

- ▶ Performance overview January - June 2018



# ATLAS T2 VO REPORT

---

Gianfranco Sciacca

AEC - Laboratory for High Energy Physics, University of Bern, Switzerland

CHIPP-CSCS face 2 face - 21 June 2018

## Availability - HammerCloud AFTs and PFTs

### Historic view for "panda\_queues\_all" from 00:00 01.01.2018 to 00:00 15.06.2018

Show 100 entries

Search:

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 330 hours	offline		brokeroff		online		NoQueue		test	
					%	count	%	count	%	count	%	count	%	count
ANALY_CSCS	CSCS-LCG2	T2D	DE		2.59	4	0.19	2	94.64	22	0	0	2.05	10
ANALY_CSCS-HPC	CSCS-LCG2	T2D	DE		3.67	6	0.29	3	90.26	22	0	0	5.26	10
ANALY_UNIBE-LHEP	UNIBE-LHEP	T2	ND		1.27	2	0.1	1	94.94	34	0	0	3.16	23
ANALY_UNIBE-LHEP-UBELIX	UNIBE-LHEP	T2	ND		1.27	2	0.1	1	94.46	27	0	0	3.64	16
CSCS-LCG2	UNDEFINED	UNDEFINED	UNDEFINED		1.53	2	24.33	3	73.34	14	0	0	0.27	5
CSCS-LCG2-HPC	CSCS-LCG2	T2D	DE		2.42	4	24.33	3	70.03	14	0	0	2.69	5
CSCS-LCG2-HPC_MCORE	CSCS-LCG2	T2D	DE		4.26	6	0	0	92.52	18	0	0	2.69	5
CSCS-LCG2_MCORE	UNDEFINED	UNDEFINED	UNDEFINED		2.99	4	0	0	95.15	19	0	0	1.34	7
UNIBE-LHEP	UNIBE-LHEP	T2	ND		1.09	2	24.33	3	71.25	14	0	0	2.81	5
UNIBE-LHEP-UBELIX	UNIBE-LHEP	T2	ND		0.03	1	24.33	3	74.62	20	0	0	0.49	12
UNIBE-LHEP-UBELIX_MCORE	UNIBE-LHEP	T2	ND		0.03	1	0	0	98.95	22	0	0	0.49	12
UNIBE-LHEP_MCORE	UNIBE-LHEP	T2	ND		1.09	2	0	0	95.53	18	0	0	2.86	7

Showing 1 to 12 of 12 entries

First Previous 1 Next Last



## Availability

- **CSCS-LCG2**: 96.58% (rank 40 / 76 T2s - was 46)
- **UNIBE-LHEP**: 98.01% (rank 32 / 76 T2s - was 32)

**Algorithm:**  
(CREAM-CE + ARC-CE +  
HTCONDOR-CE +  
GLOBUS) \* (all SRMv2 +  
all SRM + all GRIDFTP)

## Site Availability using ATLAS\_CRITICAL


From 2018/01/01 to 2018/06/15



## Service Availability

### Service Availability using ATLAS\_CRITICAL

From 2018/01/01 to 2018/06/15





## Reliability

- **CSCS-LCG2:** 96.95% (rank 43 / 76 T2s - was 36)
- **UNIBE-LHEP:** 98.62% (rank 33 / 76 T2s - was 24)

**Algorithm:**  
 (CREAM-CE + ARC-CE +  
 HTCNDOR-CE +  
 GLOBUS) \* (all SRMv2 +  
 all SRM + all GRIDFTP)


## Site Reliability using ATLAS\_CRITICAL

From 2018/01/01 to 2018/06/15



## WallClock HS06 - CH-CHIPP-CSCS Federation

- **CSCS-LCG2 HS06 ATLAS installed (\*)**: 43970 (66%) - **pledged: 36000** (67%) (2018)  
of total
- **UNIBE-LHEP HS06 ATLAS installed (\*\*)**: 22000 (34%) - **pledged 18000** (33%) (2018)  
of total



**total installed**

**total pledged**

**CSCS +15%**  
(why not +27% ?)

**Bern +20%**

■ CSCS-LCG2


■ UNIBE-LHEP

(\*) - ATLAS share

(\*\*) - also serving t2k.org and fermilab/uboone  
- some opportunistic usage on Ubelix

## WallClock HS06 - CSCS-LCG2

- **CSCS-LCG2 HS06 ATLAS installed: 43970 - pledged: 36000** (2018)



installed

pledged

Phoenix 67%

Daint 33%  
(incl. T0)

■ grid      ■ hpc  
(\* ) - ATLAS share

Maximum: 41,844 , Minimum: 0.00 , Average: 29,584 , Current: 31,048



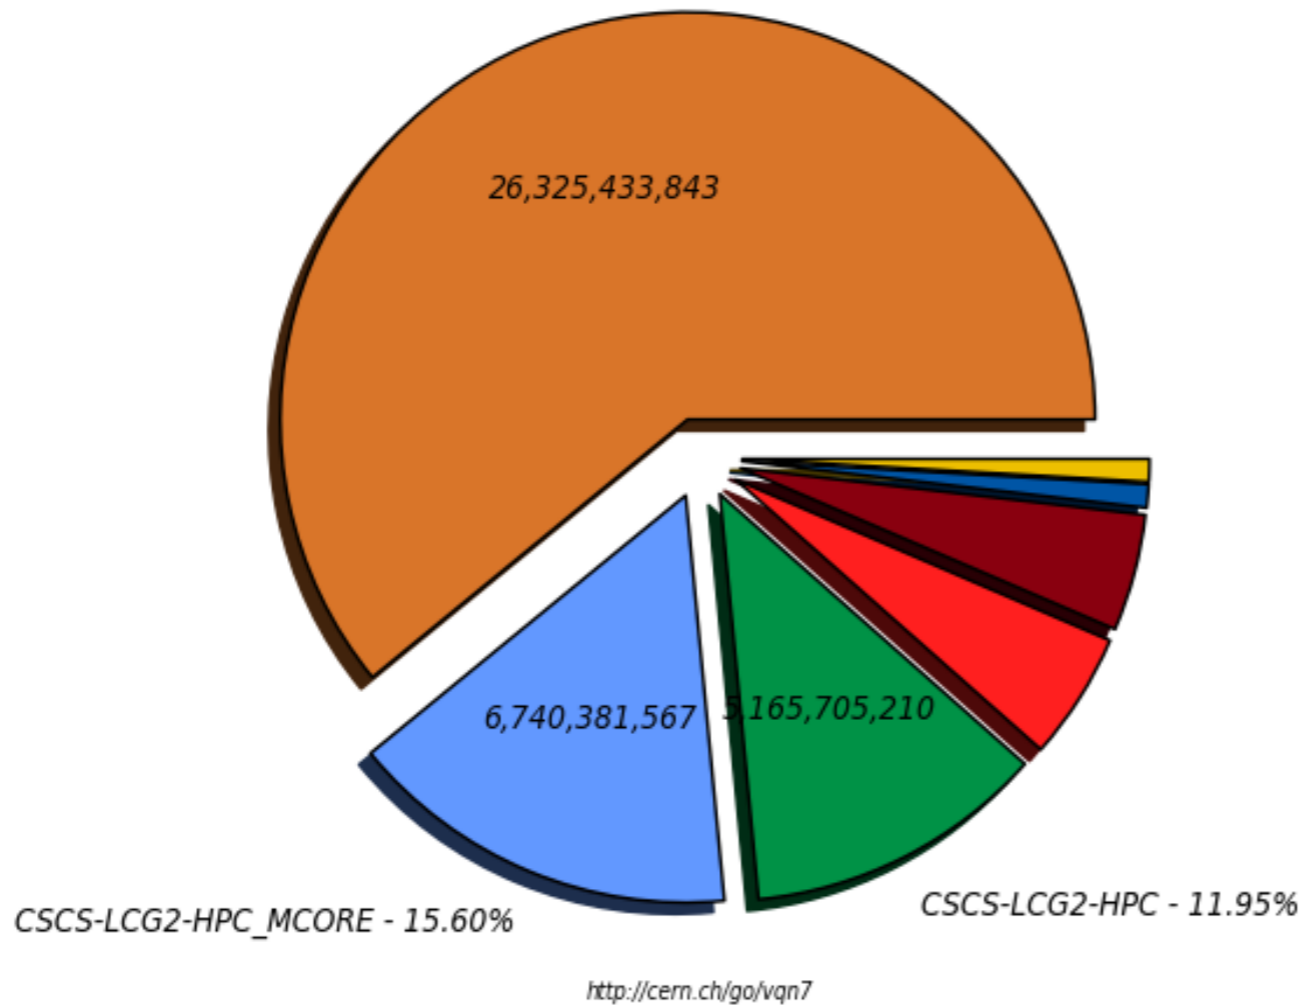
## WallClock HS06 - CSCS-LCG2

- **CSCS-LCG2 HS06 ATLAS installed: 43970 - pledged: 36000** (2018). (\*)



Wall Clock consumption All Jobs in seconds (Sum: 43,214,388,603)

CSCS-LCG2\_MCORE - 60.92%



**HPC capacity: 31% of total**

**HPC delivered: 28% of total**  
(after subtracting the T0 activity)


- CSCS-LCG2\_MCORE - 60.92% (26,325,433,843)
- CSCS-LCG2-HPC - 11.95% (5,165,705,210)
- CSCS-PIZDAINT - 4.64% (2,004,767,596)
- ANALY\_CSCS-HPC - 0.90% (388,765,679)
- CSCS-LCG2-HPC\_MCORE - 15.60% (6,740,381,567)
- ANALY\_CSCS - 5.06% (2,188,588,429)
- CSCS-LCG2 - 0.93% (400,746,279)

(\*) - ATLAS share



## WallClock HS06 - UNIBE-LHEP

- **UNIBE-LHEP HS06 ATLAS installed (\*\*): 22000 - pledged: 18000 (2018)**




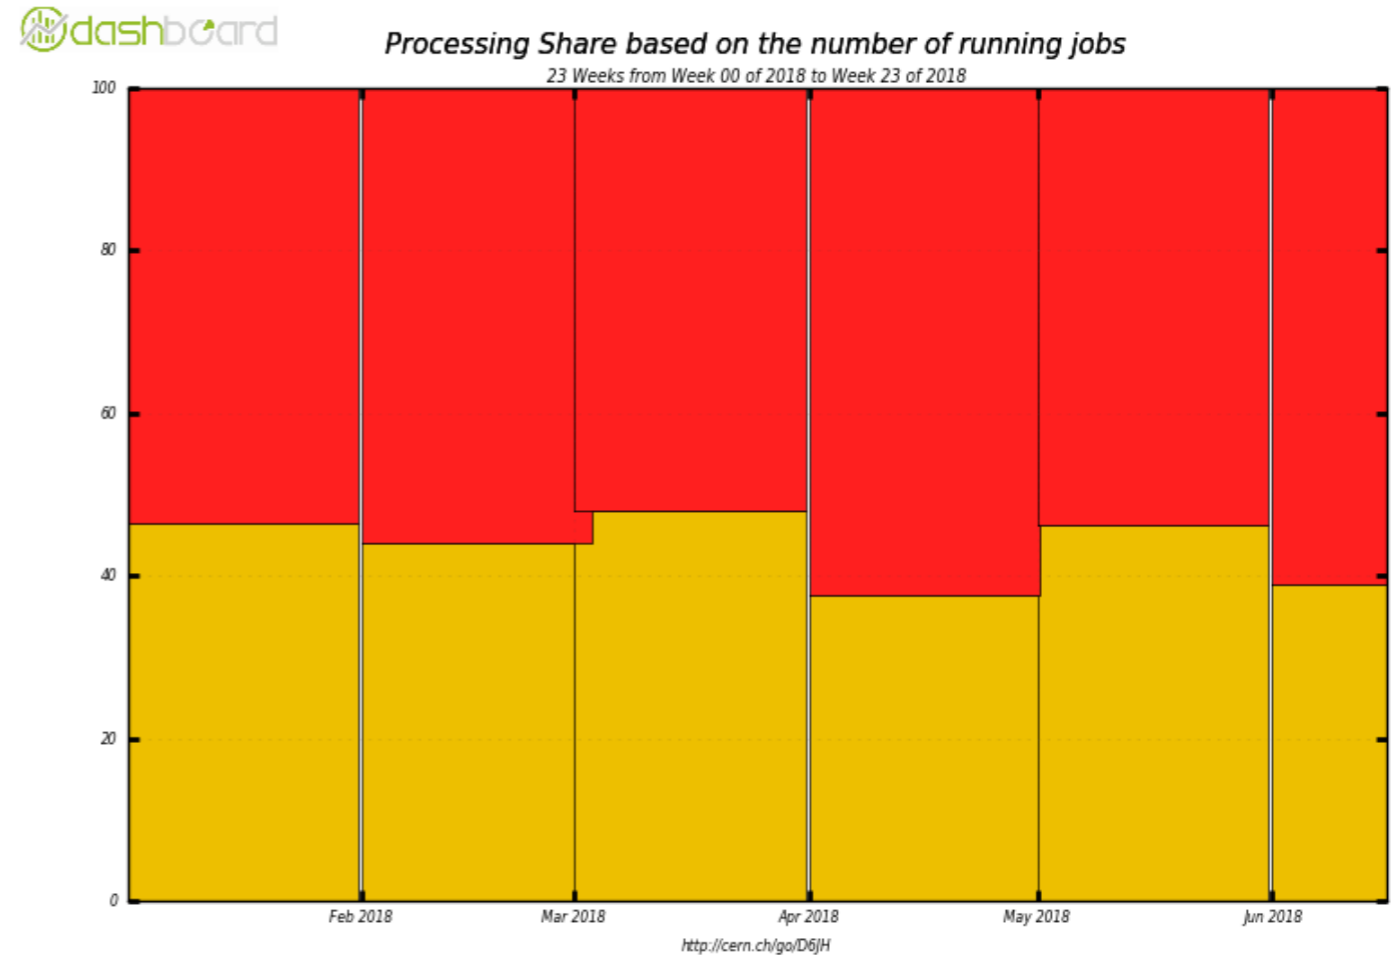

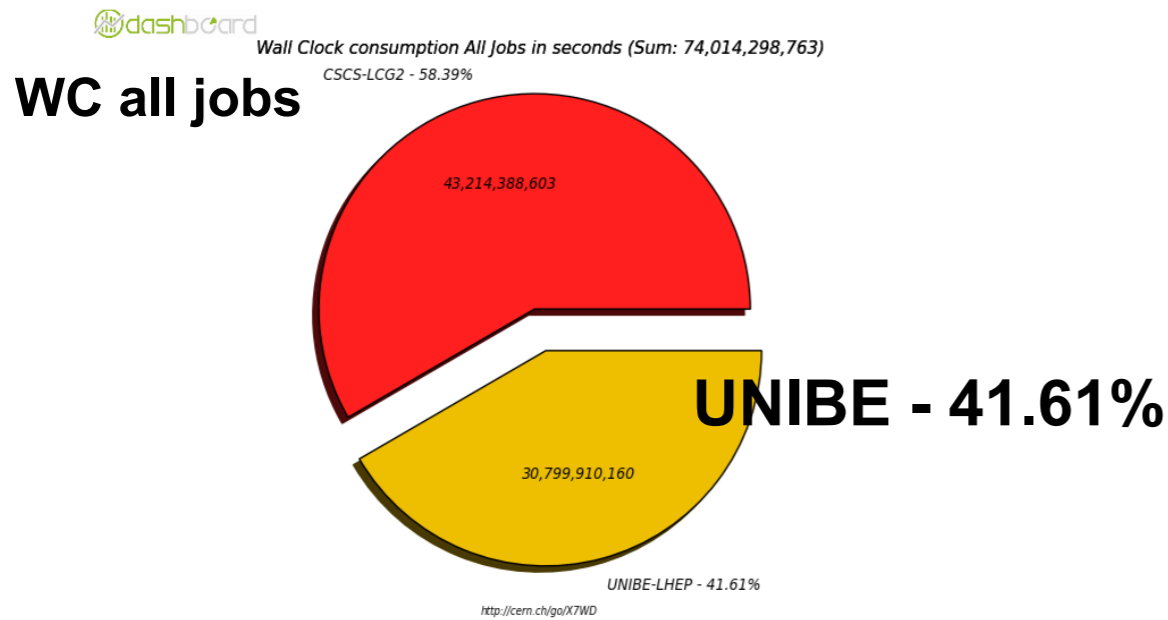
installed  
pledged

■ UNIBE-LHEP

Maximum: 24,512 , Minimum: 0.00 , Average: 18,781 , Current: 13,992

## Relative shares

- **UNIBE-LHEP installed estimated in about 37% of the total installed capacity**

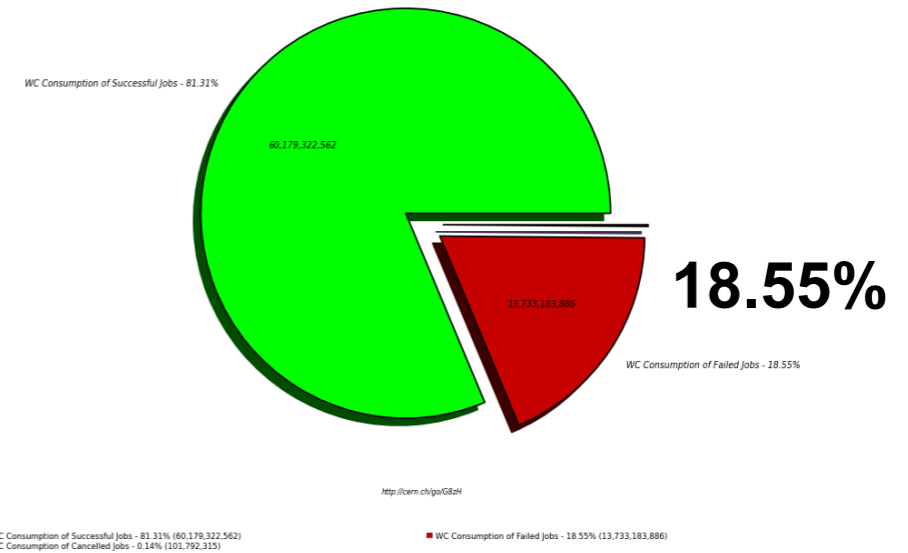


## Success vs fail WallClock efficiency


- **CSCS-LCG2: 76%**
- **UNIBE-LHEP: 90%**



WC Consumption for Successful and Failed Jobs (Sum: 74,014,298,763)



WallClock Consumption for Successful and Failed Jobs  
23 Weeks from Week 00 of 2018 to Week 23 of 2018




WallClock Consumption of Successful Jobs  
WallClock Consumption of Cancelled Jobs

WallClock Consumption of Failed Jobs

Maximum: 15,143,557,217, Minimum: 0.00, Average: 10,573,471,251, Current: 5,222,572,535



WallClock Efficiency over time based on success/all accomplished jobs  
23 Weeks from Week 00 of 2018 to Week 23 of 2018



UNIBE-LHEP (89.44) CSCS-LCG2 (76.28)

Total: 145.52, Average Rate: 0.00 /s




## Success vs fail WallClock efficiency

- **Phoenix: 81%**
- **UNIBE-LHEP: 71%**



WallClock Efficiency over time based on success/all accomplished jobs  
165 Days from Week 00 of 2018 to Week 23 of 2018



■ grid (80.87)

■ hpc (71.11)

Total: 150.26 , Average Rate: 0.00 /s


## CPU / WallClock efficiency

- **CSCS-LCG2: 80%**
- **UNIBE-LHEP: 85%**



Efficiency Good Jobs

23 Weeks from Week 00 of 2018 to Week 23 of 2018




■ CSCS-LCG2 (0.80) ■ UNIBE-LHEP (0.85)

Total: 1.58 , Average Rate: 0.00/s



Average Efficiency Good Jobs




<http://cern.ch/go/sK8K>

## CPU / WallClock efficiency by PanDA queue




Efficiency Good Jobs  
165 Days from Week 00 of 2018 to Week 23 of 2018




## CPU / WallClock efficiency by ADC activity



**Efficiency Good Jobs**  
165 Days from Week 00 of 2018 to Week 23 of 2018



<http://cern.ch/go/sp9v>

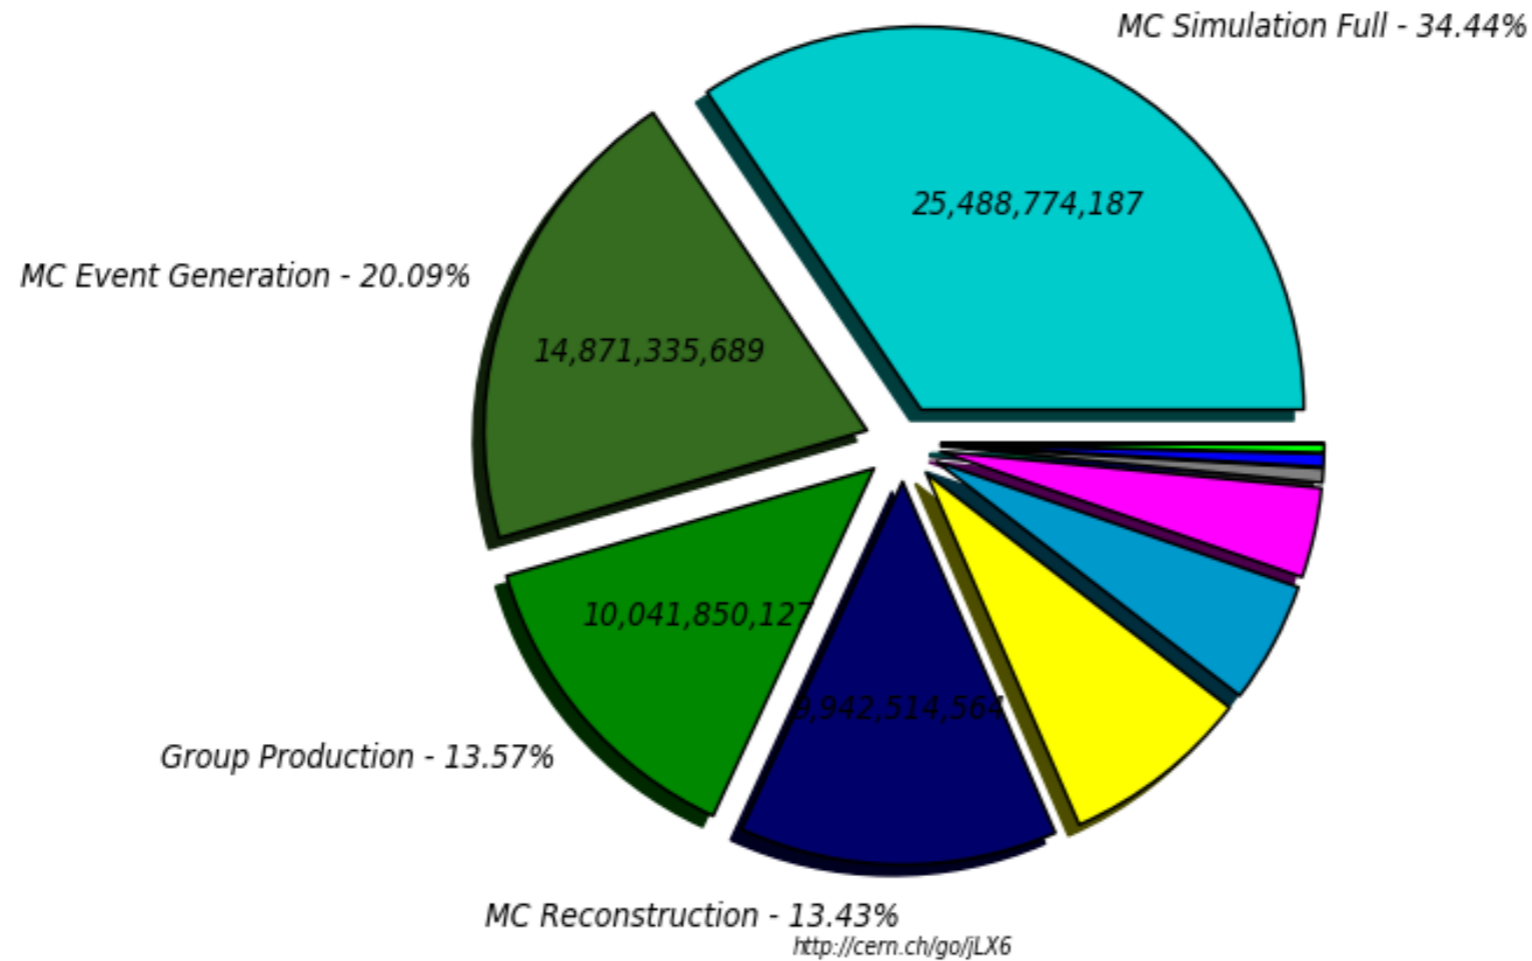




## WallClock by ADC activity



Wall Clock consumption All Jobs in seconds (Sum: 74,014,298,763)




- MC Simulation Full - 34.44% (25,488,774,188)
- Group Production - 13.57% (10,041,850,127)
- Data Processing - 8.00% (5,924,672,460)
- User Analysis - 3.80% (2,812,282,029)
- MC Simulation - 0.54% (402,934,091)
- Others - 0.01% (3,899,970)
- MC Event Generation - 20.09% (14,871,335,689)
- MC Reconstruction - 13.43% (9,942,514,564)
- MC Simulation Fast - 5.15% (3,808,808,465)
- Group Analysis - 0.61% (447,969,150)
- Testing - 0.36% (269,258,030)

## CPU / WallClock efficiency by type of job



Wall Clock consumption All Jobs in seconds (Sum: 43,214,388,603)

**CSCS**  
MC reco 63%  
Group prod 69%




- MC Simulation Full - 38.71% (16,729,939,649)
- MC Reconstruction - 15.06% (6,506,578,149)
- Group Production - 11.82% (5,108,606,364)
- Data Processing - 10.50% (4,537,773,734)
- MC Event Generation - 11.74% (5,074,783,839)
- MC Simulation Fast - 5.38% (2,325,035,455)
- User Analysis - 5.03% (2,171,556,921)
- MC Simulation - 0.49% (212,558,975)
- Group Analysis - 0.84% (363,234,283)
- Testing - 0.42% (180,740,561)
- Others - 0.01% (3,580,672)



Efficiency Good Jobs

23 Weeks from Week 00 of 2018 to Week 23 of 2018




- MC Event Generation (0.92)
- MC Simulation Full (0.96)
- MC Simulation (0.64)
- User Analysis (0.77)
- MC Simulation Fast (0.88)
- Group Analysis (0.81)
- Group Production (0.69)
- Others (0.54)
- MC Reconstruction (0.63)
- Data Processing (0.45)
- Testing (0.56)

Total: 6.72, Average Rate: 0.00 /s



Wall Clock consumption All Jobs in seconds (Sum: 30,799,910,160)

**UNIBE**  
MC reco 64%  
Group prod 73%




- MC Event Generation - 31.81% (9,796,551,850)
- MC Simulation Full - 28.44% (8,758,834,539)
- Group Production - 16.02% (4,933,243,763)
- MC Reconstruction - 11.16% (3,435,936,415)
- MC Simulation Fast - 4.82% (1,483,773,010)
- Data Processing - 4.50% (1,386,898,725)
- User Analysis - 2.08% (640,725,108)
- MC Simulation - 0.62% (190,375,116)
- Group Analysis - 0.28% (84,734,867)
- Testing - 0.29% (88,517,469)
- Others - 0.00% (319,298)



Efficiency Good Jobs

23 Weeks from Week 00 of 2018 to Week 23 of 2018



- MC Simulation Full (0.95)
- MC Event Generation (0.91)
- User Analysis (0.79)
- MC Simulation Fast (0.88)
- MC Simulation (0.84)
- Group Production (0.73)
- Group Analysis (0.72)
- MC Reconstruction (0.64)
- Data Processing (0.61)
- Testing (0.59)

Total: 6.51, Average Rate: 0.00 /s

## Summary

- ▶ **WC HS06 delivered vs. Pledged on target**
  - ▶ 2018 pledged lower than it should have been
- ▶ **A bit short in delivered vs. Installed since April**
- ▶ **Ranking similar to last semester report, still room for improvement**
- ▶ **Improved on some recurring problems**
  - ▶ GPFS performance and stability, CPU/WC efficiency
- ▶ **Needs addressing:**
  - ▶ **Failed WC still high**
  - ▶ **Is arc04 stable? Or just a monitoring artefact?**

## Outlook 1/2

- ▶ **Future** (discussions, r&d, etc):
  - ▶ **Joining the NDGF data-“fjord”?** (*nordic version of a data-lake*)
    - ▶ **Both UNIBE and CSCS**
    - ▶ **dCache based**
    - ▶ **dCache pools under the NDGF head nodes**
    - ▶ **Currently 4 Tier-1 storage sties and 1 Tier-2 storage site (Slovenia)**
    - ▶ ***dCache pools can be managed remotely by an unprivileged user***
    - ▶ **Could open ports to tape usage for CSCS**
    - ▶ *What about politics?*

## Outlook 2/2

- ▶ **Future** (discussions, r&d, etc):

- ▶ **Object stores**

- ▶ **For pure object store, need for ATLAS development for protocols (s3, swift)**
- ▶ **First step in that direction would be add a CEPH FS pool to dCache (LHC protocols)**
- ▶ **Better QoS for data loss, but it has to be tested vs. offering less disk space**

- ▶ **Tape @ CSCS**

- ▶ **There could be a concrete use case for a new storage model for data lakes**
- ▶ **More storage tiers with different QoS**
- ▶ **Again, this is new r&d for ATLAS, discussions ongoing**
- ▶ **The tape@CSCS case will be brought up next week at the ATLAS s&c week**
- ▶ *Currently tape needs dCache, unless we can provide another interface based on https, root, s3 that is usable by FTS*

- ▶ **Event service with arc05?**