


ALICE @ CNAF

Stefano Piano – INFN sez. Trieste

Operations Organization

- ALICE-IT Computing Coordination:
 - Stefano Piano
- Deputy:
 - Stefano Bagnasco
- Tier-2 Operations Coordination:
 - Francesco Noferini (assisted by Nicolo Jacazio at CNAF)
 - Monthly phone conference for coordination and performance monitoring
 - Yearly face-to-face workshop (2012 @ Catania, 2013 @ Trieste, 2014 @ Frascati, 2015 @ CNAF, 2016 @ PD, 2017 @ BA)
- Monthly Tier-1 Management Board at CNAF

CNAF incident & recovery

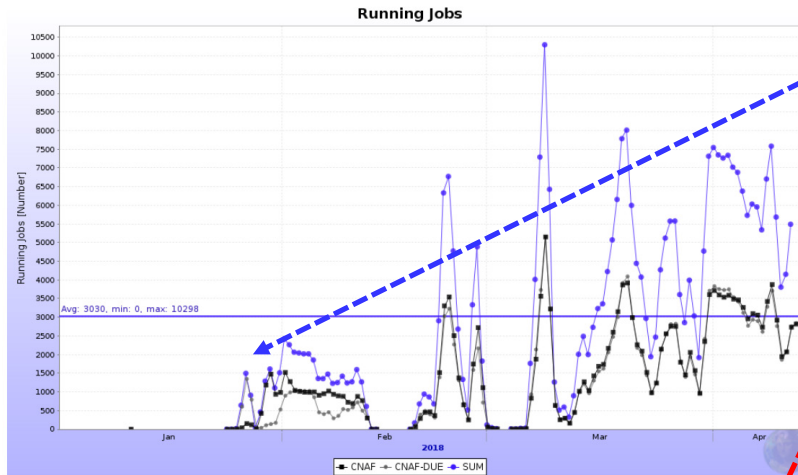


- Down Nov 9 until late Jan due to a huge flooding incident
 - More details in INFN sites presentation
- First jobs this year ran late Jan
- 68 M disk SE files were unavailable until mid Feb
 - None were lost !
- Recovery of the wet tapes in progress, looking good so far (not critical)
- *We thank and congratulate CNAF with the success of their careful efforts!*

- Capacity loss (temporary)
 - 4.6PB of disk (90% used, 7% of total)
 - 38 kHS06 CPU cores (4% of total, 22% of T1 capacity)
 - 6.2PB of RAW data replicas on tape (20% of replicated data)
 - 22 tapes damaged from the flooding (about 300 TB of data)
 - 4 xrootd servers
- Data loss
 - No RAW data loss, one more copy exist at CERN
 - Disk – max 10% (MC) to 15% (RAW) of single-replica ESDs
- Performance loss
 - CNAF is an excellent T1, used for RAW data replication and reconstruction and has very good efficiency for analysis
 - The single copy ESDs are distributed ~uniformly across all computing centres => ESD-based analyses lose ~10% statistics

- No large detrimental effect from the loss of CPU thanks to additional CPU offered at GridKA and CERN (and Bari-RECAS)
- The loss of 10% of ESDs (unique files) resulted in reduction of statistics for certain analyses (possible to re-generate ESDs)
- The AODs were replicated at other storages, in the order of importance for analysis, to preserve the turnaround speed of the analysis and to avoid bottlenecks
- The data on the damaged tapes has one replica at CERN:
 - high cost of tape restoration → ALICE will replicate again the data at CNAF
- In addition about 1 PB of 2017 RAW data were not replicated to CNAF and the replication of the data was postponed to 2018.
 - no replication at other T1s to keep balanced RAW at the T1s



Prompt ALICE reaction to CNAF restart




Since Jan 25th ALICE jobs have been running on CNAF queues! At the beginning only simulation jobs, reconstruction and analysis jobs turned on March (full functionality of SE is needed).

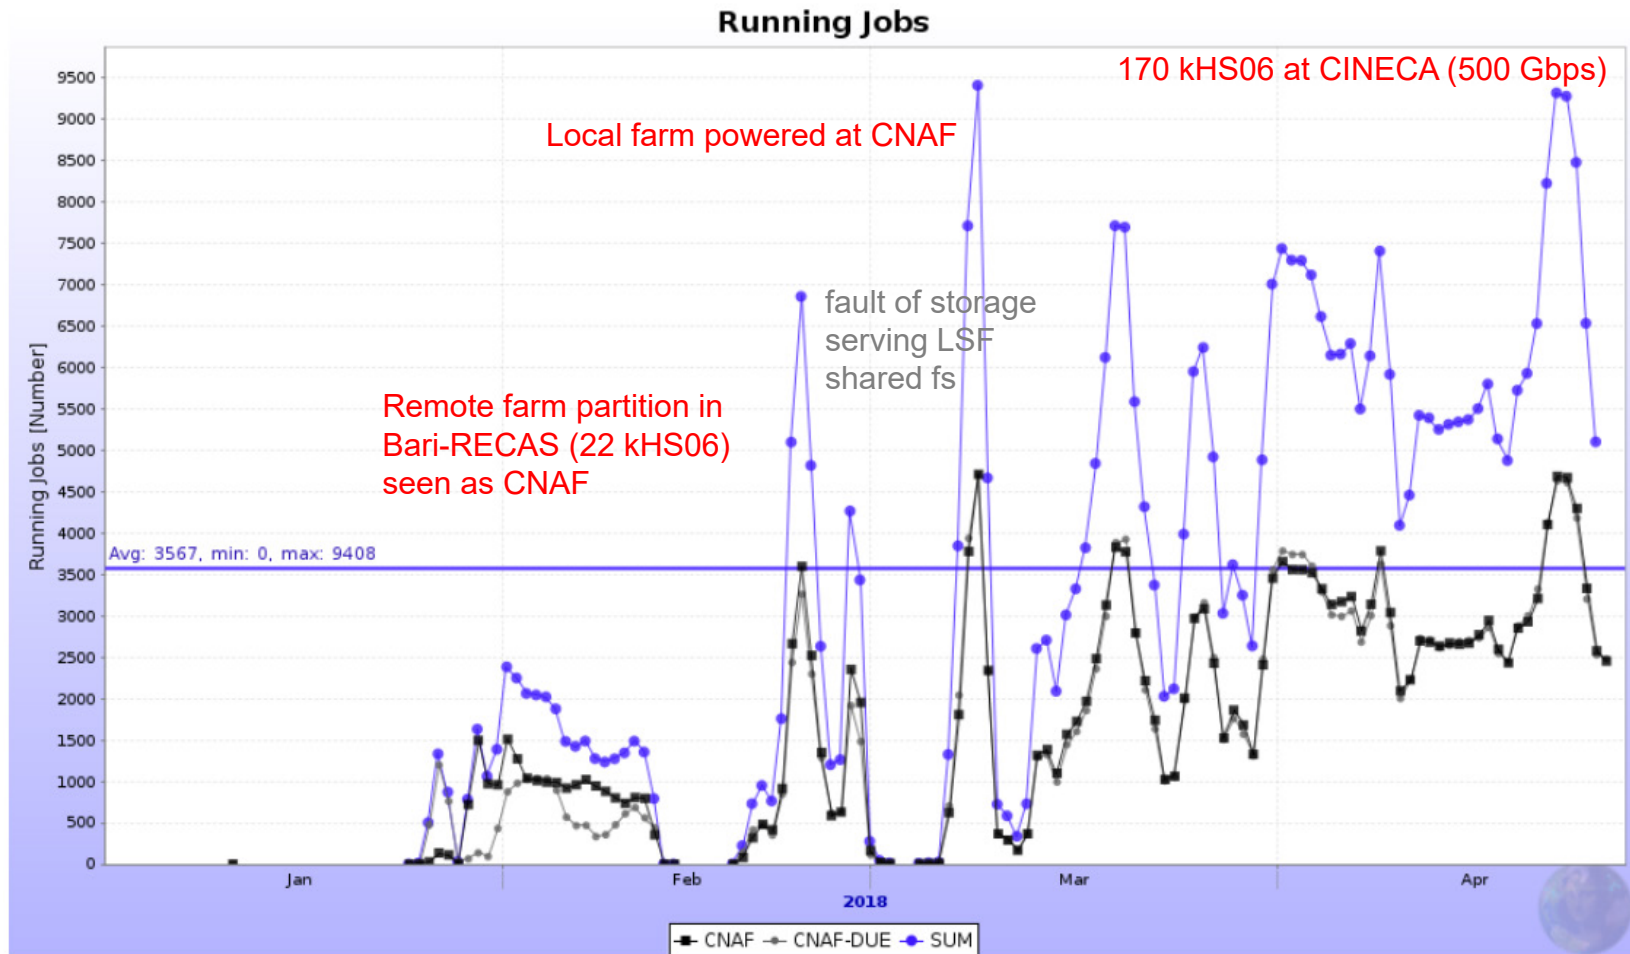
Tape available for ALICE on January 31st

Since Feb 23th SE accessible on R/W



Running Jobs





Global financial budget corresponding to pledge (as in the previous years) for 2017 and 2018

Resource procurement for 2017:

- INFN pushed in 2017 for (as much as possible) centralized/grouped tenders
- new procurement code on May 2017 slowed down our purchase processes
 - **CPU for Tier1:** 170 kHS06 at CINECA already deployed
 - **CPU for Tier2:** delivery expected on **April 2018**
 - **DISK for Tier1:** disk pledged for 2017 expected on **May** and disk pledged for 2018 expected on July
 - **DISK for Tier2** (grouped “ALICE+CMS” procedure for all Tiers2 but Torino): received offers on December, not possible to proceed with awarding procedure because “the incompleteness and the irregularity of the company declarations”... delivery **expected on September** [Torino already installed - extension of a previous purchase]
 - **TAPE for Tier1:** pledged on 2017 is **already installed**, 2018 coming soon


Resource procurement for 2018:

- Tier1 CPU (at CINECA) and DISK covered by previous tenders
- Tier2 back to dedicated tenders only for ALICE hoping for “faster” purchase process

RAW data volume expected in 2018

pp:
 $430 \text{ Hz (readout rate)} * 1.7 \text{ MB (event size)} * 86400 \text{ (seconds)} * 0.57 \text{ (combined efficiency)} * 150 \text{ (days)} = 5.4 \text{ PB}$

Pb-Pb
 With the increased HLT compression, data rate from 10 GB/s to 7.1 GB/s



(Scenario 1):
 Data taking at 7.1GB/s
 100M MB events, 250 M central events plus several triggered samples, including muon arm events ($750 \mu\text{b}^{-1}$).

(Scenario 2):
 Data taking at 7.6GB/s
 additional 100M MB events
 increasing the size of the tape buffer at CERN
 increase the number of events to be processed offline

$7.1\text{GB/s [7.6GB/s]} * 86400 \text{ (seconds)} * 0.57 \text{ (combined efficiency)} * 24 \text{ (days)} = 8.3\text{PB [9PB]}$


ALICE resources request 2019

ALICE		2017			2018		2019		
		CRSG recomm.	Pledged	Used	CRSG recomm.	Pledged	Request	2019 req. /2018 CRSR	C-RSG recomm.
CPU	Tier-0	292	292	389	350	350	430	123%	430
	Tier-1	256	235.5	295	307	279.5	365	119%	365
	Tier-2	366	279.6	299	312.9	312.9	376	120%	376
	HLT	n/a	n/a	26	n/a	n/a	n/a	n/a	0
	Total	914	807.1	1010	969.9	942.4	1171	121%	1171
	<i>Others</i>			39					
Disk	Tier-0	22.4	22.4	19.3	26.2	26.2	34.3	131%	34.3
	Tier-1	25.4	21.8	18.245	30.5	30.4	37.9	124%	37.9
	Tier-2	31.4	22.7	20.06	29	29	33.9	117%	33.9
	Total	79.2	66.9	57.6	85.7	85.6	106.1	124%	106.1
Tape	Tier-0	36.9	36.9	29.7	49.1	49.1	44.2	90%	44.2
	Tier-1	30.9	30.6	22.3	40.9	42.2	37.7	92%	37.7
	Total	67.8	67.5	52	90	91.3	81.9	91%	81.9



Thank you a lot!

Stefano Piano – INFN sez. Trieste



Pledge 2018

	Spring request (*)		Fall request (*)		PLEDGE 2018 (REBUS)	
	Tier1	Tier2	Tier1	Tier2	Tier1	Tier2
CPU (HS06)	52020	74460	52190	67660	52020	61050
DISK (TB)	5440	6970	5185	5967	5140	6659
TAPE (TB)	13500		13510		13530	

(*) ALICE requests scaled for INFN share (excluding CERN):
~17% for CPU and disk, ~33% for tape