Tier-2 @ UCSD

Frank Würthwein
UCSD

Role of T2
Infrastructure Requirements
Supporting MC production
Supporting Analysis

Role of T2

- Production of all the MC the CMS experiment needs.
- Support of Data Analysis
  - Host official Data
  - Accept jobs from all CMS users
  - Host inofficial data for “local” community
    - “local” may mean many things …

Infrastructure Requirements

- Operate a compute cluster
- Operate a storage cluster
- Operate grid interfaces to both
- Operate a Squid Server
- Operate PhEDEx installation
- May or may not operate a login facility for your local users.
- May or may not maintain cvs, twiki, etc. for your local users.

Effort Required to do this well

- Ideally, 2FTE
  - One system administrator
    - All hardware & grid interfaces etc.
  - One computing savy physicist
    - All CMS services, incl. Data Management, and debugging support.
- At least at startup of T2 it is very important to have the necessary expertise and focus !!!

Supporting MC Production

- CMSSW installation
  - Is done centrally by CMS
  - Need local person to
    - follow up on issues !!!
    - Configure TrivialFileCatalog
- SRM for local stage-out of MC produced
  - /store/unmerged = small files locally produced
  - /store/mc = datasets of merged or downloaded files
- EGEE or OSG Compute Element
- FroNtier Squid
- PhEDEx & upload link to regional T1.

CMSSW Installation

- CMS has a list of officially supported releases.
- We distribute all of them to all T2s.
- You may find that some of your local developers need additional releases, e.g. pre-releases.
- Do yourself and us a favor and NEVER mess with the official release area.
- Instead have an additional local release area for your local people, and maintain that yourself as you see fit.

TFC

- CMS has a global namespace, prefixed by a local namespace. Any physical filename is thus something like:
  - /pnfs/sdsc.edu/data3/cms/phedex/store/mc/blabla.root
  - /pnfs/sdsc.edu/data3/cms/phedex/store/ is UCSD specific.
  - /mc/blabla.root would be the logical file name (lfn).
- TFC is a set of rules to convert LFN to PFN and back
- xml file: SITECONF/$site/PhEDEx/storage.xml
- tags:
  - lfn-to-pfn
  - pfn-to-lfn
  - protocol=["direct"|"srm"|"dcap"]
- Examples of a $site can be checked out of CVS:

```
export CVSROOT=:pserver:anonymous@cmscvs.cern.ch:/cvs_server/repositories/CMSSW
cvs -d `echo $CVSROOT | awk -F@ '{print $1":98passwd\@"$2}'` login
cvs co SITECONF/$site
```

- After you create your TFC, check it into CVS

TFC

```
<storage-mapping>
  <lfn-to-pfn protocol="direct"
    path-match=".*/LoadTest07_UCSD_(.*)_.*_.*"
    result="/pnfs/sdsc.edu/data3/cms/phedex/store/PhEDEx_LoadTest07/LoadTest07_Prod_UCSD/LoadTest07_UCSD_$1"/>
  <lfn-to-pfn protocol="srm"
    path-match=".*/LoadTest07_UCSD_(.*)_.*_.*”
    result="srm://t2data2.t2.ucsd.edu:8443/srm/managerv1?SFN=/pnfs/sdsc.edu/data3/cms/phedex/store/
PhEDEx_LoadTest07/LoadTest07_Prod_UCSD/LoadTest07_UCSD_$1"/>

  <lfn-to-pfn protocol="dcap" destination-match=".*"
    path-match="/+store/(.*)" result="dcap://dcopy-1.local:22138//pnfs/sdsc.edu/data3/cms/phedex/store/$1"/>
  <lfn-to-pfn protocol="direct" destination-match=".*"
    path-match="/+store/(.*)" result="/pnfs/sdsc.edu/data3/cms/phedex/store/$1"/>
  <lfn-to-pfn protocol="srm" destination-match=".*"
    path-match="/+store/(.*)"
     result="srm://t2data2.t2.ucsd.edu:8443/srm/managerv1?SFN=/pnfs/sdsc.edu/data3/cms/phedex/store/$1"/>

  <pfn-to-lfn protocol="dcap" destination-match=".*"
    path-match="dcap://dcopy-1.local:22138//pnfs/sdsc.edu/data3/cms/phedex/store/(.*)" result="/store/$1"/>
  <pfn-to-lfn protocol="direct" destination-match=".*"
    path-match="/+pnfs/sdsc.edu/data3/cms/phedex/store/(.*)" result="/store/$1"/>
  <pfn-to-lfn protocol="srm" destination-match=".*"
    path-match=".*\?SFN=/pnfs/sdsc.edu/data3/cms/phedex/store/(.*)" result="/store/$1"/>

</storage-mapping>
```

Supporting Analysis

- Need everything that MC Prod needs
- In addition, need download links from ALL T1s !!!
  - No one T1 will have all the data!
  - Future, roughly 40% of data at FNAL, rest distributed across other 6 T1s.
  - Today most but not all the data available at both CERN & FNAL.
- Need one person on team who knows how to use CRAB !!! It is unacceptable for a T2 support team to not know how to submit a user analysis job to your site !!!
- Host data => under control of site data manager !!!

Name Space Organization

- /unmerged
  - Used by MC production team only
  - They should clean it up automatically!
  - If they don't, please complain !!!
- /users
  - Use for storage area for local users
  - In future also likely to be used for user analysis stageout.
    - We're not yet ready to do this.
- All other areas in namespace under the control of the T2 data manager at your site.

Access to Storage

- All users of CMS need read access to the full CMS namespace.

  voms-proxy-init -voms cms
- MC Production team needs write access to /mc & /unmerged
  - This is identified via /cms/Role=production or /cms/Role=cmsprod
- Local PhEDEx agent needs write access to / (except maybe /users and /unmerged)
- CMS namespace should not be world readable.

Data Manager Responsibility

- PhEDEx is based on a request system.
  - Any CMS user may request data to be moved to your site.
  - Either central data ops or your site's data manager can accept such a request.
- You can move either full datasets or full blocks. I.e. you can satisfy part of a request.
- Let's see how this works:

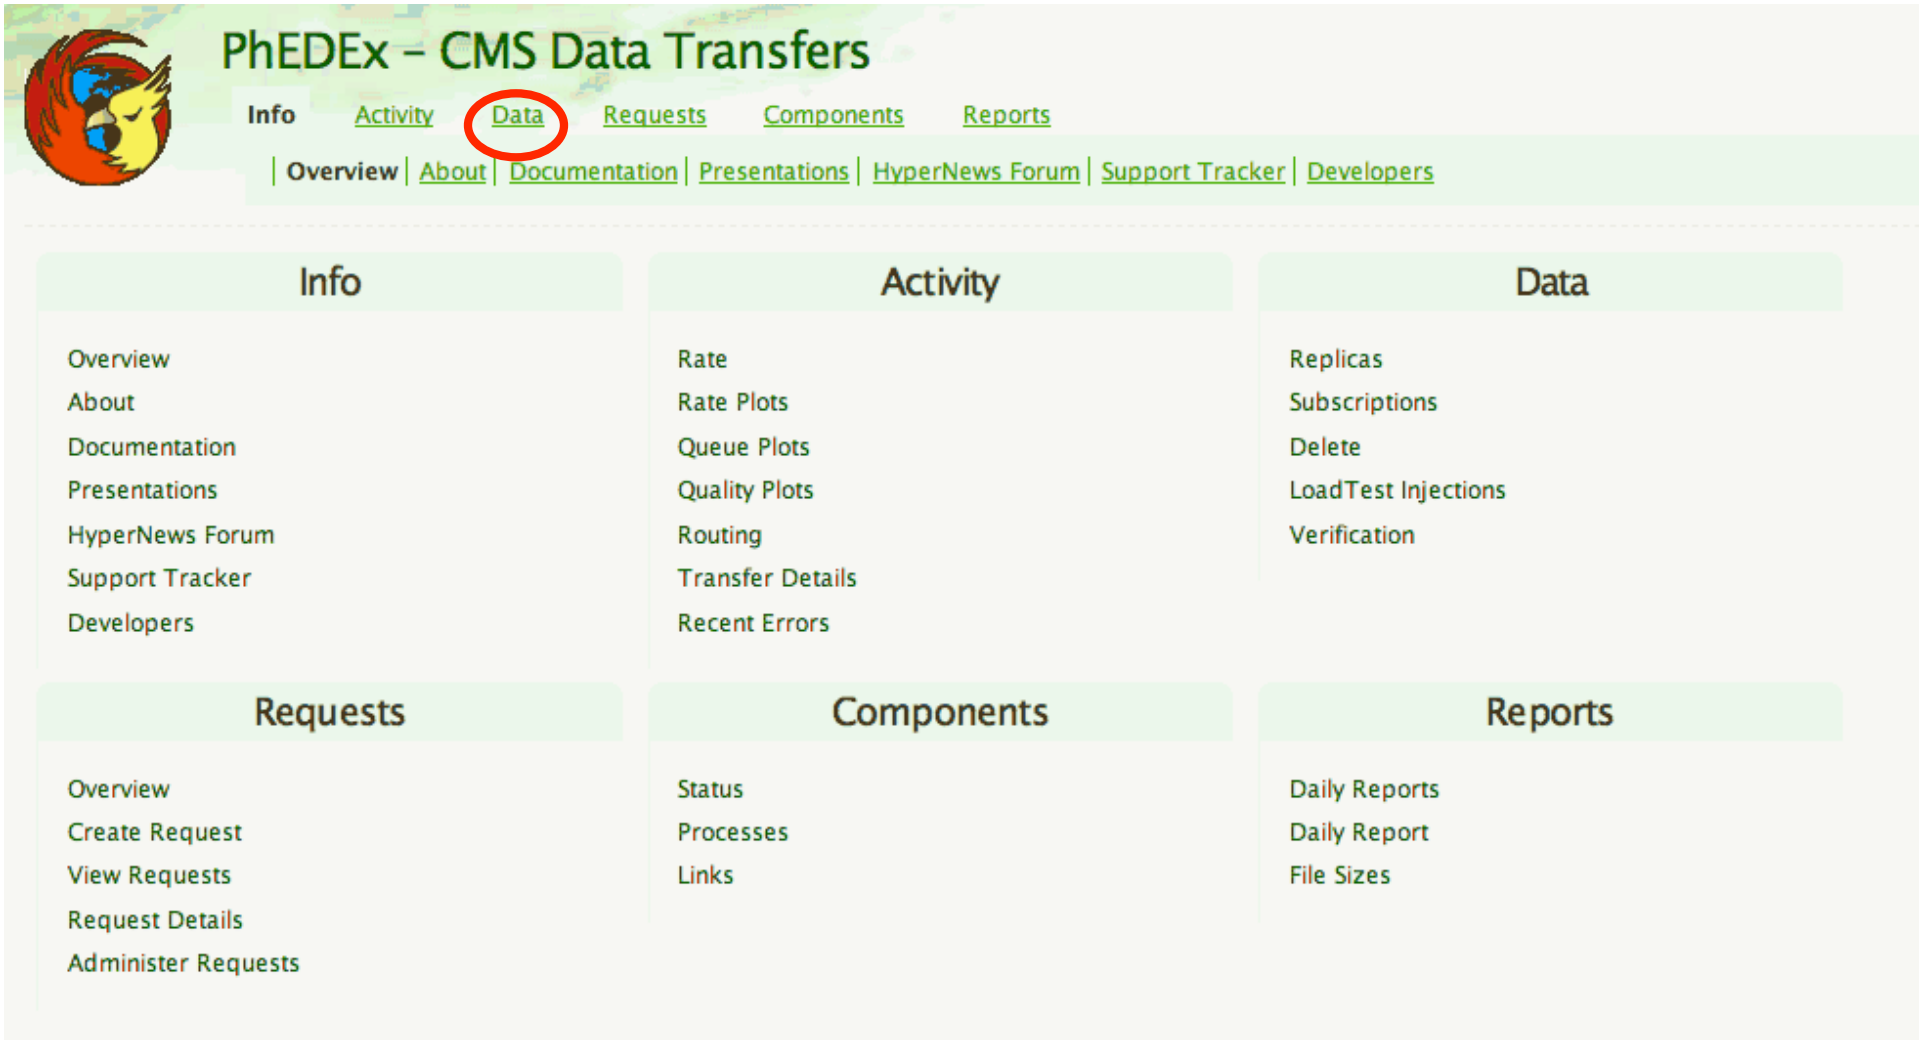

http://cmsdoc.cern.ch/cms/aprom/phedex/

Click here

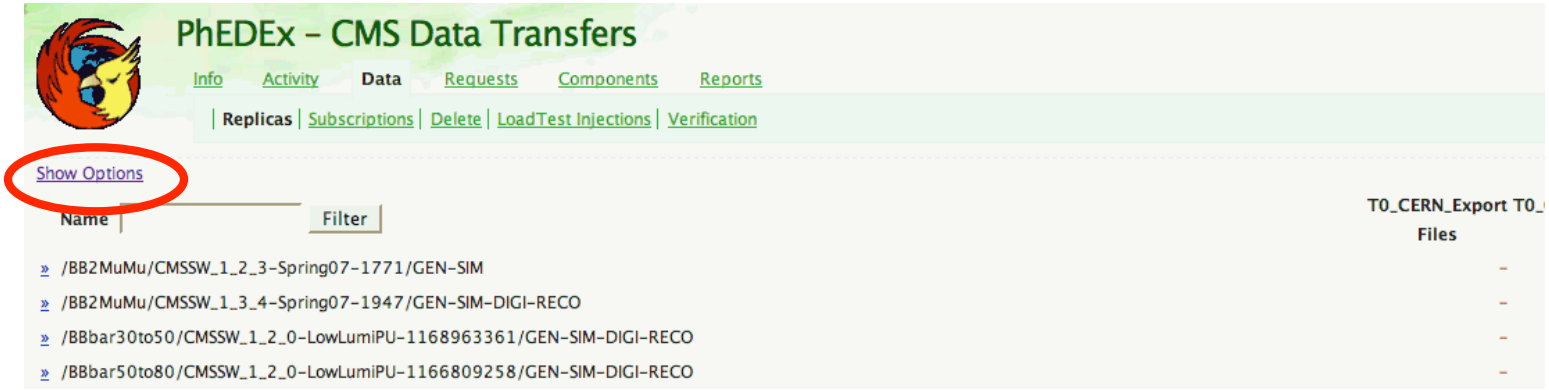

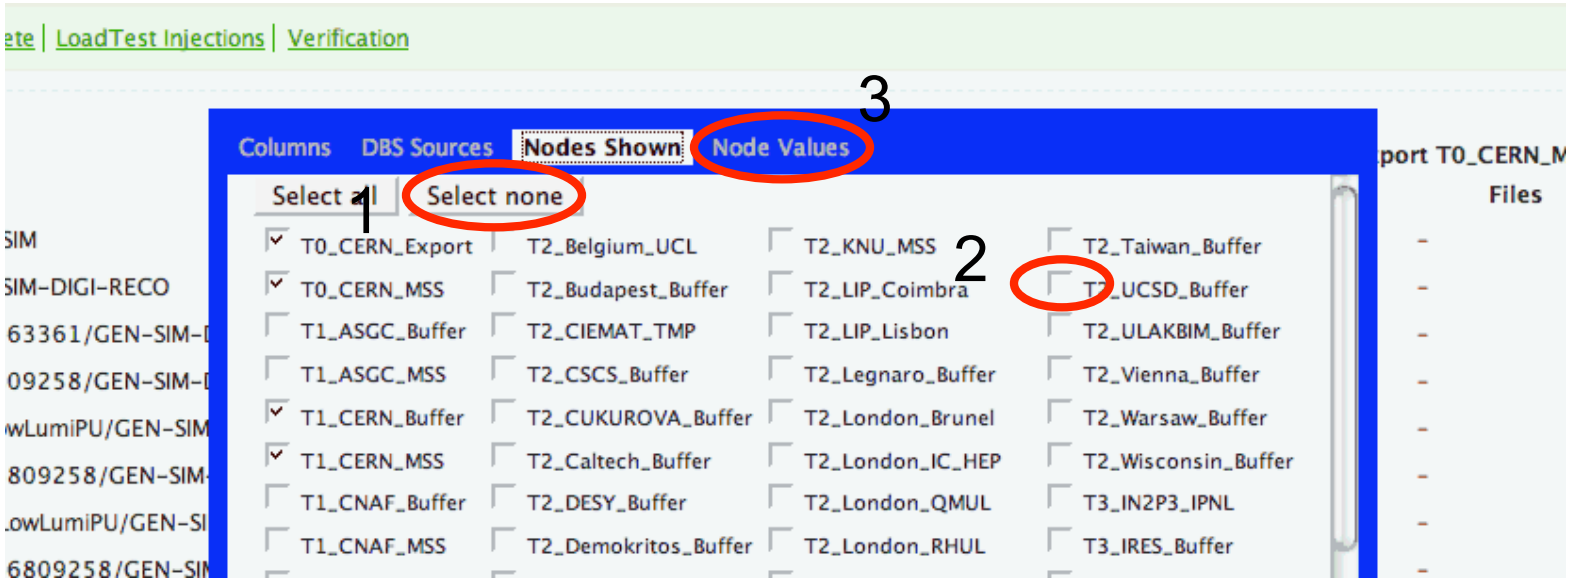

Select whichever sites you want to look at.
This is useful to find out which sites have your dataset,
As well as how much, as well as, ….

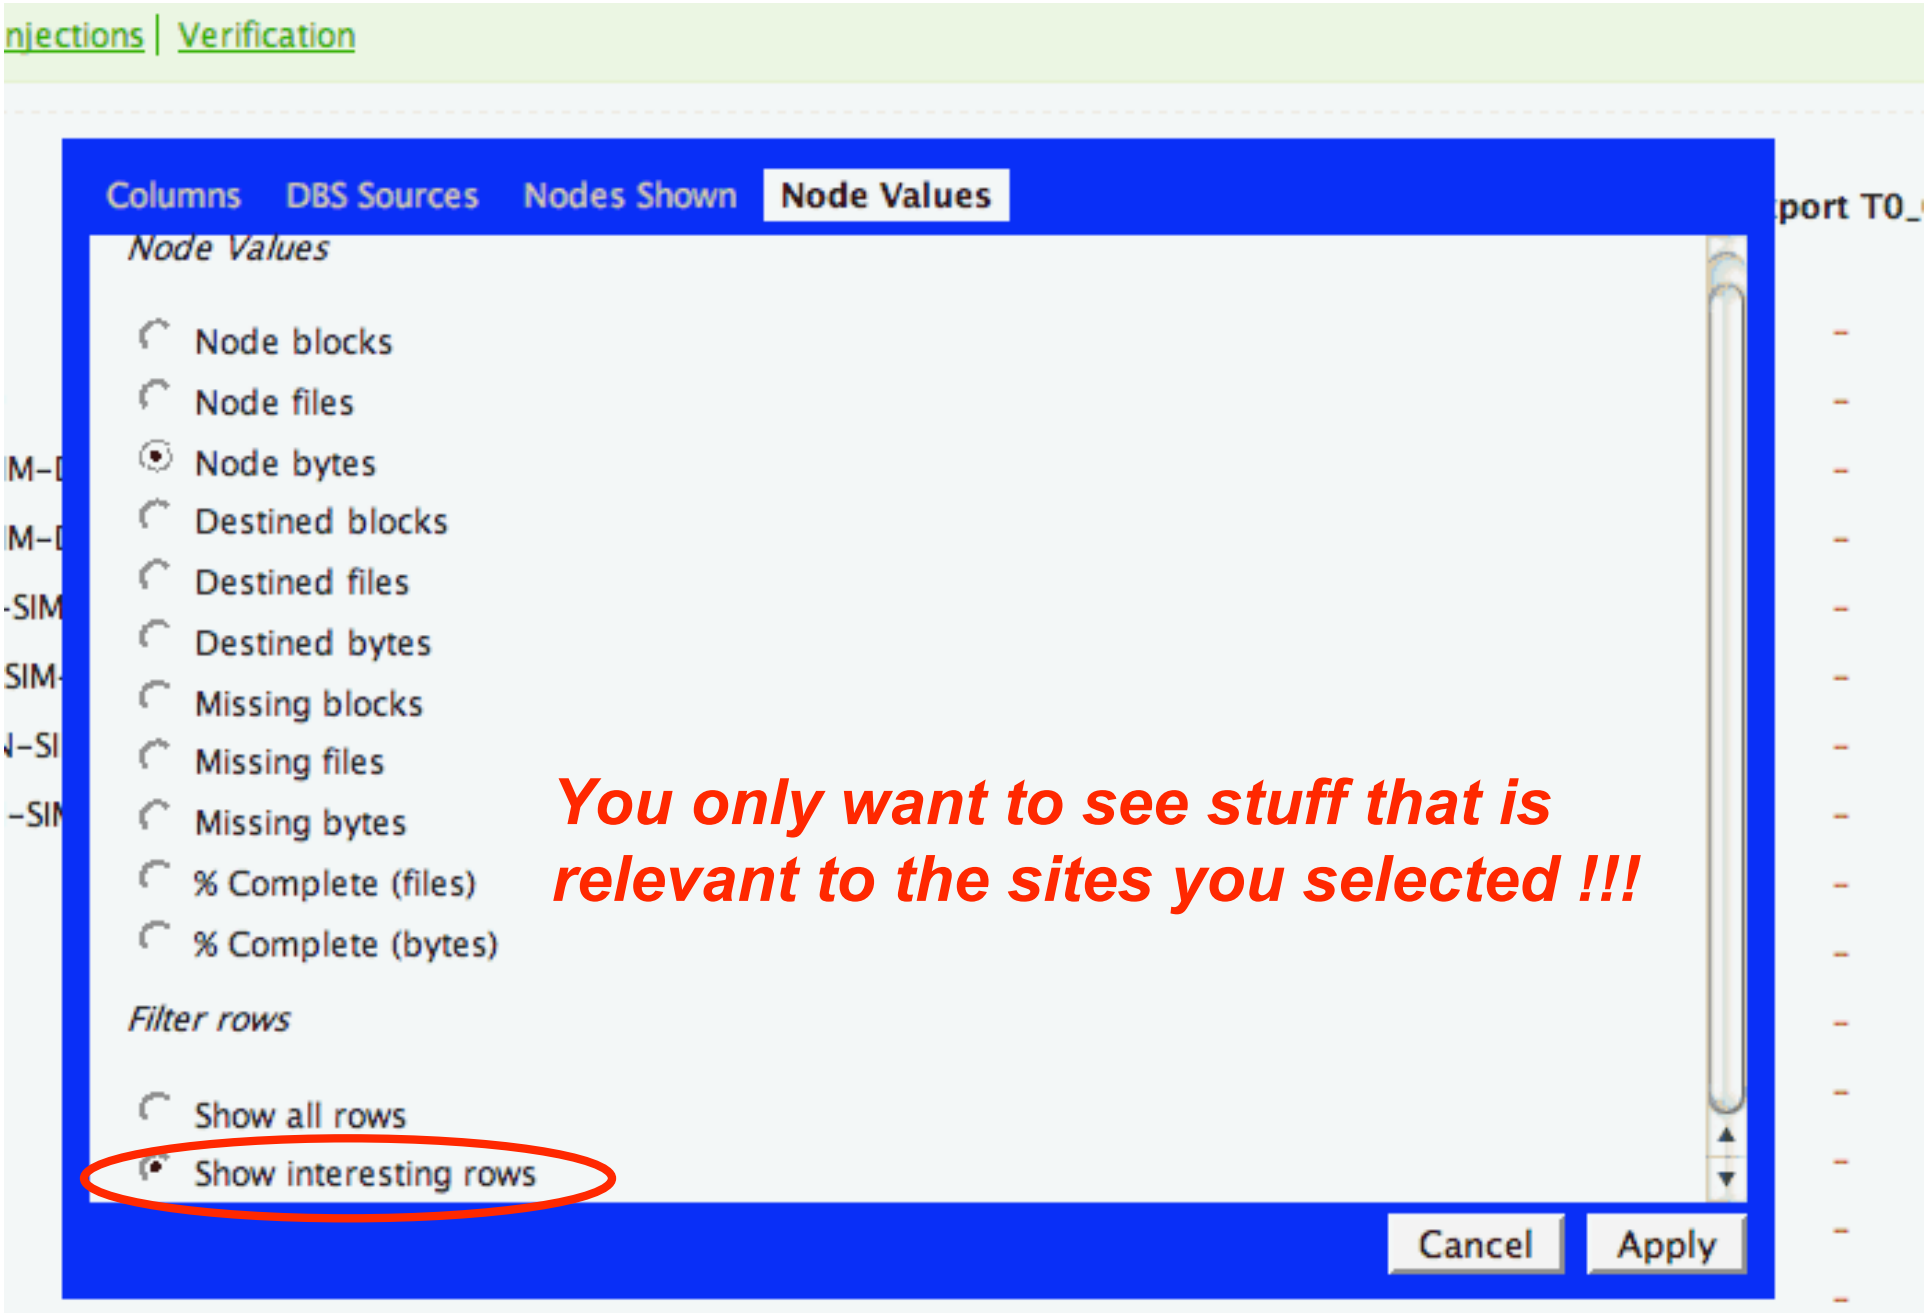
njections | Verification
Columns
DBS Sources
Nodes Shown
Node Values
Node Values
Node blocks
Node files
Node bytes
Destined blocks
Destined files
Destined bytes
Missing blocks
Missing files
Missing bytes
% Complete (files)
% Complete (bytes)
Filter rows
Show all rows
Show interesting rows
You only want to see stuff that is relevant to the sites you selected !!!
Cancel
Apply

You can dig down for the specific files you are missing.

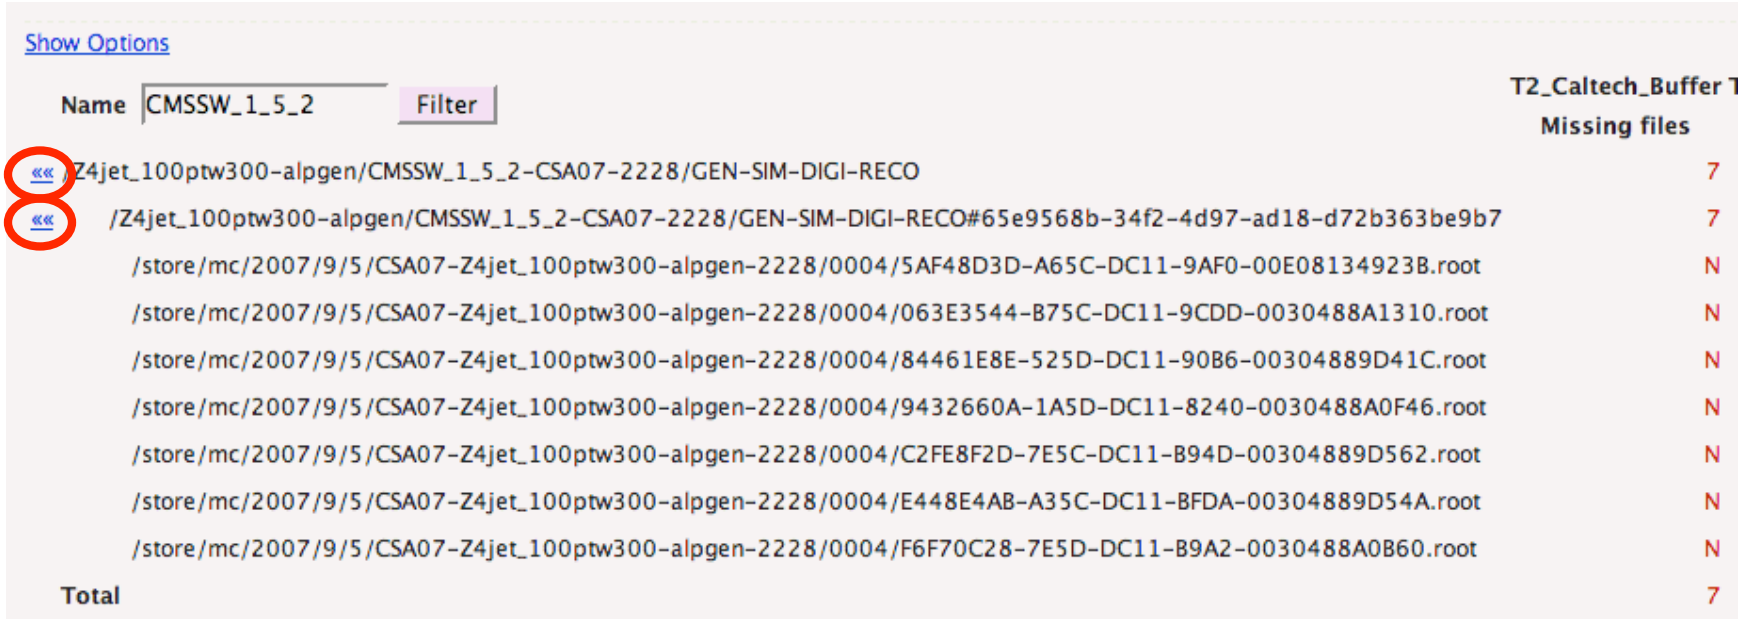

Central Data Ops recently cleaned house, deleting a lot of files that were lost, or simply bookkeeping errors.
If a dataset doesn't complete -> complain to phedex ops !!!
When complaining, provide detailed info, including url's to these phedex monitoring pages.

You can filter datasets.

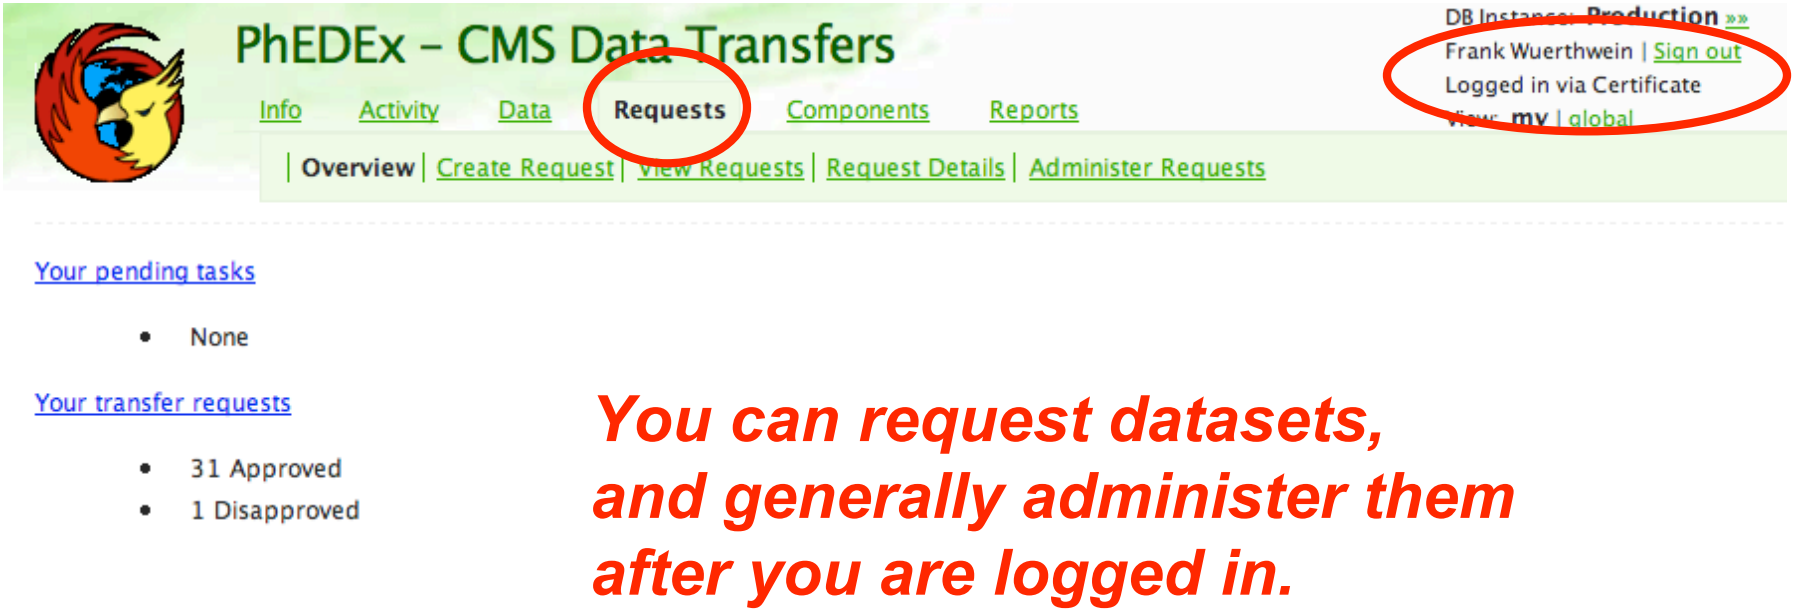

You can request datasets, and generally administer them after you are logged in.

Summary

- It is possible to have a fully functional T2.
- It's not easy yet.
- And it probably will always be a significant commitment of effort.
- Be prepared to invest 2FTE at least at the beginning.