

US Tier3 Storage and DDM

Dr. Hironori Ito

Brookhaven National Laboratory


How to set up the storage

- Instruction to set up the storage at T3
 - <https://atlaswww.hep.anl.gov/twiki/bin/view/Tier3Setup/SetupSE>
 - For the throughput test,
 - <https://atlaswww.hep.anl.gov/twiki/bin/view/Tier3Setup/ThroughputCleanup>
 - Help?:
 - RT-RACF-USAtlasTier3@bnl.gov


How ATLAS DDM works




Users request datasets at T3




Datasets request page




DQ2 Central Services (CERN)




BNL US T3 DQ2 Site Service



BNL T3 LFC



Source Site's LFC



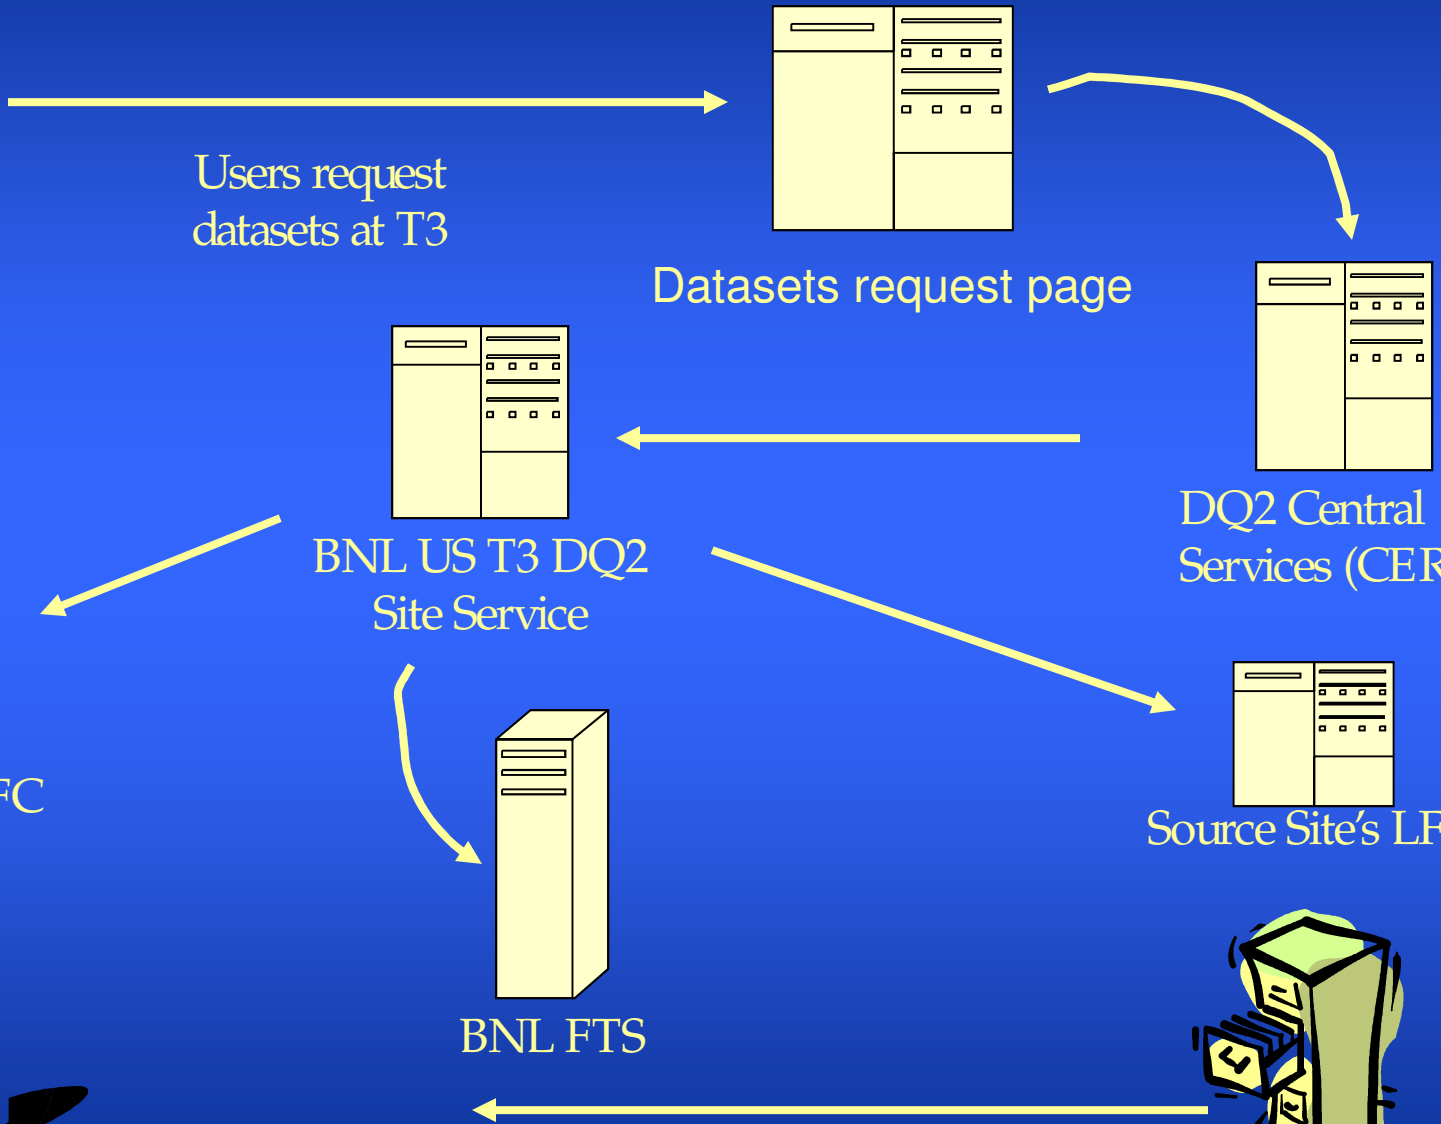
BNL FTS



T3 Storage



Source Site Storage



Current T3 Status

- Production sites
 - WISC, INNILNOISHEP and UTD
 - Have all necessary production areas
 - PRODDISK, MCDISK, etc...
- Non-production sites
 - ANL and DUKE
 - Has LOCALGROUPDISK area only

US T3's DDM Services

- US T3 DQ2 site service is running at BNL.
 - It will serve an entire US T3s.
 - host: `t3dq2v01.usatlas.bnl.gov`
- US T3 LFC service is running at BNL.
 - It will serve an entire US T3s.
 - Host: `t3lfcv01.usatlas.bnl.gov`

ANL_ and DUKE_LOCALGROUPDISK have been completely integrated into the ATLAS DDM system. The functionality of the DDM service has been tested by subscribing datasets from BNL.

Performance test

- Scheduled throughput-test program
 - Use FTS to simultaneously transfer multiple files of 3600MB in size
 - It is not designed to measure the absolute maximum, but rather to measure what usually available in the single-shot transfer.
 - T1->T2s 20 files
 - T2->T2s 10 files
 - T1/T2s->T3 7 files
 - ANL and DUKE have been already integrated.
 - Monitor
 - All tests are recorded. And, results can be seen at <http://www.usatlas.bnl.gov/dq2/throughput>
 - Useful to get general idea of transfer rates.
 - Useful to identify the problem.
- Maximum throughput for T1 and T2s have been also tested.
 - T1. at least: 1.5GB/s read and 1GB/s to write via external SRM.
 - T2s. Write at least 400MB/s (some go 1.2GB/s).

Throughput test for DUKE

Number of successful Transfers from BNL-OSG2_MCDISK

7 ●●●●●●●●

Transfer time in seconds from BNL-OSG2_MCDISK

●●●●●●●● ●●●●●●●● ●●●●●●●● ●●●●●●●●

Transfer time in seconds from NET2_MCDISK

Number of successful Transfers from NET2_MCDISK

7 ●●●●●●●●
6
5
4
3
2
1
0

10/13 01:49

Transfer time in seconds from NET2_MCDISK

■ Aver

440 ●●●●●●●●
330 ●●●●●●●●
220 ●●●●●●●●
110 ●●●●●●●●
0 ●●●●●●●●

10/13 01:49 10/1

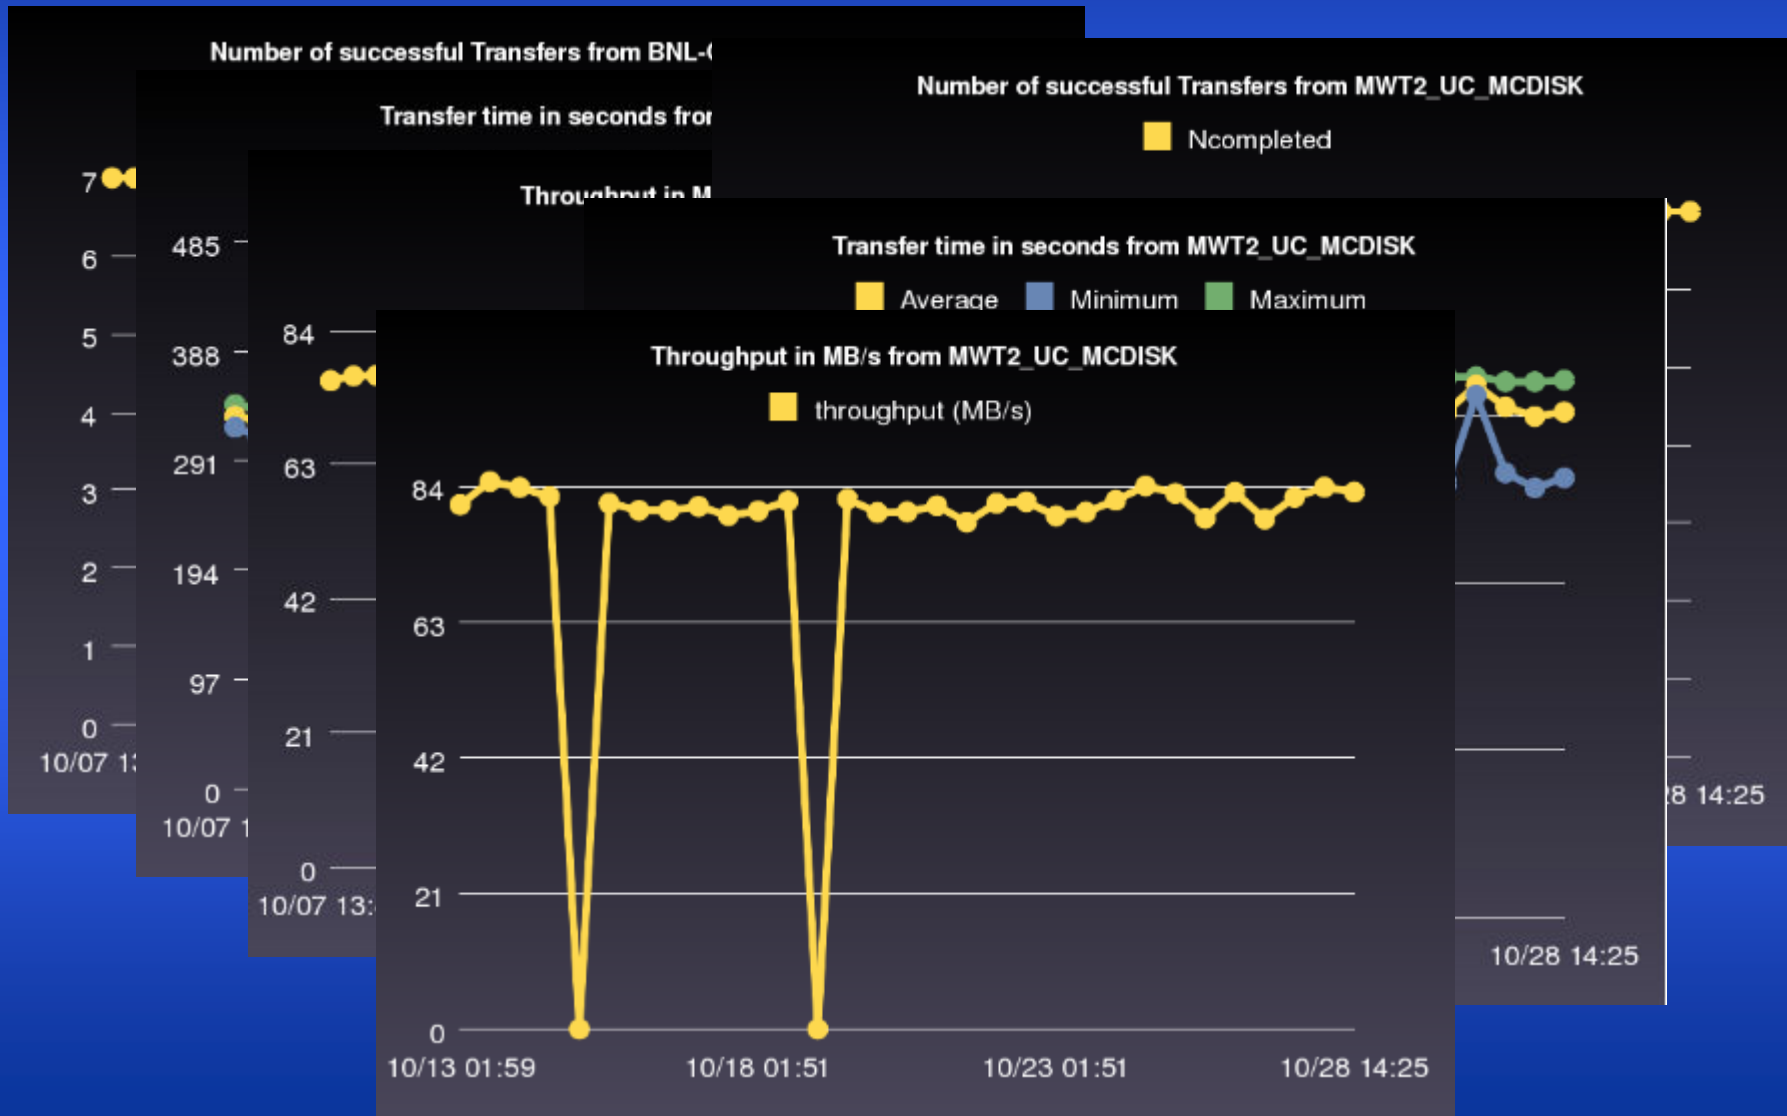
Throughput in MB/s from NET2_MCDISK

■ throughput (MB/s)

60 ●●●●●●●●
45 ●●●●●●●●
30 ●●●●●●●●
15 ●●●●●●●●
0 ●●●●●●●●

10/13 01:49 10/18 01:48 10/23 01:48 10/28 14:22

Throughput test for ANL



Expected Performance at T3

- Needs reasonable performance.
 - What is reasonable? 10MB/s? 50MB/s?
 - Might/probably need to categorize them?
 - FTS channel is shared.
 - Slow sites could interfere with faster sites.
 - Can make more channels if necessary.
 - By geography: Associate to a T2?
 - By performance
 - Should have some minimum requirement?

How to monitor DQ2 transfers

- The service will have the log monitor.
 - No USE of DDM dashboard to avoid messages from ADC operation.
 - Shows all transfers: Fail or Successful

BNL DQ2 logs

View and search BNL DQ2 log.

Search words(one string/word per each line. Can use * to match.)

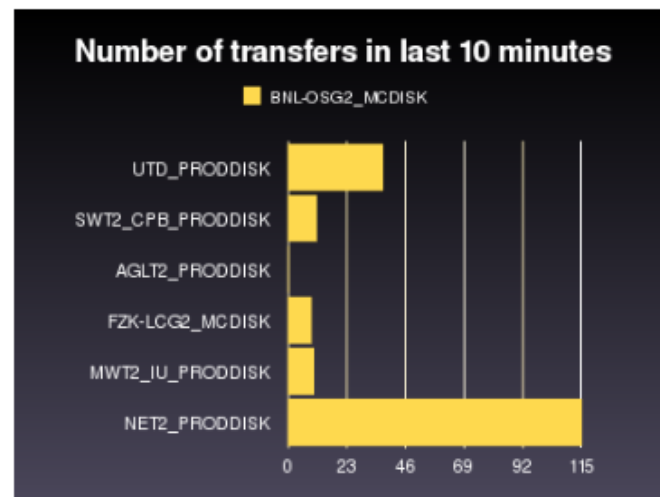
Time Range

Start Time(eg. 2009/10/10 19:54) :

End time (eg. 2009/10/11 20:00) :

Max entries per page: 200

Click to search



Showing the last 100 lines of the log. Updated every 10s. Automatically stops after 10 times. You must click "Start tail" again to restart tailing if you want to see it.

[Start tail](#) [Stop tail](#)

```
srm://head01.aglt2.org/pnfs/aglt2.org/atlasproddisk/mc09_10TeV/log/e469_s595/mc09_10TeV.105802.JF17_pythia_jet_filter.simul.log.e469_s595_tid09528704
/log.095287._041243.job.log.tgz.1 srm://dcsrm.usatlas.bnl.gov:8443/srm/managerv2?SFN=/pnfs/usatlas.bnl.gov/MCDISK/mc09_10TeV
/log/mc09_10TeV.105802.JF17_pythia_jet_filter.simul.log.e469_s595_tid09528704_sub03740835/log.095287._041243.job.log.tgz.1
srm://head01.aglt2.org/pnfs/aglt2.org/atlasproddisk/mc09_10TeV/log/e469_s595/mc09_10TeV.105802.JF17_pythia_jet_filter.simul.log.e469_s595_tid09528704
```

Managing datasets

- Support from T1/T2s will be minimum.
 - Shared knowledge: Tier3 RT/wiki/hypernews
 - Ultimately T3 itself
- Getting datasets:
 - Currently, use standard dataset request page.
 - Should/will have an own page.
- Finding what you have in your site
 - T3 LFC and T3 storage.
 - Can you the BNL DDM dataset page
 - www.usatlas.bnl.gov/dq2/monitor
- Deleting datasets
 - Need automated program.
 - Should/will have a request page.
 - Don't use the production method since it won't work.
 - Don't just delete files from the storage. All DDM catalog must be synchronized.

Conclusion

- T3 DDM services, LFC and DQ2, are in the operation.
- ANL and DUKE are in the operation.

New T3s just need to bring SRM storage online to join the party!!!

If you build it, files will come!