

56th Ticket Meeting

HPC User Support @ CINECA
September, 14th 2021


Content



- Status of the clusters - main events affecting production [July 13 – September 13]
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to SchedMD support
 - escalated to Intel support
 - escalated to NVIDIA support
 - other active tickets
- Ticket statistics on queue [July 13 – September 9]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL

Status of the clusters July 13 – September 13

Main events affecting production



July 16: login issue and recommendations on Marconi100 cluster when using hpc-sdk suite compilers

https://www.hpc.cineca.it/center_news/marconi100-login-issue-0

https://www.hpc.cineca.it/center_news/marconi100-login-issue-solved-and-recommendations-about-hpc-sdk-suite-usage

July 21: high occupancy on Marconi SCRATCH filesystem

https://www.hpc.cineca.it/center_news/marconi-scratch-almost-full-quota-imposed-0

https://www.hpc.cineca.it/center_news/marconi-scratch-quota-removed-1

July 21: scheduled maintenance operations on Marconi on July 28th

https://www.hpc.cineca.it/center_news/maintenance-marconi-july-28th

https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi-tomorrow-july-28th

https://www.hpc.cineca.it/center_news/marconi-back-production-42

July 28: reduced production on Marconi100 cluster due to extraordinary heatwave

https://www.hpc.cineca.it/center_news/marconi100-reduced-production-due-extraordinary-heatwave

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

July 30: reminder for potential login issues on Marconi100 cluster when using hpc-sdk suite compilers

https://www.hpc.cineca.it/center_news/marconi100-no-use-hpc-sdk-compilers-login-nodes

Aug 4: scheduled maintenance operations on Marconi on August 10th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-next-tuesday-august-10th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-tomorrow-august-10th

https://www.hpc.cineca.it/center_news/marconi-back-production-43

Status of the clusters July 13 – September 13

Main events affecting production



Aug 6: high occupancy on Marconi100 SCRATCH filesystem

https://www.hpc.cineca.it/center_news/marconi100-scratch-almost-full-0

https://www.hpc.cineca.it/center_news/marconi100-scratch-quota-removed-0

Aug 13: reduced Help Desk service on August 16th to 20th week

https://www.hpc.cineca.it/center_news/help-desk-service-next-week

Aug 18: scheduled maintenance operations on Marconi100 on August 24th

https://www.hpc.cineca.it/center_news/marconi100-scheduled-maintenance-next-tuesday-august-24th

https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi100-tomorrow-august-24th

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

Sept 1: filesystem issues on Marconi cluster

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

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

Sept 2: scheduled maintenance operations on Marconi on September 8th

https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-next-wednesday-september-8th


https://www.hpc.cineca.it/center_news/marconi-back-production-44

Sept 7: technical issues on the ticketing system

https://www.hpc.cineca.it/center_news/technical-issues-ticketing-system-0


https://www.hpc.cineca.it/center_news/technical-issues-ticketing-system-solved

Tickets escalated to SchedMD support



Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
16506	Allocation of GPU jobs in M100	Leo.Ma@ukaea.uk	2021-03-17 16:43:09	M100	2nd SchedMD	<p>The allocation of 1 GPU card job to nodes is pretty random, and not packing tightly; It may be a good idea to pack jobs with some ordering, and occupy less nodes</p> <p>---</p> <p>After a long exchange of messages with SchedMD, we found the way to have four 1-gpu jobs packed on the same node; It is simply needed to discard the directive:</p> <pre>#SBATCH --ntasks-per-node=1</pre> <p>or modify it as follows:</p> <pre>#SBATCH --ntasks=1</pre> <p>If not set the ntasks-per-node value will be set to the default (1).</p> <p>A "bug" (or a desired "feature", the SLURM support is still investigating on the matter) is triggered by the --ntasks-per-node directive, forcing the selection of the first free cores per node, and slurm (correctly) reacts choosing a different node to enforce the cpu-gres binding.</p> <p>SchedMD reported: «Fixing this requires a notable change in the select plugin. I am afraid that a proper fix will be available only in slurm 21.08»</p> <p>slurm 21.08 has been released on August 26th. This has been installed on Marconi cluster on September 8th and will be installed also on Marconi100 cluster during next scheduled maintenance operations on September 21st.</p>

Tickets escalated to **SchedMD** support




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
18895	Unclear error message on MARCONI ...	joerg.riemann@ip p.mpg.de	2021-06-24 14:54:01	SKL	2nd SchedMD	<p>The user report randomly observed errors in his jobs as in «srun: error: Unable to create step for job 9552876: Memory required by task is not available».</p> <p>After some extensive checks we made, we got an additional information reported for the second mpirun call: “srun: error: Unable to create step for job 3130006: Memory required by task is not available”</p> <p>It appears that the first mpirun (or, better, the first srun called by the first mpirun) isn't able to free the allocated memory fast enough for give room to the second one.</p> <p>Hence the observed behavior is caused by a slurm bug that was addressed to the slurm support team: “Second mpirun in a job fails after Slurm upgrade”.</p> <p>SchedMD suggested a workaround while waiting for the bug to be fixed in 20.11.8 version. We indicated the user to put a "sleep 5" between the two mpirun commands in the jobscript.</p> <p>We planned to perform tests using slurm 21.08, that has been installed on Marconi cluster on September 8th to check for the presence of this bug fix.</p>

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.mochalsky@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 show 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:</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.</p> <p>In order to generate the right code, it can run in a non-vector mode.</p> <p>How ever, 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.</p> <p>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>




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. Additional information retrieved to estimate the correct value of the Inter-Socket bandwidth on SKL in terms of GB/s. Performed the osu_mbr_mr benchmark (uni-directional test) that supposes to test the inter socket bandwidth using different number of MPI pairs for simultaneous data transfer.</p> <p>Intel MPI team recommended the usage of a different tool for intra-node bandwidth evaluation, Intel MLC (Intel Memory Latency Checker), that is a low-level test suite which does not use all the optimizations that can be found in Intel MPI PingPong or Uniband tests, such as switch between temporal/non-temporal stores, AVX/AVX2/AVX512 based implementations, which may work on different frequencies and that may lead to CPU drive copy slowdowns. A report has been prepared and provided with results obtained from tests performed with this tool.</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




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 to recompile the paraview version (in progress), we will include the installation of the nvidia plugin index required also by this user.
19157	Error on marconi	axel.koenies@ipp.mp.g.de	2021-07-07 11:32:02	SKL	2°	The user reported two jobs running euterpe code that wrote no output and asked for a budget refund of the lost computing resources. We performed checks for those jobs and we did not find any filesystem issue involved. We informed the user about this behaviour previously observed with euterpe code that has been fixed in other situations by recompiling the code; also, about our policy that do not foresees the budget refund due to system issues etc.

Other active 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 informed the user that the present installation of this module is a serial implementation to be used through RCM sessions.
19640	Data on Marconi100	Stanislas.Pamela@ukaea.uk	2021-08-05 13:52:02	M100	2nd	A group of users working on the same activity, AI related, report "disk quota exceeded" errors when using the FUAC5_LOWPRIO account work area. We suggested these users to work on the m100_scratch area but they need to share data as in a regular project (these users are associated only to this account on this cluster). This problem has been reported by the user to B.Duarte to ask for the possibility of activate a new account on this cluster with no budget to make use only of the account work area.
19953	Complain on the low priority queue of m100	adisi@ipp.mpg.de	2021-08-23 23:08:01	M100	1st OC	The user asks for the possibility of implementing a maximum number of running jobs per user on the QOS qos_lowprio on Marconi100 cluster.


Other active tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
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.
20018 19594	Tuning of GKNET code for MARCONI 100	imadera.kenji.7z@kyoto-u.ac.jp	2021-08-26 09:22:01	M100	2nd	The user asked for help to resolve the following issues: 1) how to run multi-GPU this code that at present run only on a single GPU. 2) how to implement GPU-GPU data transfer, as the current version does not use GPU-GPU data transfer so that MPI one-to-one communication can become the bottleneck.
20019	Segfault before code launch with srun	Patrick.TAMAIN@cea.fr	2021-08-26 09:28:03	SKL	2nd	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. Investigations are underway.

Ticket Statistics July 13 – September 9


HPC-US-FIRST & HPC-US-SECOND



Received tickets by date




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



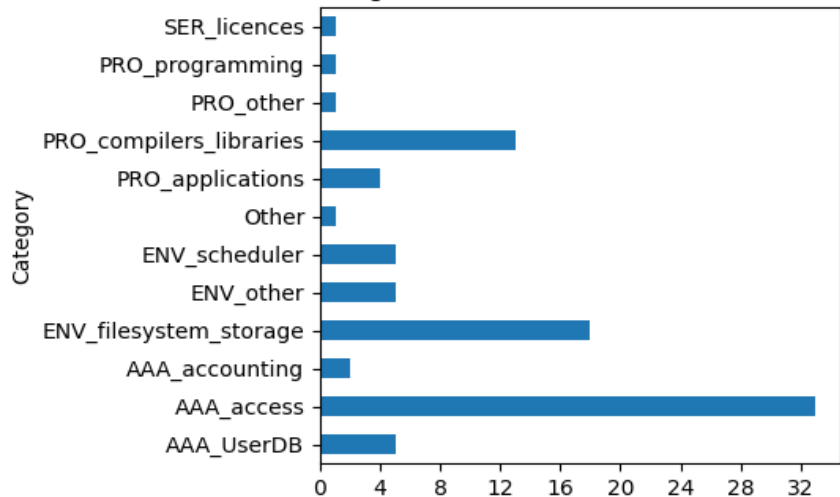
TOTAL → 89 tickets
 - HPC-US-FIRST → 60
 - escalated to HPC-US-SECOND → 29

Ticket statistics July 13 – September 9


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 July 13 – September 9

HPC-US-FIRST & HPC-US-SECOND

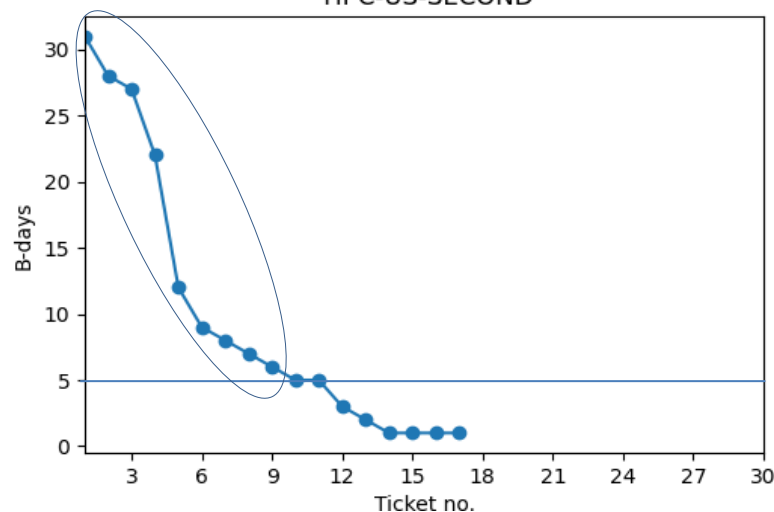


time from ticket creation to ticket resolution
 HPC US FIRST



54 / 60 resolved tickets

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




17 / 29 resolved tickets




Ticket FIRST	Days	Creation time	Requestor	Subject	Notes
19304	3	15/07/2021 11:28:01	David.Coster@ipp.mpg.de	Could you please reset my password (uid = dcoster0)?	The password reset was performed the on day 1st, ticket closed after 3 days.

Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
19431	31	2021-07-23 17:16:01	qian.xia@ukaea.uk	install mayavi	The user reported issues when installing this tool on Marconi. We provided the user all required steps/info for its installation.
19386	28	2021-07-21 15:00:02	rmh@ipp.mpg.de	ifort: command line warning #10006: ignoring unknown option '-fallow-argument-mismatch'	With Intel we have found and installed a patch in the intel mpi 2020 to the mpif90 wrapper that adds the flag -fallow-argument-mismatch only when using a gfortran compiler. We also checked that this bug has been already solved in the evolution of the intel parallel studio, intel oneapi.
19404	27	2021-07-22 13:10:01	ihor.holod@ipp.mpg.de	MUMPS library	We provided the user instructions on how to install the library on Marconi100 by using spack. The user replied back after some time to allow the closure of the ticket.



Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
19362	22	2021-07-20 08:54:01	ihor.holod@ipp.mpg .de	xl libraries	The user asked for the possibility of having new modules for fftw, hdf5 and zlib libraries on Marconi100 cluster based on xl compiler and spectrum_mpi libraries.
19372	12	2021-07-20 15:42:01	bocharov@latnet.lv	FUA35_ID2DRM: to add my collaborator	The user asked us to allow access to the PI of the project to his data on scratch and work areas. We informed the user on how to proceed, the ticket was closed late waiting for user's feedback.
19274	9	2021-07-14 10:16:01	christoph.slaby@ipp .mpg.de	Unable to submit jobs on Marconi	The user reported a job batch submission error that could not be reproduced in further tests, but it seemed related to a conflict when using the old env-skl module together with the tcsh shell.
19256	8	2021-07-13 12:50:01	jalc@ipp.mpg.de	Data Transfer	The user asked for information on how to proceed to transfer his project data from Marconi to a protected location, as due to security restrictions he is not allowed to use scp or rsync protocols. We proposed the user alternative transfer protocols, the ticket was closed late waiting for user's feedback.


Ticket statistics July 13 – September 9




Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
19391	7	2021-07-21 17:30:01	giovanni.bonny@sc kcen.be	disk quota exceeded	The user reported « disk quota exceeded errors » when using marconi_scratch area. We informed the user that this was due to the user's quota imposed during high global occupancy on this area. We also informed the user when this quota was removed.
20100	6	2021-08-31 14:14:02	rmh@ipp.mpg.de	new password	We provided to reset the password as a consequence of issues observed with the ssh key when trying to connect to the cluster.

Ticket statistics July 13 – September 9


HPC-US-FIRST



Ticket categories on HPC-US-FIRST




Classification of tickets by platform
HPC-US-FIRST



Ticket statistics July 13 – September 9

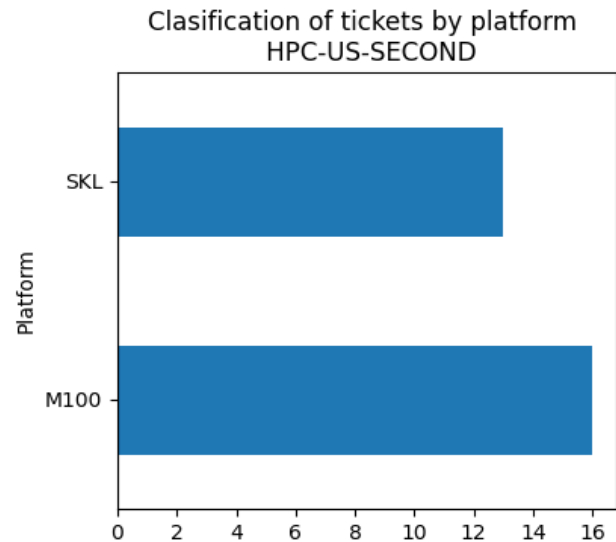


HPC-US-FIRST



HPC-US-FIRST		By status			By platform			
		open	resolved		SKL	M100	SKL&M100	Other
Information	AAA_accounting		1		1			
	ENV_filesystem_storage		6		5			1
	ENV_scheduler		1		1			
	PRO_compilers_libraries		1		1			
Problem	AAA_access	2	4		1			
	ENV_filesystem_storage		1		6			
	ENV_other		2		1			
	PRO_applications		1		1	1		
	PRO_compilers_libraries	2	3		1			
Service Request	AAA_UserDB	1	3		1	4		
	AAA_access	1	19		4			
	AAA_accounting		1		15		5	
	ENV_filesystem_storage		7		1			
	ENV_scheduler		1		4	1	1	1
	Other		1			1		
	SER_licences		1		1			
Other	AAA_access		1		1			
total	60	6	54		45	7	6	2


Ticket statistics July 13 – September 9

HPC-US-SECOND



Ticket statistics July 13 – September 9


HPC-US-SECOND



HPC-US-SECOND		By status		By platform	
		open	resolved	M100	SKL
Information	ENV_filesystem_storage		2		2
	ENV_scheduler		1	1	
	PRO_applications		1		1
Problem	ENV_filesystem_storage	1		1	
	ENV_other	1			1
	ENV_scheduler	1	1	1	1
	PRO_applications	1	1	1	1
	PRO_compilers_libraries	3	1	3	1
	PRO_programming		1	1	
Service Request	AAA_UserDB	1		1	
	AAA_access	2	4	2	4
	ENV_filesystem_storage		1		1
	ENV_other	1	1	1	1
	PRO_compilers_libraries	1	2	3	
	PRO_other		1	1	
total	29	12	17	16	13

Module usage on Marconi-SKL


Batch jobs on July 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
compiler	intel	skl_fua_dbg	normal	4	0,04
		skl_fua_prod	normal	30	2,86
	intelmpi	skl_fua_dbg	normal	4	0,04
		skl_fua_prod	normal	28	2,8
	python	skl_fua_prod	normal	69	12,83
skl_qos_fuabprod			29	8,09	
library	fftw	skl_fua_dbg	normal	1	0,15
		skl_fua_prod	normal	85	15,6
			skl_qos_fuabprod	32	8,13
	hdf5	skl_fua_dbg	normal	1	0,15
		skl_fua_prod	normal	76	13,58
			skl_qos_fuabprod	32	8,13
	mkl	skl_fua_dbg	normal	4	0,04
		skl_fua_prod	normal	28	2,8
	netcdf	skl_fua_prod	normal	70	12,84
			skl_qos_fuabprod	32	8,13
	netcdf	skl_fua_prod	normal	69	12,83
			skl_qos_fuabprod	30	8,11
	petsc	skl_fua_prod	normal	72	12,84
			skl_qos_fuabprod	32	8,13
	popt	skl_fua_prod	normal	122	41,34
skl_qos_fuabprod			52	10,42	

Module usage on Marconi-SKL


Batch jobs on August 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum	
library	fftw	bdw_all_serial	normal	1	1	
		skl_fua_prod	normal	6	2,4	
	skl_qos_fuabprod		5	2		
	hdf5	skl_fua_prod	normal	6	2,4	
			skl_qos_fuabprod	5	2	
	mkl	skl_fua_prod	bdw_all_serial	normal	1	1
			skl_fua_dbg	normal	100	8,44
			skl_fua_prod	normal	13	4,26
	netcdf	skl_fua_prod	normal	6	2,4	
			skl_qos_fuabprod	5	2	
	netcdf	skl_fua_prod	normal	6	2,4	
			skl_qos_fuabprod	5	2	
	petsc	skl_fua_prod	normal	6	2,4	
			skl_qos_fuabprod	6	2,32	

Module usage on Marconi-SKL

Batch jobs on August 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
compiler	intel	bdw_all_serial	normal	1	1
		skl_fua_dbg	normal	100	8,44
		skl_fua_prod	normal	13	4,26
	intelmpi	bdw_all_serial	normal	3	1,04
		skl_fua_dbg	normal	100	8,44
		skl_fua_prod	normal	13	4,26
	python	bdw_all_serial	normal	1	1
			normal	6	2,4
		skl_fua_prod	skl_qos_fuabprod	5	2
tool	idl	bdw_all_serial	normal	1	1