


58th Ticket Meeting

HPC User Support @ CINECA
November, 22nd 2021

Content



- Status of the clusters - main events affecting production [October 15 – Nov 19]
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to SchedMD support
 - escalated to Intel support
 - escalated to NVIDIA support
 - EUTERPE jobs
 - other active tickets
- Ticket statistics on queue [October 13 – November 17]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL

Status of the clusters [Oct 15 – Nov 19]

Main events affecting production



Oct 19th: scheduled maintenance operations on Marconi and Marconi100 clusters on October 26.

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-and-marconi100-next-tuesday-october-26th

https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi-and-marconi100-tomorrow-october-26th

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

https://www.hpc.cineca.it/center_news/update-marconi-back-production-m100-still-under-maintenance

https://www.hpc.cineca.it/center_news/marconi100-back-production-10

Oct 21st: issues when trying to login to some CINECA HPC services through OpenID.

https://www.hpc.cineca.it/center_news/issue-about-login-some-cineca-hpc-services

https://www.hpc.cineca.it/center_news/issues-openid-login-update


Oct 27th: HPC Gitlab instance (gitlab.hpc.cineca.it) not available on November 3 due to a planned maintenance.

https://www.hpc.cineca.it/center_news/hpc-gitlab-unavailable-november-3

https://www.hpc.cineca.it/center_news/reminder-hpc-gitlab-unavailable-afternoon

Status of the clusters [Oct 15 – Nov 19]

Main events affecting production



Oct 29th: Help desk service not available on November 1 due to All Saints' Day festivity.

https://www.hpc.cineca.it/center_news/help-desk-service-closed-november-1st-0

Nov 5th: consumed budgets recalculated from October 1 due to issues in the accounting procedures on Marconi100 cluster.

https://www.hpc.cineca.it/center_news/g100-and-m100-accounting-issue-solved

Nov 15th: activation of 2021 CINECA Help Desk service satisfaction survey.

https://www.hpc.cineca.it/center_news/2021-cineca-hpc-help-desk-service-satisfaction-survey

Nov 17th: scheduled maintenance operations on Marconi100 on November 23


https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-next-tuesday-november-23rd

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 serhiy.mochalskyi@ipp.mpg.de	18/09/2019 17:24:02	31/07/2021 (Intel issue opened by CINECA)	<p>Intel originally found a bug on the compiler front end and solved it. They provided us an archive with the sources, BUILD script and the outputs of the reproducer that did not showed the issue. 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:</p> <p>Intel bug report CMPLRIL0-33599</p> <p>Intel support has provided a resolution for the problem reported in the second bug:</p> <p>“” We cannot use simd for a loop that has F90 pointer assignment inside. For every iteration of the loop, it is updating the same dope vector for f4a. That means there is a loop carried dependency preventing vectorization. If it is vectorized with vecLen =2, for iteration 1 and 2, it is storing into the same location of the dope vector and the address code of f4a is picked up incorrectly. In order to generate the right code, it can run in a non-vector mode. However, attached is fixed.f90 which shows the right way to do it. We need to declare a structure of arrays to store the F90 pointers. In that way, every iteration is storing into different dope vectors. And, of course, the result from fixed.f90 is different because I used different computations in the loop.””</p> <p>We reported to Intel support users' feedback: “” Since I specifically made the f4a pointer private in my omp simd pragmas, it is very unexpected for them to still produce a loop carried dependency. Marking a variable as private is specifically defined as getting around it being a loop carried dependency. The struct of pointers, which fixes the problem, is exactly the kind of solution, the compiler should generate when encountering a pointer variable in a private clause. I hope this is the understanding of the Intel engineers as well, and they'll provide a fix in an upcoming release.””</p> <p>The bug fix will be merged into the compiler, it has been approved its release on Intel oneAPI 2022.0 release</p>

Tickets escalated to NVIDIA support




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
18851	Problem with the nvidia compiler and -lnvc on m100	nilsm@ipp.mpg.de serhiy.mochalsky@ipp.mpg.de	2021-06-23 11:00:02	M100	2nd NVIDIA	The user report a strange behaviour referring to CUDA devices detection when adding (or not) the «-lnvc» flag for the compilation of a test code available on the latest hpc-sdk compiler suite on Marconi100 cluster. We performed some investigations that show that with and without the flag -lnvc the same libraries are linked, the only difference is in the order of the linkage. We are in contact the nvidia support for discussion.
19978	OpenACC compiler problem with function intermediates	nilsm@ipp.mpg.de serhiy.mochalsky@ipp.mpg.de	2021-08-24 17:14:02	M100	2nd NVIDIA	<p>The user provided a test code that is compiled using hpc-sdk/2021--binary module: "The compiler generates an intermediate, called get_arr1 , to store the result of the multiplication in line 26. This intermediate is copied to the device, as seen in the compiler output. Unfortunately, this intermediate needs to be private, as each thread needs its own copy. It is impossible to achieve that, as its name is not known, and it doesn't exist during compile time. Array b has the correct result as the result of the function call is saved in the variable c , so no intermediate is necessary. This problem emerges for all function calls to a device function that returns an array."</p> <p>The issue reported has been escalated to nvidia support for discussion.</p>

Tickets escalated to NVIDIA support



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
19982	OpenACC compiler problem with writing back from the GPU	nilsm@ipp.mpg.de serhiy.mochalsky @ipp.mpg.de	2021-08-24 17:44:02	M100	2nd NVIDIA	The user provided a test code that is compiled using hpc-sdk/2021--binary module; it do not generate expected results related to a copyout data clause. This problem seems related to the syntax used in the copyout data clause. To complete the analysis the issue has been reported to the nvidia support to ask for further clarifications.
20303	OpenACC compiler problem, unexpected implicit wait	nilsm@ipp.mpg.de serhiy.mochalsky @ipp.mpg.de	2021-09-09 15:58:01	m100	2nd	This issue is still under investigation. We will also report it to nvidia if necessary as soon as preliminary checks will be completed.


Other active tickets on HPC-US-SECOND EUTERPE jobs



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
19157	Error on marconi	axel.koenies@ipp.mpg.de	2021-07-07 11:32:02	SKL	2nd	<p>The users reported a problem with some EUTERPE jobs labelled "RUNNING" by slurm and that did never start, no output was produced.</p> <p>We could reproduce the same error reported by the users, using mpiexec for the job execution as in the users' job scripts.</p> <p>Investigations done once the job gets stuck show that mpiexec, that spawns processes on the nodes via ssh instead of slurm (as srun does), could not connect to one of the requested nodes because the slurmd on the node was not yet aware of being part of the job and therefore refuses access (ssh failed).</p>
21746	Jobs fail to start on MARCONI	christoph.slaby@ipp.mpg.de	2021-11-03 17:00:01			

Other active tickets on HPC-US-SECOND

EUTERPE jobs



@r171c15s01 ~]\$ ps awx-forest

```

[...]  

7824 ?   Sl   0:00 slurmstepd: [9835953.extern]  

7829 ?   S    0:00 \_ sleep 100000000  

7831 ?   Sl   0:00 slurmstepd: [9835953.batch]  

7839 ?   S    0:00 \_ /bin/tcsh /var/spool/slurmd/job9835953/slurm_script  

7899 ?   S    0:00 \_ mpiexec -prepend-rank-np 192 -genv F_UFMTENDIAN big ./linux_marconi_euterpe_c -input_file input  

7900 ?   S    0:00 \_ /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy --control-port r171c15s01.marconi.cineca.it:31816  

--pmi-connect alltoall --pmi-aggregate -s0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1 --retries 10 --control-code 1891130188 --usize -2 --proxy-id 0  

7907 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

7908 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

[...]  

7963 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

7964 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

7965 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

7966 ?   S    0:00 | \_ ./linux_marconi_euterpe_c -input_file input  

7901 ?   Z    0:00 \_ [ssh] <defunct>  

7902 ?   S    0:00 \_ /usr/bin/ssh -x -q r171c15s03 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy  

--control-port r171c15s01.marconi.cineca.it:31816 --pmi-connect alltoall --pmi-aggregate -s0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1  

--retries 10 --control-code 1891130188 --usize -2 --proxy-id 2  


7903 ?   S    0:00 \_ /usr/bin/ssh -x -q r171c15s04 /cineca/prod/opt/compilers/intel/pe-xe-2018/binary/impi/2018.4.274/bin64/pmi_proxy  

--control-port r171c15s01.marconi.cineca.it:31816 --pmi-connect alltoall --pmi-aggregate -s0 --rmk slurm --launcher ssh --demux poll --pgid 0 --enable-stdin 1  

--retries 10 --control-code 1891130188 --usize -2 --proxy-id 3  


14417 ?   Ss   0:00 /usr/libexec/sss/sssd_pam --debug-to-files --socket-activated
  
```

Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
18205	paraview on M100	mattwi@fysik.dtu.dk	2021-05-28 11:46:01	M100	Int.	<p>The user reported issues when running paraview version available as module through a RCM session. To fix this is required a recompilation of paraview (we will include, if possible, the installation of the nvidia plugin index required also by this user).</p> <p>The installation of paraview is underway but it is showing issues. Also, we are in contact with nvidia to clarify the terms of the license required to install/run nvidia plugin index on M100 cluster.</p>
20015	Device to device MPI communication with OpenMP4.5	asahi.yuichi@jaea.go.jp	2021-08-26 08:08:03	M100	2nd	<p>The user report errors when trying to compile his code. This first error refers to a XLC++ version of the code that could compile and run correctly before the major update of the cluster. We helped the user into the compilation of the code by including a missing linking flag. The user confirmed that he can go on with the work.</p> <p>The second error reported refer to the usage of hpc-sdk compiler suite for the compilation of the same code (FATAL ERROR: "data in use_device clause was not found on device 1: host:0x2140dbf02), this is still under investigation.</p>
20019	Segfault before code launch with srun	Patrick.TAMAIN@cea.fr	2021-08-26 09:28:03	SKL	2nd	<p>The user is experiencing issues when trying to run his code. For the code compilation are used both modules available on the system and custom libraries available at the user personal areas. The error seems related to the petsc libraries (module) used for the compilation of the code. We provided a new installation of petsc library to check again for compatibility issues.</p>

Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
20676	QdstrmlImporter	huw.leggate@dcu.ie	2021-09-30 13:04:01	M100	2nd	The user reports that using nsys the QdstrmlImporter step fails for the qdstrm files. He tried also the execution of QdstrmlImporter from the command line obtaining the same result. He is attempting to use openmp offloading with the default gcc 8.4.0, that appears to have been built without offloading enabled, and he suspects that may be causing the crash.
20832	GNU/8.4.0 Offloading support	huw.leggate@dcu.ie	30/09/2021 13:04:01	M100	2nd	The user reports that the default gnu version 8.4.0 on Marconi100 appears to have been compiled without offloading support and asks for the possibility to provide an 8.4 build with offloading enabled as this is the required version to build other libraries.
21294	Switching MPI compilers	huw.leggate@dcu.ie	2021-10-18 14:16:01	M100	2nd	The user asks for a way to use different underlying compilers in the same environment when using mpif90, specifically, to compile with the spectrum mpi mpif90 command using gfortran and also using nvfortran, without unloading sdk-hpc module. We have suggested the usage of the environment variable OMPI_FC so to override the default fortran compiler.

Other active tickets on HPC-US-SECOND



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
21940	gcc/10 on marconi?	nicola.varini@epfl.ch	2021-11-10 18:04:01	SKL	2nd	The user has requested the installation on Marconi cluster of netcdf-fortran and hdf5 libraries, that need to be compiled with gcc/10 and mpi.
21950	-llapacke compilation error	markus.held@chalmers.se	2021-11-11 10:08:02	SKL	2nd	The user is experiencing a library linking error when compiling one code, in particular with lapack library (lapack/3.8.0--intel--pe-xe-2018--binary). As an alternative to solve the issues he has asked for the installation of lapack for gnu/8.3.0. Work in progress.
22039	Diskspace on WORK area of Marconi	jdominguezpalacios@us.es	2021-11-15 13:58:02	SKL	1st	The user, collaborator of the FUA35_MEGAEDGE project, has requested a work area quota increase to a total of 120TB. The quota has been increased as requested, the PI of the project was informed of this operation.


Other resolved tickets on HPC-US-SECOND



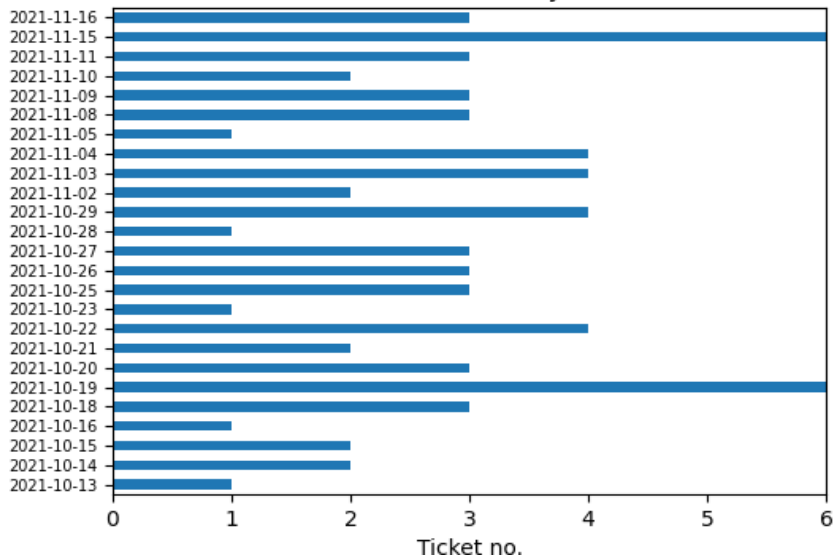
Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
18908	Paraview not running with SLURM	louis.stenger@epfl.ch	2021-06-25 07:32:02	SKL	2nd INT	The user reported errors when trying to use paraview in parallel mode. We provided the user detailed information on how to use this tool, that allows to run several pvservers on a Marconi compute node (or more), and on how to reach those from the local workstation by opening a tunnel to the assigned compute node.

Ticket Statistics Oct 13 – Nov 17


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 → 70 tickets
 - HPC-US-FIRST → 47
 - escalated to HPC-US-SECOND → 23

Ticket statistics Oct 13 – Nov 17


HPC-US-FIRST & HPC-US-SECOND



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




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




Ticket statistics Oct 13 – Nov 17

HPC-US-FIRST & HPC-US-SECOND




time from ticket creation to ticket resolution
HPC US FIRST




46 / 47 resolved tickets

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



16 / 23 resolved tickets

Ticket statistics Oct 13 – Nov 17



Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
21236	13	2021-10-14 19:44:02	adisi@ipp.mpg.de	Marconi proposal 2022	The user asked for extra budget on M100 to run some benchmarks for the finalization of his proposal for the next EUROfusion cycle. We concated the OC that agreed to provide compute hours. As this user belongs to IPP institution, we suggested to be included as a collaborator of the Research Unit account, FUSIO_ru5IPP_0, that has still budget to be consumed.
21443	10	2021-10-22 10:03:01	Matthias.Borchardt@ipp.mpg.de	Problems compiling in interactive session on M100	The user has experienced code's compilation issues when trying to link compiled objects from login node and from interactive node (same modules loaded on both these nodes). Reported errors indicate a mismatch with CUDA runtime variables. Apparently, while on login nodes default cuda version of hpc-sdk compiler is cuda/11.0, on compute nodes taken with an interactive session with no GPU request the default version is cuda/10.2. This can be overcome by using the -Mcuda=cuda/11.0 flag during compilation.
21749	10	2021-11-03 17:28:02	patrikollus@aalto.fi	Issue with pip in virtual environment	The user reported issues when trying to install some python packages on a virtualenv on Marconi cluster. We provided the user some indications on how to proceed to install these correctly.


Ticket statistics Oct 13 – Nov 17




Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
21458	8	2021-10-22 14:54:02	huw.leggate@dau .ie	Compiler error	The user report an error using nvidia hpc-sdk/2021 fortran compiler error when compiling a fortran module.
21503	7	2021-10-25 12:52:01	jordy.trilaksono@i pp.mpg.de	Issues with OpenMP	The user report an error with his job on Marconi cluster. This was due to a conflict with modules loaded in the .bashrc env file and modules loaded in the job script.
21655	6	2021-10-29 14:54:02	tzok@man.pozna n.pl	Multi-node MPI does not work	The user report an error when using interactive srun to run his MPI code. This is a known behavior with srun that allocates all the resources that are requested and, in case of a multinode request, it does so in a way that impedes other resource requesting commands such as mpirun or a second srun to start properly. This can be overcome by using the --overlap flag with srun.

Ticket statistics Oct 13 – Nov 17


HPC-US-FIRST




Ticket categories on HPC-US-FIRST



Clasification of tickets by platform
HPC-US-FIRST






Ticket statistics Oct 13 – Nov 17 *HPC-US-FIRST*



HPC-US-FIRST		By status		By platform		
		open	resolved	M100	SKL	SKL&M100
Information	AAA_UserDB		2		1	1
	ENV_scheduler		3		3	
Problem	AAA_access		7	1	5	1
	AAA_accounting		1		1	
	ENV_modules		1		1	
	ENV_scheduler		5		5	
	PRO_applications		1		1	
	PRO_compilers_libraries		2		2	
	SER_licences		1		1	
	SER_other		1	1		
Service Request	AAA_UserDB		4	1	2	1
	AAA_access	1	12	1	10	2
	AAA_accounting		1		1	
	ENV_filesystem_storage		3		3	
	ENV_scheduler		1		1	
	PRO_compilers_libraries		1		1	
total	47	1	46	4	38	5


Ticket statistics Oct 13 – Nov 17

HPC-US-SECOND



Ticket statistics Oct 13 – Nov 17

HPC-US-SECOND



HPC-US-SECOND		By status		By platform		
		open	resolved	M100	SKL	M100&SKL
Information	AAA_access	1				1
	ENV_scheduler		1	1		
	PRO_compilers_libraries	1		1		
Problem	AAA_access		2		2	
	ENV_scheduler	2	1		3	
	PRO_compilers_libraries	1	6	3	4	
	SER_interfaces		2		1	1
Service Request	AAA_UserDB		1			1
	AAA_access		1			1
	AAA_accounting		1	1		
	ENV_filesystem_storage		1		1	
	PRO_applications	1			1	
	PRO_compilers_libraries	1			1	
total	23	7	16	6	13	4

Module usage on Marconi-SKL



Note:

Module usage data on Marconi cluster has been recalculated due to a issue found on a procedure.

New generated data for 5th cycle projects (since March until October 2021) has been provided in a separate document.