



Tata Institute of Fundamental Research **टाटा मूलभूत अनुसंधान संस्थान**

ATC5 Meeting, TIFR Mumbai
Oct 24 to 26, 2019

Status Report of CMS Tier2 at TIFR

[Puneet Kumar Patel](#), [Brij Kishor Jashal](#), [Kajari Mazumdar](#), [Gobinda Majumder](#)

Agenda

- Overview & Resources
- TIFR T2 site infrastructure
- Networking
- TIFR HEP Cloud
- Upgradation and Future plan

Tier 2 LHC grid computing center at TIFR

- TIFR has played lead role for Indian participation in CMS experiment at LHC: Hardware, Computing, Physics
- National contribution to CMS international collaboration at LHC
- Indian funding agencies DAE and DST
- Average availability and reliability $\sim 95\%$
- One of the largest CMS T2 centre, active since 2008
- Crucial contribution for Higgs discovery in 2012
- 24x7 operations... sites gets red marked if tests fall below 80%
- Multiple tests and probes monitoring the health of system
- Multi agency monitoring and enforcement with GGUS ticketing system

Tier 2 LHC grid computing center at TIFR (cont.)

- Two commissioned sites for CMS central computing (Production and analysis)
 - **T2_IN_TIFR**
 - Part of the global CMS resources.
 - ~ 100 active users from collaborating Indian Institutes
 - 20% of storage for Indian community
 - Dynamic resources site – **T3_IN_TIFRCloud** since 2017
 - Dedicated resources for T3 for analysis of Indian users

Evolution of grid computing centre at TIFR

Year	CPU: No of HT cores	Pledge: CPU HEP-Spec06	Storage pledge (TB)	Power utilization
2009	150	2400 (2%)	450 (8%)	35 KVA
2010	240	2800 (2%)	600 (7%)	35 KVA
2011	275	3000 (3%)	700 (4%)	35 KVA
2012	300	3200 (1%)	750 (3%)	35 KVA
2013	300	3800 (1%)	900 (~3%)	35 KVA
2014	600	6500 (2%)	940 (~3%)	35 KVA
2015	650	7100 (~1%)	970 (~3%)	35 KVA
2016	1200	12,288 (~2%)	1980 (~3%)	40 KVA
2017	1400	20,000 (~2%)	2000 (~3%)	50 KVA
2018	2500	25,000 (~2%)	3000 (~3%)	60 KVA
2019	3500	35,000 (~3%)	5000 (~3%)	80 KVA
2020	9500	95,000 (~3%)	9000 (~3%)	120 KVA

Nos. in brackets are in fraction of CMS total resources

Grid computing at TIFR : Evolution of Network

- **Major force behind the development of NKN and Indian R&E Network.**
- 1 G dedicated P2P link from TIFR ↔ CERN (2009)
- Upgraded to 2G in 2012 → Upgraded to 4G in 2014
- Implemented fall-back path using 10G shared TEIN link to Amsterdam. (2015)
- CERN P2P link Upgraded to 8G (2015)
- Implemented LHCONE peering and L3VRF over NKN, all collaborating Indian institutes. (2015-2016)
- Upgraded to full 10G dedicated circuit till CERN (2017)
- NKN implemented CERN PoP with 10G link (2018)
- At present (10G + 10G) active links to LHC network
- TIFR first institute to have 10G end point
- Dedicated L3 peering to US West coast via Singapore and Amsterdam
- **Network for Run-III => ~40 G International circuit**

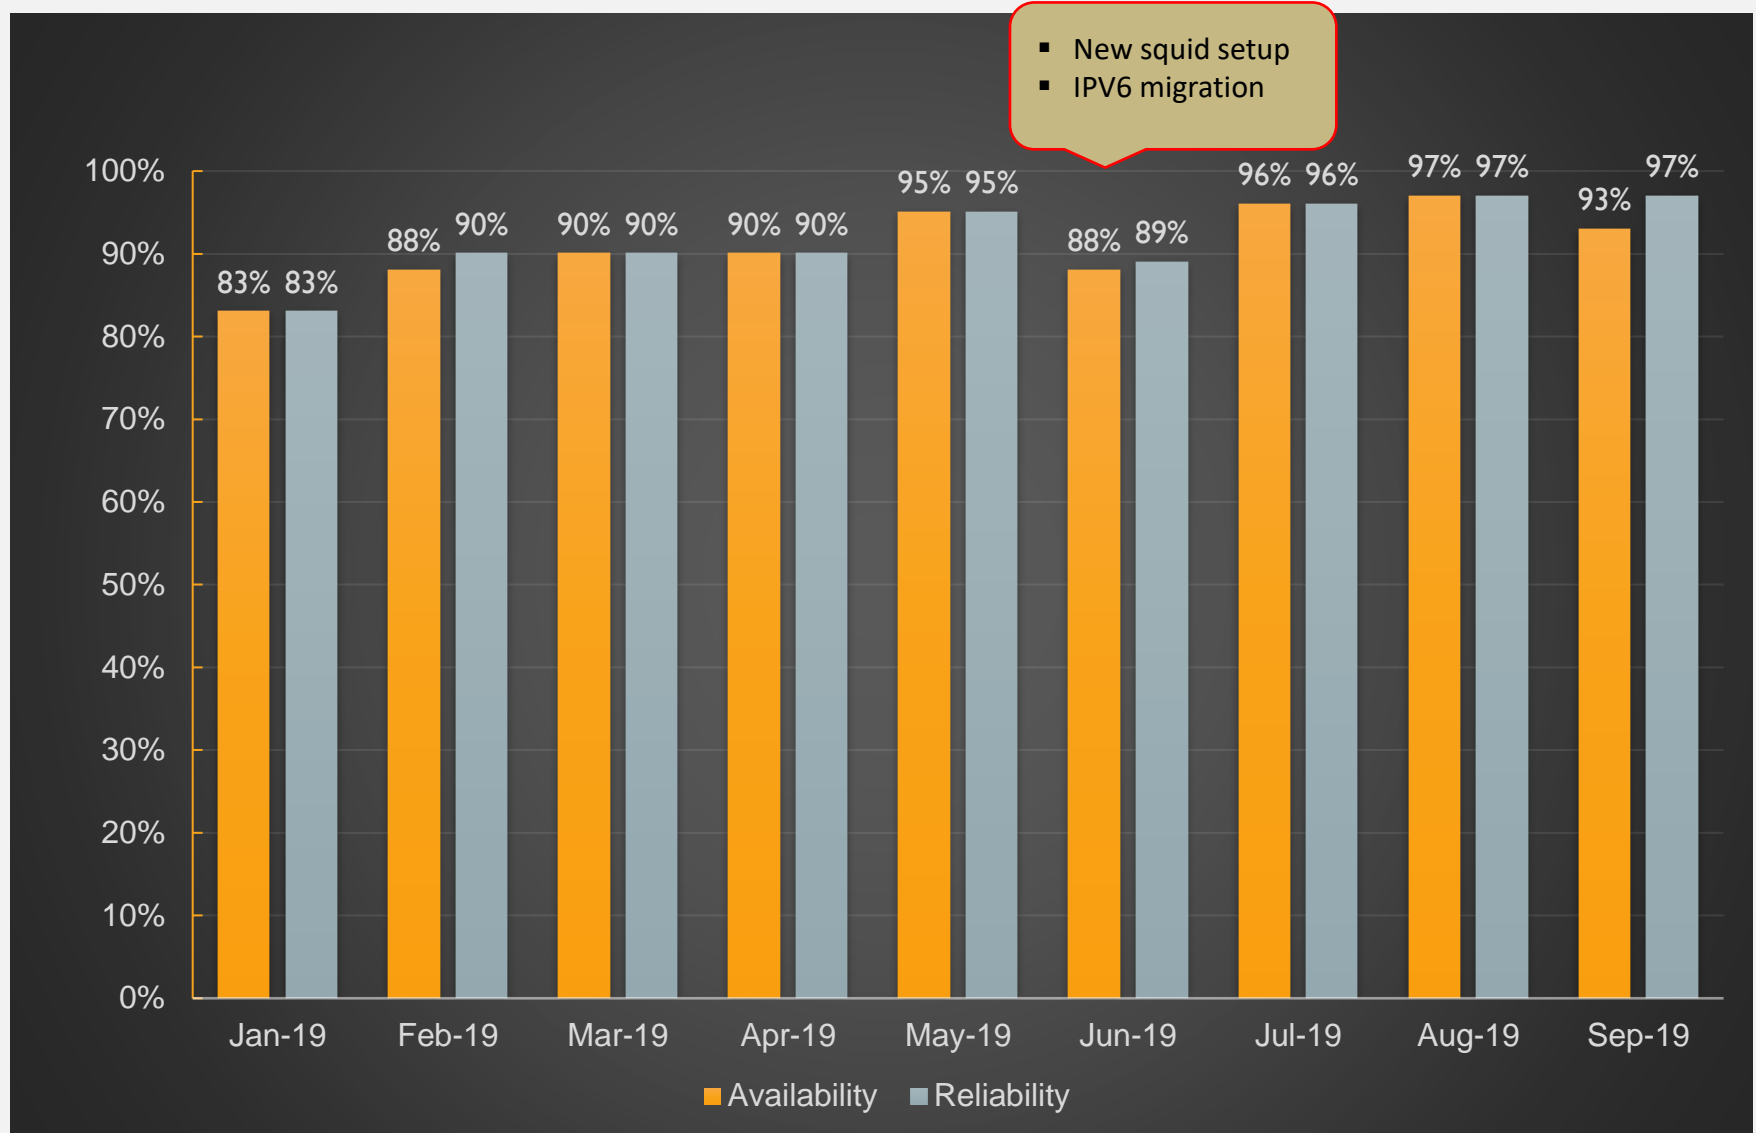
Computing

- Local storage 10 GB per core
- Standard 2 GB memory per CPU core (HT enabled)
- Using HtCondor-CE batch system
- Sufficient space for CVMFS (v2.5.1)
- Singularity - 2.6.1

Storage

- More 1.7 PB storage will be online soon
- 36 drives capacity 4U rack mountable boxes
- Keeping 4 GB per CPU core (HT enabled)
- Combination of 1, 3, 6 and 8TB SAS drives
- Using DPM as storage technology
- Supported protocol – xrootd, gridftp, lcg-cp
- Supportable 20G network via port binding

Availability and Reliability Report



Recent Status

T2_IN_TIFR																							
LifeStatus:	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Site Readiness:	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	
HammerCloud:	100%	100%	100%	100%	100%	100%	100%	100%	98%	99%	100%	100%	100%	100%	100%	99%	98%	98%	99%	100%	100%		
SAM Availability:	99%	100%	100%	100%	100%	100%	100%	100%	86%	92%	92%	94%	85%	91%	97%	92%	100%	100%	92%	92%	92%		
Good T2 links from T1s:	9/14	12/14	14/14	12/14	13/14	10/14	11/14	11/14	14/14	14/14	13/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	12/14	12/14	13/14		
Good T2 links to T1s:	12/13	12/12	12/12	12/12	14/14	12/12	12/12	12/12	11/11	11/12	11/12	13/14	10/12	11/12	11/11	12/12	10/11	11/12	11/12	11/11	10/11		
Active T2 links from T1s:	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14		
Active T2 links to T1s:	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14	14		
	30	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	
	Sep	Oct																					

Report made on 2019-10-21 10:38:04 (UTC)

* = Due to operational errors, the metric has been corrected manually (!=SSB).

-- = Errors on weekends are ignored on Site Readiness computation for T2s [+info]

"Site Readiness Status" as defined in Site-Readiness Twiki (click me):

O = The site is OK for the day

W = WARNING

Life Status: is the site functional, in the waiting room or morgue?:


O = OK

WR = The site is in the Waiting Room

Grid Job Status (Jan to Oct) 2019


1372458 completed

Completed jobs




1844 running jobs

Running jobs



616 pending jobs

Pending jobs



Total completed jobs



	total	percentage
Analysis	1276149	93.0%
Processing	46970	3.4%
Production	24514	1.8%
Unknown	8182	0.6%
Merge	6761	0.5%
LogCollect	6325	0.5%

Currently running jobs



	current	percentage
Analysis	1432	77.2%
Unknown	413	22.2%
Processing	8	0.4%
Production	1	0.1%
LogCollect	0.8	0.0%
Cleanup	0.2	0.0%

Currently pending jobs




	current	percentage
Analysis	612	99%
Processing	3	0%
Unknown	1.2	0%
Cleanup	0.4	0%
LogCollect	0.2	0%
Merge	0.2	0%

Network Utilization

Monthly

TIFR Total Traffic




Last update: Tue Mar 05 2019 10:38:21

CERN

Weekly

TIFR Total Traffic




Last update: Tue Mar 05 2019 10:43:28

CERN

Daily

TIFR-Geant Traffic




Last update: Tue Mar 05 2019 10:43:28

CERN

Weekly

TIFR-Geant Traffic




Last update: Tue Mar 05 2019 10:43:28

CERN


Total Cumulative Transfers ~11.7 PB of data (1 yr.)

CMS PhEDEx - Cumulative Transfer Volume
52 Weeks from Week 42 of 2018 to Week 42 of 2019




Total: 2,022 TB, Average Rate: 0.00 TB/s

CMS PhEDEx - Cumulative Transfer Volume
52 Weeks from Week 42 of 2018 to Week 42 of 2019




Total: 118.32 TB, Average Rate: 0.00 TB/s

CMS PhEDEx - Cumulative Transfer Volume



Total: 5,527 TB, Average Rate: 0.00 TB/s

CMS PhEDEx - Cumulative Transfer Volume



Total: 3,975 TB, Average Rate: 0.00 TB/s

TIFR-HEP Cloud

Planning for the future (cont.)

Near Future: To remain significant and fulfil the needs of Indian users

Year	CPU: No of HT cores	Pledge: CPU HEP-Spec06	Storage pledge (TB)	Power utilization
2018	2500	25,000 (3%)	3000 (4%)	60 KVA
2019	3500	35,000 (3%)	5000 (4%)	80 KVA
2020	9500	95,000 (3%)	9000 (4.5%)	120 KVA

- In 2025 for HL-LHC India need to provide 20,000 cores of CPU+GPU cluster and 15 PetaByte of storage.
- Total power needed for Grid in TIFR - ~500 KVA.
- Dedicated cooling solution for server room essential

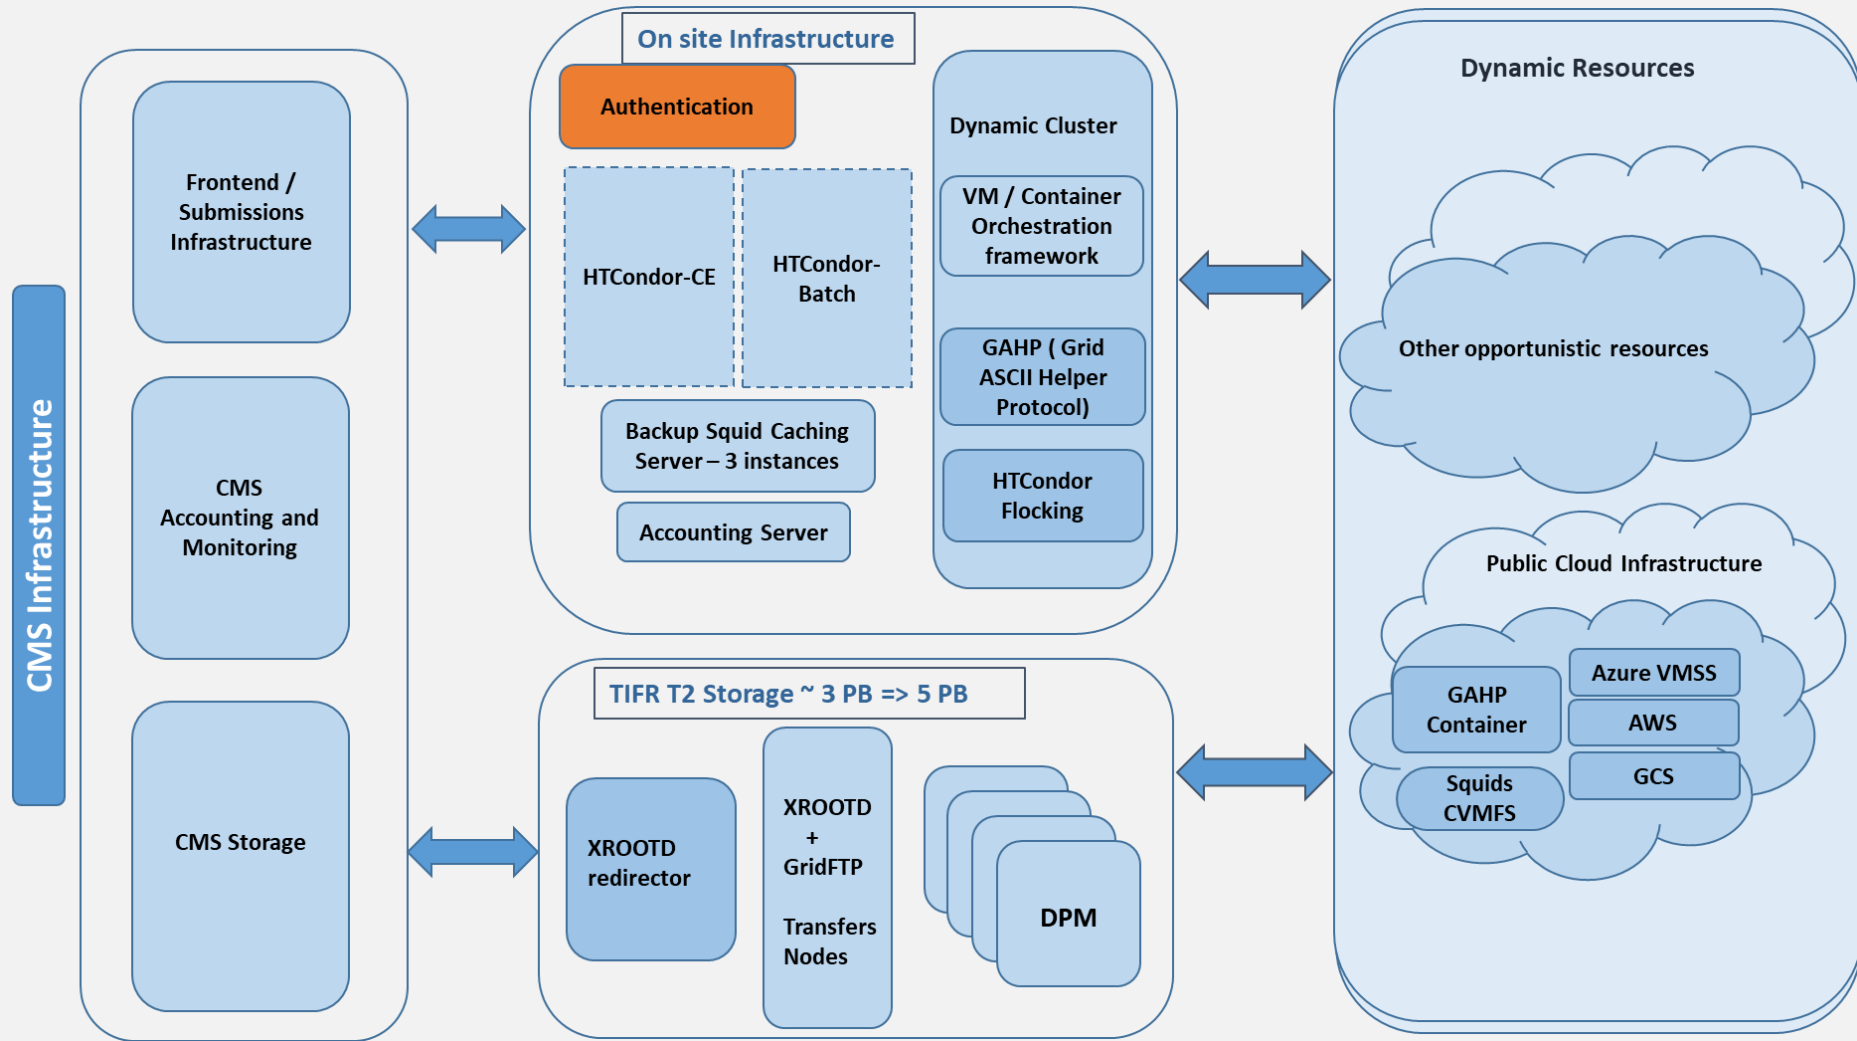
In the existing facility it is difficult to go beyond 9500 cores and 9 petabyte of storage. Even with the latest hardware after retiring old hardware.

This is a common problem

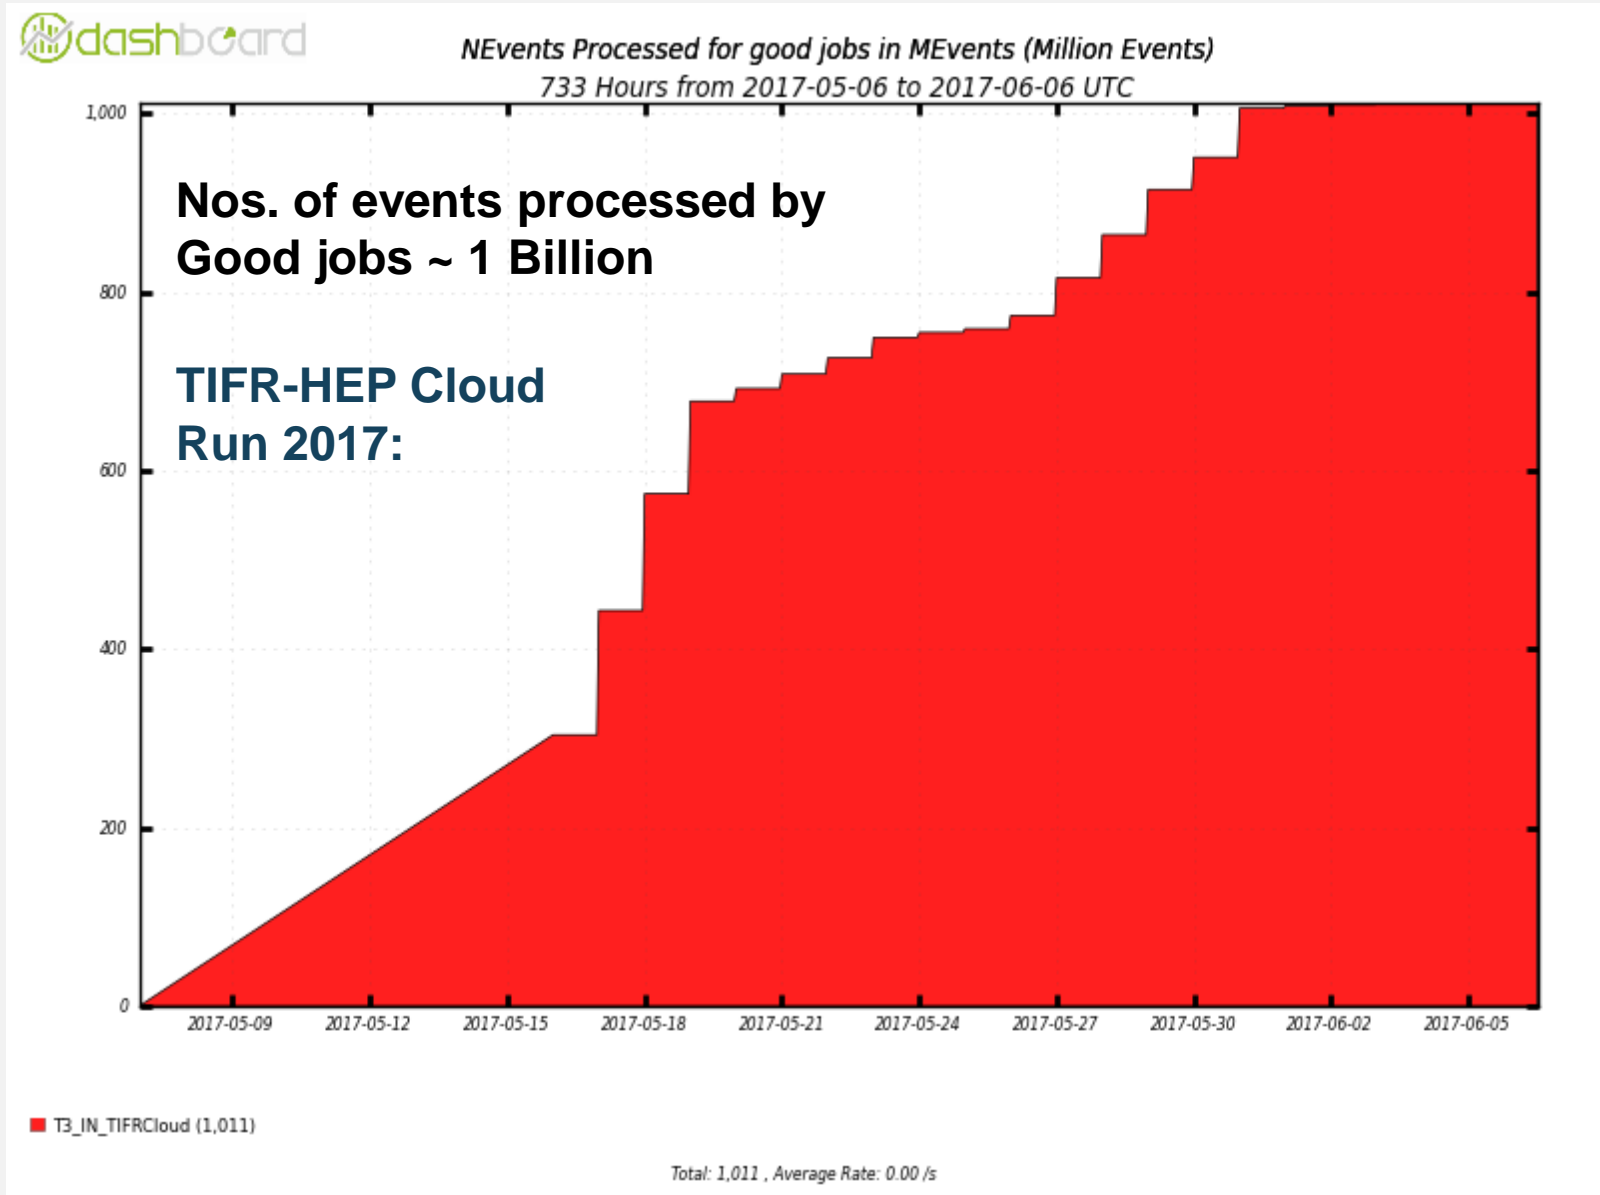
TIFR-HEP Cloud

- Dynamic Resources for WLCG Commissioned in May-2017.
- Collaboration with Microsoft Azure: Azure infrastructure with Grant of USD \$100,000/- in terms of resources.
- MS Cloud Datacentres, three in India (Mumbai, Chennai, Pune)
- Development of tools and technologies for interfacing WLCG Grid with Azure (Grid ASCII Helper Protocol and Condor Annex)
- Successfully processed **1 Billion Physics events in 30 days run.**
- **TIFR earned additional service credits from CMS**
- Resources seamlessly integrated with WLCG
- **Capable of Adding 0 to 10K cores in global pool under 10 minutes.**

TIFR-HEP Cloud (2016-2018) cont..




TIFR-HEP Cloud (Cont..)



T3_IN_TIFRCloud Job Status (Jan to Oct)


74806 completed

Completed jobs




20 running jobs

Running jobs



0 pending jobs

Pending jobs



Total completed jobs




	total	percentage
Analysis	73625	98.4%
Unknown	1170	1.6%
LogCollect	8	0.0%
Cleanup	2	0.0%
Merge	1	0.0%

Currently running jobs



	current	percentage
Analysis	20	95.3%
Unknown	1	4.7%

Currently pending jobs



	current	percentage
Unknown	1	83%
Analysis	0.2	17%

New Technologies adopted at TIFR

New Technologies adopted at TIFR.

- New T3 environment and cluster with **JupyterHub** and HTTPs based access to data
- Completely automated orchestration with true un-attended installation and commissioning using **Puppet and Foreman**.
- Moving complete configuration templates and production files on **Git**
- **Kubernetes cluster** for Docker management and virtualization
- **EOS** storage for T3.

Upgradation Phase and Future Plans

- Application servers and Grid jobs are in CentOS 7, **except storage**
- Computing nodes on 10G dual stack and 20G in storage
- IPv6 fully supported
- NKN new fallback network link
- **1.7 PB additional storage will be online in Nov.**
- Best efforts to get 20 Gbps
- **New hardware procurement (6000 job slots, 4PB storage, 7 x V100 GPU cards, 2 new PCIe based GPU server with 8 cards each, 2 x 10G switch)**
- Local Tier-3 CPU cores expansion
- Single Sign-On and certificate-based authentication

Thank you !

Questions ?