

## 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.

- General description of the Tier-0 workflow [▢](#) 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
  - ◆ Please check the Computing Plan of the Day [☞](#) if the LHC is colliding beams and we are taking data. In this case, any Tier-0 error should be reported to avoid delaying processing of new runs.
- A drawing of the main Tier-0 components can be found below ( \*clickable boxes on the picture may be used) \*)

The status of the ongoing data taking can be seen at the 2 following locations : NEW

- 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.

- P5 Elog URL [↗](#) (for information only, see instructions)
  - ◆ Show INSTRUCTIONS [▶](#) Hide [▣](#) INSTRUCTIONS: This elog is not required to be followed by the CSP in "normal" situations, however it may be useful in cases of very special events at the CMS detector. You may use it to find out who is the current 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

### #SLS

Checks of the *\*the most urgent alarms of the Tier-0\**. The main causes of runs getting stuck are monitored: *\*Failed jobs\**, *\*Stuck jobs\** and in case of overload, *\*Backlog\**. NEW

- 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.
    - 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 experts fixes/resubmits the job. If the failure is not recoverable, the field will change to back to green automatically green after 24h passed by and if there have been no new job failures. If this alarm goes yellow, it can stay yellow during all your shift, but you should always check if the counters [↗](#) are increasing (take a look in the plots [↗](#)). If you see new unknown failures, open an Elog (or answer an existing one). If you see *\*new\** Express/Repack failing with no known reason, you should open an Elog and Call the CRC.
  - ◇ 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. Express and Repack should not accumulate a backlog, PromptReco is ok. If you see a backlog above the thresholds on Express and Repack open an Elog, PromptReco above threshold is still fine. The historical information can be found here [↗](#), so you can see how the backlog developed over time.
  - ◇ 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.cern.ch/T0Mon/>
  - ◆ Show INSTRUCTIONS ▢ Hide ▢ INSTRUCTIONS:
    - ◇ Check that all 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.

---

Check the CMS Online status:

- \*Incoming Runs \* --- %COMPLETE5%

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

\*

- ◆ Storage Manager URL:  
<http://cmsonline.cern.ch/portal/page/portal/CMS%20online%20system/Storage%20Manager>
  - ◇ Show INSTRUCTIONS ▢ Hide ▢ INSTRUCTIONS: Check on the DAQ page (<http://cmsonline.cern.ch/daqStatusSCX/aDAQmon/DAQstatusGre.jpg>) that the latest run number in Storage Manager matches the current run number (if T0\_Transfer is on). If T0\_Transfer is on and the run number is not the latest, open an Elog. The Storage Manager also shows all the transfers/preliminary 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.

---

Check the \*Castor pools/subclusters used by CMS: focus on t0export and t1transfer \* --- %COMPLETE5%

- URL1s: load average summary and network utilization : [t0export](#), [t1transfer](#)
  - ◆ Show INSTRUCTIONS ▢ Hide ▢ 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 10, 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=week&title=](https://sls.cern.ch/sls/history.php?id=CASTORCMS_T0EXPORT&more=nv:Total+Space+TB&period=week&title=)
    - 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<sup>↗</sup>/Queued Transfers<sup>↗</sup> 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<sup>↗</sup>/Queued Transfers<sup>↗</sup> 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 WN's \* --- %COMPLETE5%

- URL1: CPU and Network utilization of cmst0 batch cluster<sup>↗</sup>
- URL1bis: CPU and Network utilization of cmsproc batch cluster<sup>↗</sup>
- URL2: load-average history of cmst0 hosts<sup>↗</sup>
- URL2bis: load-average history of cmsproc hosts<sup>↗</sup>
- 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% and/or that some hosts have a load average higher than 12 in sustained manner (see example of a problematic node in image below), then open an ELOG in the "T0" category.

. Check also the network utilization plot on the bottom right-hand side of URL1.

---

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

- URL1: Queued and running Jobs on cmst0<sup>↗</sup>
- Show INSTRUCTIONS  Hide  The cmst0 queue is using ~3200 job slots. If you see on URL1 a sustained saturation of running (green at 2.8k) or pending (blue) jobs like in the example shown in the picture below, it might not be an issue, however it is worth notifying the DataOps team via the ELOG

in the "T0" category.

---

This topic: Sandbox > SamirCurySandboxCSP2

Topic revision: r2 - 2012-03-19 - SamirCury



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](#)