CMS - CdG T1

Daniele Spiga
INFN-PG

20.06.2025

Tape write rate

Tape overview

- Write rates at some tapes aren't enough to timely consume their backlog.
  - Caution: Transfer load generated on tapes aren't identical, thus the write rates can't be compared
- Tier-0 has only two tapes (FNAL & RAL) to use for the second copy for the rest of the year!
- Should CMS formulate a expected throughput from each tape system?

Tape	Write rate (GB/s)	Current backlog (PB)	Free Space (PB)	Notes
FNAL	4.01	5.1	82	Consumption of backlog got faster after CTA migration
CNAF	0.69	6.6	17.8	Consumption of backlog is slow due to low throughput
JINR	1.88	0	12.7	It was good last month, but having issues nowadays
RAL	0.92	0.01	4.9	
MIT	0.35	1.25	3.75	
IN2P3	1.42	0.03	0.9	
KIT	1.62	0.08	0.8	
PIC	0.5	0.03	0.4	

Il rate più basso è "capito" e spiegato

- **C'è margine di fare qualcosa per alzarlo ?**
- **CMS sarebbe molto contento**

RUN3 targets

VO	Reads (DT) GB/s	Writes (DT) GB/s	Reads (A-DT) GB/s	Writes (A-DT) GB/s
ALICE		0.8	0.3	0.8
ATLAS	0.2	0.9	0.8	0.5
CMS	0.1	1.2	1.9	0.2
LHCb		~~2.24~~ 1.72	~~0.86~~ 1.35	
Total	0.3	4.62	4.35	1.5

spiga@pg.infn.it

Strane requests verso StoRM-WebDAV servers

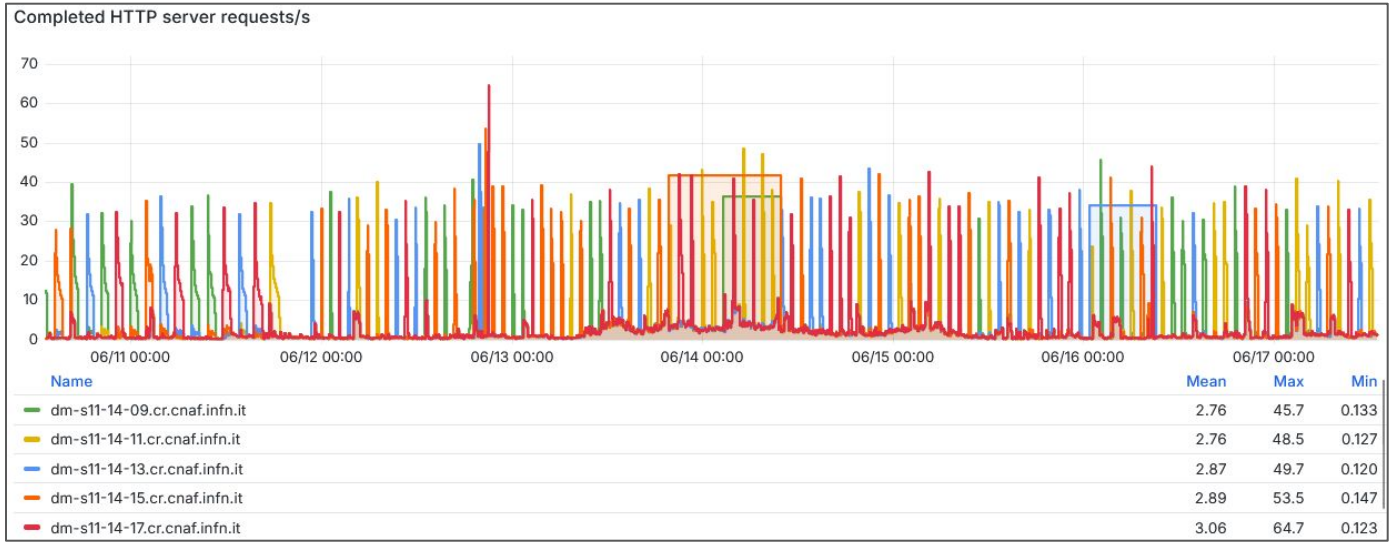

"*I am not exactly sure on what the problem was. However, I briefly deployed some changes that I have in the pipeline, which take a different approach. In case it is not able to delete a directory, it lists and deletes all the content in that directory (sub-dirs and files). Then it gives another try at deleting the desired directory. That is my best suspicious, even though any root file should be listed in the unmerged scan…*"

Periodo di preventivi

Nuove Richieste per il Tier1

Considerando il 13%

RICHIESTE 2026				
T1	pledge 2025	pledge 2026	Incremento	Eur Delta
CPU (kHS06)	143	156	13	130000
DISK (TBN)	18460	21320	2860	286000
TAPE (TB)	57850	70200	12350	123500
				539500

Le dismissioni 2026 per il T1 non sono considerate in queste stime

Resource provisioning a OTF

WLCG Open Technical Forum (OTF) #5

- 24-25 Giugno

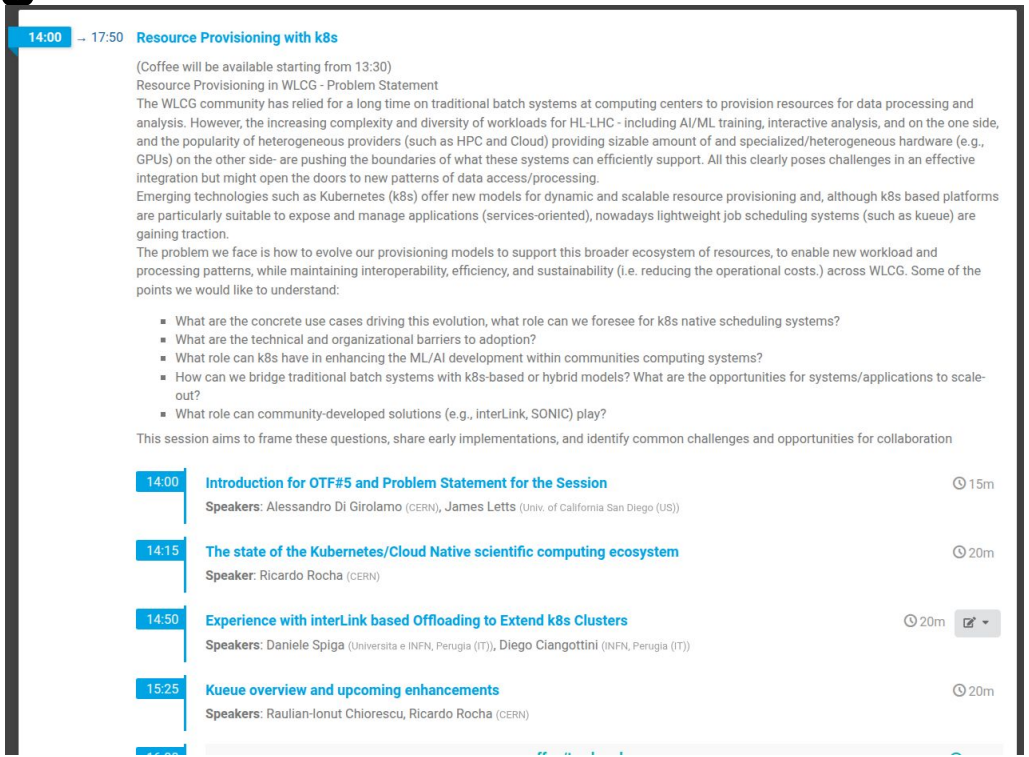

14:00 → 17:50 **Resource Provisioning with k8s**

(Coffee will be available starting from 13:30)
Resource Provisioning in WLCG - Problem Statement
The WLCG community has relied for a long time on traditional batch systems at computing centers to provision resources for data processing and analysis. However, the increasing complexity and diversity of workloads for HL-LHC - including AI/ML training, interactive analysis, and on the one side, and the popularity of heterogeneous providers (such as HPC and Cloud) providing sizable amount of and specialized/heterogeneous hardware (e.g., GPUs) on the other side- are pushing the boundaries of what these systems can efficiently support. All this clearly poses challenges in an effective integration but might open the doors to new patterns of data access/processing.
Emerging technologies such as Kubernetes (k8s) offer new models for dynamic and scalable resource provisioning and, although k8s based platforms are particularly suitable to expose and manage applications (services-oriented), nowadays lightweight job scheduling systems (such as kueue) are gaining traction.
The problem we face is how to evolve our provisioning models to support this broader ecosystem of resources, to enable new workload and processing patterns, while maintaining interoperability, efficiency, and sustainability (i.e. reducing the operational costs.) across WLCG. Some of the points we would like to understand:

- What are the concrete use cases driving this evolution, what role can we foresee for k8s native scheduling systems?
- What are the technical and organizational barriers to adoption?
- What role can k8s have in enhancing the ML/AI development within communities computing systems?
- How can we bridge traditional batch systems with k8s-based or hybrid models? What are the opportunities for systems/applications to scale-out?
- What role can community-developed solutions (e.g., interLink, SONIC) play?

This session aims to frame these questions, share early implementations, and identify common challenges and opportunities for collaboration

14:00 **Introduction for OTF#5 and Problem Statement for the Session** 15m
Speakers: Alessandro Di Girolamo (CERN), James Letts (Univ. of California San Diego (US))

14:15 **The state of the Kubernetes/Cloud Native scientific computing ecosystem** 20m
Speaker: Ricardo Rocha (CERN)

14:50 **Experience with interLink based Offloading to Extend k8s Clusters** 20m
Speakers: Daniele Spiga (Universita e INFN, Perugia (IT)), Diego Ciangottini (INFN, Perugia (IT))

15:25 **Kueue overview and upcoming enhancements** 20m
Speakers: Raulian-Ionut Chiorescu, Ricardo Rocha (CERN)

https://indico.cern.ch/event/1552799/#b-621402-resource-provisioning