


52nd Ticket Meeting

HPC User Support @ CINECA
April, 14th 2021


Content



- Status of the clusters - main events affecting production [March 18 – April 13]
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to Intel support
 - other tickets
- Ticket statistics on queue [reference period: March 15 – April 8]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL [March 2021]

Status of the clusters March 18 – April 13

Main events affecting production



Mar 17: Marconi100 upgrade March 23/24

https://www.hpc.cineca.it/center_news/marconi100-upgrade-march-2324

https://www.hpc.cineca.it/center_news/reminder-marconi100-upgrade-march-2324

https://www.hpc.cineca.it/center_news/update-marconi100-maintenance

https://www.hpc.cineca.it/center_news/marconi100-upgrade-completed

Mar 30: increased the number of nodes available on skl_fua_prod partition

https://www.hpc.cineca.it/center_news/marconi-academic-partition-reduced

Apr 1: announced issues on Marconi filesystems and login nodes

https://www.hpc.cineca.it/center_news/marconi-issues-filesystems-and-login

https://www.hpc.cineca.it/center_news/marconi-back-production-38

Apr 2: HelpDesk closed on Monday, April 5th due to the Easter holidays

https://www.hpc.cineca.it/center_news/helpdesk-closed-monday-april-5th


Apr 8: scheduled maintenance operations on Marconi on April 13th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-april-13

https://www.hpc.cineca.it/center_news/marconi-will-be-stopped-tomorrow-scheduled-maintenance


https://www.hpc.cineca.it/center_news/marconi-back-production-39

Tickets escalated to Intel support




Ticket	Subject	Creation date	Last Updated by Intel	Comments
3932	Pointers to module arrays not working with SIMD nilsm@ipp.mpg.de	18/09/2019 17:24:02	27/01/2021	<p>Intel originally found a bug on the compiler front end and solved it.</p> <p>They provided us an archive with the sources, BUILD script and the outputs of the reproducer that did not show ed the issue.</p> <p>The first bug had "hidden" a second bug on the vectorizer that was triggered by the definition of the macro "USE_ARR_IN_MODULE" in the reproducer. We reported this issue to Intel that confirmed the problem and opened a new bug: Intel bug report CMPLRIL0-33599</p>

Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
9462	Marconi100 issue 3 from the pre-production testing: Low X Bus bandwidth during Nvidia P2P benchmark	serhiy.mochalsky @ipp.mpg.de	2020-05-12 16:48:02	M100	2nd IBM Nvidia	<p>Results from tests show values close to the theoretical value bandwidth during the communication of 2 GPUs inside one socket.</p> <p>During the communication between two GPUs on different sockets the bandwidth decreases to ~39 GB/s for Bi-directional communication having the theoretical value of 64 GB/s.</p> <p>To discuss with IBM and Nvidia to reach a better understanding of how the communications intersocket work whether p2p is enabled or not.</p> <p>Nvidia tested the p2p communication tool coming with the Cuda suite, and they also notice the problem when the communication involves extra-socket GPUs.</p> <p>We are working with nvidia to reproduce the tests on different clusters and to get more info about datapaths.</p>
9677	Marconi100 problem with GPU OpenMPI library	serhiy.mochalsky @ipp.mpg.de	2020-05-18 11:00:01	M100	2nd IBM	<p>The user reported problems with GNU OpenMPI on Marconi100. To avoid conflicts with Spectrum MPI pmix support, we had to rebuild slurm against the same pmix libraries used by Spectrum MPI. This forced us to rebuild OpenMPI as well, but we still have some problems to make it properly work. We asked assistance to our IBM contact people to properly configure OpenMPI to provide the desired advanced GPU features and mellanox optimization support. After the system upgrade the module has been reinstalled.</p>

Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
13850	nvidia hpc-sdk + cuda	thomas.hayward@ipp.mpg.de	2020-11-19 14:54:02	M100	2nd nvidia	<p>The user is been supported for the compilation of his code (orb5) using hpc-sdk/2020--binary, gnu/8.4.0, and cuda/11.1 modules available on M100 cluster. The compilation of this code fails (but it could be correctly compiled w with pgi/19.10 + default cuda module). Nvidia reported that apparently the installation of the hpc-sdk module was somehow bugged. Following nvidia indications we upgraded the hpc-sdk suite to version 20.11, that it is accessible via the same module, hpc-sdk/2020--binary.</p> <p>The compilation of this code using this new release still fails; a combination of two compiling flags triggers a memory leak exhausting the memory node, causing the failure in the build. Unfortunately this is a new problem with respect to the original one, but nvidia recommends to use the last version of the suite because of some bug fixes which might cure the original problem. Working w with the user we are meanw hile looking for a workaround, and with our nvidia support to understand the origin of this behaviour. The user has provided in the meanwhile a copy of the code to nvidia person following this issue.</p>
15867	nvc++ compiler in M100	Leo.Ma@ukaea.uk	2021-02-21 11:20:01	M100	2nd nvidia	<p>The user reports that compiling his code with nvc++ compiler on hpc-sdk module it is not possible to perform offload to GPU. Investigations on this malfunctioning showed that the same code runs perfectly on another cluster that do not have Power processors. The reason of the malfunction is probably a bug in the hpc-sdk compiler version for Power nodes. We have reported this to Nvidia.</p>

Other active tickets on HPC-US-SECOND



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
16019	Marconi SKL inter socket connection	serhiy.mochalsky@ipp.mpg.de	2021-03-01 09:10:02	SKL	2nd	<p>Information about interconnection between two sockets of the same node has been retrieved and also published on the Users Guide.</p> <p>Additional information retrieved to estimate the correct value of the Inter-Socket bandwidth on SKL in terms of GB/s.</p> <p>Performed the <code>osu_mbr_mr</code> benchmark (uni-directional test) that supposes to test the inter socket bandwidth using different number of MPI pairs for simultaneous data transfer.</p> <p>To perform this or another benchmark (for example, the Intel IMB benchmark has similar test to check the aggregate bandwidth between two sockets), so to compare all results and the theoretical value.</p>
16039	paraview on m100	mattwi@fysik.dtu.dk	2021-03-01 13:56:02	M100	Int	Paraview is available as a module from a RCM session. The user asked also for nVIDIA Index plugin that show issues during the installation.
16154	c++ boost error on marconi100	markus.held@chalmers.se	2021-03-04 15:48:02	M100	2nd	The user reports an error related to the c++ boost library when running his code on a single GPU on Marconi100 cluster. Investigations are underway.


Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
16819	Torch can not find MAGMA libraries	koenstrien@gmail.com	2021-03-26 15:54:01	M100	Int	The user reports an error due to magma libraries missing when trying to train a network using pytorch. We have provided the user a new module with a Deep Learning stack, open-ce, that contains pytorch build with magma support.
16960	problem when using marconi VASP to calculate big supercells	qigui@kth.se	2021-03-31 21:02:02	SKL	2nd	The user reports systematic errors when using VASP on Marconi to calculate a big supercell (8x8x8 bcc). We suspect this problem may be due to the number of cores used in the simulations, as this error sometimes can be solved changing ALGO (used IALGO=48 that could be substituted with IALGO=38). Tests are underway.
17204	WG: ask help (job issues)	xueli.luo@kit.edu	2021-04-09 18:19:13	SKL	2nd	The user reports an error shown when running his job on 64 nodes, whilst the same error is not shown asking for 16 nodes. Additional tests are underway to understand the source of the error (to test also for a possible error on the code or input configuration).

Ticket Statistics March 15 – April 8


HPC-US-FIRST & HPC-US-SECOND



Received tickets by date



Business hours from ticket creation to initial address
HPC US FIRST & HPC US SECOND




TOTAL	→ 59 tickets
- HPC-US-FIRST	→ 43
- escalated to HPC-US-SECOND	→ 16

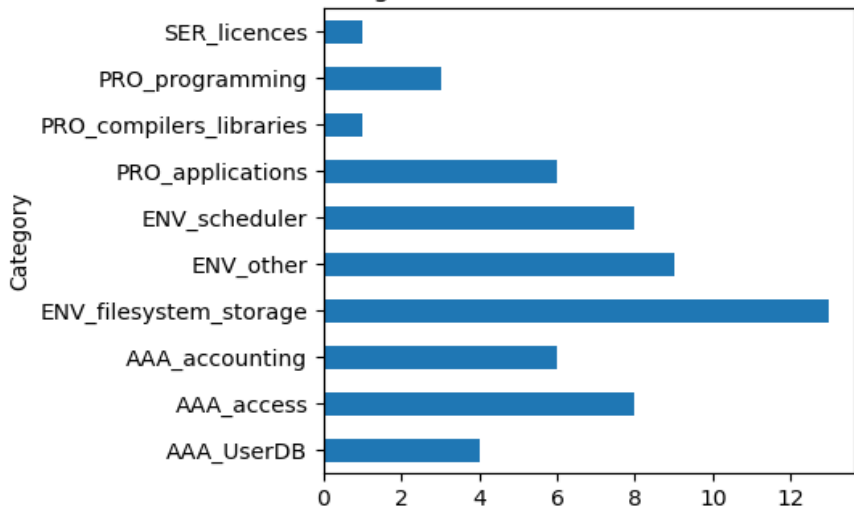
Ticket statistics

March 15 – April 8

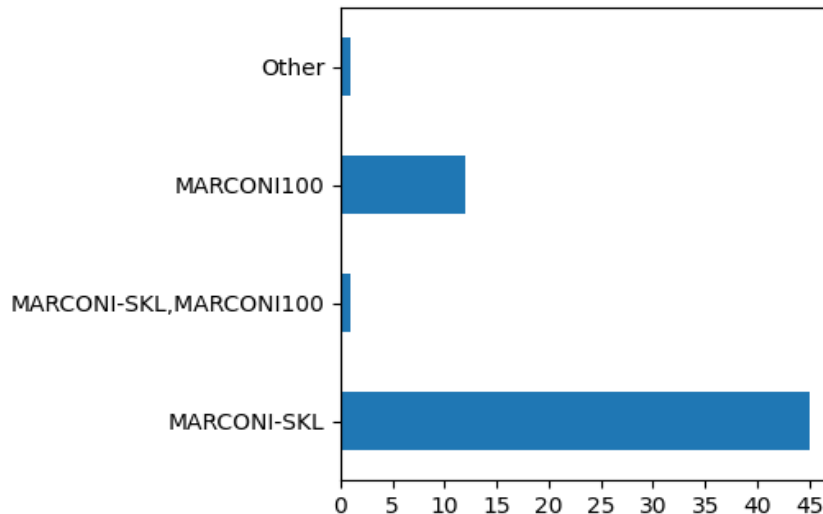
HPC-US-FIRST & HPC-US-SECOND



Ticket categories on HPC-US-FIRST & HPC-US-SECOND




Classification of tickets by platform
HPC-US-FIRST & HPC-US-SECOND




Ticket statistics March 15 – April 8

HPC-US-FIRST & HPC-US-SECOND




time from ticket creation to ticket resolution
HPC US FIRST



41/43 resolved tickets

time from ticket creation to ticket resolution
HPC-US-SECOND



12/16 resolved tickets

Ticket statistics March 15 – April 8






Ticket FIRST	Days	Creation time	Requestor	Subject	Notes
16469	7	16/03/2021 11:26:01	jsantos@ipfn.tecnico.ulisboa.pt	RCM access problems	The user reported issues when using RCM to connect to Marconi cluster that were then fixed, we informed the user the day after. This ticket was closed with some days delay.

Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
16439	16	2021-03-15 15:26:02	marco.veranda@igi.cnr.it	migrazione progetti PIXIE3D FUA34 -> FUA35	The user asked for the migration of the old project area and also for the preservation of a petsc installation in the old project work area. We found together with the user a solution to preserve the old environment for his code.
16818	7	26/03/2021 14:44:01	thomas.hayward@pp.mpg.de	archive/pgi compiler no longer working	The pgi installation available in the archive modules profile has been fixed.
16711	6	24/03/2021 19:54:02	Leo.Ma@ukaea.uk	lammps	The user asked for help to compile the newest lammps version on Marconi100 cluster. We provided to install it as a module including support for kokkos and manybody.


Ticket statistics March 15 – April 8

HPC-US-FIRST



Ticket statistics March 15 – April 8

HPC-US-FIRST




HPC-US-FIRST		By status		By platform			Other
		open	resolved	SKL	SKL&M100	M100	
Information	AAA_accounting	1				1	
	ENV_other		1	1			
	ENV_scheduler		2	2			
Problem	AAA_UserDB		1				1
	AAA_access		3	2		1	
	AAA_accounting		2	2			
	ENV_filesystem_storage		2	2			
	ENV_other		7	7			
	ENV_scheduler		4	3		1	
	PRO_applications		2	2			
PRO_programming		2	1			1	
Service Request	AAA_UserDB		3	2	1		
	AAA_access		3	3			
	AAA_accounting		2	1		1	
	ENV_filesystem_storage	1	4	5			
	ENV_scheduler		1			1	
	PRO_applications		1	1			
SER_licences		1				1	
total	43	2	41	34	1	7	1


Ticket statistics

March 15 – April 8


HPC-US-SECOND



Ticket categories on HPC-US-SECOND




Classification of tickets by platform
HPC-US-SECOND



Ticket statistics March 15 – April 8


HPC-US-SECOND



HPC-US-SECOND		By status		By platform	
		open	resolved	SKL	M100
Information	AAA_accounting		1	1	
	AAA_access		1		1
Problem	ENV_filesystem_storage	1	2	2	1
	ENV_other		1	1	
	ENV_scheduler		1	1	
	PRO_applications	2		1	1
	PRO_compilers_libraries		1		1
	PRO_programming	1		1	
Service Request	AAA_access		1	1	
	ENV_filesystem_storage		3	3	
	PRO_applications		1		1
total	16	4	12	11	5

Module usage on Marconi-SKL


Batch jobs March 2021



category	modulename	partition	qos	module_load_count	elaped_timelimit_sum
application	openfoam-ext	skl_fua_prod	normal	2	2,00
compiler	gnu	skl_fua_dbg	normal	37	5,25
		skl_fua_prod	normal	4	1,00
	openmpi	skl_fua_prod	normal	3	0,00
	python	skl_fua_dbg	normal	2	0,67
		skl_fua_prod	normal	16	8,40
skl_qos_fuabprod			5	2,45	
library	blas	skl_fua_dbg	normal	4	0,30
tool	anaconda	skl_fua_dbg	normal	1	0,43
	cmake	bdw_all_serial	normal	1	0,80
		skl_fua_dbg	normal	1	0,24
		skl_fua_prod	normal	1	1,00
	cubegui	skl_fua_prod	normal	1	1,00
	scalasca	skl_fua_prod	normal	1	1,00
	scorep	skl_fua_prod	normal	1	1,00
valgrind	skl_fua_dbg	normal	37	5,01	

Module usage on Marconi-SKL

Batch jobs March 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	blas	skl_fua_dbg	normal	4	0,30
		skl_fua_prod	normal	53	15,70
	cubelib	skl_fua_prod	normal	1	1,00
		skl_fua_dbg	normal	1	0,24
	fftw		normal	17	11,06
		skl_fua_prod	skl_qos_fuabprod	6	3,03
	hdf5	bdw_all_serial	normal	8	3,04
		skl_fua_dbg	normal	1	0,24
		skl_fua_prod	normal	16	10,89
	lapack	skl_fua_prod	skl_qos_fuabprod	6	3,03
		skl_fua_dbg	normal	4	0,30
	lapack	skl_fua_prod	normal	53	15,70
		bdw_all_serial	normal	8	3,04
	netcdf	skl_fua_dbg	normal	1	0,24
			normal	13	8,40
		skl_fua_prod	skl_qos_fuabprod	5	2,45
	netcdf	bdw_all_serial	normal	8	3,04
		skl_fua_prod	normal	11	6,82
	netcdf		skl_qos_fuabprod	5	2,45
		skl_fua_prod	normal	12	7,40
	petsc		skl_qos_fuabprod	6	3,03
		skl_fua_prod	normal	1	1,00
qt	skl_fua_prod	normal	1	1,00	
	skl_fua_dbg	normal	1	0,24	
scalapack	skl_fua_prod	normal	2	1,21	
	skl_fua_prod	normal	5	4,21	
scipy	skl_fua_prod	normal	5	4,21	
slepc	skl_fua_prod	normal	5	4,21	
wsmp	skl_fua_prod	normal	8	1,68	