This work is supported by projects Research infrastructure CERN (CERN-CZ, LM2015058) and OP RDE CERN Computing (CZ.02.1.01/0.0/0.0/16013/0001404) from EU funds and MEYS.

EUROPEAN UNION
European Structural and Investment Funds
Operational Programme Research,
Development and Education

Prague Site Report

Martin Adam
J.Chudoba, A.Mikula, P.Horák, P.Vokáč, J.Uhlířová, M.Svatoš, D.Adamová

Computing Center of Institute of Physics CAS

Almost 8000 cores published

~6PB of storage (DPM, XrootD, NFS)

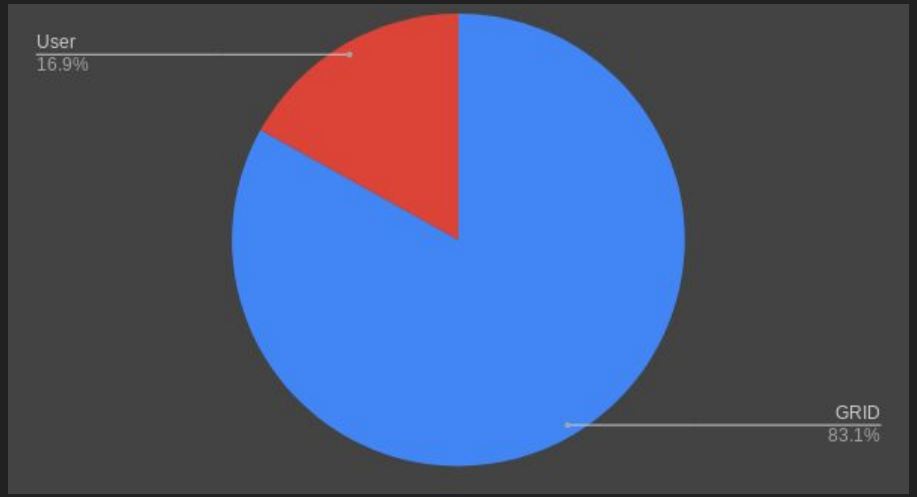

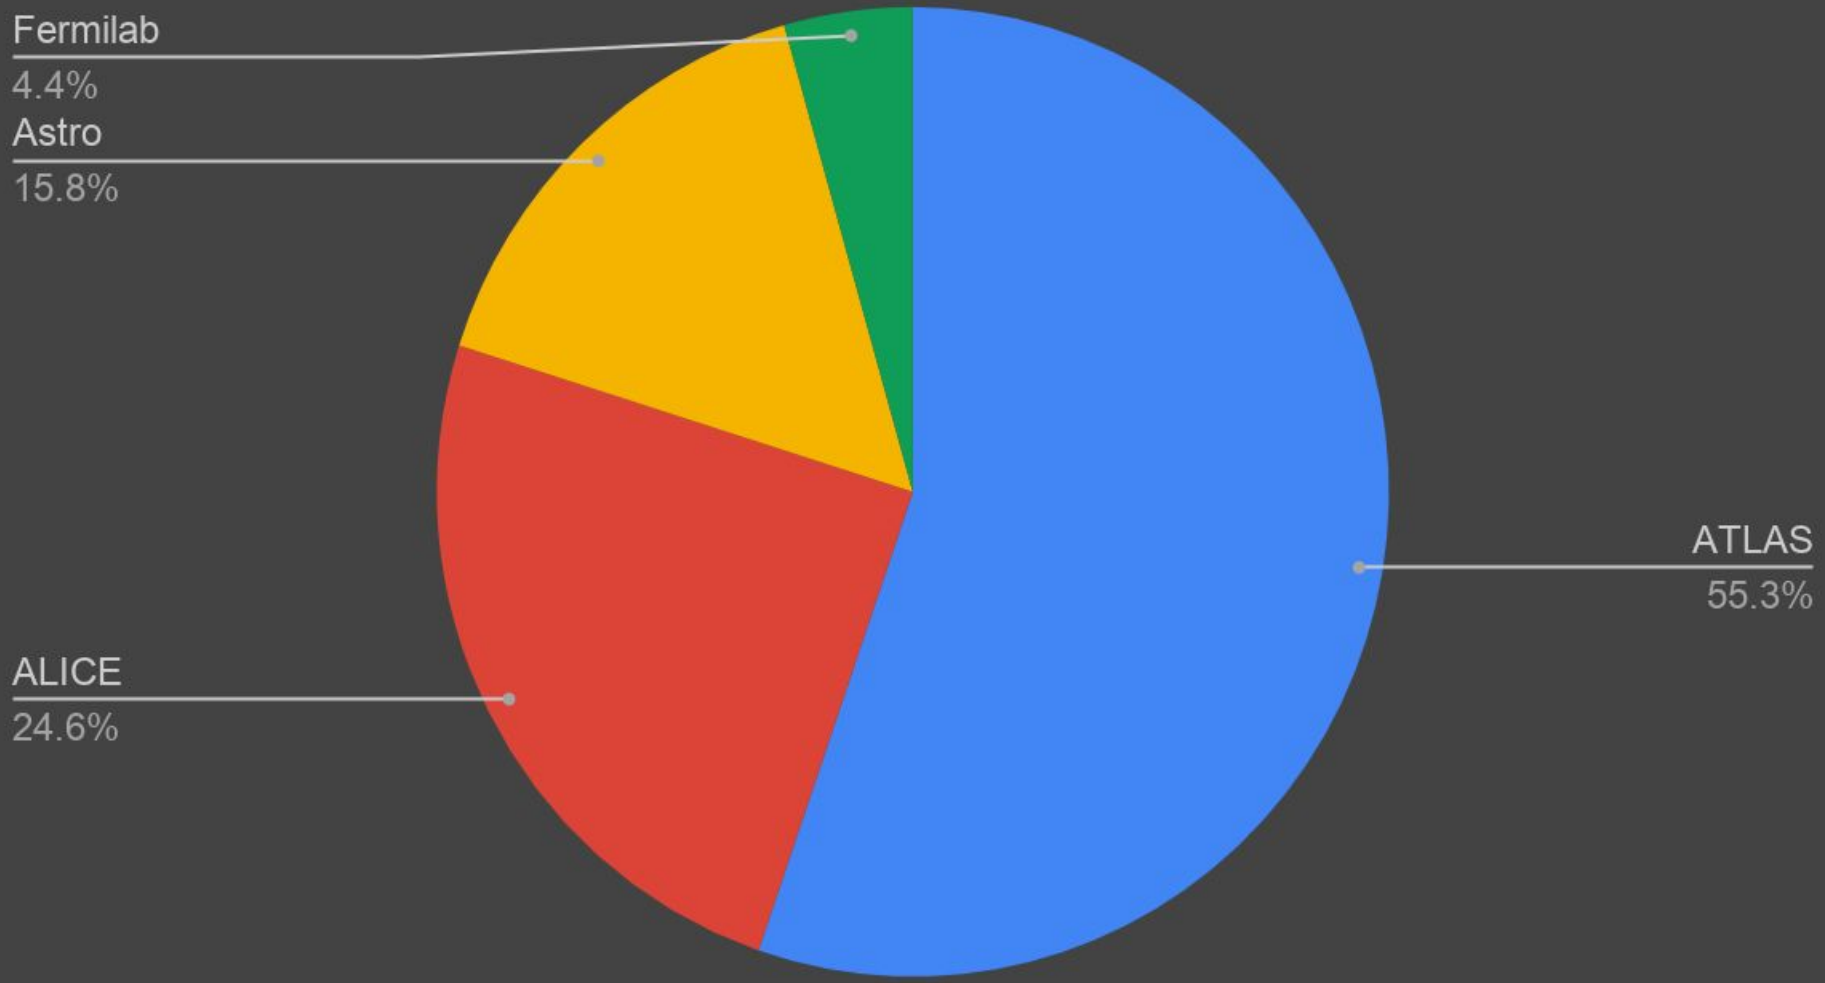
Fermilab
4.4%
Astro
15.8%
ATLAS
55.3%
ALICE
24.6%

Computing Center of Institute of Physics CAS

Distributed site located in Prague

Institutions involved:

- Institute of Physics CAS (IoP)
- Institute of Nuclear Physics CAS (NPI)
- Charles University (Faculty of Mathematics and Physics – FMP)
- IT4Innovations
- CESNET

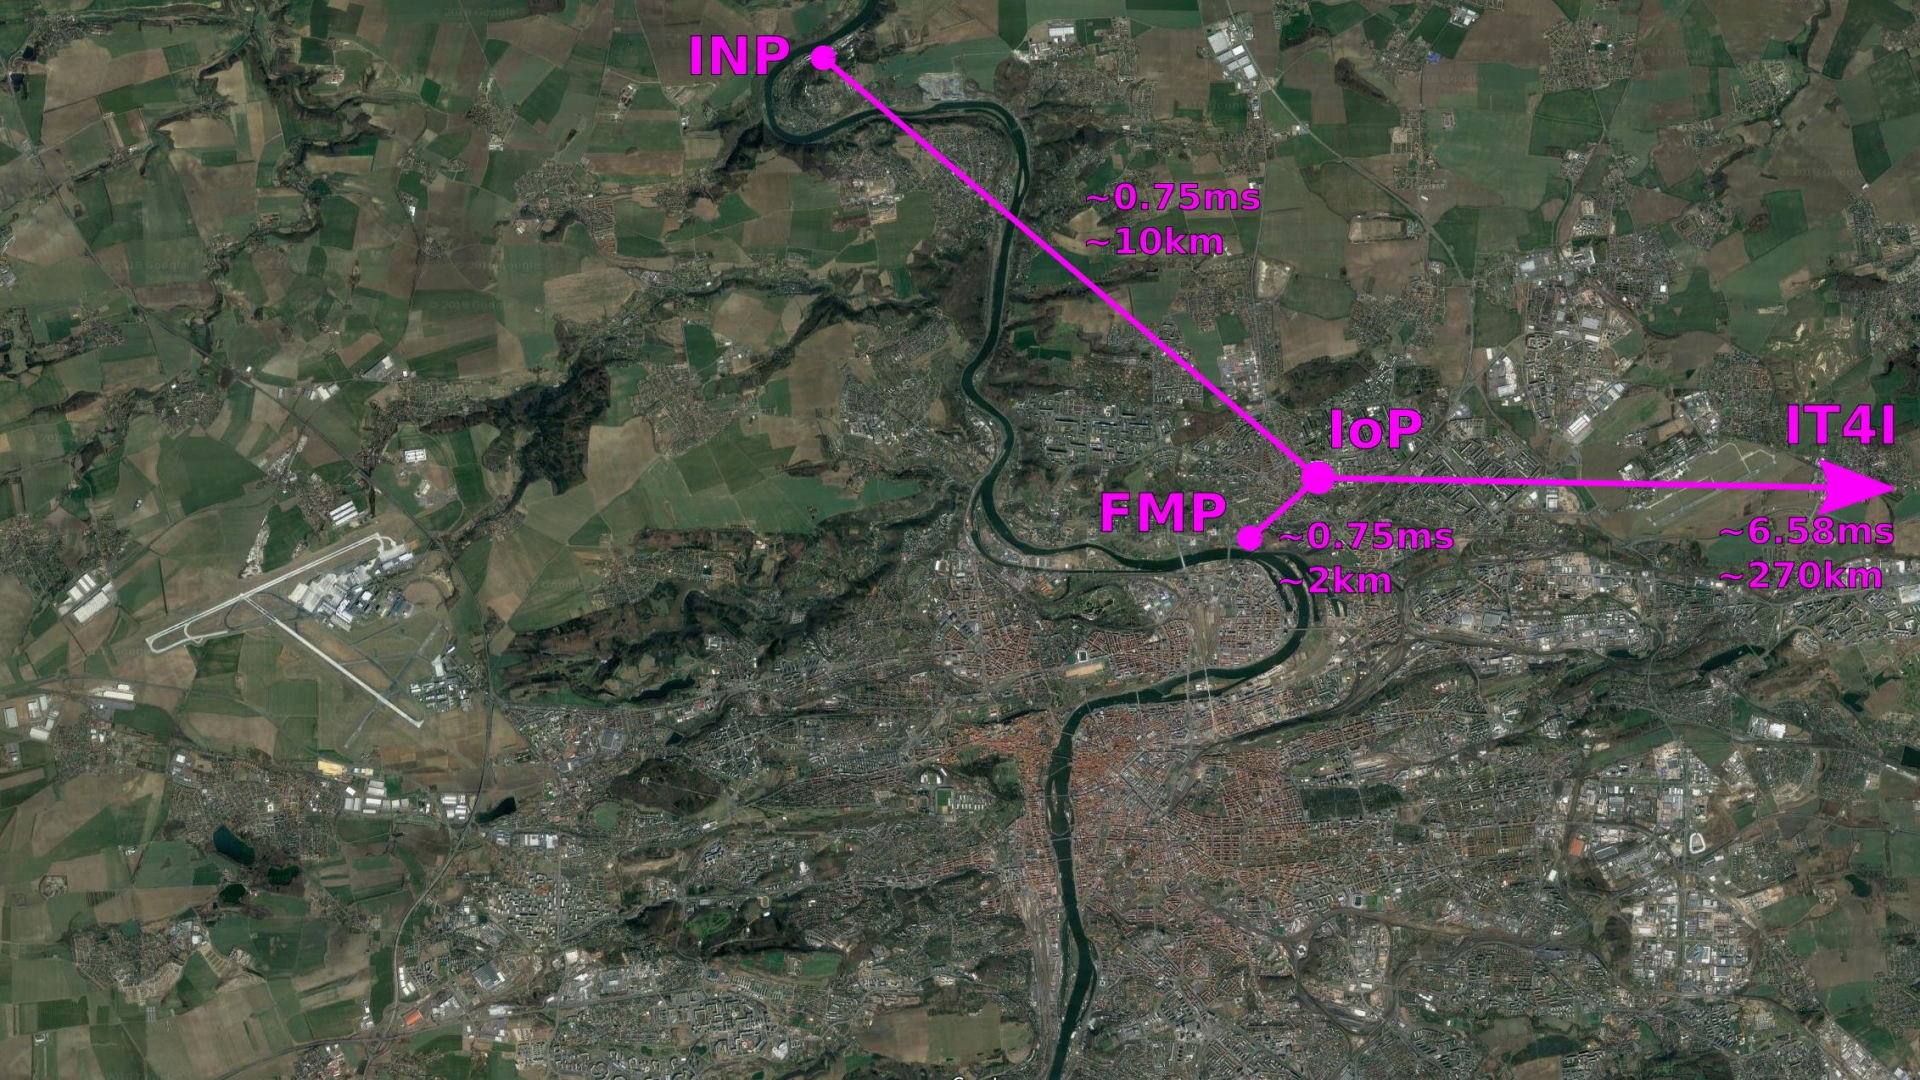
INP
~0.75ms
~10km
IoP
IT4I
FMP
~0.75ms
~2km
~6.58ms
~270km

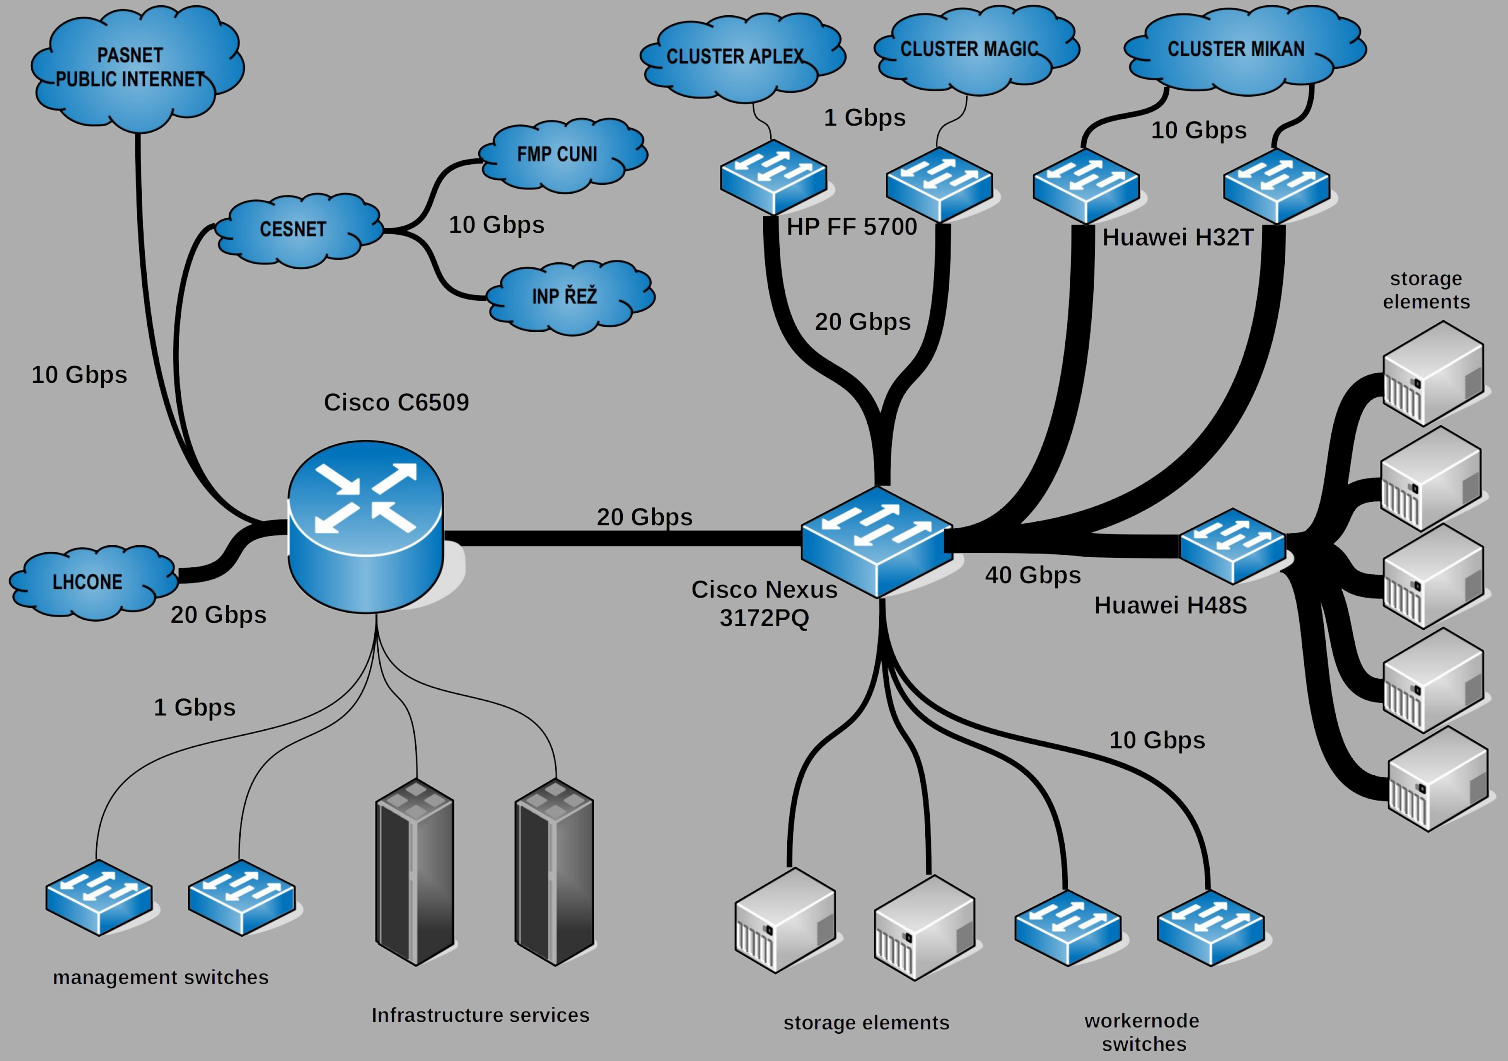
PASNET
PUBLIC INTERNET
CESNET
FMP CUNI
INP ŘEŽ
10 Gbps
CLUSTER APLEX
CLUSTER MAGIC
CLUSTER MIKAN
1 Gbps
10 Gbps
HP FF 5700
Huawei H32T
20 Gbps
storage
elements
10 Gbps
Cisco C6509
LHCONE
20 Gbps
20 Gbps
Cisco Nexus
3172PQ
40 Gbps
Huawei H48S
1 Gbps
10 Gbps
management switches
Infrastructure services
storage elements
workernode
switches

IPv6 status

Most services dual-stack

Internal mostly IPv6

Remote DPM traffic:
71 MBytes/s IPv6
131 MBytes/s IPv4

ALICE and Fermilab mostly IPv4

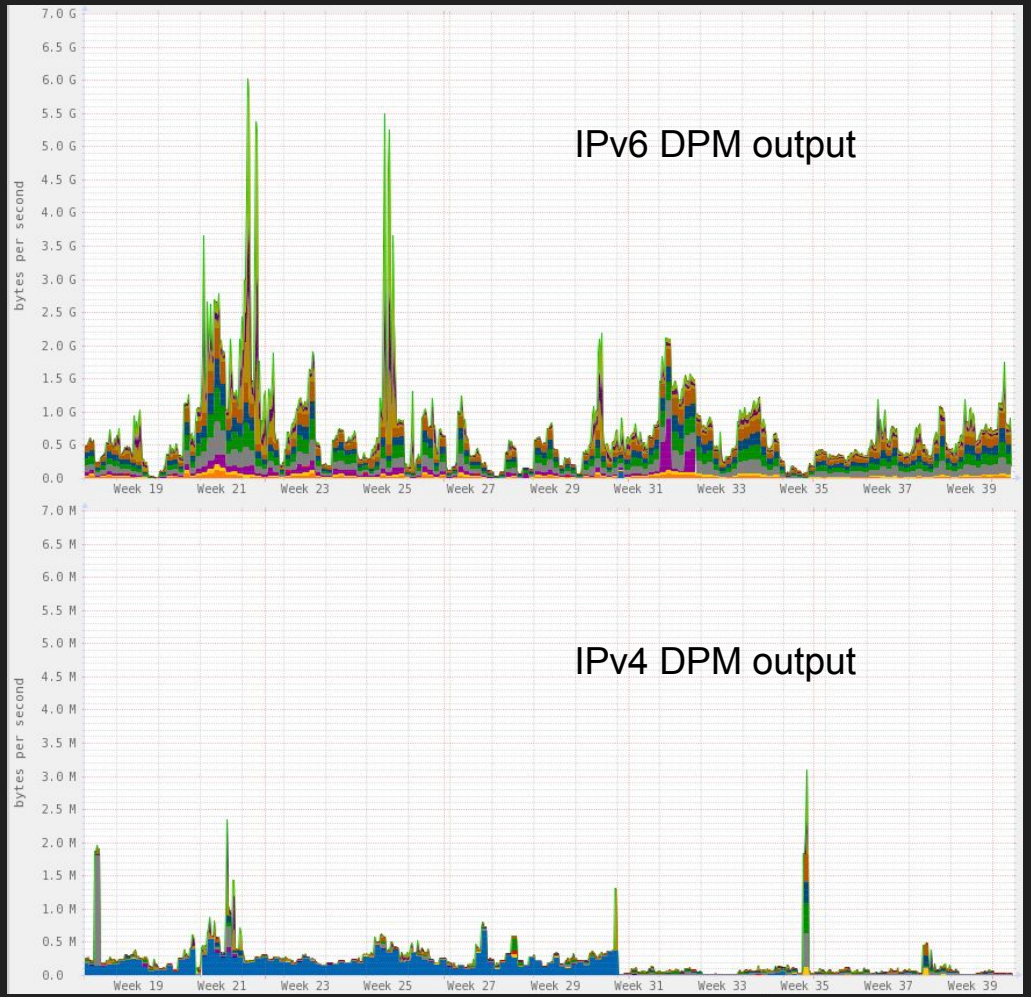

HT-Condor

ARC CE + HTC CE for OSG stack

Own accounting data processing for APEL

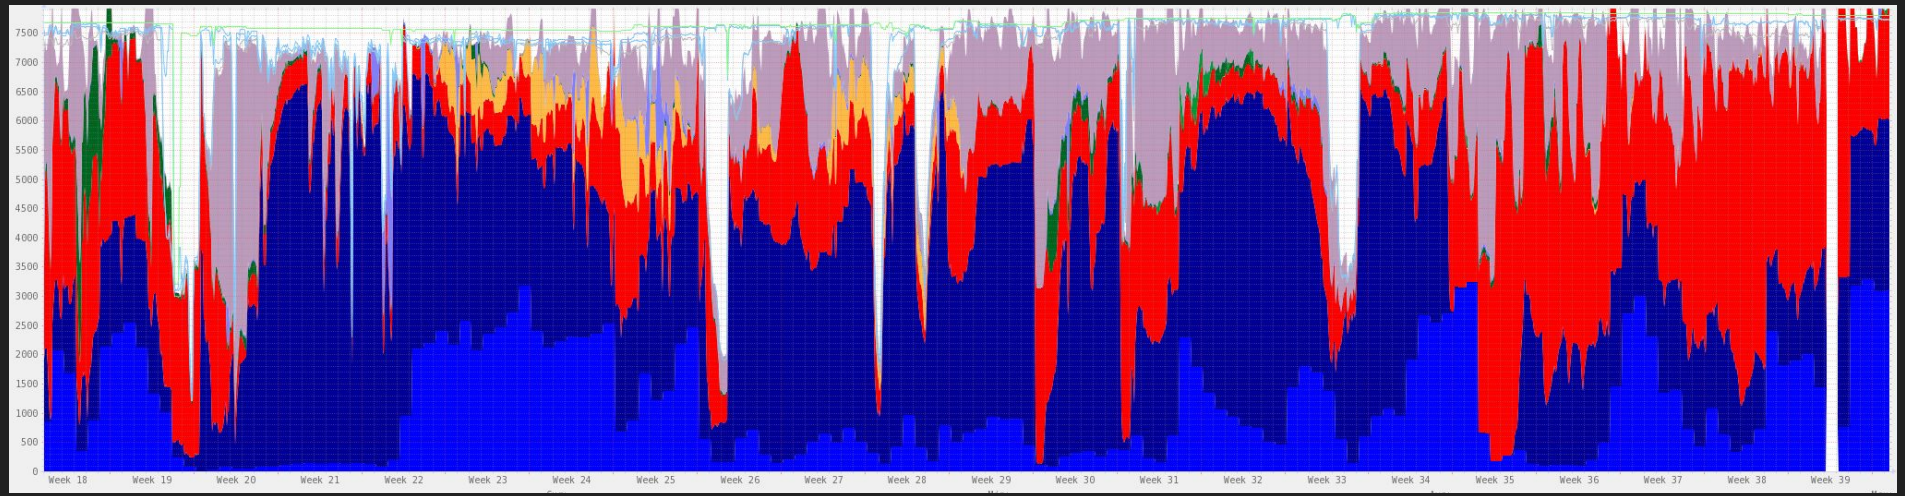

HT-Condor configuration (issues) – fairshare

Considering CPU performance

↳Cpu scaling from machine classAd to job classAd (also for accounting)

Confusing SLOT_WEIGHT and SCHEDD_SLOT_WEIGHT behavior

Hard to check correct behavior

HT-Condor configuration (issues) – fairshare

Group			
90% Hi-prio			
9% VO			
1% test			

HT-Condor configuration (issues) – fairshare

Group	VO		
90% Hi-prio			
9% VO	58% ATLAS		
	20% Astro		
	10% Fermilab		
	11% ALICE		
1% test			

HT-Condor configuration (issues) – fairshare

	VO	User/prod	
	58% ATLAS	50% User	
		50% Prod	
	20% Astro	50% User	
		50% Prod	
	10% Fermilab		
	11% ALICE		

HT-Condor configuration (issues) – fairshare

	VO	User/prod	Prod types
	58% ATLAS	50% User	
		50% Prod	80% multi-core
			20% single-core

HT-Condor configuration (issues) – fairshare

Group	VO	User/prod	Prod types
90% Hi-prio			
9% VO	58% ATLAS	50% User	
		50% Prod	80% multi-core
			20% single-core
	20% Astro	50% User	
		50% Prod	
	10% Fermilab		
	11% ALICE		
1% test			

HT-Condor configuration (issues) – limits

Default: no limit enforcement, requirements used only for scheduling

Ideal: limits work only when needed (WN running out of memory/disk)

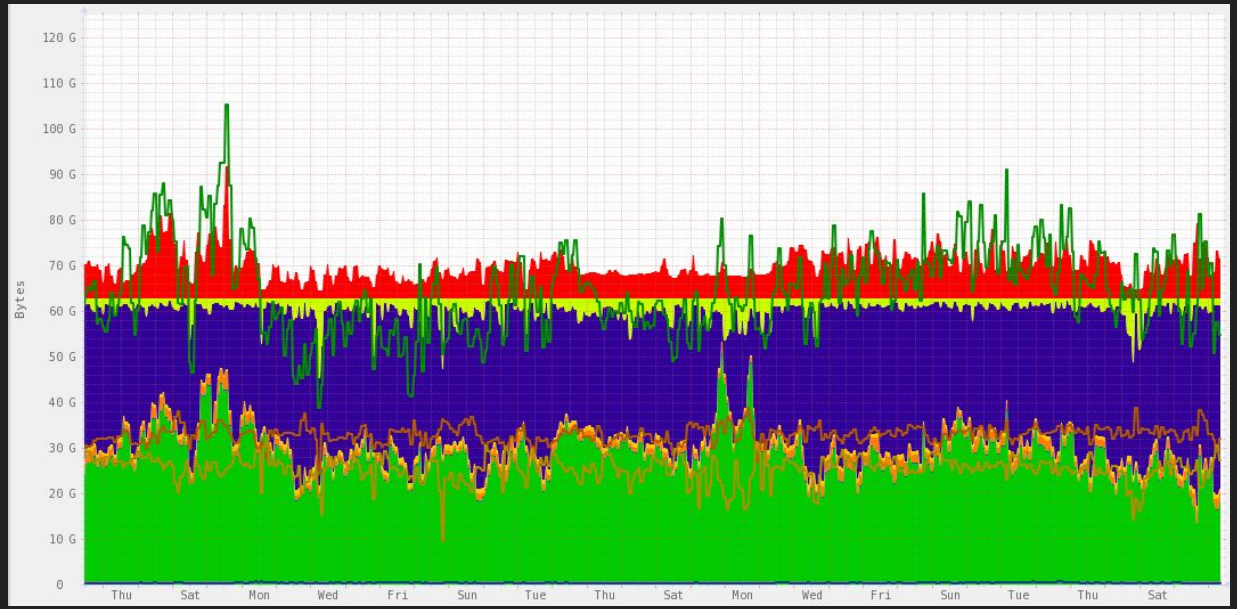

HT-Condor configuration (issues) – draining

Now using standard condor defrag daemon

+ when reboot draining: backfill with pre-emtible jobs

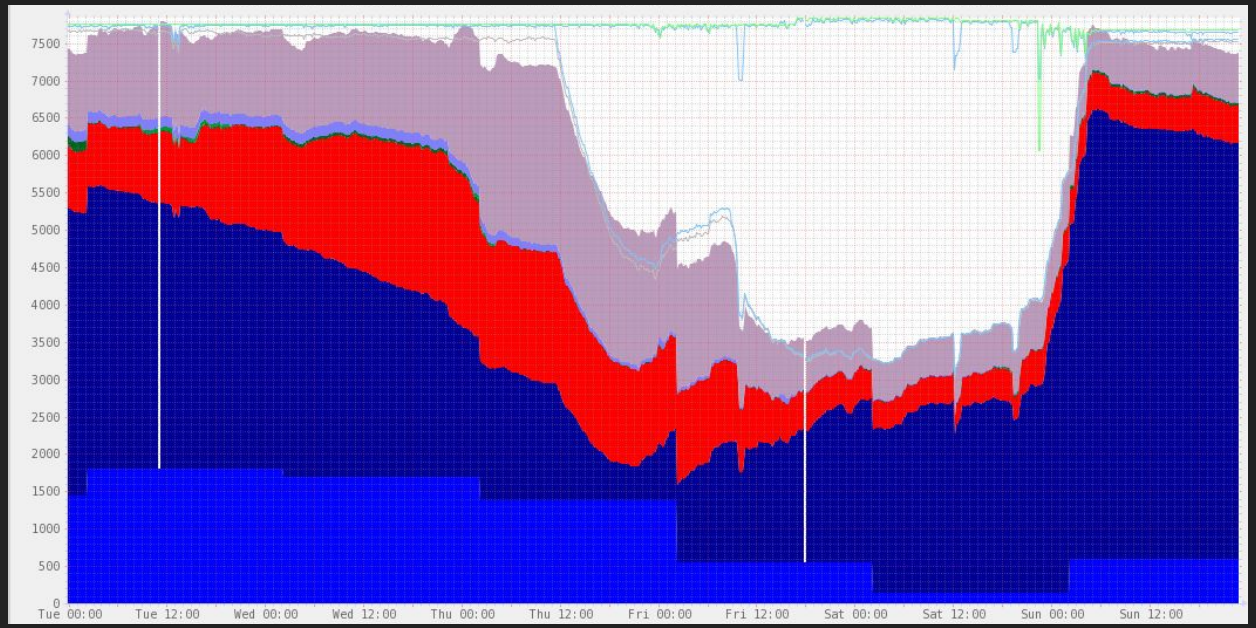

HT-Condor configuration (issues) – misc

- WN at CUNI only accepting some VOs:
  - Custom machine classAd 'FilterJobs = "x509UserProxyVOName =?= "Fermilab" ||...
- WN can be set offline:
  - On WN: classAd "StartJobs = False"
  - On negotiator: 'OFFLINE_NODES = "name1,name2,..."'
- Dedicated interactive jobs node
- Automatic accounting group for local user

Resources Extension – IT4I

15% of ATLAS walltime delivered by Salomon supercomputer in IT4I

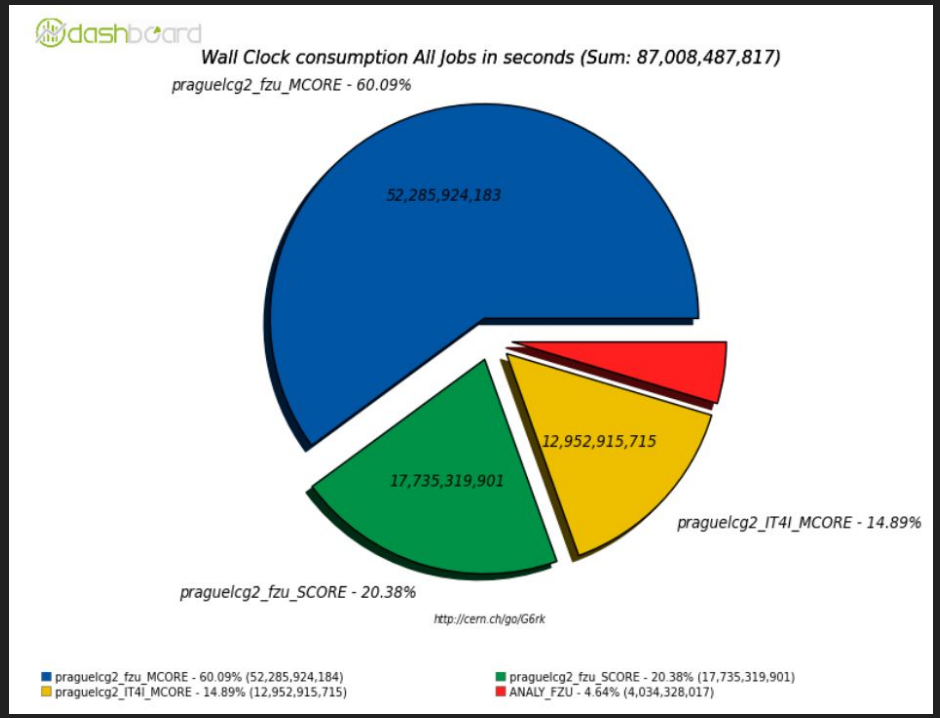

DPM legacy ⇒ DOME Upgrade

Reason: Performance & Bug fixes & Staying up to date

Operation	Protocol	Legacy [s]	DOME [s]
gfal-copy DPM ⇒ FILE	XROOT	1.71	0.37
	DAVS	1.72	0.38
	SRM	2.71	1.71
gfal-rm	XROOT	0.73	0.27
	DAVS	0.82	0.33
	SRM	0.97	0.93

Bugs in DPM DOME 1.10.3

- Host not known error on el7
- SRM/DOME reported space discrepancy
  - Own script for fixing the difference
- Many gfal-ls requests exhaust the MySQL handles
- dpm-xrootd does not create dir structure in the disk server
- HTTP authentication on dual stack
- The periodic cache cleanup may cause crashes

All bugs reported and fixed in 1.10.4 (released 1st Oct), kudos to P.Vokáč

DPM legacy ⇒ DOME Upgrade

Good upgrade guide

Update dmlite-mysql-dirspaces.py took 40mins and <10GB RAM

Problems synchronizing quota- and space- tokens

dpm-xrootd 3.6.6 in EPEL testing contains additional bug fixes

Also database migration and SL6⇒CC7 headnode upgrade

The whole operation needed one day of downtime

Foreman

Kickstarting, software installation + upgrades, remote execution
Foreman (1.19) + Katello plugin (3.8) + remote execution plugin (not yet installed)

Setup in progress now

Issues: known Katello bugs (certs, package search,...), not installing EPEL,...

Questions?

Thanks to my colleagues:

- From IoP: J.Chudoba, A.Mikula, P.Horák, J.Uhlířová, M.Svatoš
- From Czech Technical University: Petr Vokáč
- From Nuclear Physics Institute: Dagmar Adamová