

Rapporto delle attività del Tier1 CNAF

Maggio 2018

**NB: NON TUTTI I PLOT SONO COMPLETI A
CAUSA DI PROBLEMI CON IL SISTEMA DI
MONITORING A CUI STIAMO LAVORANDO!**

Run Coordinator: Marcelo Soares

7 giugno 2018

Indice

1	Utilizzo delle risorse - Storage	2
1.1	Disk usage - LHC	2
1.2	Disk usage - No LHC	4
1.3	Consistenza fra accounting CNAF e cataloghi di esperimento nell'uso dello storage	12
2	Tickets	13
2.1	Aperti	13
2.2	Chiusi nell'ultimo mese	13
3	ALICE	14
4	ATLAS	16
5	CMS	23
6	LHCb	32
7	Esperimenti di Gruppo 2	34
8	Altri esperimenti	39


1 Utilizzo delle risorse - Storage

1.1 Disk usage - LHC




CMS

CMS





LHCB


LHCb







1.2 Disk usage - No LHC




















1.3 Consistenza fra accounting CNAF e cataloghi di esperimento nell'uso dello storage

ALICE		
	Catalogo	Accounting CNAF
Disco + Buffer (PB)	3.984	3.943
ATLAS		
	Catalogo	Accounting CNAF
Disco(PB)	-	-
Tape (PB)	-	-
CMS		
	Catalogo	Accounting CNAF
Disco(PB)	3.567	3.694
Nastro (PB)	12.01	14.806
LHCb		
	Catalogo	Accounting CNAF
Disco(TB)	2.661	3.821
Nastro + Buffer (TB)	7.361	7.094

2 Tickets

2.1 Aperti

Ticket-ID	VO	Subject
134325	CMS	xrootd xonfig at T1.IT.CNAF
134929	CMS	Lost files on wet tapes
135207	CMS	checksums do not match at T1.IT.CNAF_Buffer

2.2 Chiusi nell'ultimo mese

Ticket-ID	VO	Subject
135329	CMS	T1.IT.CNAF SAM CONDOR-JobSubmit tests
135303	ATLAS	Atlas file transfers failing with destination to INFN-T1
135207	CMS	checksums do not match at T1.IT.CNAF_Buffer
135166	CMS	Merge issues T1.IT.CNAF
135153	none	This TEST ALARM
135112	ops	Issues detected at INFN-T1
135098	CMS	T1.IT.CNAF SAM SRM failures
135094	ATLAS	INFN-T1 DATATAPE and MCTAPE: transfer errors due to source and destination checksum mismatch
135045	CMS	Posix access inside singularity at CNAF
134989	CMS	T1.IT.CNAF SAM SRM failures
134978	CMS	Automatic approval of the transfer requests
134976	LHCb	MCSimulation jobs Failed at INFN-T1
134960	ATLAS	INFN-T1 failing transfers with DESTINATION SRM_PUT_TURL error on the turl request
134949	CMS	T1.IT.CNAF storm-fe-cms.cr.cnaf.infn.it SRM SAM test critical for 2 hours
134934	CMS	tape request approval
134749	CMS	T1.IT.CNAF SAM CONDOR-JobSubmit tests CancelledPurged
134639	ATLAS	Deletion errors at INFN-T1-MWTEST DATADISK

3 ALICE




Figura 1: Number of jobs and CPU efficiency - **Alice** (pledge 29045 HS06)

- Site Availability using ALICE_CRITICAL




Figura 2: Availability

Report:

- 07/05: Unscheduled downtime on tape library due to several hardware failures during the weekend

4 ATLAS




Figura 1: Number of jobs and CPU efficiency - **Atlas** (pledge 46800 HS06)




Figura 2: Number of jobs and CPU efficiency - **Atlas_himem** (pledge 46800 HS06)




Figura 3: Number of jobs and CPU efficiency - **Mcore** (pledge 46800 HS06)

- Site Availability using ATLAS_CRITICAL






Figura 3: Availability


• ATLAS Throughput




• T0 Export



- Disco ATLAS (datadisk -scratchdisk)



- Nastro ATLAS (datatape - mctape)



Report:

- 30/05: Transfer fails from Florida to CNAF: ipv6 issue
- 28/05: Tape Errors: impossible to write on them
- 24/05: Atlas file transfers failing with destination to INFN-T1 related to a mis-configuration of the StoRM spacetoken quota management: scratchdisk space token got full and it blocked also the other endpoints
- 13/05: Errors between INFN-T1 DATATAPE and MCTAPE: transfer errors due to source and destination checksum mismatch
- 07/05: INFN-T1 failing transfers with DESTINATION SRM.PUT_TURL error on the turl request due to a mis-configuration of the storage pool placement policy

5 CMS




Figura 1: Number of jobs and CPU efficiency - CMS **Single Core** (pledge 48000 HS06)




Figura 2: Number of jobs and CPU efficiency - CMS **Multi Core** (pledge 48000 HS06)

- Site Availability using CMS_CRITICAL_FULL


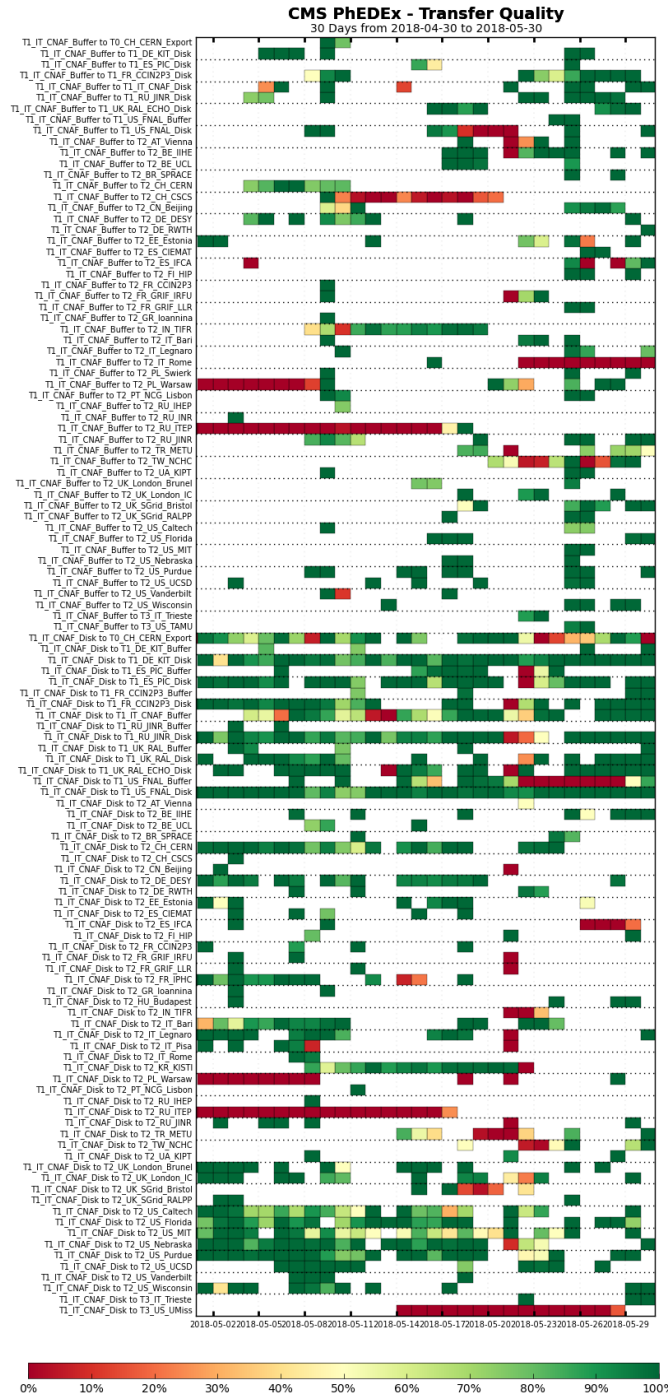



Figura 4: Availability


- Phedex Transfers Quality (CNAF → All)




- **Phedex Transfers Quality (All $\rightarrow$ CNAF)**




- Phedex Transfers Volume(CNAF → All)



- Phedex Transfers Volume (All → CNAF)



- Completed jobs status @ T1



• Site Readiness

T1_IT_CNAF																					
LifeStatus:	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Site Readiness:	O	O	O	E	E	O	O	E	E	E	E	E	E	E	E	E	O	O	O	O	E
Maintenance:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
HammerCloud:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
SAM Availability:	100%	100%	100%	79%	71%	100%	100%	79%	75%	100%	86%	100%	74%	77%	54%	40%	100%	100%	100%	100%	72%
Good T1 links from T0:	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28	28
Good T1 links from T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links from T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links to T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links to T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Active T1 links from T0:	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
Active T1 links from/to T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Active T1 links to T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Report made on 2018-05-22 16:38:04 (UTC)																					

T1_IT_CNAF																					
LifeStatus:	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Site Readiness:	E	E	O	O	O	O	E	E	O	O	O	E	E	O	O	O	E	O	E	E	E
Maintenance:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
HammerCloud:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
SAM Availability:	44%	40%	100%	100%	100%	100%	75%	84%	100%	100%	100%	40%	6%	8%	100%	100%	100%	8%	53%	71%	53%
Good T1 links from T0:	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32	32
Good T1 links from T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links from T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links to T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Good T1 links to T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Active T1 links from T0:	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
Active T1 links from/to T1s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Active T1 links to T2s:	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000	10000
Report made on 2018-06-05 10:38:04 (UTC)																					

Report:

- 04-05/05: Errors in SAM tests due a storm problem
- 07/05: Unscheduled downtime on tape library due to several hardware failures during the weekend
- 08-11/05: Errors in SAM tests due to tests to try to find the better Singularity
- 12-13/05: Transfer errors due to gridftp servers dead, resolved after restart
- 14-15/05: SAM tests errors, resolved after a intervention on CEs
- 16/05: Errors due to an unreachable gridftp server
- 21/05: SAM tests errors
- 22/05: Transfer failures due to all dead gridftp servers
- 26-27/05: SAM Errors due to CEs problems

- 31/05: All transfers from world to CNAF failed due to proxy expired on Phedex ui

6 LHCb








Figura 5: Number of jobs and CPU efficiency - LHCb + LHCb TIER2 (pledge 28080 HS06)



Report:

- 10-15/05: Data migration from Huawei disk server to newer Huawei disk server



7 Esperimenti di Gruppo 2

- AMS-O2 - Number of jobs and CPU efficiency (pledge 12800 HS06)

Report:

- Setup istanza MongoDB su macchina cloud
- Cambiato owner di /storage/gpfs_ams/ams/users/nmasi

- richiesta di xrootd
- richiesta di condor_submit sulle ui
- **ARGO-YBJ - Number of jobs and CPU efficiency (pledge 200 HS06)**
Report:
 - NTR
- **Auger - Number of jobs and CPU efficiency (pledge 3090 HS06)**
Report:
 - Cambio filesystem da da gpfs01 a data
 - Ripartono ad usare le risorse al CNAF: 30 TB storage e tutti i nodi quando disponibili. Lato CNAF tutto pronto
- **Borexino - Number of jobs and CPU efficiency (pledge 1500 HS06)**
Open Actions:
 - Richiesta GPUs
 - richiesta di tenere overpledge fino a fine giugno
 - rsync dalla ui-borexino a LNGS per 12 file da 5GB al giorno
 Report:
 - Eseguito move di file tra folder con owner diversi
- **COSMO_WNEXT - Number of jobs and CPU efficiency (pledge 2500 HS06)**
Report:
 - Trasferimento dati (400TB) da Cineca a CNAF (tape)
- **CTA - Number of jobs and CPU efficiency (pledge 4000 HS06)**
Open Actions:
 - Richiesta tool per trasferimento dati da Torino
- **CUORE - Pledge 1400 HS06**
Open Actions:

- Esigenze particolari per trasferimento dati. Soluzione proposta (da discutere): user-interface ad hoc non molto carrozzata ad uso esclusivo di cuore per gli rsync (100 GB/giorno)
- Per i 30 TB iniziali usano globus-url-copy/dataclient
- Controllare slot assegnati a una coda a partire dallo shares perchè ci sono pochi job running contemporaneamente (una decina) a fronte di 1900 HS06 pledged (190 slots)

Report:

- NTR

- **CUPID-0 - Number of jobs and CPU efficiency (pledge 100 HS06)**

Report:

- NTR

- **DarkSide-50 - Pledge 4000 HS06**

Open Actions:

- NTR

Report:

- NTR

- **Fermi-LAT - Number of jobs and CPU efficiency (pledge 900 HS06)**

Report:

- NTR

- **Gerda - Pledge 40 HS06**

Report:

- NTR

- **ICARUS - Pledge 0 HS06**

Open Actions:

- Vorrebbero i dati da una tape bagnata mandata in Germania
- Configurare area disco

Report:

- Crea una mailing list con riferimenti esperimento

- **Juno - Pledge 1000 HS06**

Open Actions:

- NONE

Report:

- NTR

- **LHAASO - Number of jobs and CPU efficiency (pledge 300 HS06)**

Open Actions:

- Richiesta tool per trasferimento dati da Torino

Report:

- NTR

- **MAGIC - Number of jobs and CPU efficiency (pledge 296 HS06)**

Report:

- NTR

- **OPERA - Number of jobs and CPU efficiency (pledge 200 HS06)**

Report:

- NTR

- **PAMELA - Number of jobs and CPU efficiency (pledge 650 HS06)**

Report:

- NTR

- **Virgo - Number of jobs and CPU efficiency (pledge 30000 HS06)**

Open Actions:

- Installare autoconf su ui di virgo
- In corso test utilizzando i sottogruppi al posto dei role nel proxy voms

Report:

- NTR

- **XENON - Number of jobs and CPU efficiency (pledge 700 HS06)**

Report:

- Rate di trasferimento abbassato. Problemi precedenti riportati da Vladimir scomparsi

8 Altri esperimenti

- **BELLEII - Number of jobs and CPU efficiency (pledge 10000 HS06)**

Report:

- NTR

- **CDF - Number of jobs and CPU efficiency (pledge 1300 HS06)**

Report:

- NTR.

- **COSINUS - Number of jobs and CPU efficiency (pledge 0 HS06)**

Report:

- NTR

Open Action:

- NTR

- **CSES-Limadou - Pledge 400 HS06**

Report:

- Assegnati 2 TB di extrapledge
- Proposto di usare gfal per i trasferimenti, mandate istruzioni su come installarlo e usarlo
- Aumentato il numero di inode a 10M perchè non si riusciva a scrivere nell'area /storage/gpfs_data/limadou/data

- **DAMPE - Pledge 3000 HS06**

Open Action:

- Utente lamenta job che non vanno in running in dampe_usr. Causa: utente non appartiene a gruppo dampeuser
- Problemi di dati riportati sbagliati nel monitoring mon-tier1

Report:

- NTR

- **FAMU - Pledge 250 HS06**

Report:

- NTR

- **Gruppi Fisica Teorica - Number of jobs and CPU efficiency**

Report:

- NTR

- **KLOE - Pledge 0 HS06**

Report:

- spostata sui gridftp di archive e tutto va

- **KM3Net - Pledge 0 HS06**

Open Actions:

- Richieste per fornire un servizio web di aiuto alla integrazione dell'esperimento Km3:
 - * high availability
 - * poche risorse di calcolo (CPU/ RAM 4 core/8 GB)
 - * server LAMP (linux, apache, mysql/mariadb, php)
- Le proposte del CNAF:
 - * server virtuale ospitato al CNAF
 - * autenticazione al servizio web tramite IAM, esteso a INFN, CNRS, NIKHEF ed eventualmente altri enti di ricerca europei)
- Esigenza ancora aperta, ma il codice non è ancora pronto da aprte dell'esperimento

Report:

- Settato tape per KM3 per trasferimento dati (tape) Lione-CNAF. Trasferimento funziona
- Per il trasferimento dati da Tolone al CNAF verrà messo in opera un sistema per il trasferimento diretto usando globus-toolkit.
- Il trasferimento è iniziato ma procede a rilento a causa del fatto che la procedura con cui vengono esposti i file a CC-Lyon non è ottimale. Ad oggi, circa 12TB su circa 14 sono stati copiati su tape/disco.

- **COMPASS - Pledge 0 HS06**

Report:

- NTR

- **LHCf - Number of jobs and CPU efficiency (pledge 4000 HS06)**

Report:

- NTR

- **NEWS - Pledge 0 HS06**

Report:

- Report phone:
 - * Emulsioni nucleari prodotte in italia e giappone
 - * storage: richiesta 60TB disk + 60TB tape
 - * cpu: 200HS
 - * 3.5GB RAM ok per simulazione, forse pochi per analisi
 - * SL6/7 OK
 - * proposto CVMFS per software
 - * FLUKA+Geant simulazione
 - * ROOT per analisi, python
 - * i job devono accedere ai 60 TB di dati su disco, quindi bisogna montare sui wn questi 60 TB con accesso POSIX
 - * configurare area storage e tape per accedere sia via WebDav con certificato sia con gridftp con e senza certificato (una soluzione potrebbe essere dataclient)
 - * Mandare istruzioni per accesso al CNAF e creazione account bastion

- **NEWCHIM - Pledge 0 HS06**

Report:

- NTR

- **FAZIA - Pledge 0 HS06**

Report:

- Dato accesso a Silvia Piantelli a 17 macchine su cloud: 16 macchine worker da 16 core e 8GB RAM ciascuna (256 core e 128GB RAM in totale) + 1 bastion-like per storage (1.2TB) e accesso alla sottorete dedicata.

- Le macchine fasia condividono lo storage di 1.2TB via nfs sulla sottorete dedicata.
- Accesso tramite LDAP+Kerberos di sezione.
- La produzione procede senza problemi.

- **PADME - Pledge 0 HS06**

Report:

- NTR

- **AGATA - Pledge 0 HS06**

Report:

Minute skype-call 11/10:

- Macchina con mac address fisso su cui installare la licenza.
- Specifiche macchina:
 - * SO: Linux
 - * accesso interattivo e grafico (x2go, per CAD)
 - * almeno una GPU NVIDIA (aumenterebbe la velocità di calcolo fino a 7 volte)
 - * compilatore CUDA ultima versione dovrebbe andare bene
 - * RAM: 64 GB
 - * Disk: 1-2 TB
 - * mac address fisso per legarci la licenza
- due o più utenti
- previsto un anno di lavoro e poi non dovrebbe più essere necessario
- Licenza non compatibile con cloud, da capire con macchine virtuali.

- **NON-LHCf - Pledge 30 CPU**

Report:

- Creata coda GENERIC su ui-tier1 con share a disposizione 30 CPU
- Lo spazio riservato a LHC-f non si potrebbe utilizzare perchè partiranno a breve delle simulationi che lo occuperanno la gran parte. Richiesta di avere un proprio spazio disco di 1-2 TB.










Figura 1




Figura 2




Figura 3




Figura 4




Figura 5




Figura 6