


60th Ticket Meeting

HPC User Support @ CINECA
January, 17th 2022


Content



- Status of the clusters - main events affecting production [Dec 15 – Jan 14]
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to Intel support
 - escalated to NVIDIA support
 - other tickets
- Ticket statistics on queue [Dec 13 – Jan 13]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL

Status of the clusters [Dec 15 – Jan 14]

Main events affecting production



Dec 23: opening time of the Help Desk service with the operator on duty from 9-13 a.m on December 24th and 31st
https://www.hpc.cineca.it/center_news/help-desk-service-christmas-time

Dec 28: restart of licserverhpc on December 29th
https://www.hpc.cineca.it/center_news/restart-licserverhpc

Jan 5: Help Desk Service not available on January 6th
https://www.hpc.cineca.it/center_news/help-desk-service-january-6


Jan 5: scheduled maintenance operations on Marconi on January 11th
https://www.hpc.cineca.it/center_news/marconi-scheduled-maintenance-january-11th
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi-tomorrow-january-11th
https://www.hpc.cineca.it/center_news/marconi-maintenance-completed-1

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.mochalskyy@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.mochalskyy@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.mochalskyy@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	niism@ipp.mpg.de e 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	niism@ipp.mpg.de e 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




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
18205	paraview on M100	mattwi@fysik.dtu.dk	2021-05-28 11:46:01	M100	Int.	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). 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.
20015	Device to device MPI communication with OpenMP4.5	asahi.yuichi@jaea.go.jp	2021-08-26 08:08:03	M100	2nd	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. 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. The user is able to work in any case that to the solution of the first problem.
20676	Qdstrmlmporter	huw.leggate@dcu.ie	2021-09-30 13:04:01	M100	2nd	The user reports that using nvprof the Qdstrmlmporter step fails for the qdstrm files.He tried also the execution of Qdstrmlmporter 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. While gcc 8.4.0 has been updated and now supports offloading, it is unclear if this solved the problem with the user. We will also suggest to use nsys that is the most recommended nVIDIA profiler.

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.
22068	ottimizzazione job sottomesso	giovanni.digianatale@epfl.ch	2021-11-16 12:02:01	M100	2nd	The user is observing out of memory errors when running an ORB5 simulation, he is asking for 32 nodes for the run. The case he is running foresees 1600M particles, 32 nodes is on the edge. This ticket is related to #23189
23189	problema memoria M100	giovanni.digianatale@epfl.ch	2022-01-07 10:50:06	SKL	2nd	Continuation of ticket #22068. The user states that he can work with restarts, but has memory errors after a certain number of them.


Other resolved tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
22039	Diskspace on WORK area of Marconi	jdominguezpalacios@us.es	2021-11-15 13:58:02	SKL	2nd	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. Before proceeding with the quota increase asked for a motivation that has been provided to the OC for approval.
22416	Compilation of PARALLAX	Andreas.Stegmeier@ipp.mpg.de	2021-11-29 17:00:01	SKL	2nd	The user requires to compile and run his Fortran code, PARALLAX, gnu/8 and an openmpi library that is compatible with gnu/8 on Marconi cluster. We have installed openmpi-4.1.1 compiled with gnu-8.3.0 (available as a module in the advanced profile). It is also required for this code the netcdf 4.x library (under installation together with dependencies needed by the library).

Ticket Statistics Dec 13 – Jan 13


HPC-US-FIRST & HPC-US-SECOND



Received tickets by date




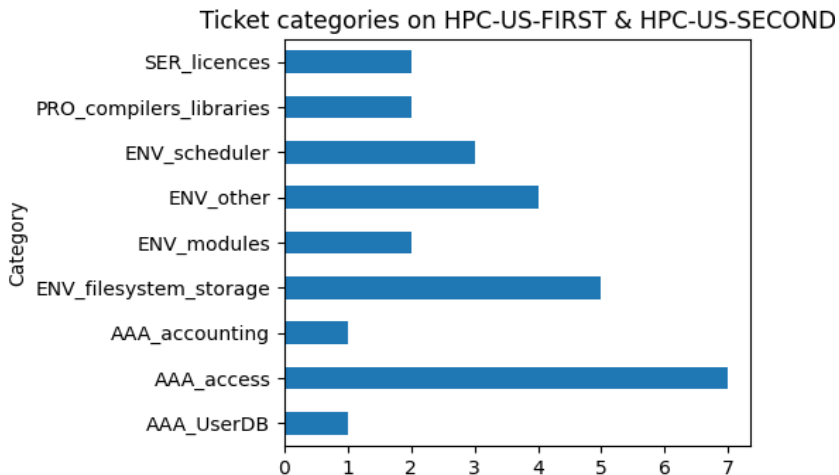

Business hours from ticket creation to initial address
HPC US FIRST & HPC US SECOND



TOTAL	→ 27 tickets
- HPC-US-FIRST	→ 21
- escalated to HPC-US-SECOND	→ 6


Ticket statistics

Dec 13 – Jan 13
HPC-US-FIRST & HPC-US-SECOND




Ticket statistics Dec 13 – Jan 13

HPC-US-FIRST & HPC-US-SECOND




time from ticket creation to ticket resolution
HPC US FIRST



20 resolved tickets

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



5 resolved tickets

Ticket statistics Dec 13 – Jan 13







Ticket FIRST	Days	Creation time	Requestor	Subject	Notes
23171	3	05/01/2022 18:28:01	Feng.LIU@cea.fr	Cineca technical portal login	User requests for informations about resetting his/her UserDB password. NOTE: this was actually solved in two business days, because January 6th was an holiday

Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
23015	15	2021-12-23 16:38:02	nicola.varini@epfl.ch	AMGX with xlc?	User asked for an installation of AMGX libraries with XL compiler suite on MARCONI100.
22784	6	2021-12-15 11:50:01	cas@ipp.mpg.de	batch job MARCONI failed	User reported a job that failed because of an OPA instability. The issue was solved via a reboot of the OPA switches.
22786	6	2021-12-15 12:00:01	cas@ipp.mpg.de	CINECA_SCRATCH stale file handle	User reported a job that failed because of an OPA instability. The issue was solved via a reboot of the OPA switches.
22791	6	2021-12-15 12:58:01	alexey.mishchenko@ipp.mpg.de	problems to sbatch jobs (128 nodes)	User reported a job that failed because of an OPA instability. The issue was solved via a reboot of the OPA switches.

Ticket statistics Dec 13 – Jan 13

HPC-US-FIRST






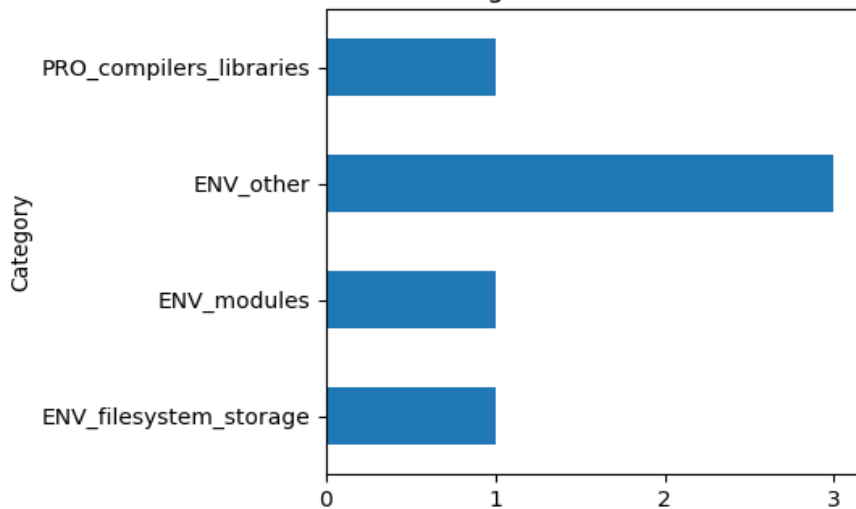
HPC-US-SECOND		By status		By platform		
		open	resolved	M100	SKL	SKL&M100
Information	AAA_accounting		1		1	
	ENV_modules		1	1		
	ENV_scheduler		1		1	
	SER_licences		1		1	
Problem	AAA_UserDB		1		1	
	AAA_access		2	1		1
	ENV_filesystem_storage		2		2	
	ENV_other		1		1	
	ENV_scheduler		2		2	
	SER_licences		1		1	
Service Request	AAA_access		5		4	1
	ENV_filesystem_storage		2	1	1	
	PRO_compilers_libraries	1			1	
total	21	1	20	3	16	2

Ticket statistics Dec 13 – Jan 13

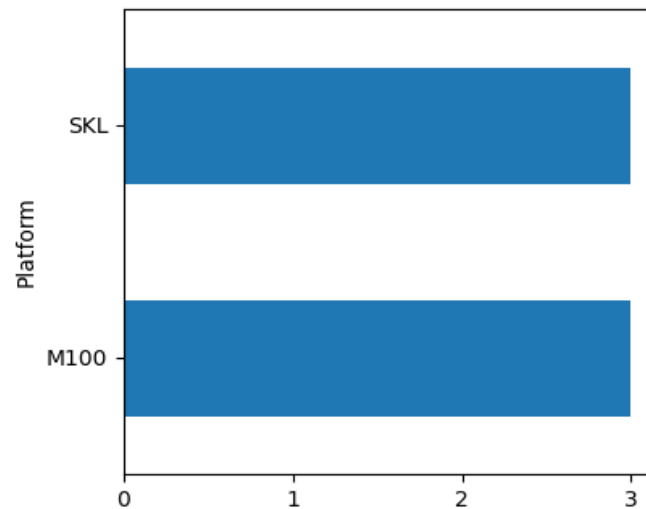
HPC-US-SECOND



Ticket categories on HPC-US-SECOND




Classification of tickets by platform
HPC-US-SECOND



Ticket statistics

Dec 13 – Jan 13


HPC-US-SECOND



HPC-US-SECOND		By status		By platform	
		open	resolved	M100	SKL
Problem	ENV_filesystem_storage	1		1	
	ENV_other		3		3
Service Request	ENV_modules		1	1	
	PRO_compilers_libraries		1	1	
total		1	5	3	3

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
application	openfoam	skl_fua_prod	normal	10	1,19
	openfoam-ext	skl_fua_prod	normal	5	4
	qe	skl_fua_prod	normal	836	10,96
	starccm+	bdw_all_serial	normal	1	0,02
		skl_fua_dbg	normal	9	5,96
		skl_fua_prod	normal	265	215,46
vasp	skl_fua_prod	normal	484	247,88	

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
compiler	gnu	bdw_all_serial	normal	1	1,01
		skl_fua_dbg	normal	13	1,15
		skl_fua_prod	normal	414	243,27
	intel	skl_fua_dbg	normal	143	13,3
		skl_fua_prod	normal	145	70,18
	intelmpi	skl_fua_dbg	skl_qos_fualowprio	1	0,11
			normal	143	13,3
		skl_fua_prod	normal	145	70,18
			skl_qos_fualowprio	1	0,11
	openmpi	skl_fua_prod	normal	405	234,32
	python	bdw_all_serial	normal	81	74,37
		skl_fua_dbg	normal	41	24,24
		skl_fua_prod	normal	873	167,36
			skl_qos_fuabprod	79	40,42
skl_qos_fualprod			1	1	

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
tool	anaconda	bdw_all_serial	normal	70	64,39
		skl_fua_dbg	normal	10	9,33
		skl_fua_prod	normal	22	10,73
			skl_qos_fuabprod	19	1,77
	cmake	bdw_all_serial	normal	27	25,93
		skl_fua_dbg	normal	8	6,94
		skl_fua_prod	normal	135	92,6
			skl_qos_fuabprod	13	1,57
	skl_fua_prod	skl_qos_fualprod	50	14,64	
		cubegui	skl_fua_prod	normal	221
	idl	bdw_all_serial	normal	52	48,11
		skl_fua_dbg	normal	9	3,85
			normal	28	9,9
		skl_fua_prod	skl_qos_fualowprio	1	0,11
			skl_qos_fualprod	1	1
	scalasca	skl_fua_prod	normal	221	192,49
	scorep	skl_fua_prod	normal	221	192,49
totalview	skl_fua_dbg	normal	1	0,11	
valgrind	skl_fua_dbg	normal	13	1,15	
vtune	skl_fua_dbg	normal	31	6,79	
	skl_fua_prod	normal	36	7,05	

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	blas	bdw_all_serial	normal	52	48,11
		skl_fua_dbg	normal	239	40,24
		skl_fua_prod	normal	1699	505,31
			skl_qos_fuabprod	13	8,75
			skl_qos_fualowprio	52	14,8
	boost	skl_fua_prod	normal	12	3,19
	cubelib	skl_fua_prod	normal	221	192,49
	fftw	bdw_all_serial	normal	27	25,93
			normal	376	73,12
		skl_fua_prod	normal	2784	872,06
			skl_qos_fuabprod	181	128,87
			skl_qos_fualowprio	53	14,91
			skl_qos_fualprod	84	39,03
	hdf5	skl_fua_dbg	normal	149	26,4
		skl_fua_prod	normal	2618	845,01
			skl_qos_fuabprod	127	69,7
skl_qos_fualprod			94	33,96	

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	lapack	bdw_all_serial	normal	52	48,11
		skl_fua_dbg	normal	239	40,24
		skl_fua_prod	normal	946	385,97
			skl_qos_fuabprod	13	8,75
			skl_qos_fualowprio	52	14,8
	metis	skl_fua_prod	normal	2	2
			skl_qos_fuabprod	8	4,96
	mkl	skl_fua_dbg	normal	143	13,3
		skl_fua_prod	normal	145	70,18
			skl_qos_fualowprio	1	0,11
	mumps	skl_fua_prod	normal	1	0,66
			skl_qos_fuabprod	8	4,96
	nag	bdw_all_serial	normal	22	19,94
	netcdf	skl_fua_dbg	normal	94	13,96
		skl_fua_prod	normal	1904	368,1
			skl_qos_fuabprod	91	41,95
			skl_qos_fualprod	43	18,32
	netcdf-cxx4	skl_fua_dbg	normal	24	4,66
		skl_fua_prod	normal	275	147,68
			skl_qos_fualprod	43	18,32

Module usage on Marconi-SKL


December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	netcdf	skl_fua_dbg	normal	70	9,3
		skl_fua_prod	normal	1629	220,41
			skl_qos_fuabprod	91	41,95
	numpy	skl_fua_prod	normal	1	0
	parmetis	skl_fua_dbg	normal	3	2,11
		skl_fua_prod	normal	120	88,28
			skl_qos_fuabprod	8	4,96
	petsc	skl_fua_dbg	normal	38	4,55
		skl_fua_prod	normal	1147	256,21
			skl_qos_fuabprod	66	38,85
			skl_qos_fualprod	1	1
	qt	skl_fua_prod	normal	221	192,49
	scalapack	bdw_all_serial	normal	52	48,11
		skl_fua_dbg	normal	35	0,99
		skl_fua_prod	normal	1270	243,61
			skl_qos_fuabprod	13	8,75
			skl_qos_fualowprio	52	14,8
scipy	skl_fua_prod	normal	9	8,95	

Module usage on Marconi-SKL

December 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	gzip	bdw_all_serial	normal	4	4,03
		skl_fua_prod	normal	132	75,92
			skl_qos_fuabprod	5	4,25
	zlib	bdw_all_serial	normal	4	4,03
		skl_fua_prod	normal	19	13,56
			skl_qos_fuabprod	3	2,88