

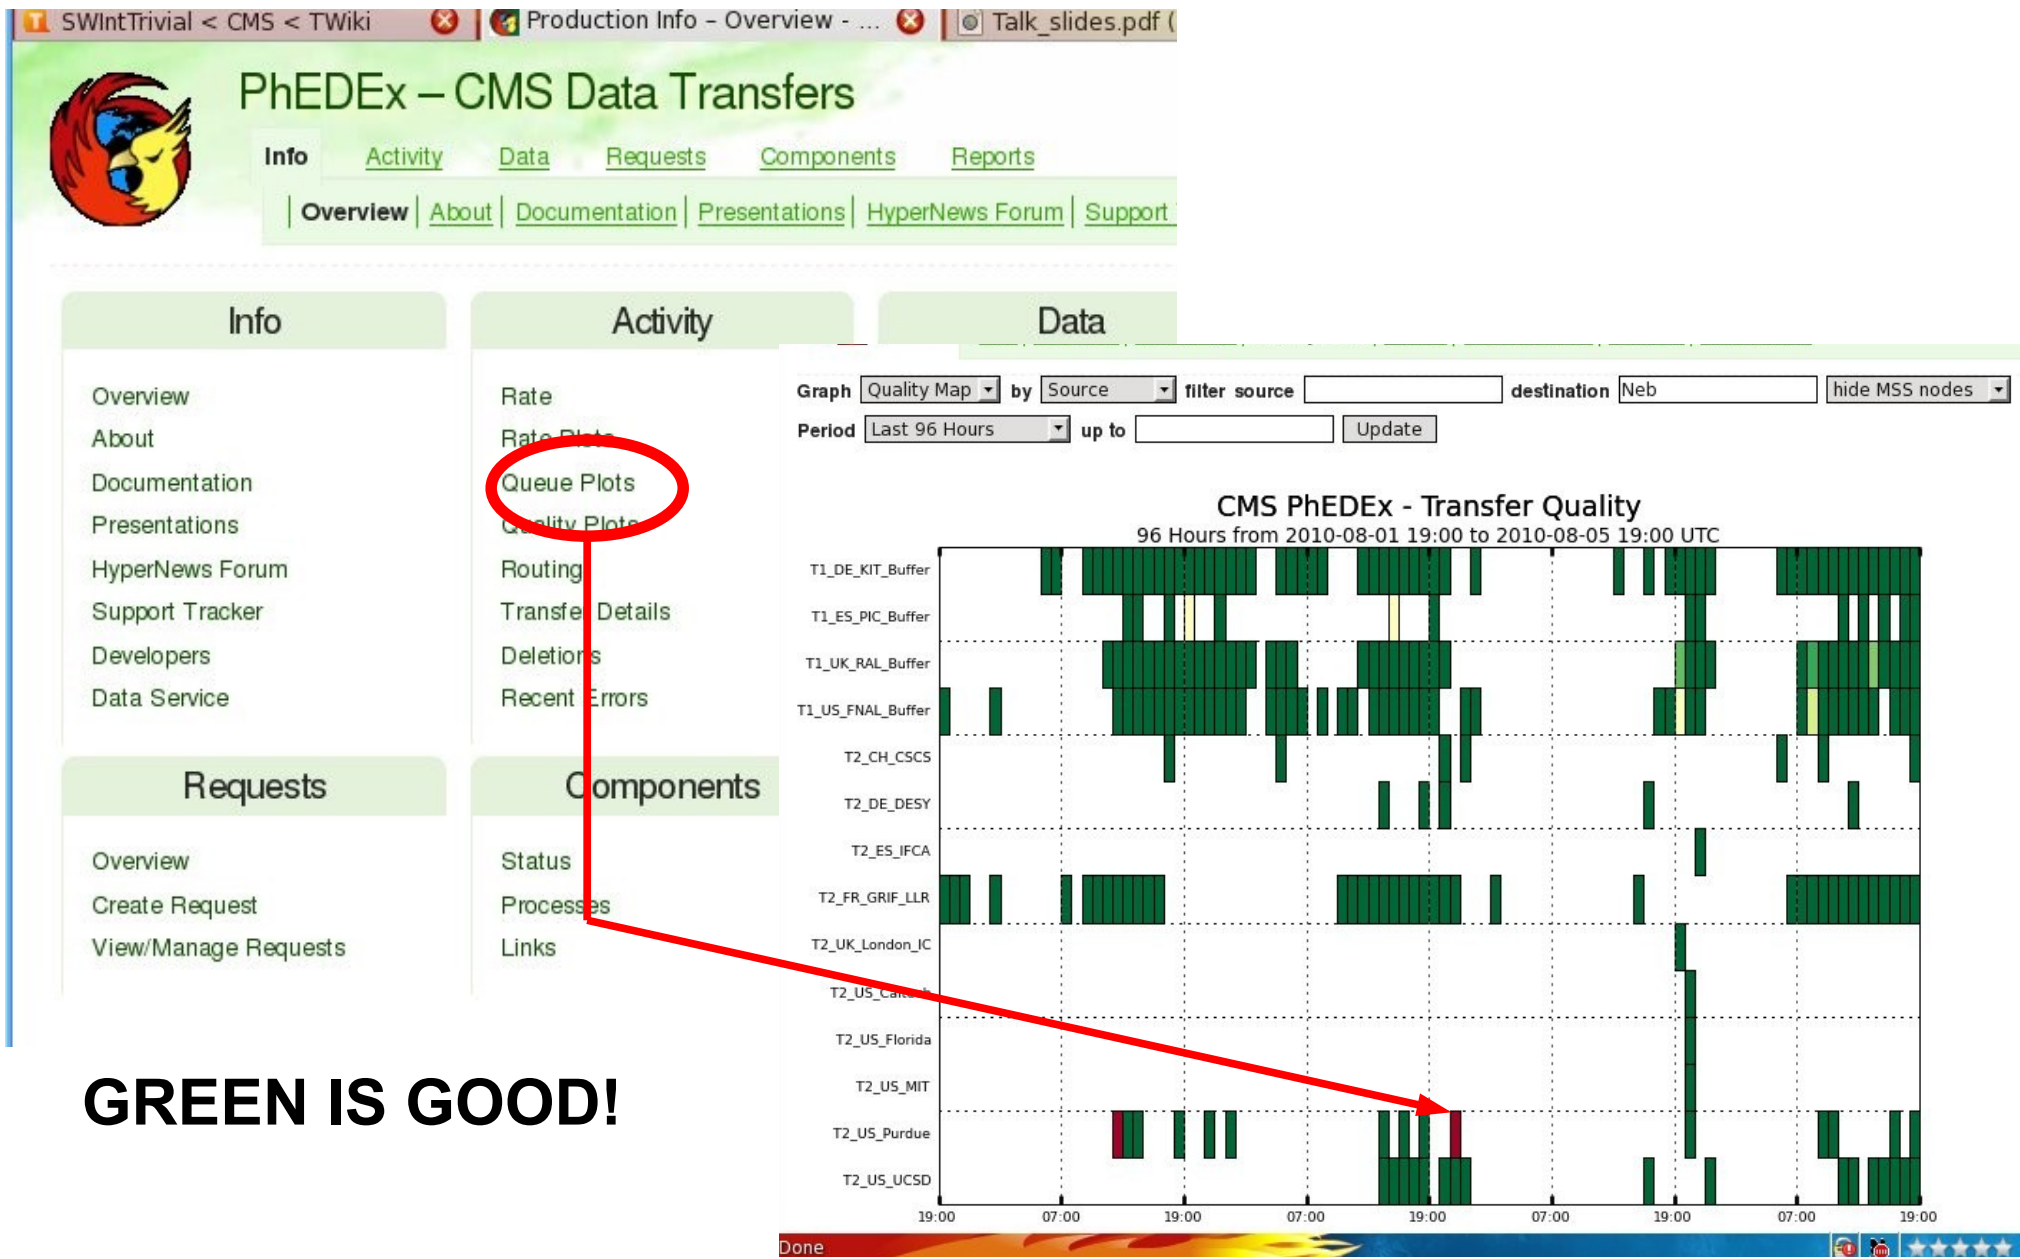
PhEDEx Troubleshooting Tutorial

T-3 Workshop August 2010
Carl Lundstedt



PhEDEx Monitoring

<https://cmsweb.cern.ch/phedex/prod/Info::Main>



PhEDEx Monitoring

SWIntTrivial < CMS < TWiki Production Info - Overview - ... Talk_slides.pdf



PhEDEx – CMS Data Transfers

[Info](#) [Activity](#) [Data](#) [Requests](#) [Components](#) [Reports](#)
[Overview](#) [About](#) [Documentation](#) [Presentations](#) [HyperNews Forum](#) [Support](#)

Info

[Overview](#)
[About](#)
[Documentation](#)
[Presentations](#)
[HyperNews Forum](#)
[Support Tracker](#)
[Developers](#)
[Data Service](#)

Activity

[Rate](#)
[Rate Plots](#)
[Queue Plots](#)
[Quality Plots](#)
[Routing](#)
[Transfer Details](#)
[Deletions](#)
[Recent Errors](#)

Data

[Replicas](#)
[Subscriptions](#)
[LoadTest Injections](#)

Requests

[Overview](#)
[Create Request](#)
[View/Manage Requests](#)

Components

[Status](#)
[Processes](#)
[Links](#)

Error 1

To Node: T3_US_OSU	From Node: T1_US_F
Time Assigned: 2010-08-05 18:21:31 UTC (0m00 since assigned) (-33m16 from now)	Time Transfer Start: 2010-08-0
Time Exported: 2010-08-05 18:28:53 UTC (7m21 since assigned) (-26m38 from now)	Time Transfer Done: 2010-08-0
Time Pumped: 2010-08-05 18:50:41 UTC (29m09 since assigned) (-4m26 from now)	Time Transfer Expires: 2010-08-0

Report Code: 1

Transfer Code: 92

To PFN: srm://cms-0.mps.ohio-state.edu:8443/srm/v2/server?SFN=/data/se/store/mc/Spring10/Wmunu/GEN-SIM-RECO/START3X_V26_S
From PFN: srm://cmssrm.fnal.gov:8443/srm/manager2?SFN=/11/store/mc/Spring10/Wmunu/GEN-SIM-RECO/START3X_V26_S09-v1/0018/
Space Token: (none)

Transfer Log:
(20 lines)
(2449 chars)
2010-08-05 18:52:36 srm-copy (30448): Executing: /sharesoft/osg/ce-1.2/srm-client-lbnl/bin/srm
2010-08-05 18:52:36 srm-copy (30448): srm-copy 2.2.1.3.6 Mon Nov 30 09:10:48 PST 2009
2010-08-05 18:52:36 srm-copy (30448): SRM-Clients and BeStMan Copyright(c) 2007-2009,
2010-08-05 18:52:36 srm-copy (30448): Lawrence Berkeley National Laboratory. All rights reserv
2010-08-05 18:52:36 srm-copy (30448): Support at SRM@LBL.GOV and documents at http://datagrid.
2010-08-05 18:52:36 srm-copy (30448):
2010-08-05 18:52:36 srm-copy (30448): #####
2010-08-05 18:52:36 srm-copy (30448): SRM_HOME is /sharesoft/osg/ce-1.2/srm-client-lbnl
2010-08-05 18:52:36 srm-copy (30448): JAVA_HOME is /sharesoft/osg/ce-1.2/jdk1.5

PhEDEx Monitoring

- Logs...logs...logs....

2010-03-24 16:10:37: FileDownload[32418]: xstats: task=82805824 file=31881284
from=T1_US_FNAL_Buffer to=T2_US_Nebraska priority=5 report-code=1 xfer-code=1
size=2760810314 t-expire=1269457170.66698 t-assign=1269427403.55378 t-
export=1269427876.15998 t-inxfer=1269429294.00402 t-xfer=1269447006.40915 t-
done=1269447037.13745
lfn=/store/PhEDEx_LoadTest07/LoadTest07_Debug_FNAL/Nebraska/2540/LoadTest07_FNAL
_DD_tOOL6cG64RasJtPa_2540 from-pfn=srm://cmssrm.fnal.gov:8443/srm/managerv2?
SFN=/11/store/PhEDEx_LoadTest07/LoadTest07_Prod_FNAL/LoadTest07_FNAL_DD to-
pfn=srm://dcache07.unl.edu:8443/srm/v2/server?
SFN=/mnt/hadoop/user/uscms01/LoadTestDownload/LoadTest07_FNAL_DD_tOOL6cG64Ras
JtPa_2540 detail=(SOURCE error during TRANSFER_PREPARATION phase:
[CONNECTION_ERROR] failed to contact on remote SRM
[httpg://cmssrm.fnal.gov:8443/srm/managerv2]. Givin' up after 3 tries) validate=(Executing:
/grid_home/cmsphedex/SITECONF/T2_US_Nebraska/PhEDEx/FileDownloadVerify -d 1
srm://dcache07.unl.edu:8443/srm/v2/server?
SFN=/mnt/hadoop/user/uscms01/LoadTestDownload/LoadTest07_FNAL_DD_tOOL6cG64Ras
JtPa_2540 2760810314 cksum:1137249321 0 ~~ **[Errno 2] No such file or directory:**
'/mnt/hadoop/user/uscms01/LoadTestDownload/LoadTest07_FNAL_DD_tOOL6cG64RasJtPa_
2540' ~~ **Job exited with status code 1 (256)** after 1.776 seconds ~~) job-
log=/grid_home/cmsphedex/Debug_T2_US_Nebraska/state/download-fts-
t3//archive/job.1269446925.0

PhEDEx Monitoring

- Logs...logs...logs....
- Agent logs are located in a directory name
Debug_T3_SITENAME/logs/<agent name>
- Often, when waiting to if things are going ok I'll do a
tail -f Debug_T2_US_Nebraska/logs/download-fts-all

PhEDEx Monitoring

Common Tasks – How do I

- Determine what datasets are replicated at my site
(DBS client or) phedex pages → Data → Replicas → “/”
- Determine if I'm having any download errors
phedex pages → Activity → Quality Plots (by source) dest yoursite
- See if phedex thinks my agents are running
phedex pages → components
- Create a subscription (request a dataset)
Requests → Create Request
View/Manage Requests
- Check if I have a link to a desired site
Components → Links

PhEDEx Triage

- Try to determine what's failing and document it. You may have to invoke an expert and they'll require a great deal of information. Try to determine reproducible outcomes. “When I do this, this happens, but when I do that that doesn't happen.”
- Check low hanging fruit →
- Is your Voms/Grid proxy still valid? The VOMS proxy extension is hard coded to expire in 8 days at the VO level. Do a **voms-proxy-info -all**
- Can you complete a transfer manually?
- Are transfers from a single site succeeding/failing? Is it a problem with your site or the source site?
- Is there a phedex “link” enabled from the site your trying to pull from?

Common solutions, and sometimes strange solutions:

- Restart the agents.
- Move the 'work' directory out of the way in
Debug_T3_<SITECONF>/state/, restart the agents

The Debug_T3_<SITENAME>/state directory has a number of files which tell you what commands the service is trying.

For instance, in my 'ftsjob.dmp has a line:

```
glite-transfer-submit -s https://cmsfts1.fnal.gov:8443/glite-data-transfer-fts/services/FileTransfer  
-f /grid_home/cmsphedex/Debug_T2_US_Nebraska/state/  
download-fts-all/work/job.1281451463.912/copyjob
```

Where copyjob is a single line:

```
srm://dcache-se-cms.desy.de:8443/srm/managerv2?SFN=/pnfs/desy.de/cms/tier2/store  
/phedex_monarctest/monarctest_DESY/MonarcTest_DESY_8A  
srm://dcache07.unl.edu:8443/srm/v2/server?SFN=/mnt/hadoop/user/uscms01  
/pnfs/unl.edu/data4/cms/store/PhEDEx  
_LoadTest07/LoadTest07_Debug_DE_DESY/US_Nebraska/44/LoadTest07_DESY_8A_  
CM5GzAaMe M1yyp69_44
```

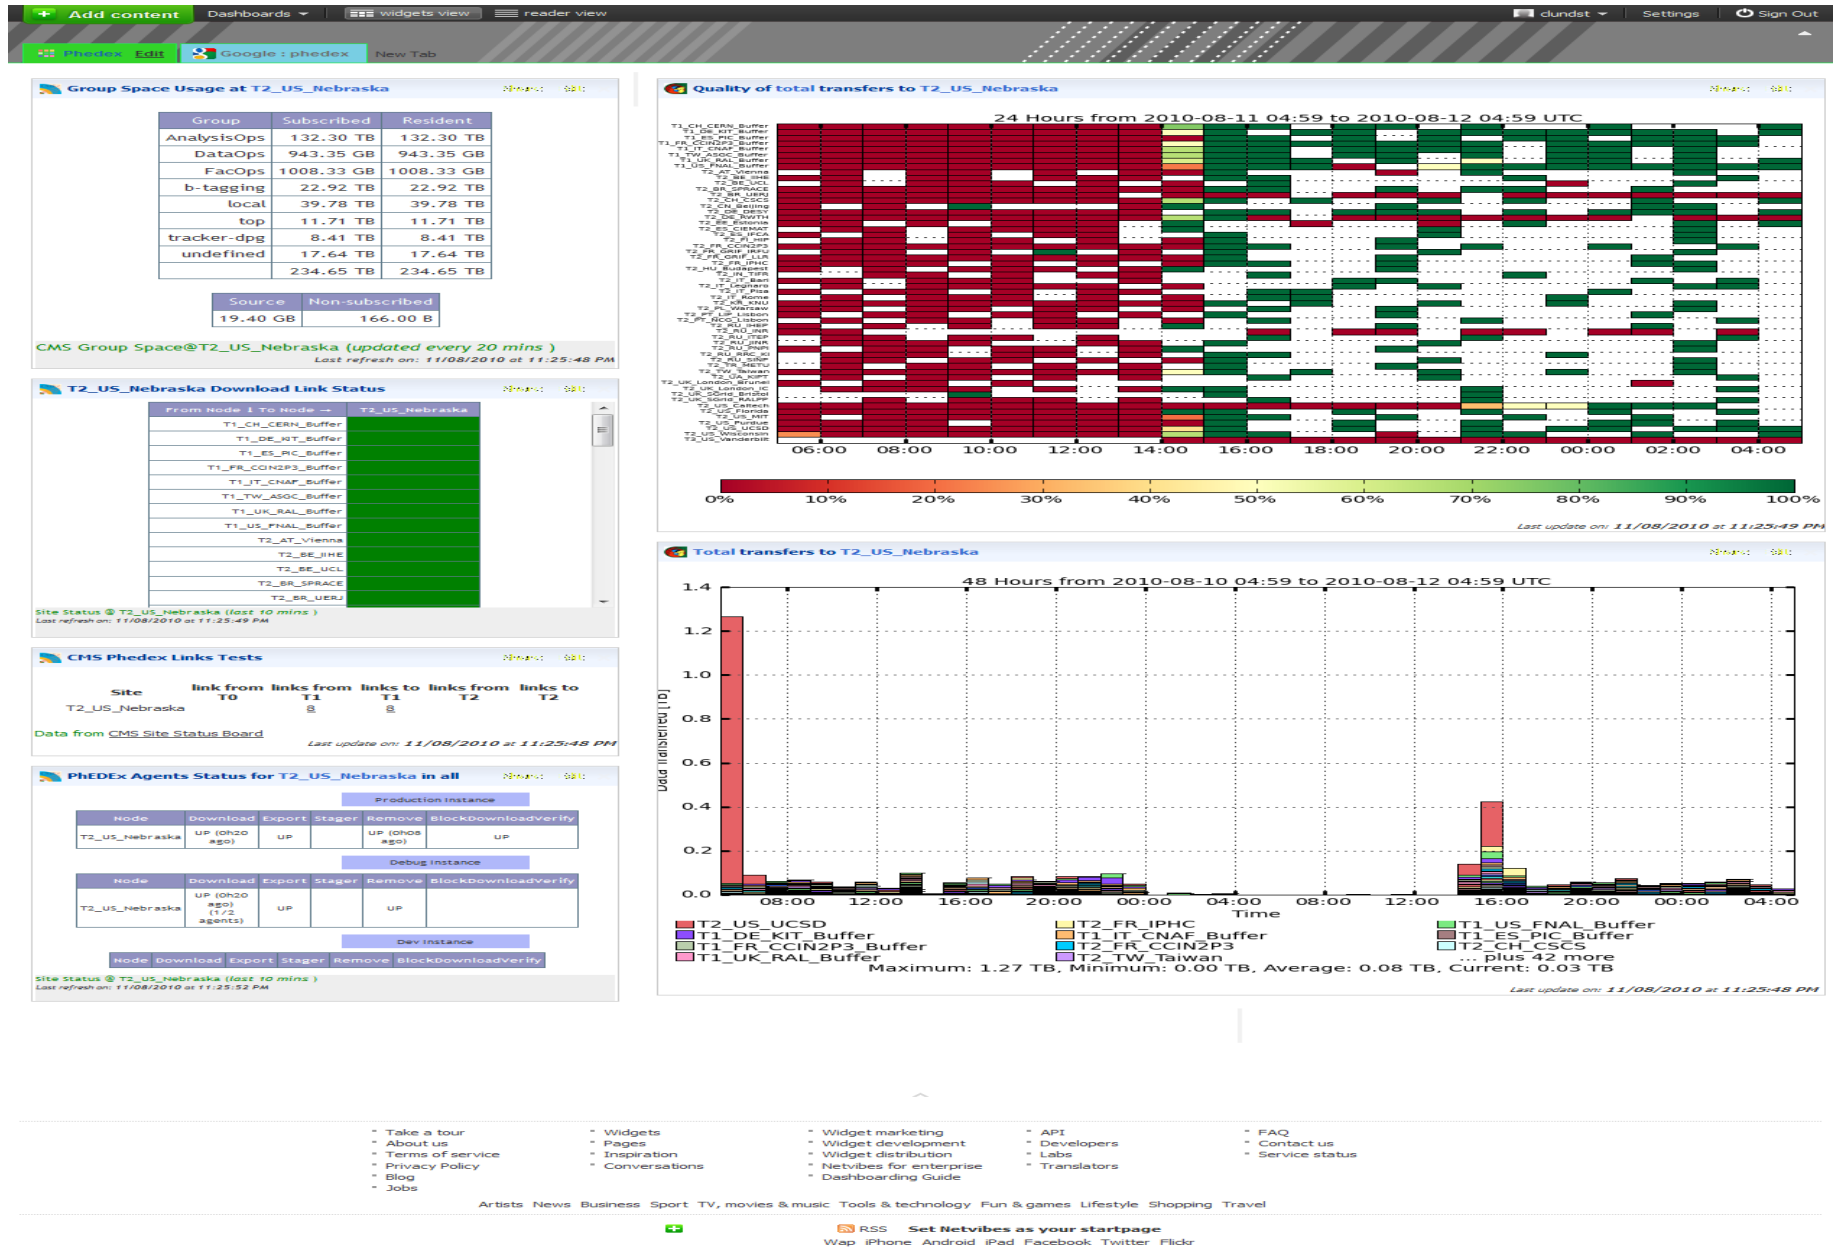
**I could issue this command by hand as a way of starting to debug the transfer if it fails.
Transfers states are kept in an archive directory.**

Common solutions, and sometimes strange solutions:

- Restart the agents.
- Move the 'work' directory out of the way in
Debug_T3_<SITECONF>/state/, restart the agents

Create a dashboard – Netvibes

<http://eco.netvibes.com/widgets/tag/phedex>



cmsSync --- Am I missing files, have I lost files that Phedex thinks I have?

<http://osg-test2.unl.edu/documentation/cmssync>

To install cmsSync and use it there a couple of requirements.

- install the Nebraska Yum repository
- The node you run cmsSync on has to 'mount' your SE in some fashion. Hadoop+Fuse work, dCache + pnfs work, NFS certainly works. The node you plan to run cmsSync from does not need to be 'special' in any other way. It needs not have phedex installed etc.

Installing the Nebraska Repo:

For 32-bit linux installs: rpm -i <http://t2.unl.edu/store/rpms/SL5/i686/Nebraska-repo-0.1-1.noarch.rpm>

For 64-bit linux installs: rpm -i http://t2.unl.edu/store/rpms/SL5/x86_64/Nebraska-repo-0.1-1.noarch.rpm

Installing cmsSync

[yum install cmsSync](#)

(I had to download the rpm and install it by hand, see next slide)

```
wget http://t2.unl.edu/store/rpms/SL5/x86\_64/cmsSync-0.3.1-0.noarch.rpm
```

```
rpm -i cmsSync-0.3.1-0.noarch.rpm
```

Running the client

CmsSync (Cont.)

The client does the following things:

- * Gets a list of all blocks registered for your SE.
- * Retrieve your site's TFC.
- * Builds a list of file names from these blocks (long process).
- * Spiders a given base directory (long process).
- * Compares the contents of dCache and the list of files which should be at your site.

This requires PNFS to be mounted at your site and for you to know the SE name you use with PhEDEx. The command line usage goes like this:

```
cmsSync --se=srm.unl.edu /mnt/hadoop/user/uscms01/pnfs/unl.edu/data4/cms/store
```

I run this in an empty directory so that the output files don't get lost.

Some info that you'll get back (stored in files in the directory you run in):

All files phedex thinks are at your site.

Files that are 'missing' from the filesystem

Files in the files system that phedex doesn't know about (User files, cruft?)

If you find you have missing files you can submit a savannah ticket, or if you're having good luck with transfers just submit a deletion->"don't remove subscription" and redownload the whole dataset. (Probably not what you want to do, but easy)