEOROLOGICAL FORECAST

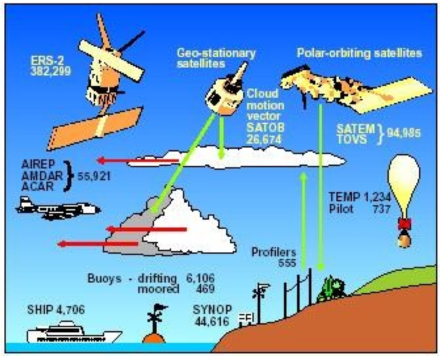

Analysis

Every 6 hours

Global Model

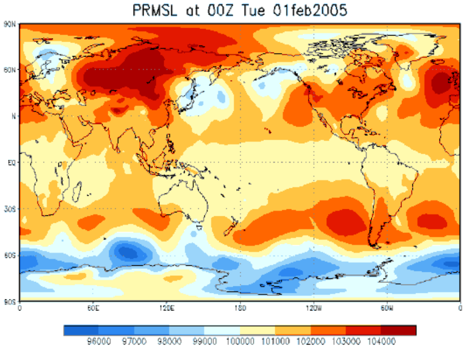

LAM Nesting

Limited Area Model

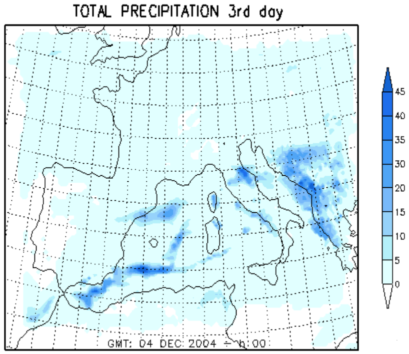

Local data

Assimilation

- MOLOCH on HPC Custer
- 8 core in 9 hours
- 128 cores in 45 min
- Excellent scaling!

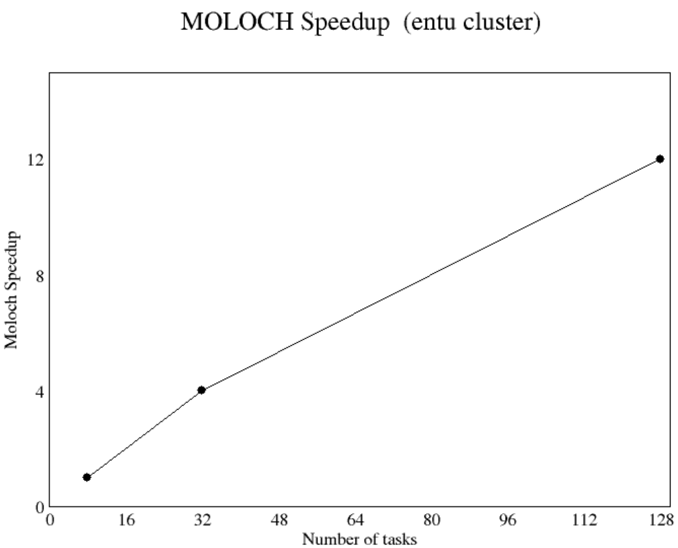

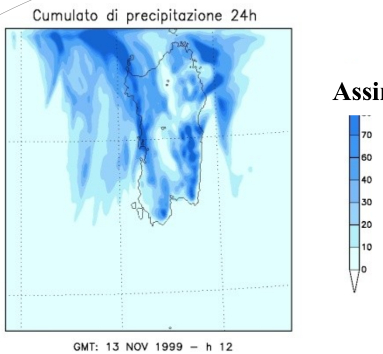

Assimilation

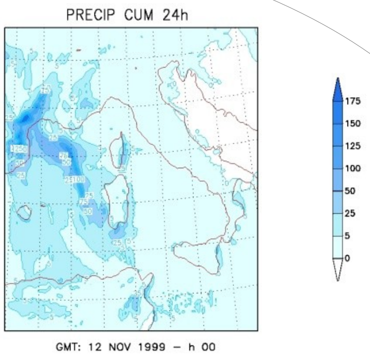

MOLOCH

BOLAM

Models running @ CRS4

Computing portal for near surface imaging and remote collaboration in geoscience and archeology

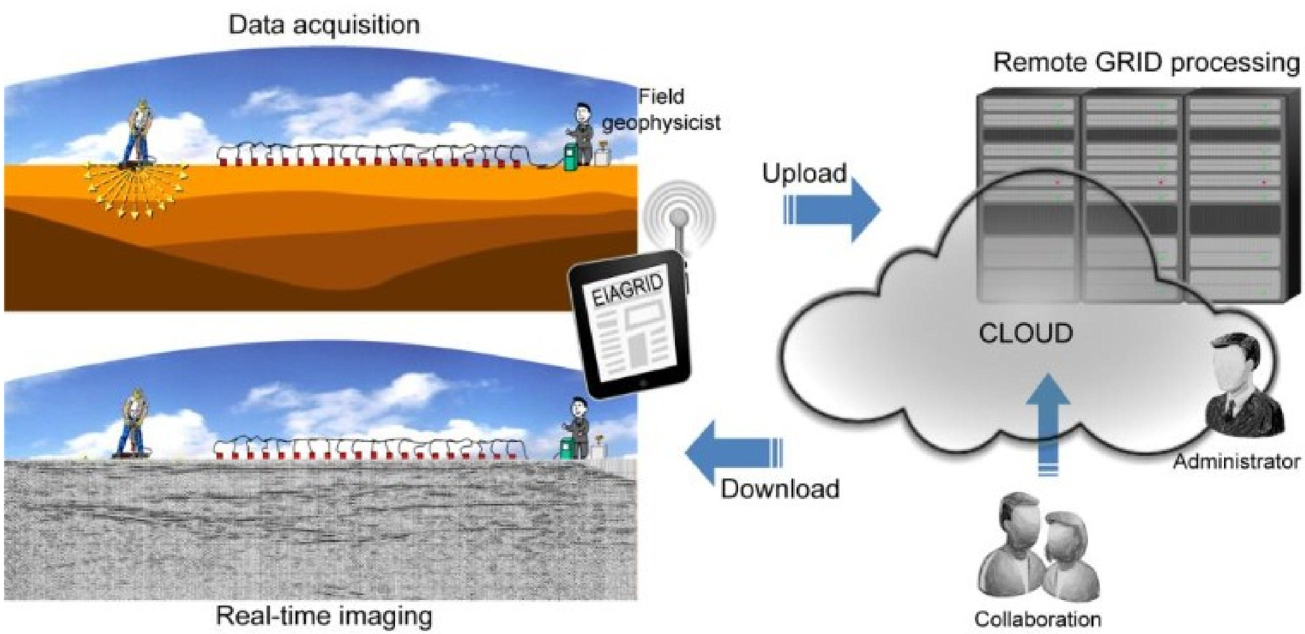

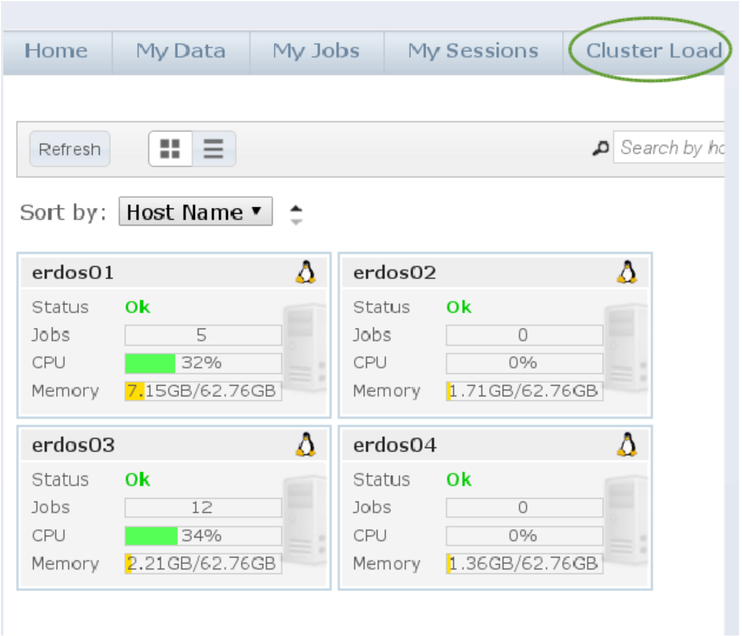

- The Portal allows location independent access to CRS4 HCP infrastructure
- Parallel computing on remote Cluster permits the immediately creation of high quality images
- Remote visualization is done via VNC or Remote Desktop

Solar field optimization in solar tower systems

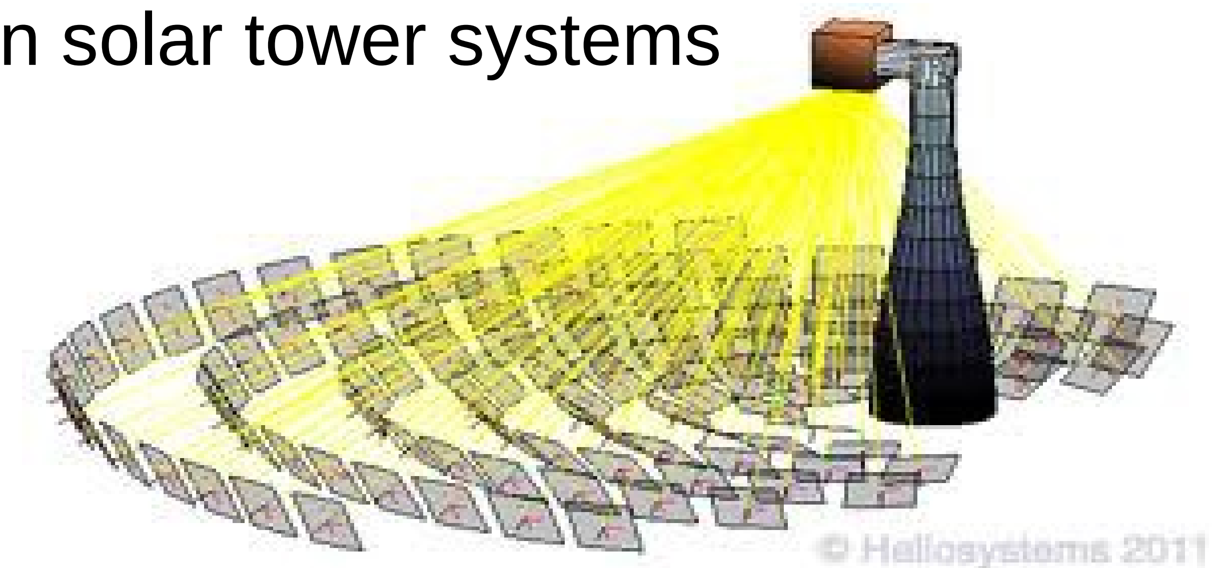

•Three main components: ground heliostats, a tower, and a central receiver at the top of the tower

•The heliostats captures solar radiation from the sun and re-direct it to the central receiver

•Heliostat is moved following the sun in its daily movements

•The heliostats should be as close as possible to the tower, have a very high reflectivity (>95%), and interfere each other as less as possible

•In this way, maximum amount of solar radiation can reach the receiver, where it is converted into heat of a transfer fluid and then into electrical energy

•The optimization of system implies the calculation of the path of each solar ray from the sun to the receiver

Solar field optimization in solar tower systems

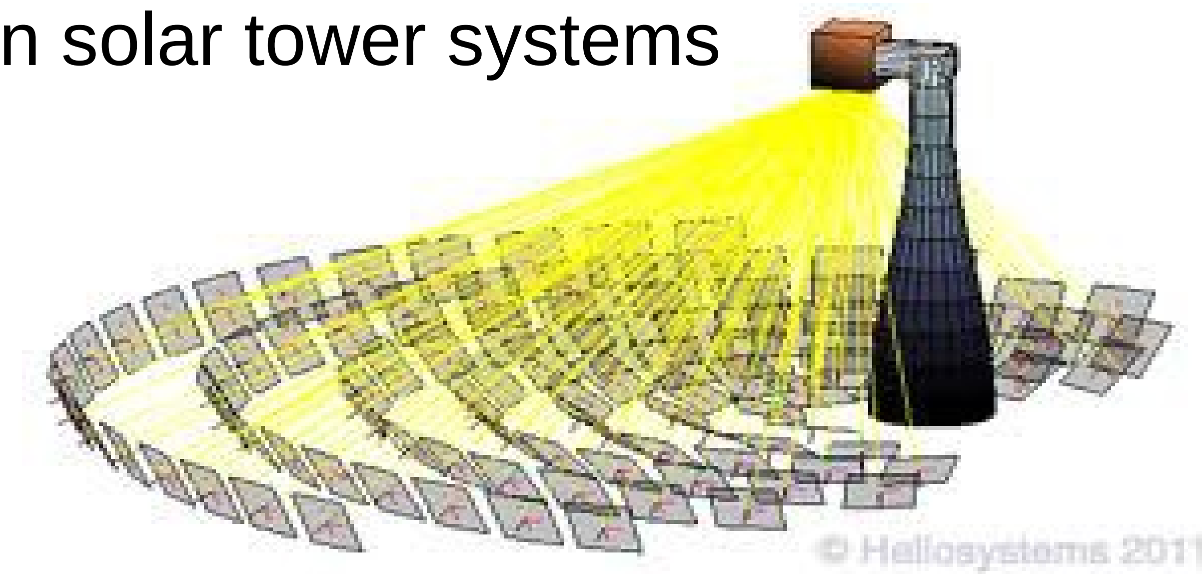

•Some numbers:

- N. of heliostats about 10000
- N. of rays composing the solar radiation: about 50
- N. of considered solar coordinates: 1000
- Thus, the path of about 20×10^9 solar rays must be computed!

- The software CRS4-2 (CRS4 Research Software for Central Receiver Solar System Simulations), entirely developed in our laboratories, performs such simulations on a cluster machine.

- Typically, CPU time is of order 2-3 weeks, using about 300 cluster cores

Molecular Dynamic (MD) @CRS4

NAMD Scalable Molecular Dynamics

Nanoscale Molecular Dynamics program

- Nanoscale simulation of Biological systems (Protein, DNA, Lipids, small molecules)

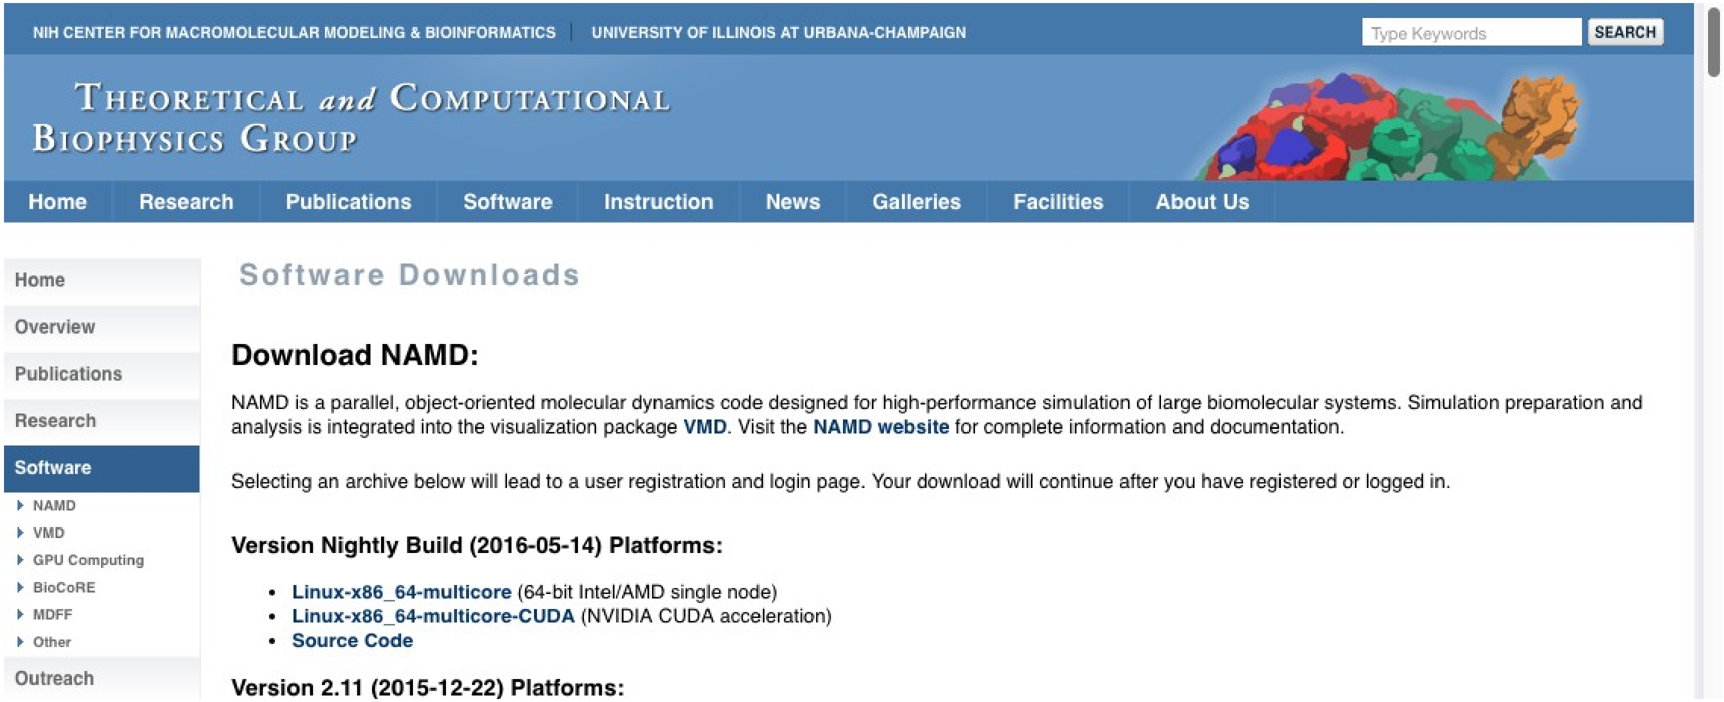

NIH CENTER FOR MACROMOLECULAR MODELING & BIOINFORMATICS | UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Type Keywords SEARCH

THEORETICAL and COMPUTATIONAL BIOPHYSICS GROUP

Home | Research | Publications | Software | Instruction | News | Galleries | Facilities | About Us

Home
Overview
Publications
Research
Software
- NAMD
- VMD
- GPU Computing
- BioCoRE
- MDFF
- Other

Outreach

Software Downloads

Download NAMD:

NAMD is a parallel, object-oriented molecular dynamics code designed for high-performance simulation of large biomolecular systems. Simulation preparation and analysis is integrated into the visualization package **VMD**. Visit the **NAMD website** for complete information and documentation.

Selecting an archive below will lead to a user registration and login page. Your download will continue after you have registered or logged in.

Version Nightly Build (2016-05-14) Platforms:

- **Linux-x86_64-multicore** (64-bit Intel/AMD single node)
- **Linux-x86_64-multicore-CUDA** (NVIDIA CUDA acceleration)
- **Source Code**

Version 2.11 (2015-12-22) Platforms:

NAMD Scalable Molecular Dynamics

Molecular Dynamic (MD) @CRS4

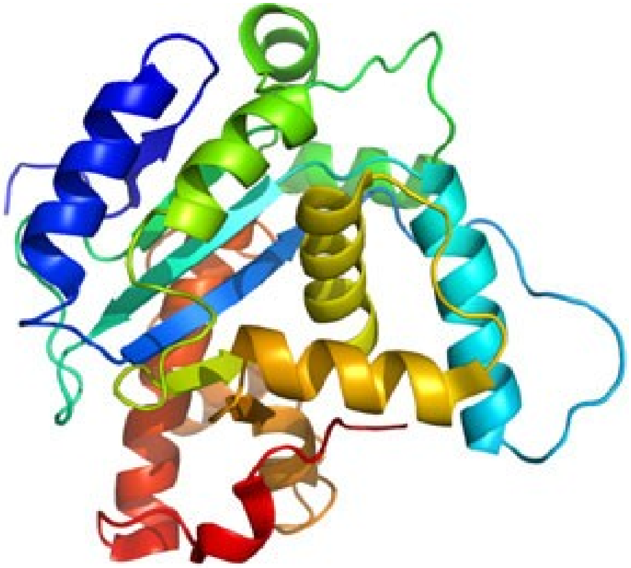

Protein structure 3d

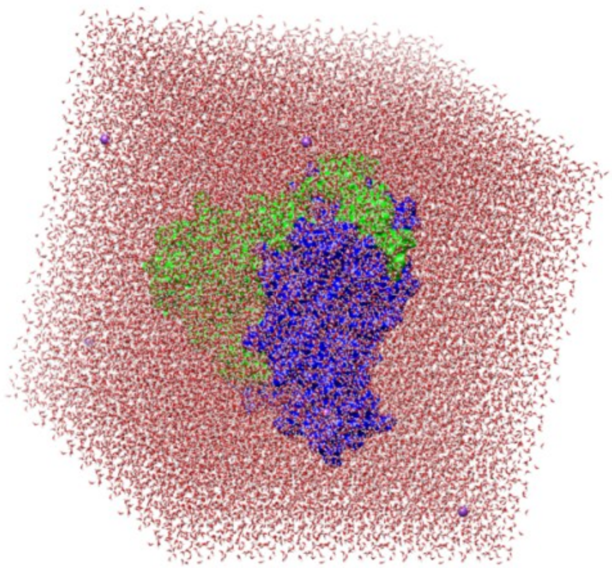

MD: Protein enclosed in a box of water and ions

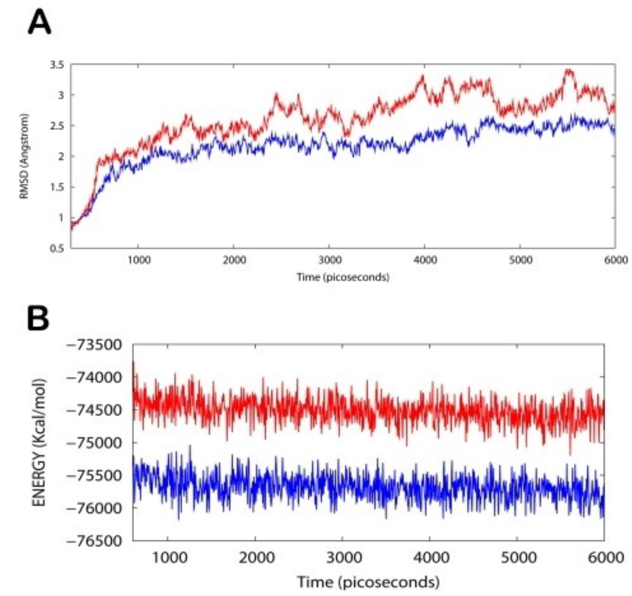

Analysis

- Simulations have been done with NAMD on 128 nodes of HPC cluster
- A run produces about 50-100 ns: the duration of the simulation depends on the size of the system protein and box of water (a dynamic average of 100 ns lasts 8 days)

NAMD Scalable Molecular Dynamics @CRS4

Study protein-peptide interactions relevant to Multiple Sclerosis Disease

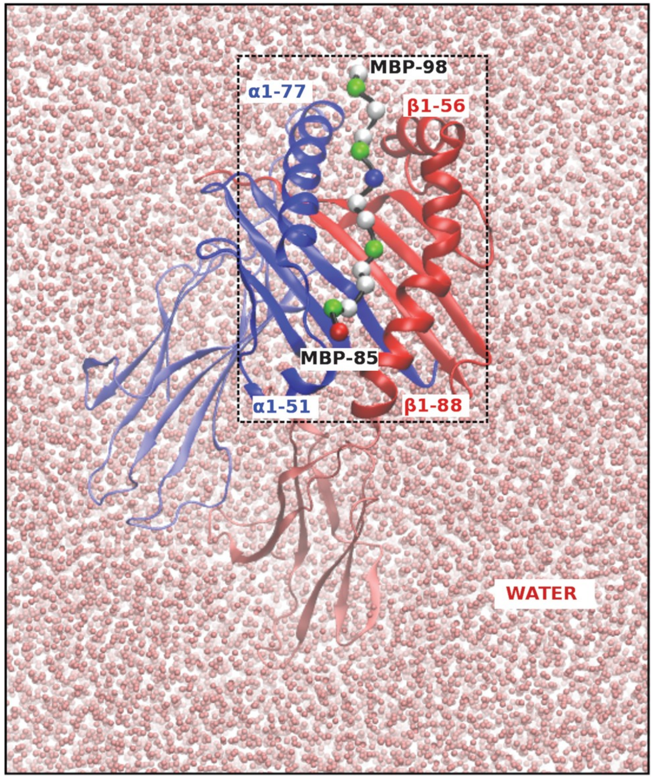

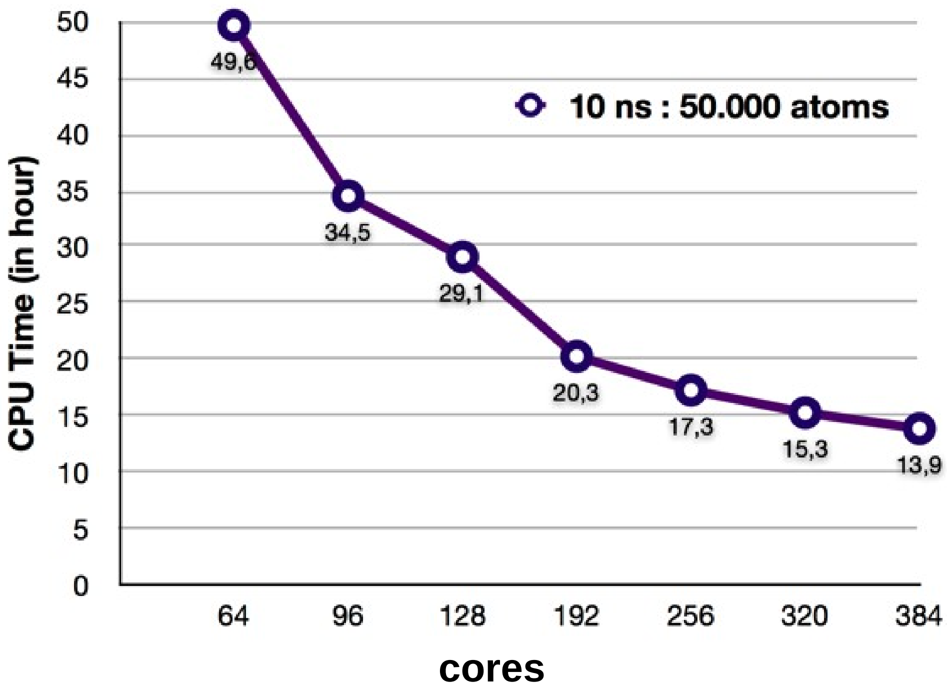

NAMD Scalable Molecular Dynamics @CRS4

Excellent Scaling up to 512 cores, for large biological system

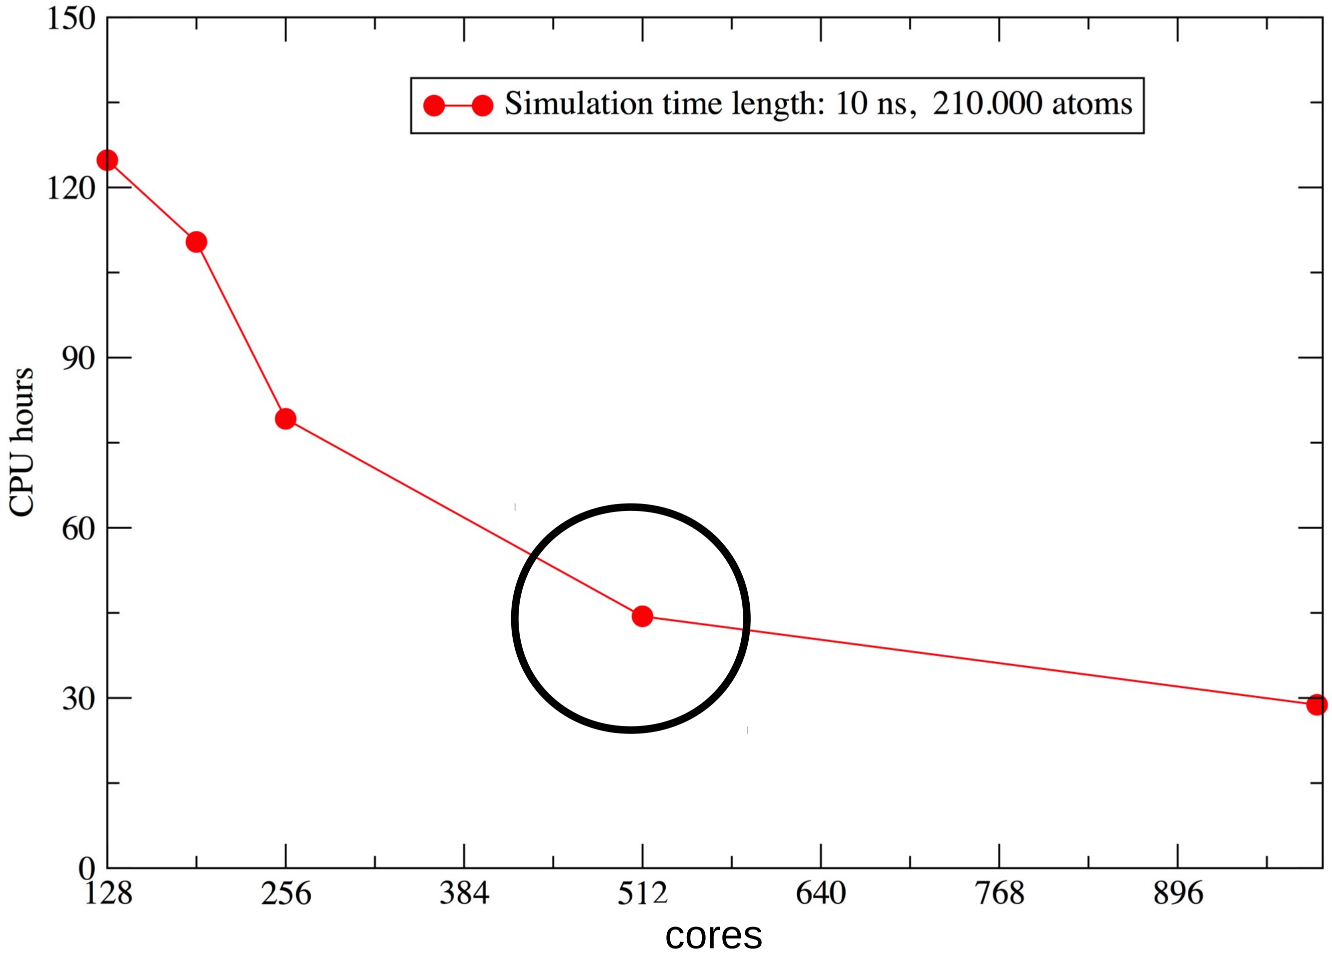

- Total no. of atoms: **210.000**
- Simulation time length: 10 ns

CFD Simulation of the Control Rod Emergency Insertion in the MYRRHA Nuclear Facility

MYRRHA = **M**ulti-purpose h**Y**brid **R**esearch **R**eactor for **H**igh-tech **A**pplications

- Normal operation: reactivity control function
- Emergency operation: safety function, insertion time < 1sec

Computation @ CRS4:

- Average mesh is rapresented by 10 mln cells
- Mesh is generated in about 30 min using 8 Intel cores
- Use of STARCCM+11 and Infiniband FDR connection
- Control rod insertion: a physical time of 0.6 sec is simulated in 6 hours on 400 Intel cores

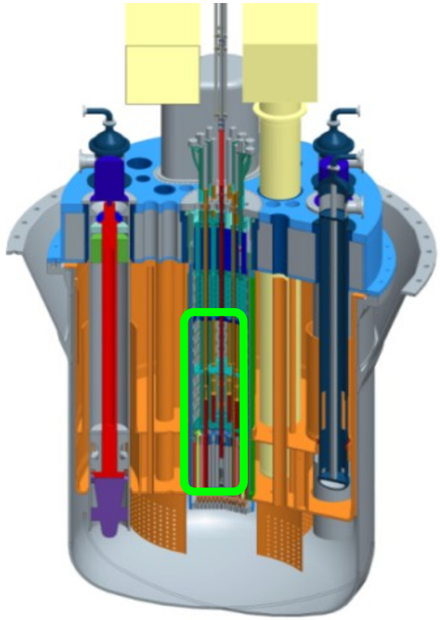

HPC research at CRS4

LUNA ROSSA

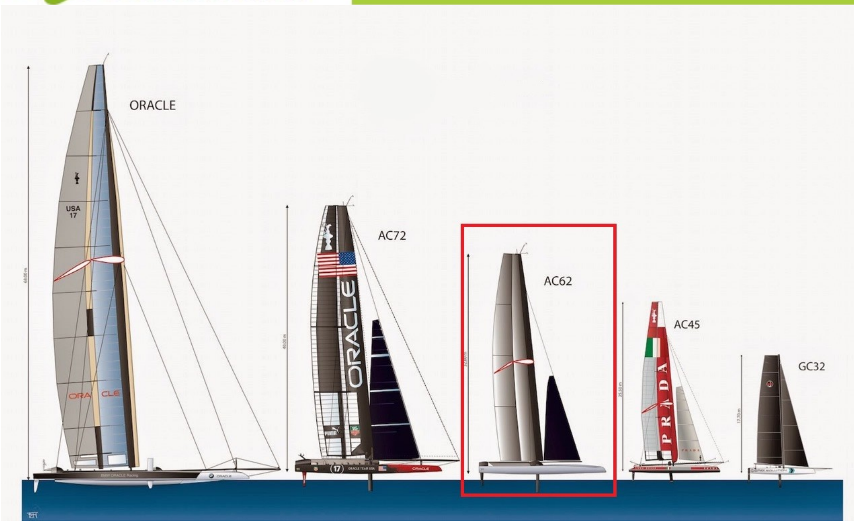

Hull length	19.0 m
Maximum beam	11.8 m
Wing height	29.9 m
Draft	3.0 m
Weight	4300 kg
Sail area:	
Wing	175 m²
Jib	60 m²
Estimated maximum speed:	
Windward	30 kn
Downwind	50 kn
Crew	8

Figure1: AC62 yacht

- CRS4 was official supplier of Luna Rossa Challenge for the preparation of the 35th America's Cup, providing HPC resources and technological support to the team
- Cagliari was chosen as the base of Team
- The computing resources were mainly used for the hydrodynamic and aerodynamic studies and development of the AC62 and AC45 yacht
- The development of the boats was done in the "virtual towing tank" and "virtual wind tunnel" through CFD (computational fluid dynamic) STARCCM+ v9.04.009 software
- Infiniband FDR connection was used

HPC research at CRS4

LUNA ROSSA

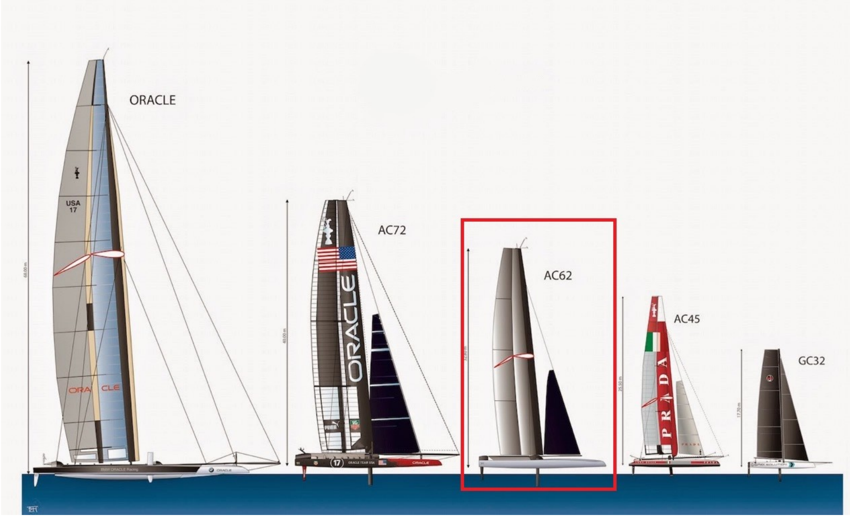

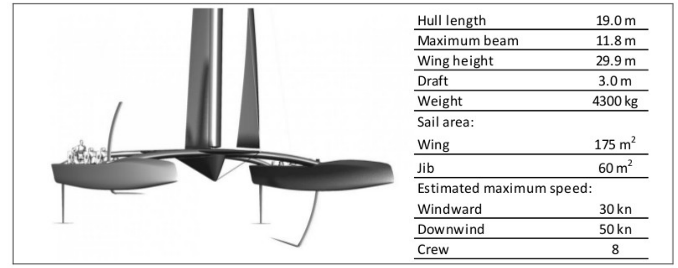

Figure1: AC62 yacht

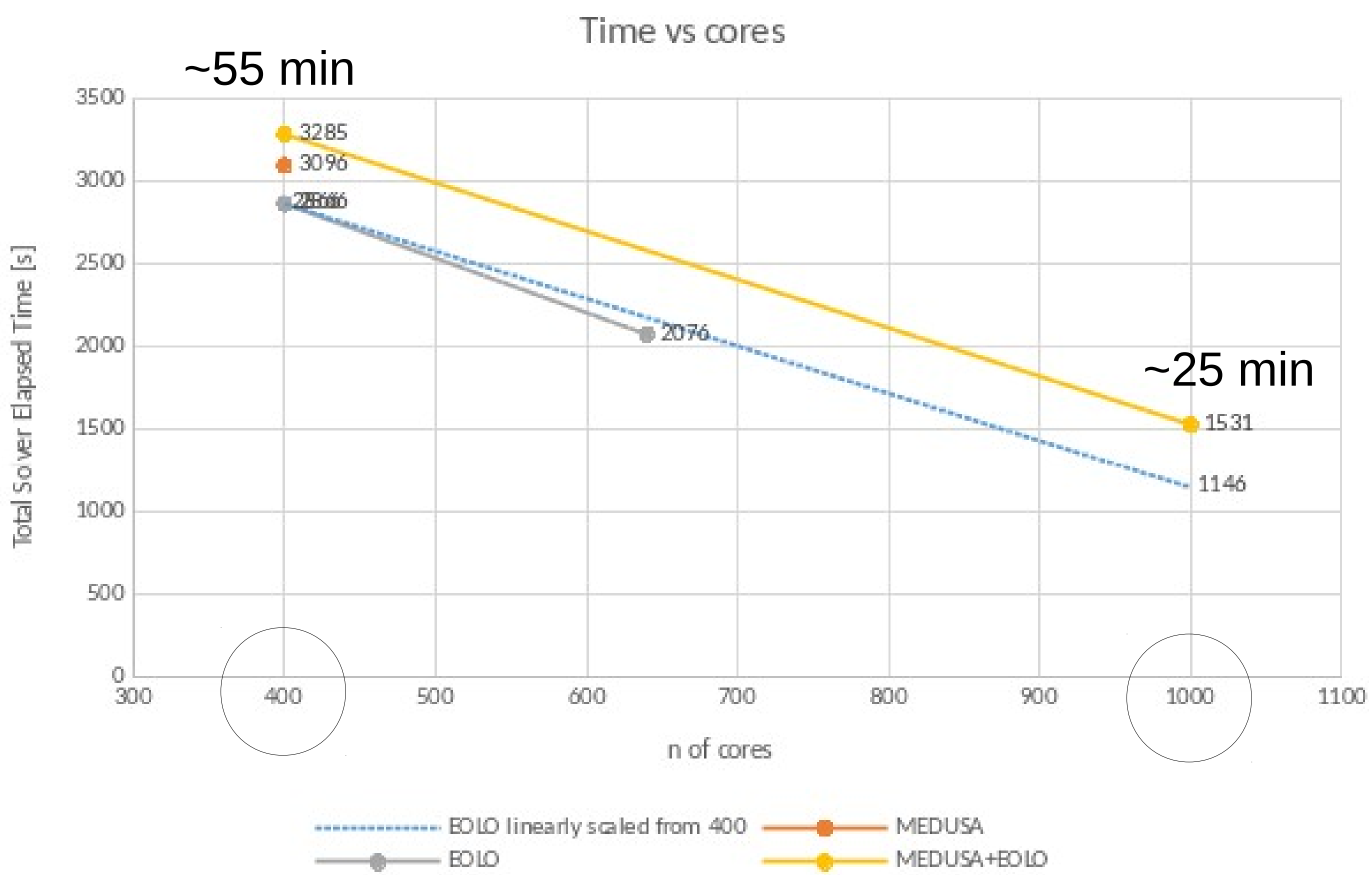

HPC research at CRS4

L U N A R O S S A

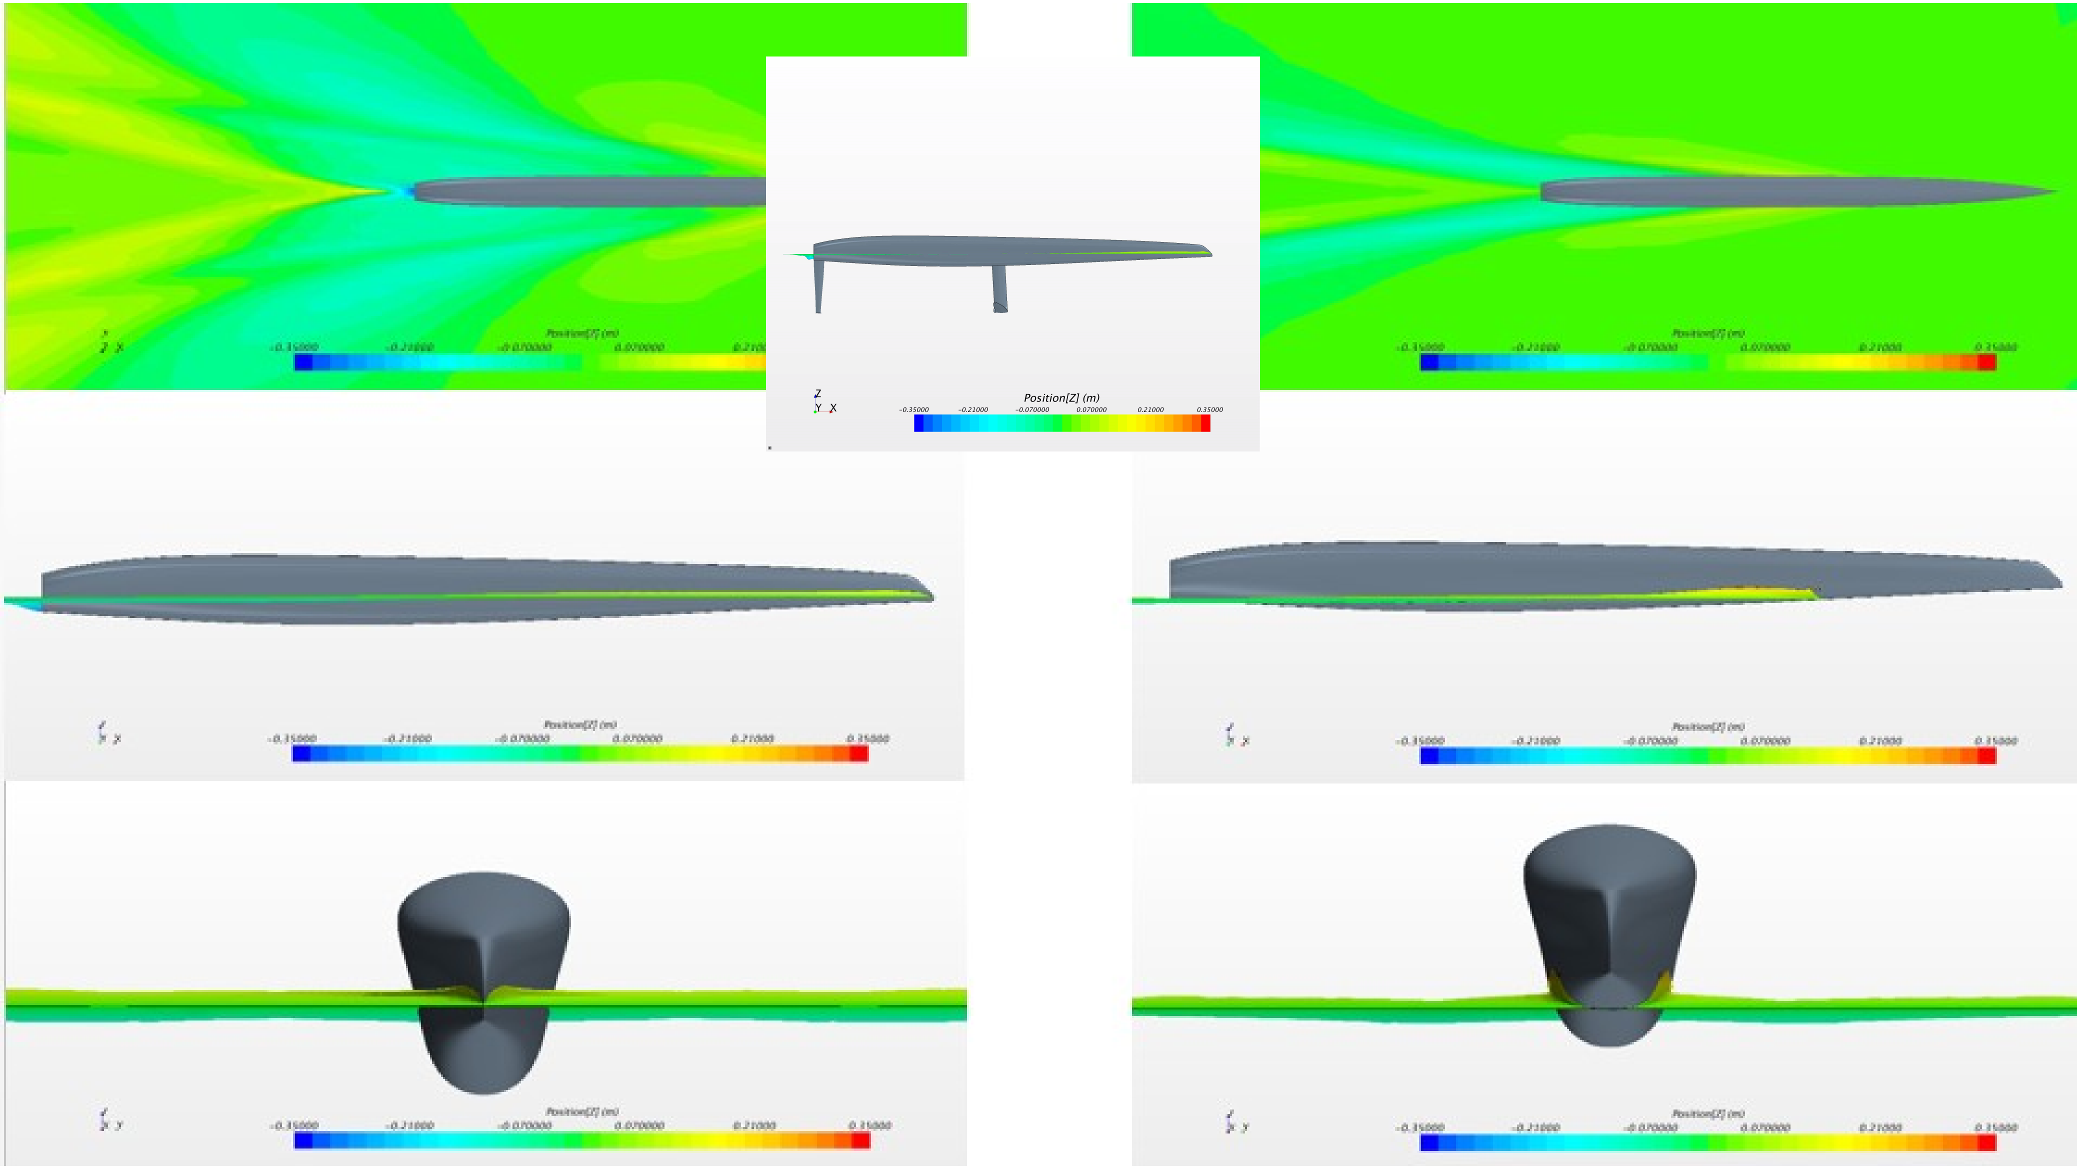

BOTIN PARTNERS
NAVAL ARCHITECTURE
botinpartners.com

SAILING YACHT PHYSICS / COMPUTATIONAL FLUID DYNAMICS

- Hydrodynamic and aerodynamic resistance computations on racing yacht hulls, appendages and sails

Cluster In house – Colocation service

32 nodes HP sl380s gen8+1 node dl380p gen8
14,5 Tflops
4,4 TB RAM
20 TB disks
656 Intel Xeon E5-2680 v2 @ 2.80GHz cores
66 1Gbps ports
33 Infiniband FDR 56Gbps low latency ports
Redundant power supplies
Redundant link network to the data center resources

Hewlett Packard Enterprise

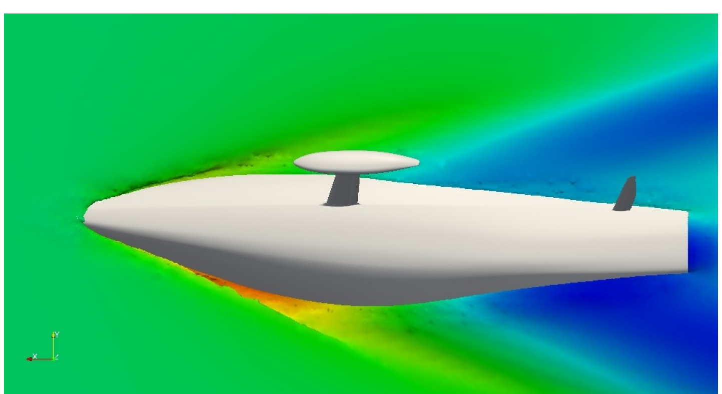

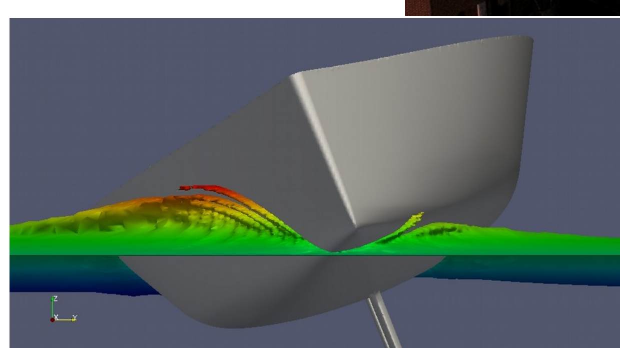

A.Carrau – Botin Partners

HPC research at CRS4

BOTIN PARTNERS
NAVAL ARCHITECTURE
botinpartners.com

SAILING YACHT PHYSICS / COMPUTATIONAL FLUID DYNAMICS

- OpenFoam CFD software and a custom solver based on OPF libraries are used
- Meshes may be composed by 3.000.000 up to 50,60 million of hexahedral elements
- Depending on the number of elements 100, 200 or 600 cores are needed
- Simulation runs takes from 18 to 28 hours
- Computation without Infiniband (low latency connection) is 3 times slower
- Aerodynamic simulations use 60mln of elements and 600 cores taking around 24hours
- Hydrodynamic simulations use 3-6 mln of elements and 100 cores in 18-24 hours
- Postprocessing is done using Paraview and custom software

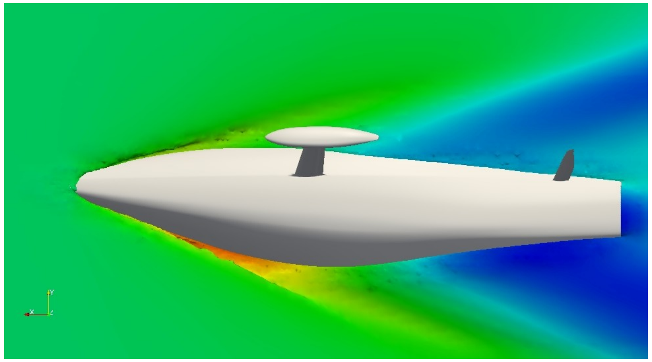

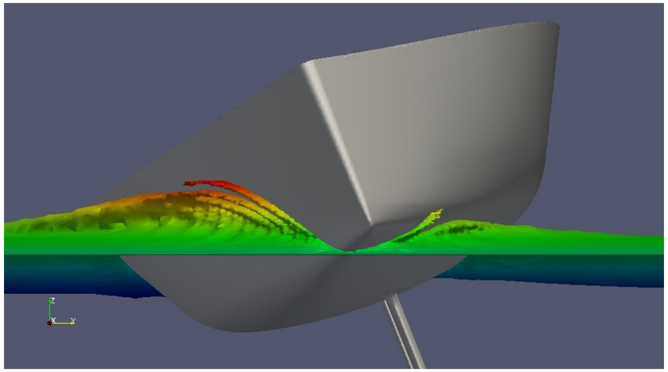

A.Carrau – Botin Partners

CRS4 Next Generation Sequencing Lab

illumina

- Currently the **largest** sequencing center in Italy
- The infrastructure, fully automated, unique in Italy and among the few in Europe, enables CRS4 to perform large scale sequencing projects, from the biological sample to complete analysis
- Equipment:
  - 2 Illumina HiSeq2000
  - 1 Illumina HiSeq2500
  - 1 Illumina HiSeq3000
  - 1 Illumina MiSeq

- Capacity: about 10 Tbases/month
- NGS applications: Whole DNAseq, exome seq, RNAseq, ChIP-Seq
- Includes also 2 Affymetrix Genechip 3000 for microarray gene expression and genotyping

CRS4 Next Generation Sequencing Lab

The sequencing platform is directly interconnected to CRS4 HPC resources, equipped with over 3000 cores and a 4.5PB storage system

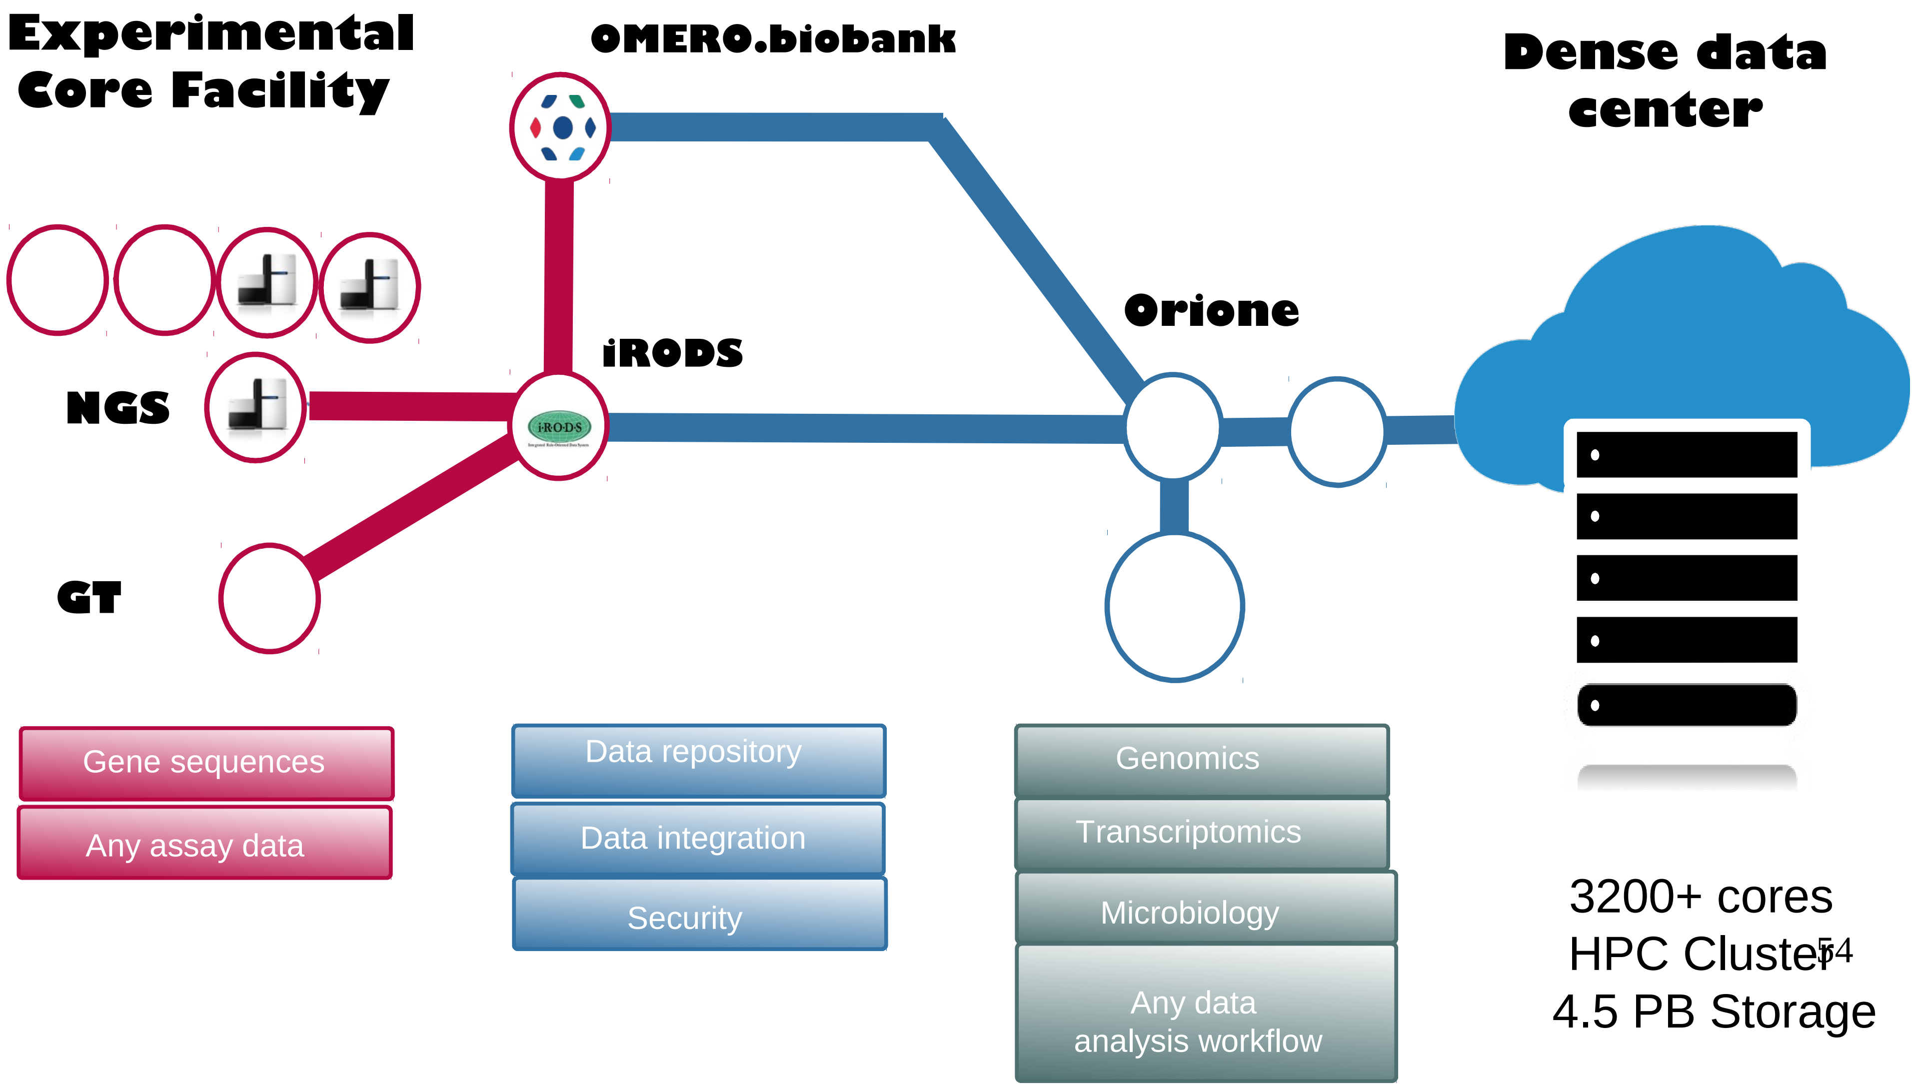

3200+ cores
HPC Cluster
4.5 PB Storage

RGB -CNR @ CRS4 Next Generation Sequencing Lab

Istituto di Ricerca Genetica e Biomedica

Determine the genetic basis of human pathologies and qualitative traits

DATA PRE-PROCESSING ANALYSIS

Sequencing of the DNA (3,000 Sardinians) and the RNA (2,000 samples) of individuals using the Next-Generation Sequencing machines of the CRS4 facility

Clinical info:
- Type 1 Diabetes Database (2,500 cases and relative controls)
- Multiple Sclerosis Database (3,500 cases and relative controls)
- ProgeNIA database (7000 samples)

Genome-wide association study (GWAS)

Correlation of DNA, RNA clinical information and diseases to determine genetic basis of human features and pathologies

Genomics and Evolution

Genetics and Diseases

Genetics and Phenotypes

Marongiu Pala Sanna Sidore Steri @ CNR

RGB -CNR @ CRS4 Next Generation Sequencing Lab

Istituto di Ricerca Genetica e Biomedica

CLUSTER USAGE

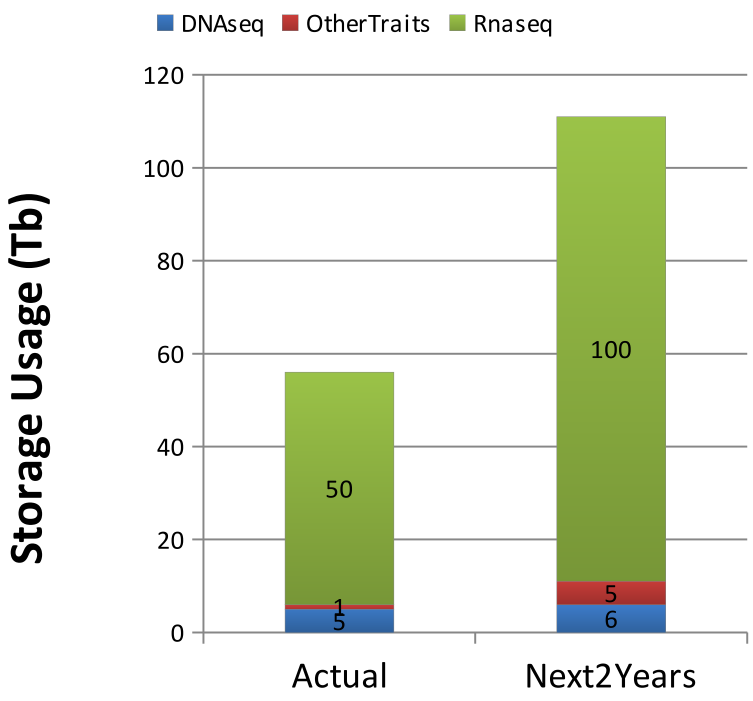

Requirement for a single analysis:

Most common computation burdens
25 CPU days
173 CPU days

Most common RAM burdens
2 Gb/job
8 Gb/job

Seldom RAM burdens
30 Gb/job
64 Gb/job

For more information about

research projects people

@CRS4

visit

http://www.crs4.it