

# Tier-1 status

Luca dell'Agnello

18 dicembre 2020

# Stato Tier1

- Stato risorse “stazionario”
  - CPU a plegde 2019
    - CPU power Installato 373 kHS06 (pledge 2020: ~405 kHS06)
    - CPU 2019 in consegna a gennaio
    - CPU 2020-21 disponibili fine Q1 2021 (+240 kHS06)
  - Disco 2020 in consegna a gennaio
    - 40.5 PB netti installati (pledge 2020: ~45.5 PB netti)
    - Lotto 1 in consegna a gennaio
    - Lotto 2 (CEPH) in consegna a febbraio (?)
  - Tape a pledge 2020
- Intervento su router di frontiera: 12/1/2021 (19:00-22:00 CET)
  - Ripristino ridondanza router di accesso
  - Non impatta LHCOPN/ONE

Lotto	HS06
2015	49.000
Bari-ReCaS	10.940
2016	92.400
2017	42.000
CINECA1	178.880
2019	35.000
CINECA2	240.000

# News

- Cambio policy Centos → slide Andrea
- Termine supporto Globus toolkit (2022)
  - Gridftp ☹️
- “HTTP as the WLCG baseline protocol for TPC”
  - Xrootd per Alice
  - Migrazione da gridftp entro 2021
    - StoRM in pre-produzione
  - Data challenge in Q2 2021
- Storm (srm) ancora importante per accesso a tape
- Migrazione a token-based AAI (IAM 😊)
  - “Fully X509-free WLCG is a LS3 (tight) target. Tokens and X509 need to interoperate in the next years”

## Proposed Timelines

We all work better with a strawman to argue against:

- **March 2021:** Baseline services for WLCG sites includes HTTP-TPC endpoint.
- **July 2021:** IAM services available, including VOMS endpoint.
- **October 2021:** All WLCG pilot factories have the ability to submit to CE using WLCG tokens.
- **December 2021:** VOMS-Admin shutoff; IAM becomes authoritative identity provider endpoint (including VOMS endpoint).
- **January 2022:** OSG ends support for Grid Community Toolkit. Globus GridFTP support no longer required at WLCG sites.
- **March 2022:** Baseline services for WLCG sites include token support for HTTP endpoints.
- **October 2022:** Rucio transfers performed with token auth in production.
- **March 2023:** Experiments stageout performed via tokens.
- **March 2024:** X.509 client auth becomes optional.
- 2025 serves for schedule contingency.
- **Token transition must be done by 2025** to allow for experiments to complete their other HL-LHC activities.

I don't claim this must be our schedule – but it is along the lines of what we need to coordinate the community.

Brian Bockelman  
GDB 9/12/2020

8 FEARLESS SCIENCE



MORGRIDGE  
INSTITUTE FOR RESEARCH

# Datalakes

Definition of “datalake” in our context: an infrastructure where CPU and storage capacity are loosely coupled (not necessarily co-located).

- Storages in the datalake needs to be connected by a fast and reliable network
- Storage might offer different QoS classes
- A content delivery system might help serving data from storage to CPUs

Mixed statements and experiences from the experiments about caching, latency hiding and remote data access.

All experiments declare interest in QoS but today we prototyped very little (except for disk/tape)

Regional datalake implementation ideas have been presented at DOMA meetings. Focus in 2021 is to prototype those ideas. DOMA ACCESS and QoS WG will merge into a “datalakes WG”, looking after that.



# Archive Storage

## Tape Storage:

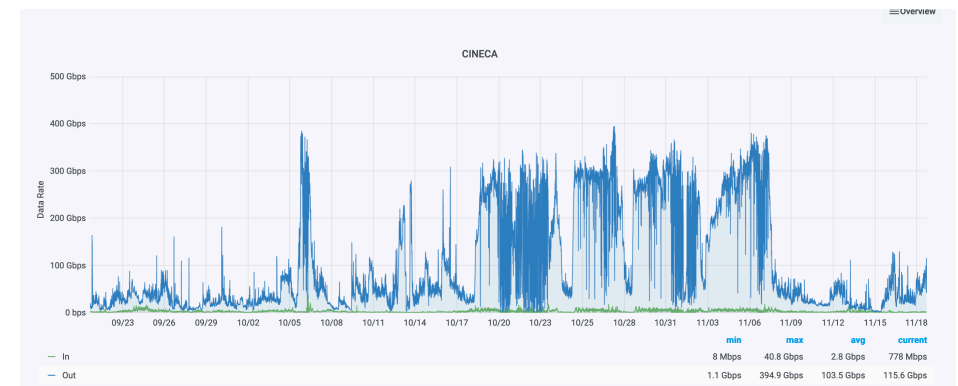
- Three frontend solutions in WLCG: CTA, dCache, StoRM
- In the short term SRM will continue playing a role. FTS should hide the complexity of “stage+transfer” via SRM(dCache,StoRM) or xrootd(CTA)+HTTP
- In the medium term, harmonise the tape access through a common REST API. The dCache bulk request API seems a good way to standardize. Will be followed up in DOMA general meeting in January

## Archive storage does not need to be tape

- Disk-based solution presented by KISTI
- Storage TCO needs to be considered, particularly if the usage will increase (e.g. tape carousels)
- The Archive Storage working group should be revamped. A discussion at one of the next GDBs?

# Stato risorse: CPU

- A pledge 2019
- CPU gara 2019 (35 kHS06) in consegna a Dicembre
- Risorse 2020-21: accordo con CINECA per aggiunta di altri 4 rack analoghi a quelli già in uso (2018)
  - ~240 kHS06 in totale
    - In parte per rimpiazzi gara CPU 2016 (~100 kHS06 in scadenza Q1 2021)
    - In parte per coprire delta pledge 2021 (100 kHS06)
    - In parte per coprire (eventuali) guasti CPU 2018 (180 kHS06 @CINECA)
  - Necessario upgrade infrastruttura di collegamento con CINECA
    - 400 Gbit/s (→ 600 Gbit/s) → 800 Gbit/s
  - Ridondanza percorso ottico non facile ☹️





# Prospetto risorse 2021

Lotto	HS06 tot (HS06)
2015	49.000
<b>Bari-ReCaS</b>	<b>10.940</b>
2016	92.400
2017	42.000
CINECA1	178.880
2019	35.000
CINECA2	240.000

- 2014 da dismettere subito
- Risorse totali 630 KHS06
- 540 KHS06 pledged

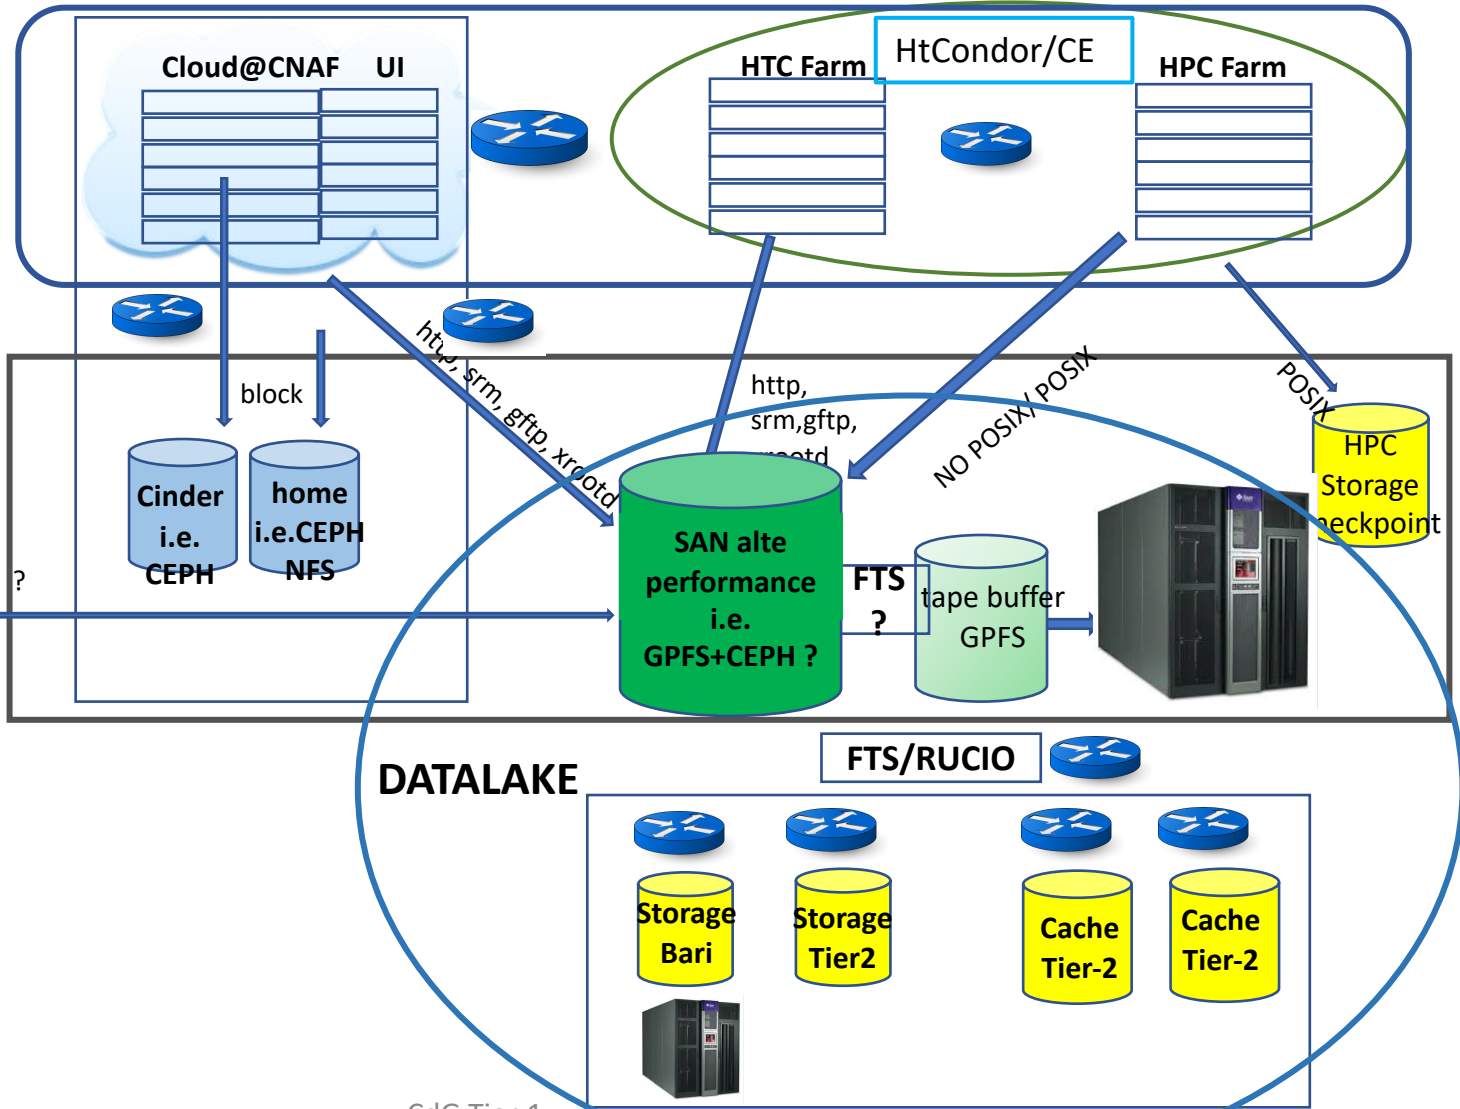

# Problema router di frontiera

- Edge router (GPN)
  - Installato 8/7
  - Aggiornato 1/10 su consiglio di Cisco per risolvere problema di sporadica perdita di pacchetti
    - Durante aggiornamento riscontrati problemi peggiori (perdita di pacchetti quasi totale).
    - Riscontrato problema HW sul componente sospetto (dopo lungo troubleshooting con la casa madre).
    - Componente guasto fornito in sostituzione consegnato e configurato (ora in scorta fredda).
    - In attesa di secondo sw-router da affiancare a quello sostituito per fare la transizione più smooth possibile.
  - Non interessa LHCOPN/ONE
  - Impatto su Tier-1 limitato ad accessi su bastion e trasferimenti via GPN

# Preparazione per il Tecnopolo

- Verifica requirement esperimenti
  - Eventuale discussione (anche) al prossimo CdG
- Risorse di calcolo
  - Necessita' connettivita' (outbound) nodi calcolo
    - Throghput medio/max richiesto per job
  - Prevedete cambiamenti nel modo di utilizzare la farm con il trasferimento al tecnopolo?
    - Es. accesso cloud vs HTC vs HPC
  - pensate possa essere utile un piccolo cluster per job interattivi?
    - jupyter notebook?
  - Uso di GPU
- Accesso ai dati
  - Protocollo Posix?
  - Object storage?

HTC TIER2 HTC Napoli HTC BARI

Experiment	2020			2021		
	CPU	DISCO	TAPE	CPU	DISCO	TAPE
	HS06	TB-N	TB	HS06	TB-N	TB
ALICE	73000	8823	11687	71400	8228	11948
ATLAS	95130	7920	19890	105300	9450	21150
CMS	84500	8840	28600	87100	10010	29900
LHCb	55760	6868	13362	97580	7633	12903
<b>Total LHC TIER1</b>	<b>308390</b>	<b>32451</b>	<b>73539</b>	<b>361380</b>	<b>35321</b>	<b>75901</b>
Belle2	16300	650	0	27133	650	250
CDF	0	0	4000	0	0	4000
FCC	0	0	0	833	100	0
KLOE	0	33	3075	0	33	3075
LCHF	11000	90	0	11000	100	0
Muon Collider	3300	275	200	2500	150	150
NA62	4417	90	1780	3300	275	200
PADME	31708	0	0	4417	90	1780
LHCb Tier2				50976	0	0
<b>TOTALE GRUPPO I</b>	<b>66725</b>	<b>1138</b>	<b>9055</b>	<b>100159</b>	<b>1398</b>	<b>9455</b>
AMS2	19300	2200	1000	21833	2400	1100
ARGO	0	120	1000			
AUGER	2908	615	0	3833	800	150
BOREX	2069	250	51	2069	338	56
CTA	5080	1500	400	5260	1565	400
CUORE	3650	400	0	3942	400	0
CUPID	433	15	10	906	25	10
CYGN0				167	20	20
DAMPE	13834	450	200	17306	550	200
DARKSIDE	4917	1590	300	4917	2290	1770
DUNE				500	15	10
ENUBET	500	10	0	500	10	0
EUCLID	0	1000	1000	0	1000	1000
FERMI/GLAST	833	15	40	833	15	40
GAPS				833	40	0
Gerda	40	50	45	40	60	90
Herd	4167	100	0	6111	200	0
Icarus	4000	1000	2000	4000	1000	2000
JUNO	3000	270	0	3833	620	100
KM3	300	250	200	300	250	200
LHAASO	300	60	0	300	60	0
LIMADOU	483	30	2	747	60	4
LITEBIRD				0	5	5
LSPE	1000	21	0	0	21	0
MAGIC	0	65	150			
NEWS	284	110	110	284	240	110
Opera	200	15	15	200	15	15
PAMELA	733	110	150	747	110	150
TRISTAN	500	10	0	583	40	0
Xenon100	1458	300	2000	1458	300	3000
<b>TOTALE GRUPPO II</b>	<b>69989</b>	<b>10556</b>	<b>8673</b>	<b>81502</b>	<b>12449</b>	<b>10430</b>
Virgo	40000	756	2468	40000	756	8673
<b>TOTALE GRIP+Virgo</b>	<b>109989</b>	<b>11312</b>	<b>11141</b>	<b>121502</b>	<b>13205</b>	<b>19103</b>
ASFIN	542	12	0	1042	212	0
EIC-NET				250	10	0
Famu	2932	23	15	2932	33	15
Fazia		50	0	0	50	0
FOOT	354	60	0	514	80	0
GAMMA/AGATA			1660	0	0	1710
JLAB12	5000	50		4167	50	0
n-TOF	7967	5		8390	10	0
NEWCHIM/FARCOS		30	510	0	30	610
<b>Totale GRUPPO III</b>	<b>16795</b>	<b>230</b>	<b>2185</b>	<b>17295</b>	<b>475</b>	<b>2335</b>
<b>All experiments</b>	<b>501899</b>	<b>45131</b>	<b>95920</b>	<b>600336</b>	<b>50399</b>	<b>106794</b>
<b>All w/ overlap factor</b>	<b>418249</b>	<b>45131</b>	<b>95920</b>	<b>500280</b>	<b>50399</b>	<b>106794</b>
<b>CNAF TOTAL (PLAN)</b>	<b>460000</b>	<b>45131</b>	<b>95920</b>	<b>530000</b>	<b>50399</b>	<b>106794</b>