Previsioni uso rete in CSN4

Roberto.Alfieri@pr.infn.it

Applicazioni rilevanti per la rete

Le applicazioni in fisica teorica che hanno un impatto rilevante sula rete sono le **simulazioni su reticolo** (elaborate su sistemi HPC) quali

Lattice QCD, Cosmologia Computazionale, Fluidodinamica Computazionale,..

Modello computazionale:

- Elaborazione **in centri HPC** di configuazioni (**data-set) di notevoli dimensioni (anche decine di GB)** in Input e Output
- I dati vengono generalmente ospitati dai centri HPC per il tempo necessario per l'elaborazione e il trasferiemento e vengono mantenuti a **casa dell'utente** o condivisi in **repository specifici**

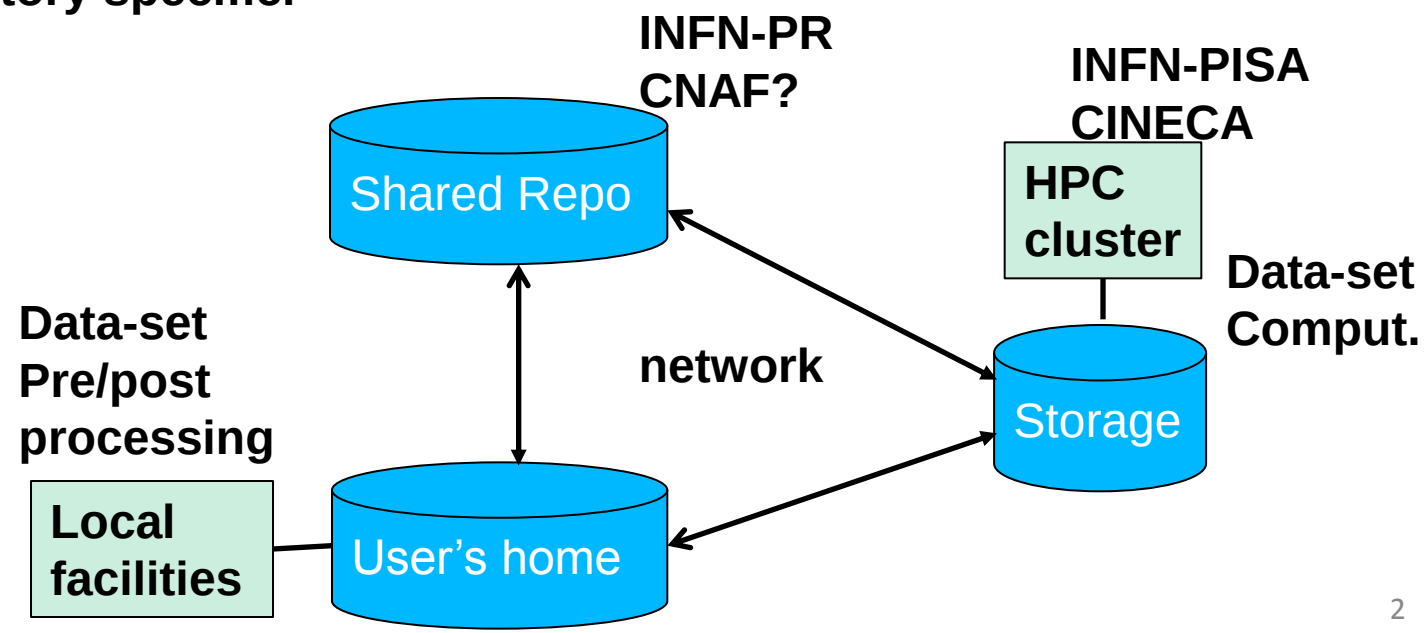

Calcolo HPC medio ZEFIRO (INFN-Pisa)

Calcolo medio (+ fast lane calcolo high end) , triennio 2014-2018

Finanziamento: SUMA (140kE), CSN4 (130KE + 20KE anno), CCR (contr. Switch IB)

THEOCLUSTER ZEFIRO

1 nodo 64 cores AMD, 512GBram
2K cores, 20 Tflops
Accesso via AAI

Zefiro è il theocluster (cluster per il calcolo di fisica teorica) di nuova installazione relativo al progetto SUMA SUMA Project. E' costituito da 32 macchine ciascuna con 512 Gb di ram e 4 processori (ciascuno da 16 core); conta quindi 2048 core di calcolo AMD Opteron 6380 (2,5GHz) e collegamento Infiniband QDR gestiti da uno switch Mellanox IS5100 a 108 porte. L'utilizzo previsto è quello del calcolo locale e per l'accesso ad esso è stata dedicata una User Interface specifica. Gli accessi sono regolati tramite lo scheduler LSF nella versione 9.

ACCESSO AL CLUSTER
USO DEL CLUSTER

STORAGE AREA RISERVATA AI FISICI TEORICI

---

english version

Zefiro is a cluster (for theoretical physics) of the SUMA project SUMA Project, It's made up of 32 computer each with 512 GB of RAM, 4 processors (each with 16 cores). then there will be a total of 2048 computing cores AMD Opteron 6380 (2,5GHz) and Infiniband connection QDR managed by a switch Mellanox IS5025 with 36 ports. It's provided a use for local calculation ane was made a specific User Interface. The accesses are regulated by the scheduler LSF version 9.

ACCESS
UTILIZATION

RESERVED STORAGE AREA

Calcolo HPC medio
ZEFIRO: utilizzo

AcctGroup	Total Jobs	Succ Jobs	Succ Rate(%)	Slots	Walltime (sec)	CPU Time (sec)	CPU Eff(%)	Walltime Share(%)	Avg Wait (sec)
npqcd	20067	9974	49.70	8030	45337621	40226169	88.73	22.71	2100
teorici	3925	1908	48.61	3442	33367920	31563478	94.59	16.72	6240
sissa	956	407	42.57	782	27755719	21781231	78.47	13.90	9064
qcdlat	808	19	2.35	17	24508573	23845276	97.29	12.28	37850
fluent	3791	1647	43.45	7	17206115	14862382	86.38	8.62	377
aae	307	267	86.97	280	14357194	11559639	80.51	7.19	17096
fieldturb	441	147	33.33	340	11737982	9663026	82.32	5.88	10393
tasp	525	450	85.71	493	7522945	6493703	86.32	3.77	1841
ninpha	208	71	34.13	75	5949714	3650891	61.36	2.98	4178
indark	646	304	47.06	29	2664321	2460504	92.35	1.33	758
manybody	231	61	26.41	100	2651078	2560805	96.59	1.33	1962
ftecp	164	78	47.56	60	1588642	1116666	70.29	0.80	36667
thdevel	821	258	31.43	52	1554245	1249619	80.40	0.78	6275
abacus	157	21	13.38	34	1459943	1173224	80.36	0.73	1060
biophys	94	73	77.66	93	926500	926068	99.95	0.46	789
wsip	60	20	33.33	34	880764	743240	84.39	0.44	3253
teongrav	61	9	14.75	29	130339	95570	73.32	0.07	1598
cms	73	47	64.38	22	13815	1612	11.67	0.01	1524
staff	1	1	100.00	1	224	86	38.39	0.00	5

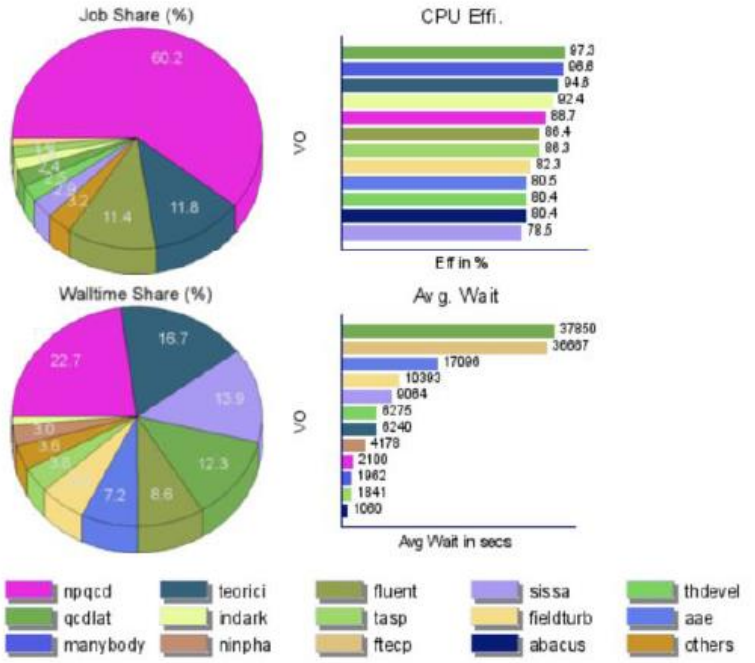

APR-GIU 2014

13 IS + SISSA -> 3.2 Mcore-h su 3.8 disponibili (Circa 85% del tempo macchina)
Disco Home 18 TB saturati. Upgrade a breve a 35 TB

Calcolo HPC high-end
FERMI (CINECA)

IBM BG/Q FERMI

FERMI, the new Italian HPC system

This new system, based on IBM Blue Gene Q architecture, replaced the IBM SP6 supercomputer and was installed and became fully operational in summer 2012.

- FERMI is composed of 10.240 PowerA2 sockets running at 1.6GHz, with 16 cores each, totaling 163.840 compute cores and a system peak performance of 2.1 PFlops.
- Each processor comes with 16 Gbyte of RAM (1Gbyte per core).
- The BG/Q system is equipped with a performant scratch storage system with a capacity of 2.6 Pbyte and a bandwidth in excess of 100 GByte/s.

Access to FERMI is on the basis of national and European calls for proposals. ISCRA (Italian SuperComputing Resources Allocation) and PRACE manage the access to the Tier-0 supercomputer by the way of international peer-review procedures ensuring world-class research.

Architecture: 10 BGQ Frames
Model: IBM-BG/Q
Processor Type: IBM PowerA2, 1.6 GHz
Computing Cores: 163840
Computing Nodes: 10240
RAM: 1GByte / core
Internal Network: Network interface with 11 links -> 5D Torus
Disk Space: 2.6 PByte of scratch space
Peak Performance: 2PFlop/s

160 Kcores
IBM PowerA2
1.7 Pflops

Accesso:
- **Call PRACE (2 all'anno, circa 1 Gcore-hours per call, procedura complessa)**
- **Convenzione INFN/CINECA 2013 e 2014**

Calcolo HPC Hig-end
EURORA (CINECA)

EURORA

EURORA, the Cineca-Eurotech PRACE prototype, record for energy efficiency

Eurora is the new prototype system installed at Cineca, **Eurora** is the world's most green supercomputer: **1st in the Green500 List of June 2013**. The its energy efficiency improves upon the previous greenest supercomputer in the world by nearly 30%.

Each one of the 64 Eurora nodes is equipped with two Intel Xeon Sandybridge processors and either two NVIDIA K20s accelerators or two intel Intel Xeon Phi 5120D co-processors and, based on a hot water cooling system, is capable of delivering 3,150 megaflops per watt of sustained performance.

Funded by PRACE 2IP, Eurora is the result of the collaboration between Eurotech and Cineca and, equipped with NVIDIA accelerators and Intel Xeon PHI co-processors, the prototype is now fully functional and accessible to the members of PRACE and to the major Italian research entities.

Architecture: 1 rack
Model: Eurora prototype
Processor Type:

- Intel Xeon CPU E5-2658 @ 2.10GHz
- Intel Xeon CPU E5-2687W @ 3.10GHz

Accelerator Type:

- Nvidia Tesla K20s
- Intel Xeon-Phi 5120D

Computing Cores: 1024
Computing Nodes: 64
RAM: 16GByte DDR3 1600MHz per node (5 nodes with 32 GByte)
Internal Network:

- 1 FPGA (Altera Stratix V) per node
- IB QDR interconnect
- 3D Torus interconnect

Disk Space: 90 GByte SSD per node

Sustained Performance: 3,150MFlop/w

1 Kcores Intel +
64 GPU +
64 PHI

150 TFlops

Calcolo HPC Hig-end CINECA: utilizzo

Convenzione INFN/CINECA (Fermi + Aurora)
-- 2013: 150 Mcore-hours
-- 2014: 100 Mcore-hours

IS	assegnate (Mcore-hours)	consumate (Mcore-hours)	consumate (%)
INF14_biophys	3.0	1.559	0.1
INF14_fbs	2.0	90.695	4.5
INF14_fldturb	12.0	224.264	1.9
INF14_ftecp	8.0	694.615	8.7
INF14_lqcd123	17.5	2.921.737	16.7
INF14_manybody	1.0	0	0
INF14_npqcd	18.0	7.846.216	43.6
INF14_qcdlat	24.5	3.967.021	16.2
INF14_teongrav	14.0	98.335	0.7
	100.0	15.844.446	15.8

situazione consumo primo trimestre 2014

ILDG : LQCD repository

ACCELERATORS | PHOTON SCIENCE | PARTICLE PHYSICS

Deutsches Elektronen-Synchrotron
A Research Centre of the Helmholtz Association

Google™ Custom Search

DESY HOME | RESEARCH | NEWS | ABOUT DESY | CONTACT

Home / LDG

LDG

Latfor DataGrid (LDG)

The International Lattice DataGrid (ILDG) was started with the aim of making gauge field configurations available to an international group of scientists using Grid technologies. Each configuration submitted to ILDG will consist of a set of meta-data and a set of binary files. The meta-data will be made available as XML documents which conform to the XML schema currently being developed by ILDG's meta-data working group. The meta-data will be searchable via the planned ILDG web interface. It contains an URI referring to the configurations that will be made available within a Grid-of-Grids.

The Latfor DataGrid (LDG) is the regional grid for continental Europe. The following institutions are operating the LDG services:

- DESY (Hamburg and Zeuthen, Germany)
- IN2P3 (Lyon, France)
- INFN (Parma, Italy)
- JSC (Jülich, Germany)
- RUG (Groningen, Netherlands)
- SARA (Amsterdam, Netherlands)

Status

For current status of ILDG services see Inca monitor at CSSM.

Server nazionale a INFN-PR

Tecnologia Grid (Storm, 17 TB)

In fase di riorganizzazione

Evoluzione del calcolo e future necessita' di rete

CALCOLO HPC medio

- **ZEFIRO** (2 Kcores AMD, 20 TFlops)
- **Nuovo Cluster SISSA** (Set 2014, 5 Kcores Intel, 100TFlops)

I cluster di PISA e SISSA verranno *probabilmente* integrati (Convenzione INFN/SISSA)

CALCOLO HPC High End (CINECA)

– ~~FERMI~~ Dismissione (2015) -- ~~EURORA~~ Dismissione (2016?)

– **Next Tier1 (**Nov 2014, Cofin. CINECA/SUMA/ENI, 1 PFlops, 512 nodi Intel+MIC/GPU)

-- **Next Tier0** (late 2015, order of magnitude 50 PFlops)

STORAGE

ILDG (repository configurazioni LQCD). Server da INFN-PR a CNAF?

Repository dati CINECA al CNAF? (In fase di studio la condivisione GPFS)

NECESSITA' DI RETE (oltre l'utilizzo fisiologico)

Connettivita' veloce verso **CINECA, INFN-Pisa, CNAF**

Link potenziati **CNAF-CINECA** e **PISA-SISSA ?**

Grazie per l’attenzione!