CINECA SCAI
SuperComputing Applications and Innovation

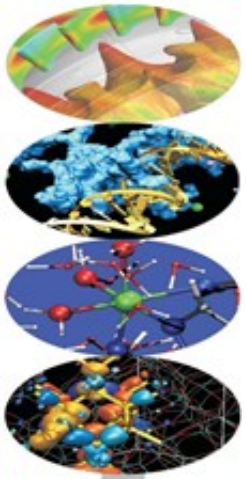

50th Ticket Meeting

HPC User Support @ CINECA
February, 18th 2021

CINECA

Content Overview

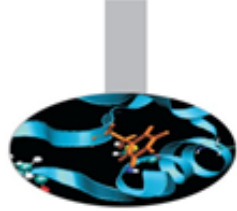

- Status of the clusters - main events affecting production [Jan 14 – Feb 17]
- Examination of active tickets on HPC-US-SECOND queue
  - escalated to Intel support
  - other tickets
- Ticket statistics on queue [reference period: Jan 11 – Feb 14]
  - HPC-US-FIRST
  - HPC-US-SECOND
- VASP usage on Marconi-SKL [March 2020 - January 2021]
- Module usage on Marconi-SKL [January 2020]

Status of the clusters Jan 14 – Feb 17

Main events affecting production

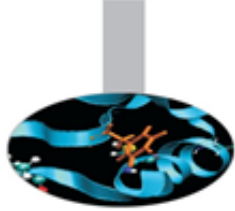

Jan 14: announced to all users a high occupancy on Marconi100 scratch area
https://www.hpc.cineca.it/center_news/marconi100-scratch-almost-full

Jan 19: reminder for the activation of automatic cleaning procedure on Marconi100 scratch filesystem on January 26th
https://www.hpc.cineca.it/center_news/marconi100-reminder-activation-cleaning-procedure-scratch-area

Jan 19: announced scheduled maintenance operations of Marconi100 on January 26th
https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-january-26th
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-and-activation-cleaning-procedure-marconi100-january-26th

Jan 19: announced to all EUROfusion users the possibility to get access to Marconi100 cluster and to submit jobs using the low priority QOS

Jan 25: reminder for scheduled maintenance operations and activation of automatic cleaning procedure on Marconi100 on January 26th
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-and-activation-cleaning-procedure-marconi100-january-26th
https://www.hpc.cineca.it/center_news/marconi100-back-production-5

Jan 29: preemption activated on Marconi100 (not announced yet to academic users): a new “m100_usr_preempt” partition (PriorityTier=1) has been defined on the nodes of “m100_fua_prod” partition (PriorityTier=100).

Feb 1: reminder for scheduled maintenance operation on Marconi on February 2nd.
https://www.hpc.cineca.it/center_news/reminder-scheduled-maintenance-marconi-tomorrow-february-2nd
https://www.hpc.cineca.it/center_news/marconi-back-production-36

Feb 16: announced scheduled maintenance operations of Marconi100 on February 23rd
https://www.hpc.cineca.it/center_news/scheduled-maintenance-marconi100-february-23th

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**	Intel had checked the issue with the last two compiler releases and reported that it worked; they provided a tar file with the source, BUILD script and output files from the BUILD script . We have reported this information to the user. We could observe that the code provided by Intel in "bp_test.tar" doesn't present the issues, as written in the code: ... ! from customer re: 19.0u5 -- Commenting any of these defines makes it work #define USE_SIMD !#define USE_WPR !#define USE_UNUSED_CODEPARTS !#define USE_ARR_IN_MODULE #define USE_POINTERS ... We reported to Intel: *Without using "USE_ARR_IN_MODULE" it works. With it doesn't work. In attachment you will find a tar file (bp_test_new.tar), keeping the same syntax of the previous source where the issue is still present. According to the code developer the "USE_ARR_IN_MODULE" must be used.* … Intel filed a bug report for the issue: CMPLRIL0-33599

Other active tickets on HPC-US-SECOND

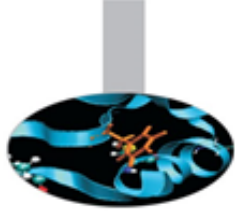

Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
9462	Marconi100 issue 3 from the pre-production testing: Low X Bus bandwidth during Nvidia P2P benchmark	serhiy.mochalskyy@ipp.mpg.de	2020-05-12 16:48:02	M100	2nd IBM Nvidia	Results from tests show values close to the theoretical value bandwidth during the communication of 2 GPUs inside one socket. 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. To discuss with IBM and Nvidia to reach a better understanding of how the communications intersocket work whether p2p is enabled or not. Nvidia tested the p2p communication tool coming with the Cuda suite, and they also notice the problem when the communication involves extra-socket GPUs. We are working with nvdia to reproduce the tests on different clusters and to get more info about datapaths.
9677	Marconi100 problem with GPU OpenMPI library	serhiy.mochalskyy@ipp.mpg.de	2020-05-18 11:00:01	M100	2nd IBM	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. We expect to perform a system stack upgrade in the forthcoming weeks. Further tests will be performed after this upgrade.

Other active tickets on HPC-US-SECOND

Ticket	Subject	Requestors	Created	Comments
15031	File system issues	axel.koenies@ipp.mpg.de	2021-01-14 10:20:02	Users report failed jobs on Marconi SKL partition. First investigations on these failed jobs showed OPA instabilities on the involved nodes, connected to a specific switch, that caused the jobs to fail due to the expulsion of a node or to a sudden unmounting of the file system on a node (without expulsion of the node). Our system administrators provided to do an electrical reboot of the involved switch but it did not fixed the problem. The switch was then replaced and the linked nodes drained, but the problem showed up again. Subsequent investigations and interventions of our system administrators could identify a problematic cable that has been substituted. After this substitution the problem has not showed up again.
15107	still problems with the file system	axel.koenies@ipp.mpg.de	2021-01-17 00:32:01	
15269	Abnormal job termination on MARCONI	christoph.slaby@ipp.mpg.de	2021-01-25 15:16:02	
15374	Re: Slurm Job_id=8784016 [...]	jason.parisi@physics.ox.ac.uk	2021-01-28 16:04:02	
15523	information about job failures	gmerlo@ices.utexas.edu	2021-02-04 00:40:01	
15543	Issue with saving files on Marconi	aylwiniantchenko@gmail.com	2021-02-04 15:18:02	

Other active tickets on HPC-US-SECOND

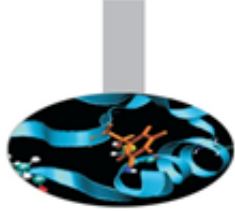

Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
13850	nvidia hpc-sdk + cuda	thomas.hayward@ipp.mpg.de	2020-11-19 14:54:02	M100	2nd nvidia	The user is been supported for the compilation of his code (orb5) using hpc-sdk/2020--binary, gnu/8.4.0, and cuda/11.1 modules available on M100 cluster. The compilation of this code fails (but it could be correctly compiled with 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. 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.

Other active tickets on HPC-US-SECOND

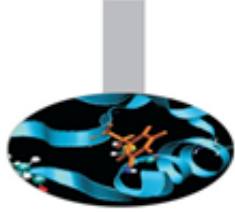

Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
15018	jsoncpp and gnu bug	mattwi@fysik.dtu.dk	2021-01-13 16:46:01	m100	2nd	The user has reported issues when trying to compile his code with nvcc (cuda/10.2) and gnu/8.4.0 including the jsoncpp library. The problem seems to lie in the nvcc compiler version and can be reproduced with cuda/10.2, using cuda/11.1 the compilation aborts. 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. We are working to make the application see the libraries for compatibility.

Other resolved tickets on HPC-US-SECOND

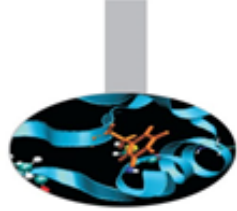

Ticket	Subject	Requestors	Created	Host	Supp. Level	Comments
14109	run SSH singularity Marconi	albert.gutierrez@bsc.es	2020-11-30 16:32:01	SKL	2nd	The user is worknig with singularity to prepare several containers to run MPI. He is experiencing issues with the configuration of SSH required for the communication between the containers. We have provided information to the user about the system for the containers configuration.
14725	errori con MPI...	gregorio.vlad@enea.it	2020-12-24 16:06:02	SKL	1st	The user has reported some failed jobs that might be failed due to some issues on the assigned compute nodes. Investigations performed on the compute nodes showed no issues; the user managed to avoid the errors by lowering the optimization level during the code compilation.
14912	Connect globus to cineca	K.L.vandePlassche@differ.nl	2021-01-08 12:24:02	User DB	Esc MW	The user is experiencing issues to add his new x509 certificate into his personal page on our UserDB site. We have applied some changes into the UserDB site procedure that manages x509 certificates information.

Ticket Statistics Jan 11 – Feb 14

HPC-US-FIRST & HPC-US-SECOND

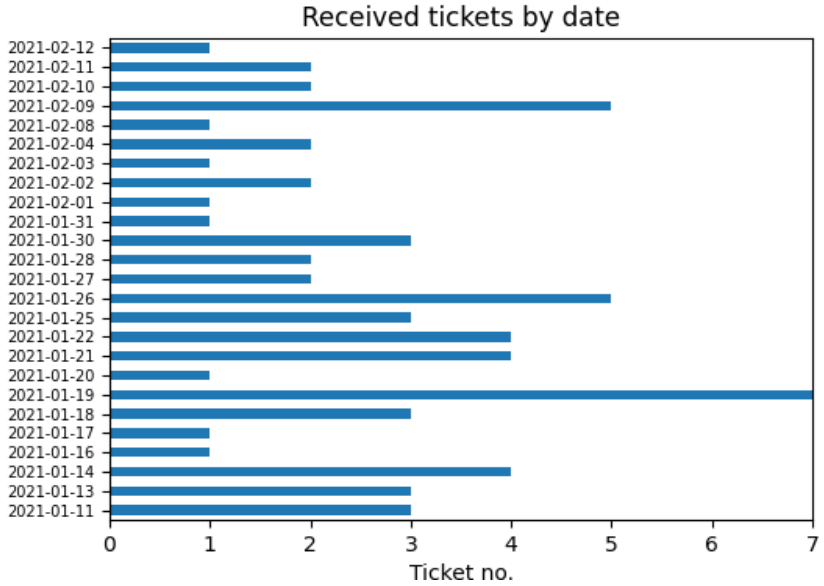

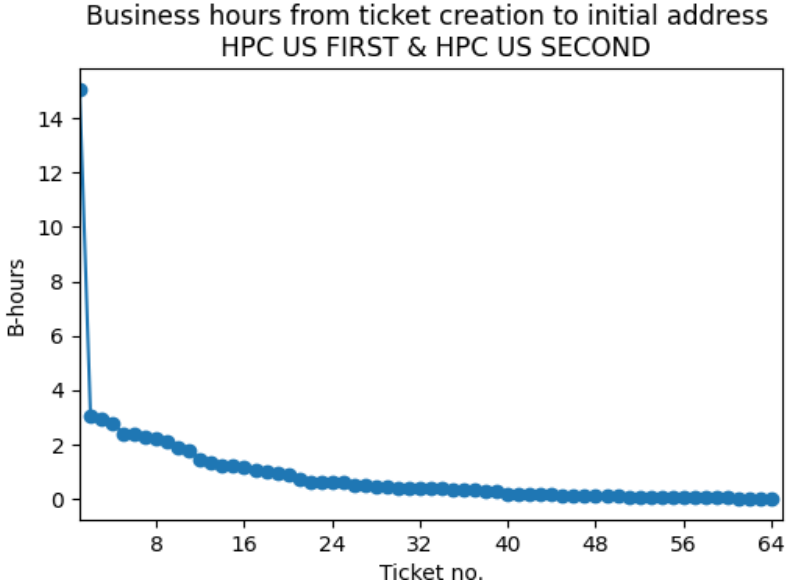

TOTAL	→ **64** tickets
- HPC-US-FIRST	→ **37**
- escalated to HPC-US-SECOND	→ **27**

Ticket statistics Jan 11 – Feb 14

HPC-US-FIRST & HPC-US-SECOND

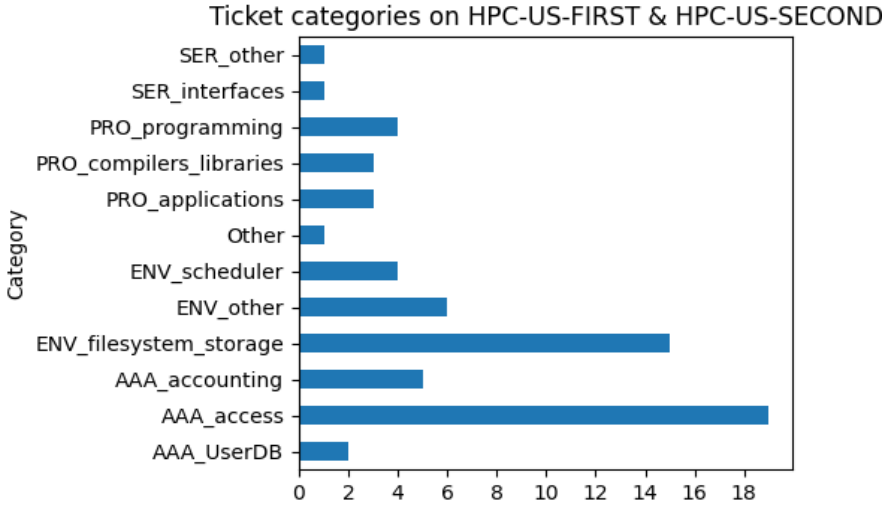

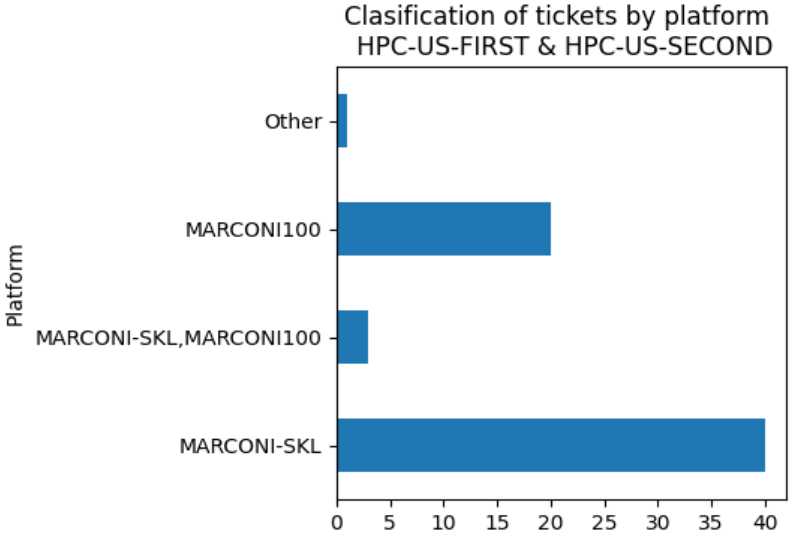

Ticket statistics Jan 11 – Feb 14

HPC-US-FIRST & HPC-US-SECOND

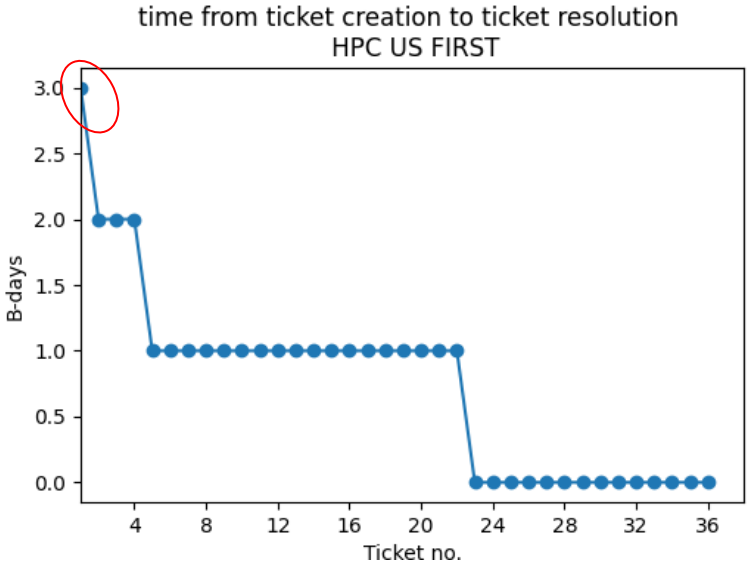

resolved tickets

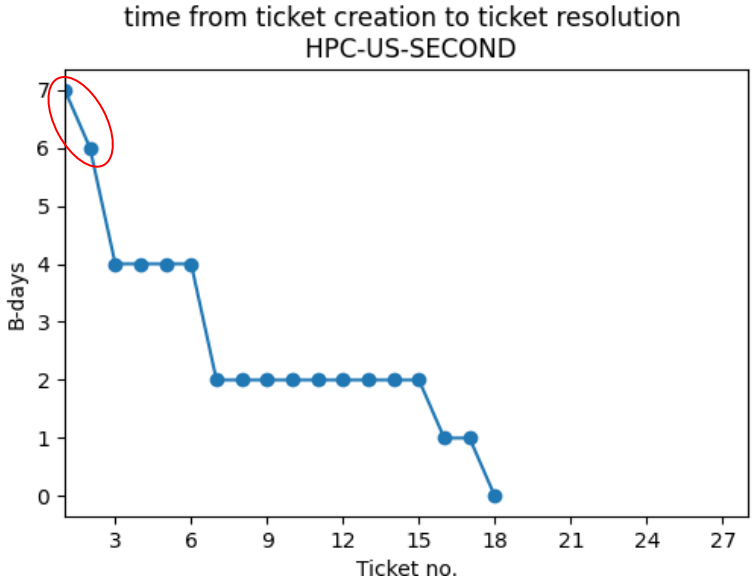

resolved tickets

Ticket statistics Dec 7 – Jan 10

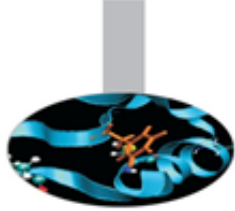

HPC-US-FIRST

Ticket	Days	Creation time	Requestor	Subject	Notes
15330	3	2021-01-27 15:10:02	Leo.Ma@ukaea.uk	increase disk quota in M100	We have increased the work area quota for FUSIO_ru4CCFE_0 account, defined on Marconi100, to 10TB.

HPC-US-SECOND

Ticket	Days	Creation time	Requestor	Subject	Notes
15337	7	2021-01-27 15:38:01	qian.xia@ukaea.uk	access issue (scp, rsync)	The user has reported an issue with rsync command, and a second issue with a failed job on Marconi SKL (because of OPA instabilities).
15015	6	2021-01-13 14:22:01	hm1234@york.ac.uk	VSCode remote - SSH	The user has asked about any known issue on Marconi that automatically shuts off VScode servers. We have informed him about our policy, a process that runs on the login nodes of our cluster will be automatically closed after 10 minutes of activity.

Ticket statistics Jan 11 – Feb 14

HPC-US-FIRST

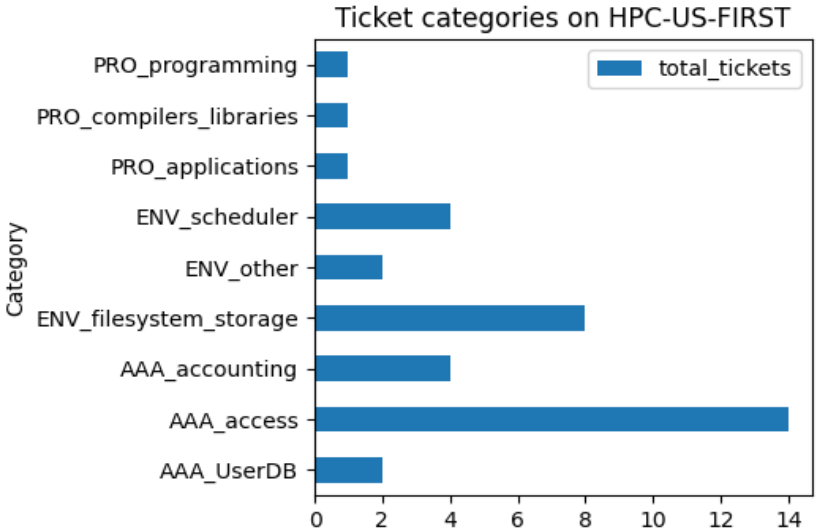

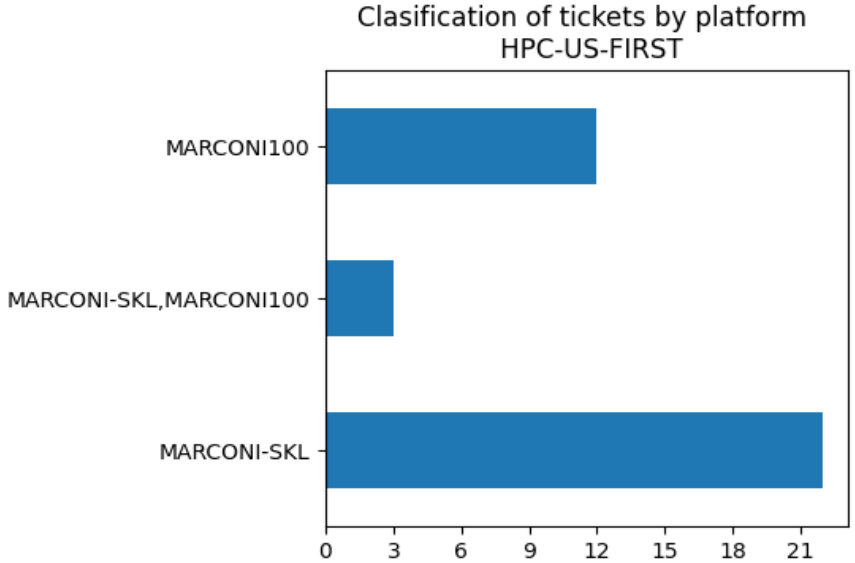

Ticket statistics Jan 11 – Feb 14

HPC-US-FIRST

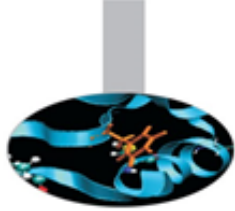

HPC-US-FIRST		By status		By platform		
		open	resolved	SKL	SKL&M100	M100
Information	AAA_accounting		1	1		
	AAA_UserDB		1		1	
	ENV_other		1	1		
	ENV_scheduler		1			1
	PRO_programming		1	1		
Problem	AAA_access		3	3		
	AAA_accounting		1	1		
	ENV_filesystem_storage		3	3		
	ENV_other	1				1
	ENV_scheduler		1			1
	PRO_applications		1	1		
	PRO_compilers_libraries		1	1		
Service Request	AAA_access		11	1		
	AAA_UserDB		1	6	1	4
	AAA_accounting		2	1		1
	ENV_filesystem_storage		5	1		4
	ENV_scheduler		2	1	1	
total	**37**	**1**	**36**	**22**	**3**	**12**

Ticket statistics Jan 11 – Feb 14

HPC-US-SECOND

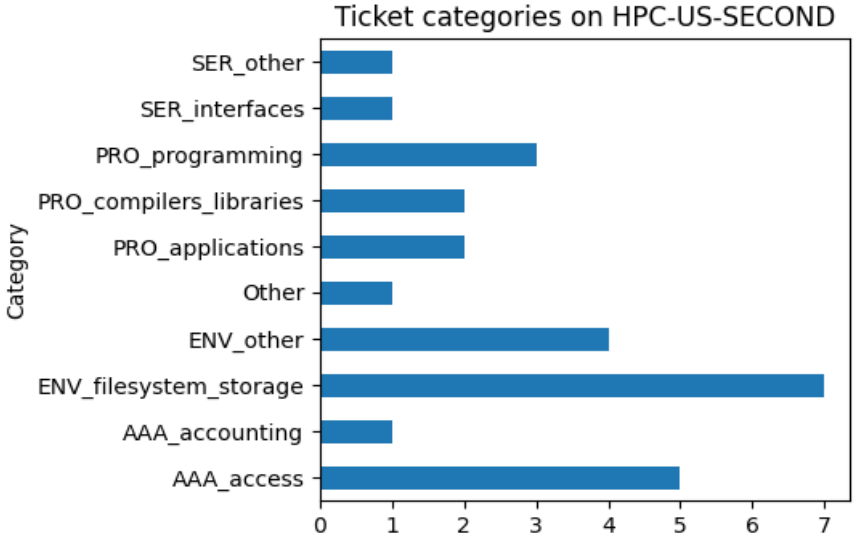

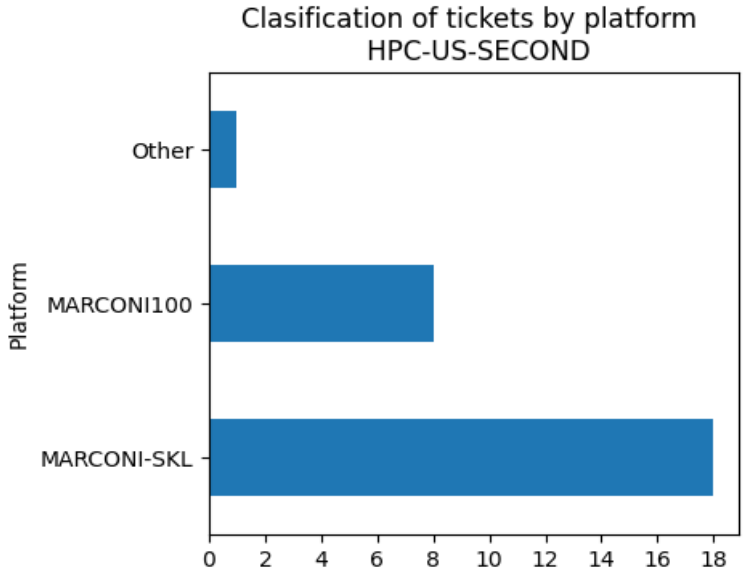

Ticket statistics Jan 11 – Feb 14

HPC-US-SECOND

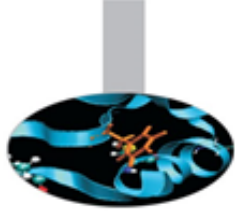

HPC-US-SECOND		By status		By platform		
		open	resolved	SKL	M100	Other
Information	Other		1			1
Problem	AAA_access		1	1		
	ENV_filesystem_storage	4	3	7		
	ENV_other		4	2	2	
	PRO_applications		2	1	1	
	PRO_programming	3		2	1	
	SER_interfaces		1	1		
	SER_other	1		1		
Service Request	AAA_access		4	1	3	
	AAA_accounting		1	1		
	PRO_compiler_libraries	1	1	1	3	
total	**27**	**9**	**18**	**18**	**8**	**1**

VASP usage on Marconi-SKL

March 2020 – January 2021

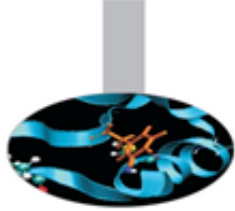

project	VASP module version	job count
fua34_befusion	6.1.0	1
fua34_cucsmo_3	5.4.4	1054
	6.1.0	17
fua34_dymarc3	5.3.5	69
	5.4.4	3899
	6.1.0	149
fua34_fecrcn	5.4.4	520
fua34_hccel	5.4.4	1801
fua34_indiloit	5.4.4	2551
fua34_marip2fm	5.4.4	3
fua34_tempis	6.1.0	20
fua34_vamedif	5.4.4	1734
fusio_ru4isspu	5.4.4	57

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

Batch jobs Mar 1 – Jan 31

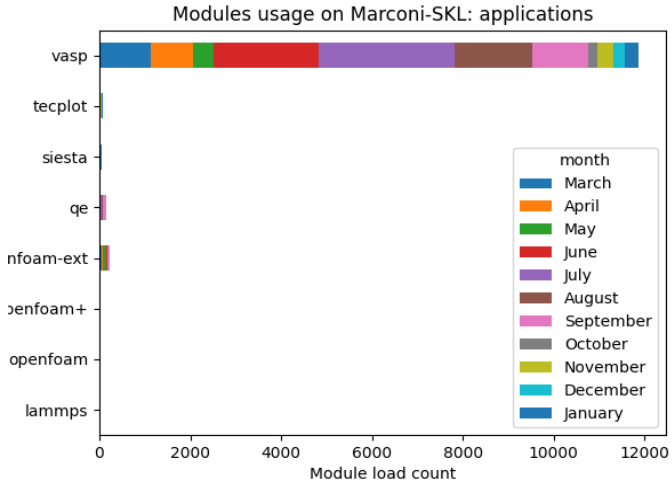

Module name	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan21
lammps	7							16			
openfoam	16	1		3		1	13				
openfoam+							2				
openfoam-ext	46	28	80	25	8		40				
qe		4		33	20	29	59				
siesta	41	3	13				3				
tecplot					3	3	10	12	22	2	23
vasp	1129	934	454	2301	2999	1722	1234	177	353	262	312

Batch jobs Jan 2021

category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
application	qe	skl_fua_prod	normal	104	46,12969
	tecplot	skl_fua_prod	normal	23	6,201921
	vasp	skl_fua_prod	normal	292	204,1965

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

Batch jobs Mar 1 – Jan 31

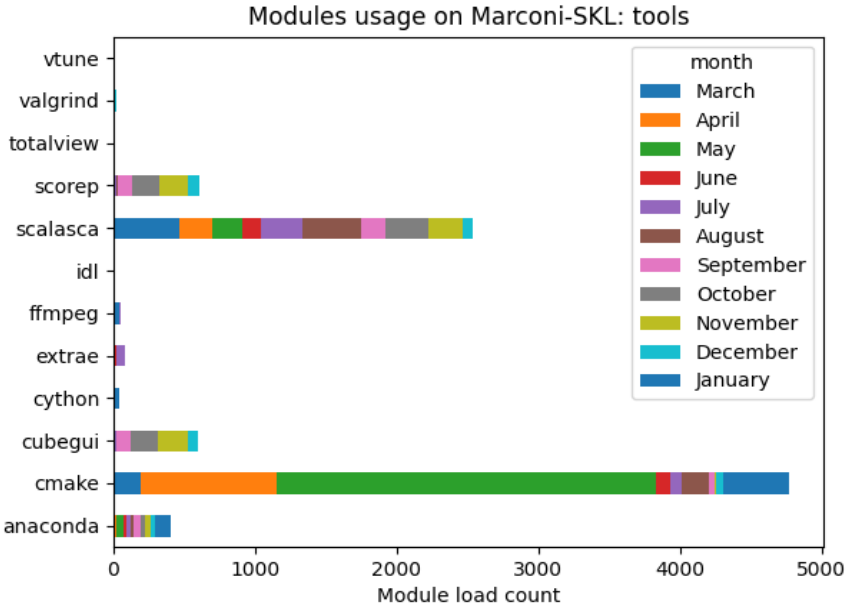

Module name	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan21
anaconda		22	50	18	29	19	51	37	38	27	115
cmake	197	951	2683	97	89	188	41	3	5	49	472
cubegui					16	4	100	194	208	77	
cython	43										
extrae				17	64						
ffmpeg	44				2			6			
idl									1		
scalasca	461	236	217	123	293	418	177	300	239	77	
scorep					25	4	100	193	208	76	
totalview			2	4							
valgrind	2	2								13	
vtune									1		

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

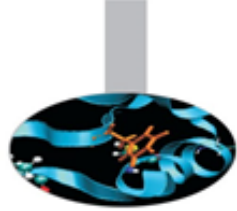

Batch jobs Jan 2020

category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
tool	anaconda	bdw_all_serial	normal	82	64,92458
		skl_fua_dbg	normal	15	14,33694
		skl_fua_prod	normal	29	20,51604
	cmake	bdw_all_serial	normal	5	5,000625
		skl_fua_dbg	normal	251	73,84048
		skl_fua_prod	normal	222	14,55631
	ffmpeg	skl_fua_prod	normal	39	5,762211
	idl	bdw_all_serial	normal	5	5,000625

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

Batch jobs Mar 1 – Dec 31

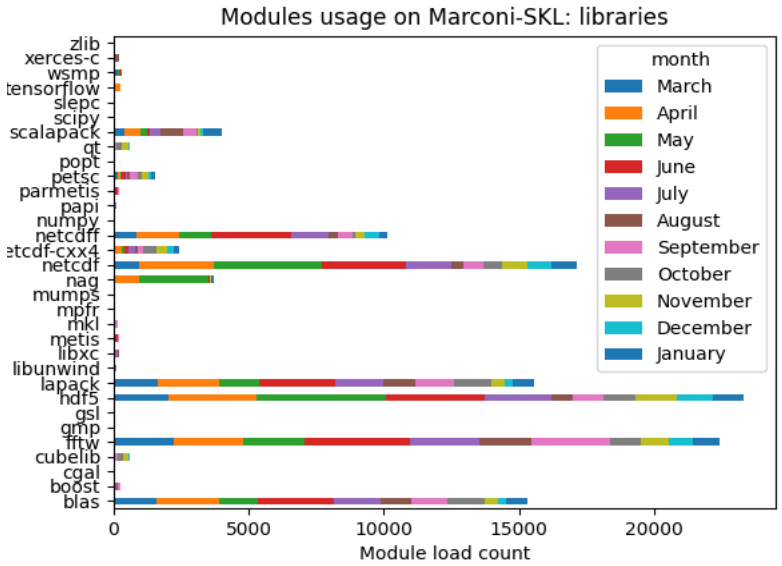

Module name	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan21
blas	1577	2308	1444	2838	1691	1151	1358	1388	494	285	770
boost	16	1		3	101	35	96				
cgal							2				
cubelib					25	4	99	196	207	77	
fftw	2229	2569	2270	3883	2592	1934	2892	1131	1047	908	950
gmp							2				
gsl					10						
hdf5	2032	3236	4815	3652	2449	821	1127	1177	1516	1365	1144
lapack	1620	2304	1458	2838	1777	1177	1444	1370	495	287	769
libunwind				17	64						
libxc					137	47		7			
metis		12	58	82			35			3	
mkl	7						140			6	16
mpfr							2				
mumps		12	14				35				
nag	2	960	2558	23	13		24	1	2		126
netcdf	946	2766	3985	3142	1679	410	753	691	932	893	934
netcdf-cxx4	58	215	127	150	260	69	189	513	391	271	178
netcdff	862	1561	1198	2949	1406	341	503	136	317	580	278
numpy	41	3	14								
papi				17	66						
parmetis		12	52	81			36			3	
petsc	162	83	74	104	58	109	314	126	241	122	161
popt											44
qt					16	4	101	196	206	77	
scalapack	379	596	252	93	415	819	522	44	105	107	648
scipy	2							9		21	31
slepc						1		9			11
tensorflow		264					1		2		
wsmp	128	7	59	86	11						
xerces-c	49	4	5	3		152					
zlib							2				

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

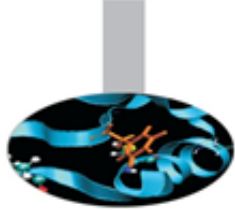

Batch jobs Jan 2020

category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	blas	bdw_all_serial	normal	5	5,000625
		skl_fua_dbg	normal	317	101,026
		skl_fua_prod	normal	2185	651,8881
			skl_qos_fuabprod	11	2,84537
	boost	skl_fua_prod	normal	56	24,24403
	fftw	bdw_all_serial	normal	5	5,000625
		skl_fua_dbg	normal	357	109,1242
		skl_fua_prod	normal	2893	896,1812
			skl_qos_fuabprod	6	2,495368
			skl_qos_fualprod	17	5,961597
	hdf5	bdw_all_serial	normal	30	18,98326
		skl_fua_dbg	normal	386	90,8409
		skl_fua_prod	normal	2773	707,2593
			skl_qos_fuabprod	20	5,924245
			skl_qos_fualprod	124	57,04277
	lapack	bdw_all_serial	normal	5	5,000625
		skl_fua_dbg	normal	317	101,0343
		skl_fua_prod	normal	2110	600,3656
			skl_qos_fuabprod	11	2,84537
	mkl	skl_fua_prod	normal	197	103,5692
			skl_qos_fualprod	4	0,068611
	nag	skl_fua_dbg	normal	126	11,15111

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

Batch jobs Jan 2020

category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
library	netcdf	bdw_all_serial	normal	30	18,98326
		skl_fua_dbg	normal	336	83,10257
		skl_fua_prod	normal	2292	511,5682
			skl_qos_fuabprod	6	2,495368
			skl_qos_fualprod	113	51,35781
	netcdf-cxx4	skl_fua_dbg	normal	46	2,741667
		skl_fua_prod	normal	18	0,715
			skl_qos_fualprod	113	51,35781
	netcdff	bdw_all_serial	normal	30	18,98326
		skl_fua_dbg	normal	39	6,520417
		skl_fua_prod	normal	1832	460,3038
			skl_qos_fuabprod	6	2,495368
	petsc	skl_fua_prod	normal	145	98,86519
			skl_qos_fuabprod	13	4,016178
	popt	skl_fua_dbg	normal	36	2,055556
		skl_fua_prod	normal	8	4,581944
	scalapack	bdw_all_serial	normal	5	5,000625
		skl_fua_dbg	normal	253	74,85715
		skl_fua_prod	normal	845	299,7234
	scipy	skl_fua_dbg	normal	20	14,96611
		skl_fua_prod	normal	11	6,353796
	slepc	skl_fua_prod	normal	11	6,353796

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

Batch jobs Mar 1 – Dec 31

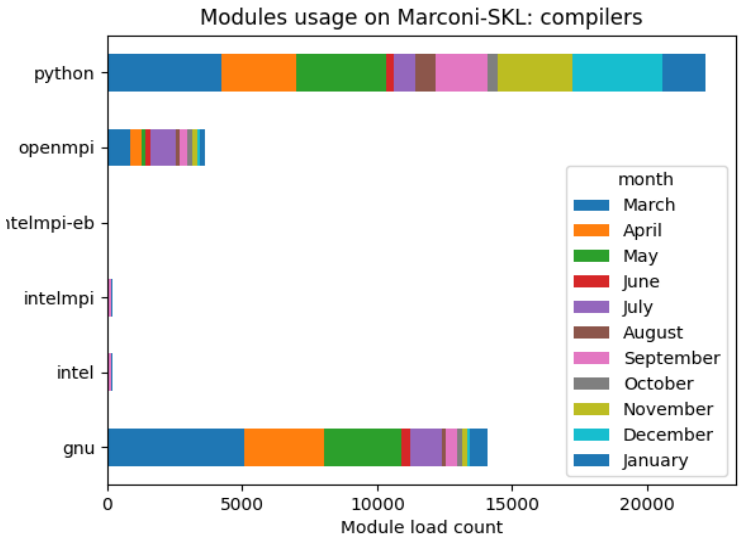

Module name	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan21
gnu	5062	2981	2871	318	1147	163	418	172	195	103	662
intel	8						140	5		6	25
intelmpi	8						140	5		6	25
intelmpi-eb			2								
openmpi	842	443	106	206	925	145	301	168	192	91	182
python	4230	2757	3334	300	787	739	1952	371	2787	3304	1626

Batch jobs Jan 2020

category	modulename	partition	qos	module_load_count	elapsed_timelimit_sum
compiler	gnu	skl_fua_dbg	normal	251	73,84048
		skl_fua_prod	normal	520	196,1527
	intel	skl_fua_dbg	normal	8	0,65
		skl_fua_prod	normal	199	104,5464
			skl_qos_fualprod	4	0,068611
	intelmpi	skl_fua_dbg	normal	8	0,65
		skl_fua_prod	normal	198	103,5702
			skl_qos_fualprod	4	0,068611
	openmpi	skl_fua_dbg	normal	1	0,148889
		skl_fua_prod	normal	415	261,5397
	python	bdw_all_serial	normal	85	67,16507
		skl_fua_dbg	normal	298	103,1437
		skl_fua_prod	normal	1292	230,9499
			skl_qos_fuabprod	6	2,495368