

51st Ticket Meeting

HPC User Support @ CINECA
March, 18th 2021


Content Overview



- Status of the clusters - main events affecting production [Feb 18 – March 17]
- 5th Cycle allocation on Marconi and Marconi100 clusters
- Examination of active tickets on HPC-US-SECOND queue
 - escalated to Intel support
 - other tickets
- Ticket statistics on queue [reference period: Feb 15 – March 14]
 - HPC-US-FIRST
 - HPC-US-SECOND
- Module usage on Marconi-SKL [February 2020]

Status of the clusters Feb 18 – March 17

Main events affecting production



Feb 22: scheduled maintenance operations on Marconi100 cluster on February 23rd
https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-february-23th
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi100-february-23rd
https://www.hpc.cineca.it/center_news/marconi100-back-production-6

Feb 24: scheduled maintenance operations on Marconi cluster on March 2nd
https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi-march-2nd
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi-march-2nd
https://www.hpc.cineca.it/center_news/marconi-back-production-37

Mar 10: announced high occupancy on Marconi SCRATCH area (users' \$CINECA_SCRATCH quota imposed to all users)
https://www.hpc.cineca.it/center_news/marconi-scratch-almost-full-quota-imposed
https://www.hpc.cineca.it/center_news/marconi-scratch-quota-removed-0

Mar 11: scheduled maintenance operations on Marconi100 cluster on March 16th
https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-march-16th
https://www.hpc.cineca.it/center_news/m100-will-be-stopped-tomorrow-scheduled-maintenance
https://www.hpc.cineca.it/center_news/marconi100-back-production-7

Mar 16: announced issue in budget accounting on Marconi100 cluster
https://www.hpc.cineca.it/center_news/marconi100-issue-budget-accounting-0

Mar 17: Marconi100 upgrade March 23/24
https://www.hpc.cineca.it/center_news/marconi100-upgrade-march-2324

5th Cycle Allocation on Marconi and Marconi100




- All projects have been defined on the UserDB site on February 25th with validity:
start date: **01/03/2021**
end date: **28/02/2022**
- **FUAC5** → 33 projects with agreed allocation on **Marconi100**
33 defined on the UserDB site
32 activated on the cluster
- **FUA35** → 132 projects with agreed allocation on **Marconi SKL**
132 defined on the UserDB site
129 activated on the cluster
2 new projects activated on March 17th
- **FUSIO_ru5** → 20 research units' projects defined on both clusters
20 defined on the UserDB site
17 activated on the clusters
- **FUA35_LOWPRIO** and **FUAC5_LOWPRIO** accounts activated on the clusters

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.mochalskyy@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 nvdia to reproduce the tests on different dusters and to get more info about datapaths.</p>
9677	Marconi100 problem with GPU OpenMPI library	serhiy.mochalskyy@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>We expect to perform a system stack upgrade in the forthcoming weeks. Further tests will be performed after this upgrade.</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 resolved tickets on HPC-US-SECOND




Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
15018	jsoncpp and gnu bug	mattwi@fysik.dtu.dk	2021-01-13 16:46:01	m100	2nd	<p>The user has reported issues when trying to compile his code with nvcc (cuda/10.2) and gnu/8.4.0 including the jsoncpp library. The problem seems to lie in the nvcc compiler version and can be reproduced with cuda/10.2, using cuda/11.1 the compilation aborts.</p> <p>A switch from cuda to hpc-sdk was tried, the compilation worked passing explicitly the version 11.0 of the compiler that is contained inside the hpc-sdk module (module is set to work with 10.2, as m100 nodes do not have yet the drivers that support Cuda11). The user confirmed the correct compilation of the code but observed driver compatibility issues when running it.</p> <p>A new cuda/11.0 module is available at the candidate modules profile, using this it is possible to overcome the driver compatibility issue.</p>

Ticket Statistics Feb 15 – Mar 14


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	→ 104 tickets
- HPC-US-FIRST	→ 82
- escalated to HPC-US-SECOND	→ 22

Ticket statistics Feb 15 – Mar 14


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 Feb 15 – Mar 14

HPC-US-FIRST & HPC-US-SECOND




time from ticket creation to ticket resolution
HPC US FIRST




80 / 82 resolved tickets

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



17 / 22 resolved tickets

Ticket statistics Feb 15 – Mar 14




Ticket FIRST	Days	Creation time	Requestor	Subject	Notes
15745	3	2021-02-16 09:48:01	rmh@ipp.mpg.de	deleting of files	Our system administrators provided to delete a users' directory in the FUSIO_HLST work project area defined on Marconi-SKL.
16041	3	2021-03-01 14:04:02	komm@ipp.cas.cz	Lost password	The user requested a password reset that was immediately sent. The ticket was closed with some delay.


Ticket SECOND	Days	Creation time	Requestor	Subject	Notes
15744	16	2021-02-16 09:14:01	rmh@ipp.mpg.de	disk space on MARCONI	A new project collaborator has been allowed to access data in FUSIO_HLST/psverma area.
15919	10	2021-02-23 22:10:01	jan.wrobel@pw.edu.pl	request for an extension of my project FUA34_FeCrCN	The user asked for the possibility of extend the FUA34 project as he did not managed to ask for a new FUA35 allocation.
15965	9	2021-02-25 15:54:01	axel.koenies@ipp.mpg.de	EUROfusion project EUGY	The user asked for the transformation on the FUA34 work project area into the newest FUA35 work project area. Once the new project was activated on the cluster and the new work area created we provided to this transformation.
15993	7	2021-02-26 13:44:02	antonio.froio@polito.it	Richiesta versione software	We have installed the requested starccm+ version on Marconi-SKL and allowed this to reach the external license server.
16133	7	2021-03-03 20:10:02	J.GonzalezMunos@diffen.nl	Issue with SOLPS-ITER installation in new account	The user reported compilation issues to permission denies issues to /marconi_work/FUSIO_ALL/SOLPS-ITER repository. We allowed this user as a collaborator of the FUSIO_ALL project.

Ticket statistics Feb 15 – Mar 14

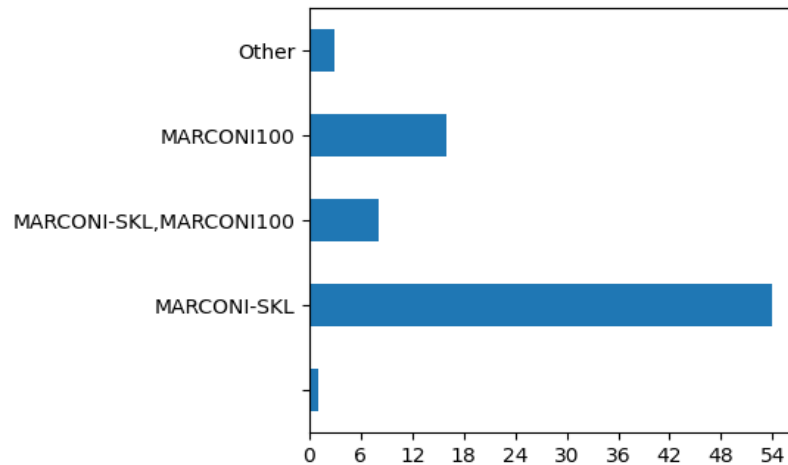
HPC-US-FIRST



Ticket categories on HPC-US-FIRST




Classification of tickets by platform
HPC-US-FIRST



Ticket statistics Feb 15 – Mar 14

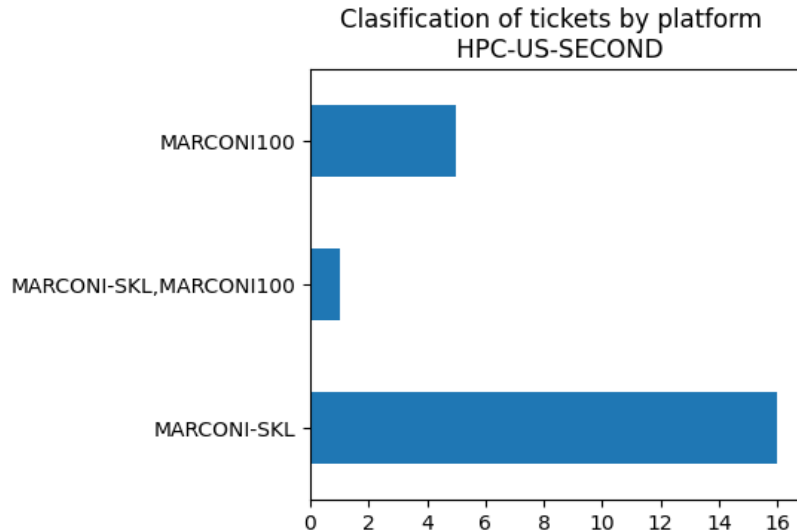


HPC-US-FIRST



HPC-US-FIRST		By status		By platform			
		open	resolved	SKL	SKL&M100	M100	Other
Information	AAA_UserDB		1	1			
	AAA_accounting		11	8	2	1	
	ENV_filesystem_storage		2	2			
	ENV_other		1			1	
	PRO_compilers_libraries		2	2			
	SER_other		1	1			
Problem	AAA_UserDB		6	3			3
	AAA_access		6	4		2	
	AAA_accounting		1	1			
	ENV_filesystem_storage		3	2		1	
	ENV_other		1	1			
	ENV_scheduler		5	3		2	
	PRO_applications		1			1	
	PRO_programming		1				1
Service Request	AAA_UserDB		5	4	1		
	AAA_access		7	5	1	1	
	AAA_accounting	2	12	4	4	5	
	ENV_filesystem_storage		13	13			
	PRO_compilers_libraries		1			1	
total	82	2	80	54	8	16	3


Ticket statistics Feb 15 – Mar 14

HPC-US-SECOND



Ticket statistics Feb 15 – Mar 14

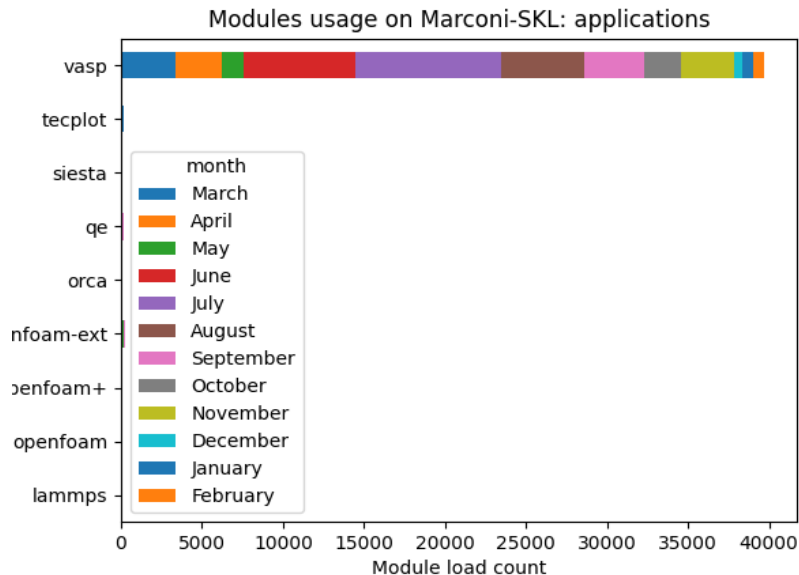

HPC-US-SECOND



HPC-US-SECOND		By status		By platform		
		open	resolved	SKL	M100	Other
Information	AAA_accounting		1	1		
	ENV_filesystem_storage		1	1		
	PRO_applications		1	1		
	PRO_compilers_libraries		1			1
	ENV_filesystem_storage		2	2		
	ENV_other	1	1	2		
	ENV_scheduler	1	1	2		
	PRO_applications		1			1
	PRO_compilers_libraries	1	2	2		1
	PRO_programming	1				1
Service Request	AAA_accounting		3	3		
	ENV_filesystem_storage		2	2		
	PRO_applications	1	1		1	1
total	22	5	17	5	1	16

Module usage on Marconi-SKL: applications


Batch jobs Mar 1 – Feb 28 (4th cycle)




	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
lammps	21							44				
openfoam	16	1		3		1	13					
openfoam+							2					
openfoam-ext	46	28	80	25	8		40					
orca												22
qe		4		33	20	29	59					
siesta	41	3	13				3					
tecplot					9	9	30	24	44	4	46	12
vasp	3387	2802	1362	6903	8997	5166	3702	2270	3262	524	624	728

Module usage on Marconi-SKL: tools

Batch jobs Mar 1 – Feb 28 (4th cycle)




Modules usage on Marconi-SKL: tools




	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
anaconda		66	150	54	75	49	143	86	84	54	230	86
cmake	295	2827	7805	235	233	536	93	14	30	98	954	6170
cubegui					40	12	300	388	416	154		116
cython	43											
extrae				17	64							
ffmpeg	46				2			12	74	28	78	130
idl									38	12		
scalasca	1385	708	651	369	871	1254	531	600	478	154		116
scorep					49	12	300	386	416	152		116
totalview			6	12								
valgrind	2	2								26		28
vtune									2			28

Module usage on Marconi-SKL: libraries

Batch jobs Mar 1 – Feb 28 (4th cycle)




Modules usage on Marconi-SKL: libraries




	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
blas	2187	3434	1902	2990	2431	2799	3154	3068	1358	744	1756	6902
boost	16	1		3	113	43	96					
cgal							2					
cubelib					49	12	297	392	414	154		116
fftw	4247	4025	3952	5569	4448	4922	6980	3708	3310	1664	1924	5436
gmp							2					
gsl					10							
hdf5	3902	6544	12073	5568	4363	1935	2405	2224	3062	3066	2452	9242
lapack	2230	3430	1916	2990	2517	2821	3242	3032	1360	748	1754	6880
libunwind				17	64							
libxc					411	141		14				
metis		12	58	82			105			6		
mkl	7						140			12	24	54
mpfr							2					
mumps		12	14				105					
nag	2	2876	7674	69	39		72	2	4		252	1760
netcdf	1118	5358	9803	4070	2577	768	1375	1234	1866	1786	1882	8342
netcdf-cxx4	88	645	381	450	780	207	567	1026	782	542	356	820
netcdfff	988	1745	1440	3535	1758	561	693	200	636	1160	574	1342
numpy	41	3	14						20	40		
papi				17	66							
parmetis		12	52	81			108			6		
petsc	472	249	222	290	170	327	866	338	482	244	322	214
popt											88	200
qt					40	12	299	392	412	154		116
scalapack	757	1352	488	121	975	2243	1322	230	232	214	1426	4714
scipy	2							18		42	62	48
slepc						3		18			22	44
tensorflow		792					3		4			
wsmc	384	21	177	258	33							
xerces-c	147	12	15	9		456						28
zlib							2					

Module usage on Marconi-SKL: compilers

Batch jobs Mar 1 – Feb 28 (4th cycle)




Modules usage on Marconi-SKL: compilers



	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
gnu	15176	8943	8603	742	3421	489	1248	910	986	228	1544	17714
intel	10						140	10		12	42	54
intelmpi	10						140	10		12	42	54
intelmpi-eb			6									
openmpi	2526	1329	308	406	2775	435	897	904	980	206	584	1876
python	12494	8115	9754	798	1915	1747	5396	838	5600	6648	3252	16126

Module usage on Marconi-SKL: applications


Batch jobs Feb 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
application	openfoam-ext	skl_fua_prod	normal	1	0,28
	qe	skl_fua_prod	normal	9	7,93
	tecplot	skl_fua_prod	skl_qos_fualowprio	6	2,65
	vasp	skl_fua_prod	normal	360	137,12

Module usage on Marconi-SKL: tools


Batch jobs Feb 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
tool	anaconda	bdw_all_serial	normal	42	35,4
		skl_fua_dbg	normal	1	1,01
		skl_fua_prod	normal	36	26,78
	cmake	bdw_all_serial	normal	41	34,23
		skl_fua_dbg	normal	381	198,89
		skl_fua_prod	normal	1545	316,2
	cubegui	skl_fua_prod	normal	36	26,64
	cython	bdw_all_serial	normal	4	0,13
	ffmpeg	bdw_all_serial	normal	4	0,13
		skl_fua_prod	normal	64	36,86
	idl	bdw_all_serial	normal	36	34,09
		skl_fua_prod	normal	7	0,04
	scalasca	skl_fua_prod	normal	36	26,64
	scorep	skl_fua_prod	normal	36	26,64
	valgrind	skl_fua_dbg	normal	14	13,14
vtune	skl_fua_dbg	normal	14	1,42	

Module usage on Marconi-SKL: libraries I


Batch jobs Feb 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	blas	bdw_all_serial	normal	37	34,1
		skl_fua_dbg	normal	525	236,73
		skl_fua_prod	normal	3705	1214,38
		skl_fua_prod	skl_qos_fuabprod	38	21,51
	boost	skl_fua_prod	normal	12	9,12
	cubelib	skl_fua_prod	normal	36	26,64
	fftw	bdw_all_serial	normal	41	34,23
		skl_fua_dbg	normal	550	260,59
			normal	3078	860,27
		skl_fua_prod	skl_qos_fuabprod	6	2,53
			skl_qos_fualowprio	23	9,72
	skl_qos_fualprod		55	17,31	
	hdf5	bdw_all_serial	normal	14	6,25
		skl_fua_dbg	normal	466	212,54
			normal	3044	793,91
		skl_fua_prod	skl_qos_fuabprod	18	6,88
			skl_qos_fualowprio	23	9,72
	skl_qos_fualprod		394	136,85	
	lapack	bdw_all_serial	normal	41	34,23
		skl_fua_dbg	normal	523	235,31
		skl_fua_prod	normal	3379	1194,4
			skl_qos_fuabprod	38	21,51
	mkl	skl_fua_dbg	normal	27	9,36
		skl_fua_prod	normal	61	21,9
nag	skl_fua_dbg	normal	15	13,28	
	skl_fua_prod	normal	1	0,28	

Module usage on Marconi-SKL: libraries II


Batch jobs Feb 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum	
library	blas	skl_fua_dbg	normal	525	236,73	
		skl_fua_prod	normal	3705	1214,38	
			skl_qos_fuabprod	38	21,51	
	boost	skl_fua_prod	normal	12	9,12	
	cubelib	skl_fua_prod	normal	36	26,64	
	fftw	bdw_all_serial	normal	41	34,23	
			skl_fua_dbg	normal	550	260,59
		skl_fua_prod	normal	3078	860,27	
			skl_qos_fuabprod	6	2,53	
			skl_qos_fualowprio	23	9,72	
			skl_qos_fualprod	55	17,31	
		hdf5	bdw_all_serial	normal	14	6,25
	skl_fua_prod		skl_fua_dbg	normal	466	212,54
			normal	3044	793,91	
			skl_qos_fuabprod	18	6,88	
			skl_qos_fualowprio	23	9,72	
	skl_qos_fualprod	394	136,85			
	lapack	bdw_all_serial	normal	41	34,23	
		skl_fua_prod	skl_fua_dbg	normal	523	235,31
			normal	3379	1194,4	
			skl_qos_fuabprod	38	21,51	
	mkl	skl_fua_dbg	normal	27	9,36	
		skl_fua_prod	normal	61	21,9	
	nag	skl_fua_dbg	normal	15	13,28	
		skl_fua_prod	normal	1	0,28	

Module usage on Marconi-SKL: **compilers**

Batch jobs Feb 2021



category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
compiler	gnu	bdw_all_serial	normal	5	0,14
		skl_fua_dbg	normal	344	189,48
		skl_fua_prod	normal	2461	554,3
	intel	skl_fua_dbg	normal	27	9,36
		skl_fua_prod	normal	61	21,9
	intelpi	skl_fua_dbg	normal	27	9,36
		skl_fua_prod	normal	61	21,9
	openmpi	skl_fua_dbg	normal	2	1,1
		skl_fua_prod	normal	952	280,36
	python	bdw_all_serial	normal	47	35,54
		skl_fua_dbg	normal	373	198,77
		skl_fua_prod	normal	1707	403,47
			skl_qos_fuabprod	6	2,53