

Tier-0 workflows					
CMS Online Services: DAQstatus <a href="#">↗</a> , SMAalarms <a href="#">↗</a> , P5 Elog <a href="#">↗</a>	Tier-0 Service : T0 T0Mon <a href="#">↗</a>	T0export Castor Pool : Load&Network <a href="#">↗</a> , Total Space <a href="#">↗</a> , Active Transfers <a href="#">↗</a> , Queued Transfers <a href="#">↗</a>	T1transfer Castor Pool : Load&Network <a href="#">↗</a> , Active Transfers <a href="#">↗</a> , Queued Transfers <a href="#">↗</a>	cmst0 LSF Pool : Load&Network <a href="#">↗</a>	Tier-0 Jobs : cmst0 Jobs <a href="#">↗</a>

---

## Tier-0 workflows monitoring

### Introduction / Machine status

**Please read this once at the start of your shift** to learn about the Tier-0 workflow. This introduction will help you to understand the importance of the different components of the workflow and which problem to look for.

- A general T0 description is provided here [▢ Hide](#)
  - ◆ The T0 is one of the most important computing systems of CMS, it's responsible for creating the RAW datasets out of the data stream sent from the pit. It also handles the first reconstruction pass of the RAW data called PromptReco. It runs many kinds of jobs against the collision data, among them the most important types of Jobs are EXPRESS and REPACK. EXPRESS jobs speedily reconstruct a special portion of the RAW data coming from the detector and are supposed to finish within 1 hour of recording of this data. REPACK jobs process all the data coming from the pit and convert the data into RAW files and split them into Primary datasets. These jobs should run in real time, a delay impacts all subsequent workflows. For example online shifters for the detector subsystems can't work if these jobs get delayed. The main problems that can be encountered are stuck transfers from P5 or Express and/or Repack jobs failing causing runs to get stuck within the T0 processing chain. NEW
  - ◆ As you probably have read in the Computing Plan of the Day [↗](#), you already know if we are during data-taking period or not. When we are, any error in the T0 should be reported. We should not have runs delayed.
- A drawing of the main Tier-0 components can be found below ( **clickable boxes on the picture may be used**)

## CMS Online Services

Check the CMS Online status related to **incoming runs** --- `%COMPLETE5%`

During LHC collisions, periodically checkpoint the T0Mon monitoring below against the ongoing data taking status seen here.

- DAQ Status URL: <http://cmsonline.cern.ch/daqStatusSCX/aDAQmon/DAQstatusGre.jpg> (shift reload this periodically)
  - ◆ Show INSTRUCTIONS  Hide  INSTRUCTIONS: This overview indicates whether a run is currently being ongoing and specifies the data taking mode and if data is being transferred from the CMS detector to the Tier-0. In the top center of the page you can find the on-going Run Number; a green field on the bottom right indicates "TIER0\_TRANSFER\_ON" when data is sent to the Tier-0. The first line under "Data Flow" on the top right specifies the data taking mode (or "HLT key", which is a string containing a tag such as "interfill", "physics", ...). The tag "physics" is the most relevant data. The bottom left histograms shows the run history in the last 24h; the run numbers specified on the graph should be reflected in the T0Mon page, see Monitoring links below.
- Storage Manager URL: <http://cmsonline.cern.ch/portal/page/portal/CMS%20online%20system/Storage%20Manager>
  - ◆ Show INSTRUCTIONS  Hide  INSTRUCTIONS: Check on the DAQ page above (see URL on previous bullet) that the latest run number in Storage Manager matches the current run number (if T0\_Transfer is on and some data has been logged). If T0\_Transfer is on and the run number is not the latest, open an Elog. The Storage Manager also shows all the transfer/preliminar processing chain. close -> inject -> transfr -> check -> repack is what matters, if you see missing files ( next number is lower than the previous ) and it is red, open an elog. (if it's blue, it's expected). If you see anything red around the page, ex: Server down, open an Elog.
- Prompt Calibration Loop URL [URL](#)
  - ◆ Show INSTRUCTIONS  Hide  INSTRUCTIONS: This monitoring page (look only at **Latency since the end of the run**) controls whether Express workflows delivered the calibration payload to the online systems(black dot correspondent to the run on the axis X). This has to happen before PromptReco starts (red dotted line). Runs before (left side) the red

line, can be ignored. If the red line gets too close (3 runs before) to the last run that uploaded conditions (has the black dot), open an Elog to warn the Tier-0 Experts.

- P5 Elog URL [↗](#) (for information only, see instructions)
  - ◆ Show INSTRUCTIONS  Hide  INSTRUCTIONS: This elog is not needed in "normal" situations, however it may be useful in cases of very special events at the CMS detector. Or you may use it to simply find out who is the online shifter (== the shift role corresponding to yours, but for every thing related to online data taking). You will need to log in with your AFS credentials.

## Tier-0 Service

Checks of the **the most urgent alarms of the Tier-0** --- %COMPLETE5%

The main causes of runs getting stuck are monitored: **Failed jobs, Stuck jobs** and in case of overload, **Backlog**.

- Tier-0 Alarms URL : <http://cmsprod.web.cern.ch/cmsprod/sls/> [↗](#)
- Show INSTRUCTIONS  Hide  INSTRUCTIONS:
  - ◆ General assumptions
    - ◇ Green bar - check is fine
    - ◇ Yellow bar - check is failing, there's no "red", or worse state. It's binary.
    - ◇ There are no defective plots, values = 0 are fine
    - ◇ Report if a monitoring element shows a yellow bar in the Elog in the T0 Processing category.
  - ◆ Permanent Failures
    - ◇ This is the most important check in this category. It checks whether there have been any failed jobs in the last 24h. The information corresponds to the counters in the T0mon page (see below), it will turn yellow once a single job failure was detected, it will not change back to green until an expert fixes/resubmits the job. Even if it's yellow, check if the number of failures are increasing (look at the plots [↗](#)). If you see new failures, and there's no known reason(check Elog), open an Elog.
  - ◆ Backlog
    - ◇ This shows how much jobs are queued in the Tier-0 system and are \*still not submitted\* to the batch system, so you won't see them in the "running/pending" jobs plot [↗](#). Thresholds have been specified in the alarm page [↗](#), for each kind of job. If you see any job type accumulating above the thresholds, open an Elog.
  - ◆ Long jobs
    - ◇ This is mostly an indicator if everything is normal or not. Jobs can take longer than expected, but if they take a lot longer, it is possible that they are stuck due to an infrastructure problem. If you see the alarm going yellow, open an Elog.
  - ◆ Cluster health
    - ◇ Skip for now, this is a new alarm under commissioning. Experts are taking care of this alarm for now.

Check the **status of Tier 0 components and workflows through T0Mon** --- %COMPLETE5%

- T0Mon URL: <https://cmsweb-testbed.cern.ch/T0Mon/> [↗](#)
  - ◆ Show INSTRUCTIONS  Hide  INSTRUCTIONS:
    - ◇ Check that T0 components are up: Look at the "Time from component's last heartbeat" section near the top of the page. If one or more of the components show up in red as below, open an elog in the T0 category.
    - ◇ Check for new runs appearing: New runs should appear on T0Mon in the "Runs"

table within a few minutes of appearing on the Storage Manager page above (excluding TransferTest or local runs with labels such as privmuon). If a new run doesn't appear on T0Mon, or shows up with an HLTKey of "UNKNOWN", open an elog.

INSTRUCTIONS:

- For each of the URL1 links above, check the "load average" pie chart on the left-hand side : if you see that some hosts have a load average higher than 20, please open an ELOG in the "T0" category.

Check also the network utilization plot on the bottom right-hand side of URL1s : if you see a sustained throughput at the resource-saturation plateau (that is # hosts x 100 Mbyte/s) open an ELOG in the "T0" category.

◆ URL2s: Total space on t0export

[https://sls.cern.ch/sls/history.php?id=CASTORCMS\\_T0EXPORT&more=nv:Total+Space+TB&period=](https://sls.cern.ch/sls/history.php?id=CASTORCMS_T0EXPORT&more=nv:Total+Space+TB&period=)

◇ Show INSTRUCTIONS ▢ Hide ▢ INSTRUCTIONS:

- The total space available on t0export should be constant. Watch out for sudden drops, as these indicate a missing disk server as shown below. If this happens, make an ELOG post in the T0 category.

• URL3s: Active Transfers↗/Queued Transfers↗ on t0export pool .

◆ Show INSTRUCTIONS ▢ Hide ▢ INSTRUCTIONS:

- ◇ There can be up to a few thousand active transfers but the number of queued transfers should not go above 200. If this happens, please make an ELOG post in the T0 category.

• URL5s: Active Transfers↗/Queued Transfers↗ on t1transfer pool.

• Show INSTRUCTIONS ▢ Hide ▢ INSTRUCTIONS:

- ◆ There can be up to a thousand active transfers but the number of queued transfers should not go above 200. If this happens, please make an ELOG post in the T0 category.

---

Check the **activity on the Tier-0 LSF WNs** --- %COMPLETE5%

• URL1: Load average summary and network utilization of cmst0 batch cluster↗

- ◆ Show INSTRUCTIONS ▢ Hide ▢ For URL1 above, check the CPU Utilization and the "load average" pie chart on the left-hand side : if you see a sustained CPU utilization above 95%, open an ELOG in the "T0" category.

---

Check the **Tier-0 Jobs** --- %COMPLETE5%

• URL: Queued and running Jobs on cmst0↗

- Show INSTRUCTIONS ▢ Hide ▢ The cmst0 queue currently has ~3500 job slots. We should use all of them, if there are jobs to run(backlog↗). If there are more than 3500 pending+running jobs, but much less than 3500 running, it might be a LSF issue, report in the T0 section of the Elog.

---

This topic: Sandbox > CompOpsTest

Topic revision: r1 - 2013-01-11 - DiegoBallesterosVillamizar



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback