

CMS CSCS Status Report

CHIPP Face-to-Face Meeting
13 September 2019

Vinzenz Stampf





Overview




News


- ▶ April: Thomas out, Vinzenz in
- ▶ March/April: Decommissioning CEs (arc 01, 2 & 3), moving fully to HPC
 - ▶ merging T2_CH_CSCS_HPC (CE) and T2_CH_CSCS (SE) — End of April
 - ▶ except for not really: monitors are sometimes merged, mostly not
 - ▶ take following slides with a grain of salt
- ▶ Performance of the last six months: ok but could be improved
- ▶ Most important issues: Certificates (as always), full pools (user quota exceeded) & dCache (recently)



Sam Ranking




**Last F2F:
20Jun18 -
21Jan19**



**CSCS
89.71%
36/48**


**Now:
21Jan18 -
13Sep19**



**CSCS
87.68%
36/49**




Site Availability



**only merged (?)
CSCS
stats available**

Site Availability using CMS_CRITICAL_FULL

From 2019/01/21 to 2019/09/09



Ticket-ID	Creation Date	Subject	Site	Scope	Type	VO	Priority	Responsible Unit	Status	Last Update
142850	2019-08-27 08:32:00	All transfers are failing to CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	in progress	2019-09-03 09:23:00
141991	2019-07-01 09:13:00	SAM Test (CE) failing for T2_CH_CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-07-11 13:12:00
141753	2019-06-17 14:13:00	T2_CH_CSCS xrootd tests are failing	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-06-28 11:50:00
141383	2019-05-24 09:06:00	Transfers failing to T2_CH_CSCS - pools are full	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-06-29 08:51:00
140228	2019-03-15 18:18:00	Transfers failing from T2_CH_CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-04-13 23:04:00
139246	2019-01-21 09:02:00	All transfers failing to CSCS - expired proxy	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-02-01 08:32:00


Not really clear in this statistics if former arc 01/2/3 included!



Site Availability



Site Availability using CMS_CRITICAL_FULL



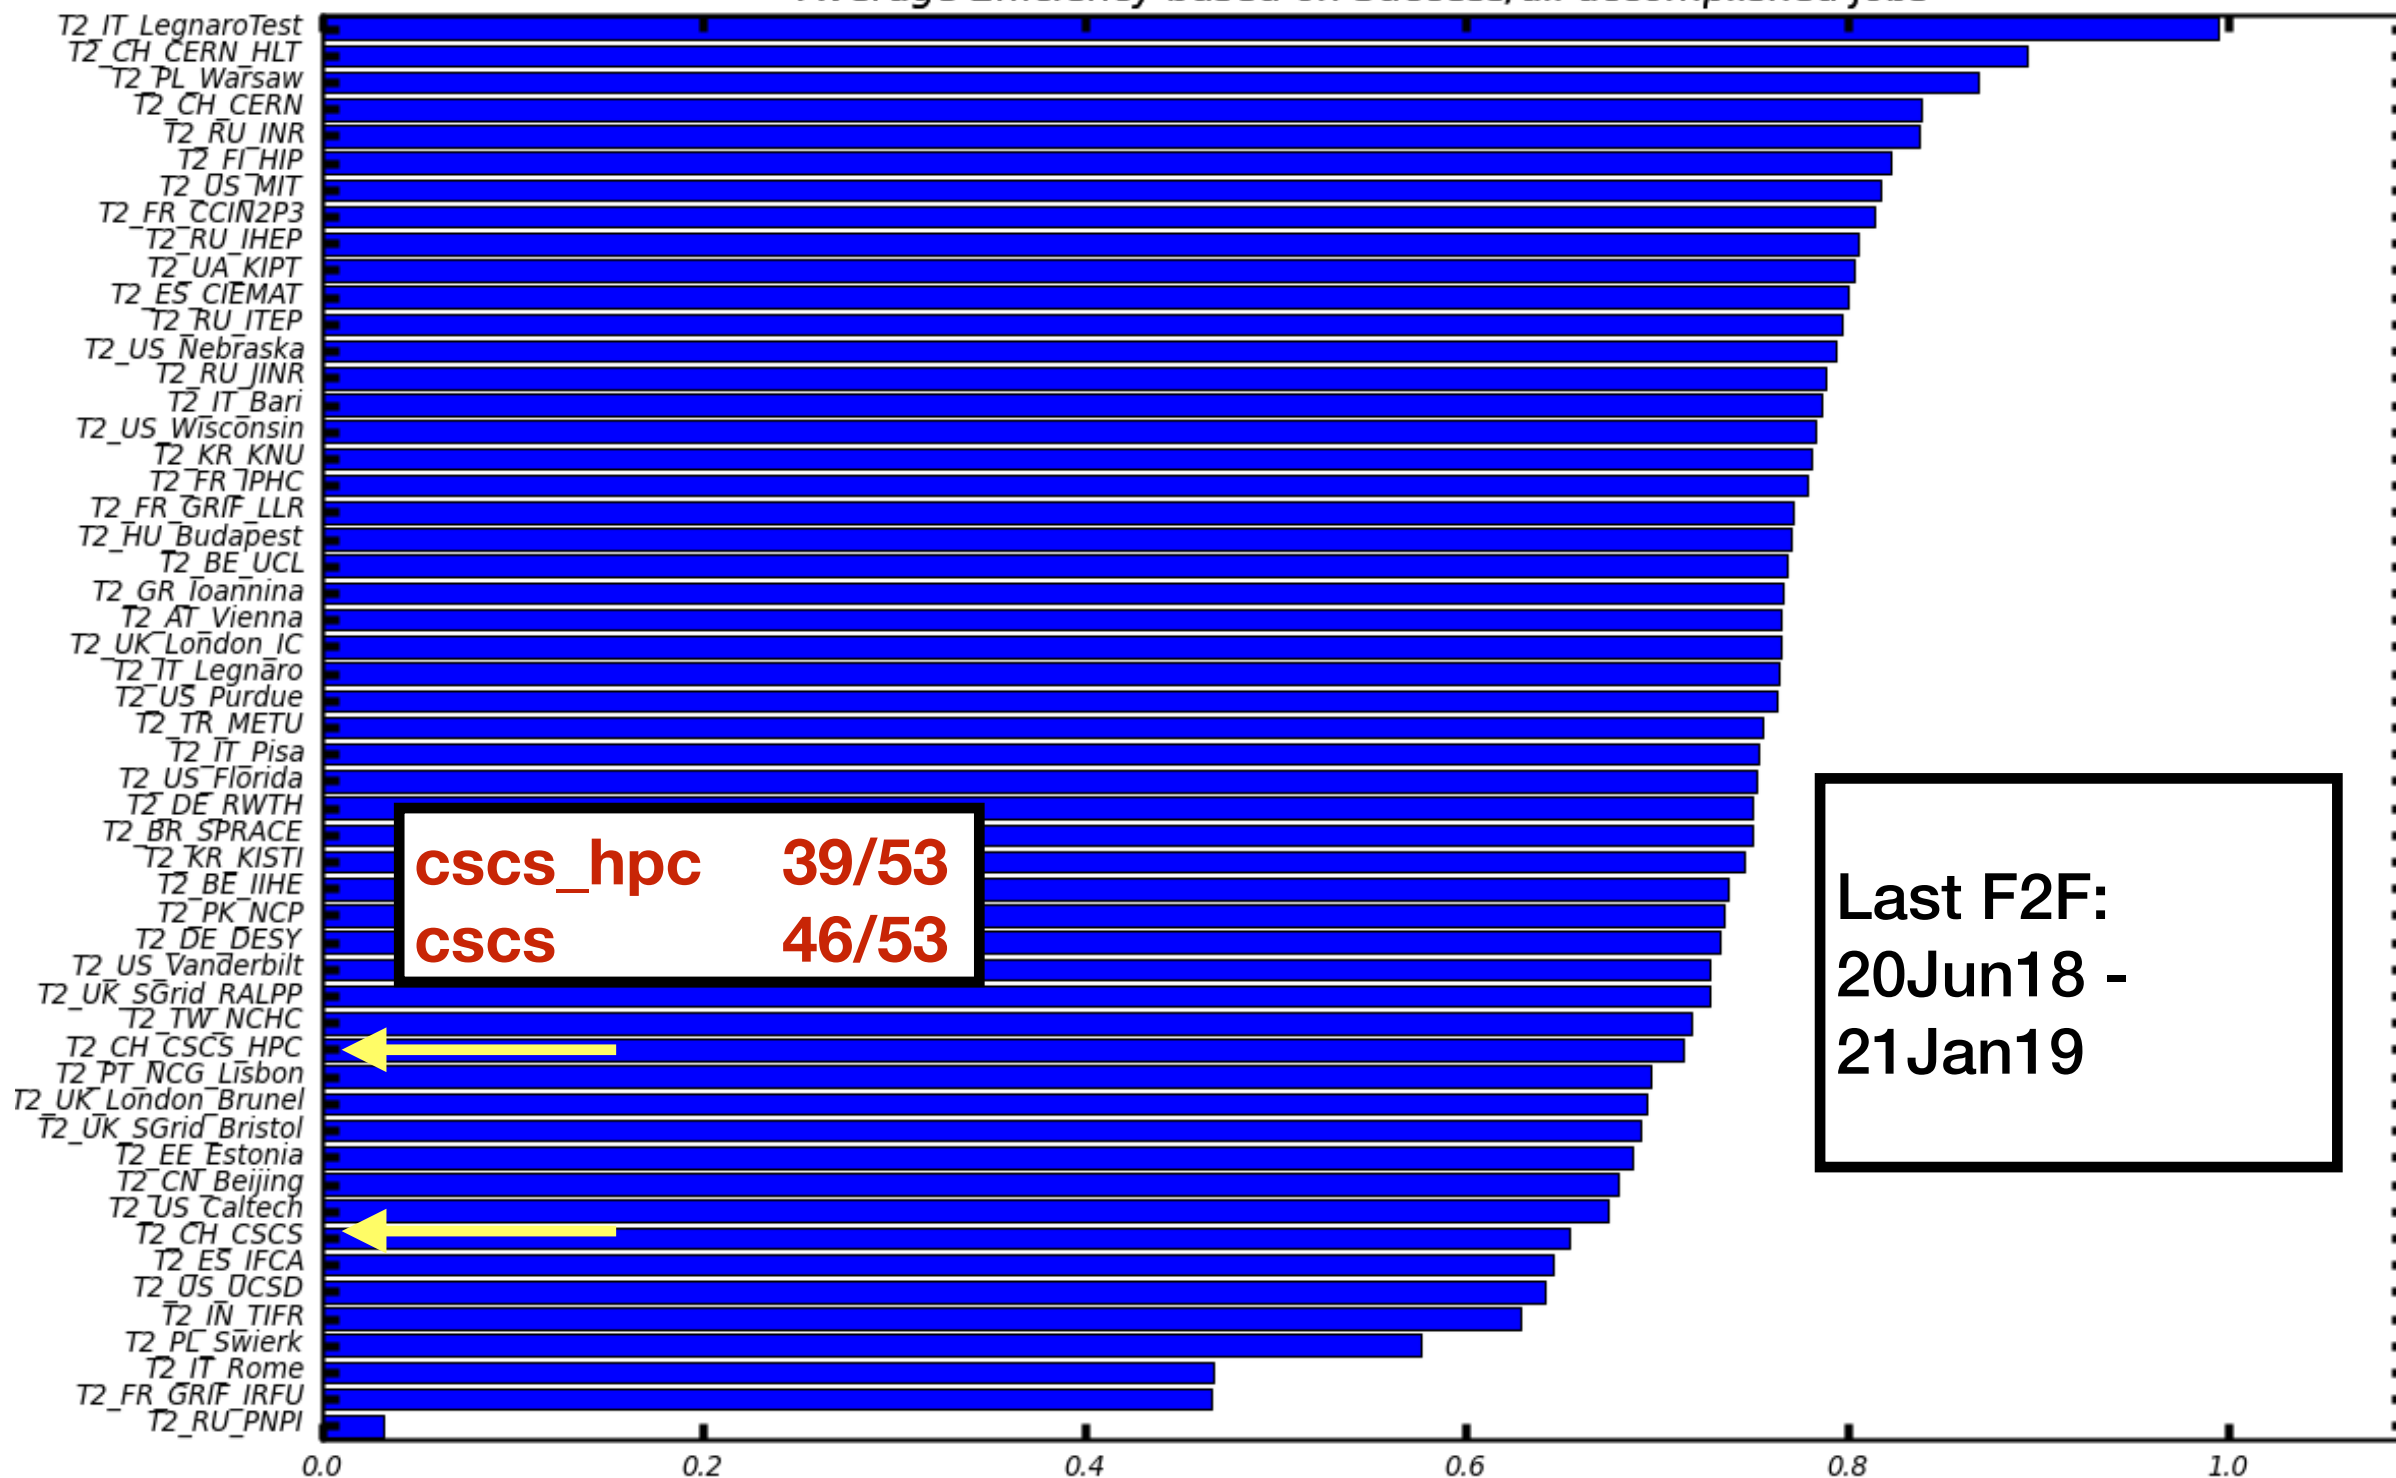
Ticket-ID	Creation Date	Subject	Site	Scope	Type	VO	Priority	Responsible Unit	Status	Last Update
142850	2019-09-27 08:32:00	All transfers are failing to CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	in progress	2019-09-03 09:23:00
41991	2019-07-01 09:13:00	SAM Test (CE) failing for T2_CH_CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-07-11 13:12:00
41752	2019-06-17 14:13:00	T2_CH_CSCS xrootd tests are failing	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-06-28 11:50:00
141383	2019-05-24 09:06:00	Transfers failing to T2_CH_CSCS - pools are full	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-06-29 08:51:00
140228	2019-03-15 18:18:00	Transfers failing from T2_CH_CSCS	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-04-13 23:04:00
139246	2019-01-21 09:02:00	All transfers failing to CSCS - expired proxy	CSCS-LCG2	WLCG	USER	cms	urgent	NGI_CH	closed	2019-02-01 08:32:00



Job Efficiency – last period



Average Efficiency based on Success/all accomplished jobs




cscs_hpc 39/53
cscs 46/53

Last F2F:
 20Jun18 -
 21Jan19

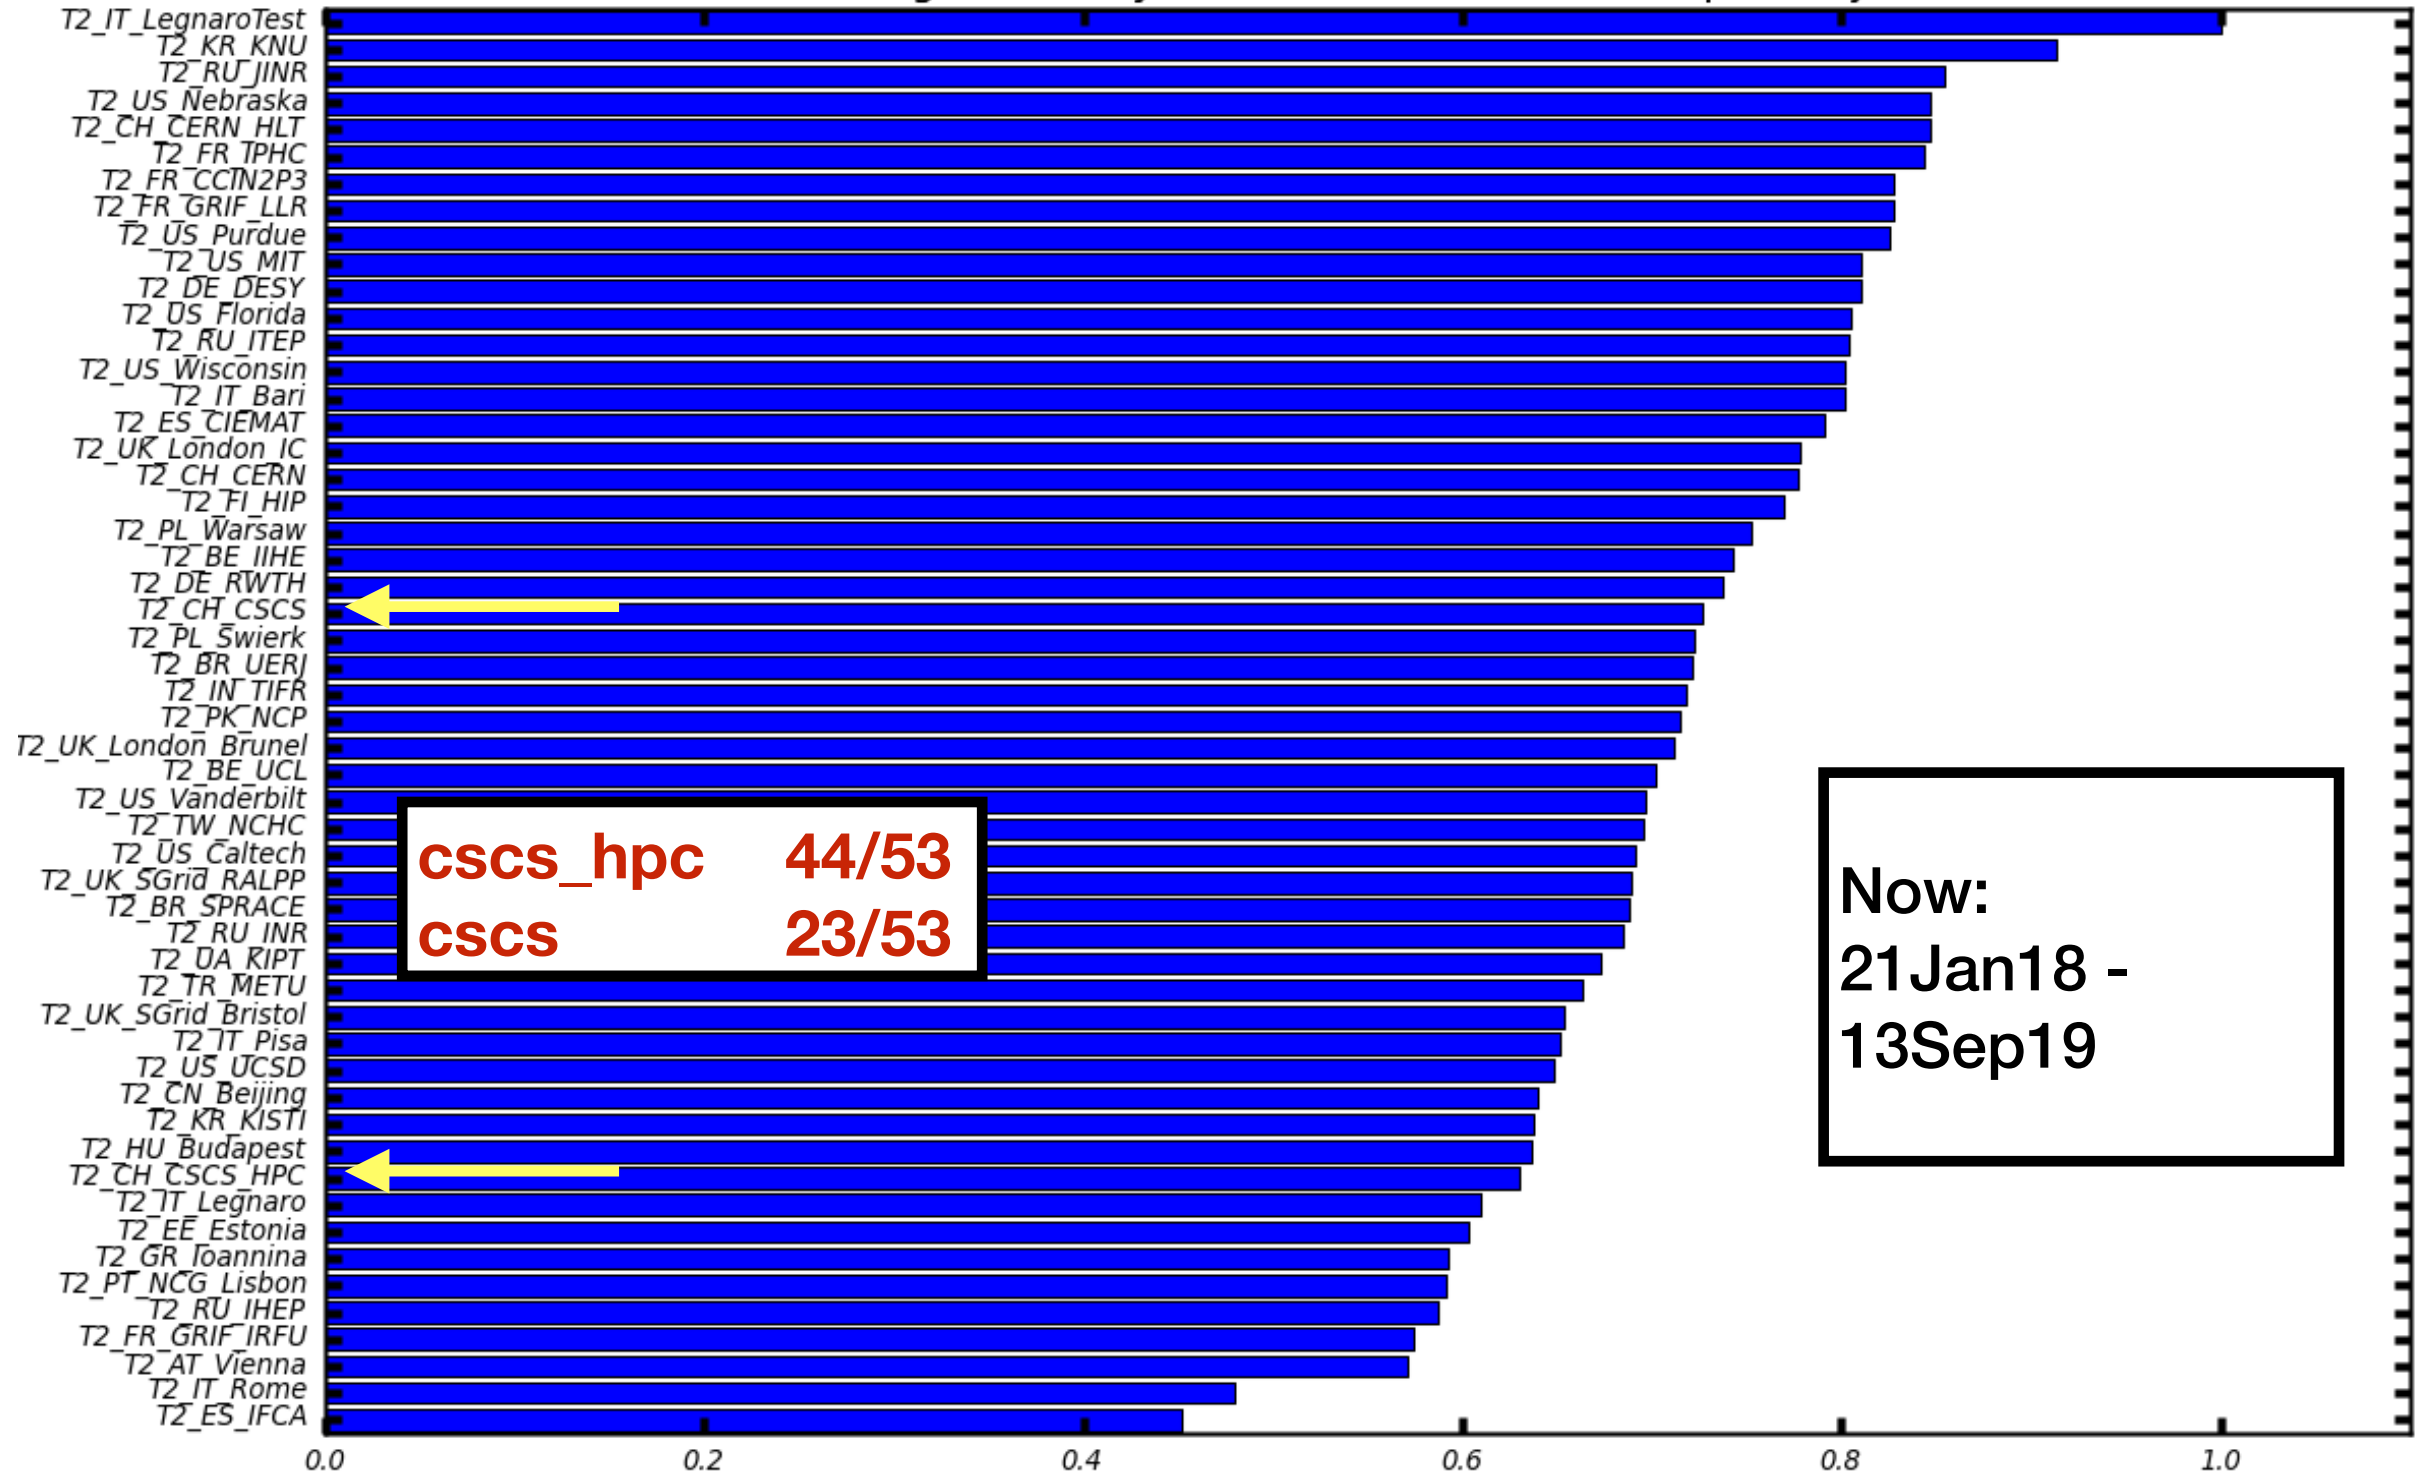


Job Efficiency – this period



dashboard

Average Efficiency based on Success/all accomplished jobs




cscs_hpc 44/53
cscs 23/53

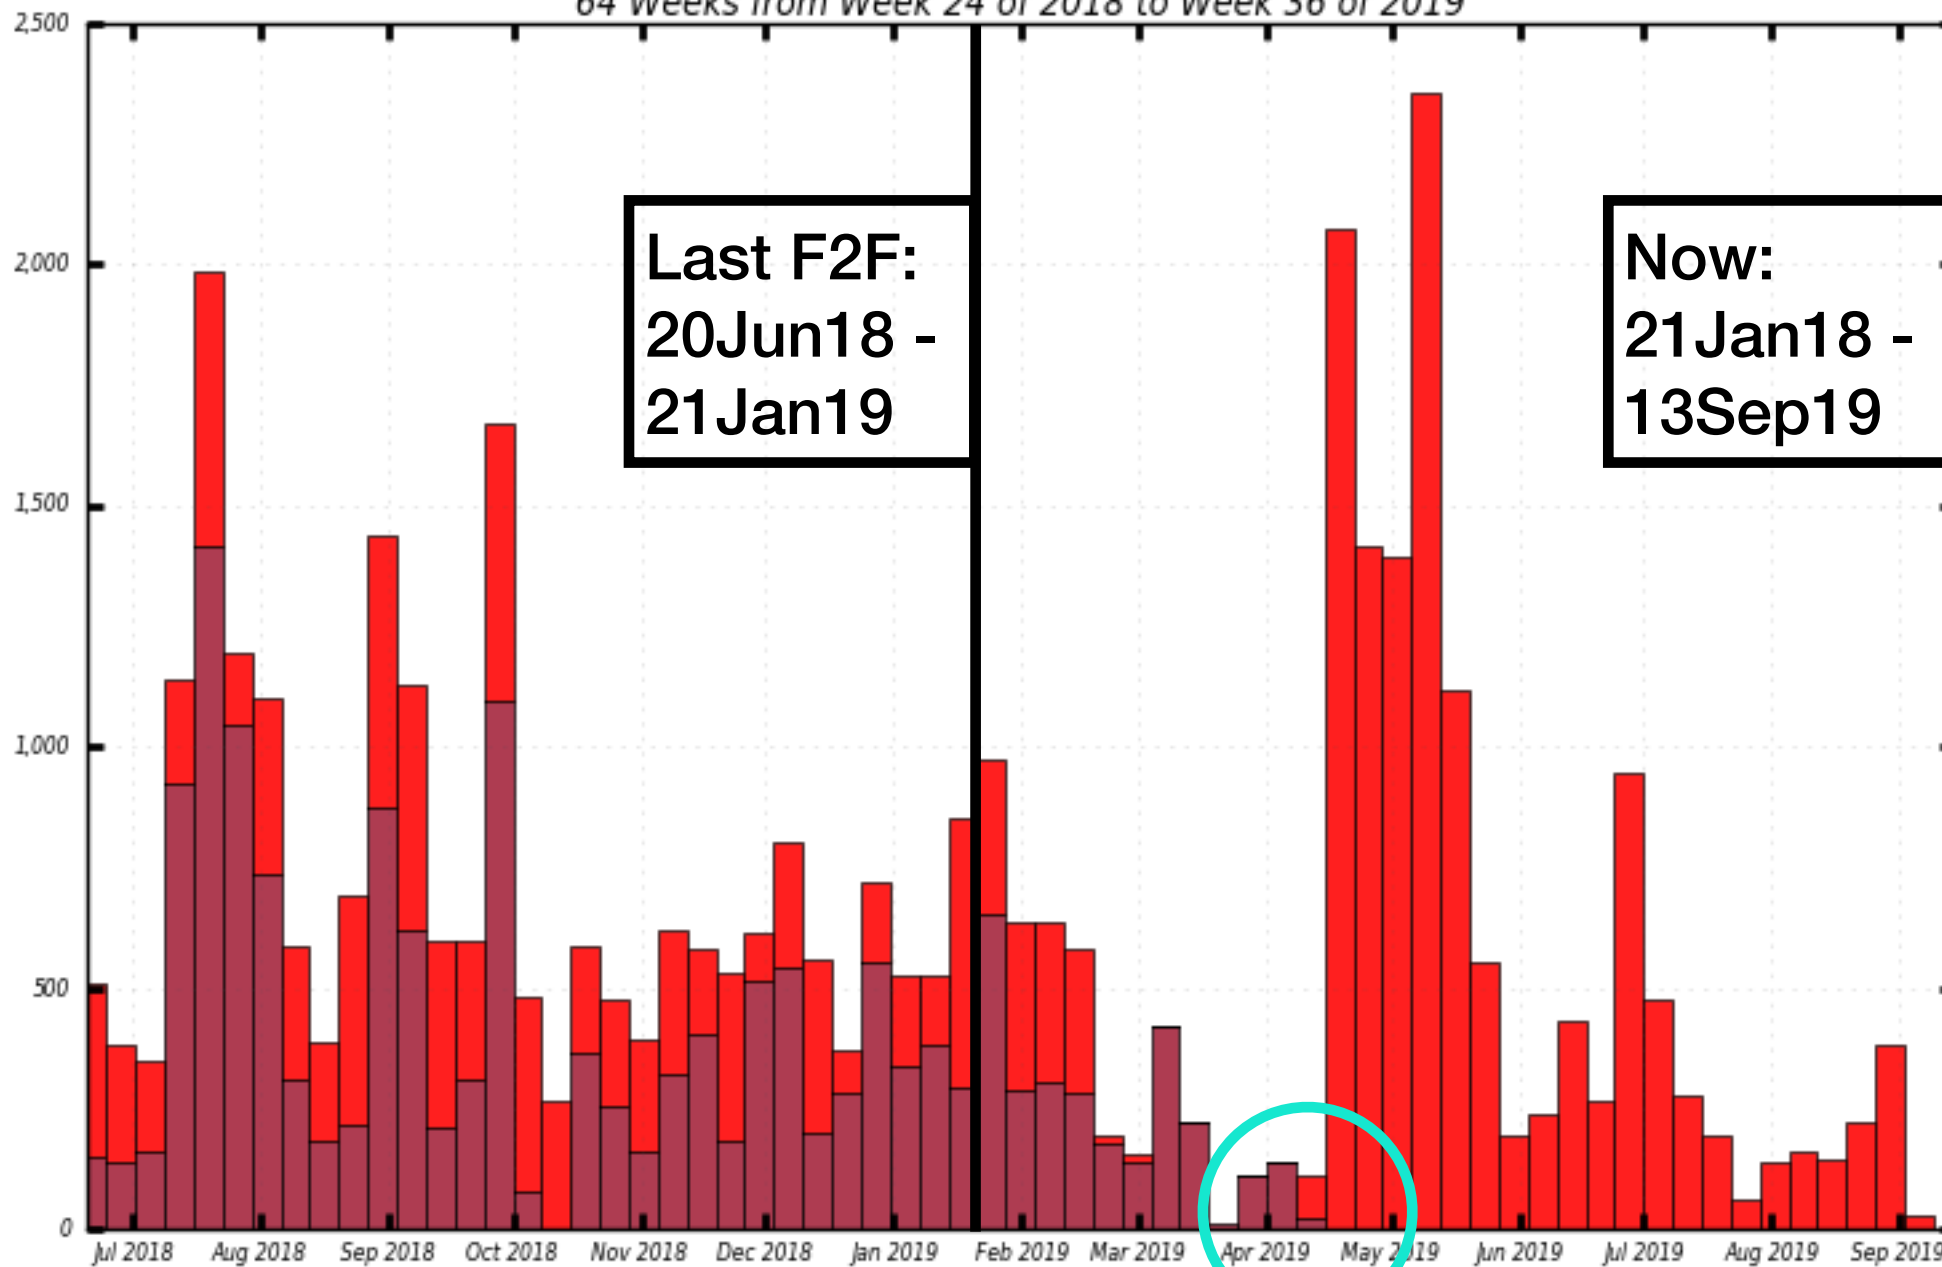
Now:
 21Jan18 -
 13Sep19



CPU Consumption



weeks/week: CPU consumption Good Jobs
64 Weeks from Week 24 of 2018 to Week 36 of 2019



Last F2F:
20Jun18 -
21Jan19

Now:
21Jan18 -
13Sep19

merging HPC & CSCS


T2_CH_CSCS

T2_CH_CSCS_HPC

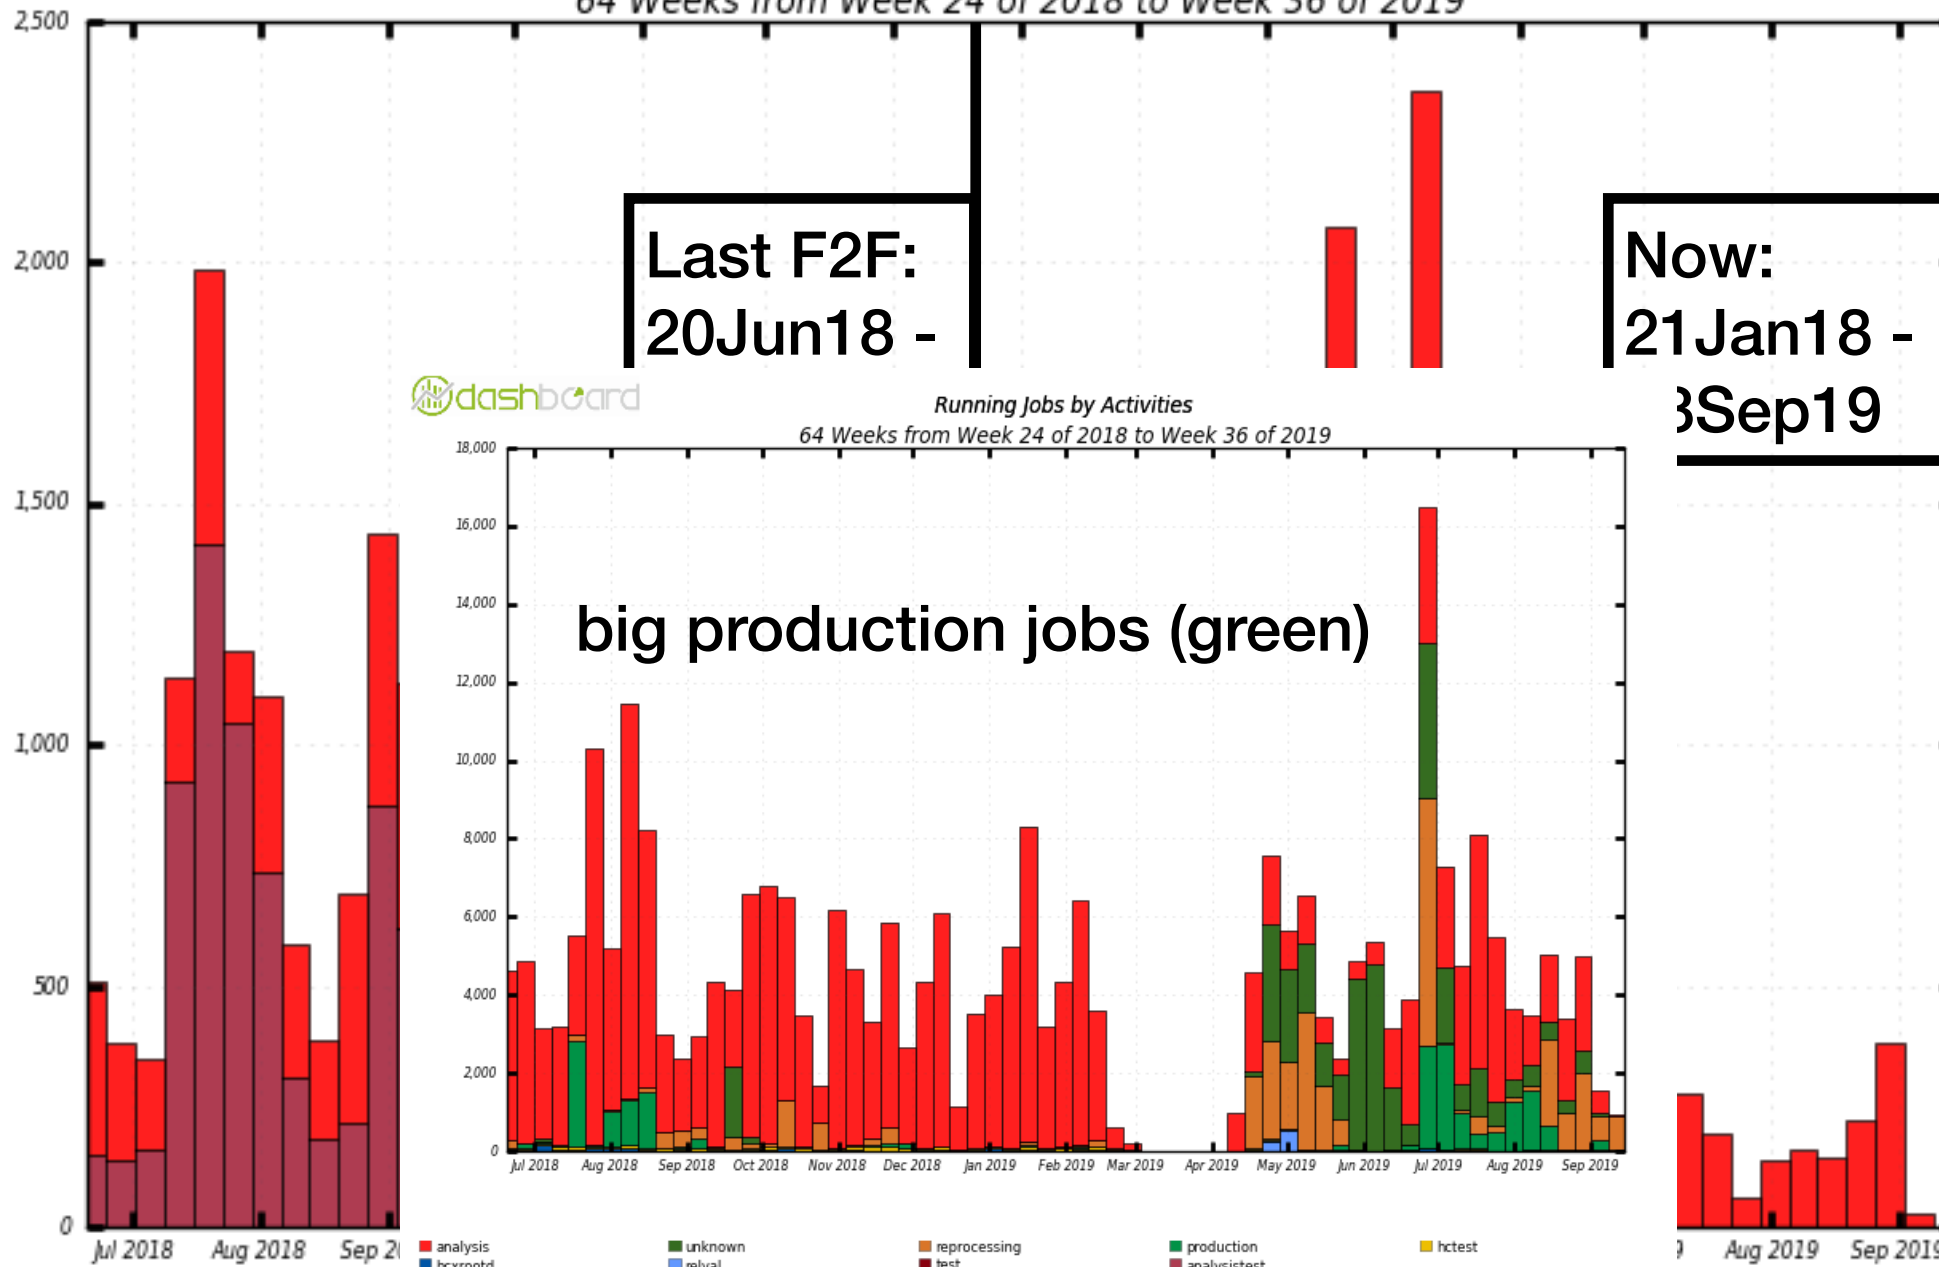
Maximum: 2,355 , Minimum: 0.00 , Average: 604.41 , Current: 0.00



CPU Consumption

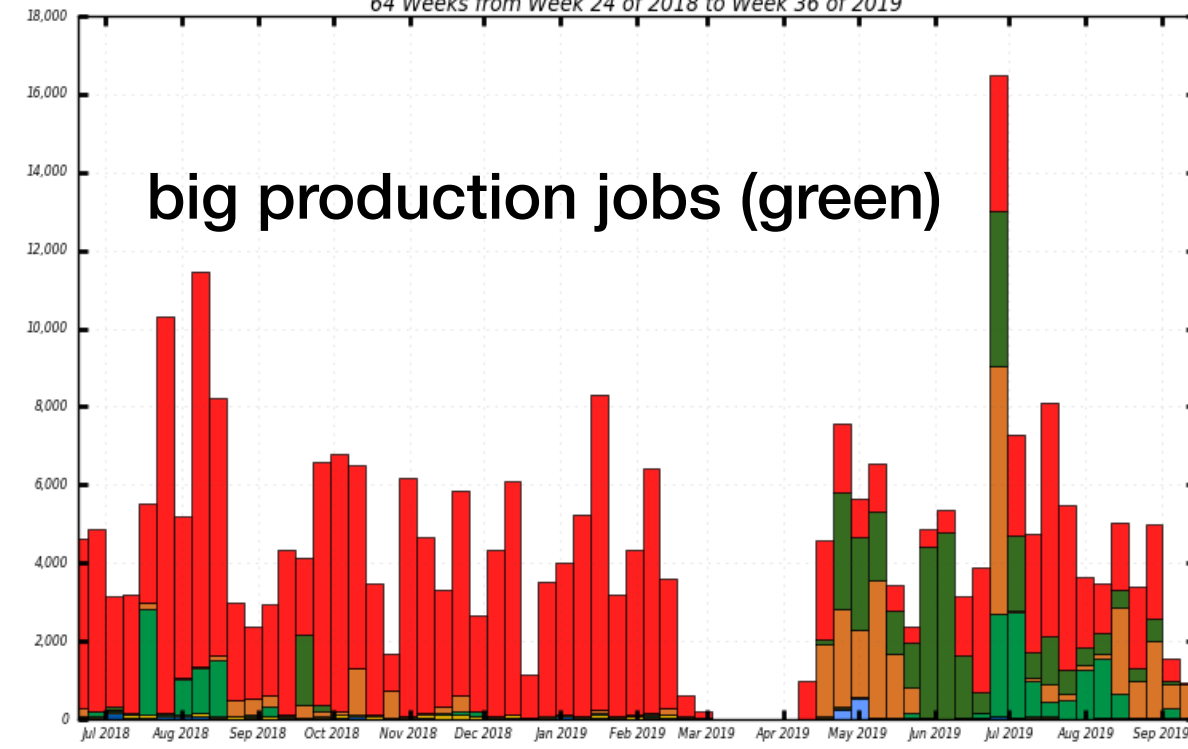


weeks/week: CPU consumption Good Jobs
64 Weeks from Week 24 of 2018 to Week 36 of 2019



Running Jobs by Activities
64 Weeks from Week 24 of 2018 to Week 36 of 2019

big production jobs (green)



Maximum: 16,501, Minimum: 0.00, Average: 4,394, Current: 922.00


T2_CH_CSCS

T2_CH_CSCS_HPC

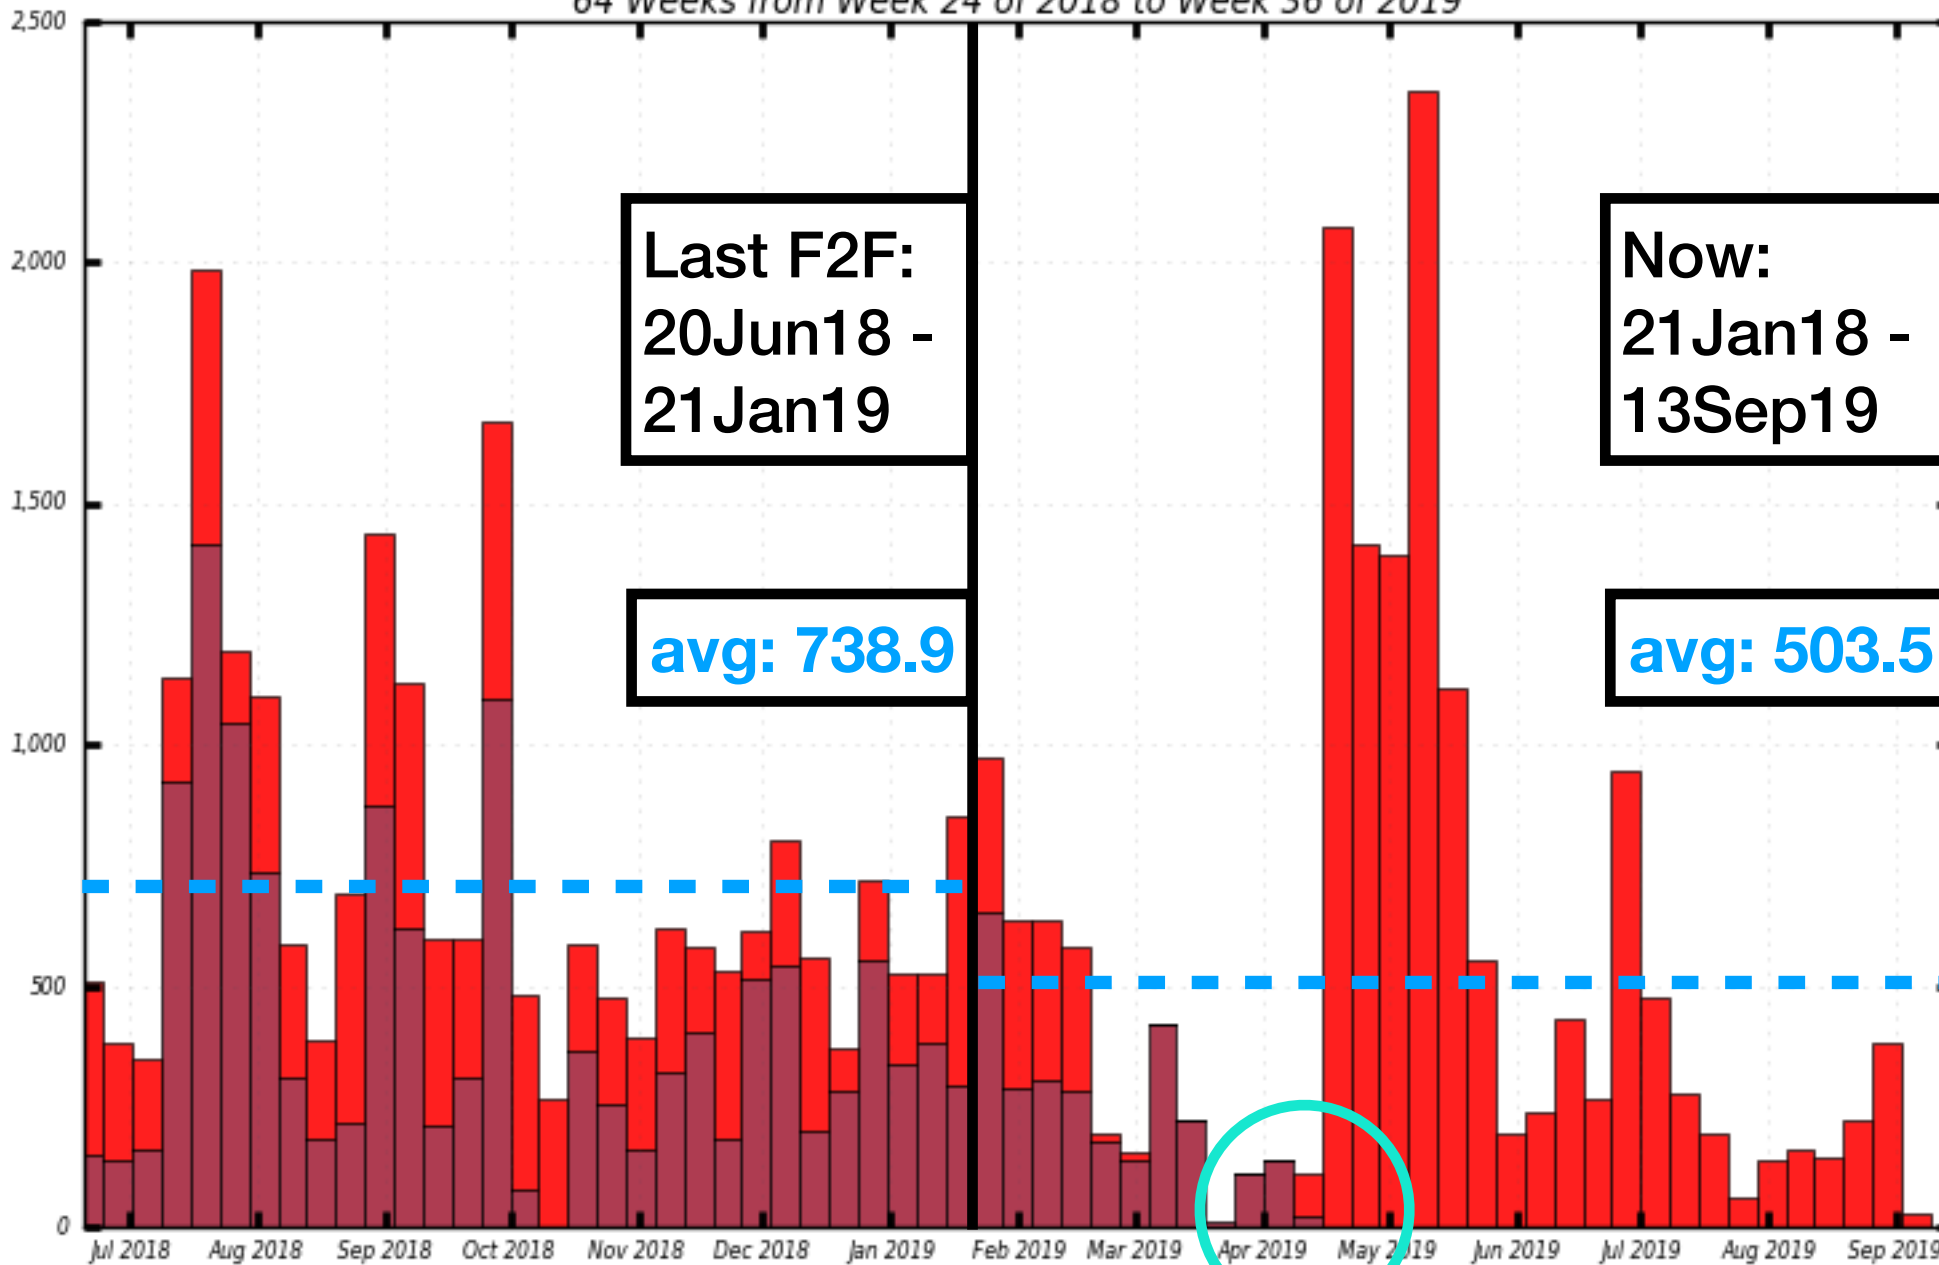
Maximum: 2,355, Minimum: 0.00, Average: 604.41, Current: 0.00



CPU Consumption



weeks/week: CPU consumption Good Jobs
64 Weeks from Week 24 of 2018 to Week 36 of 2019



T2_CH_CSCS


T2_CH_CSCS_HPC

Maximum: 2,355 , Minimum: 0.00 , Average: 604.41 , Current: 0.00

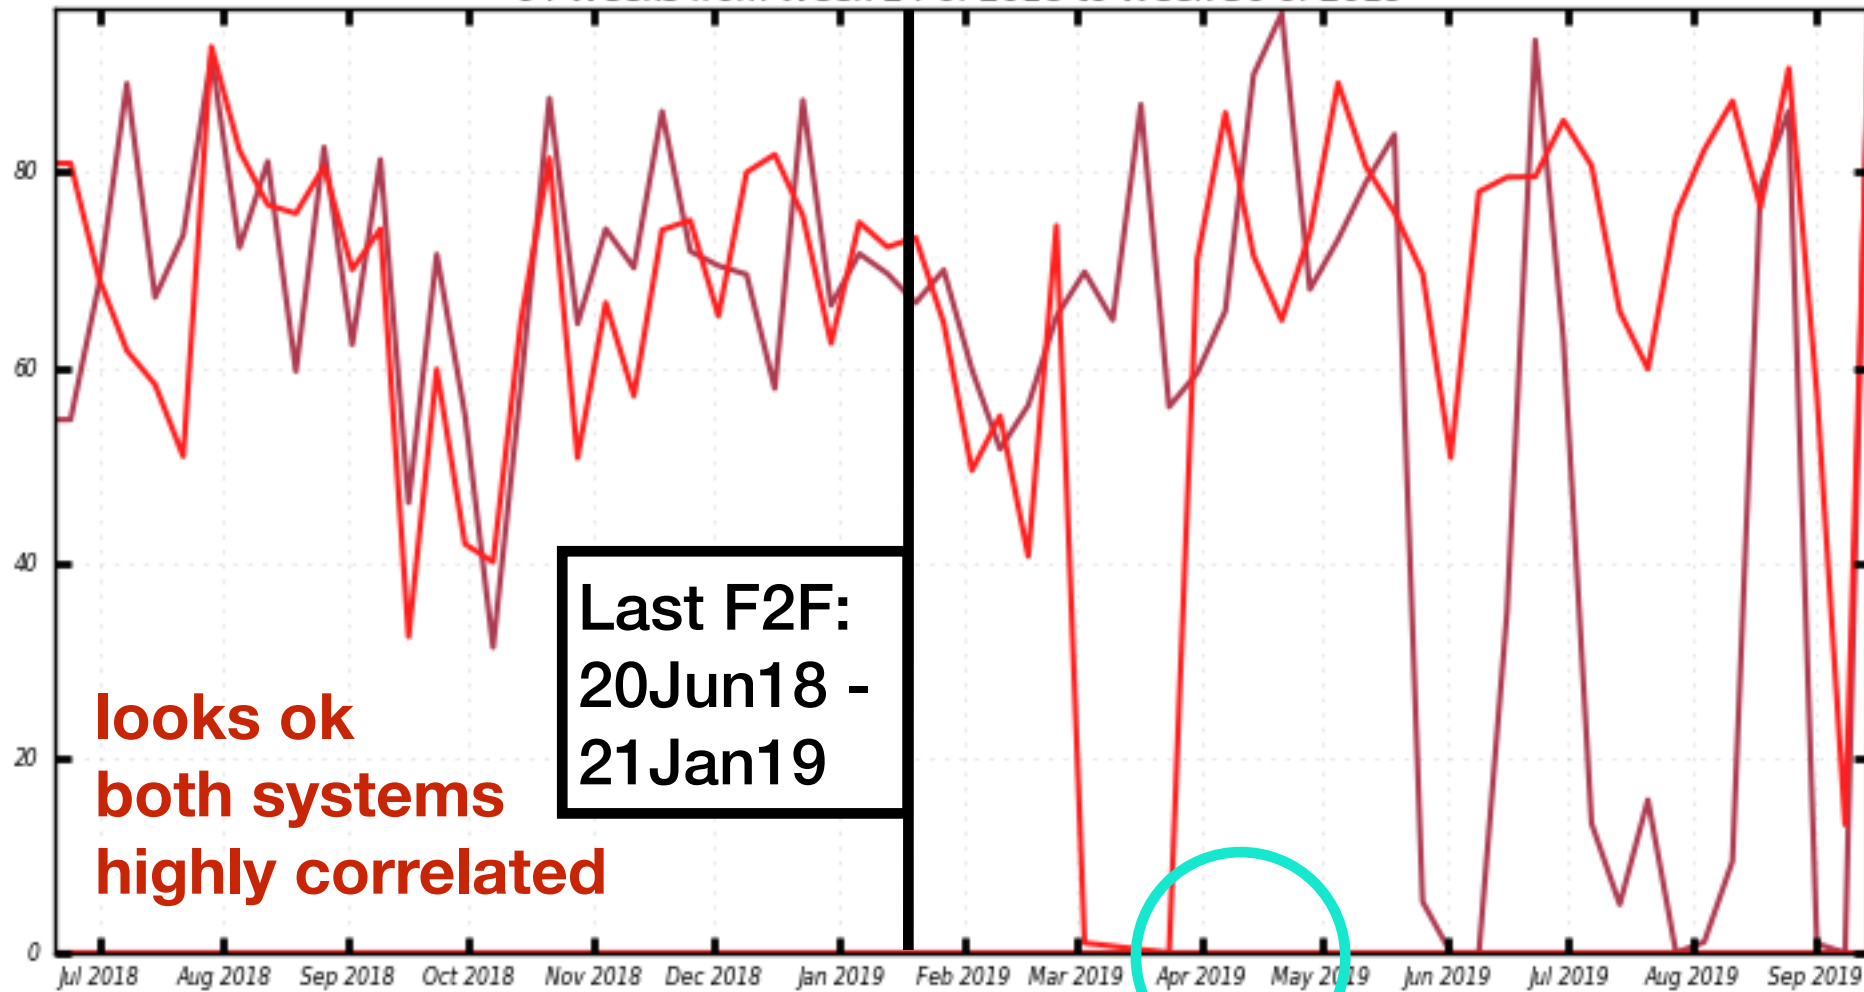
merging HPC & CSCS



Job Efficiency — over time



Efficiency over time based on success/all accomplished jobs
64 Weeks from Week 24 of 2018 to Week 36 of 2019



looks ok
both systems
highly correlated

Last F2F:
20 Jun 18 -
21 Jan 19

Now:
21 Jan 18 -
13 Sep 19

merging HPC & CSCS
why still two contributions
to this plot?!


T2_CH_CSCS (66.66)

T2_CH_CSCS_HPC (58.78)

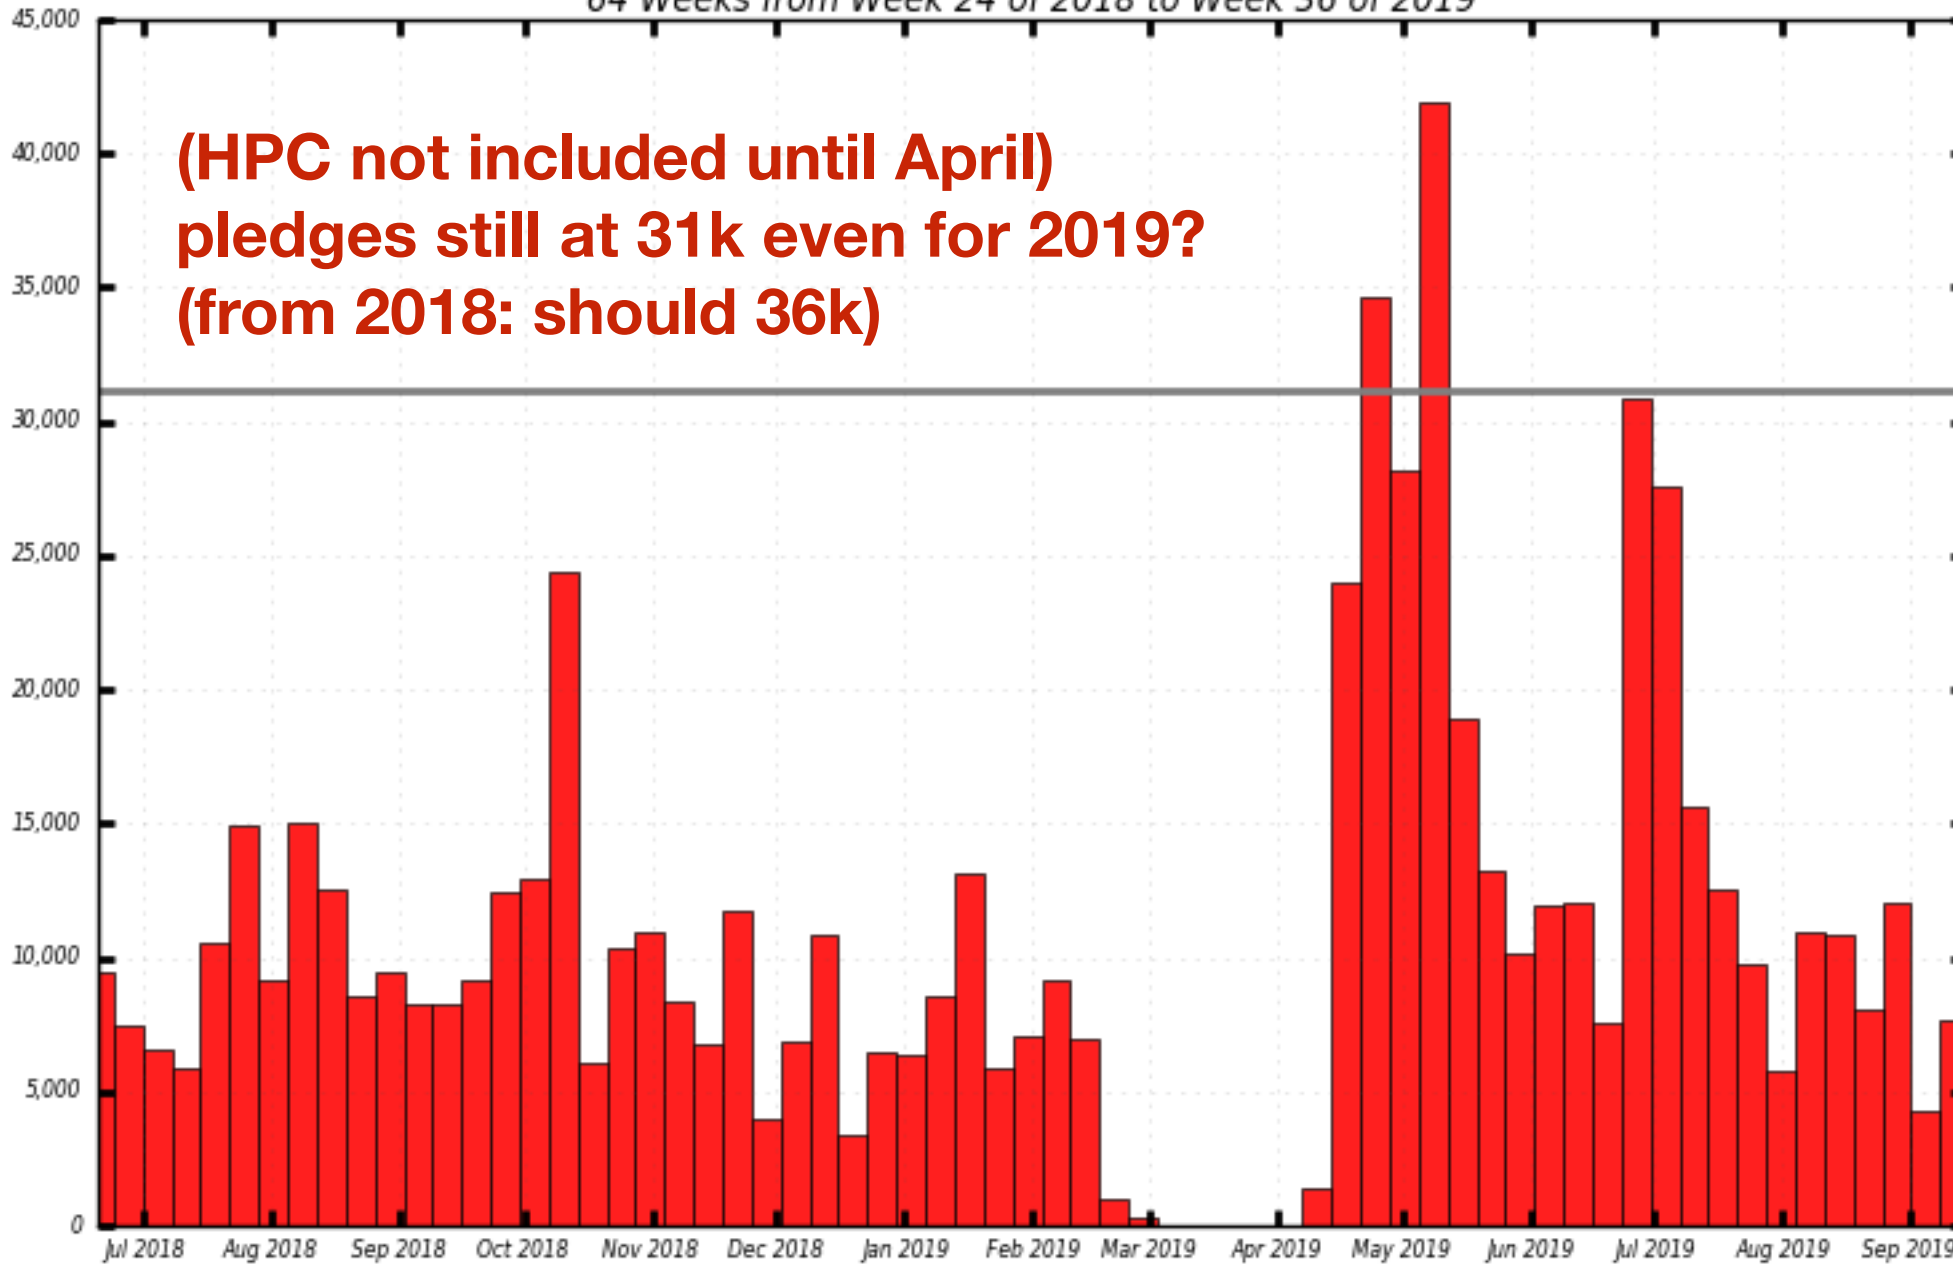
Total: 13.14 , Average Rate: 0.00 /s



Pledges



WallClock HEPSP06 Hours
64 Weeks from Week 24 of 2018 to Week 36 of 2019



T2_CH_CSCS

T2_CH_CSCS_HPC

Maximum: 41,958 , Minimum: 0.00 , Average: 10,449 , Current: 7,646

Back Up

CSCS status report from CMS

25-01-2018 - 20-06-2018

21 June 2018

Thomas Klijnsma
F2F meeting Zurich




ETH zürich

SAM Ranking: T2_CH_CSCS_HPC

Site Availability using CMS_CRITICAL_FULL

From 2018/01/25 to 2018/06/20




- Mostly the same reason for outage, except the long singularity problem for PizDaint

SAM Ranking: T2_CH_CSCS

Site Availability using CMS_CRITICAL_FULL


From 2018/01/25 to 2018/06/20



Ticket-ID	Creation	Subject	Comment
135130	15-05-18	Transfers failing to T2_CH_CSCS - All pools are full	First linkgroups, then broken dCache head node (~28 May)
134821	30-04-18	T2_CH_CSCS SAM SRM and stageout failures	Outdated certificates of Ses
134661	22-04-18	T2_CH_CSCS : SAM3 CE is critical	Required update of link groups
133578	19-02-18	Transfers failing from T2_CH_CSCS	Required update of link groups
133350	08-02-18	T2_CH_CSCS SAM JobSubmit failures	Outdated certificates of Ses

SAM Ranking: T2_CH_CSCS

Site Availability using CMS_CRITICAL_FULL




Ticket-ID	Creation	Subject	Comment
135130	15-05-18	Transfers failing to T2_CH_CSCS - All pools are full	First linkgroups, then broken dCache head node (~28 May)
134821	30-04-18	T2_CH_CSCS SAM SRM and stageout failures	Outdated certificates of Ses
134661	22-04-18	T2_CH_CSCS : SAM3 CE is critical	Required update of link groups
133578	19-02-18	Transfers failing from T2_CH_CSCS	Required update of link groups
133350	08-02-18	T2_CH_CSCS SAM JobSubmit failures	Outdated certificates of Ses

Average efficiency, all jobs

CSCS_HPC:
67.6% 14/53


CSCS: →
67.3% 15/53



Last F2F:
03-10-2017
to
24-01-2018

CSCS: →
66.2% 8/53

CSCS_HPC: →
61.6% 23/53




Now:
25-01-2018
to
20-06-2018

CPU Consumption Good jobs



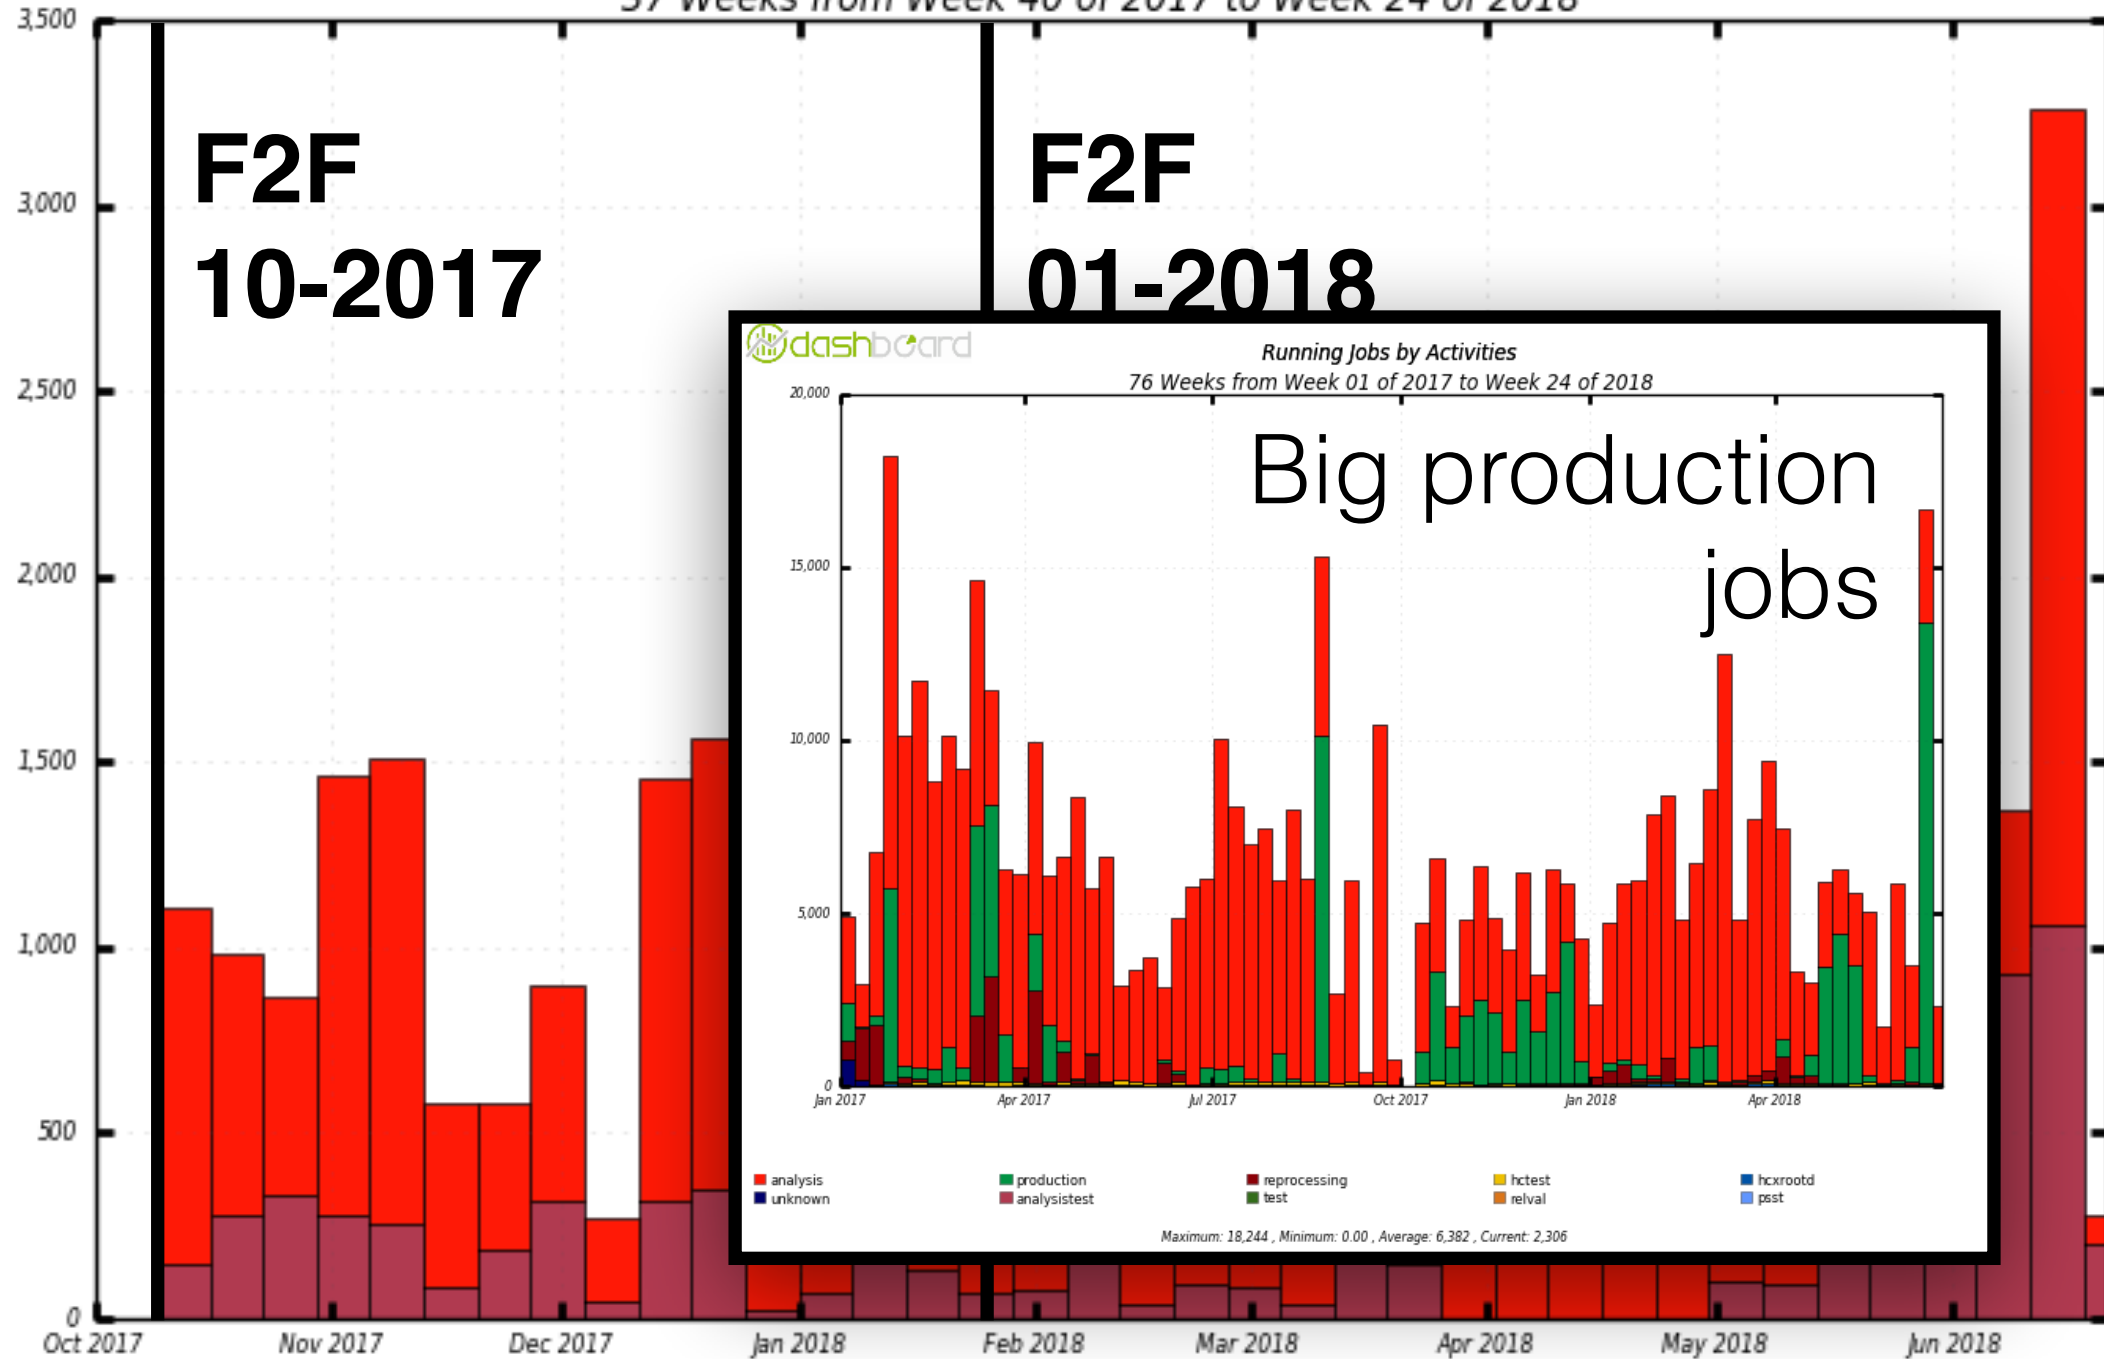
weeks/week: CPU consumption Good Jobs
37 Weeks from Week 40 of 2017 to Week 24 of 2018



CPU Consumption Good jobs



weeks/week: CPU consumption Good Jobs
37 Weeks from Week 40 of 2017 to Week 24 of 2018



T2_CH_CSCS

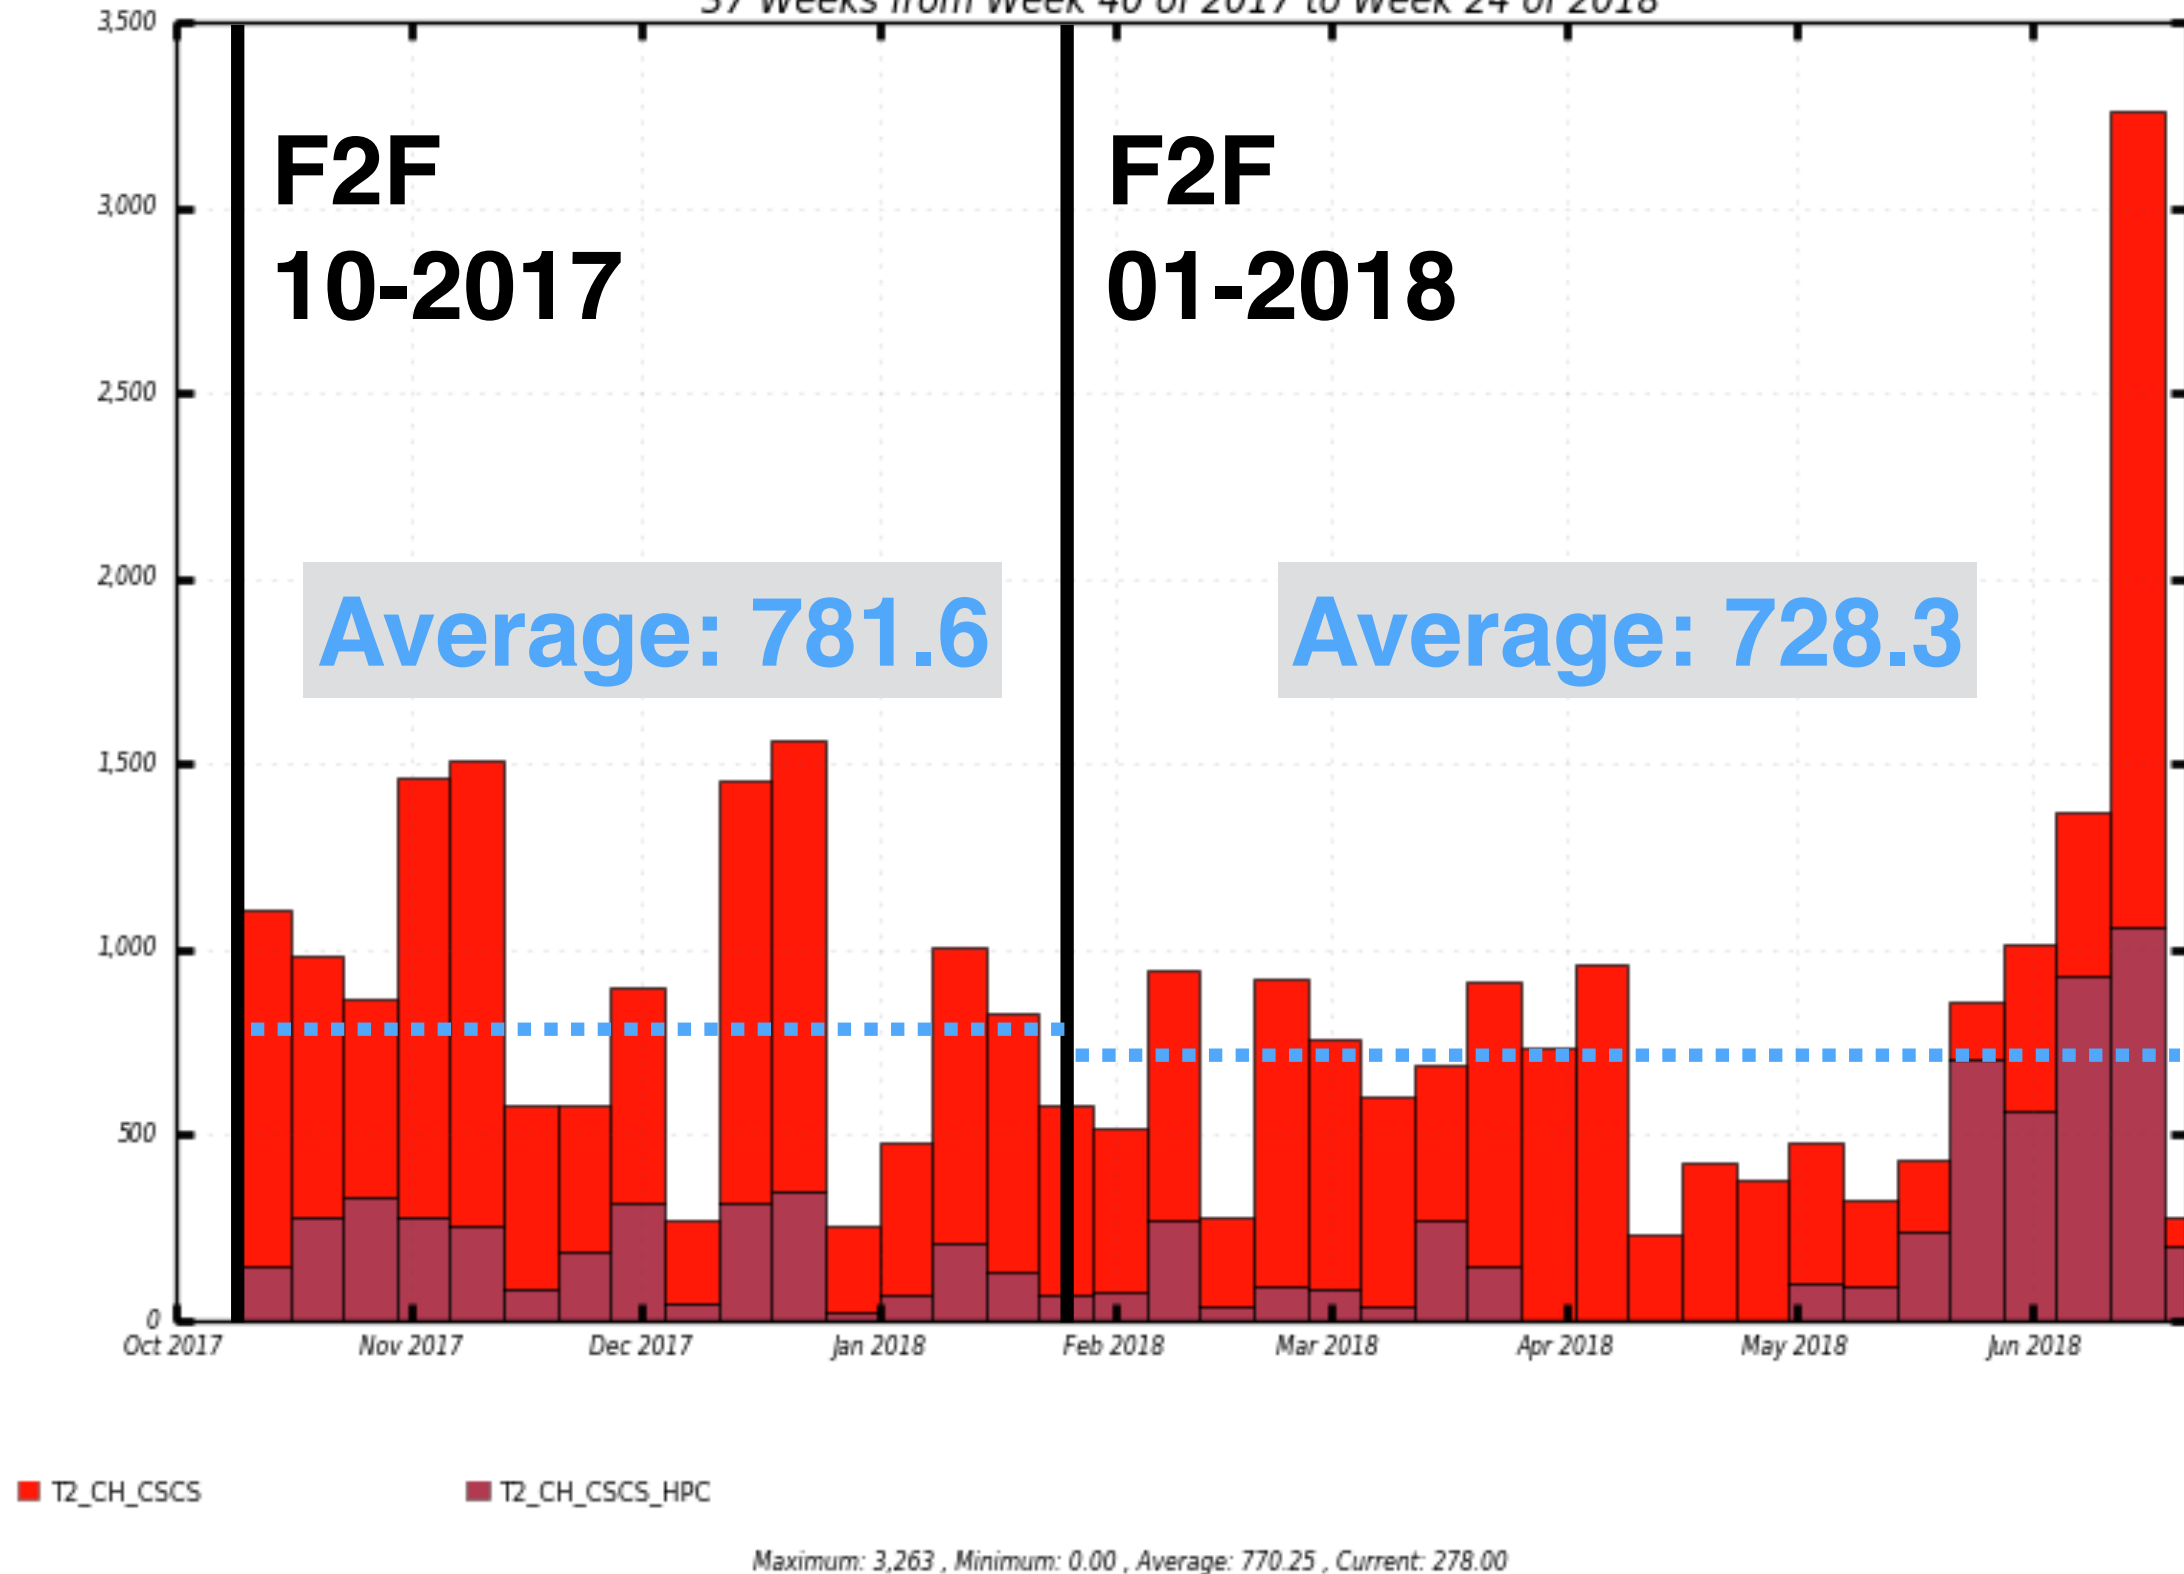
T2_CH_CSCS_HPC

Maximum: 3,263 , Minimum: 0.00 , Average: 770.25 , Current: 278.00

CPU Consumption Good jobs



weeks/week: CPU consumption Good Jobs
37 Weeks from Week 40 of 2017 to Week 24 of 2018

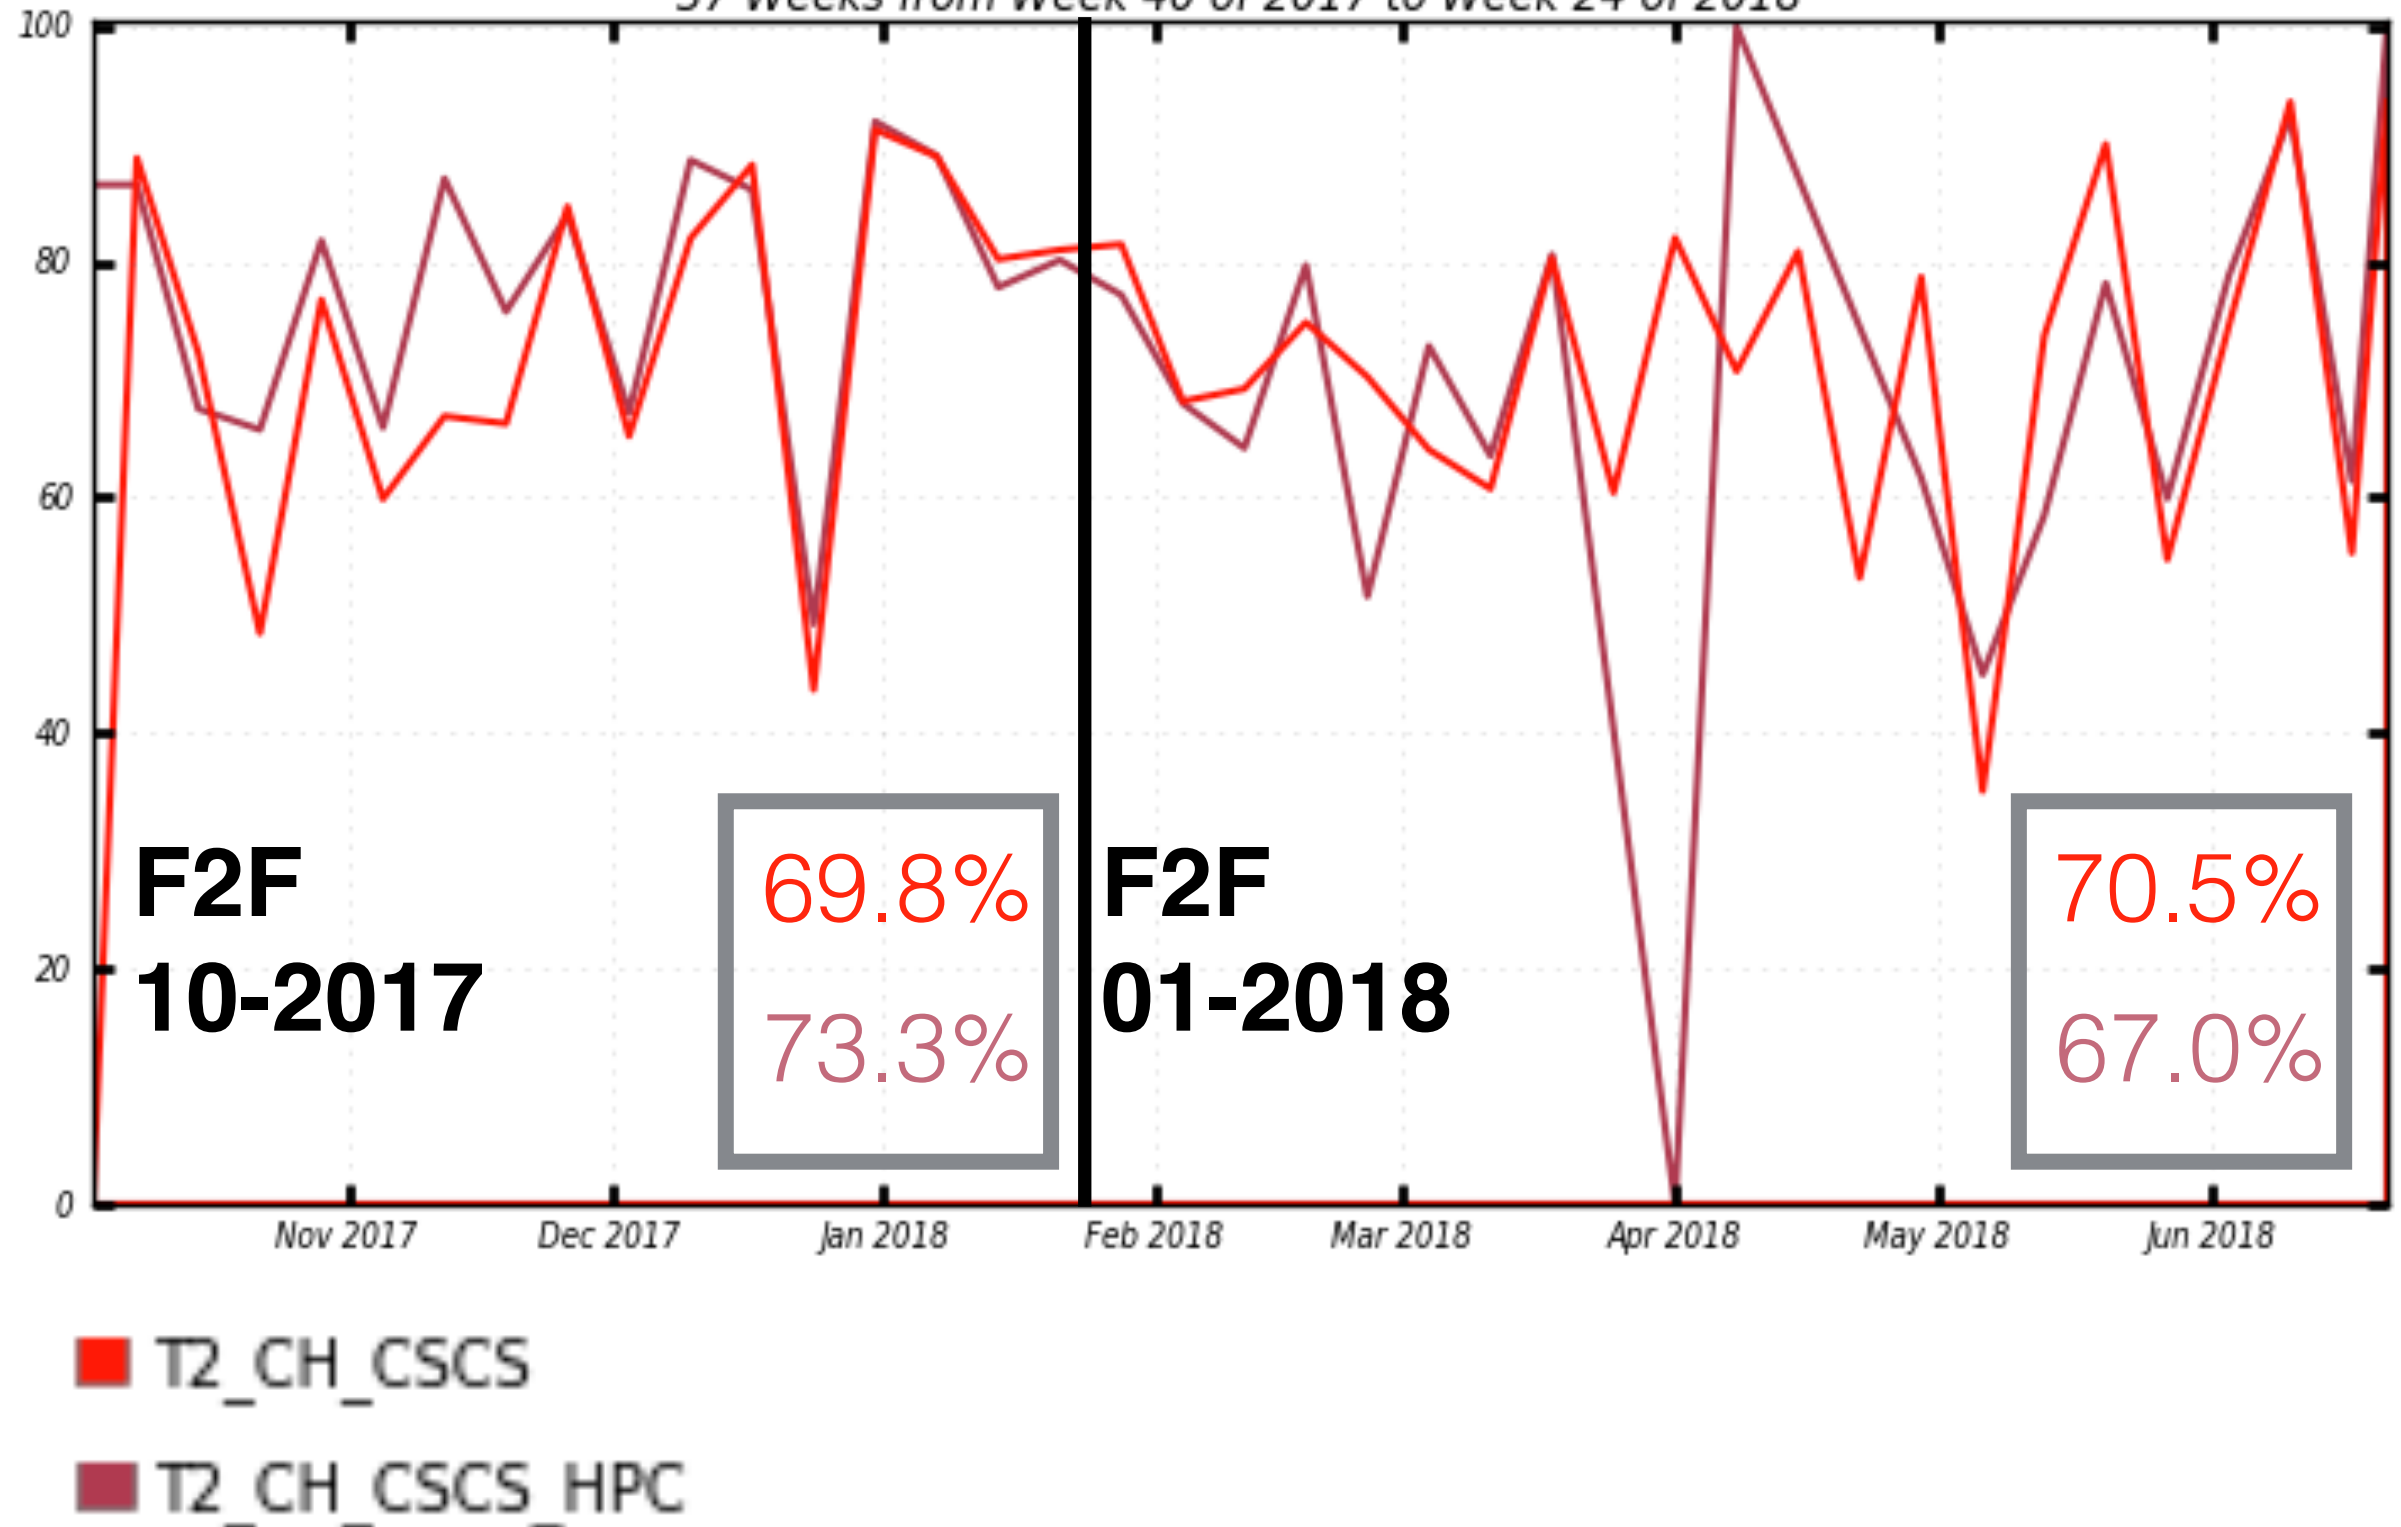


Successful jobs / # of all jobs

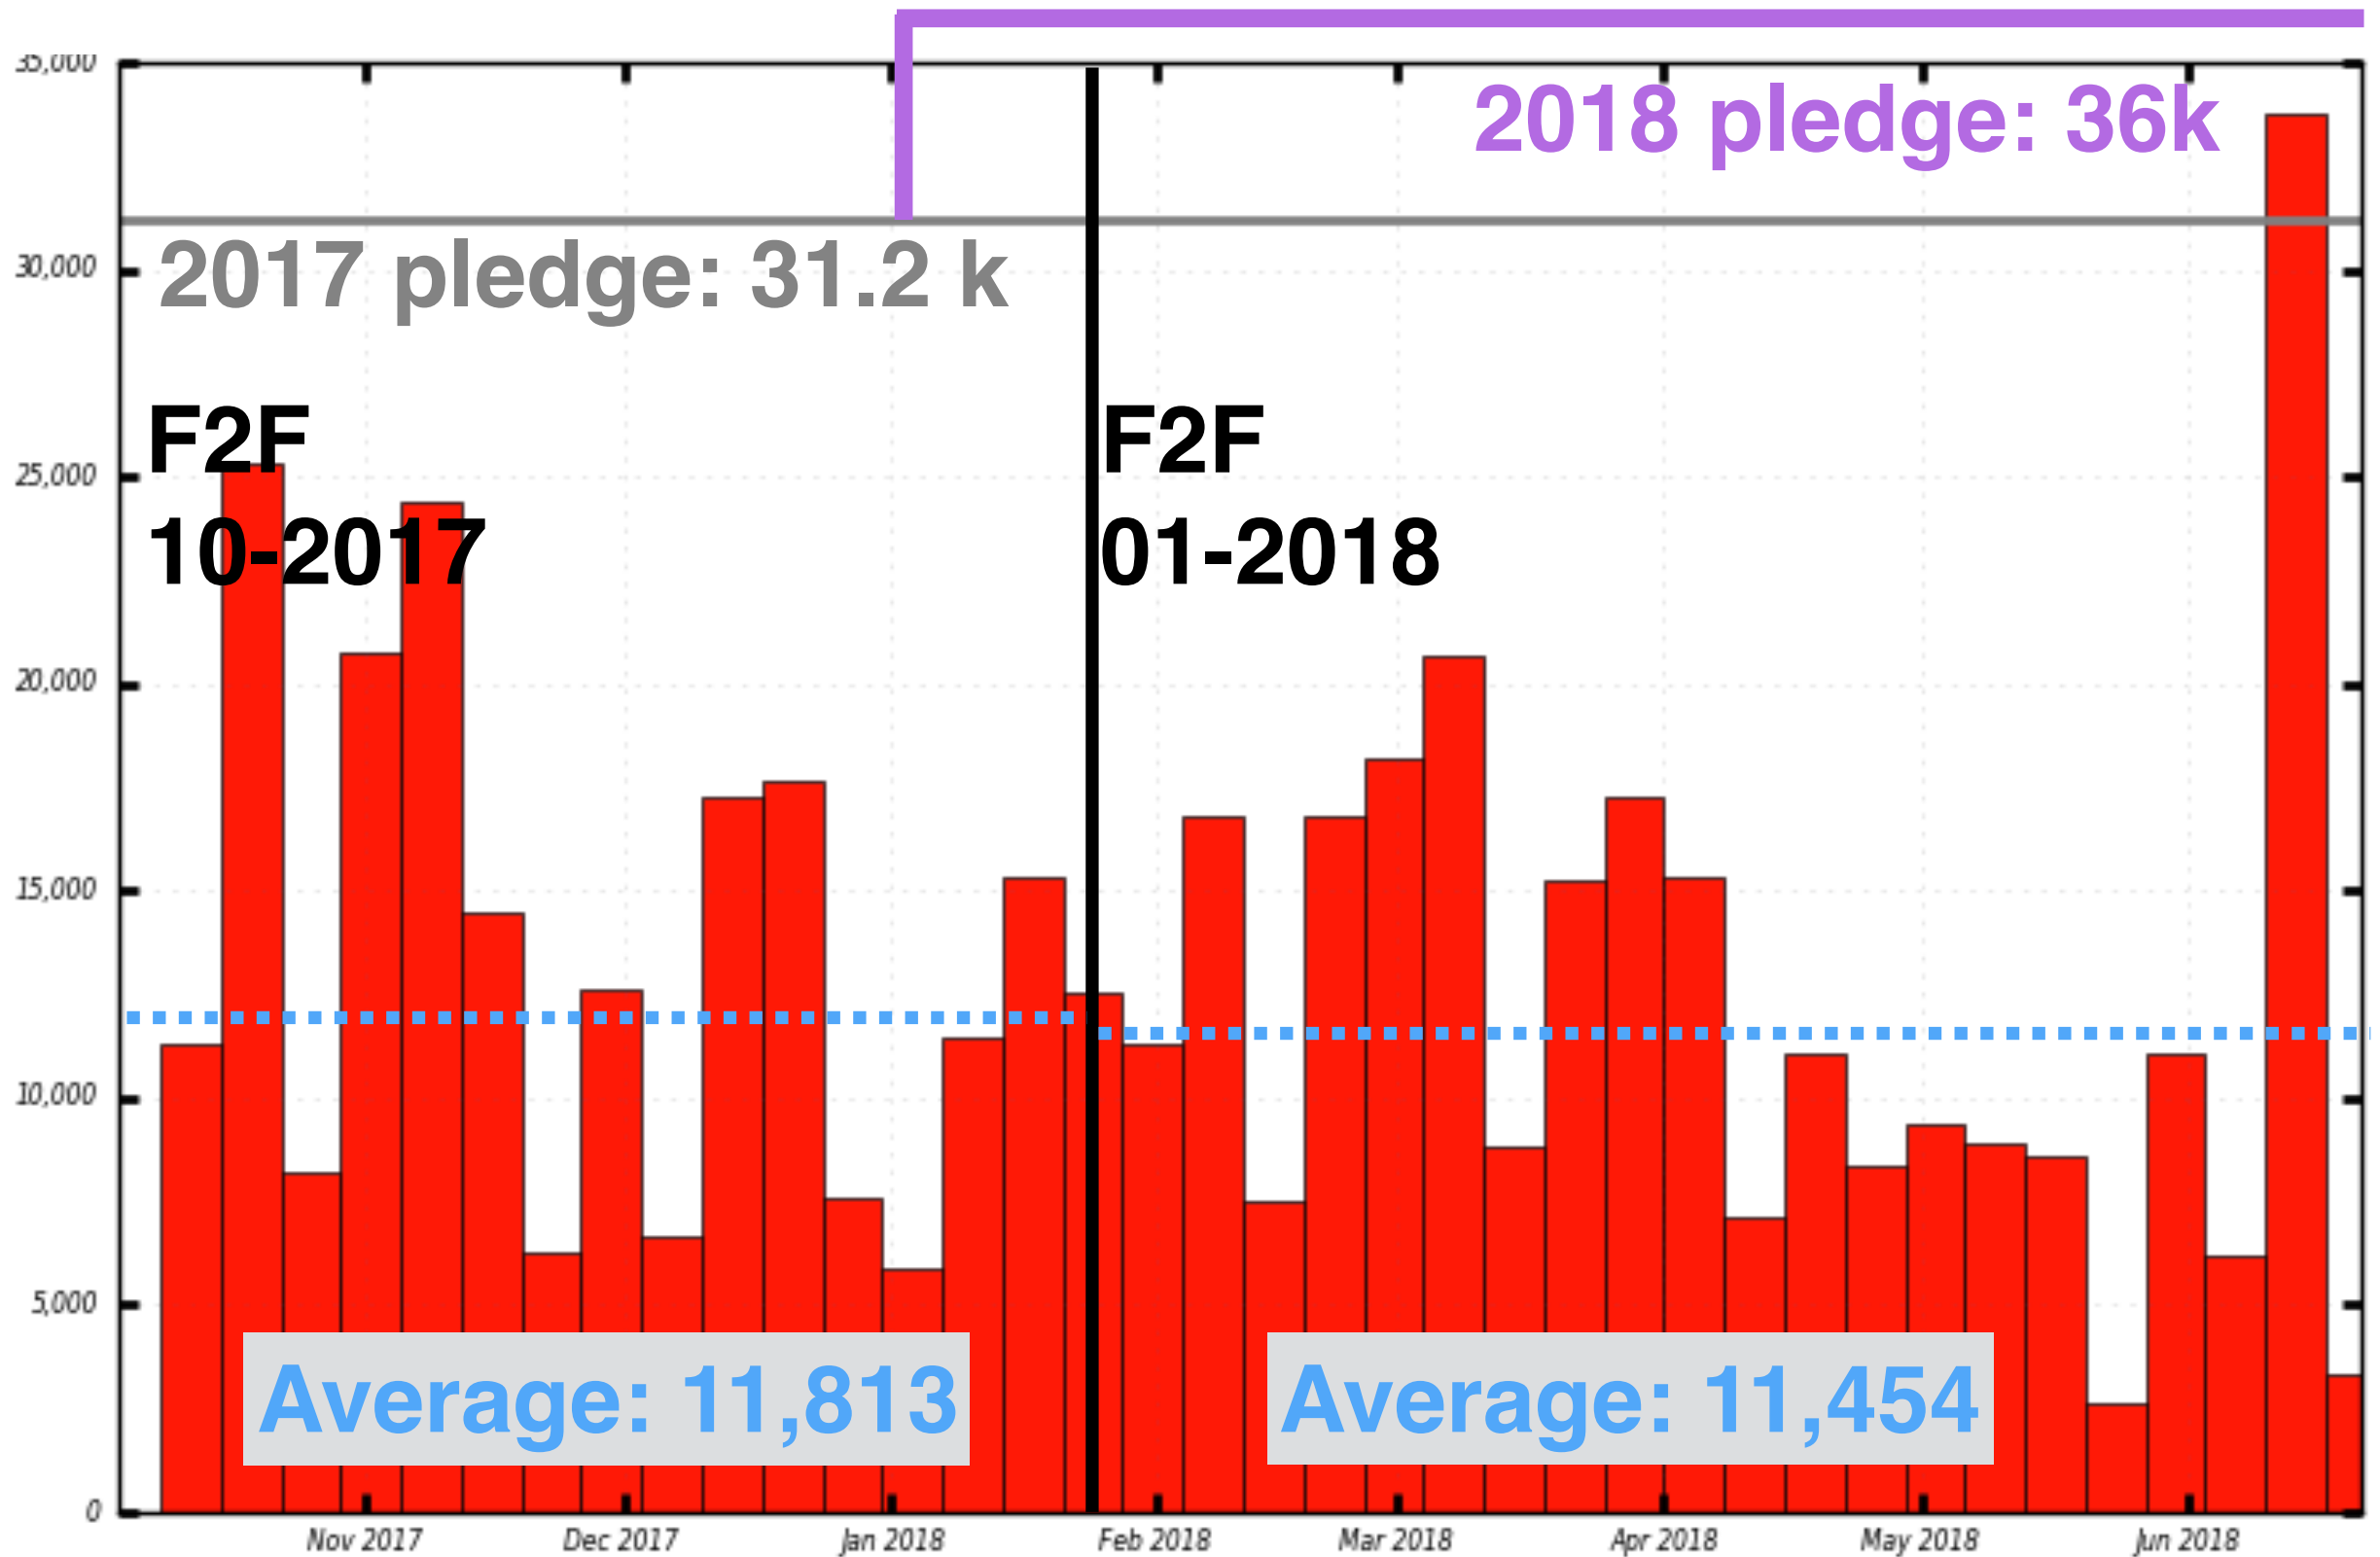


Efficiency over time based on success/all accomplished jobs


37 Weeks from Week 40 of 2017 to Week 24 of 2018



WC HEPSPREC06 Hours: pledged vs. delivered



WC HEPSPEC06 Hours: pledged vs. delivered



Conclusion

- SAM Ranking has decreased since last F2F
 - It looks to me like the tickets arrive too late; some unavailability has already happened by the time we are made aware
 - Look into some automatic warning system
 - If not existent, I will build some code that sends a warning email every time we go below a threshold
- When the system is not down, it works as stable as ever
- Overall delivery seems to have not increased yet, whereas the pledge has gone up from 31.2k to 36k

CSCS status report from CMS 01-10-2017 - 24-01-2018

CHIPP-CSCS F2F meeting in Lugano

Thomas Klijnsma

25th of January 2018

ETH zürich

GGUS tickets since 01-10-2017

Ticket #	Date	Title	Comment
127100	17/01/18	Upgrade spacemon client	Ongoing multi-party issue. Not very urgent, but taking a long time now.
132927	24/01/18	Problem with APEL Accounting for all of ...	Accounting issue; beyond my knowledge
132201	10/12/17	T2_CH_CSCS SAM srm failures	Quickly solved after updating certificates on certain storage elements
132125	09/12/17	T2_CH_CSCS_HPC SAM xrootd ...	Some errors due to CVMFS slowness/lack of cache. Solved by itself.
132017	08/12/17	T2_CH_CSCS SAM failures	Was already fixed before we read the ticket
131999	02/12/17	apply for EOS space for T2_CH_CERN	Related to a CERN user asking for EOS space in the wrong place.
131352	19/01/18	need for info on a specific WN	Related to high Slurm workload
131315	10/11/17	Transfer failing to T2_CH_CSCS - ...	Phedex certificate renewal scripts needed updating
131309	18/11/17	T2_CH_CSCS SAM CONDOR-JobSubmit failures	Slurm issues, but resolved itself.
131111	26/10/17	T2_CH_CSCS_HPC SAM tests cancelled	Watchdog was reenabled. SAM tests recovered.
131067	02/11/17	T2_CH_CSCS SAM JobSubmit failures	a-rex was in bad state; Once reenabled, Slurm overload. Caused a short WR status.

Summary of tickets

- No automatic ticket notification for me; still checking manually
- Mostly minor issues that lead to short unavailability:
 - Phedex certificates, certificates on SEs, SLURM overload, an issue with a-rex being in a bad state

Overview of important times in particle physics

- Follow-up to the previous F2F: Some important conferences of particle physics and their dates can be found at: <https://wiki.chipp.ch/twiki/bin/view/LCGTier2/BlogEphemeridesVOs>

VO Calendar of events/conferences


The idea is to maintain the following list of conferences/events where it may be important NOT to do scheduled maintenances, if possible.

2018

- **Moriond 2018**
 - Conference dates: 10-03 to 24-03
 - Avoid downtime between: 19-02 to 09-03
- **LHCP 2018**
 - Conference dates: 04-06 to 09-06
 - Avoid downtime between: 14-05 to 01-06
- **ICHEP 2018**
 - Conference dates: 04-07 to 11-07
 - Avoid downtime between: 18-06 to 02-07

SAM ranking since 01-10-2017

- Apparently rather widely accepted as 'the' ranking at management level
- In the lowest quartile during the last F2F, but now around the average
- Very green for the last month




CSCS_HPC:
96.50% 26/54

CSCS:
95.12% 29/54

Site Availability using CMS_CRITICAL


From 2017/10/01 to 2018/01/24



CPU Consumption per week



weeks/week: CPU consumption Good Jobs
55 Weeks from Week 01 of 2017 to Week 03 of 2018



T2_CH_CSCS

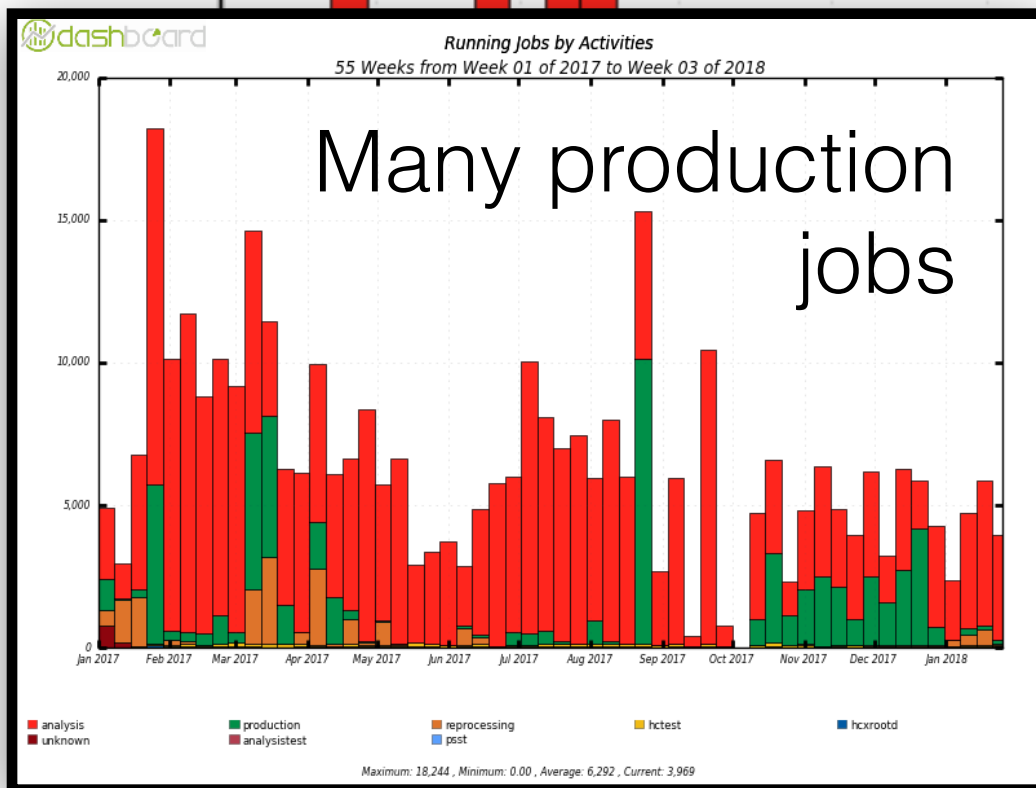
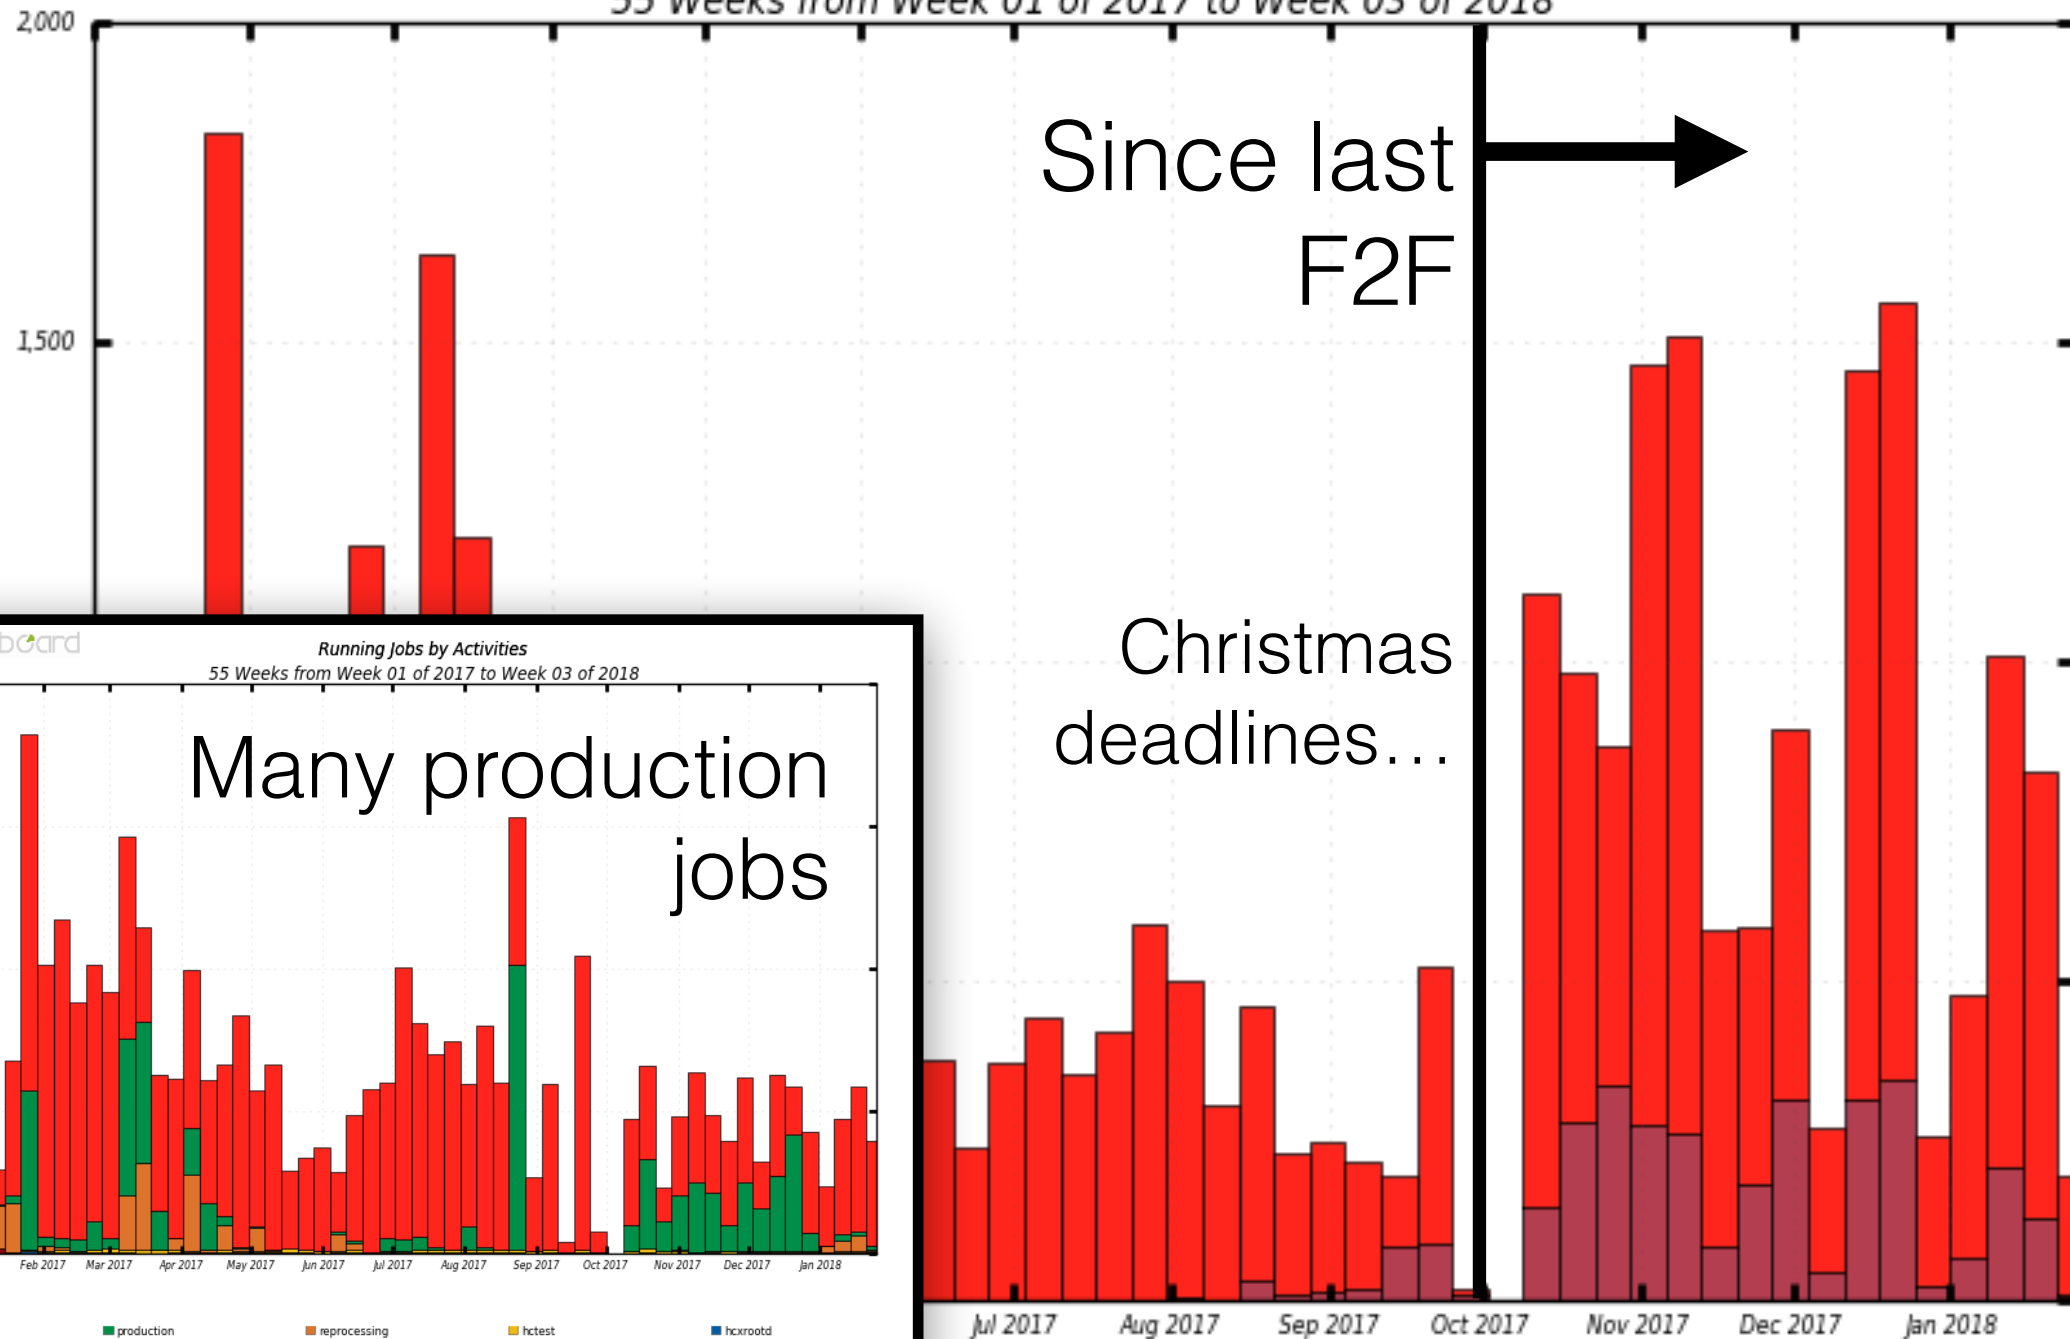
T2_CH_CSCS_HPC

Maximum: 1,828 , Minimum: 0.00 , Average: 608.95 , Current: 195.00

CPU Consumption per week



weeks/week: CPU consumption Good Jobs
55 Weeks from Week 01 of 2017 to Week 03 of 2018



T2_CH_CSCS

T2_CH_CSCS_HPC


Maximum: 1,828 , Minimum: 0.00 , Average: 608.95 , Current: 195.00

CPU Efficiency (all jobs)



Average Efficiency All Jobs

since 01-10-2017



May be *due* to the high number of (efficient) production jobs

During most of 2017, CSCS was closer to the average


Average: 62.3%

Success / total jobs



Efficiency over time based on success/all accomplished jobs

55 Weeks from Week 01 of 2017 to Week 03 of 2018




- Both systems stable since 10-2017, ~70% success/all
- Systems are strongly correlated

Failed jobs exit codes



Since 01-10-2017

Application failure by exit code (Sum: 67,739)
 terminated by wrapper because using too much RAM (RSS)





- ExitCode: 8028 - FileNotFoundError with fallback (14,763)
- ExitCode: 7002 - Config File Read Error (6,284)
- ExitCode: 0 - Application finished properly (4,061)
- ExitCode: 8021 - FileReadError (3,513)
- ExitCode: 139 - Segmentation violation (ANSI) (3,254)
- ExitCode: 1 - Hangup (POSIX) (1,779)
- ExitCode: 8002 - StdLib? Exception (694.00)
- ExitCode: 8009 - ConfigurationProblem (668.00)
- ExitCode: 50669 - Application terminated by wrapper for not defined reason (615.00)
- ExitCode: 255 - nondefined (499.00)
- ExitCode: 50660 - Application terminated by wrapper because using too much RAM (RSS) (10,437)
- ExitCode: 8019 - FileNotFoundError (5,464)
- ExitCode: 50664 - Application terminated by wrapper because using too much Wall Clock ti (642.00)
- ExitCode: 8001 - CMS exception (CMSSW) (3,447)
- ExitCode: 10040 - failed to generate cmsRun cfg file at runtime (2,927)
- ExitCode: 8020 - FileOpenError (777.00)
- ExitCode: 60302 - Output file(s) not found (694.00)
- ExitCode: 60311 - Stage Out Failure in ProdAgent job (642.00)
- ExitCode: 8006 - ProductNotFound (513.00)

WC Success/Failed

Since 01-10-2017



WallClock Consumption for Successful, Failed and/or Aborted Jobs
16 Weeks from Week 40 of 2017 to Week 03 of 2018



WallClock Consumption of Successful Jobs

WallClock Consumption of Failed and/or Aborted Jobs

Maximum: 570,075,663 , Minimum: 0.00 , Average: 319,399,962 , Current: 314,878,965

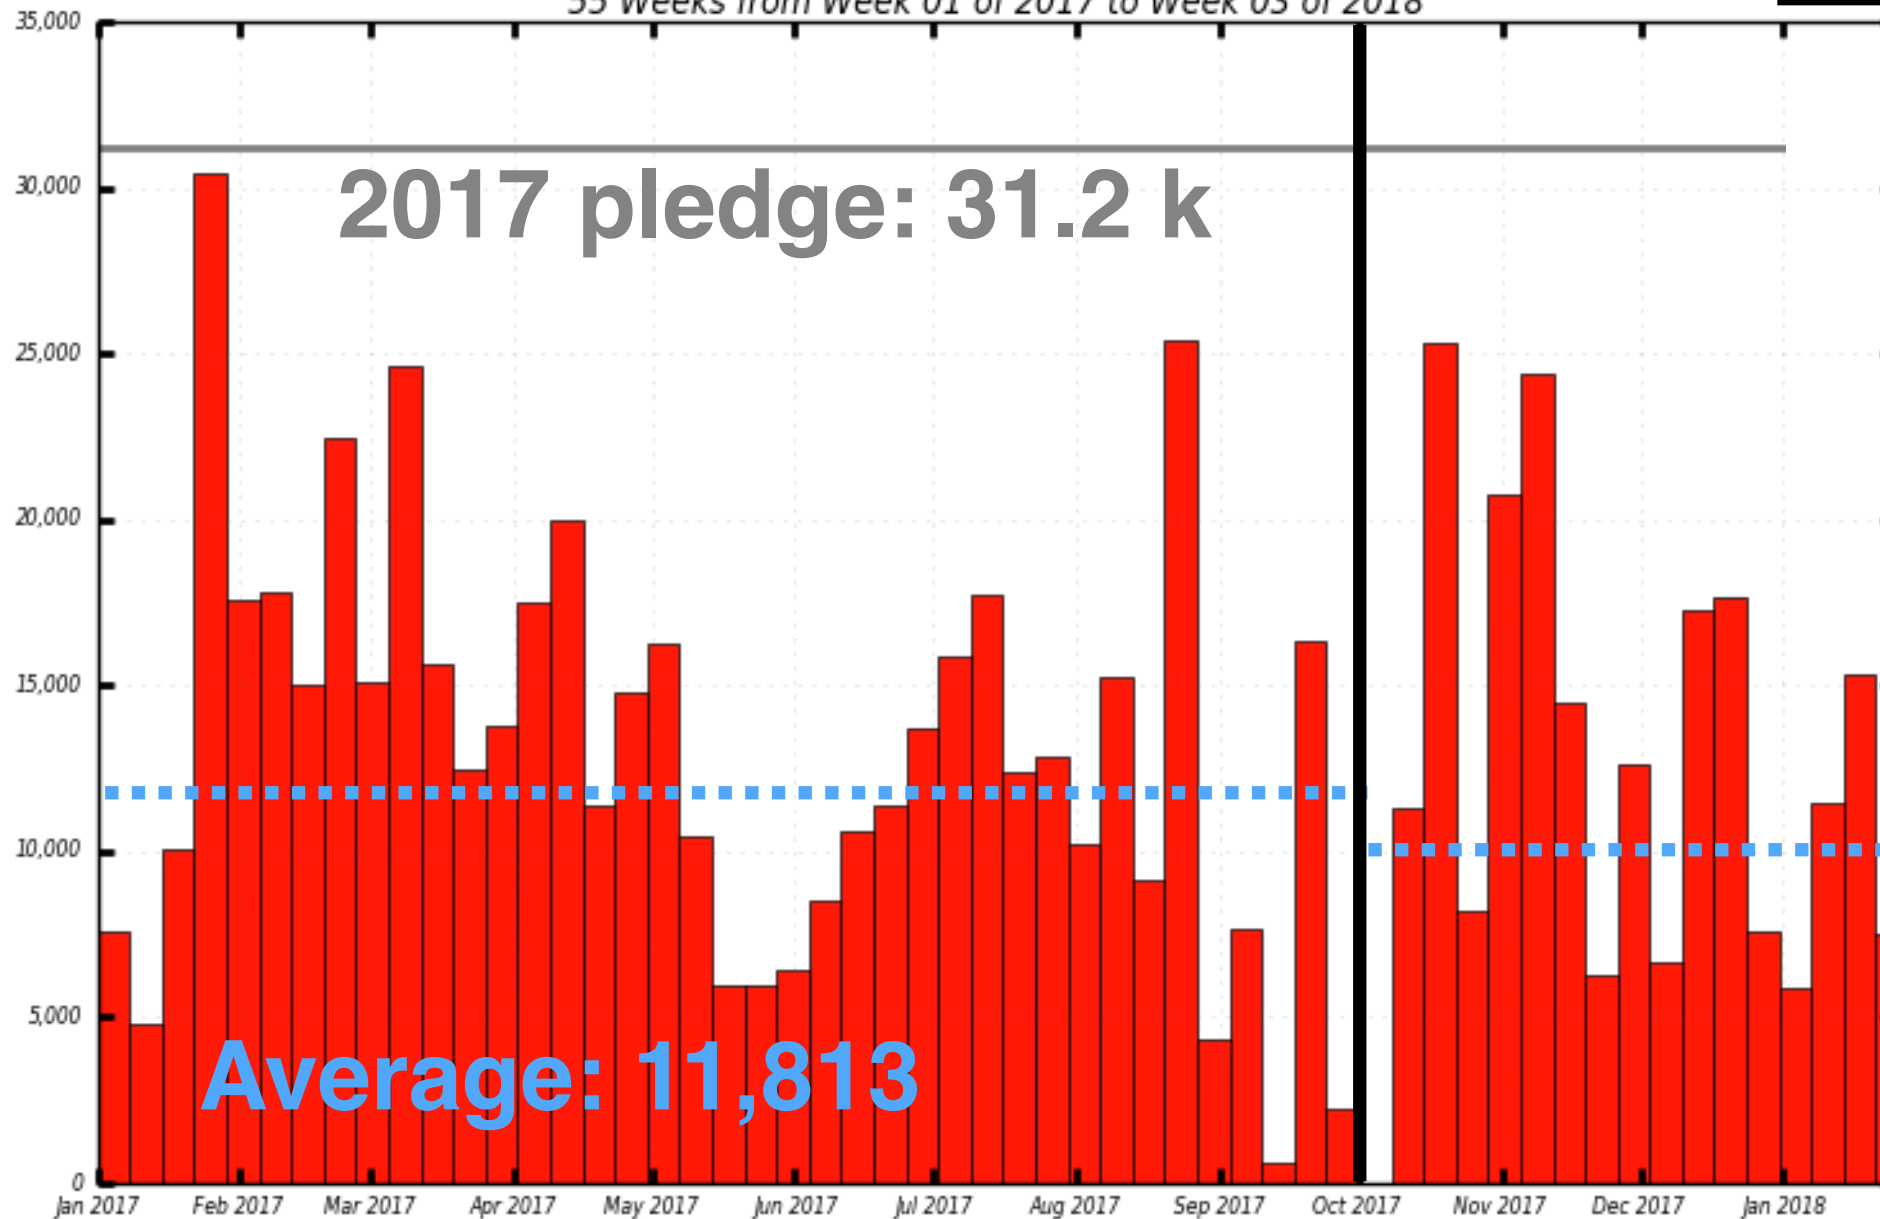
Resource pledge vs. delivered

- NB:** Due to accounting difficulties, take the info shown here with a grain of salt

**2018
pledge:
36k**




WallClock HEPSPROC6 Hours compared to pledges where available
55 Weeks from Week 01 of 2017 to Week 03 of 2018



SAM Availability overlaid with # of jobs

[Link to data](#)

Site Availability using CMS_CRITICAL_FULL



Maximum: 3,500 , Minimum: 0.00 , Average: 1,262 , Current: 1,780