



MARCONI

Installation roadmap

Elda Rossi

Alessandro Grottesi





A new Supercomputer (codename: MARCONI) has been installed at CINECA, available for Italian and European research community.


It is a Lenovo NeXtScale system based on Intel technology, with a final peak performance around 20PFlop/s.

Deployment of Marconi is started July 2016, the complete delivery expected as July 2017.


The first part (2 PFlop/s conventional "Broadwell" based) is in full production as July 2016.

The second part (11 Pflop/s "Knights Landing" based) is in full production as January 2017.

The third part (5+ Pflop/s "SkyLake" based) is in full production as August 2017




A short story



- In 2015 the computing resources in Cineca were:
 - ✦ Tier-0: **FERMI** (acquired in summer 2012)
 - ✦ Tier-1: **GALILEO** (acquired in Jan 2015)
 - ✦ Front-end, Viz, BigData: **PICO** (acquired in Nov 2014)
- FERMI arrived at the expected end of its production.
- Galileo is nearly arrived at the expected end of its production too.
- The Cineca governing bodies, aimed at supporting scientific research, approved a development plan, with an investment of Euro 50 million in two phases, from 2016 to 2020:
 - 2 x 5 \Rightarrow 10 Pflops in 2016-2017
 - 10 x 5 \Rightarrow 50 Pflops in 2019-2020


 **Marconi**

MARCONI: the new Tier-0 system



- A tender was issued in 2015 and assigned Jan 2016 to **lenovo**
- The system will be delivered in three phases:
 - A1: April 2016 (BRD 2 PFs)
 - A2: Sept 2016 (KNL 11 PFs)
 - A3: Aug 2017 (SKL 5+ PFs)
- In total:
 - 18 PFs peak performance
 - 17PB raw storage
 - 3MWatt power consumption

MARCONI: New Tier-0 system



Technical Features:

- Intel based
- Architecture: Lenovo NeXtScale
- Fabric: Intel OmniPath

A2

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


A1

BRD 2x18 cores, 2.3GHz
1500 nodes, 2PFs

A3

SKL 2x20 cores, 2.3 GHz;
1500 nodes, 5 (+2) PFs


Marconi ranking on Top500



A1

64	CINECA Italy	Fermi - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	163,840	1,788.9	2,097.2	822
65	BASF Germany	QURIOSITY - Apollo XL230k, Xeon Gold 6148 20C 2.4GHz, Intel Omni-Path HPE	35,280	1,750.2	2,709.5	1,169
66	CINECA Italy	Marconi Intel Xeon - Lenovo NeXtScale nx360M5, Xeon E5-2697v4 18C 2.3GHz, Omni-Path Lenovo	54,432	1,723.9	2,003.1	1,361
67	Internet Service A China	Inspur TS10000, Xeon E5-2682v4 16C 2.5GHz, 10G Ethernet, NVIDIA Tesla P100 Inspur	39,680	1,714.0	3,149.6	967
68	Government United States	SwiftLucy - Cluster Platform 3000 BL460c Gen9, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR HPE	57,600	1,703.3	2,304.0	2,880
69	Japan Agency for Marine -Earth Science and Technology Japan	Gyoukou - ZettaScaler-2.0 HPC system, Xeon D-1571 16C 1.3GHz, Infiniband EDR, PEZY-SC2 ExaScalar	3,176,000	1,677.1	3,207.7	164

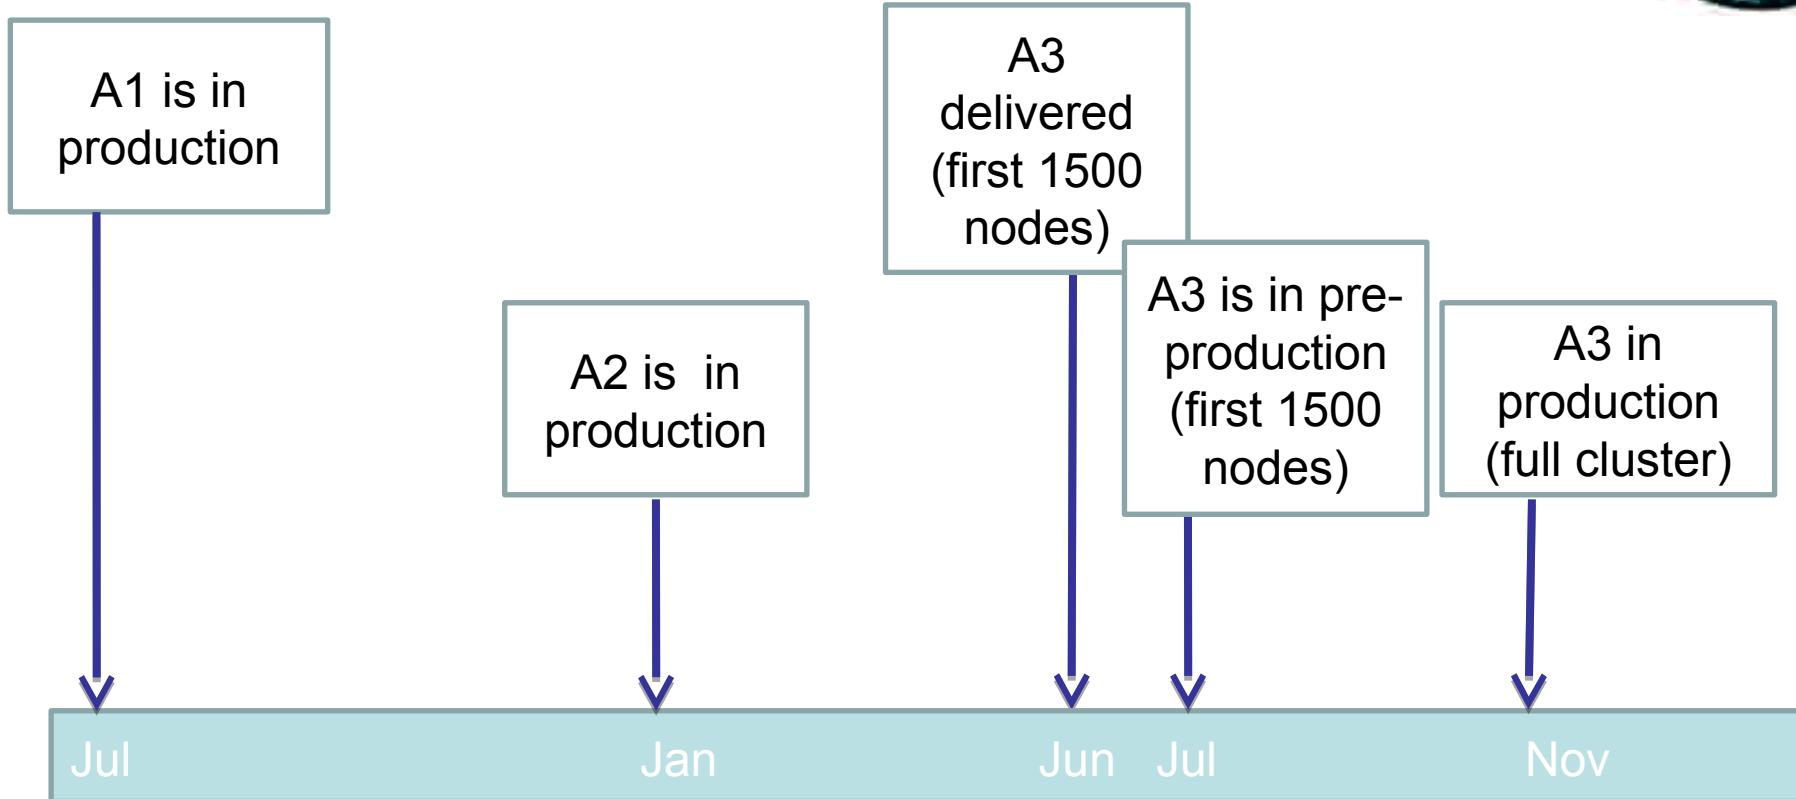
Marconi ranking on Top500




A2

10	DOE/NNSA/LANL/SNL United States	Trinity - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	301,056	8,100.9	11,078.9	4,233
11	United Kingdom Meteorological Office United Kingdom	Cray XC40, Xeon E5-2695v4 18C 2.1GHz, Aries interconnect Cray Inc.	241,920	7,038.9	8,128.5	3,629
12	Texas Advanced Computing Center/Univ. of Texas United States	Stampede2 - PowerEdge C6320P, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path Dell	285,600	6,807.1	12,794.9	1,890
13	Barcelona Supercomputing Center Spain	MareNostrum - Lenovo SD530, Xeon Platinum 8160 24C 2.1GHz, Intel Omni-Path Lenovo	148,176	6,227.2	9,957.4	1,380
14	CINECA Italy	Marconi Intel Xeon Phi - CINECA Cluster, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path Lenovo	241,808	6,223.0	10,833.0	1,600
15	NASA/Ames Research Center/NAS United States	Pleiades - SGI ICE X, Intel Xeon E5-2670/E5-2680v2/E5-2680v3 /E5-2680v4 2.6/2.8/2.5/2.4 GHz, Infiniband FDR HPE	241,108	5,951.6	7,107.1	4,407
16	DOE/SC/Argonne National Laboratory United States	Theta - Cray XC40, Intel Xeon Phi 7230 64C 1.3GHz, Aries interconnect Cray Inc.	231,424	5,884.6	9,627.2	1,087
17	HLRS - Höchstleistungsrechenzentrum Stuttgart Germany	Hazel Hen - Cray XC40, Xeon E5-2680v3 12C 2.5GHz, Aries interconnect Cray Inc.	185,088	5,640.2	7,403.5	3,615
18	King Abdullah University of Science and Technology Saudi Arabia	Shaheen II - Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect Cray Inc.	196,608	5,537.0	7,235.2	2,834
19	Total Exploration Production France	Pangea - SGI ICE X, Xeon Xeon E5-2670/ E5-2680v3 12C	220,800	5,283.1	6,712.3	4,150

Timings





A1 (half reserved to EUROfusion)

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
2PFs	1512	2x Intel Broadwell 18cores @2.3GHz	128 GB	Intel OmniPath 2:1 100Gb/s	21	8 Front End Nodes (2xBDW 18c +128GB RAM)+ MGMT nodes

Core tot: 54.432


Core-h/anno=476.824.320

A2

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
11 PFs	3600	Intel KnightsLanding 68cores @1.4 GHz	96 GB	Intel OmniPath 2:1 100Gb/s	50	Share login nodes with A1

Core tot: 244.800

Core-h/anno=2.144.448.000



A3 (great part reserved to EUROfusion)

Peak Perf.	Comp. Nodes	Socket	RAM/CN	Interconnect	Rack #	Service & Mgmt nodes
5PFs	1512 (+ 1.000)	2x Intel SkyLake 20cores @2.3GHz	192 GB	Intel OmniPath 2:1 100Gb/s	21	Share login nodes with A1 and A2

Core tot: 60.480

Core-h/anno=529.804.800 (500 M core-h/y)



