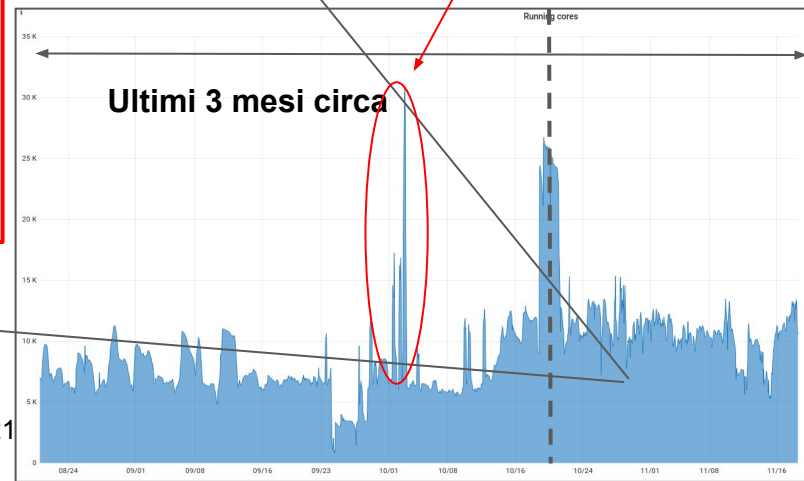



CMS - CdG T1

Daniele Spiga
INFN-PG

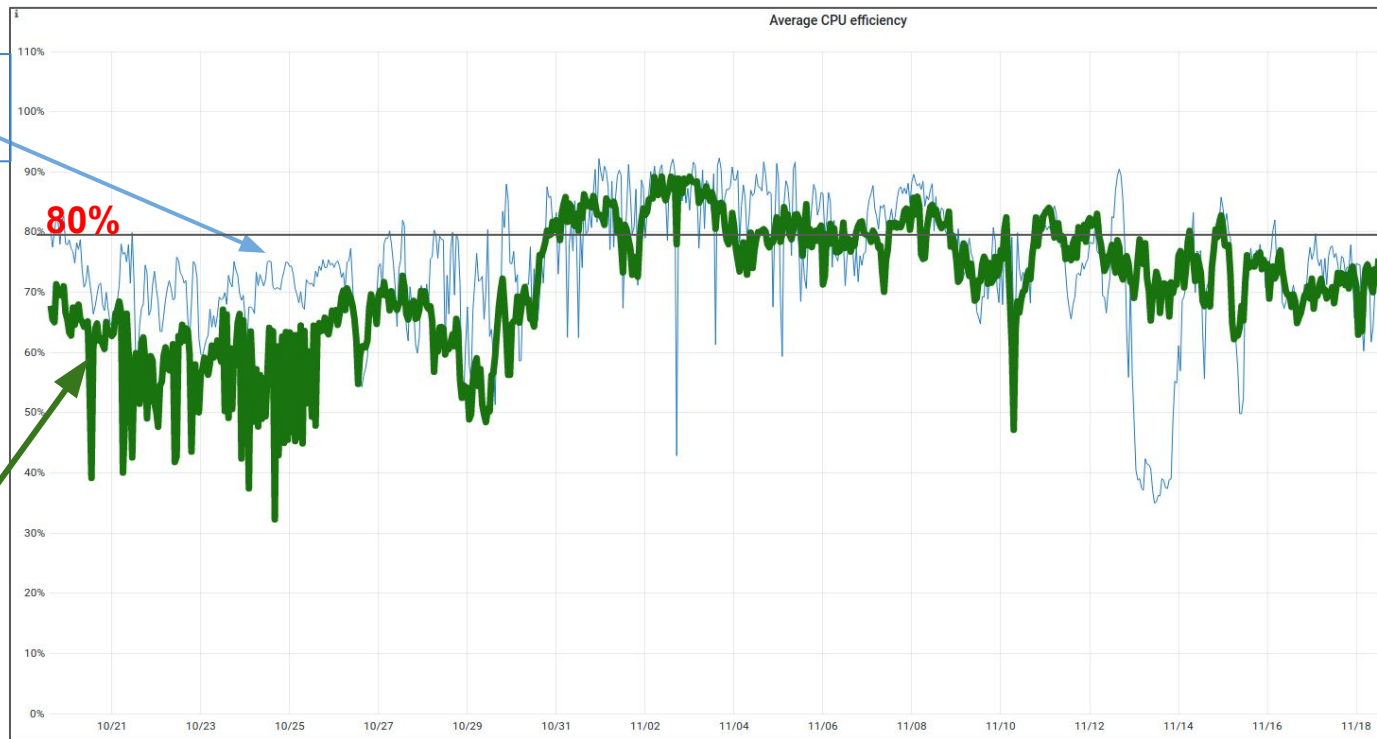
19.11.2021

Utilizzo Risorse Tier1





Efficienza di CPU @CNAF





HTCondor e migrazione a Token: Primo Step

SCITOKEN library per la token based AuthN è stata abilitata sui CE e **CNAF è stato il primo sito EU a fare la validazione** (insieme ad un sito USA) con i pilots di CMS [Grazie a StefanoDP!!]

- Al momento il token si utilizza solo per l'AuthN con il pilot e **i payload sul WN lavorano con i proxy x509**

I test fatti da CMS sono ancora “manuali”, ora Submission Infrastructure configurerà l'integration TB (ITB) creando un gruppo per testare sistematicamente tutti i siti pronti

- Sarà necessario iniziare con una situazione ibrida (non tutti sono pronti)
- **CNAF sarà inserito nella lista**

Executive: come **T1_IT_CNAF è pronti**. Quindi ora l'obiettivo è vedere dove è il prossimo problema :)



Debugging XRootD setup

Reminder: Abbiamo riscontrato un problema “un po subdolo” da debuggare che si manifesta nella lettura remota di dati via xrootd pool.

Ci sono tre effetti lato client:

- Max Thread reached
- Connection reset by peer
- Timeout

Alzare banalmente le soglie di protezione sposta il problema sul load dei server di xrootd





- Teoricamente dovremmo capire perchè (non banale ed è work in progress)
- E' anche in corso un confronto tecnico con il CERN. [Grazie Lucia/Vladimir!! E tutto storage]

```

botd_log-20210824-210826 02:31:22 891 XrdLink: Unable to receive from tboccali.114525:138@login.m100.cineca.it; connection reset by peer
botd_log-20210824-210826 02:38:19 24570 XrdLink: Unable to receive from tboccali.117013:96@login.m100.cineca.it; connection reset by peer
botd_log-20210824-210826 02:40:04 24606 XrdLink: Unable to receive from tboccali.121358:97@login.m100.cineca.it; connection reset by peer
botd_log-20210824-210823 04:06:24 10881 XrdLink: Unable to receive from cms001.3048:139@wna520.jlnr-t1.ru; connection reset by peer
botd_log-20210824-210823 04:08:54 31260 XrdLink: Unable to receive from tardis.2390:65@f03-001-163-e.gridka.de; connection reset by peer
botd_log-20210824-210823 05:15:10 24422 XrdLink: Unable to receive from ?:253@lxslc713.lhep.ac.cn; connection reset by peer
botd_log-20210824-210823 06:11:01 27747 XrdLink: Unable to send to cmspilot.271:349@[::134.158.73.173]; connection reset by peer
botd_log-20210824-210823 13:32:57 32591 XrdLink: Unable to receive from pilcms01.2358:409@grid311.kfkl.hu; connection reset by peer
botd_log-20210824-210823 21:54:39 3351 XrdLink: Unable to send to cms001.4278:491@[::159.93.230.2]; connection reset by peer
botd_log-20210824-210823 23:05:58 32538 XrdLink: Unable to send to siluo.227:437@[2001:1458:301:7d::100:6d]; connection reset by peer
botd_log-20210824-210823 23:06:00 7319 XrdLink: Unable to receive from siluo.227:274@[2001:1458:301:f9::100:6e]; connection reset by peer
botd_log-20210824-210823 23:06:12 980 XrdLink: Unable to receive from siluo.227:316@b7s14n1799.cern.ch; connection reset by peer
botd_log-20210824-210823 23:20:21 1083 XrdLink: Unable to receive from siluo.227:338@b7s04p4014.cern.ch; connection reset by peer
botd_log-20210824-210824 00:41:11 24422 XrdLink: Unable to receive from cmspilot.375:50@[::129.59.197.42]; connection reset by peer
botd_log-20210825-210824 09:05:50 23650 XrdLink: Unable to receive from scc-sdm-.1796:111@[2a00:1398:4:182a::a00:1d7]; connection reset by peer
botd_log-20210825-210824 11:01:48 18481 XrdLink: Unable to receive from pcms02.2333:54@[2a00:139c:4:cd:0:1:5:71]; connection reset by peer
botd_log-20210825-210824 11:23:07 13349 XrdLink: Unable to receive from pcms02.2398:428@c01-122-114.gridka.de; connection reset by peer
botd_log-20210825-210824 17:29:19 31495 XrdLink: Unable to receive from pilcms12.375:690@t2-fw-01-w.lnl.infn.it; connection reset by peer

```

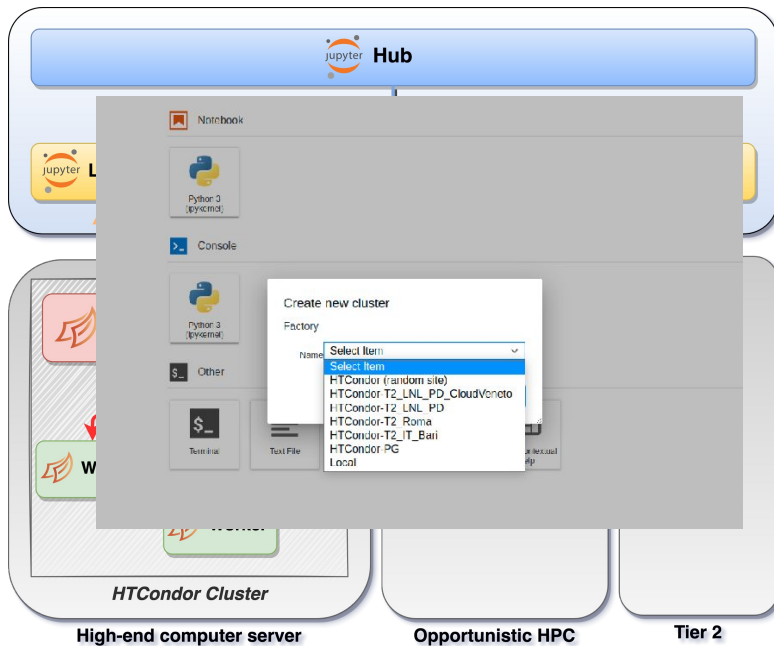
Abbiamo fatto un po di analisi dei logs per capire l'entità, non banale e forse non troppo affidabile



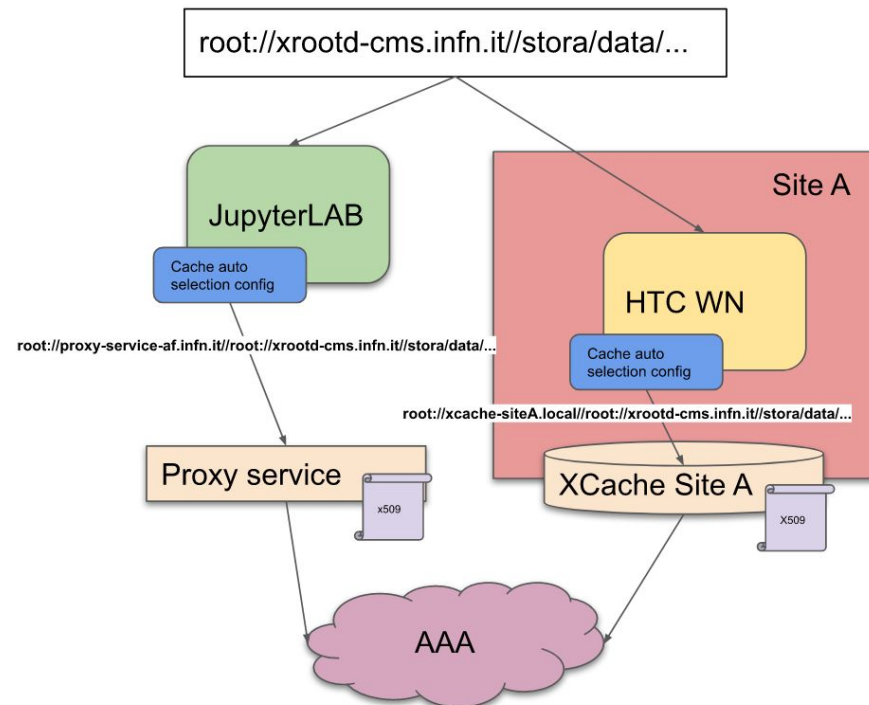
R&D su infrastruttura per l'analisi

Cosa: abilitare un sistema per l'analisi interattiva e fare misure!

<https://cms-it-hub.cloud.cnaf.infn.it/>



Data Flow : work in progress





Sottomesso Proposal PRACE

“PRACE 24th Call for Proposals for Project Access - proposal No. 2021250016 “
per 8MCoreH (1M GPU hours)

- Sottomesso, superato la validazione amministrativa e ora siamo in attesa della review Tecnica e Scientifica

processed. The PowerAtCMS project has a twofold ambition:

- use Marconi 100 for the CMS Experiment, by replicating the infrastructure we used for the PRACE Project Access Grant #2018194658 on Marconi A2 [lhc@hpc];
- implement a mechanism to logically partition single nodes with “mostly CPU bound” workflows with “mostly GPU bound” ones; the two workflows collaborate towards a common goal, fully exploiting the node capabilities;
- an End-to-End Deep Learning reconstruction workflow is trained to substitute the most time consuming classical algorithms in CMS: the reconstruction of particle showers in the Endcap Hadron Calorimeter Reconstruction (HGCal), their accurate simulation when generating synthetic datasets for physics studies, and the linking of these clusters to those from other detector components, in the global event reconstruction of CMS events, based on the concept of particle flow.



M100 sul fronte tecnico

Stiamo integrando la parte GPU. Per ora stiamo lavorando al setup del pilot e siamo vicini alla meta

Site	Entry	AssignedGPUs	CUDACapability	CUDALockMhz	CUDAComputeUnits	CUDAcoresPerCU	CUDADeviceName	CUDADriverVersion	CUDAEECCEnable	CUDAGlobalMemoryMb	Machine
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r241n05.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r246n08.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r226n14.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r217n14.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r243n03.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r226n16.m100.cine.
T2_US_Florida	CMSHTPC_T2_US_Florida_condor...	CUDA7	8.0	1410	108	64	A100-SXM-80GB	11.0	true	81252	c1000a-s11.uhfp.c
T2_US_Wisconsin	CMSHTPC_T2_US_Wisconsin_cm...	CUDA0	8.0	1410	108	64	A100-SXM4-40GB	11.2	true	40536	oconnorgpu2000.ch
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r250n01.m100.cine.
T1_IT_CNAF	CMSHTPC_T1_IT_CNAF_CINECA...	CUDA0,CUDA...	7.0	1530	80	64	Tesla V100-SXM2-16...	11.0	true	16128	r247n17.m100.cine.
T2_US_Purdue	CMSHTPC_T2_US_Purdue_Hamm...	CUDA0	7.5	1590	40	64	NVIDIA Tesla T4	11.3	true	15110	hammer-f015.rcac...
T2_US_Caltech	CMSHTPC_T2_US_Caltech_cit2_g...	CUDA0,CUDA...	6.1	1734	20	128	NVIDIA GeForce GTX...	11.4	false	8120	imperium-sm.hep.c.
T3_US_NotreDa...	CMS_US_ND-CAML_gpu	CUDA0,CUDA3	7.5	1620	72	64	Quadro RTX 6000	11.2	true	22699	qa-rtx6k-034.crc.nd.

Abbiamo preparato tutto per la “validazione della fisica”

- oramai è alle porte, siamo in attesa che CMS ci faccia l'injection del workflow

spiga@pg.infn.it

WHAT:

- all the standard 0 PU relvals
- processing data should be ok
- you need to specify the proper SCRAM_ARCH, and the release should exist for PPC. We are using now CMSSW_11_2_0 and slc7_ppc64le_gcc9 so it would be a good starting point

WHERE:

- the testbed should be used at this point (see * for the complete command we use)
- target should be T1_IT_CNAF

CdG - team should be testbed-vocms0263