

# EUROfusion HPC Operations Committee Meeting #61

23 July 2021 – 9:30 - 13:00

## Minutes

<b>1</b>	<b>Attendees .....</b>	<b>1</b>
<b>2</b>	<b>Location .....</b>	<b>1</b>
<b>3</b>	<b>Agenda .....</b>	<b>2</b>
<b>4</b>	<b>Documents distributed for discussion.....</b>	<b>2</b>
<b>5</b>	<b>Summary of the meeting .....</b>	<b>2</b>
5.1	Welcome and introduction .....	2
5.2	General information .....	3
5.3	Marconi-Fusion.....	3
5.3.1	Report for June 2021.....	3
5.3.2	Feedback from Ticket Meetings #54 and update on network(s) status .....	6
5.3.3	Discussion on current issues .....	6
5.4	Gateway .....	6
5.4.1	Report for June 2021.....	6
5.4.2	Discussion on current issues .....	7
5.5	AOB.....	7
<b>6</b>	<b>Next meetings .....</b>	<b>8</b>

### **1**     **Attendees**

Members and experts: Richard Kamendje (EUROfusion chair), David Coster (IPP), Jacques David (CEA), Roman Hatzky (IPP), Leo Ma (CCFE), Norbert Meyer (PSNC), France Boillod-Cerneux (CEA), Gilles Fourestey (EPFL), David Vicente (BSC), Fredric Granberg (Univ. Helsinki), François Robin (CEA secretary)

Guests: Elda Rossi (CINECA), Maria Montagna (CINECA), Francesco Iannone (ENEA)

Excused: Christophe Calvin (CEA), Denis Kalupin (RO for the Work Package Advanced Computing, EUROfusion), Michal Owsiak (PSNC)

### **2**     **Location**

by video

### **3 Agenda**

1. Introduction (R. Kamendje)
  - a. Welcome and confirmation of the agenda
  - b. Review and confirmation of the minutes of the previous meeting (OC-60)
  - c. General information
2. Marconi-Fusion
  - a. Monthly report from ENEA to the Operations Committee (June 2021) (E. Rossi)
  - b. Monthly report of the Ticket Meeting (June/July 2021) (F. Boillod-Cerneux)
  - c. Discussion on current issues
3. Gateway
  - a. Monthly report from ENEA to the Operations Committee (June 2021) (F. Iannone)
  - b. Discussion on current issues
4. AOB
5. Next meetings

### **4 Documents distributed for discussion**

1. Agenda of the meeting
2. Minutes of the previous meeting (OC-60 held on 21 June 2021)
3. Monthly report for May 2021 prepared by E. Rossi (CINECA)
4. Slides prepared for the meeting by R. Kamendje (EUROfusion), F. Robin (CEA), F. Boillod-Cerneux (CEA), Roman Hatzky (IPP), J. David (CEA), F. Iannone (ENEA)

### **5 Summary of the meeting**

*(see slides attached for more details) (actions are highlighted with a ▲)*

#### **5.1 Welcome and introduction**

Richard opens the meeting and welcomes the participants. He explains that the ACHs are now represented in the Operations Committee with one full member per ACH. Therefore, in addition to Roman, new members nominated by the EUROfusion Programme Manager are:

- Michal Owskiak (PSNC)
- Gilles Fourestey (EPFL)
- David Vicente (BSC)

- Fredric Granberg (Univ. Helsinki)

He welcomes the new members and summarizes the role of the Operations Committee and the organization of the meetings of the committee. He mentions also the important role of the former HLST in terms of test, validation, exploration, reporting of errors... and explains that in the future it is expected that all ACHs will contribute to this important activity.

The agenda of the meeting is then approved and the minutes of the previous meeting, OC-60, are approved as well.

## 5.2 General information

Richard addresses two topics:

- Regarding the presentation to the EUROfusion General Assembly of the final report of the HPC expert group (done by François), Richard notes that the presentation was well received and the conclusions were considered as reasonable. At this stage, the remaining unknown is the budget available for this operation.
- Regarding the dialogue between EUROfusion and Nvidia, it was decided based on the first meeting to organize the discussions along two tracks: one track about hardware, one track about software. For the first track a NDA is necessary, not for the second one which is open to all interested people (although the purpose is to discuss, the audience should not be too large). It is noted that a similar contact with other major technology providers (like AMD and Intel) should take place before the launch of the Call for EoI in order to further calibrate the requirements. Name of possible contact persons should be provided to Richard.

## 5.3 Marconi-Fusion

### 5.3.1 Report for June 2021

Elda presents the monthly report for June 2021.

During this month the configuration for EUROfusion, unchanged from the previous month, included 2912 SKL nodes (A3 partition) and 99 Marconi100 nodes (C1 partition).

- Maintenance:
  - Maintenance of Marconi was performed on June 15 (duration 4 hours). Reboot of all nodes and sanity checks were performed before and after maintenance – several slow nodes were identified during the post-maintenance checks and were removed from production. The maintenance scheduled for July 8 was cancelled, therefore, the next maintenance will take place on August 10.
  - Maintenance of Marconi100 was performed on June 29 (duration 4 hours). Sanity checks were performed before and after maintenance – several nodes with slow GPU

were identified during the post-maintenance checks and were removed from production. The maintenance scheduled for July 20 was cancelled, therefore, the next maintenance will take place on August 24.

- No maintenance took place on the Gateway.
- Main events
  - Major incidents: none during the month. However it was noted that several incidents affecting the Marconi100 login nodes occurred possibly in relation with the usage of RedHat 8.1 and NVidia compilers on heavily loaded nodes. The issue was submitted to IBM.
  - Summary of still open system problems: none.
  - Security news: none.
- Availability and usage: the table below summarizes the figures for June 2021 as well as the maintenance time and the number of major incidents for the month.
  - Regarding the A3 partition, the availability and the usage were excellent. It is noted that 114 projects out of 133 already started (usage equal or more than 1% of the allocation). Statistics on the number of nodes used are presented. This month, as usual, the largest usage (25.4%) is made by jobs in the range from 34 to 65 nodes.
  - Regarding the C1 partition, the availability (due to the fact that some nodes had to be removed from production to be fixed) is lower than the previous months and slightly lower than the SLA. Elda mentions that most of the problems come from overheating; CINECA is considering decreasing the frequency of the nodes in order to reduce the number of problems – this point will be discussed during the next OC meeting ▲. The usage (87.0%) is much better than in previous months, which is very good. This usage mostly stems from the low priority queue (74% of the overall usage) which is higher than the previous month (64%). It is done by 4 users listed by Elda, the top usage by far being attributed to Gabriele Merlo from the Swiss Plasma Center (83% of the total usage of the low priority queue, 62% of the total usage of the C1 partition). Roman comments that this was to be expected as some codes like ORB5 are ready to run on GPU nodes. It is agreed that Richard will contact Gabriele in order to ask him what usage he is doing and what his plans are in terms of usage of the low priority queue in the future ▲. It is noted that the usage of the partition by academic users (with preemption in case nodes are needed by EUROfusion users) is around 5% leading to a total usage of the partition of close to 92%. This usage is lower than the previous month due to the increase of usage from EUROfusion. Only 8 projects out of 32 already started (usage equal or more than 1% of the allocation). Regarding the 5 Japanese projects for FY2021 (April 2021–March 2022) only two projects made a usage of the C1 partition (P\_TRANS and GKNET\_DL). Statistics on the number of nodes used are presented, it is noted that most of the usage (77%) comes from jobs using 33 to 64 GPUs. The number of jobs using no GPU is larger than the previous month but remains acceptable (1552 out of 4532) – this needs to be monitored during the next meetings ▲.

			SKL (A3)	Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
			<i>Raw</i>						
Availability	A3+	Corrected		99,1%	98,7%	99,1%	98,7%	99,4%	99,7%
	login	Target (SLA)		97,0%	97,0%	97,0%	97,0%	97,0%	97,0%
Usage	A3	Real		87,9%	96,2%	96,4%	88,3%	83,0%	86,9%
		Target		85,0%	85,0%	85,0%	85,0%	85,0%	85,0%
			M100 (C1)	Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
			<i>Raw usage</i>						
Availability	C1+	Corrected		100,0%	98,7%	98,8%	98,1%	99,8%	96,4%
	login	Target (SLA)		97,0%	97,0%	97,0%	97,0%	97,0%	97,0%
Usage	C1	Real		10,6%	49,4%	34,2%	33,1%	51,7%	87,0%
		Target		none	none	none	none	none	none
				Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
Major incidents (from monthly re				1	1	3	3	1	0
				Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
Maintenance	A3	Month		9,5	4	10,5	5,5	8	4
	C1	Month		5,5	4,5	41,5	8	5,5	4
	A3	Remaining		166,5	166,5	164	164	164	164
	C1	Remaining		168	168	134,5	134,5	134,5	134,5

- Gateway:
  - The availability is excellent (99.3%) – and compliant with the SLA (97%).
  - The usage is rather high for an interactive system (70.3%) as it has been the case during the last months, showing that additional resources are needed for the Gateway.
- File systems including long-term archive (LTS): no problem to report, no additional usage by EUROfusion of the LTS.
- IT infrastructure: no problem to report.
- Elda presents detailed statistics regarding tickets. It is noted that this month there was a lower number of tickets than the previous month. The service provided is compliant with SLA in terms of % of tickets solved in 5 business days (87% for a SLA of 50%) and in 30 business days (94% for a SLA of 80%). Only a few minor deviations with no noticeable impact on users are noted.
- Summary of the KPIs for the month: all KPIs are compliant with the SLA except for a minor deviation regarding the availability of the C1 partition and minor deviations regarding tickets with no impact on users.

### 5.3.2 Feedback from Ticket Meetings #54 and update on network(s) status

France reports that the Ticket Committee noted that the systems were stable and that the handling of tickets by CINECA was good – no critical situation. The number of tickets submitted since the previous Ticket Meeting remained low and the overall situation is good. One ticket escalated to Intel remains open, investigations are on-going.

Sanity checks were successfully performed before and after maintenance for Marconi100 as agreed (no maintenance occurred on Marconi during the period of interest for the last ticket meeting). These checks have shown that the situation before and after maintenance is rather similar.

Two jobs failed on Marconi during the application performance tests performed by the ACH in Garching on July 6. Regarding the performances, the usual fluctuation is observed on Marconi while the performances are very stable on Marconi 100.

For likwid some additional tests are in progress, CINECA is working with Nils. For HPCMD, the security issue reported during the previous OC meeting was solved and the installation is in progress.

Roman comments that the installation of HPCMD is a big step forward and thanks CINECA for this excellent work even if the time it took was long. He notes that with such a tool it is possible to assess how efficient a code on the machine is which is very important.

Richard comments that it might be good to ask for the next systems the availability of such tools from the start. François suggest that one or two sentences explaining the target functionalities be prepared in order to be inserted in the future call. It is agreed that Roman, France and Jacques will prepare this short text ▲.

### 5.3.3 Discussion on current issues

Elda explains that the GPU allocation problem with Slurm reported by Leo will be fixed in the next release of Slurm (August or September 2021).

## 5.4 Gateway

### 5.4.1 Report for June 2021

Francesco presents the Gateway monthly report for June 2021. There were changes neither in the system configuration nor in the Slurm configuration. The list of the virtual nodes provided by CINECA for the Gateway remains also unchanged since the last OC meeting.

He notes that the situation of the GPFS storage is good with a reasonable utilization (74%). He also notes that 10 users account for 77% of this usage. Currently there are 417 user accounts on the Gateway; however, 118 were never logged in.

The current issues and actions are reported. The two most important ones are the installation of a new version of HDF5 and the study of the Gateway extension – no update at this time except regarding the extension of the Gateway.

The availability of the Gateway compute nodes (as already reported by Elda) was excellent (99%) and the usage (70%) normal for an interactive machine. The availability of the login nodes was excellent also (100%) and the usage (9–14%) normal for this kind of nodes even if higher than the previous months due to one user (g2tjons). David will contact this user to discuss with him this situation ▲. Two minor incidents occurred during the month: one issue with NAG license, one period of instability of LPAD/SSSD until June 28.

The tickets are all solved in a timely manner, 13 were opened, 13 were closed, 84% in 1 week, 100% in 9 days.

Regarding the extension of the Gateway, several points are discussed:

- For the DevOps tools, it is agreed that the request from EUROfusion is only to provide the necessary infrastructure (including virtual machines).
- For the hardware extension, ENEA prepared with CINECA a plan for moving 64 nodes of the A3 partition to the Gateway. This operation could be completed by January 1, 2022 and involves minor costs. Francesco confirms that after this extension, the Gateway will remain an independent system with no single point of failure as it was before.
- For the software update and additional services, including new VM, the question raised is the operating system version to use – possibly Redhat 8.x. This point is discussed, it seems important to avoid using different operating systems on the supercomputer and on the Gateway as in both cases the stability is very important. It is agreed that the choice of the future operating system will be done once the strategy for the upgrade of the system software of Marconi is proposed by CINECA ▲.

#### 5.4.2 Discussion on current issues

No point was discussed.

### 5.5 AOB

France presents the current status of the webinar series and the first ideas for the third IFERC GPU workshops (possibly a face to face or at least more interactive event in June 2022).

Regarding the webinars, several points are discussed:

- France explains that the number of participants remains rather small and that the impact on the overall community seems limited. Richard comments that the goal is not to have a large number of participants but to provide useful information.

- François suggests organizing from time to time a webinar more “simple” like the one done by CINECA at the beginning of Marconi100. David suggests organizing such seminar around the start of a new cycle.
- Richard adds that there is no pressure to have a specific frequency; the webinar should be organized when useful. France thinks that more hands-on session could be nice.
- A questionnaire to user does not seem useful as people will certainly ask if they have ideas or needs.
- The involvement of the ACHs in the preparation of the workshop would be useful. Richard suggests ACHs to discuss this question during regular ACHs meetings. David adds that ACHs could also propose speakers such as computer scientists involved in the porting to GPU of GENE or ORB5.

France concludes her presentation on the need to have a stable platform for the webinars (the last webinar had to be rescheduled). Richard will study this question ▲.

François reminds Francesco that the deliverable for Q2-2021 is expected to be prepared by ENEA soon ▲.

## 6 Next meetings

The next meetings will be held on:

- OC-62: September 20 (9:30–12:30) – by video
- OC-63: October 22 (9:30–12:30) – by video end/of physical meeting? It is agreed that, if possible, this meeting should be a face to face meeting since it would be a good opportunity to get to know the expected new participants to the OC meetings at that time (representatives from the ACHs) and also to visit Marconi100.

There will be no OC meeting in August. However, Elda and Francesco will provide the usual monthly reports for July by August 23 ▲.

François will prepare a doodle poll for the following meeting OC-64 (November 2021<sup>1</sup>) ▲.

---

<sup>1</sup> The date of November 26 morning (9:30–12:30) was decided after the meeting





# EUROfusion HPC Operation Committee

**R. Kamendje (Chair)**  
**F. Robin (Secretary)**



This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

## Agenda



1. Introduction (R. Kamendje)
  - a. Welcome and confirmation of the agenda
    - i. New members of the Operation Committee
  - b. Review and confirmation of the minutes of the previous meeting (OC-60)
  - c. General information
    - i. Feedback from the last EUROfusion GA meeting regarding computing resources for 2023-2029
    - ii. Status of discussions between EUROfusion and NVidia
2. Marconi-Fusion
  - a. Monthly report from ENEA to the Operation Committee (June 2021) (E. Rossi).
  - b. Monthly report of the Ticket Meeting (June 2021) (F. Boillod-Cerneux, R. Hatzky)
  - c. Discussion on current issues
    - i. Update on the allocation of GPU jobs in M100 (E. Rossi)
    - ii. Top10 users of the low priority queue on M100 (E. Rossi)
3. Gateway
  - a. Monthly report from ENEA to the Operation Committee (June 2021) (F. Iannone)
  - b. Discussion on current issues
    - i. Discussion about extension of the current Gateway (number of nodes, role of the OC,...) (All)
4. AOB
  - a. Update on Eurofusion Webinars on GPUs (F. Boillod-Cerneux, R. Hatzky, J. David)
5. Next meetings
  - a. 23 August 2021 – Written report from E. Rossi and F. Iannone
  - b. 20 September 2021 (9:30-12:30): OC meeting #62 – by video
  - c. 22 October 2021 (9:30-12:30): OC meeting #63 – by video and/or face-to-face meeting



- **ACH-Cat. 1: High Performance Computing**

- scalable algorithms, code parallelization and performance optimization, code refactoring, GPU-enabling, ...

- 3 ACHs now in place

- *IPP-Garching* **Roman Hatzky**

- **BSC-Barcelona** **David Vicente**

- **EPFL/SPC- Lausanne** **Gilles Fourestey**



- **ACH-Cat. 2: Integrated Modelling and Control**

- code adaptation to IMAS, IMAS framework development, code integration

- 1 ACH now in place

- **PSNC-Poznan** **Michal Owsiak**



- **ACH-Cat. 3: Data Management**

- open access, data management, data analysis tools, aspects of AI and VVUQ

- 1 ACH now in place

- **VTT-Helsinki** **Fredric Granberg**



## General information



## **Marconi Fusion**

# EUROfusion HPC: Monthly report to the Operation Committee (June2021)

Elda Rossi, Claudia Truini (Cineca) July 7, 2021

## Content

EUROfusion HPC: Monthly report to the Operation Committee (June2021) .....	1
The operation of the HPC equipment .....	2
Main events: Scheduled Maintenance .....	2
Main events: Incidents & other relevant news.....	3
Main events: summary of (still opened) system problems .....	3
Main events: security news .....	3
Main events: HPC-news .....	4
Availability of the dedicated equipment .....	5
Availability of Marconi-Fusion partition - SKL (A3) .....	6
Availability of accelerated partition C1– M100 .....	7
Availability of the Gateway cluster .....	8
Usage of the dedicated equipment .....	9
Usage of the MARCONI-fusion partition SKL (A3).....	9
Usage of the accelerated partition - M100.....	15
Usage of the Gateway partition.....	20
Usage of the MARCONI file systems .....	21
Usage of the M100 file systems.....	21
Usage of the long-term archive (DRES).....	22
Characteristics of Usage .....	23
Job parallelism on Marconi-fusion (Marconi – A3).....	23
Job parallelism on Marconi-100 (GPUs usage) .....	24
Availability of the IT infrastructure .....	25
External network .....	25
Scratch filesystem MARCONI & M100 .....	26
Long-term Archive Marconi & M100 .....	27
Summary IT infrastructure.....	28
Service provided to EUROfusion .....	29
Statistics on the ticketing system .....	29
Provisioning: statistics on the username/project creation (KPI 8) .....	31
KPI .....	32
Appendix .....	34
Update on the allocation of GPU jobs in M100 (E. Rossi).....	34
Top10 users of the low priority queue on M100 (E. Rossi).....	34

# The operation of the HPC equipment

We report here major events affecting the availability and usage of Marconi-fusion for the period, as well as data related to the availability of the system.

## Main events: Scheduled Maintenance

### Jun 29 (M100)

- The SLURM prolog/epilog was modified to manage DCGMI, a tool by nVidia that performs a GPU health check. This feature checks that the GPUs are correctly working and generates a GPU statistics report in the job workDir.
- run sanity checks both pre and post maintenance. Post maintenance found several nodes with a slow GPU, drained from production.

### Jun 15 (MARCONI)

- Reboot of all nodes
- Sanity checks before the reboot OK except for two nodes slow
- Sanity checks after the reboot passed except for 4 nodes slow drained from production.

Gateway maintenance this month not planned

Next Maintenance: (<https://www.hpc.cineca.it/calendario-completo/month>)

- Jul 8 (Marconi - CANCELLED)
- Jul 20 (M100 - CANCELLED)

## Main events: Incidents & other relevant news

In colour: changes to the infrastructure

In white: incidents to the infrastructure (in grey: not affecting HPC production)

Data	Where	What	Notes
Jun 11	M100	FileSystem issue	Due to the depletion of the number of inodes on WORK FileSystem, some users reported error of "No space left on device". Number of inodes have been increased to solve the problem.
Jun 15	Marconi	Schedule maintenance	
Jun 29	M100	Schedule maintenance	

During all the month, several incidents on the M100 login nodes occurred. A preliminary analysis suggests a possible involvement of RedHut 8.1 and Nvidia compilers, together with crowded login nodes (an Hackathon event on M100 caused the first incident). According to Nvidia, no similar issues have been detected before, an official call to IBM has been opened.

## Main events: summary of (still opened) system problems

Nothing to report

## Main events: security news

Nothing to report

## Main events: HPC-news

29/06/2021

### **Marconi100 back to production**

Dear users, the maintenance operations on Marconi100 have been completed and the cluster is now...

28/06/2021

### **Reminder: scheduled maintenance of Marconi100 tomorrow, June 29th**

Dear Users, This is to remind you that Marconi100 will be stopped for scheduled maintenance...

24/06/2021

### **Scheduled maintenance of Marconi100 on June 29th**

Dear Users, This is to inform you that Marconi100 will be stopped for scheduled maintenance...

18/06/2021

### **MARCONI100: scratch quota removed**

Dear Marconi users, After the emergency of yesterday, the occupation quota of the scratch area of...

17/06/2021

### **Marconi100: scratch is almost full - quota imposed**

Dear Marconi100 Users, we inform you that the scratch space has reached the occupation of more...

15/06/2021

### **Marconi100 - login issue update**

Dear Users, this is an update to inform you that the login issue on Marconi100 is related just to...

15/06/2021

### **Marconi100 - login issue**

Dear users, this is to inform you that we are experiencing some login issues on Marconi100 cluster...

15/06/2021

### **Marconi back in production**

Dear users, the maintenance operations on Marconi have been completed and the cluster is now back...

14/06/2021

### **Reminder: scheduled maintenance of Marconi tomorrow, June 15th**

Dear Users, This is to inform you that Marconi will be stopped for scheduled maintenance...

08/06/2021

### **Scheduled maintenance of Marconi on June 15th**

Dear Users, This is to inform you that Marconi will be stopped for scheduled maintenance...

01/06/2021

### **Help Desk service not available on June 2nd**

Dear Users, This is to inform you that, due to the Italian Republic Celebration Day, the Help...

# Availability of the dedicated equipment

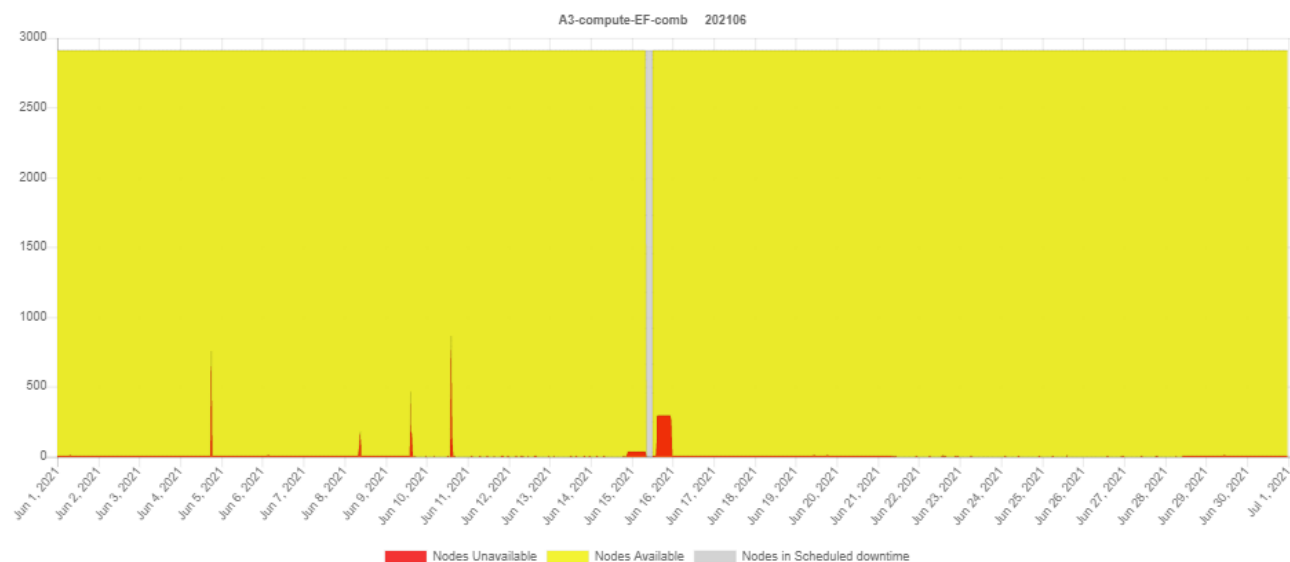
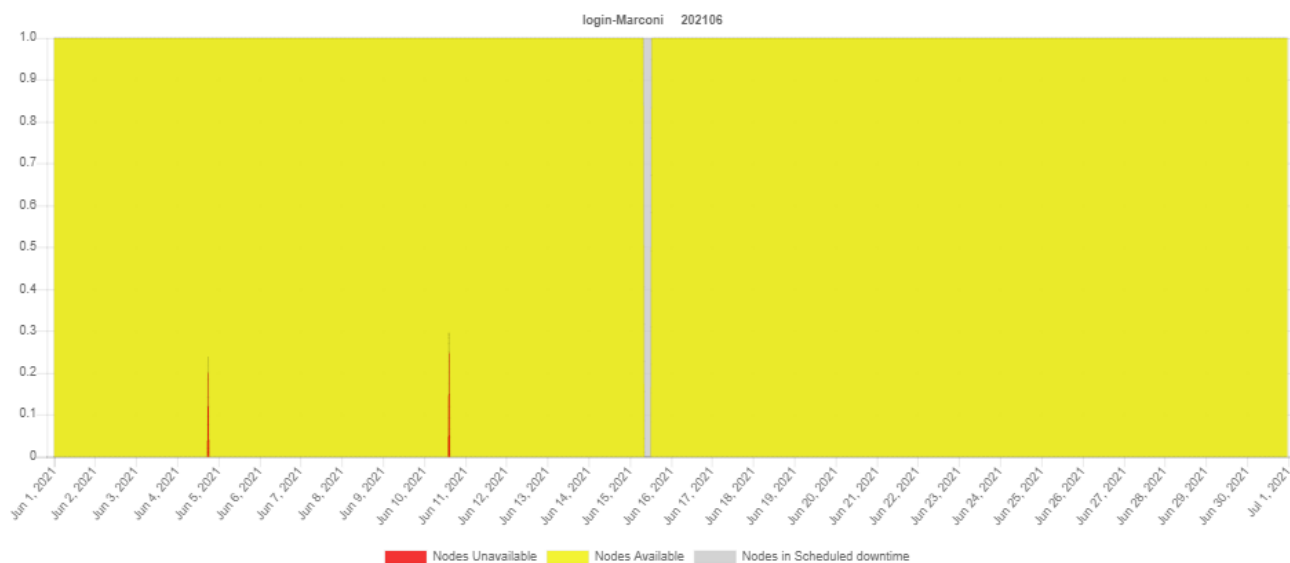
The equipment dedicated to EUROfusion in Cineca is made of:

- **Marconi-fusion:** 2912 nodes of the A3 partition of Marconi (SKL) distributed into 41 racks. Two of those nodes are part of the "test" partition (4 nodes in total, 2 from EUROfusion partition and 2 from CINECA partition)
- **Accelerated partition:** 99 nodes on Marconi100 (M100).
  
- **Gateway:** 28 SKL nodes in a dedicated cluster (EFGW). Four of those nodes are used as "permanent" login nodes, so only 24 nodes are monitored as part of the EFGW Cluster.



# Availability of Marconi-Fusion partition - SKL (A3)

The data are collected by dedicated monitoring logs whose architecture is compliant with the “policy” rules agreed upon. (chiamategdc.cineca.it) – 2408 nodes

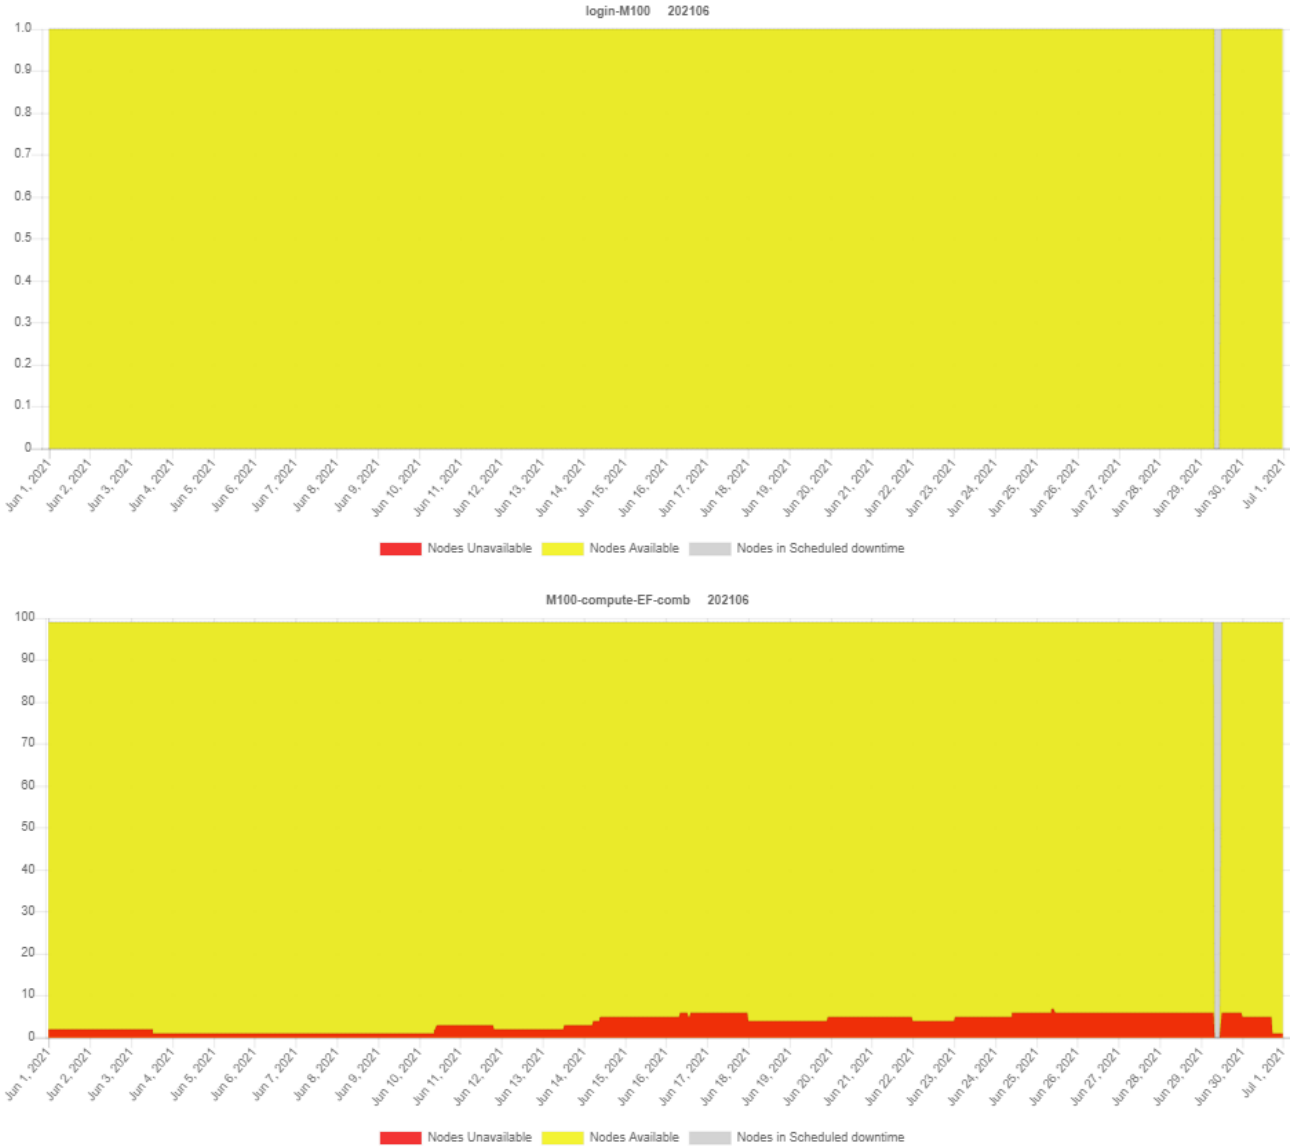


Maintenance

Coverage=(100)% what	Value	KPI
A3-Availability/ (Theor-maint)	99.7%	KPI-1a
A3-Scheduled maintenance	4.0 h (0 extra-h)	KPI-2a

# Availability of accelerated partition C1- M100

In the following figure, we report the availability of M100 (EUfus partition) in the period of reference (80 nodes):




The low value for availability (<97%) is due to several nodes out of production since mid-June. This is due mainly to node overheating cause by seasonal high temperature. We are considering the possibility to decrease the GPU frequency to make the nodes more resilient to the overheating.

From this figure the KPIs were computed:

Coverage=(100.0 – 1.03)% what	Value	KPI
C1 - Availability/ (Theor-maint)	96.4 %	KPI-1b
C1- Scheduled maintenance	4 h (0 extra-h)	KPI-2b

# Availability of the Gateway cluster

Availability on the 24 nodes devoted to computing (4 nodes user as login nodes)




Coverage=(100)% what	Value	KPI
Availability/ (Theor-maint)	99.3 %	KPI 1d
Scheduled maintenance	0 h (0 extra maintenance)	KPI 2 d

# Usage of the dedicated equipment

## Usage of the MARCONI-fusion partition SKL (A3)

### General Usage A3

In the following, the usage of Marconi-fusion, on Marconi SKL – A3, is reported.



Maint

what	Used node-h/Avail node-h	Value	KPI
Usage/(available) – A3	<b>1'804'927 / 2'076'947</b>	86.9%	KPI-4a

### Usage per queue (A3)

In the table the “usage” (in node-h units) is reported in terms of the SLURM partition/Quality of Service.

Queues	njobs	node-h	core-avr	Q-h avr
skl_fua_dbg/normal	2'427	767	78	0.02
skl_fua_prod/normal	44'526	1'416'814	257	1.95
skl_fua_prod/skl_qos_fuabprod	390	278'116	5'379	1.44
skl_fua_prod/skl_qos_fualprod	333	24'544	169	1.92
skl_fua_prod/qos_special	-	-	-	-
skl_fua_prod/qos_prio	-	-	-	-
<b>skl_fua_prod/skl_qos_fualowprio</b>	<b>417</b>	<b>81'743</b>	<b>770</b>	<b>82.16</b>

Queues/QOS on A3

The usage of the “lowPriority” queue is represented by the correspondent line of the above table (in red)

The special queue was required by no one

### Usage per projects (A3)

As indicated by the Allocation Committee, several classes of projects have been defined on the UserDB of Cineca on the A3 (conventional) partition:

- FUSIO\_: service projects.
- FUSIO\_ru5\*: Research Unit projects, 5th cycle, from Mar 1 2021 to Feb 28, 2022
- FUSIO\_ja\*: projects assigned for the EU-JP collaboration
- FUA35\_\* : 5th cycle on A3, from Mar 1 2021 to Feb 28, 2022

Project-A3	defined	on-cluster	active	budget (node-h)	used (node-h)	used%	> 100%	from	to	status
FUA35	134	133	123	28'799'444	6'993'929	24%	7	01/03/2021	28/02/2022	running
<i>FUA35_LOWPRIO not included</i>										

	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
1	FUSIO_HLST_1	272	1'254	742	188'646	800'000		
2	FUSIO_ALL_1	1	9'216	0	8'003	10'417		
3	FUSIO_TEST_1	10	48	1	4'225	6'250		
	<b>FUA35_LOWPRIO</b>	<b>381</b>	<b>686</b>	<b>59'678</b>	<b>83'607</b>	<b>1</b>		

	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
1	FUSIO_ru5CCFE	-	-	-	2'038	1'000	204%	34%
2	FUSIO_ru5IPP	-	-	-	1'170	1'000	117%	34%
3	FUSIO_ru5IPPLM	-	-	-	1'034	1'000	103%	34%
4	FUSIO_ru5IST	-	-	-	1'024	1'000	102%	34%
5	FUSIO_ru5NCSRD	-	-	-	931	1'000	93%	34%
6	FUSIO_ru5DTU	-	-	-	201	1'000	20%	34%
7	FUSIO_ru5FOM	-	-	-	2	1'000	0%	34%
8	FUSIO_ru5CU	-	-	-	0	1'000	0%	34%
9	FUSIO_ru5CEA	-	-	-	-	1'000	0%	34%
10	FUSIO_ru5CIEMA	-	-	-	-	1'000	0%	34%
11	FUSIO_ru5ENEA	-	-	-	-	1'000	0%	34%
12	FUSIO_ru5EPFL	-	-	-	-	1'000	0%	34%
13	FUSIO_ru5FZJ	-	-	-	-	1'000	0%	34%
14	FUSIO_ru5IPPCR	-	-	-	-	1'000	0%	34%
15	FUSIO_ru5ISSPU	-	-	-	-	1'000	0%	34%
16	FUSIO_ru5JSI	-	-	-	-	1'000	0%	34%
17	FUSIO_ru5VTT	-	-	-	-	1'000	0%	34%

	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
1	FUA35_GKEMHEL	19	3'072	21'534	73'240	53'000	138%	34%
2	FUA35_ETENMIBE	550	632	107'034	107'034	81'900	131%	34%
3	FUA35_MW_EMIT	111	240	3'024	8'979	7'000	128%	34%
4	FUA35_MULTAUG2	-	-	-	56'190	50'000	112%	34%
5	FUA35_FusNeu_5	75	1'036	25'879	110'548	106'600	104%	34%
6	FUA35_DEMOPPR	3'118	48	4'694	32'846	32'500	101%	34%
7	FUA35_REFMULx	23	3'349	7'475	25'161	25'000	101%	34%
8	FUA35_XTurb	14	2'304	6'715	44'813	46'000	97%	34%
9	FUA35_GSNTITE	50	1'021	18'037	218'875	234'000	94%	34%
10	FUA35_STATS	34	3'269	9'637	52'539	72'000	73%	34%
11	FUA35_KIPPSOL	1	6'144	1'217	15'115	25'000	60%	34%
12	FUA35_OPTSTEL5	34	1'355	7'955	30'997	52'000	60%	34%
13	FUA35_DCLLMHD	44	78	60	71'580	132'600	54%	34%
14	FUA35_MEGAEDGE	27	172	558	62'421	117'000	53%	34%
15	FUA35_JSELM1	9	2'064	508	16'685	32'500	51%	34%
16	FUA35_SUERTE	77	587	11'407	29'581	59'800	49%	34%
17	FUA35_DIVturb	88	555	5'899	27'050	56'000	48%	34%
18	FUA35_VAC_ND	5	1'920	1'366	94'113	200'000	47%	34%
19	FUA35_EST3D	61	384	8'412	36'607	79'000	46%	34%
20	FUA35_FH2CPS	-	-	-	441	1'000	44%	34%
21	FUA35_ALARA2	67	1'014	9'325	86'036	196'300	44%	34%
22	FUA35_ROBIN	32	6'516	31'716	175'938	406'000	43%	34%
23	FUA35_QLTURB	1'072	199	15'091	66'828	156'000	43%	34%
24	FUA35_MHD	666	711	144'097	723'225	1'800'000	40%	34%
25	FUA35_MFP_KIT	15	998	6'384	15'340	39'000	39%	34%
26	FUA35_F4E	74	1'390	23'593	78'281	200'000	39%	34%
27	FUA35_GYS_RSST	-	-	-	41'070	109'000	38%	34%
28	FUA35_TSVVT421	3'203	166	80'913	493'210	1'325'000	37%	34%
29	FUA35_DEFGEN	414	553	56'628	237'593	650'000	37%	34%
30	FUA35_EMCAD	1'424	48	21'797	55'999	153'400	37%	34%
31	FUA35_SiSteel	4'125	371	23'041	53'954	149'000	36%	34%
32	FUA35_TSVV3	256	690	43'508	301'792	835'000	36%	34%
33	FUA35_MEGAFILD	185	1'446	42'533	162'862	453'000	36%	34%
34	FUA35_JOREKITG	165	521	14'261	72'971	204'000	36%	34%
35	FUA35_VENUSLEV	114	1'518	7'263	34'905	100'750	35%	34%
36	FUA35_MAXJNEGS	43	953	11'502	28'905	84'000	34%	34%
37	FUA35_CUCsMo_4	61	62	1'772	9'592	28'000	34%	34%
38	FUA35_TURBGENE	276	641	34'958	126'037	375'000	34%	34%
39	FUA35_UQMWA3	-	-	-	32'168	96'000	34%	34%
40	FUA35_TRASIMEX	10'404	49	7'739	39'923	120'900	33%	34%
41	FUA35_NKUAGY21	366	63	3'866	24'641	75'000	33%	34%
42	FUA35_GBSSTEL	411	339	20'662	60'621	188'000	32%	34%
43	FUA35_MEGAPR	431	528	27'977	100'292	313'000	32%	34%
44	FUA35_GNKEW7X	79	3'832	29'308	138'501	438'000	32%	34%
45	FUA35_TSVV7	449	553	89'349	162'048	516'000	31%	34%
46	FUA35_SOLPSITE-	2	48	38	9'101	30'000	30%	34%
47	FUA35_CluSTab	1'485	210	88'819	100'387	331'500	30%	34%
48	FUA35_KinEdgeV	104	1'783	18'914	98'324	325'000	30%	34%
49	FUA35_MONASRID	943	381	63'903	125'879	418'600	30%	34%
50	FUA35_ORB5ELE	87	1'594	22'955	53'616	178'750	30%	34%

	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
51	FUA35_STELTURB	2'513	348	58'083	377'488	1'260'000	30%	34%
52	FUA35_JSELM	29	599	6'769	64'560	216'000	30%	34%
53	FUA35_GKNN	2'560	70	12'822	40'395	136'000	30%	34%
54	FUA35_WPJET1	138	1'368	33'298	165'506	600'000	28%	34%
55	FUA35_ORBFAST	222	854	6'536	89'616	325'000	28%	34%
56	FUA35_UHPDI	6	72	143	18'364	68'000	27%	34%
57	FUA35_MHD3D_2	145	135	4'176	8'255	32'500	25%	34%
58	FUA35_SPICE	366	55	4'205	4'440	18'000	25%	34%
59	FUA35_GBS_CZ	108	1'143	11'133	25'541	108'000	24%	34%
60	FUA35_JT-LANE2	643	192	4'282	16'605	71'000	23%	34%
61	FUA35_SOLDyn4	26	480	15'841	21'389	94'250	23%	34%
62	FUA35_SOLBOUT5	1'130	238	30'807	95'209	424'000	22%	34%
63	FUA35_SPIM3D	7	438	97	18'010	81'250	22%	34%
64	FUA35_SNOX4	17	4'608	20'089	47'066	215'000	22%	34%
65	FUA35_LDP-KIT	117	451	5'666	21'932	105'300	21%	34%
66	FUA35_OrbZONE	33	795	3'499	54'985	269'100	20%	34%
67	FUA35_PICLS	64	384	10'876	56'564	281'000	20%	34%
68	FUA35_NAIBOL	-	-	-	10'809	53'760	20%	34%
69	FUA35_ASCOT-EP	21	3'072	4'397	25'618	128'000	20%	34%
70	FUA35_REDIS	627	359	46'126	136'531	700'000	20%	34%
71	FUA35_LHPED21	775	1'611	119'570	252'796	1'350'000	19%	34%
72	FUA35_2DPICMCC	-	-	-	1'811	10'000	18%	34%
73	FUA35_WPDCNeu2	64	502	2'789	5'562	32'500	17%	34%
74	FUA35_Gyro2021	46	382	2'998	10'459	62'000	17%	34%
75	FUA35_FELTOR	155	48	2'369	7'833	48'000	16%	34%
76	FUA35_MHDwcll	-	-	-	10'528	65'000	16%	34%
77	FUA35_NEDEHE	3	720	1'035	5'168	32'500	16%	34%
78	FUA35_MODE2	53	372	425	3'276	20'833	16%	34%
79	FUA35_OXGK	1'841	111	8'669	35'080	234'000	15%	34%
80	FUA35_TEMPIS	864	141	26'485	44'636	300'000	15%	34%
81	FUA35_NAITEC2	-	-	-	7'959	54'000	15%	34%
82	FUA35_GYO-GYS2	403	478	9'550	13'641	94'000	15%	34%
83	FUA35_TTMI	190	49	1'909	3'484	25'000	14%	34%
84	FUA35_RF_WAVES	-	-	-	661	5'000	13%	34%
85	FUA35_HYMHGK4	33	142	12	7'052	54'600	13%	34%
86	FUA35_ASCOT-FI	26	978	2'763	14'643	117'000	13%	34%
87	FUA35_EUGY	57	1'333	7'351	33'908	283'000	12%	34%
88	FUA35_SOLF	46	91	285	3'823	32'000	12%	34%
89	FUA35_MOTISANI	1	288	48	54'131	466'700	12%	34%
90	FUA35_TDSLTPLD	171	209	6'157	8'553	81'900	10%	34%
91	FUA35_TSVV2	105	960	19'853	66'010	650'000	10%	34%
92	FUA35_QBiT	158	637	252	3'220	32'500	10%	34%
93	FUA35_LoGy	30	1'753	8'405	15'220	162'500	9%	34%
94	FUA35_MCHIFI	141	712	26'357	44'545	500'000	9%	34%
95	FUA35_FKMR	322	347	7'537	22'212	250'000	9%	34%
96	FUA35_neutralG	66	206	12	4'754	55'000	9%	34%
97	FUA35_BOUTFLOW	113	53	740	1'757	25'000	7%	34%
98	FUA35_PIXIE3D	48	208	2'449	11'686	213'000	5%	34%
99	FUA35_REFTOP	22	4'747	1'780	1'782	32'500	5%	34%
100	FUA35_ID2DRM	112	355	2'756	13'889	263'900	5%	34%

	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
101	FUA35_ASCOT-ESP	156	350	1'960	3'424	78'000	4%	34%
102	FUA35_MCINS	199	339	5'280	25'602	650'000	4%	34%
103	FUA35_DIVGAS-K	-	-	-	16'603	466'700	4%	34%
104	FUA35_UNEDJET	-	-	-	1'009	32'500	3%	34%
105	FUA35_HiFi-NSD	44	89	112	12'572	588'900	2%	34%
106	FUA35_ALVES3	165	48	220	669	34'000	2%	34%
107	FUA35_HEAT	424	174	241	2'127	114'400	2%	34%
108	FUA35_FAST7X	260	478	1'846	3'537	194'000	2%	34%
109	FUA35_GKINNER	41	357	4'604	4'642	359'000	1%	34%
110	FUA35_ELM-UK	12	132	15	1'072	109'850	1%	34%
111	FUA35_NNFT2021	37	125	20	506	55'900	1%	34%
112	FUA35_APLEPY	-	-	-	80	12'300	1%	34%
113	FUA35_ERO20	2	6'192	1'834	1'834	297'000	1%	34%
114	FUA35_STELLMHD	27	392	380	395	78'000	1%	34%
115	FUA35_ConfTrans	29	768	313	421	90'000	0%	34%
116	FUA35_BeTi2021	80	96	0	768	184'600	0%	34%
117	FUA35_PC-CFD	-	-	-	647	304'200	0%	34%
118	FUA35_InZights	-	-	-	293	192'000	0%	34%
119	FUA35_NTMTHR	5	48	-	84	165'000	0%	34%
120	FUA35_REFRAMP	38	197	9	9	32'500	0%	34%
121	FUA35_TSVV10	-	-	-	214	1'500'000	0%	34%
122	FUA35_ITEDGE	-	-	-	1	195'000	0%	34%
123	FUA35_EGGWPCD	-	-	-	-	2'000	0%	34%
124	FUA35_ENEATOP	-	-	-	-	25'000	0%	34%
125	FUA35_EP12WAVS	-	-	-	-	30'000	0%	34%
126	FUA35_ETG_PED	-	-	-	-	25'000	0%	34%
127	FUA35_KNEE2	-	-	-	-	25'000	0%	34%
128	FUA35_MHD-BMG	-	-	-	-	58'500	0%	34%
129	FUA35_MULTTCV	-	-	-	-	281'000	0%	34%
130	FUA35_TEMFracT	-	-	-	-	37'000	0%	34%
131	FUA35_TOPICA5	-	-	-	-	20'000	0%	34%
132	FUA35_UNEDONES	-	-	-	-	30'000	0%	34%
133	FUA35_WCLLPDR	-	-	-	-	36'400	0%	34%





### Average number of nodes requested by jobs in the wait queue (A3)

Queued nodes (nodes required by queued jobs) are reported as a function of time.

Data come from a cron-driven data collection. Data from log files correctly describes all “pending” jobs (in red), but not “queued” jobs (jobs in Pending state ready to execute). In the figure, negative values mean unavailability of data.

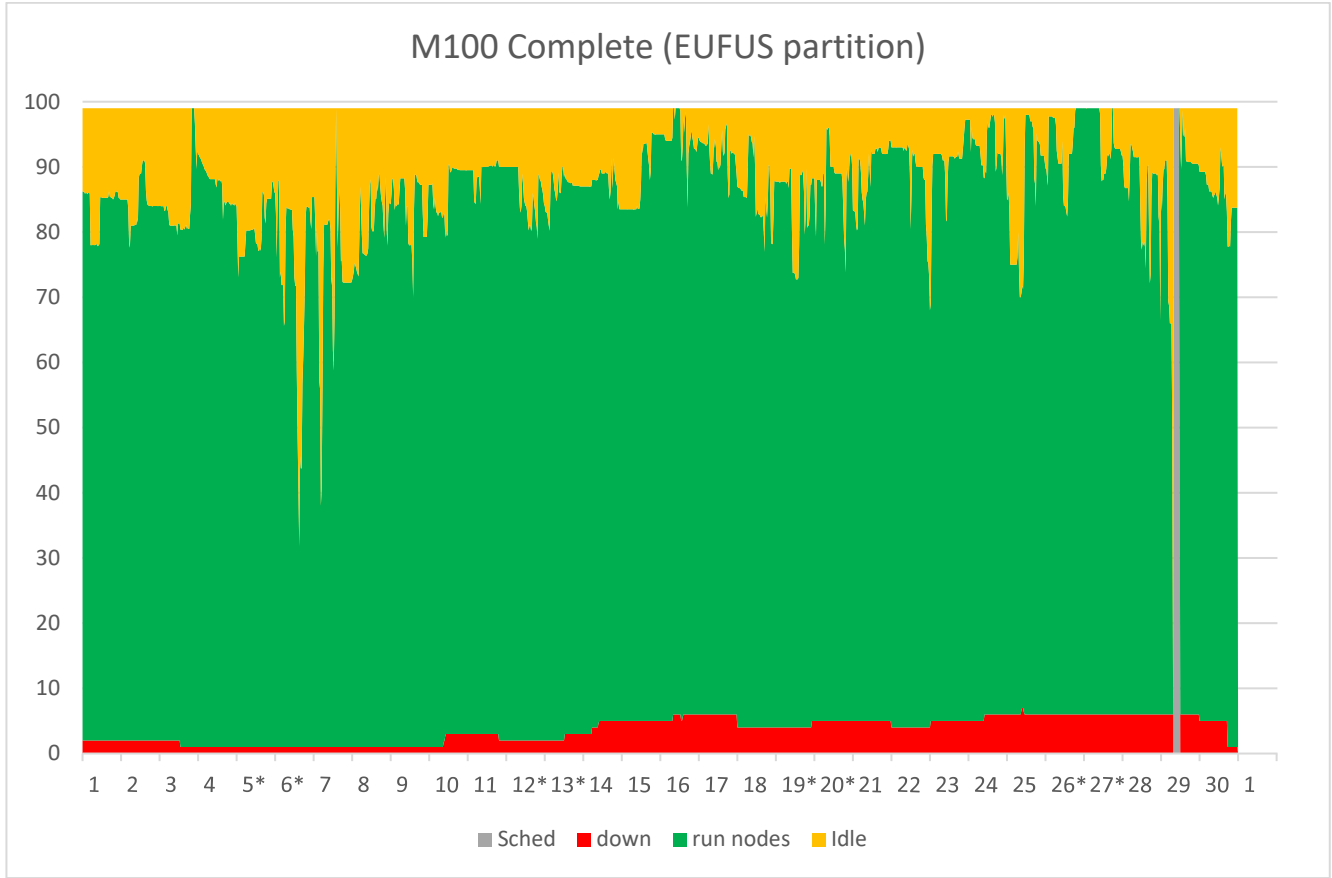
By comparing the “usage” figure with the blue curve (Q nodes), it is easy to see that low level of usage are mainly due to not availability of jobs in the Q state, or system draining in preparation of sched maintenance.



# Usage of the accelerated partition - M100

## General Usage M100

The following picture report the usage of the accelerated partition in the month, for the EUfus partition (99 nodes):



Maint

what	Used node-h/Avail node-h	value	KPI
Usage/(available) – M100	59'445 / 68'324	87.0%	KPI-4b

(\*) we report here the real node-h value used in the month, taking care of the use of the GPUs.

## Usage per queue (M100)

In the table, the “usage” (in node-h units) is reported in terms of the SLURM partitions/QoS:

Queues	njobs	node-h	core-avr	Q-h avr
m100_fua_prod/m100_qos_fuadbg	84	15	31	0.01
m100_fua_prod/normal	3'151	14'819	42	1.53
m100_fua_prod/m100_qos_fuabpro	14	263	1'024	6.27
m100_fua_prod/qos_lowprio	906	43'832	172	8.39

Queues on M100

## Usage per projects (M100)

As indicated by the Allocation Committee, several classes of projects have been defined on the UserDB of Cineca on the accelerated partition:

- FUSIO\_: service projects.
- FUSIO\_ru5\*: Research Unit projects, 4th phase, from March 1 to Feb 28, 2022
- FUAC5\_\* : cycle 5 on Marconi100, from May 1 to Feb 28, 2022

Project-M100	defined	on-cluster	active	budget (node-h)	used (node-h)	used%	> 100%	from	to	status
FUAC5	33	32	12	838'001	136'508	16%	-	01/03/2021	28/02/2022	running
<i>FUAC5_LOWPRIO not included</i>										
project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time			
1 FUSIO_ALL_3	-	-	-	104						
2 FUSIO_HLST_3	131	65	108	2'416						
3 FUSIO_TEST_3	12	144	0	1						
4 FUAC5_LOWPRIO	902	173	43'832	92'308						

project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time		
1 FUSIO_ru5CCFE_0	4	1	0	393	200	196%	34%		
2 FUSIO_ru5CIEMA_0	-	-	-	13	200	6%	34%		
3 FUSIO_ru5IPPLM_0	-	-	-	2	200	1%	34%		
4 FUSIO_ru5IPP_0	2	24	0	0	200	0%	34%		
5 FUSIO_ru5CEA_0	-	-	-	-	200	0%	34%		
6 FUSIO_ru5CU_0	-	-	-	-	200	0%	34%		
7 FUSIO_ru5DTU_0	-	-	-	-	200	0%	34%		
8 FUSIO_ru5ENEA_0	-	-	-	-	200	0%	34%		
9 FUSIO_ru5EPFL_0	-	-	-	-	200	0%	34%		
10 FUSIO_ru5FOM_0	-	-	-	-	200	0%	34%		
11 FUSIO_ru5FZJ_0	-	-	-	-	200	0%	34%		
12 FUSIO_ru5IPPCR_0	-	-	-	-	200	0%	34%		
13 FUSIO_ru5ISSPU_0	-	-	-	-	200	0%	34%		
14 FUSIO_ru5IST_0	-	-	-	-	200	0%	34%		
15 FUSIO_ru5JSI_0	-	-	-	-	200	0%	34%		
16 FUSIO_ru5NCSRDC_0	-	-	-	-	200	0%	34%		
17 FUSIO_ru5VTT_0	-	-	-	-	200	0%	34%		

project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time		
1 FUSIO_ja5P_TRA	16	46	9	9	10'000	0%	17%		
2 FUSIO_ja5GKNET	-	-	-	-	15'600	0%	17%		
3 FUSIO_ja5MLPWI	-	-	-	-	5'000	0%	17%		
4 FUSIO_ja5PPGYS	-	-	-	-	6'000	0%	17%		



	project	njobs	core avr	node-h this month	tot node-h	budget	%use	%time
1	FUAC5_FELTOR	1'419	43	8'659	18'648	23'000	81%	34%
2	FUAC5_OrbZONE	59	359	4'785	10'385	23'000	45%	34%
3	FUAC5_GKNN	1'306	3	191	287	1'000	29%	34%
4	FUAC5_TSVV10	13	512	138	2'577	23'000	11%	34%
5	FUAC5_LHPED21	23	265	717	4'624	60'000	8%	34%
6	FUAC5_ORBFAST	9	213	2	3'847	80'000	5%	34%
7	FUAC5_TSVV2	43	435	336	3'665	80'000	5%	34%
8	FUAC5_GBSSTEL	228	87	157	165	6'000	3%	34%
9	FUAC5_ASCOT-EP	-	-	-	2	6'000	0%	34%
10	FUAC5_BeTi2021	-	-	-	0	23'000	0%	34%
11	FUAC5_TSVV3	-	-	-	0	50'000	0%	34%
12	FUAC5_ALVES3	-	-	-	-	23'000	0%	34%
13	FUAC5_APLEPY	-	-	-	-	2'000	0%	34%
14	FUAC5_DCLLMHD	-	-	-	-	40'000	0%	34%
15	FUAC5_DEFGEN	-	-	-	-	80'000	0%	34%
16	FUAC5_HiFi-NSD	-	-	-	-	5'000	0%	34%
17	FUAC5_HYMHGK4	-	-	-	-	3'000	0%	34%
18	FUAC5_InZights	-	-	-	-	9'000	0%	34%
19	FUAC5_JSELM	-	-	-	-	5'000	0%	34%
20	FUAC5_LoGy	-	-	-	-	2'000	0%	34%
21	FUAC5_MCHIFI	-	-	-	-	2'000	0%	34%
22	FUAC5_MCINS	-	-	-	-	7'000	0%	34%
23	FUAC5_MEGAEDGE	-	-	-	-	9'000	0%	34%
24	FUAC5_MHD	-	-	-	-	9'000	0%	34%
25	FUAC5_MOTISANI	-	-	-	-	13'000	0%	34%
26	FUAC5_ORB5ELE	-	-	-	-	73'000	0%	34%
27	FUAC5_PICLS	-	-	-	-	18'000	0%	34%
28	FUAC5_PIXIE3D	-	-	-	-	3'000	0%	34%
29	FUAC5_SOLDyn4	-	-	-	-	4'000	0%	34%
30	FUAC5_TEMPIS	-	-	-	-	67'000	0%	34%
31	FUAC5_TSVVT421	-	-	-	-	75'000	0%	34%
32	FUAC5_UQMWA3	-	-	-	-	14'000	0%	34%

### Average number of nodes requested by jobs in the wait queue (M100)

Queued nodes (nodes required by queued jobs) are reported as a function of time.

Data come from a cron-driven data collection. Data from log files correctly describes all “pending” jobs (in red), but not “queued” jobs (jobs in Pending state ready to execute). In the figure, negative values mean unavailability of data.

By comparing the “usage” figure with the blue curve (Q nodes), it is easy to see that low level of usage are mainly due to not availability of jobs in the Q state, or system draining in preparation of sched maintenance.



### Usage of the Low Priority queue

EUROfusion users with no projects defined on M100 can get access to the cluster for testing and benchmarks just asking to the User Support team.

They can run only in lowprio mode, using the project FUAC5\_LOWPRIO.

Users with active project on M100 too can use the lowprio project if they wish. In the following the usage for the month.

FUAC5\_LOWPRIO: (902 jobs)

Pui Wai Ma Culham Centre for Fusion Energy - UK	497 jobs	(also collaborator of FUSIO_ru5CCFE)
Koen Strien Technische Universiteit Eindhoven	1 jobs	(also collaborator of FUAC5_GWNN)
Gabriele Merlo Swiss Plasma Center	368 jobs	(also collaborator of FUAC5_LHPED21)
Alessandro di Siena Max Plank Institute	36 jobs	(also collaborator of FUSIO_ru5IPP)

### Usage of the Pre-emption queue

The preemption queue insists on the EUfus partition and is opened to users not in the EUROfusion community.

To access this queue (that is free of charge but with the drawback of possible kill of the job) users must ask for the *m100\_usr\_preemp* partition in the SLURM job. In the following the usage in the month:

- Num of jobs: 8361
- 436 of them by the Cineca staff (for testing)
- 111.329 core-h (3.479 node-h) in total
- 536 jobs were preempted

Queue/partition	Node-h	%
M100_fua	59'445	87.0 %
M100_usr_preempt	3'479	5.1 %
Total on fua partition	62'924	92.1 %
Available	68'324	

## Usage of the Gateway partition

what	Used node-h/Avail node-h	value	KPI
Usage/(available) – GW	12'065 / 17'161	70.3%	KPI-4d

*Queues on Gateway partition 17'613*

Partition	Num. of jobs	Used node-h
gw	3000	12'063
gwdbg	404	3

## Usage of the MARCONI file systems

	<i>size</i>	<i>used</i>	<i>used%</i>
home	101T	35T	35%
CINECA_scratch	2.2P	1.7 P	80%
work	5.9P	3.8 P	64%

General figures (MARCONI)

GB	Jan21	Feb-21	Mar-21	Apr 21	May21	Jun21	Available for EF
home	4'368	4'670	4'731	5'028	5'085	5'211	<b>7.00P</b>
CINECA_scratch	320'485	324'352	285'515	329'718	347'317	335'268	
work	852'524	936'031	983'851	1'103'139	1'133'040	1'178'887	
Gateway	631'808	631'808	631'808	631'808	631'808	631'808	<b>0.600P</b>
<b>Gran tot</b>	<b>1'488'700</b>	<b>1'572'509</b>	<b>1'620'390</b>	<b>1'739'975</b>	<b>1'769'933</b>	<b>1'815'906</b>	<b>7.600P(***)</b>

(\*) since sept 19, the gw storage is set to the total storage available on the EFGW server.

(\*\*) Starting from Apr 2017 "scratch" filesystem is not considered in the "Gran tot" figures

(\*\*\*) From Jul17: "Scratch" area = 10PB raw = 7.0PB net, 600TB grated to Gateway

## Usage of the M100 file systems

	<i>size</i>	<i>used</i>	<i>used%</i>
home	205 T	12 T	6%
CINECA_scratch	1.8 P	1.3 P	72%
work	2.3 P	689 T	31%
Work_AI	109 T	1.9TG	2%

General figures (M100)

	Jan21	Feb21	Mar21	Apr21	May21	Jun21
home	243	294	322	415	452	469
CINECA_scratch	8'276	12'950	22'678	27'202	28'277	26'489
work	13'661	14'904	22'269	16'966	16'770	18'269
<b>Gran tot</b>	<b>13'904</b>	<b>15'198</b>	<b>22'591</b>	<b>17'381</b>	<b>17'222</b>	<b>18'738</b>

Usage by EUROfusion users (MARCONI-100) - GB



## Usage of the long-term archive (DRES)


	<i>size</i>	<i>quota</i>	<i>used%</i>
Archive on-line	(5PB raw)	42 T	
Archive on-tape	10P	903 T	

DRES	Disk/tape	Created	quota
DRES_ruISSPUL	tape	Jan 20, 2017	4TB
DRES_FU_Hatch	Tape	Oct 15, 2018	10TB
DRES_Gateway	Tape	Nov 14, 2018	400TB
DRES_FU_ORBFA	Tape	Feb 21, 2019	50TB
DRES_FU_Chris	Disk	March 20, 2019	2TB
DRES_FU_Merlo	Tape	June 28, 2019	5TB
DRES_FU_g2dsa	Tape	Aug 26, 2019	200TB
DRES_FU_beeke	Tape	Jan 20, 2020	2TB
DRES_FUA34_GG	Tape	Aug 4, 2020	19 TB
DRES_FU_gonza	Disk	Aug 31, 2020	20 TB
DRES_MEGAEDGE	Disk	Nov 6, 2020	20 TB
DRES_FU_InZig	Tape	Apr 14, 2021	10 TB

*No new DRES in the month*

# Characteristics of Usage

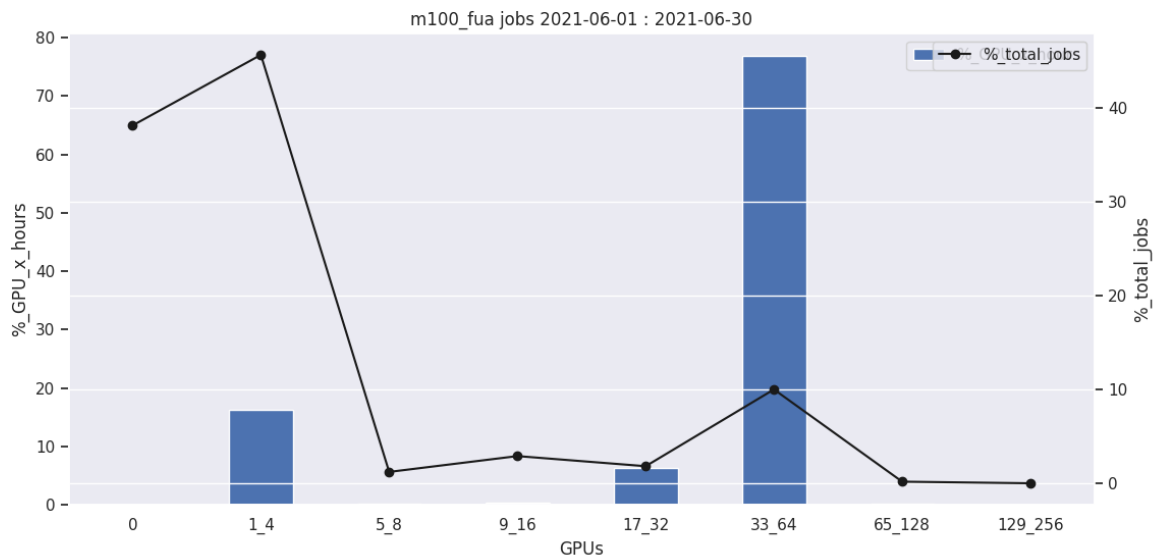
## Job parallelism on Marconi-fusion (Marconi - A3)



Node_h	Tot_jobs	Req-cpus	Alloc-cpus	%_req_alloc	
1	67'306	25'733	347'134	1'224'883	28.34%
2_9	294'379	11'975	2'117'258	2'394'626	88.42%
10_17	260'682	4'765	2'510'699	2'768'470	90.69%
18_33	406'220	2'108	2'235'151	2'661'744	83.97%
34_65	440'913	865	1'455'376	2'232'242	65.20%
66_129	226'690	335	1'471'568	1'691'364	87.00%
130_256	40'555	23	216'960	216'960	100.00%
257_512	0	0	0	0	0.00%
> 512	0	0	0	0	0.00%

# Job parallelism on Marconi-100 (GPUs usage)

Jun 2021: 4.065 jobs



GPU range	Job count	% job count	GPU-h	%GPU-h
0	1'552	38.18%	-	0.00%
1_4	1'858	45.71%	35'425	16.22%
5_8	49	1.21%	409	0.19%
9_16	118	2.90%	708	0.32%
17_32	74	1.82%	13'594	6.23%
33_64	407	10.01%	167'893	76.89%
65_128	7	0.17%	327	0.15%
129_256	-	0.00%	-	0.00%



May 2021: 4.532 jobs

# Availability of the IT infrastructure

## External network

We report the figures for the month of reference, taken from the site of GARR, the Italian Research & Education Network ([https://gins.garr.it/Statistics/xtrail\\_select.php?view\\_type=PoP](https://gins.garr.it/Statistics/xtrail_select.php?view_type=PoP)). The data refer to the 2 Points of Presence (PoP) giving access to CINECA, one of them normally active and the other one available as backup in case of failure of the primary one.

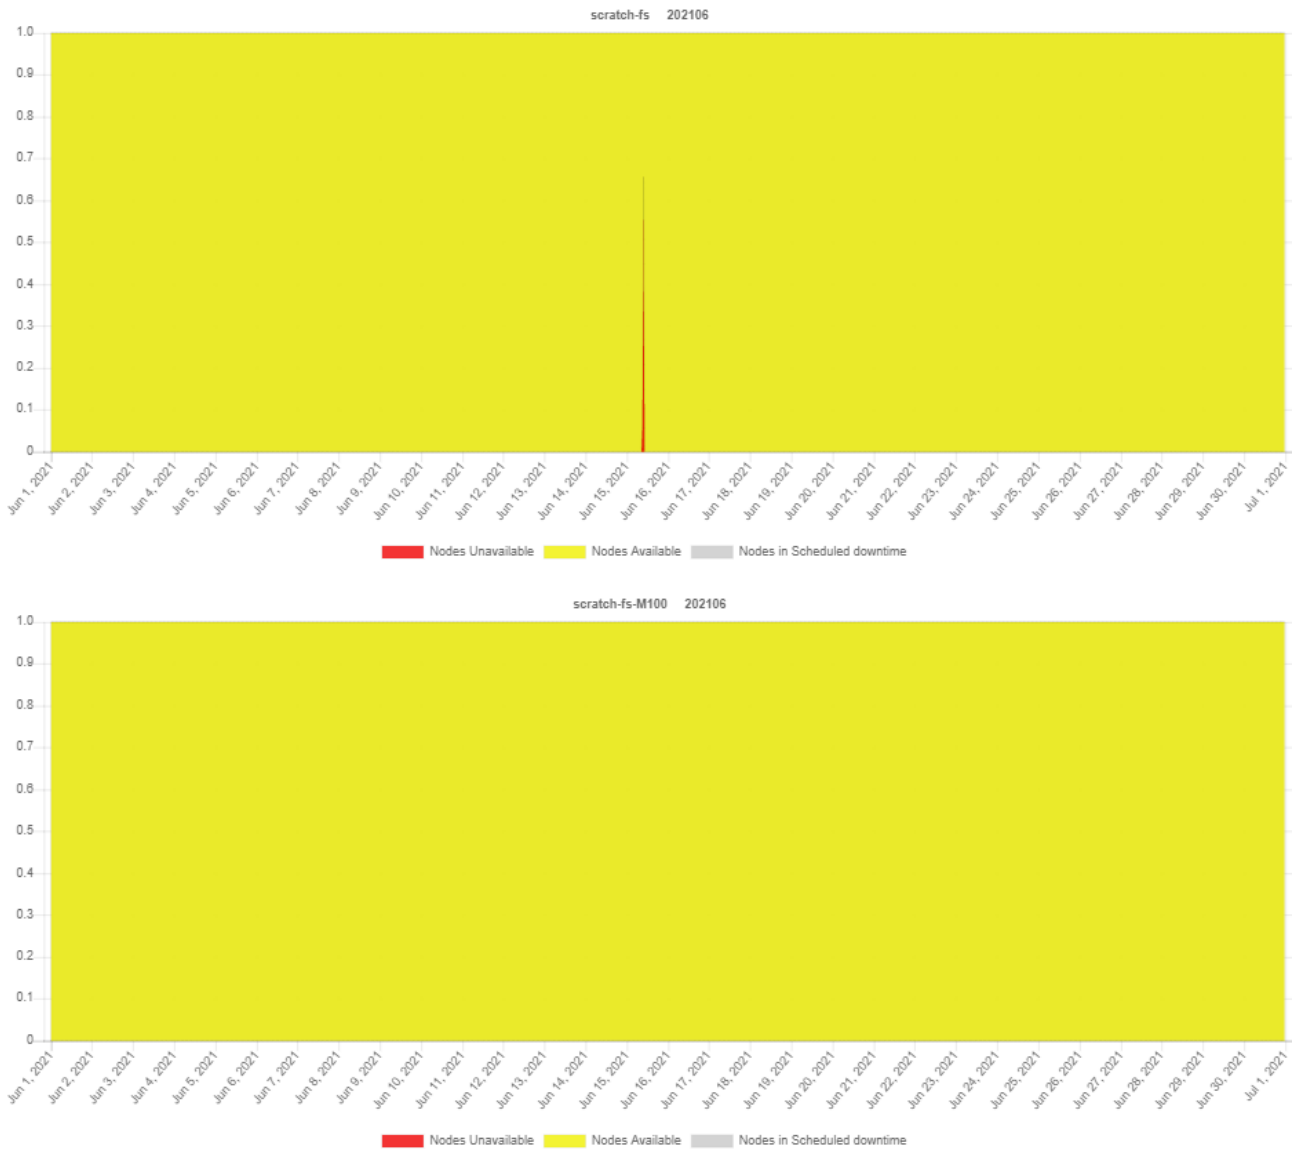
Since April 2021, this is a 20 Gbps network with a guaranteed bandwidth of 15 Gbps reserved to HPC access and data movement.



The availability of the network is complete and it has never been saturated. The average traffic (both in input and output) is around 1-3Gb/s with one peak values near 16 Gb/s.

# Scratch filesystem MARCONI & M100

The storage directly connected to MARCONI and that connected to M100 (home, scratch and work) is provided via IBM Spectrum Scale. The figures below measure the availability of the storage from the login nodes for the 2 clusters



Maintenance: 0/ 0 h

Availability: 99.9/ 100 % (coverage: 100 %)

# Long-term Archive Marconi & M100

The Archive storage is hosted on another infrastructure and exported through a couple of NFS servers in HA configuration on the login nodes of MARCONI and M100 as: /gss/gss\_work

In order to monitor the availability of the archive, we will use the same criteria, as the “scratch filesystem” but adding the availability of the NFS service (in logical AND).



Maintenance: 4/4 h

Availability: 100/100% (coverage: 100%)

# Summary IT infrastructure

Coverage= 100%	Avail	KPI
External Netw	OK	
“scratch”	99.9/ 100 %	KPI-5
Long-term archive	100 / 100 %	

Coverage100%	Sched Maint	KPI
External Netw	0 h	
“scratch”	0 / 0 h	KPI-6
Long-term archive	4 / 4 h	

# Service provided to EUROfusion

## Statistics on the ticketing system

Help Desk service available only half time in the central week of August (August 17-21).

### Tickets to the HelpDesk

We report ticket solved, arrived, and pending for EUROfusion. Moreover, we also report % of solved ticket in the month in less than 30 days and less than 5 days, as required by the KPI.

	Jan 21	Feb21	Mar21	Apr21	May21	Jun21	Jul21	Aug21	Sep21	Oct21	Nov21	Dec20
Tickets tot solved	435	445	635	527	388	443						501
arrived	499	439	631	446	369	486						476
pending	190	189	183	126	75	160						114
Ticket EUfus solved	54	60	116	86	65	77						49
arrived	58	58	114	71	54	74						47
pending/stalled	22	27	28	20	9	16						13
% solved (KPI-11)												
< 5 gg (50%)	88.89%	81.67%	88.79%	79.07%	66.15%	87.01%						79.59%
< 30 dd (80%)	98.15%	93.33%	95.69%	93.02%	86.15%	93.51%						95.92%


## EUROfusion

## June\_2021


<b>TOT (HD)</b>				<b>Ticket efficiency (TOT - HD)</b>		KPI 11	
TOT Solved tickets		77					
TOT Created tickets		74		KPI 12- %Ticket resolved within 5 gg			87.01
TOT Pending tickets		16		KPI 12- %Ticket resolved within 30 gg			93.51
<b>First Level Solved (HD) or escalated</b>				KPI 9		<b>Second Level Solved (HD)</b>	
Solved tickets			49	Solved tickets			28
Escalated to 2nd level tickets			28	resolution time (avr, max)	235		2230.65
resolution time 1st level (avr, max-h)	7.86		17.96	first answer time (avr, max-h)	60.7		447.5
escalation time to 2nd level (avr/max-h)	13.76		32.37	out of SLA (1 nbw)			4
out of SLA (1 nbd)			2				
<b>TOT (Sys-Man)</b>				KPI 7			
		tot	from HD				
TOT Solved tickets		49	16				
TOT Created tickets		57					
response time (avr-min) on ticket from HD			303				
out of SLA (30 min)			9				




### Solved tickets (all HD tickets)




### Type EUROfusion - 1st + 2nd level



### Type EUROfusion - 1st level



### Type EUROfusion - 2nd level



## Provisioning: statistics on the username/project creation (KPI 8)

The data regarding the month are not available due to a problem in the data collecting script. The cumulative data will be reported next month

Month	Users on UserDB	New HPC users in the month	Response Time (avr h)	Response Time (max h)	> 1 day
Jan 21	689	4	0.71	1.33	0
Feb	714	16	1.34	6.2	0
March	731	20	1.22	6.7	0
April	747	12	2.38	16.2	0
May	752	7	1.61	7.68	0
June	757	7	1.86	10.42	0
Jul					
Aug					
Sep					
Oct					
Nov					
Dec 2020	680	3	0.48	1.1	0

	Ref.	Value	Note
<b>HPC infrastructure</b>			
KPI-1a Availability Conventional (A3)	97%	99.7 %	OK
KPI-1b Availability Acc (M100)	97%	96.4%	OK
KPI-1d Availability Gateway	97%	99.3 %	OK
KPI-2a Maintenance Conventional (A3)	8h/m +7dd/y	4 h (0 h extra)	Remaining: 165.5 h (*)
KPI-2b Maintenance Acc-1 (M100)	8h/m +7dd/y	4 h (0 h extra)	Remaining: 134.5 h (*)
KPI-2d Maintenance Gateway	8h/m +7dd/y	0 h (0 h extra)	Remaining: 168 h (*)
KPI-3a Incidents Conventional (A3)	100		To be defined
KPI-3b Incidents Acc (M100)	100		
KPI-3d Incidents Gateway	100		
KPI-4a Usage Conventional (A3)	85% 60%	86.9 %	OK
KPI-4b Usage Acc (M100)	-	87.0 %	
KPI-4d Usage Gateway	-	70.3%	
<b>IT infrastructure</b>			
KPI-5 Availability - IT	97%	OK 99.9/ 100% 100/ 100%	OK
KPI-6 Maintenance - IT	8h/m + 7dd/y	0 h 0/ 0 h 4 / 4 h	Remaining: 134.5 h (*)

## Services

KPI-7 Sysman response time	30 min	303 min (avr)	(9 tick > 30min) out of 16
KPI-8 Provisioning	1 nbd	1.9 h (avr)	OK
KPI-9 1st level time to solve or transfer	1 nbd	7.9 h – 13.8 h (avr)	(2 tick > 1nbd) out of 77
KPI-10 2nd level time for expert report	1 business week	60.7 h (avr)	(4 tick > 1 nbw) out of 28
KPI-11 ticket global efficiency	Less 5 dd (>50%) Less 30 dd (>80%)	87.01% 93.51%	OK

(\*) remaining hours of maintenance recomputed for year 2021

# Appendix

## Update on the allocation of GPU jobs in M100 (E. Rossi)

This issue is closed.

The abnormal behaviour found by Leo was recognised, it has been reported to SLURM developers, they understand the problem and told us it will be fixed in the next SLURM release, expected in August-September. As soon as it will be released as a stable version, it will be installed and users informed

## Top10 users of the low priority queue on M100 (E. Rossi)

qos\_lowprio: (902 jobs +4)

Pui Wai Ma Culham Centre for Fusion Energy - UK	497 jobs (+4)	2 411 node-h
Koen Strien Technische Universiteit Eindhoven	1 jobs	0.01 node-h
Gabriele Merlo Swiss Plasma Center	368 jobs	36 619 node-h
Alessandro di Siena Max Plank Institute	36 jobs	4 803 node-h



# Feedback from Ticket Meeting #55 and update on network(s) status

F. Boillod-Cerneux (Chair), R.Hatzky (HLST) and J.David



This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

## Last Ticket Meeting was #54 (16th June 2021)



Summary from slides presented & Susana's minutes

### Status of systems: June 16<sup>th</sup> – July 14<sup>th</sup>

June 17: quota imposed on \$CINECA\_SCRATCH area on Marconi100 cluster

June 24: scheduled maintenance operations on Marconi100 cluster on June 29th

July 7: job level GPU usage and accounting report available on Marconi100 cluster

July 13: interruption of production on dedicated EUROfusion nodes on Marconi cluster due to power outage




Next events: Save the date!

★	9 Septembre 2021 09h00 – 10h30	Webinar #11	By video
	14 September 2021 10h00 – 12h00	TM meeting #56	By video
	20 Septembre 2021 9h30 – 12h30	OC meeting #63	By video
	15 October 2021 10h00 – 12h00	TM meeting #57	By video
	22 October 2021 9h30 – 12h30	OC meeting #64	By video
★	28 October 2021 09h00 – 10h30	Webinar #12	By video
★	22 November 2021 10h00 – 12h00	TM meeting #58	By video



Marconi and Marconi100 Eurofusion Tickets in 2020/21



**No issue and Marconi and Marconi100 are stable  
Thank you CINECA**

Feedback from TM55 to OC61

## Tickets reported to L2



### Escalated tickets:

- Intel ticket 3932: update given by Intel. Investigation is ongoing.
- SchedMD ticket 16506: pending for the slurm update in august
- Lenovo/IBM - Ticket 16019: Giacomo made a complete report to present the investigation related to this ticket and the already closed one. Ticket Committee shall receive this report when validated by CINECA.

### Discussed ticket: 19157

- The user reported two jobs running Euterpe code that wrote no output
- Observed by Serhiy, but non reproducible bugs
- Options:
  - Use sleep to test if the job is corrupted
  - Systematically use `-vvvv` option to launch Euterpe jobs
  - check with slurm support if they observed such behavior and gives us some insights

Feedback from TM55 to OC61

# Sanity checks



	Marconi		
	Production 09.07.2021	Maintenance 08.07.2021	
		Pre-maintenance	Post-maintenance
Linpack	✓	cancelled	cancelled
Stream	✓	cancelled	cancelled
Switches		cancelled	cancelled

Next maintenance for Marconi is scheduled on August 10<sup>th</sup> 2021

Feedback from TM55 to OC61

# Sanity checks



	MARCONI100		
	Maintenance 29.06.2021		Production 14.06.2021
	Pre-maintenance	Post-maintenance	
Dgemm	✓	✓	✓
Stream		✓	✓
Nvlink		✓	✓
p2p		✓	✓

Next maintenance for M100 is scheduled on August 24<sup>th</sup> 2021

Feedback from TM55 to OC61



# HLST's Applications performance curve status



Work done by HLST (Serhiy M.) and PIs on STARWALL, GENE, EUTERPE, ORB5, LAMMPS codes





time in minutes		06/07/2021			
Code name	Marconi partition	test1	test2	test3	avg
STARWALL	SKL 16 nodes	18,76	18,73	25,66	21,05
	SKL 64 nodes	26,2	28,98	24,61	26,59
GENE	SKL 16 nodes	21,68	P7	19,87	20,78
	SKL 16 nodes (iter, in sec)	1,35	P7	1,23	1,29
	SKL 128 nodes	20,36	18,35	18,44	19,05
	SKL 128 nodes (iter, in sec)	3,7	3,69	3,71	3,702
EUTERPE	SKL 16 nodes	16,45	16,58	16,61	16,55
	SKL 128 nodes	38,04	35,18	41,18	38,13
Code name	Marconi100 partition	test1	test2	test3	avg
ORB5	16 nodes, 4 MPI tasks per node, 4 GPUs per node	32,2	32,2	32,38	32,26
LAMMPS	16 nodes, 4 MPI tasks per node, 4 GPUs per node	24,28	24,13	24,03	24,15
GENE	"16 nodes, 4 MPI tasks per node, 4 GPUs per node"	20,05	20,1	19,95	20,03

P7 - Stale file handle: node r146c16s03 was bad. Everything is working fine now.  
 The wall clock time fluctuations on the SKL partition were in common range.

From : HLST own measurements (Serhiy)

Feedback from TM55 to OC61

# HLST's Applications performance curve status




Feedback from TM55 to OC61


# HLST's Applications performance curve status



EUTERPE, SKL, 16 nodes



EUTERPE, SKL, 128 nodes




Feedback from TM55 to OC61


# HLST's Applications performance curve status



GENE, SKL, 16 nodes






GENE, SKL, 128 nodes



Feedback from TM55 to OC61

# HLST's Applications performance curve status



Feedback from TM55 to OC61

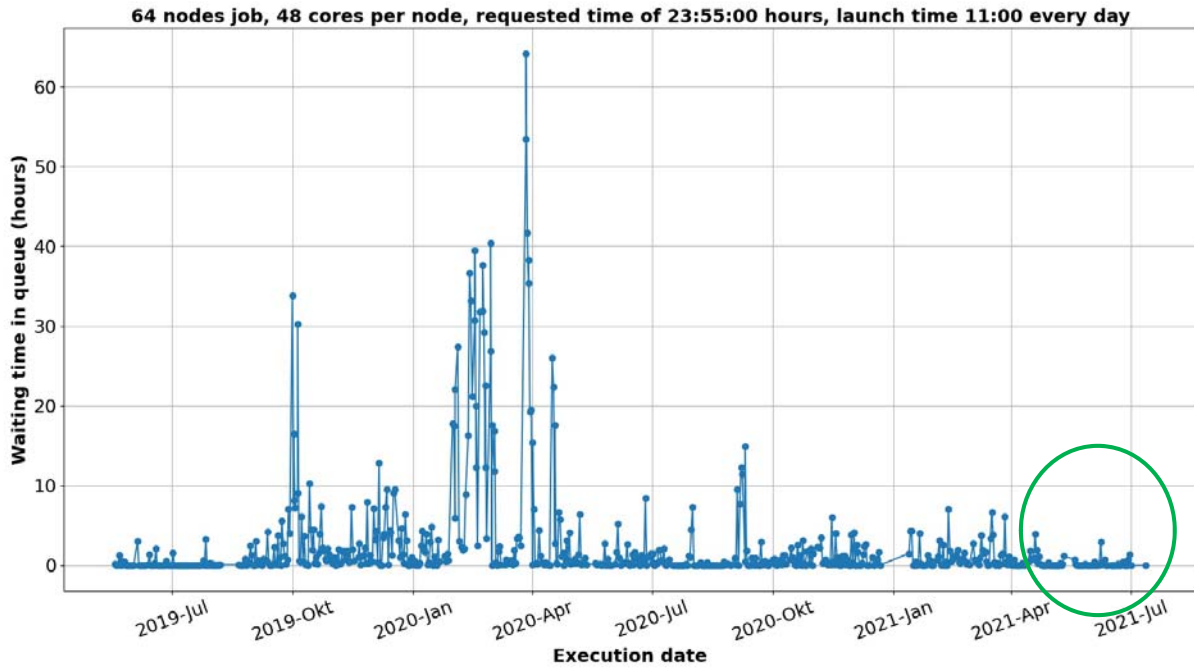
# HLST's Applications performance curve status



- The waiting time to submit all jobs on SKL was within one hour
- It is already usual to have relatively long waiting time (up to three hours) on M100 partition.
- The system starts to be heavily used (look results about node usage next point).

time in hours		06/07/2021			
Code name	Marconi partition	test1	test2	test3	avg
STARWALL	SKL 16 nodes	1,027	1,06	1,08	1,056
	SKL 64 nodes	0,885	0,89	1,08	0,952
GENE	SKL 16 nodes	0,009	0,484	0,654	0,382
	SKL 16 nodes (iter, in sec)	0,238	0,544	0,59	0,382
	SKL 128 nodes	0,755	0,925	0,959	0,88
EUTERPE	SKL 128 nodes (iter, in sec)	0,001	0,001	0,077	0,026
	SKL 16 nodes	1,027	1,06	1,08	1,056
	SKL 128 nodes	0,885	0,89	1,08	0,952
Code name	Marconi100 partition	08/06/2021			
		test1	test2	test3	avg
ORB5	16 nodes, 4 MPI tasks per node, 4 GPUs per node	0	1,553	1,779	1,111
LAMMPS	16 nodes, 4 MPI tasks per node, 4 GPUs per node	0,564	0,972	1,376	0,971
GENE	"16 nodes, 4 MPI tasks per node, 4 GPUs per node"	2,117	2,347	3,075	2,513

Feedback from TM55 to OC61



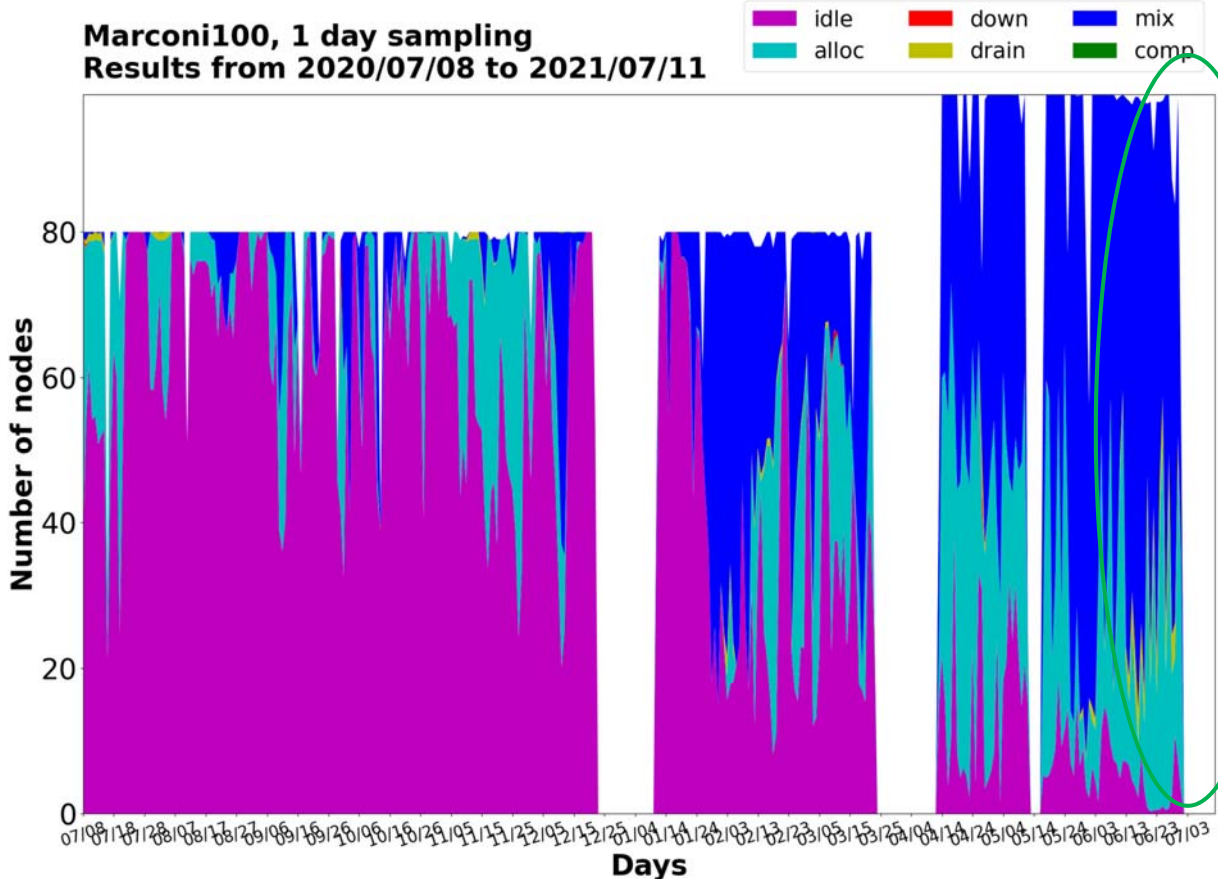
no significantly long waiting time was recorded as well during last month.



From : HLST own measurements (Serhiy)


Feedback from TM55 to OC61

Marconi100 usage 2021



In the Marconi100 node allocation very positive trend can be observed during last month.

Feedback from TM55 to OC61



Good usage of SKL partition




From : HLST own measurements (Nils)

Feedback from TM55 to OC61

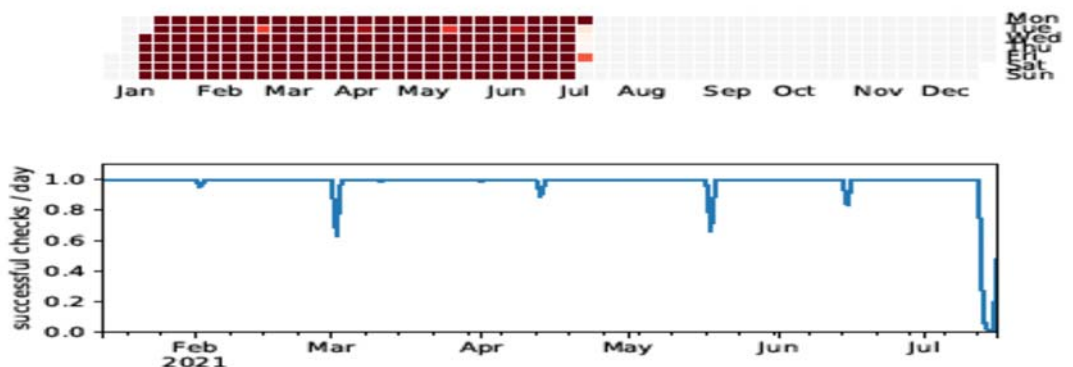
## Uptime / FS uptime



### Uptime 2021



### FS Uptime 2021



From HLST's Nils.M measurements


Feedback from TM55 to OC61

External network (ssh to marconi login nodes) :


marconi			
year	month	mean_up	#days
2021	5	98,8%	31
	6	99,4%	30
	7	99,8%	13



marconi : jdad / previous (june 2021) & current month (july, till 2021-07-13)



marconi : jdad / since year 2020




From jdad's home workstation ssh measurements.


External network (ssh to m100 login nodes) :



m100 : jdad / previous (may) & current month (july, till 2021-07-13)



For comparison : marconi : same period



From jdad's home workstation ssh measurements.



## Likwid:

- Tests are in progress (CINECA working with Nils)

## HPCMD:

- Installation is in progress.

Feedback from TM55 to OC61

Operation Committee meeting OC-61 – July 23, 2021



## Close of TM report

F. Boillod-Cerneux (Chair), R.Hatzky (HLST) and J.David






			SKL (A3)	Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
			<i>Raw</i>						
Availability	A3+	Corrected		99,1%	98,7%	99,1%	98,7%	99,4%	99,7%
	login	Target (SLA)		97,0%	97,0%	97,0%	97,0%	97,0%	97,0%
Usage	A3	Real		87,9%	96,2%	96,4%	88,3%	83,0%	86,9%
		Target		85,0%	85,0%	85,0%	85,0%	85,0%	85,0%
			M100 (C1)	Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
			<i>Raw usage</i>						
Availability	C1+	Corrected		100,0%	98,7%	98,8%	98,1%	99,8%	96,4%
	login	Target (SLA)		97,0%	97,0%	97,0%	97,0%	97,0%	97,0%
Usage	C1	Real		10,6%	49,4%	34,2%	33,1%	51,7%	87,0%
		Target		none	none	none	none	none	none
				Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
Major incidents (from monthly re				1	1	3	3	1	0
				Jan. 2021	Feb. 2021	March 2021	April 2021	May 2021	June 2021
Maintenance	A3	Month		9,5	4	10,5	5,5	8	4
	C1	Month		5,5	4,5	41,5	8	5,5	4
	A3	Remaining		166,5	166,5	164	164	164	164
	C1	Remaining		168	168	134,5	134,5	134,5	134,5


## Availability




Availability (excluding lincpack run in May and in Nov. 2018)





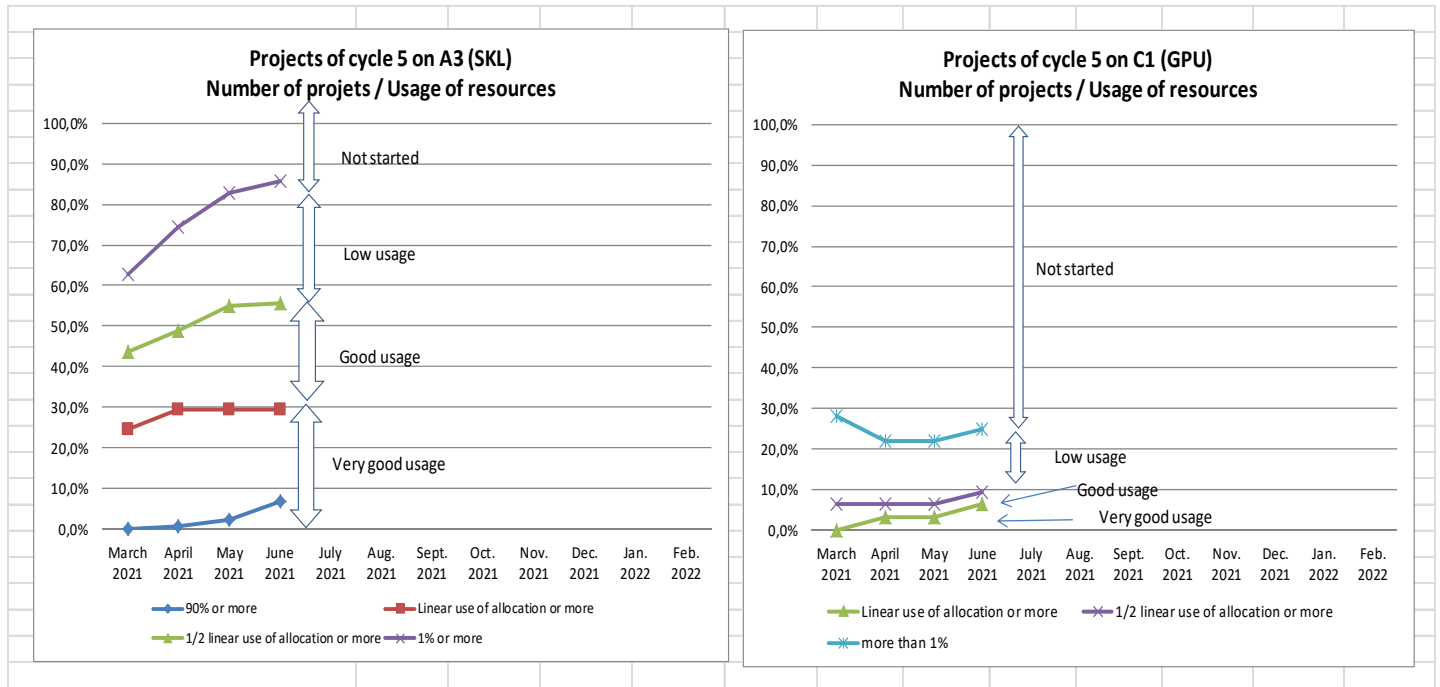


## Failed job (3 codes benchmark performed by HLST)



SKL (A3): STARWALL (16,64), GENE (16,128), EUTERPE (16,128)

GPU (C1): ORB5 (16), LAMMPS (16), GENE (16)



SKL (A3): STARWALL (16,64), GENE (16,128), EUTERPE (16,128)  
GPU (C1): ORB5 (16), LAMMPS (16), GENE (16)



## Gateway

# EUROFUSION GATEWAY

## CURRENT STATUS OF GATEWAY OPERATIONS

### Monthly report June 2021

F.Iannone, G.Guarnieri (ENEA)

A.Federico, A.Vanni (CINECA)

## GATEWAY OPERATIONS: CURRENT STATUS

### ✓ Computing resources:

#### **28 SKL nodes as independent HPC system**

- 16 compute nodes with 192 GB of RAM
- 8 compute nodes with 384 GB of RAM
- 4 permanent login nodes (s51....s54) with 192 GB of RAM and 2 NVIDIA TESLA K80 in each node

#### **Management nodes:** 12 LENOVO NeXtScale nx360 as management nodes

- 2 Nodes as SLURM controllers
- 2 nodes as OPA fabric subnet manager
- 3 AFS servers
- 1 ITM Mysql Database (itmmysql1)
- 1 Flex License Manager (s33.eufus.eu)
- 13 Virtual nodes for Gateway services (Authentication/Authorization, Portal,Gforge, Catalog,...)

### ✓ SLURM PARTITIONS & QoS:

partition	Qos	#cores	max walltime (default)	max memory	priority	Notes
gw	-	1056 (22 nodes)	48 h	unlimited	-	main partition
gw	normal	256	48 h	376 GB	0	max 256 cores per user
gw	qos_gwlong	144	144 h	376 GB	40	max 144 cores per user , max 2 job per user
gwdbg	-	96 (2 nodes)	30 m	unlimited	-	debug partition
gwdbg	normal	96	30 m	376 GB	0	max 96 cores per user

## GATEWAY OPERATIONS: CURRENT STATUS

✓ 13 Virtual Nodes including the following Gateway User Services:

node	core	ram	disk	description
s01	1	4	30	DNS master (dominio eufus.eu)/MV sysdoc/Slave krb5/AFS DB server/AFS Fileserver
s02	1	2	40	krb5/AFS, DBserver/AFS,Fileserver/AFS, Backup,scheduler/LDAP, master/Admtools server (cli)
s03	1	2	40	Mail server/MV mail.eufus.eu/MV wiki.eufus.eu/MV ticket.eufus.eu
s04	1	2	40	Identity provider
s05	1	2	40	Admtools
s07	1	2	40	Node for VM cloning
s08	2	16	40	Portal
s09	4	32	40	catalog (ex gforge6-test)
s10	4	32	40	gforge6 (still up)
s11	4	32	40	gitlab
s12	4	32	200	gforge-next: new gforge server
s13	4	32	120	gforge-test : test for gforge-next
s88	8	45	36	swimas: containers builder with fakeroot setup (g2tomz)
<b>Tot.</b>	<b>36</b>	<b>233</b>		

✓ PIA Phase II:

provide at least 10 virtual hosts for ITM Gateway services with (on average) 4 core/16 GB

✓ In red VMs under CPT management (means as root)

## GATEWAY OPERATIONS: CURRENT STATUS



✓ Storage usage:

- Home/projects (AFS): **6.0 TB** - Home (GPFS): **2.8 TB** (cindata) - GPFS: **517TB**

Tbytes	AFS	GPFS/HOME	GPFS	work	scratch
Jul.2020	4.6	2.6	483	125	355
Aug.2020	4.6	2.6	490	129	356
Sep.2020	4.7	2.6	511	136	369
Oct.2020	5.6	2.6	529	142	384
Nov.2020	5.7	2.6	544	153	389
Dec.2020	5.7	2.6	566	164	392
Jan.2021	5.7	2.6	602	176	419
Feb.2021	5.8	2.6	483	197	279
Mar.2021	5.8	2.8	483	210	265
Apr.2021	5.8	2.8	495	223	265
May 2021	5.9	2.8	505	232	265
Jun.2021	6.0	2.8	517	244	266

Quota Gateway:  
**700 TB (usage 74%)**

GPFS AREA:  
~10 users with 77% of usage




## GATEWAY USERS ACCOUNT: CURRENT STATUS

### ✓ Gateway Users

- **417** users

Gateway Users July 2021



### ✓ Accounting data:

#### Country

- **118** users never login in the Gateway (since 2017) (checking perl5 directory in homedir)

## GATEWAY CURRENT ISSUES, ACTIONS & INFOS

### ✓ GATEWAY SLURM issue (open/stand-by)

- Jobs submission in AFS
  - Gateway Users can choose to move your homedirs from AFS to GPFS and they use AFS storage areas for files sharing on WAN

### ✓ GATEWAY repositories access for data integrity (No UPDATE)

- VMs: PORTAL, GFORGE6, CATALOG and GFORGE-NEXT can access to the repository directories
  - ✓ D.Coster, R.Coelho, D.Figat, M. Owsiak, V.Pais
- AFS pts group: *itm:admins* can access to the repository directories
  - ✓ D.Coster, G.Manduchi, B.Palak, T.Zok, D.Figat, Mi. Owsiak, V. Pais
  - ✓ J.Signoret, O.Hoenen, M.Plociennik, D. Yadin, L.Fleury, T.Aniel, F.Imbeaux
- Recovery strategy for repository directory corruption
  - ✓ Restore on a specific date from backup (if someone please reports a problem)

### ✓ Scientific Library installation on Gateway

- new version of HDF5 1.10.4 for Intel18/IMPI has been installed on GUB request

### ✓ SSL certificate installation

- SSL certificate for gforge-test (VM) has been installed on D.Figat request

### ✓ Gateway Extension

- Additional 64 compute nodes detached from MARCONI Fusion A3 partition
- To assess migration to Redhat 8:
  - needs a testbed layout: 1 login nodes + 4 compute nodes + storage area (cineca software)
- CINECA has assessed feasibility costs and time (discussion on the last slides)

## GATEWAY MONTHLY REPORT: June 2021

### ✓ Data of availability & Usage of Compute Nodes :

- (node-h: T=17280, A=17160, U=12065: Avail: **99.3%**, Usage: **70.3%**)
  - Scheduled maintenance : 0 h
  - Gateway incident: 0

#### Jobs Statistics

Partitions/QOS	#job	nodes-h	core-avr	Q-h avr
gw/nomal	3000	12063	15.8	0.009
gw/qos_gwlong	0	0	0	0
gwdbg	404	3	5.9	0

	#job	nodes-h	core-avr	Q-h avr	Avail.%	Usage %
Jun-2020	3468	8574	9.1	0.003	99.9	49.6
Jul-2020	5216	11931	7.9	0.289	97.8	68.3
Aug-2020	3663	10479	18.3	1.4	100	58.8
Sep-2020	7943	9588	18	0	100	55.4
Oct-2020	4205	9477	11	0	99.3	53.7
Nov-2020	4928	10500	10	0.002	98.5	62.1
Dec-2020	6217	11076	5.9	0	99.8	62.3
Jan-2021	5463	11593	8	0.001	99.3	65.4
Feb-2021	3200	9952	11	0.001	99.7	59.7
Mar-2021	4040	12218	8	0.3	99.1	68.5
Apr-2021	3229	12413	15.9	0.003	100	71.8
May-2021	3680	12240	19.7	0.005	100	71,0
<b>Jun-2021</b>	<b>3404</b>	<b>12065</b>	<b>15.8</b>	<b>0.009</b>	<b>99.3</b>	<b>70.3</b>

## GATEWAY MONTHLY REPORT: June 2021

### Jobs Statistics for Users

Partitions/QOS	#job	nodes-h	core-avr	Q-h avr
gw/nomal	3000	12063	15.8	0.009

user	# jobs	node-h	core-avr	Q-h avr	jobname	name	surname
g2iys	109	3530	35.3	0.003	?????	Ilya	Senichenkov
g2dfan	245	2487	20.0	0.000	?????	Dongmei	Fan
g2snewto	1149	2374	7.0	0.002	?????	Sarah	Newton
g2rcoelh	613	997	24.1	0.005	CK3	Rui	Coelho
g2ccowl	71	993	16.0	0.003	SOLPS-ITER	Cyd	Cowley
g2wuhsh	222	671	8.8	0.001	SOLPS-ITER	Haosheng	Wu
g2grubi	52	291	11.4	0.000	?????	Giulio	Rubino
g2auccel	18	206	16.0	0.000	SOLPS-ITER	Andrea	Uccello
g2gfrazz	97	130	61.1	0.001	?????	Gino	Frazzoli
g2iboro	20	120	24.0	0.001	SOLPS-ITER	Irina	Borodkina
g2efran	33	93	16.0	0.000	?????	Emil	Fransson
g2gpelka	17	68	36.1	0.004	?????	Grzegorz	Pelka
g2jofe	12	56	25.3	0.001	ETS	Jorge	Ferreira
g2galbe	34	33	15.3	0.000	ERO2.0	Gabriele	Alberti
g2rdoyl	17	23	20.5	0.000	?????	Rhys	Doyle
g2subba	32	20	1.0	0.000	SOLPS-ITER	Fabio	Subba
g2mporad	23	16	16.0	0.015	?????	Michele	Poradziskin
g2SSIPIL	4	11	72.0	0.004	ASCOT4PAR	Seppo	Sipila
g2xiang8	81	6	1.0	0.000	B2.5	Lingyan	Xiang
g2diy	16	3	19.0	0.000	?????	Dmitriy	Yadikin
g2rtatsu	2	0	8.0	0.000	SOLPS-ITER	Ryoko	Tatsumi
g2sdenk	4	0	16.0	0.012	?????	Severin	Denk
g2tjohns	13	0	28.3	0.000	?????	Thomas	Johnson
g2mjabl	2	0	18.0	0.000	?????	Malgorzata	Jablczynska
g2fianno	2	0	11.0	0.000	?????	Francesco	Iannone
g2msci	9	0	2.0	0.000	?????	Massimiliano	Sciscio
g2ahnie	2	0	4.0	0.000	HESEL	Anders Henry	Nielsen

## GATEWAY MONTHLY REPORT: June 2021

### ✓ Data of Usage of 4 Login Nodes : (s51/s52/s53/s54)

▪ (node-h: T=2880, A=2880, U=271-414, Avail: 100%, Usage: 9.4-14.3%)

- Scheduled maintenance :0 h
- Gateway incidents: 2
  - NAG license issue on Jun.8<sup>th</sup>
  - LDAP/SSSD instability started on Jun.14<sup>th</sup>  
(until Jun.28<sup>th</sup>) (frequently restore of LDAP entry)

Statistics	#users	Cputime / usage (%) nodes-h	logintime [h] <sup>(1)</sup>
Jun-2020	68	218.7-265 / 7.5-9.2%	42710
Jul-2020	56	35.2-74.7 / 1.2-2.5 %	25216
Aug-2020	44	14-62.1 / 0.4-2 %	14480
Sep-2020	50	36-172 / 1.2-6 %	40803
Oct-2020	54	36-298 / 1.2-10%	51350
Nov-2020	60	60-79 / 2.1-2.4 %	19650
Dec-2020	42	56 -88 / 1.8-3 %	23465
Jan-2021	37	17-34 / 0.5-1.1 %	29359
Feb-2021	46	16-59 / 0.5-2.1 %	23127
Mar-2021	53	84-148 / 3.3-5 %	46607
Apr-2021	49	17-50 / 0.6-1.7 %	22632
May-2021	43	115-153 / 3.8-5 %	16497
Jun-2021	51	271-414 / 9.4-14.3	17997

#username	node-hours	login hours	#username	node-hours	login hours
<b>g2tjohns</b>	<b>192</b>	<b>0.48</b>	g2orm	0	0.41
g2bpalak	18	373.3	g2fdejo	0	18.79
g2past	13	0	g2snowak	0	5.27
g2jofe	12	994.99	g2xiang8	0	2.58
g2rcoelh	6	1642.12	g2ssipil	0	2338.47
g2diy	4	0.73	g2iys	0	2.28
g2dfigat	3	0.08	g2mkozei	0	1.47
g2dfan	3	999.44	g2mporad	0	20.5
g2fnapo	3	1149.16	g2maradi	0	0
g2ccowl	2	130.64	g2mjabl	0	1137.03
g2lfleur	2	8.15	g2gpouli	0	0.23
g2phuy	2	3.57	g2fianno	0	544.72
g2dpc	2	2686.67	g2chmiel	0	19.92
g2wdk	2	2.61	g2subba	0	562.14
g2tslend	2	224.57	g2thayw	0	1139.08
g2galbe	1	2.82	g2tomz	0	5.07
g2rdoyl	1	0.05	g2gpelka	0	1241.85
g2ahnie	1	0	g2tfarm	0	2.24
g2fpi	1	5.42	g2lfrass	0	0.32
g2sdenk	1	42.07	g2dmoult	0	7.98
g2wuhsh	1	811.82	g2sdixon	0	0.47
g2skorte	0	5.19	g2apopa	0	5.99
g2michal	0	61.07	g2fiann	0	2.22
g2rtatsu	0	3.99	g2oluk	0	517.81
g2iboro	0	1271.88	g2bpogo	0	0.01

/var/account/pacct daily files for cpu time (command sa -m)

/var/log/wtmp daily files for login session time (commad ac -p)

<sup>(1)</sup> More login sessions for each users not always active

## GATEWAY MONTHLY REPORT: June 2021

### ✓ Currently Ticket Statistics

Date Created	Subject	From	Priority	Help Topic	Current Status	Last Updated	Due Date	Overdue	Answered	Thread Count	Attachment Count	Response time (day)
29/06/2021 11:06	Data access	Francois Dejoie	Normal	Storage areas	Closed	07/07/2021 12:59	06/07/2021 11:06	0	1	3	0	8
24/06/2021 06:52	Certificate for gforge-test	Daniel Figat	High	Problem	Closed	29/06/2021 10:08	25/06/2021 06:52	0	1	3	1	5
27/06/2021 11:56	File system problems	Rui Coelho	Normal	User Environment	Closed	28/06/2021 11:12	04/07/2021 11:56	0	1	3	0	1
24/06/2021 09:45	session is frozen s52	Dmitriy Yadin	High	Access Issue	Closed	24/06/2021 10:00	25/06/2021 09:45	0	0	4	1	0
21/06/2021 11:58	Unable to allocate resources on partition	Malgorzata Jablczynska	High	Access Issue	Closed	23/06/2021 09:30	22/06/2021 11:58	0	0	8	2	2
18/06/2021 11:35	Renew the license on Matlab	Ryoko Tatsumi	Normal	System Software	Closed	21/06/2021 19:14	25/06/2021 11:35	0	1	2	1	3
21/06/2021 12:56	Some of the processes are defunc	Michal Owsiak	Low	General Inquiry	Closed	21/06/2021 19:12	28/06/2021 12:56	0	1	2	0	0
08/06/2021 09:59	Login shell change	Francois Dejoie	Normal	User Environment	Closed	17/06/2021 16:05	15/06/2021 09:59	0	1	16	0	9
16/06/2021 07:35	Issues on gforge-next	Daniel Figat	High	Access Issue	Closed	16/06/2021 12:19	17/06/2021 07:35	0	1	6	0	0
08/06/2021 12:03	nag library issue	Dmitriy Yadin	Normal	System Software	Closed	15/06/2021 18:10	15/06/2021 12:03	0	1	6	0	7
08/06/2021 10:56	Gforge-Next Access	Thomas Farmer	High	Access Issue	Closed	15/06/2021 18:07	09/06/2021 10:56	0	1	3	0	7
14/06/2021 14:52	SSH access fails	Thomas Hayward	High	Access Issue	Closed	14/06/2021 17:34	15/06/2021 14:52	0	1	2	0	0
03/06/2021 15:27	problems with environment s53	Dmitriy Yadin	Normal	User Environment	Closed	05/06/2021 08:45	10/06/2021 15:27	0	1	4	0	2

### Tickets:

- Some tickets depend on the LDAP/SSSD instability issue

Opened	Closed	Reopened	Service Time	Help Topic	Opened	Overdue	Closed	Reopened
<b>13</b>	<b>13</b>	<b>4</b>	84% in 1 week	General Inquiry	1	0	1	0
			100% in 9 days	Feedback	0	0	0	0
				Problem / Access Issue	5	0	5	2
				System Software	2	1	2	0
				Applications	0	0	0	0
				Storage areas	1	1	1	0
				User Environment	3	1	3	2
				Job Submission	0	0	0	0
				Problem	1	0	1	0


# GATEWAY EXTENSION

- ✓ ENEA shall reply to EF Letter TD/RK-21-55
  - A. Move to Gateway MARCONI Fusion (A3/SKL) compute nodes (64)
  - B. Provide modern software development environments:
    - i. To extend the Gateway VM Infrastructures (additional 32 VM (4 cores/32 GB RAM / 200GB HD))
    - ii. DevOps tools (Gitlab/Wiki/Jira/CI-CD, Containers builder envs) in charge to CPT/ACH ? (gforge model including provisioning)

## ITEM A (additional hardware resources)(Costs & Timescale)

- Add n.64 compute nodes detached from MARCONI Fusion A3 (SKL) partition
- Add n.1 Core Switch OPA (already available need to extend maintenance)
- Add n.2 GEth Switches (already available need to extend maintenance)
- Add Cabling: 8 QSFP link 20m , 16 CU link, 2 LC fibre 20m, 1 DAC 1m
- Add GSS-26 storage for GPFS up to 2.4 PB raw (already available need to extend maintenance)
- Installation and configuration

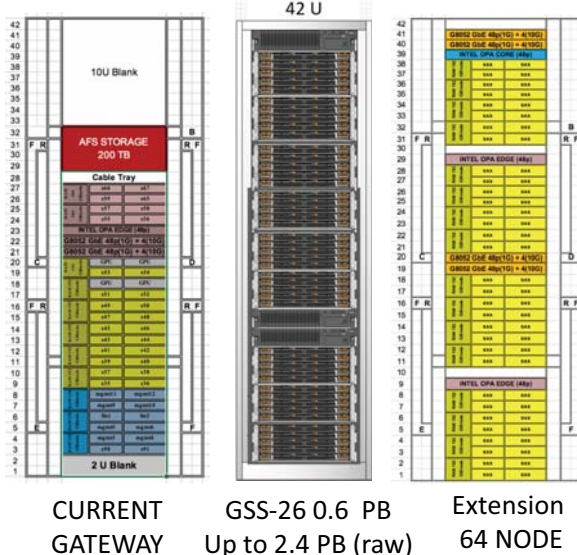
96 ports Single -rail for 96 nodes/hosts, 2 tier, 2:1 oversub fat-tree



### Costs (100% funding):

- Maintenance contract extension: **21 kEuro/year**
- New cabling: **8.5 kEuro**
- Installation & Configuration: **2 kEuro**

**Timescale: Jan.1<sup>st</sup> 2022**





## ITEM B (additional Gateway Services)(Costs & Timescale)

### Costs (100% funding):

- 32 VM with 4 cores – RAM 32 GB –HD 200GB each: **109.440 kEuro/year**
  - **PaaS** for one Linux image:
    - Centos 7.x or Redhat 8.x (excluding redhat license cost)
- Redhat license subscription:
  - n.4 Front-End + n.32 VM: **7.2 kEuro/year**  
Red Hat Enterprise Linux Server, Standard (Physical or Virtual Nodes)
  - n.2 Head Nodes n. 92 Compute Nodes : **7.8 kEuro/year**  
Red Hat Enterprise Linux Server for HPC Head Node, Standard (Physical or Virtual Nodes)  
Red Hat Enterprise Linux Server for HPC Compute Node, Self-support (Physical or Virtual Node)
- Gateway Migration from Centos 7.x to Redhat 8.x: **50 kEuro**
  - Build a new Linux Redhat Image
  - Rebuild the current CINECA software from Centos 7.x Redhat 8.x (CINECA)
  - Rebuild the PIA software list from CENTOS 7.x to 8.X (ENEA)
  - Support to rebuild ITMENV/IMAS in charge of CPT/ACH????

### Timescale:

**Gateway Migration: from Sep.15<sup>th</sup> 2021 to Dec. 31<sup>st</sup> 2021**

**Gateway Extended Deploying: Jan.1<sup>st</sup> 2022**



## **AOB**

## **AOB**



- Update on Eurofusion Webinars on GPUs and on the IFERC Workshop on GPUs in 2022 (F. Boillod-Cerneux, R. Hatzky, J. David)



# EUROfusion Webinars

F. Boillod-Cerneux (Chair), R.Hatzky (HLST) and J.David



This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

## Webinars



**Currently on YouTube (historic of the numbers: numbers are updated each month)**

Month	# of registrations and # of participants at EUROfusion webinars	Registration	Participants	Questions	Views	Subscriptions
may	Marconi100 - 05.2020	22	72	10	135	8
june	Marconi100 - 06.2020	21	50	4	52	8
july	Kokkos - 07.2020	13	39	5	119	11
september	OpenACC - 09.2020	8	36	11	102	12
october	Vulkan - 28.10.2020	3	26	4	533	19
november	GPUs Roadmap - 30.11.2020	2	18	3	64	24
december	IFERC Workshop GPUs - 17.12.2020	0	35	21	78	28
january	OpenMP Introduction part 1 - 27.01.2021	0	55	13	184	47
february	OpenMP Introduction part 2 - 16.02.2021	0	52	12	106	47
march	OpenMP optimization/SIMD - 16.03.2021	1	52	4	60	53
april	OpenMP introduction to Oaffload GPU - 21.04.2021	0	34	25	70	62
may	IFERC-CSC Workshop on JFRS-1 projects - 18.05.2021	0	25	5	27	62
may	Programming OpenMP - 24.05.2021	0	26	20	33	62
june	IFERC Workshop on GPUs Programming #2 - 14.06.2021	0	32	6	26	69



We have now completed a year of webinars. Regarding the numbers on the previous slides, here are some conclusion:

- **Positive conclusions**
  - I think we covered different topics on GPUs during a year, which is a good idea/format
  - The serie of OpenMP webinar is clearly a success (participants are really happy). The concept of serie is definitely something we should dig in
- **Negative conclusions**
  - The number of participant is clearly decreasing over the year. At the end (when we have about 30 participants), we always reach the same participants (about 15-20 people attend regularly to all the webinars)
  - Low impact on the overall community: we reach only ~50 to 30 people on ~400: we need to make progress here. The impact of video on YouTube is low either.

Feedback from TM54 to OC60



- **France proposition for ameliorations for 2022:**
  - Try to get more webinars as a serie
  - I think one webinar/event per month is too much: people are saturated with online meeting. Maybe one every two month on a regular basis? Except for serie webinar, where 1 per month is clearly required.
  - I think we should move to a mix (if covid situation allows this) of online webinars and F2F hands on sessions.
    - By experience, F2F reduced sessions with talk/hands on is more efficient than webinar online (people are "more involved").
    - I would propose to make less webinars (unless we have a clear demand from users for a webinar, or a specific topic) and try to put in place more F2F formations. I still need to think about this to get something "doable".
      - Thinking out loud, Maybe at most 2 F2F formations per year, and at most 6 webinars per year?
  - be Final point is: let's not be rigid and keep the 1 webinar per month format in 2022: having a "fixed/mandatory" schedule is, I think, not the best option.
  - Let's more flexible: trying to get something more adapted to what users need/would like to have.

## France's open question for OC:

- shall we propose a questionnaire to the users? To check what they would like to have? (format, duration and topics)
- Or do you (OC committee) think the experience is good, and format is fine as is (id est one web/per month except when we have a WK, and no web in summer)?

Feedback from TM54 to OC60



## Upcoming!

Webinar	Organisation	Academic / Industry	First Name	Last Name	Subject
9 Sept	OpenMP & RWTH	Academic	Michael & Christian	Klemm & Terboven	Optimisation on GPUs
28 Oct	OpenMP & RWTH	Academic	Michael & Christian	Klemm & Terboven	OpenMP Q&A

**Important point: we need to have a stable product/solution for upcoming webinars**



Feedback from TM55 to OC61



## Next webinars

- webinars to organize starting from November 2021:
  - Possibly 1 webinar on OpenMP: France will discuss at the end/mid June with Christian & Michael
  - 1 webinar on likwid (speakers provided by Serhyi)
  - 1 webinar on vasp on GPU (speakers to be found)
  - 1 webinar on software engineering, speaker provided by Roman (thanks!) Erik Lindahl (U of Stockholm)
  - **As a conclusion, planning of webinars for 2021 is closed**
    - Reminder: no webinar in August.

Feedback from TM55 to OC61



Feedback from TM55 to OC61

## Operation Committee meeting OC-61 – July 23, 2021



# Close of webinar report

F. Boillod-Cerneux (Chair), R.Hatzky (HLST) and J.David





# IFERC Workshop 2021

F. Boillod-Cerneux (Chair), R.Hatzky (HLST) and J.David



This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

## IFERC Workshop 2021



### 14 June 2021 Workshop on GPUs

- Deliverable
  - RP4 is approved.

**Next steps: workshop #3 before June 2022 (RP6).**

**Possibly a F2F event?**

**If so, where? (France is pending for your propositions 😊)**

**Shall we think about something more interactive than talks presentation?**



## Next meetings

## Next meetings



- Operation Committee (monthly)
  - 23 August 2021 – *Written report from E. Rossi and F. Iannone*
  - 20 September 2021 (9:30-12:30): OC meeting #62 - by video
  - 22 October 2021 (9:30-12:30): OC meeting #63 f2f and/or by video
- Project Committee (quarterly)
  - 27 September 2021 (morning) – f2f or by video (to be decided by the end of August 2021)