Dynamic resource provisioning of the CMS online cluster using a cloud overlay

Marc DOBSON
CERN Geneva, Switzerland
for the CMS Data Acquisition group (DAQ)

CHEP 2016, 10th – 14th October
San Francisco – California, USA

The CMS (Compact Muon Solenoid) Detector

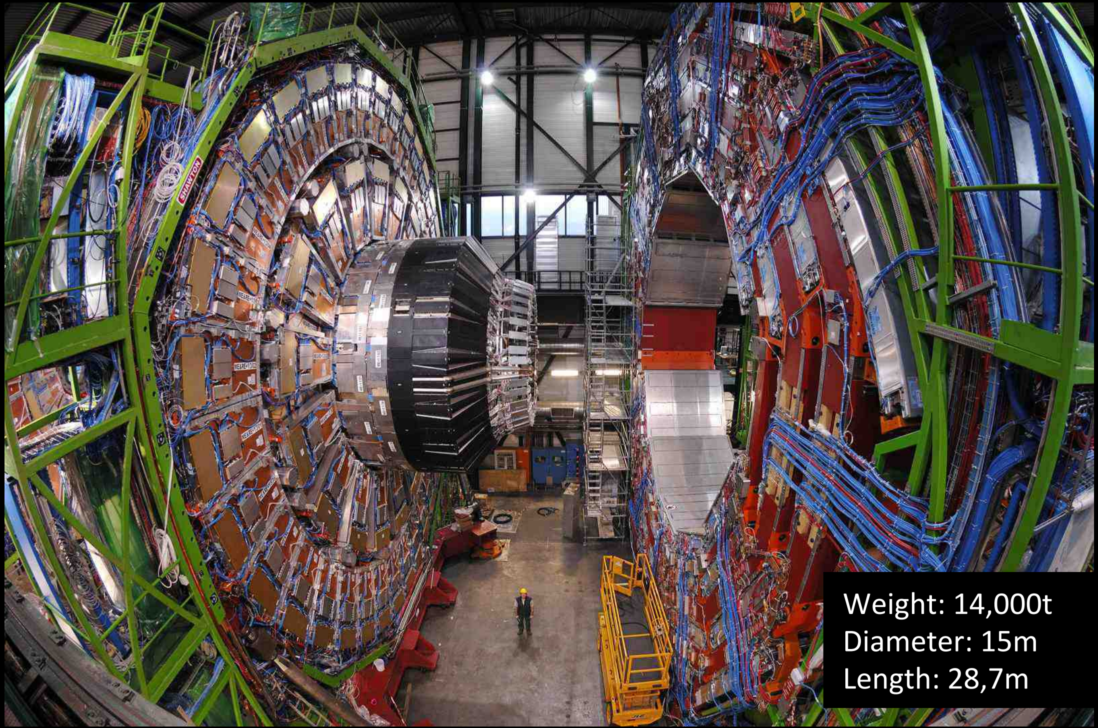

Weight: 14,000t
Diameter: 15m
Length: 28,7m

- 40 MHz collision rate
- Data from each collision is ~1 MegaByte
- First level of trigger (HW) reduces rate to 100kHz

Flow of data

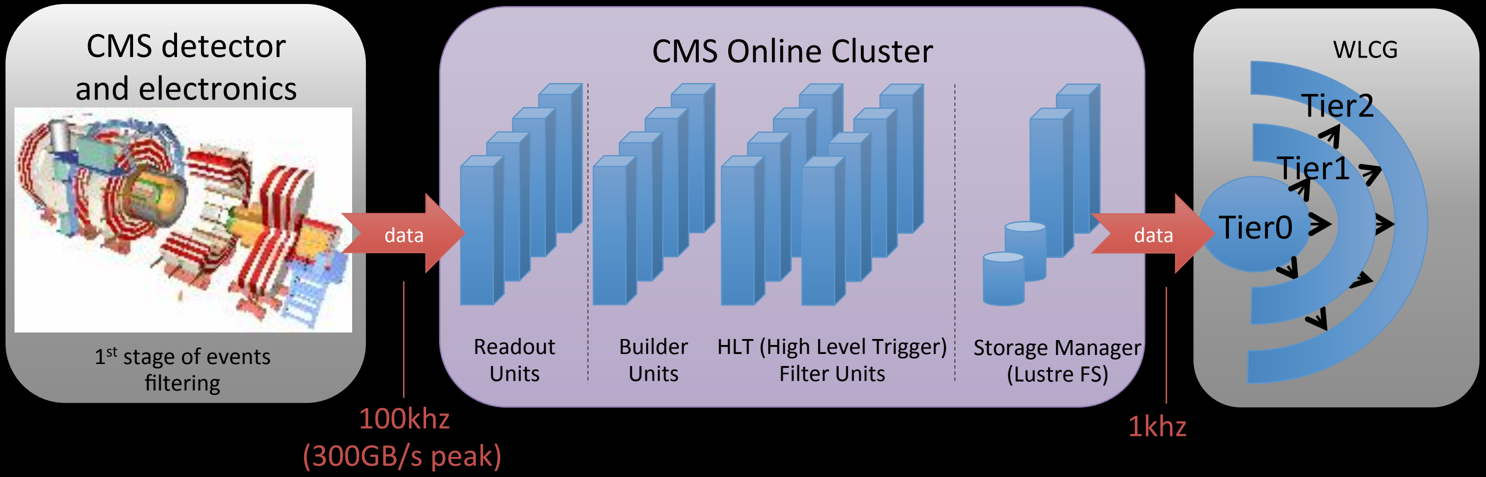

- The CMS Online Cluster is used to **filter interesting events** before shipping them to Tier0
- The data is then analysed offline by the Worldwide LHC Computing Grid (WLCG)

Cloud Purpose: Provide the computing power to the WLCG

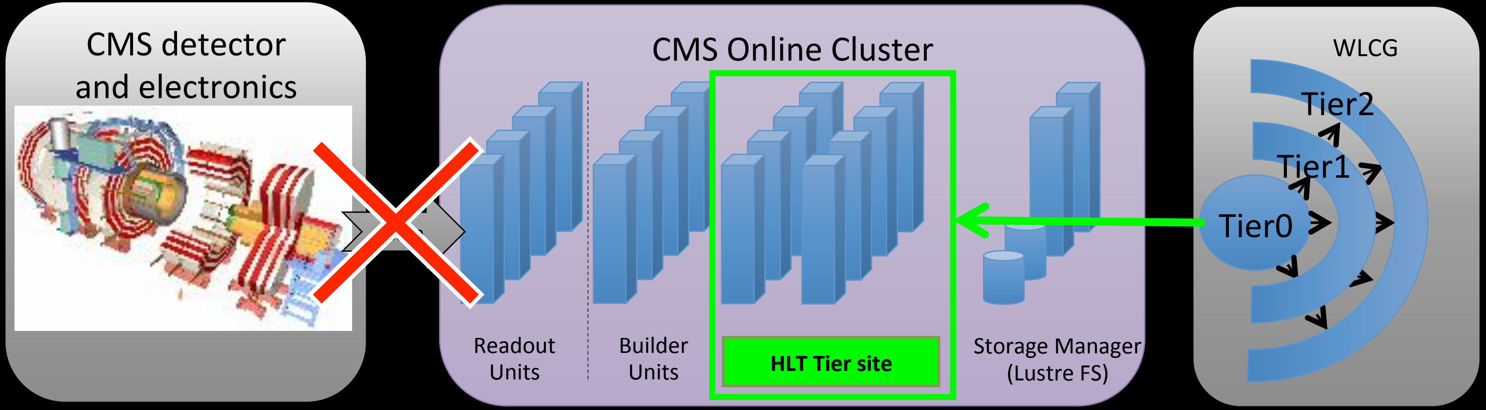

The cluster is idle during several weeks a year

Year	Technical stops	Machine development	Beam Commissioning	Total
2015	117 days	15 days	59 days	191 days
2016	112 days	21 days	27 days	160 days

CMS HLT computing power

Nodes	Cores	RAM GB	Disk TB
940	21808	51968	450

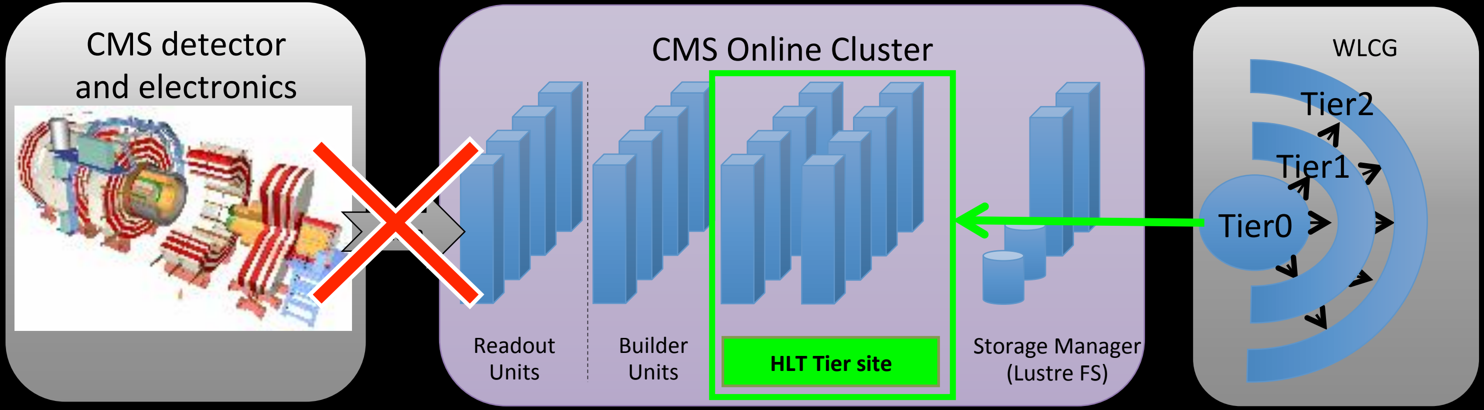

Farm size in 2016 (HEP-SPEC06)	**HLT Farm**	**HLT Cloud (68% HLT) + Cloud only**	**Tier0**	**All CMS Tier1 sites**
CMS	**500K**	**340K + 60K**	**315K**	**400K**

HLT Cloud architecture

Based on Openstack Grizzly

Only core service (nova, glance, keystone, APIs)

Corosync/Pacemaker

RabbitMQ

MySQL DB

Own mechanism for distributing image

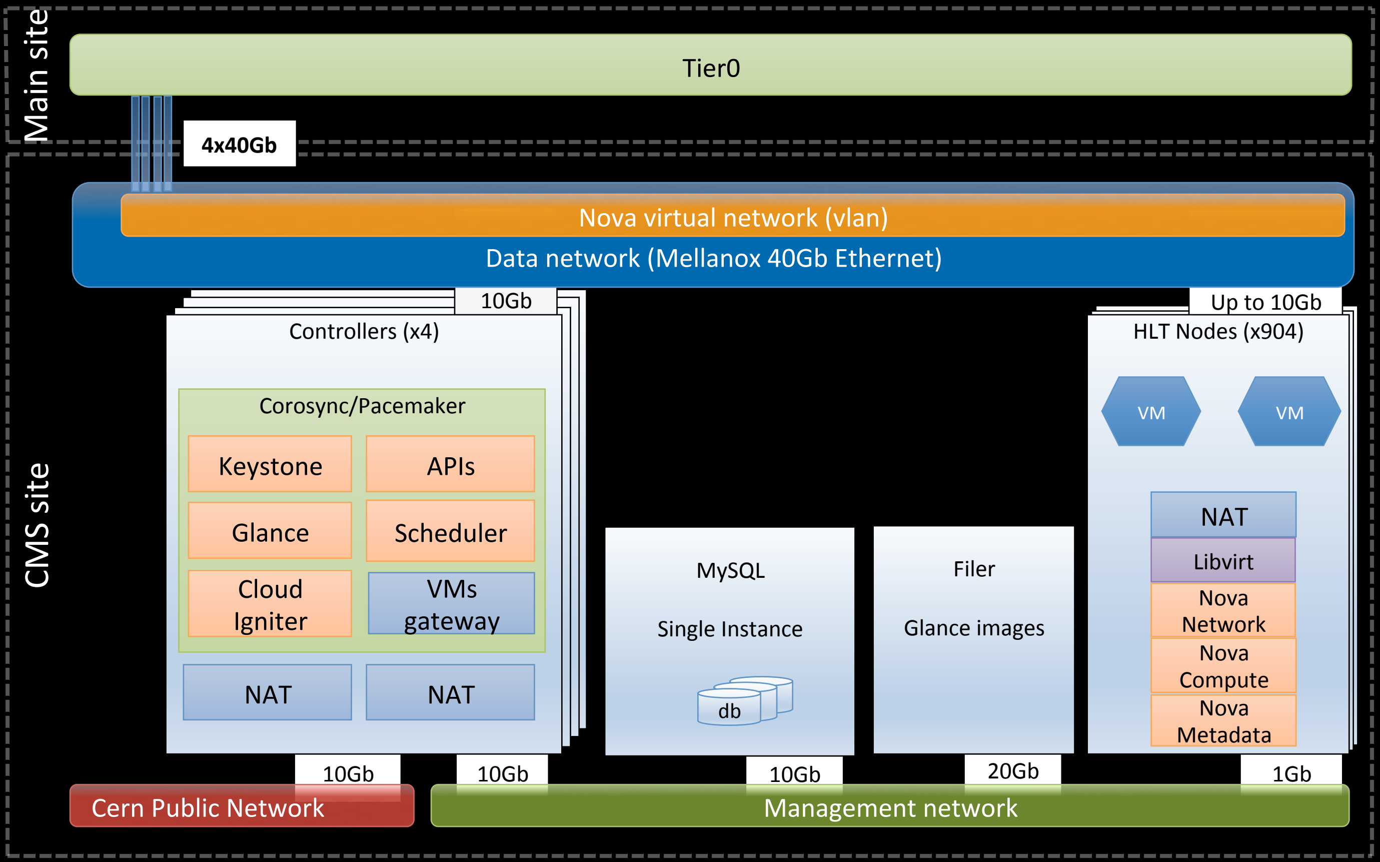

HLT Cloud specifications

HLT Compute nodes			Virtual machines		
Nodes	**Cores**	**RAM GB**	**VMs per node**	**VM Cores**	**VM RAM GB**
256 x dual E5-2670	16	32	2	8	14
360 x dual E5-2680v3	24	64	2	12	30
324 x dual E5-2680v4	28	64	2	14	30
Total: 940	**21808**	51968	**1880**	21808	48208

Up to 68% of the HLT farm is allocated to cloud at any one time.
Retain sufficient power for cosmic data taking and tests.
Currently only 54% are used (for operational reasons).

Typical WLCG job specification:
- Multi core
- 7-8 GB of memory
- 4-6 hours length

2016: 3K cores baseline Cloud only nodes (retired HLT nodes)

HLT Cloud is classed as Tier2 site (no local data storage): data is retrieved from CERN

Jobs running in the Cloud with ~54% of the HLT farm

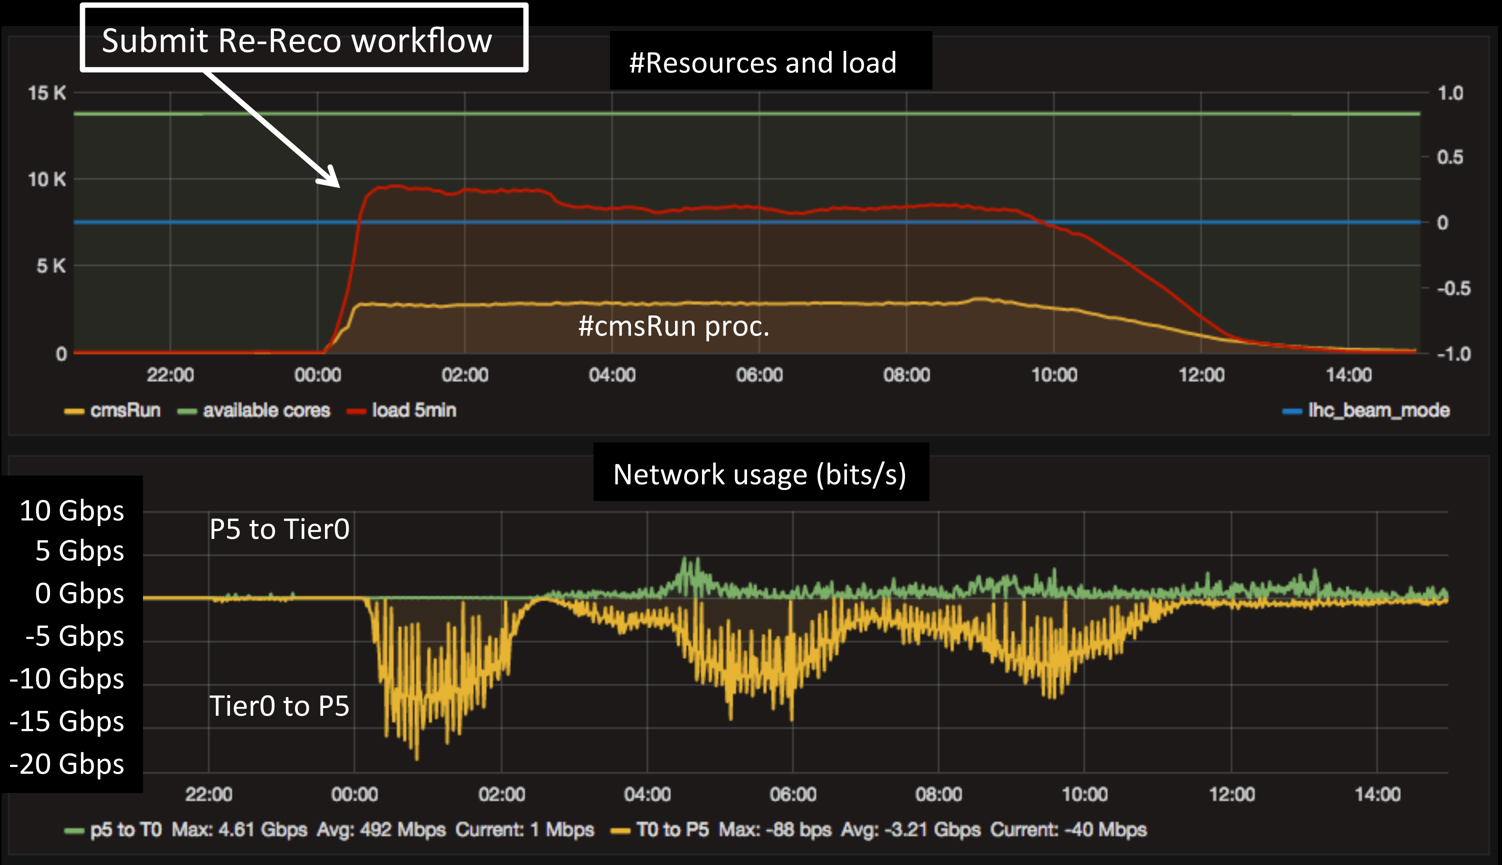

All cores fully used (multi-core cmsRun processes).

High speed network links to Tier0 used to get data and software for jobs.

Cloud network usage in reverse direction from data taking.

Link of 4x40Gb/s, so even network intensive jobs can run in the HLT Cloud.

HLT Cloud participation vs Tier1 sites

(Ramp up from ~5% to ~54% of the HLT farm used)

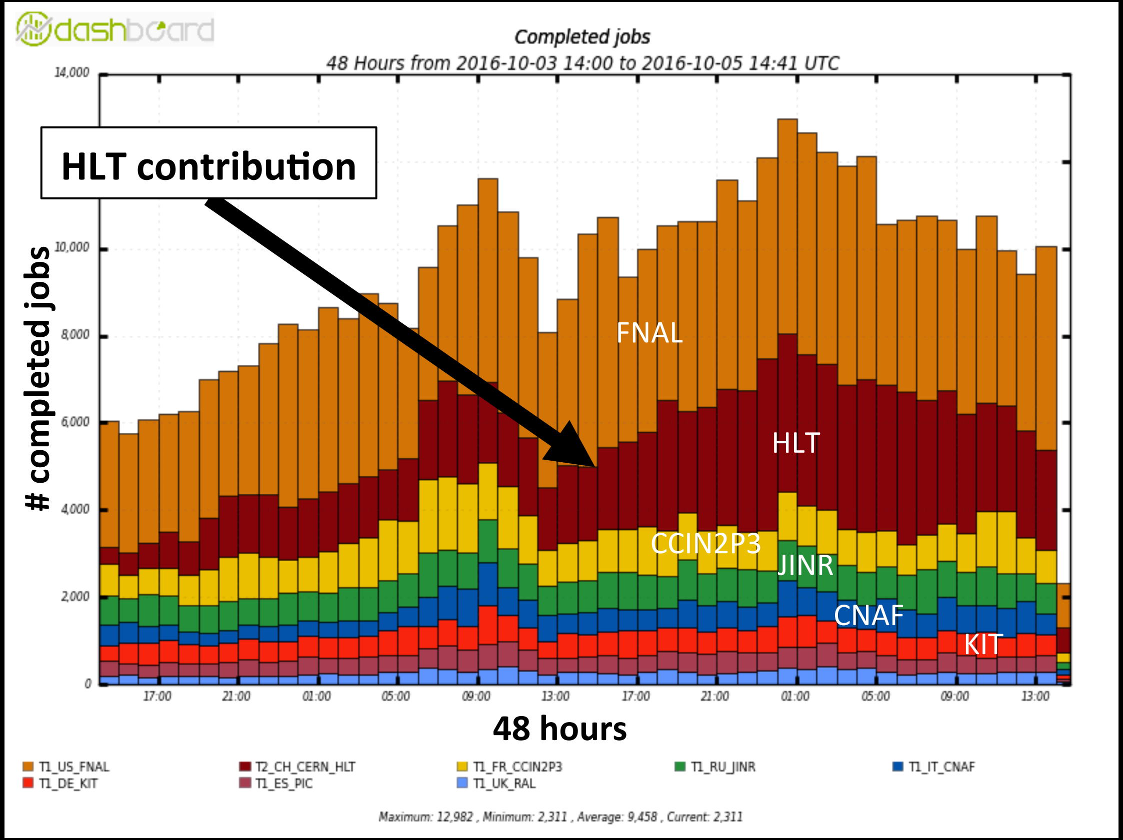

Challenge for 2016

Inter-fill periods (3h to 6h in average)

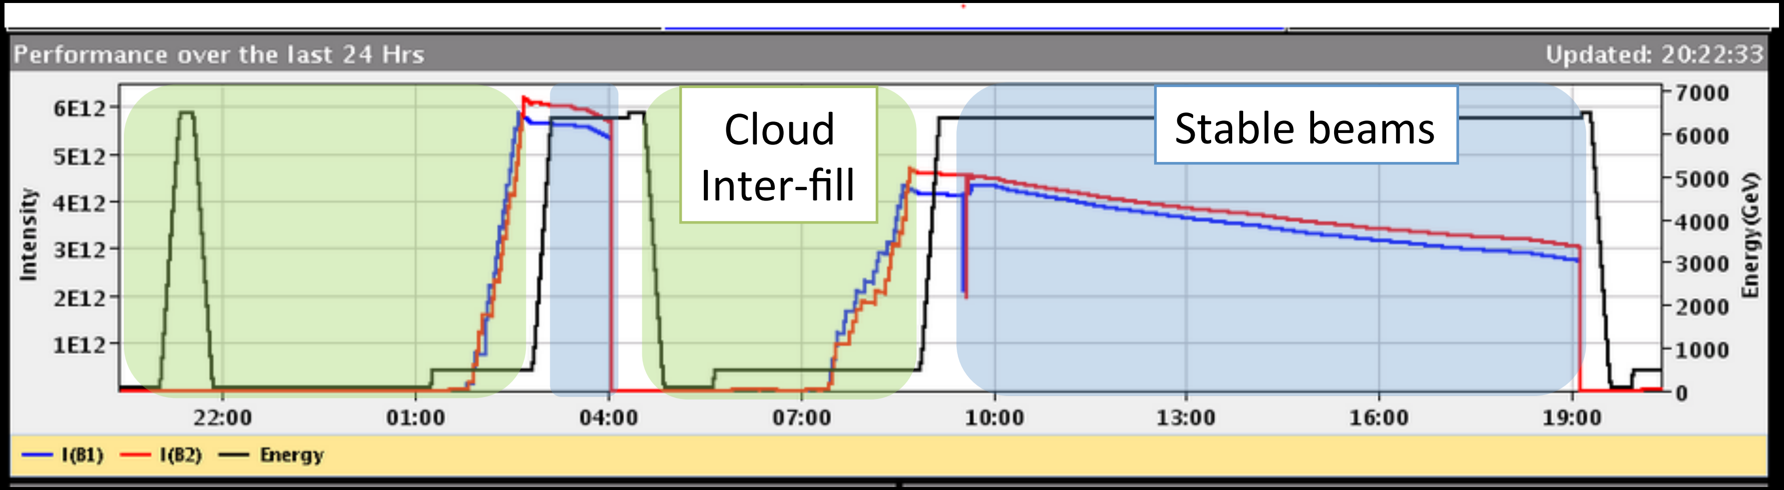

- Stable beams: beam conditions for data taking
- Cloud inter-fill = Period between ramp down and ramp up
  - No particle collisions
  - Gap with Stable Beams = safety margin

Inter-fill periods in number of days

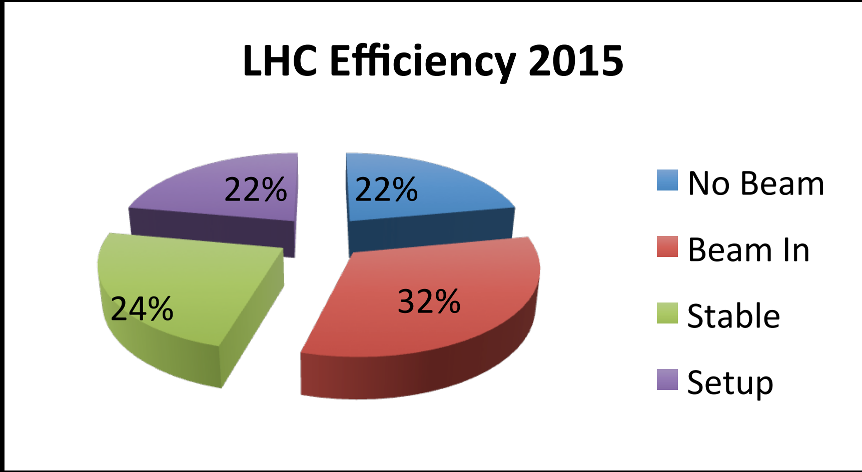

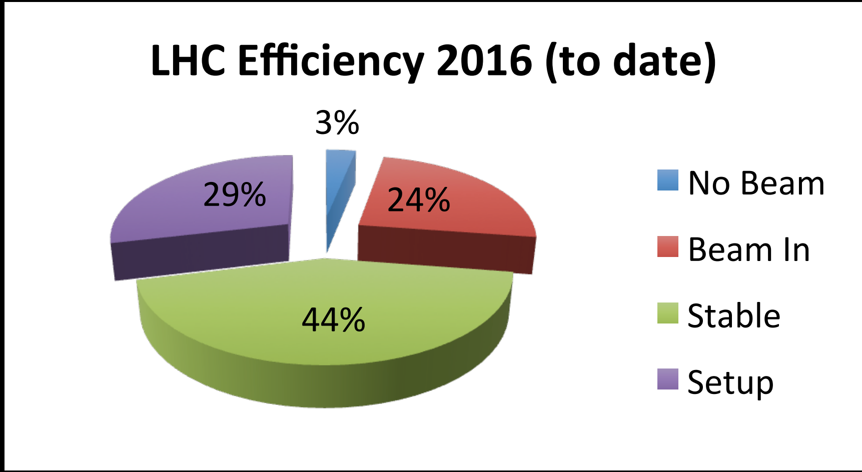

Covers 1st beam to last beam of year

No Beam = Access, long stops
Beam In = All beam modes but SB
Stable = Stable Beams
Setup = Setup, Precycle, Ramp-down

Period	Number of days (2015)	Number of days (2016 estimated)
maintenance periods	191	160
Inter-fill*	(365-191)x0.76 = **132**	(365-160)x0.56 = **115**

* Ignoring safety buffer explained in previous slide

Inter-fill periods: How To

Inter-fill: Unpredictable and short periods of few hours

Detect inter-fill period using LHC states (from Ramp Down to Ramp UP).
Allocate up to 68% of HLT to cloud (leave some for tests and cosmic data taking).

Up to mid-2016 used to:

1. Start/Stop the Cloud
2. Grid jobs with tailored length for better efficiency
3. Kill jobs at end of inter-fill

Summer 2016, increased efficiency:

1. Cloud/jobs are suspended/resumed
2. Jobs no longer tailored in length
3. 95-99% of jobs finish successfully

Inter-fill experience over the last few months

68% HLT ramp test
Validating Infrastructure handles the load.
+ suspend/resume VMs

Baseline 27/7
+ **120** HLT compute nodes during real **daytime inter-fills**

Baseline 24/7
+ **54%** most recent HLT compute nodes during real **daytime inter-fills**

Baseline 24/7
+ **54%** most recent HLT compute nodes during **ALL inter-fills**

in production:
Baseline 24/7
+ **1k cores** during **ALL inter-fills**

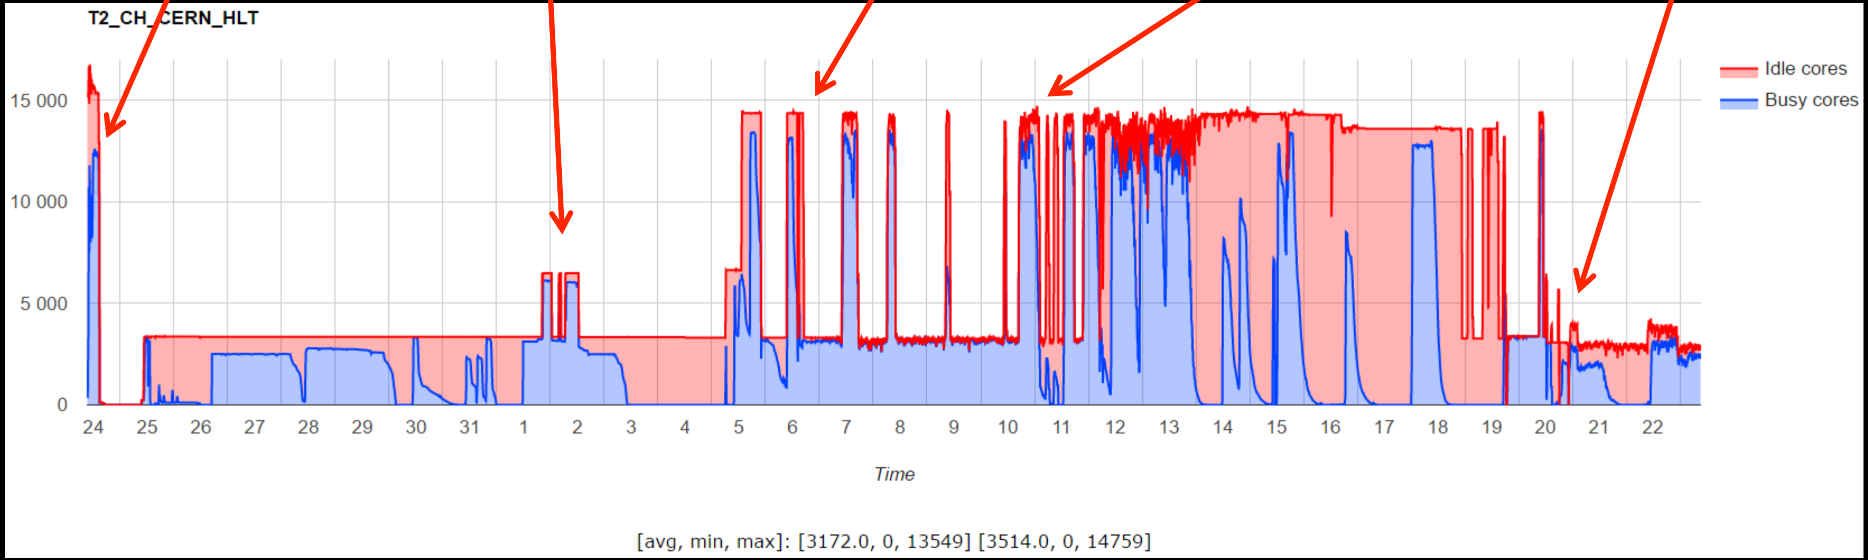

Inter-fill experience: production

in production: Baseline 24/7 + **1k cores** during **ALL inter-fills**

in production: Baseline 24/7 + **3k cores** during **ALL inter-fills**

in production: Baseline 24/7 + **6k cores** during **ALL inter-fills**

in production: Nominal System 3k + 11k cores **ALL inter-fills**

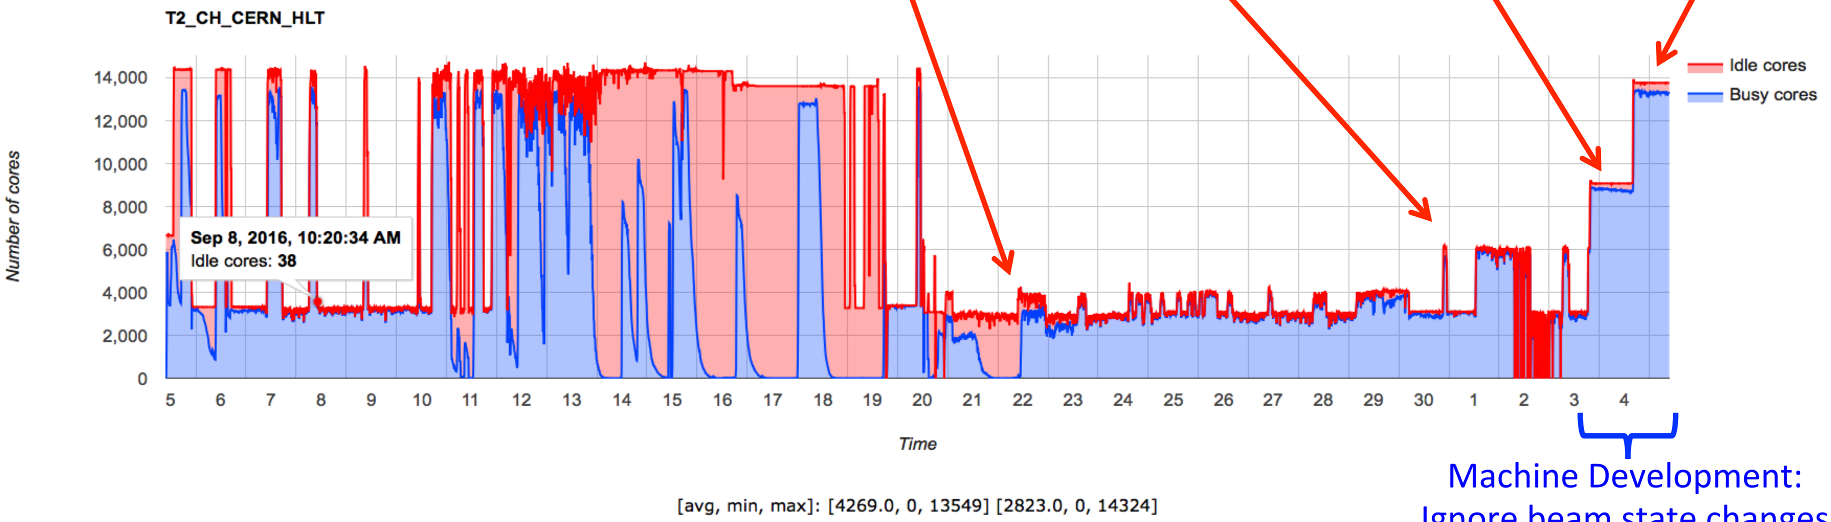

Machine Development: Ignore beam state changes

Inter-fill effect on HLT contribution

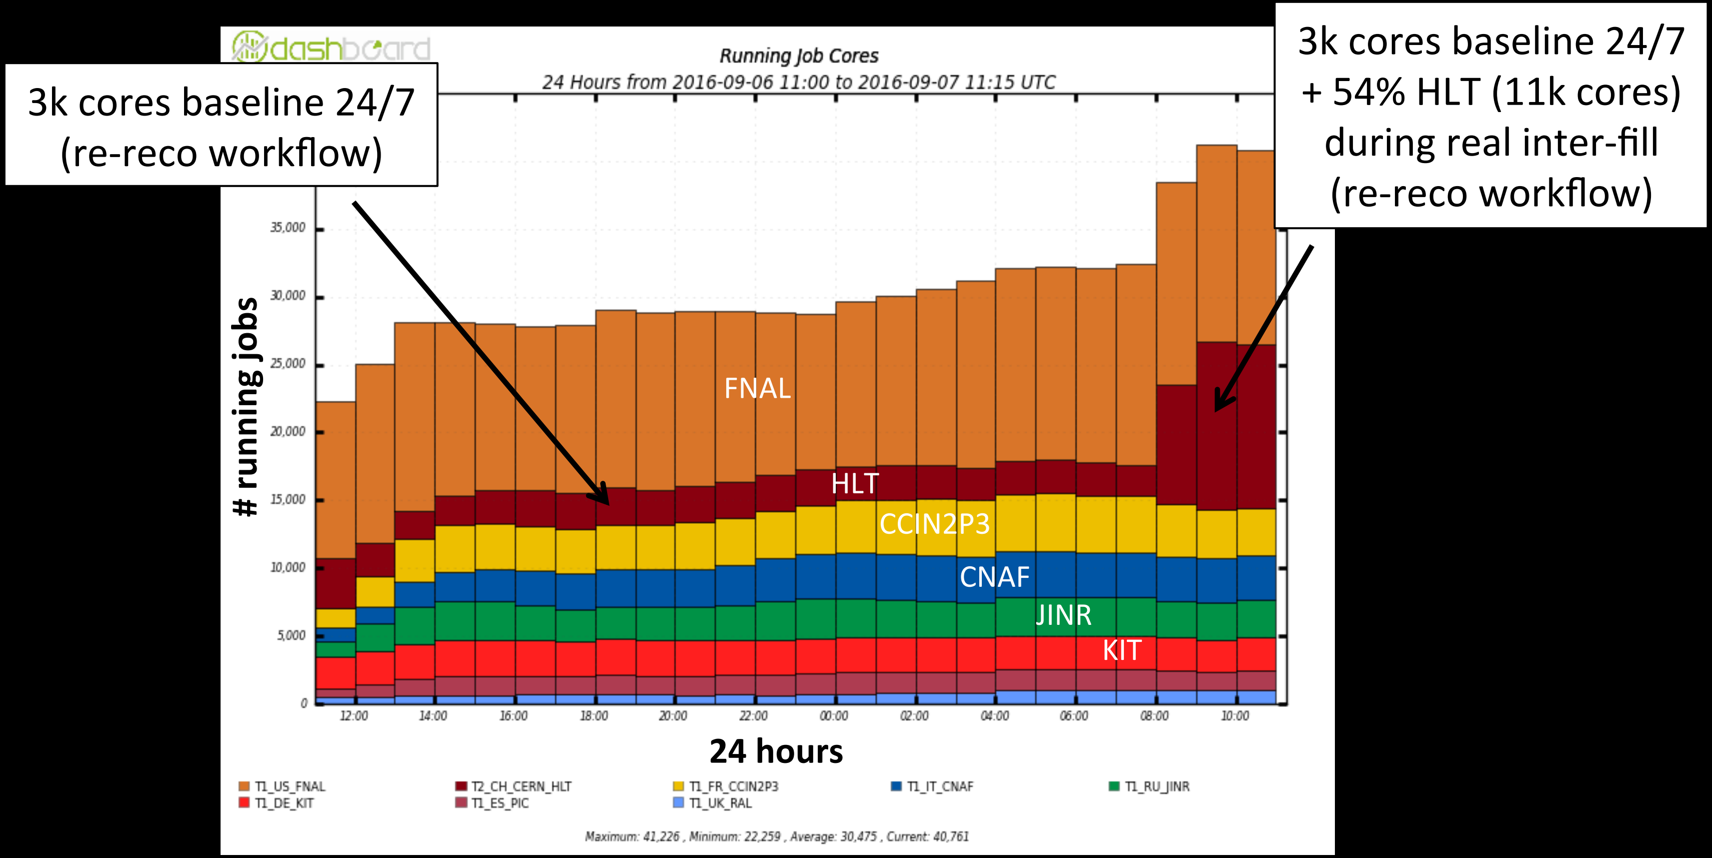

Inter-fill Cloud: the gains

- Only inefficiency is time to Suspend and Resume jobs
  - No plot as granularity of monitoring plots is 5min
  - Suspend < 12min, resume < 6min
  - Typical inter-fill ~3h => 10% inefficiency
  - Only consider 65% of HLT, not baseline 3k cores over whole year

Cloud mode	Cloud days (2016)	Efficiency	Total CPU Capacity (cores)	Average CPU Capacity over year (cores)	Capacity per year (kHS)
Cooling Maintenance	18 days	0%	0	0	0
Maintenance periods	142 days	100%	2112k	~5.8k	~130
Inter-fill periods	115 days	90%	1540k	~4.2k	~96
Total	**257 days**		**3652k**	**~10k**	**~226** (out of 340)
Baseline	347 days	100%	1299k	~3k	60

Outlook: start up cloud during fill !

Extend the HLT Cloud periods by starting as soon as the rate of data taken from the detector and the computing power needed decrease.

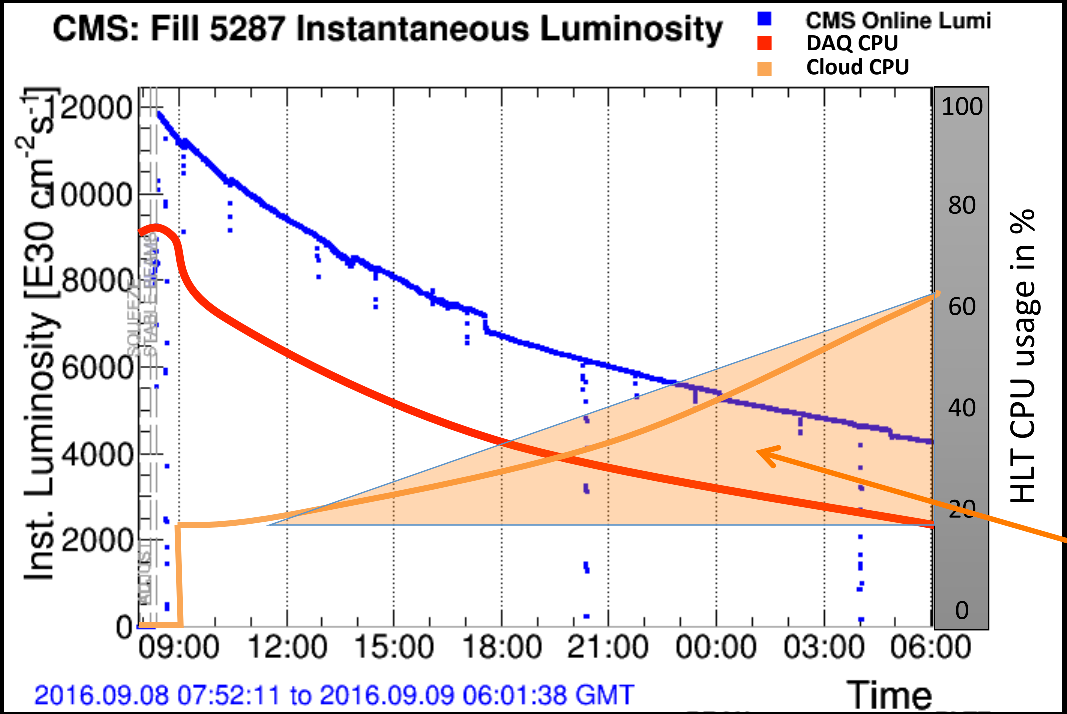

Very rough estimate of CPU allocated to the cloud according to the instantaneous luminosity (keeping 20% spare CPU for DAQ as a safe margin for glitches in luminosity or trigger rate, and hardware failures)

Integrated CPU power available:
25% of HLT Cloud CPU useable

Equivalent to doubling of 24/7 cloud
3-4k cores or 80kHS
For a total of 300kHS out of max 340kHS !

Conclusion

- Multi-role cluster
  - Data taking
  - Opportunistic Tier site (Maintenance or Inter-fill periods)
- HLT Cluster computing resources usage is optimized (226kHS used out of 340kHS)
- First Cloud resource to go production!
- Work ongoing to extend the HLT Cloud periods by starting as soon as the rate of data taken from the detector and the computing power needed decrease => could gain another ~80kHS

Thank you.

CMS Online Cluster

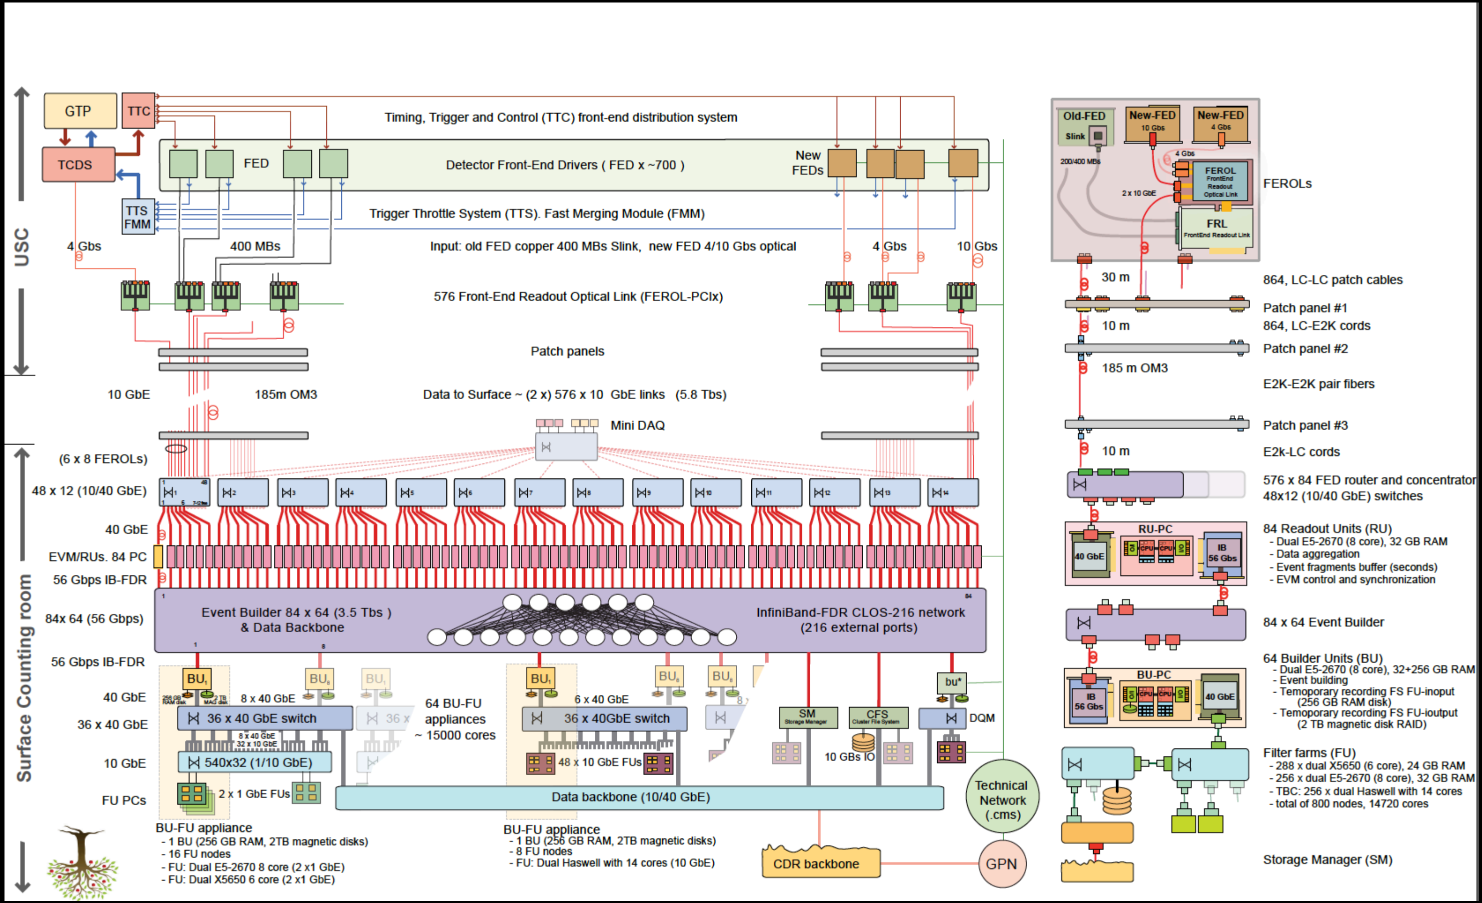

LHC cycles

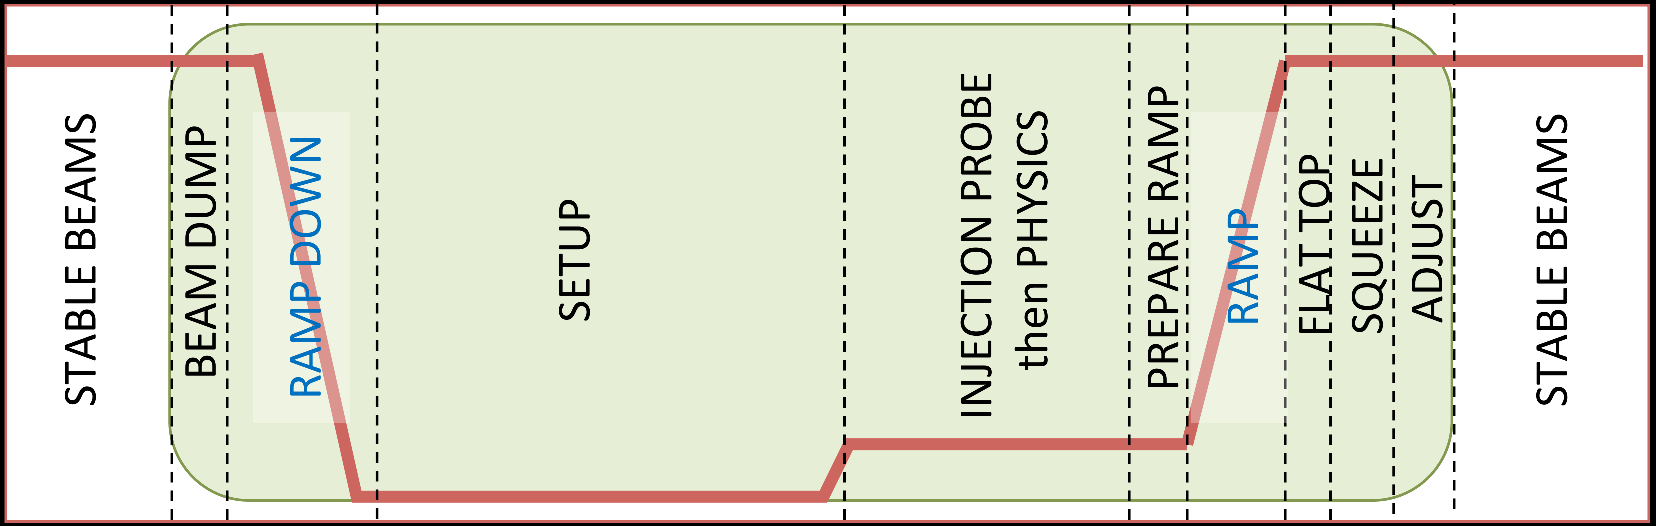
STABLE BEAMS
BEAM DUMP
RAMP DOWN
SETUP
INJECTION PROBE then PHYSICS
PREPARE RAMP
RAMP
FLAT TOP
SQUEEZE
ADJUST
STABLE BEAMS