US/BNL DDM services

Hironori Ito
Brookhaven National Laboratory

Current Status

- DQ2s
  - BNL, BNL_T3, AGLT2, MWT2_IU, MWT2_UC, NET2, SWT2, WT2
- LFCs
  - BNL, BNL-T3, AGLT2, MWT2_IU, MWT2_UC, NET2, SWT2_CPB, SWT2_UTA, SWTW_OU, WT2, WISC, UTD, ILLINOISHEP,OUHEP

Future

- DQ2
  - BNL
    - Serves BNL and all T2s.
  - BNL_T3
    - Serves all T3s
- LFC
  - BNL and T2s will keep own LFCs
  - All T3s will use BNL_T3 LFC
  - Must be version 1.7.2 or higher to do bulk deletion. Deadline is this week (2nd week of November) according to the last operation meeting.

Schedule

- US T3s have already migrated to BNL_T3 DQ2 SS
  - 1st week of November
- BNL FTS has been upgraded to V 2.2
  - 1st week of November
  - Checksum support via srm
  - Testing will start within a week (3rd week of November) via the throughput testing program.
- BNL T1/T2 DQ2 will be upgraded within a week.
- Migration of T3 LFC will start within the week.
- T2 DQ2s will migrate to BNL
  - The migration of dCache T2s (AGLT2 and MWT2) will start by the end of the month to the early next month (December) since the checksum feature in its SRM has been already tested at T1 sites.
  - Migration of Bestman SRM sites will be conducted after examining results from the throughput test.

Monitoring DDM remotely

- BNL DQ2 log monitor
  - Access it at http://www.usatlas.bnl.gov/dq2log/dq2log
  - New entries in DQ2 log is indexed every 2 minutes
  - The record will be kept for about last 2 months.
  - Show plots of successful and failed transfers in last 10 minutes
  - The extensible and fast search capabilities
    - Search is fast for ~30 Millions lines of logs.
      - Much faster than “grep”.
    - Mix of “AND” and “OR”
    - Use “*” for all matches
  - PANDA has already integrated the link to the log monitor to see why jobs are still in “transferring” state

Integration with PANDA

2009-11-06 19:53:37,108 - INFO - File log.097785._006220.job.log.tgz.1 [b3a636f2-363e-410a-98f6-ea18a6faa481] registered for site BNL-OSG2_MCDISK.

2009-11-06 19:53:37,108 - INFO - File log.097785._006013.job.log.tgz.1 [d88507f5-48f2-4791-b759-743c50655df8] registered for site BNL-OSG2_MCDISK.

2009-11-06 19:53:42,221 - INFO - FAILED GUID de6965d3-88b9-4f22-8497-0d9af6ab5657 FOR MWT2_UC_PRODDISK->BNL-OSG2_MCDISK [FTS State [Failed] FTS Retries [1] Reason [SOURCE error during TRANSFER_PREPARATION phase: [INVALID_PATH] source file doesn't exist] Source Host [uct2-dc1.uchicago.edu]]

2009-11-06 19:53:42,221 - INFO - FAILED GUID 62f98290-ac8f-4a62-b3b1-d8bac498f477 FOR MWT2_UC_PRODDISK->BNL-OSG2_MCDISK [FTS State [Failed] FTS Retries [1] Reason [SOURCE error during TRANSFER_PREPARATION phase: [LOCALITY] Source file [srm://uct2-dc1.uchicago.edu /pnfs/uchicago.edu/atlasproddisk/mc09_7TeV/log/e468_s624/mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785/log.097785._006218.job.log.tgz.1]: locality is UNAVAILABLE] Source Host [uct2-dc1.uchicago.edu]]

2009-11-06 19:54:59,805 - INFO - Subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for site BNL-OSG2_MCDISK will query sources ['MWT2_UC_PRODDISK'] (took off sources ['BNL-OSG2_MCDISK'])

2009-11-06 19:55:00,178 - WARNING - Skipped 1 files from MWT2_UC_PRODDISK due to bad source files for subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for BNL-OSG2_MCDISK.

2009-11-06 19:55:00,178 - WARNING - Skipped 1 files from MWT2_UC_PRODDISK due to failed transfers for subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for BNL-OSG2_MCDISK.

2009-11-06 19:56:58,742 - INFO - Subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for site BNL-OSG2_MCDISK will query sources ['MWT2_UC_PRODDISK'] (took off sources ['BNL-OSG2_MCDISK'])

2009-11-06 19:56:59,115 - WARNING - Skipped 1 files from MWT2_UC_PRODDISK due to bad source files for subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for BNL-OSG2_MCDISK.

2009-11-06 19:56:59,955 - INFO - Broken subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for site BNL-OSG2_MCDISK as some files do not have at least one valid replica.

2009-11-06 05:45:34,384 - INFO - Subscription mc09_7TeV.105010.J1_pythia_jetjet.simul.log.e468_s624_tid097785_sub03876186 for site BNL-OSG2_MCDISK will query sources ['MWT2_UC_PRODDISK'] (took off sources ['BNL-OSG2_MCDISK'])

DDM Monitor future plan

- Add an option to search by DSN in the page as it was already done for PANDA
- Link to FTS transfer information
  - There is already link to FTS in the throughput page to find the cause of failures in the throughput test.
  - Will be able to view FTS log of a particular transfer.
- Any other things you wish you had to find the cause of the problems?
- Some stability problem for index.
- Do we want more notification for the problems?
  - As it can do the complex search, you can do very complex algorithm to notify any conditions.
  - Web service can be added.
    - For example, it would be possible to do something like.
      - getLastTenFailedTransferBySite?site=ABC

Getting information from FTS

ATLAS Throughput Test for MWT2_IU_MCDISK

Narrow the range of time by filling the start and end date

Start Date

End Date Show plots

Failed FTS Job ID	Source Site	Time	# of successful files / # of total files
fac96559-c99b-11de-acb4-aebc95b8c4f4	BNL-OSG2_MCDISK	2009-11-04 23:44:23 UTC	0/20

[redacted]

srm://dcsrm.usatlas.bnl.gov:8443/srm/managerv2?SFN=/pnfs/usatlas.bnl.gov/MCDISK/test/hiro/user.HironoriIto.MICH.Load_001.010	srm://iut2-dc1.iu.edu:8443/srm/managerv2?SFN=/pnfs/iu.edu/atlasmcdisk/loadtest/2009/11/04/loadtest_MWT2_IU_MCDISK_1257378262.87683	
srm://dcsrm.usatlas.bnl.gov:8443/srm/managerv2?SFN=/pnfs/usatlas.bnl.gov/MCDISK/test/hiro/user.HironoriIto.MICH.Load_001.011	srm://iut2-dc1.iu.edu:8443/srm/managerv2?SFN=/pnfs/iu.edu/atlasmcdisk/loadtest/2009/11/04/loadtest_MWT2_IU_MCDISK_1257378262.87860	No transfer markers received for more than 180 seconds

FTS transfer detailed for FTSID=fac96559-c99b-11de-acb4-aebc95b8c4f4

State	Source TURL	Destination TURL	Error Message	Written/Total bytes
Failed	gsiftp://dcdoor14.usatlas.bnl.gov.:2811//pnfs/usatlas.bnl.gov/MCDISK/test/hiro/user.HironoriIto.MICH.Load_001.011	gsiftp://iut2-dc3.iu.edu:2811//pnfs/iu.edu/atlasmcdisk/loadtest/2009/11/04/loadtest_MWT2_IU_MCDISK_1257378262.87860	No transfer markers received for more than 180 seconds	3600000000/3600000000
Completed	gsiftp://dcdoor05.usatlas.bnl.gov.:2811//pnfs/usatlas.bnl.gov/MCDISK/test/hiro/user.HironoriIto.MICH.Load_001.011	gsiftp://iut2-dc1.iu.edu:2811//pnfs/iu.edu/atlasmcdisk/loadtest/2009/11/04/loadtest_MWT2_IU_MCDISK_1257378262.87860		3600000000/3600000000

close

User Dataset

- Usage of LOCALUSERDISK is monitored by Dataset monitor page
- It is categorized by DN of the dataset owner fr.om the central catalog
- Space usage is also monitored.
- Once a month, each owner will get the notification email about the deletion of datasets and instruction about them if users want to keep them.
  - All files will be deleted without exceptions! Neither I or any other person have time managing personal requests from over 1000 people with their personal choice of non-deleting dataset. If users want to keep files, they must move files to atlasgroupdisk or atlaslocalgroupdisk area.

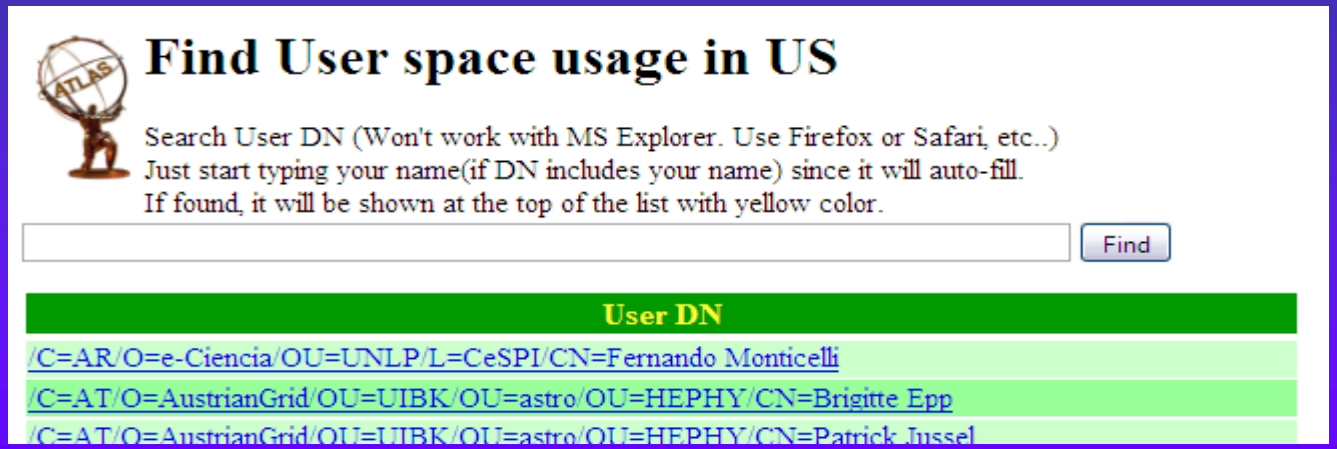

User dataset

- Datasets are color-coded.
- For those who wants the plain text list, there is a link to obtain it.

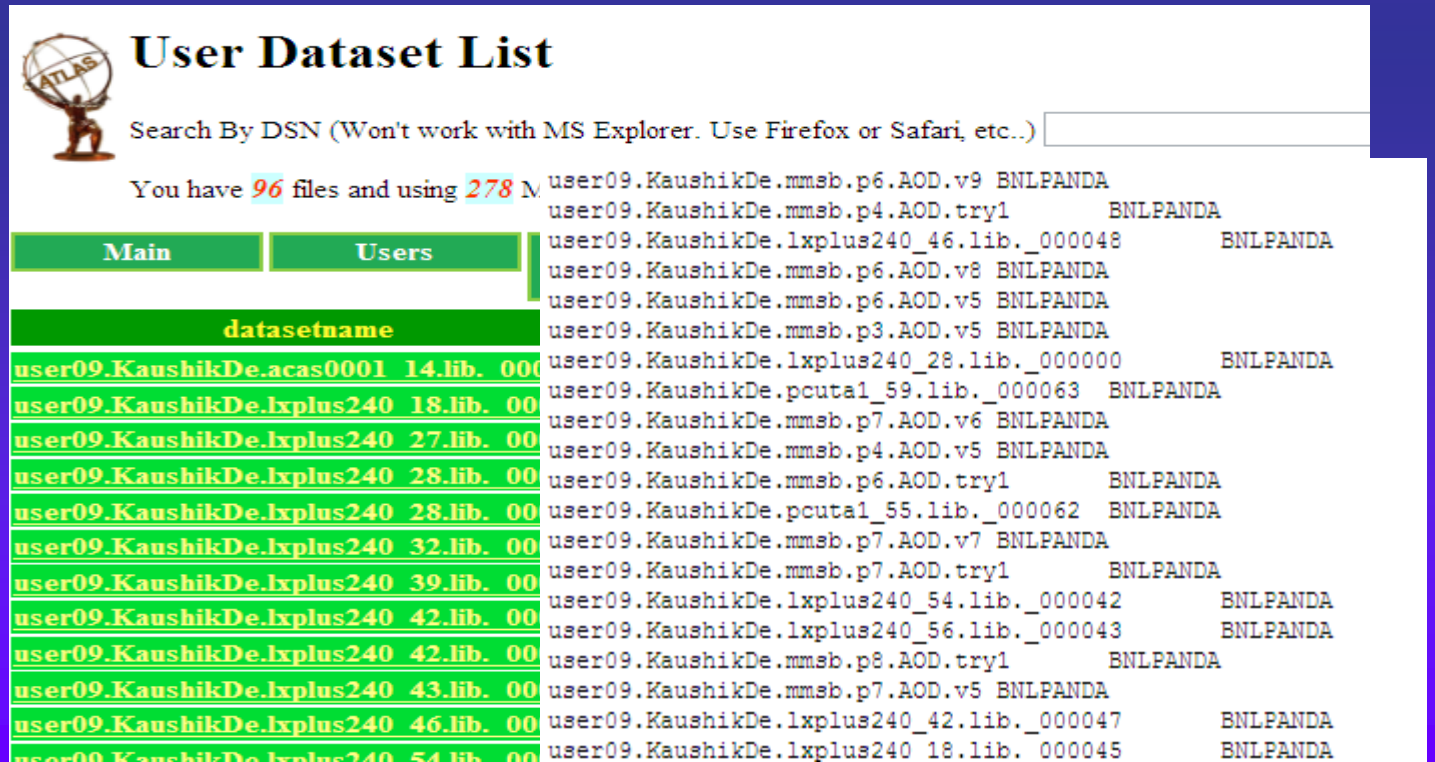

Managing space

Fast and detail information about the space

Summary of datasets at BNL-OSG2_MCDISK

Search By DSN (Won't work with MS Explorer. Use Firefox or Safari, etc..)

Number of Datasets = 10121

- **Number of files = 3520183**

Total Size of files = 241469.34 GB

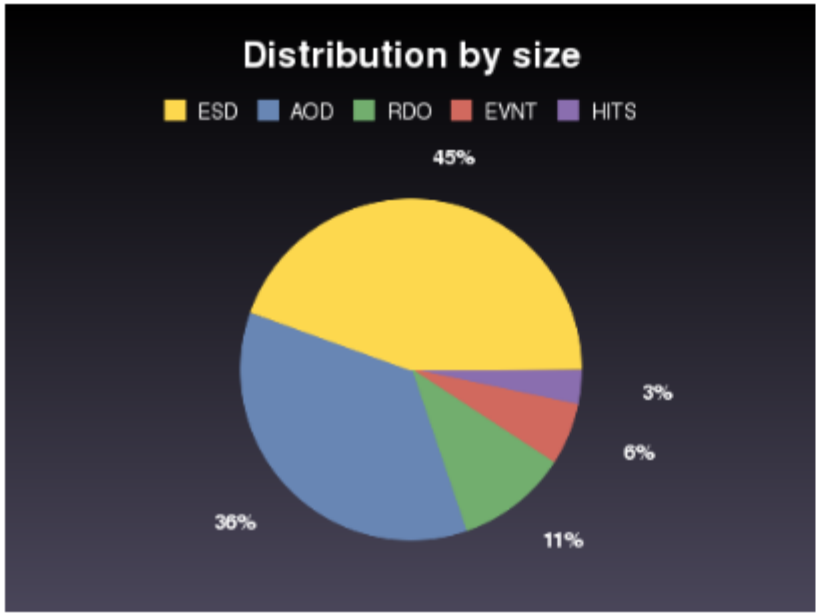

TYPE	# of Datasets	# of Files	Total Size(GB)
ESD	2472 / 24.4 %	660718 / 18.8 %	107271.34 GB/ 44.4 %
AOD	4621 / 45.7 %	2169207 / 61.6 %	86022.61 GB/ 35.6 %
RDO	620 / 6.1 %	233026 / 6.6 %	25446.91 GB/ 10.5 %
EVNT	921 / 9.1 %	224332 / 6.4 %	14151.44 GB/ 5.9 %
HITS	467 / 4.6 %	171590 / 4.9 %	7819.34 GB/ 3.2 %
NTUP	906 / 9.0 %	60429 / 1.7 %	750.68 GB/ 0.3 %
HPTV	3 / 0.0 %	50 / 0.0 %	6.72 GB/ 0.0 %
TXT	1 / 0.0 %	1 / 0.0 %	0.28 GB/ 0.0 %
TAG	3 / 0.0 %	12 / 0.0 %	0.01 GB/ 0.0 %
user	49 / 0.5 %	294 / 0.0 %	0.00 GB/ 0.0 %
HIST	42 / 0.4 %	524 / 0.0 %	0.00 GB/ 0.0 %
BS	8 / 0.1 %	0 / 0.0 %	0.00 GB/ 0.0 %
NTUP01	8 / 0.1 %	0 / 0.0 %	0.00 GB/ 0.0 %

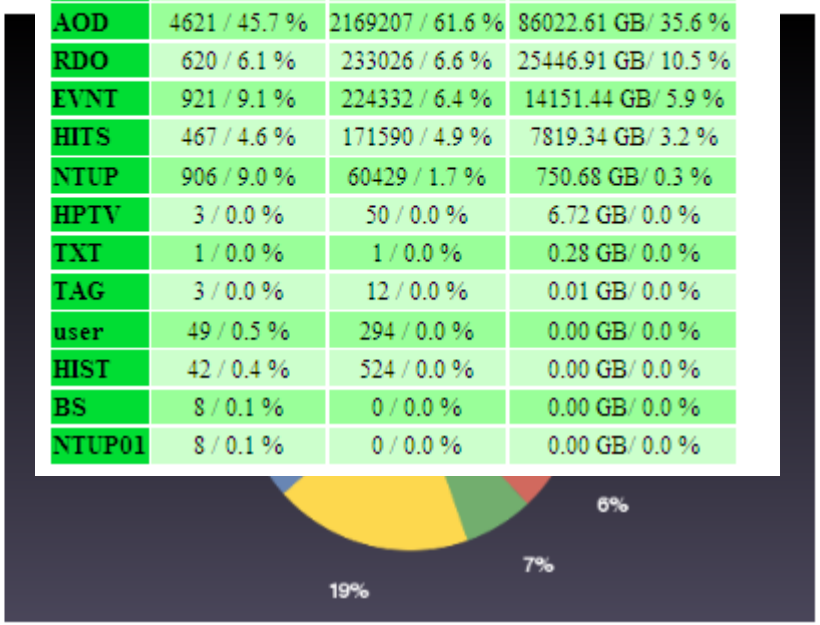

More web services

- dCache name space service
  - Filling LFC with dCache PNFSID and space token
  - BNL PILOT and PANDA Mover uses to check if files are already in read pools.
    - Pilot/Panda mover get PNFSID via getreplicax
    - It calls web API isFileInPool?pnfsid=XYZ
      - Main advantage – No load to PNFSD compared with dc_check which will check the metadata of the file
    - Copy files via dccp without name space

DDM monitor

- Not only the monitor is the DDM information page, it can act like DDM central/local catalog via its own web service.
  - Advantage
    - It has own record
    - No load to the central or local catalog.
    - It is a web service. No special client is needed.
  - Disadvantage
    - The record can be out of sync with the central or local catalog.
  - Future
    - Might get the direct connection to LFC. Won't be out of sync with the local catalog.

LFC web service

- Use LFC like LRC
- Web service is very convenient. Perfect for the local use.
  - getByGuid and getByGuids
  - getByGuidForSite
  - getByPnfsid and getByPnfsids
  - getBySfn
  - deleteBySfn
  - deleteByPnfsid