

53rd Ticket Meeting

HPC User Support @ CINECA
May, 19th 2021


Content



- Status of the clusters - main events affecting production [April 14 – May 18]
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to Intel support
 - escalated to NVIDIA support
 - other active tickets
- Ticket statistics on queue [reference period: April 9 – May 16]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL

Status of the clusters April 14 – May 18

Main events affecting production



April 21: scheduled maintenance operations of Marconi100 on April 27th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-april-27th

https://www.hpc.cineca.it/center_news/m100-will-be-stopped-tomorrow-scheduled-maintenance-0

https://www.hpc.cineca.it/center_news/marconi100-back-production-and-report-updates

April 29: slurm issues on Marconi cluster

https://www.hpc.cineca.it/center_news/marconi-slurm-issues-0

https://www.hpc.cineca.it/center_news/marconi-back-production-40


May 11: scheduled maintenance operations of Marconi on May 18th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-may-18th

https://www.hpc.cineca.it/center_news/marconi-will-be-stopped-tomorrow-scheduled-maintenance-0


https://www.hpc.cineca.it/center_news/marconi-back-production-and-report-updates

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		<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
16019	Marconi SKL inter socket connection	serhiy.mochalsky@ipp.mpg.de	2021-03-01 09:10:02	SKL	2nd Intel	<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>

Tickets escalated to **NVIDIA** support




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.mochalskiy@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> <p>We reached IBM to discuss the issue with them. They passed the problem to a specialist team in US after testing the benchmark and confirming the behaviour: "The essential point is understanding the data flux GPUtoGPU when they are hosted by different CPUs. We suspect that cache and/or memory have a play in this, unless it is possible to activate a direct link towards x-Bus, that is not implying any direct intervention from the CPU. It is possible that p2p=enabled is not enough for that purpose."</p>

Tickets escalated to **NVIDIA** support




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 with pgj/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 with the user we are meanwhile 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>

Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
16479	Marconi-Fusion automatic runs	mhoelzl@ipp.mpg.de	2021-03-16 16:32:02	SKL	Internal	The workflow proposed by the user would make use of a gitlab runner service on Marconi login nodes that is not available at present, the service is discouraged by our system administrators for possible security issues. Some other solutions may be devised, to be discussed with the production team and the sysadmins.
17735	git lfs on Marconi/M100	K.L.vandePlassche@diff er.nl	2021-05-04 12:36:02	M100 SKL	1st	The user asked for the possibility of having git lfs they use for version controlling large files. We installed a new module on M100 cluster, git/2.27.0, that contains this tool. The user asked also for the same module on Marconi cluster, this is work in progress.
17945	possible gcc upgrade	par@ipfn.ist.utl.pt	2021-05-13 10:26:02	M100	1st	The user asked for the possibility of installing as module a gcc-10 (or higher) version in M100 to take advantage of the newer c++ features made available by the c++20 standard. We indicated this user all the available gnu versions installed on the cluster but they still do not support the required features. The user is using at present a custom version of the compiler.

Other resolved tickets on HPC-US-SECOND



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
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.</p> <p>After the system upgrade the module has been reinstalled ensuring its integration with SLURM. The module name remains the same, it is available in the advanced modules profile. Also, the module help was updated by adding some basic information on the gpu binding.</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. A new hpc-sdk release (version 21.3) has been installed on M100 cluster; this version allows a correct offload on Power9 processors.</p>
16039	paraview on m100	mattwi@fysik.dtu.dk	2021-03-01 13:56:02	M100	Int	<p>Paraview is available as a module from a RCM session. The user asked also for nVIDIA Index plugin that show issues during the installation. This plugin is not available for M100 architecture.</p>


Other resolved 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	<p>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.</p> <p>The user was supported also into the setup of the sponcv package, that was required on his work environment.</p>
16960	problem when using marconi VASP to calculate big supercells	qigui@kth.se	2021-03-31 21:02:02	SKL	2nd	<p>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).</p>
17204	WG: ask help (job issues)	xueli.luo@kit.edu	2021-04-09 18:19:13	SKL	2nd	<p>The user reports an error shown when running his job on 64 nodes, whilst the same error is not shown asking for 16 nodes. The user managed to run his code by using a different size for his jobs.</p>

Ticket Statistics April 9 – May 16


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	→ 62 tickets
- HPC-US-FIRST	→ 41
- escalated to HPC-US-SECOND	→ 21

Ticket statistics April 9 – May 16

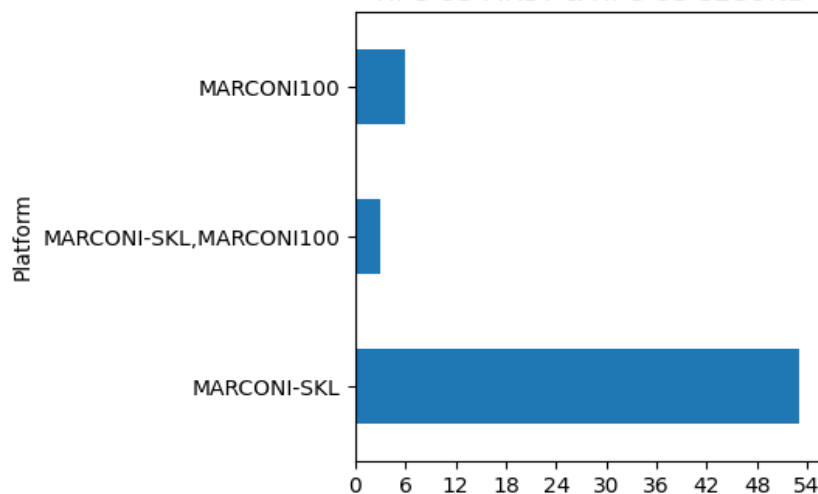
HPC-US-FIRST & HPC-US-SECOND



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




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




Ticket statistics April 9 – May 16

HPC-US-FIRST & HPC-US-SECOND




time from ticket creation to ticket resolution
HPC US FIRST




41 / 42 resolved tickets

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



12 / 21 resolved tickets

Ticket statistics April 9 – May 16

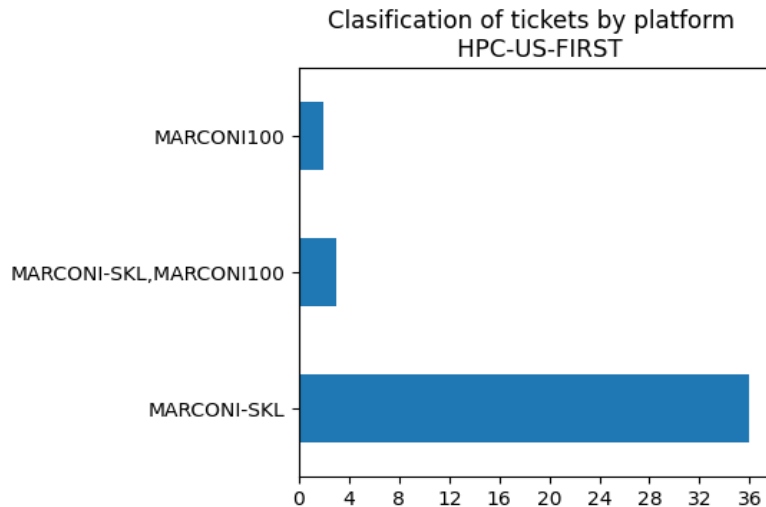




Ticket FIRST	Days	Creation time	Requestor	Subject	Notes
17690	6	2021-04-30 15:28:01	T.W.J.Driessen@diffen.nl	SOLPS-ITER	The user was linked as a collaborator of the FUSIO_ALL account on our UserDB site. This was closed with delay.
17288	3	2021-04-13 17:32:02	Klaus.Hallatschek@ipp.mpg.de	Question on available compiler	The user asked for some information about available modules compiled with latest Intel compiler release on Marconi cluster. This was closed with delay.

Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
17209	21	2021-04-10 14:36:02	louis.ad@hotmail.ch	Modules and CMake search paths	The user reported some suggestions about how to manage path to libraries when using cmake for code compilation on Marconi.
17214	10	2021-04-11 14:52:02	ralf.keiber@ipp.mpg.de	job problem on Marconi	The user reported an eutrope job that did not produce any output and end by time limit. We checked the nodes involved in the job but we could not find any issue. We suggested this user to recompile the code, after that the wrong behaviour has not been shown anymore.
17639	10	2021-04-28 14:47:44	lorenz.romaner@unileoben.ac.at	Re: Email Address Change	We supported this user into the update of his UserDB personal page information and to reset his HPC password.
17522	7	2021-04-22 09:52:01	daniele.bonfiglio@igi.cnr.it	Richiesta gestione cartelle FUA34_DISJOR e FUA34_AUGJOR su Marconi	The user asked for support on how to manage and preserve the old work areas of projects FUA34_DISJOR and FUA34_AUGJOR, as those projects have now converged into FUA35_MHD project. We supported these users into the copy of the former work areas into the new one.


Ticket statistics April 9 – May 16

HPC-US-FIRST



Ticket statistics April 9 – May 16

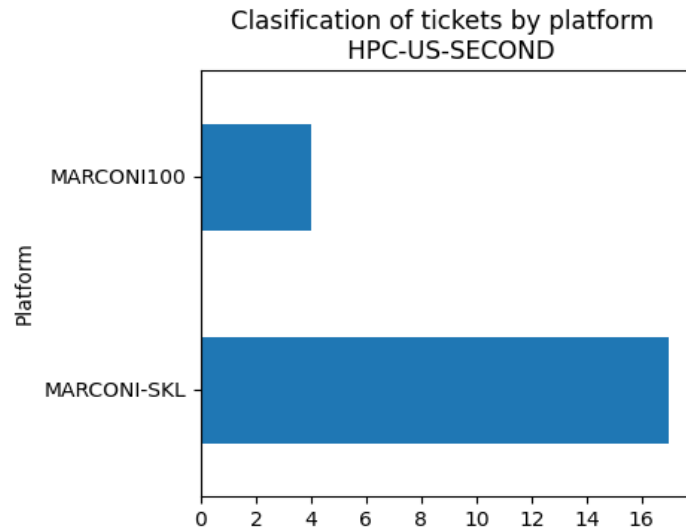


HPC-US-FIRST



HPC-US-FIRST		By status		By platform		
		open	resolved	SKL	SKL&M100	M100
Information	AAA_access		1	1		
	AAA_accounting		1	1		
	ENV_scheduler		1	1		
	PRO_compilers_libraries		2	2		
	SER_interfaces		1			1
Problem	AAA_UserDB		1	1		
	AAA_access		2	2		
	AAA_accounting		2	2		
	ENV_filesystem_storage		1	1		
	ENV_other	1		1		
	ENV_scheduler		4	4		
	SER_interfaces		1			1
Service Request	AAA_UserDB		1	1		
	AAA_access		12	11	1	
	AAA_accounting		5	3	2	
	ENV_filesystem_storage		3	3		
	ENV_scheduler		1	1		
	PRO_compilers_libraries		1	1		
total	41	1	40	36	3	2


Ticket statistics April 9 – May 16

HPC-US-SECOND



Ticket statistics April 9 – May 16


HPC-US-SECOND



HPC-US-SECOND		By status		By platform	
		open	resolved	SKL	M100
Information	PRO_applications	1		1	
	PRO_compilers_libraries		1	1	
Problem	ENV_filesystem_storage	1	1	2	
	ENV_other	2	2	3	1
	ENV_scheduler		1	1	
	PRO_applications	1	3	4	
	PRO_compilers_libraries		1		1
	PRO_other	1		1	
Service Request	AAA_access		1	1	
	ENV_filesystem_storage	1	2	3	
	PRO_applications	1			1
	PRO_compilers_libraries	1			1
total	21	9	12	17	4

Module usage on Marconi-SKL


Batch jobs March 2021



category	modulename	partition	qos	module_load_count	elapsed_tmlimit_sum
compiler	python	bdw_all_serial	normal	2	1,053263889
		skl_fua_prod	normal	5	2,127909091
			skl_qos_fuabprod	6	3,131731602
tool	anaconda	bdw_all_serial	normal	2	1,053263889
	likwid	skl_fua_prod	normal	3	2,099578544

Module usage on Marconi-SKL

Batch jobs March 2021



category	modulename	partition	qos	module_load_count	elapsed_tlimit_sum
library	blas	skl_fua_prod	normal	105	32,57898644
	fftw	skl_fua_prod	normal	8	3,712131313
			skl_qos_fuabprod	6	3,131731602
	hdf5	bdw_all_serial	normal	1	0,27
		skl_fua_prod	normal	47	11,71324242
			skl_qos_fuabprod	6	3,131731602
	lapack	skl_fua_prod	normal	105	32,57898644
	netcdf	bdw_all_serial	normal	1	0,27
		skl_fua_prod	normal	6	2,333242424
			skl_qos_fuabprod	6	3,131731602
	netcdfff	bdw_all_serial	normal	1	0,27
		skl_fua_prod	normal	6	2,333242424
			skl_qos_fuabprod	6	3,131731602
	petsc	skl_fua_prod	normal	7	2,979909091
skl_qos_fuabprod			6	3,131731602	
popt	skl_fua_prod	normal	40	8,733333333	