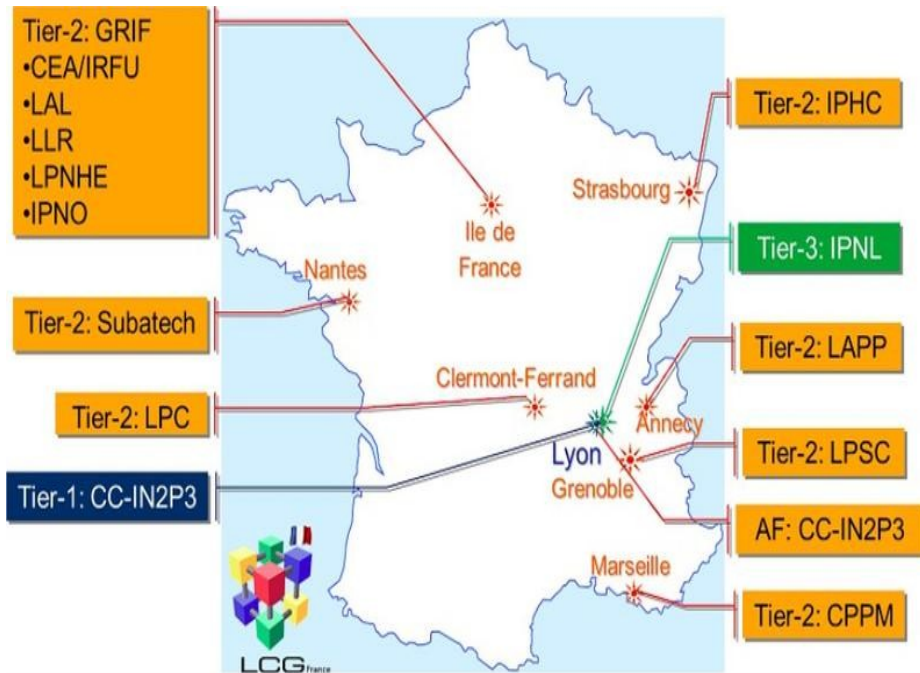


***Recent news from FR ATLAS sites***  
***Emmanouil Vamvakopoulos***

***ATLAS Sites Jamboree***  
***Jan 2016, CERN***



Role	Site	ALICE	ATLAS	CMS	LHCb
Tier-1	IN2P3-CC	✓	✓	✓	✓
	IN2P3-CC-T2 (AF)		✓	✓	
	IN2P3-CPPM		✓		✓
Tier-2	GRIF	✓	✓	✓	✓
	IN2P3-LPC	✓	✓		✓
	IN2P3-IPHC	✓		✓	
	IN2P3-LAPP		✓		✓
	IN2P3-LPSC	✓	✓		
	IN2P3-SUBATECH	✓			
Tier-3	IN2P3-IPNL	✓		✓	

## 6 sites T2s for Atlas VO

## ■ Pledges CPU (KHS)

- ▶ 2015 : 104 over **231** (total) → **T1+AF** : 57 over **124** (total)
- ▶ 2016 : 116 over **270** (total) → **T1+AF** : 60 over **138** (total)

## ■ Pledges Disk (PBytes)

- ▶ 2015 : 10 over **20** (total) → **T1+AF** : 5.1 over 11.6 (total)
- ▶ 2016 : 12 over **24** (total) → **T1+AF** : 5.8 over 13 (total)


## ■ Pledges Disk (PBytes)

- ▶ 2015: 6.7 over **18** (total)
- ▶ 2016: 11 over **25** (total)

# FR T1+T2s : Process Events → jun 2015-jan 2016

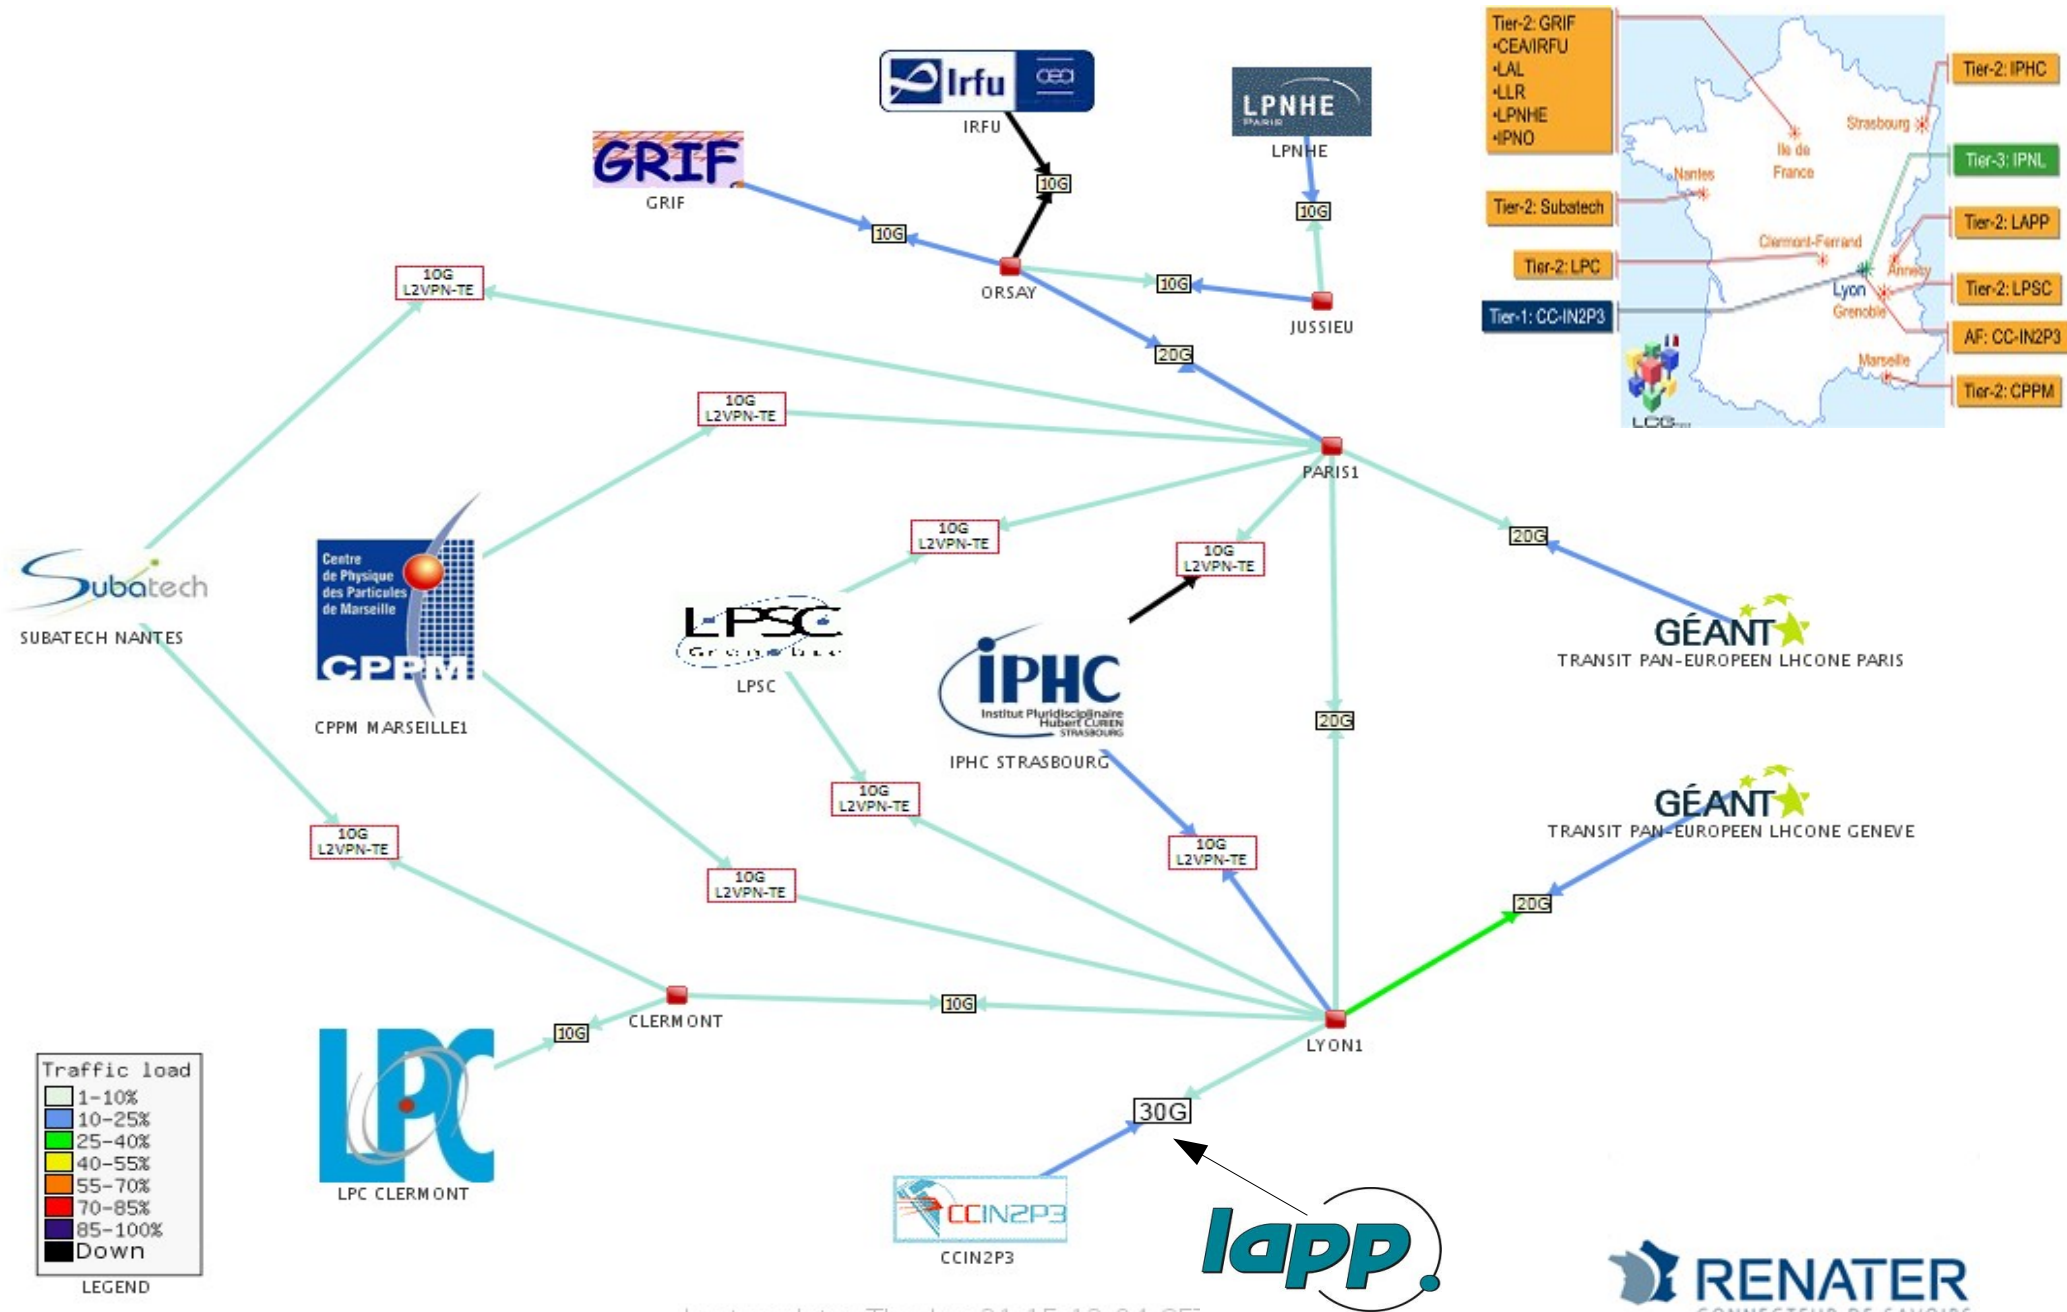


NEvents Processed in MEvents (Million Events) (Sum: 101,210)



- IN2P3-CC - 36.31% (36,749)
- GRIF-IRFU - 6.16% (6,237)
- GRIF-LPNHE - 3.97% (4,013)
- IN2P3-CC-T2 - 24.26% (24,557)
- IN2P3-CPPM - 6.02% (6,097)
- IN2P3-CC-T3 - 0.01% (12.00)
- IN2P3-LPC - 6.98% (7,069)
- IN2P3-LPSC - 5.24% (5,306)
- IN2P3-LAPP - 6.21% (6,285)
- GRIF-LAL - 4.83% (4,885)

# LCG-FRANCE – LHCONETOPLOGY : T1+T2s



Last update: Thu Jan 21 15:12:04 CET 2016

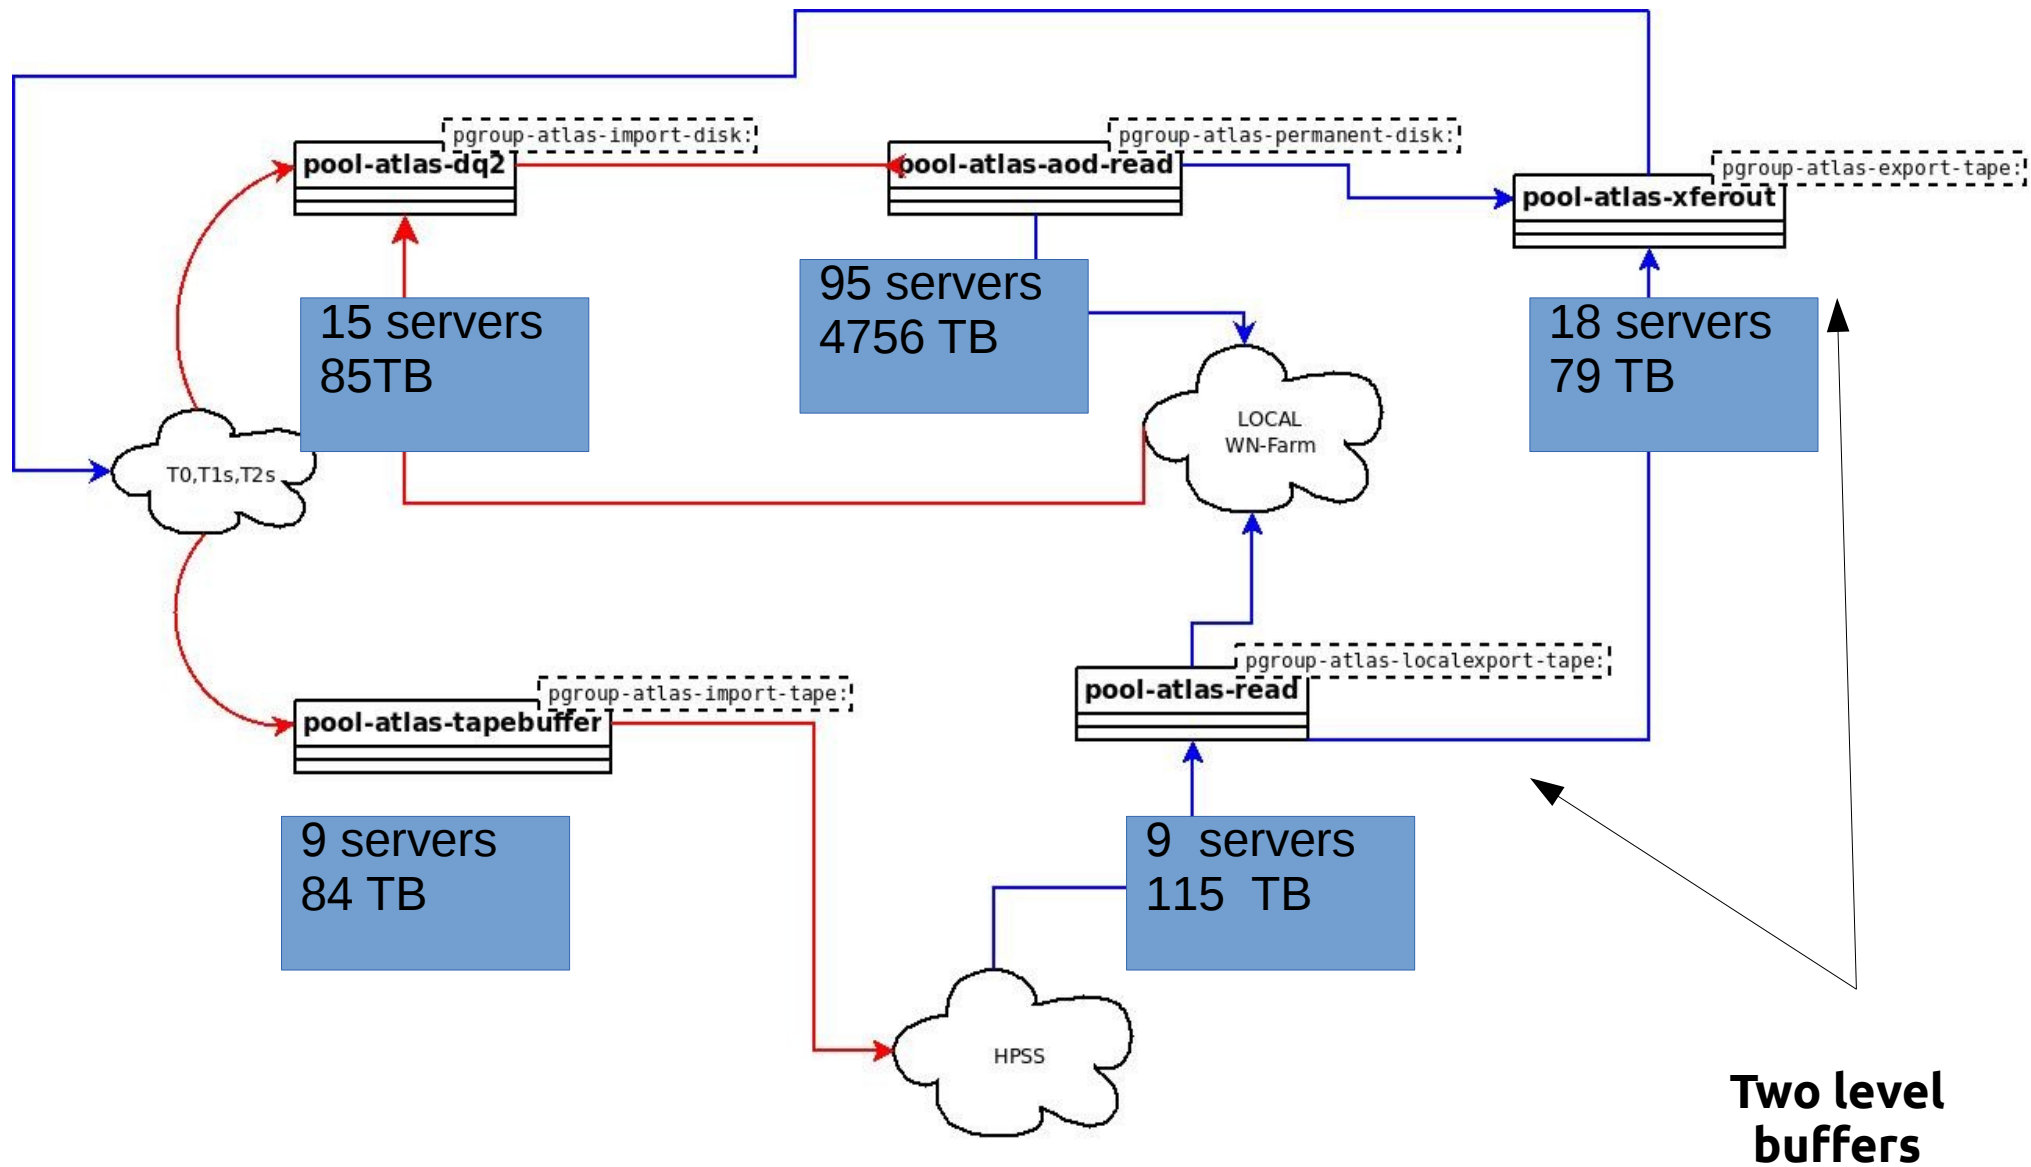
Routers IPv6 (except LPNHE)



- **Regular monthly meeting → Atlas CAF**
  - ▶ Sites Reports
  
- **Regular LCG-France workshops (2 per year)**
  - ▶ Working Group «**Landscape of the Sites**»
    - Budget
    - Quality
    - Human Effort
    - Future of T2s



# dCache Pool Diagram



172 servers  
146 atlas-pools

## ■ Dell PowerEdge R730xd


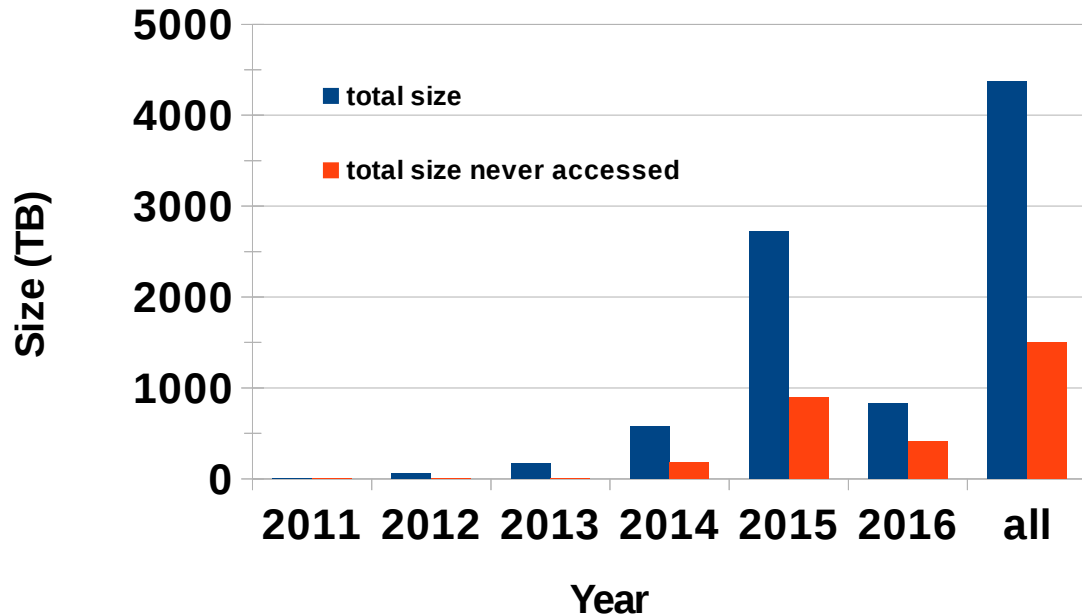
- ▶ 2 x Intel(R) Xeon(R) CPU E5-2609 v3 @ 1.90GHz ( total 12 cores, no HT)
- ▶ 32GB RAM / NIC 10Gbit/sec
- ▶ 6 TB single disks on Raid6
- ▶ 1 Disk enclosure PowerVault MD1400
- ▶ Total raw capacity 120TB (2x60TB)

## ■ 11 new nodes





# File access history



## 2011- up to now

- Total 4369TB → 29.5M of files
- never accessed 1.5PB → 17.2M files

## 2015

- Total 2751TB → 19.3M of files
- never accessed **900TB** → 11.8M files

## 2016

- Total 832TB → 2.9M of files
- Never accessed **411TB** → 1.9M of files

## ■ Protocols

- ▶ Suspend the usage of dcap doors on 2016
- ▶ Full usage of XRootD (local, direct reading and fax door)

## ■ CHIMERA db automatic dump every 30 days

- ▶ Stored in Swift object storage
- ▶ Only for ATLASDATADISK spacetoken
- ▶ <https://its.cern.ch/jira/browse/ATLDDMOPS-5124>

## ■ Incident

- ▶ atlas-read stage-in buffer buffer filled 100% 1 nov 2015
- ▶ 200TB staged-in data in 5 days
- ▶ Due to large pinning time from the side of the client (**Rucio**) → 7 days

\*<https://indico.in2p3.fr/event/11973/>

## ■ **Memory limits**

- ▶ rssmem reported by Grid Engine
- ▶ Still, we use vmem limits , we are trying to relax those limits per case (+50%)
- ▶ Fusion of Panda resources (with vmem at maximum limit)
- ▶ Preliminary test for decommissioning of IN2P3-CC-T2 site

## ■ **New HTCondor Working Group at CC-IN2P3**

- ▶ Test the product
- ▶ Investigate the potential future integration at CC

## ■ Openstack infrastructure pass to KILO (Sep 2015)

- ▶ Update the hypervisors for openstack@Cloud → Dell C6100 ~ 720 cores
- ▶ Migration to Neutron up to end 2016
- ▶ Development **synergy** with INFN for the fair share on top of openstack


## ■ Tests with CloudScheduler

- ▶ CloudScheduler uses native Nova api to instantiate VMs according with the workload (panda)
- ▶ CernVM image
- ▶ 60 Vms X 12CPU MAX 32GB 240GB disk
- ▶ Mixture of single and Mcore jobs (4 cores) → event-simul Only!
- ▶ Benchmarks base on Hammer cloud template (512)
  - *CERN-P1 StressTest - mc12 AtlasG4\_trf 17.2.2 (512)*
  - *AtlasProduction / 17.2.2*
  - *mc12\_8TeV.175590.Herwigpp\_pMSSM\_DStau\_MSL\_120\_M1\_000.evgen.EVNT.e1707\_tid0121239\_5\_00\_derHCBM*

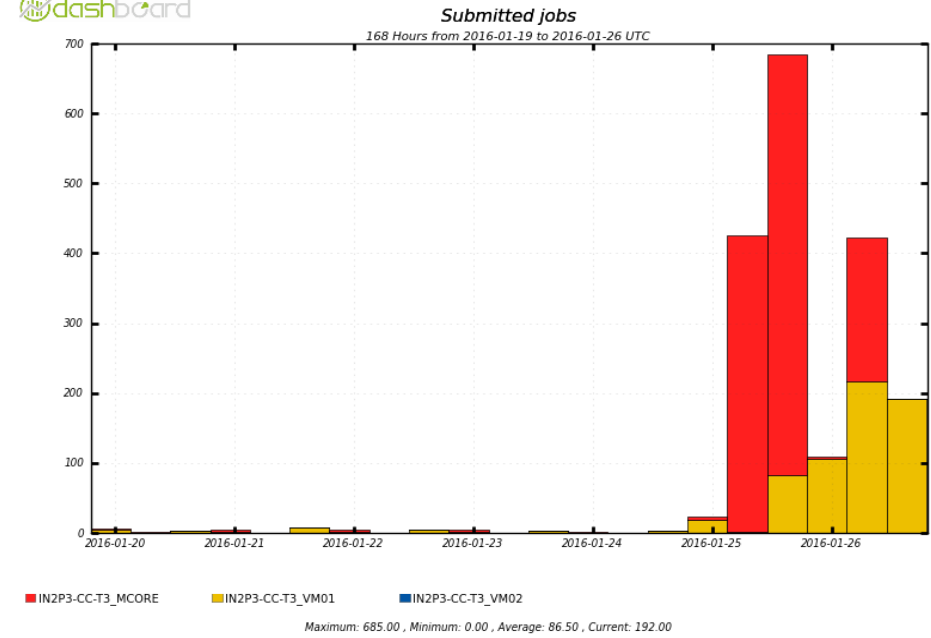
# Cloudatlas@CC Activity -> Last week

benchmark

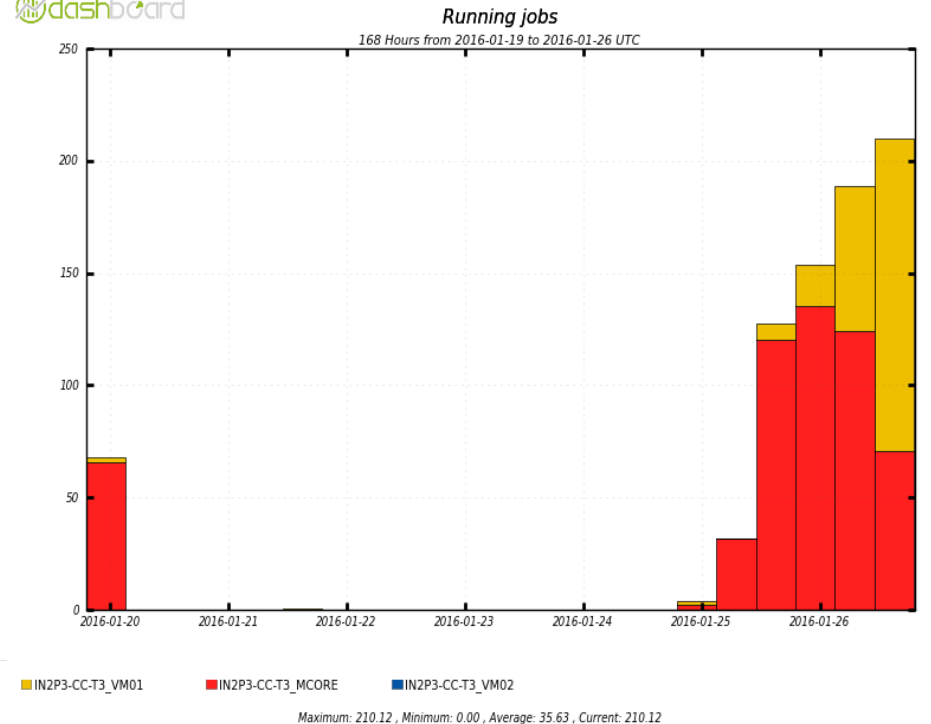
Production-jobs



dashboard



dashboard



**Thank you very much  
for your attention**

# Backup Slides



- **Batch System**

- Grid Engine by UNIVA

- **193 KHS06 on Prod Nov 2015**

- **20,000 job slots**

- **Computing farm contains 800 servers on production**

- ▶ PowerEdge M610 16 cores, 62 servers ( parallel computing)
- ▶ C6100 24 cores, 243 servers
- ▶ PowerEdge C6220 32 cores 347 servers
- ▶ PowerEdge C6220 II 40 cores, 148 servers

- **Pledge 2015 → 124 KHS06 for LHC**

- ▶ 56 KHS06 for ATLAS VO

- **Pledge 2016 → 143 KHS06 for LHC**

- ▶ ~62KHS06 for Atlas VO

- **+ 5.8 KHS06 for atlas ~ T1 2016**

- **Multicore queue is defined only on Tier1 site**

- ▶ **Static allocation of the resources ( manual repartition 24h-72h)**

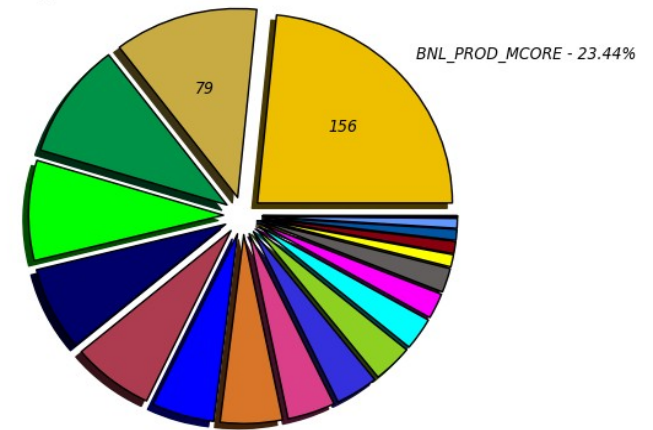
- ▶ **~5000 dedicate slots for Multicore atlas Jobs**

- **Second term of 2015:**

- ▶ **7.19 % of 668 M events (simul- on plegdes resources) for CC multicore queue**






NEvents Processed in MEvents (Million Events) (Sum: 668.00)  
TRIUMF\_MCORE - 11.95%



BNL_PROD_MCORE - 23.44% (157.00)	TRIUMF_MCORE - 11.95% (80.00)	CERN-PROD_MCORE - 9.93% (66.00)
RAL-LCG2_MCORE - 8.56% (57.00)	IN2P3-CC_MCORE_HIMEM - 7.19% (48.00)	RRC-KI-T1_MCORE - 6.81% (45.00)
NIKHEF-ELPROD_MCORE - 5.19% (35.00)	FZK-LCG2_MCORE - 5.14% (34.00)	UIO_MCORE - 3.96% (26.00)
pic_MCORE - 3.54% (24.00)	Riwan-LCG2_MCORE - 3.28% (22.00)	SARA-MATRIX_MCORE - 2.77% (19.00)
ORNL_Titan_MCORE - 2.17% (15.00)	INFN-T1_MCORE - 2.06% (14.00)	NCS_MCORE - 1.05% (7.00)
HPC2N_MCORE - 1.03% (7.00)	CERN-PROD_CLOUD_MCORE - 0.93% (6.00)	DCSC_MCORE - 0.83% (6.00)
IN2P3-CC_MCORE - 0.08% (1.00)	BNL_CLOUD_MCORE - 0.08% (1.00)	TRIUMF_ARC_MCORE - 0.01% (0.00)

# Links: North and South → Geant



## ■ dCache next update to 2.13.12+ on Dec 2015

### ▶ 44 Million files in Chimera for ATLAS

- Total 5501TB – Used 4517 TB – Free 621 TB – Buffers 363 TB

### ▶ AtlasDataDisk

- Total 4551 TB – Used 4122 TB – Free 429 TB

### ▶ AtlasLocalGroupDisk

- Total 283 TB – Used 252 TB – Free 31 TB

## ■ Pledge disk

▶ 2015 → 5100 TB

▶ 2016 → 5820 TB