

Grid data transfer tools


Tier 3(g,w) meeting at ANL ASC – May 19, 2009

Marco Mambelli – University of Chicago

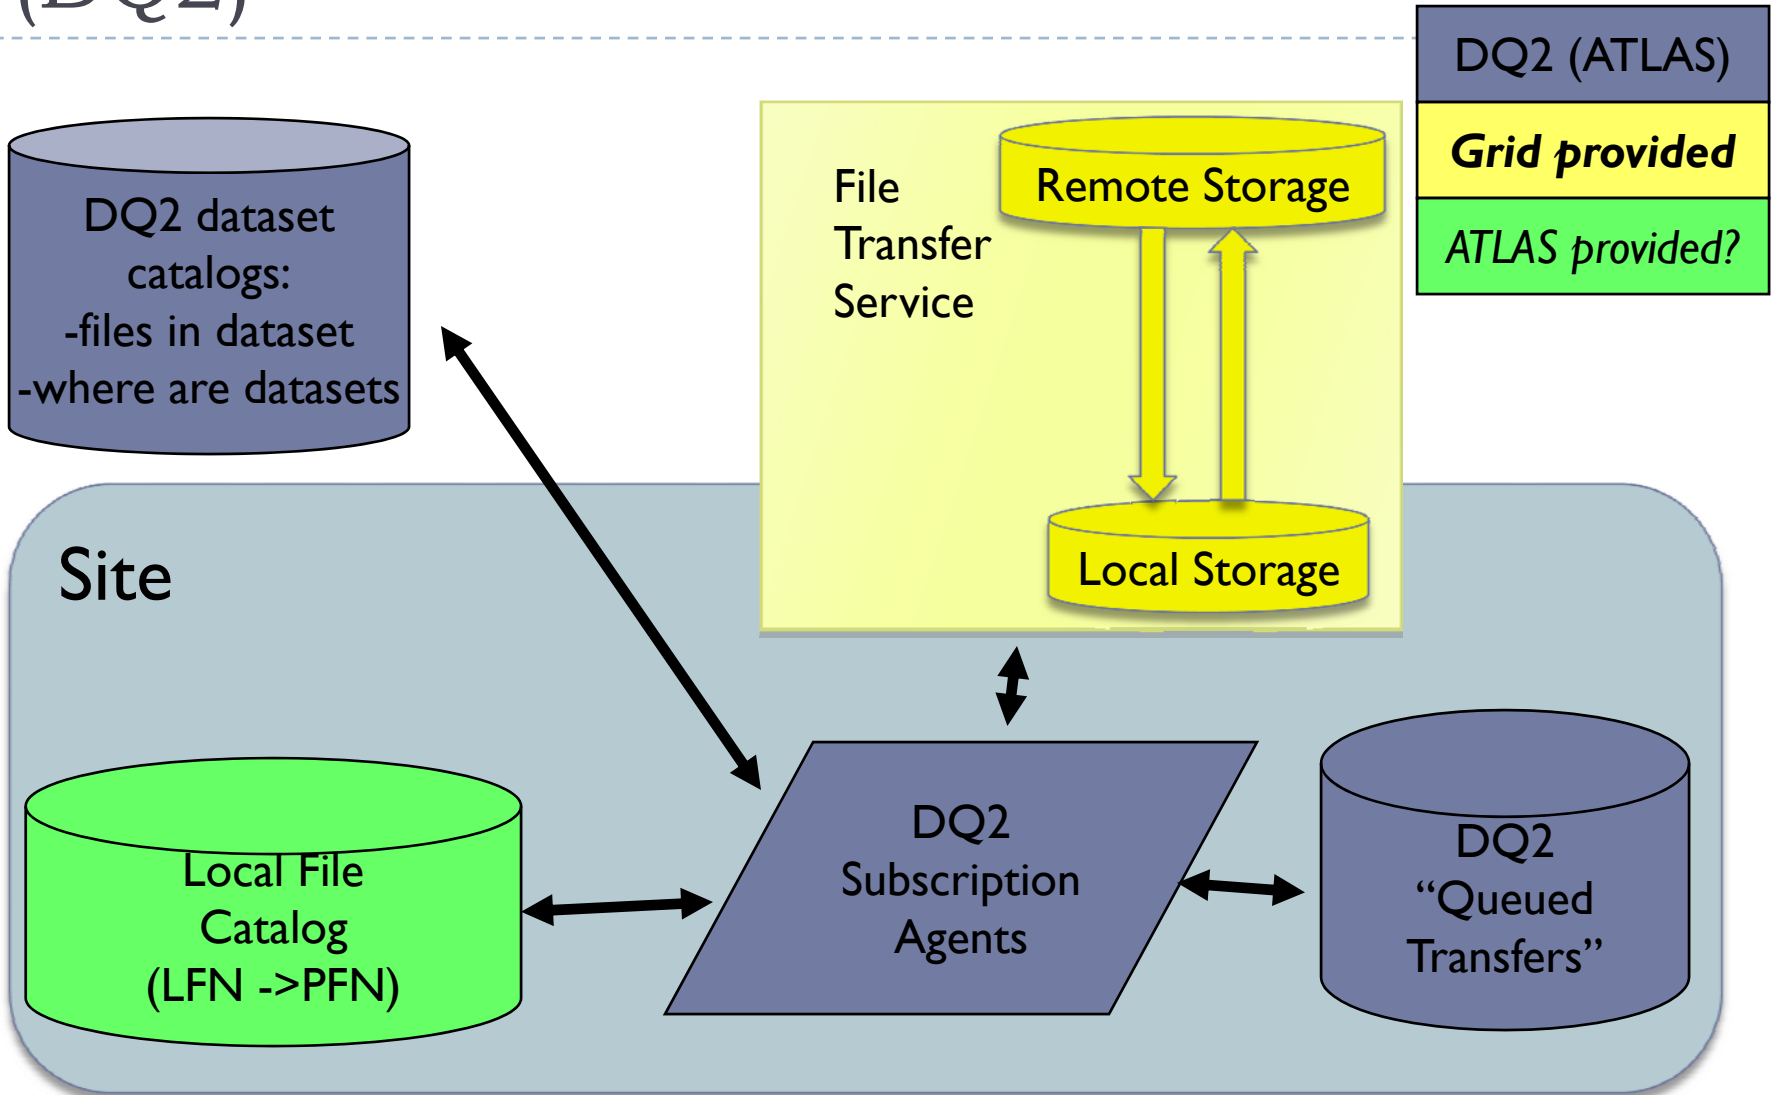
marco@hep.uchicago.edu

Grid/OSG components

- ▶ Standard installation
- ▶ Set of services or resources containing data and providing processing power
- ▶ Computing elements (CE)
- ▶ Storage elements (SE)
- ▶ Information systems
- ▶ Clients



ATLAS Distributed Data Management (DQ2)



ATLAS space token

token name	storage type	used for	@T2	@T1	@T0
ATLASDATATAPE	T1D0	RAW data, ESD, AOD from re-proc		X	X
ATLASDATADISK	T0D1	ESD, AOD from data	X	X	X
ATLASMCTAPE	T1D0	HITS from G4, AOD from ATLFAST		X	
ATLASMCDISK	T0D1	AOD from MC	X	X	X
ATLASPRODDISK	T0D1	buffer for in-and export	X		
ATLASGROUPDISK	T0D1	DPD	X	X	X
ATLASUSERDISK	T0D1	User Data	X	X *)	
ATLASLOCALGROUP DISK	T0D1	Local User Data @T3			

How data is organized at sites (SEs)



Name used to refer files

- ▶ **Logical File Name**

- ▶ lfn:/grid/dteam/vkoukis-gfal-test2

- ▶ **Grid Unique Identifier (GUID)**


- ▶ guid:4f74b453-5aaf-4af5-af2d-b438d081bb63

- ▶ **Storage URL (for a specific replica, on a specific Storage Element)**

- ▶ srm://se01.athena.hellasgrid.gr/pnfs/athena.hellasgrid.gr/data/dteam/generated/2007-04-15/file775276c2-7bbf-4f3d-92cf-5d5cdbaba254

- ▶ **Transport URL (for a specific replica, on an SE, with a specific protocol)**

- ▶ gsidcap://se01.athena.hellasgrid.gr:22128//pnfs/athena.hellasgrid.gr/data/dteam/generated/2007-04-15/file775276c2-7bbf-4f3d-92cf-5d5cdbaba254



Datasets and Containers

- ▶ Container datasets are the logical equivalent of physics datasets. They contain the same files.
 - ▶ Their name is the dataset name before the tid (or T for cosmic data). Name ends with a '/'
 - ▶ Containers can contain only DDM datasets. You cannot make containers of containers
- ▶ Dataset in DDM is a set of files
- ▶ For example, MC containers contain the files belonging to the tid datasets (containing tidxxxx with xxxx as task number).
- ▶ Most of the dq2 commands are similar for datasets and containers except few ones



Example

- ▶ Data :

data08_cosmag.□00091361.physics_RPCwBeam.recon.ESD_FILTERED.o4_f70/

- ▶ MC

:mc08.106020.PythiaWenu_1Lepton.recon.AOD.e352_s462_r541/



How to move data

- ▶ **DQ2 Client Tools**
 - ▶ command line
 - ▶ package developed by ATLAS
 - ▶ RPM and Pacman
 - ▶ uses Grid middleware
 - ▶ tested on WLCG-Client and gLiteUI



Discover and get data

- ▶ **DDM central catalog:**
 - ▶ list of sites contain at least a fraction of dataset
 - ▶ List of GUIDs per dataset + checksum + size (for consistency check)
- ▶ **LFC catalog : List of files on a site:**
 - ▶ Each catalog contains list of physical replicas for a GUID
 - ▶ Input : list of GUIDs
 - ▶ Middle : For each GUID, list of replicas managed by TI LFC
 - ▶ Output: List of GUIDs on the site
 - ▶ In addition: list of SURLS of the files (used by jobs)
- ▶ **Storage Element catalog:**
 - ▶ Map the SURL
- ▶ **WARNING** : No check of consistency between catalogs : DDM/ LFC/SE



Client installation

▶ Install DQ2 clients (and WLCG-client)

- ▶ `$ source /share/pacman/setup.sh`
- ▶ `$ mkdir /share/wlcg-client; cd /share/wlcg-client`
- ▶ `$ pacman -get http://www.mwt2.org/caches/:wlcg-client`
- ▶ (choose local installation 'l')
- ▶ `source setup.sh`
- ▶ `vi $VDT_LOCATION/vdt/etc/vdt-update-certs.conf`
- ▶ (choose OSG certificate set)
- ▶ `source $VDT_LOCATION/vdt-questions.sh;`
`$VDT_LOCATION/vdt/sbin/vdt-setup-ca-certificates`

▶ DQ2 client itself

- ▶ `mkdir /share/dq2-client; cd /share/dq2-client`
- ▶ `pacman -trust-all-caches -allow tar-overwrite -get`
`http://atlas.web.cern.ch/Atlas/GROUPS/DATABASE/project`
`/ddm/releases/pacman/cache:DQ2Clients`



Client setup

- ▶ **Setup with WLCG client (superset of OSG/VDT Client)**

- ▶ `$ source /share/osg-client/setup.sh`
- ▶ `$ source /share/dq2-client/setup.sh`
- ▶ (you need a grid certificate in ATLAS VO)
- ▶ `$ voms-proxy-init -voms atlas:/atlas`

- ▶ **Now you can:**

- ▶ move dataset
- ▶ find files
- ▶ copy files



Commands and Use

- ▶ <https://twiki.cern.ch/twiki/bin/view/Atlas/DQ2ClientsHowTo>
- ▶ `dq2-ls -fp DATASETNAME`
- ▶ `dq2-ls -L SITENAME -fp DATASETNAME`
- ▶ `dq2-get DATASETNAME`
- ▶ `dq2-get -f FILENAME DATASETNAME`
- ▶ `dq2-put -s SOURCEDIRECTORY DATASETNAME`
- ▶ `dq2-list-datasets-container CONTAINERNAME`
- ▶ `dq2-list-dataset-replicas-container CONTAINERNAME`
- ▶ `dq2-get -f FILENAMEI CONTAINERNAME`



References and other way to move data

▶ DQ2 Clients

- ▶ <https://twiki.cern.ch/twiki/bin/view/Atlas/DistributedDataManagement>
- ▶ <https://twiki.cern.ch/twiki/bin/view/Atlas/DQ2ClientsHowTo>
- ▶ <http://www.usatlas.bnl.gov/twiki/bin/view/Admins/DQ2ClientsInOSG>

▶ Subscriptions

- ▶ <https://savannah.cern.ch/projects/dq2-ddm-ops>

▶ Direct copy

- ▶ <http://www.usatlas.bnl.gov/twiki/bin/view/Admins/WlcgClient>

Examples for the first two follow



Subscription service

The screenshot shows a web browser window titled "Panda monitor and browser" with the URL <http://panda.cern.ch:25980/server/pandamon/query?mode=reqsubs0>. The page content includes a navigation menu at the top with links like "Production", "Clouds", "DDM", "PandaMover", "AutoPilot", "Sites", "Analysis", "Physics data", "Usage", "Plots", "ProdDash", and "DDMDash". A sidebar on the left contains sections for "CERN monitor", "Panda monitor", "Jobs", "Quick search", "Summaries", "Tasks", "Datasets", and "Datasets Distribution". The main content area is titled "Select Parameters For Subscription Request" and contains a section for "Information about request submission" with a list of parameters: Dataset pattern, Data Format, Software Version, Destination Tier, Validity, DN auth, and Request Approval policy. Below this is a form with "Request Parameters" and "Control Parameters" sections, each containing several dropdown menus and input fields. A "Continue" button is at the bottom of the form, followed by a note about certificate registration.

Subscription Request Parameters:

- Dataset Pattern:** "pattern*format*version", length < 100 (see [AMI](#) for details about ATLAS official dataset names)
Datasets State must be "frozen"
- Data Format:** None
- Software Version:** version of the SW
- Destination Tier:** site from TiersOfATLAS (see [ToA](#) for details) It must be **DISK**-storage according to CREM policy
- Validity:** OneTimeCopy / Periodic
- DN auth:** For registration please proceed to the [registration form](#) (check status of your account and **add information about your certificate** click [here](#))
- Request Approval policy:** [link](#)

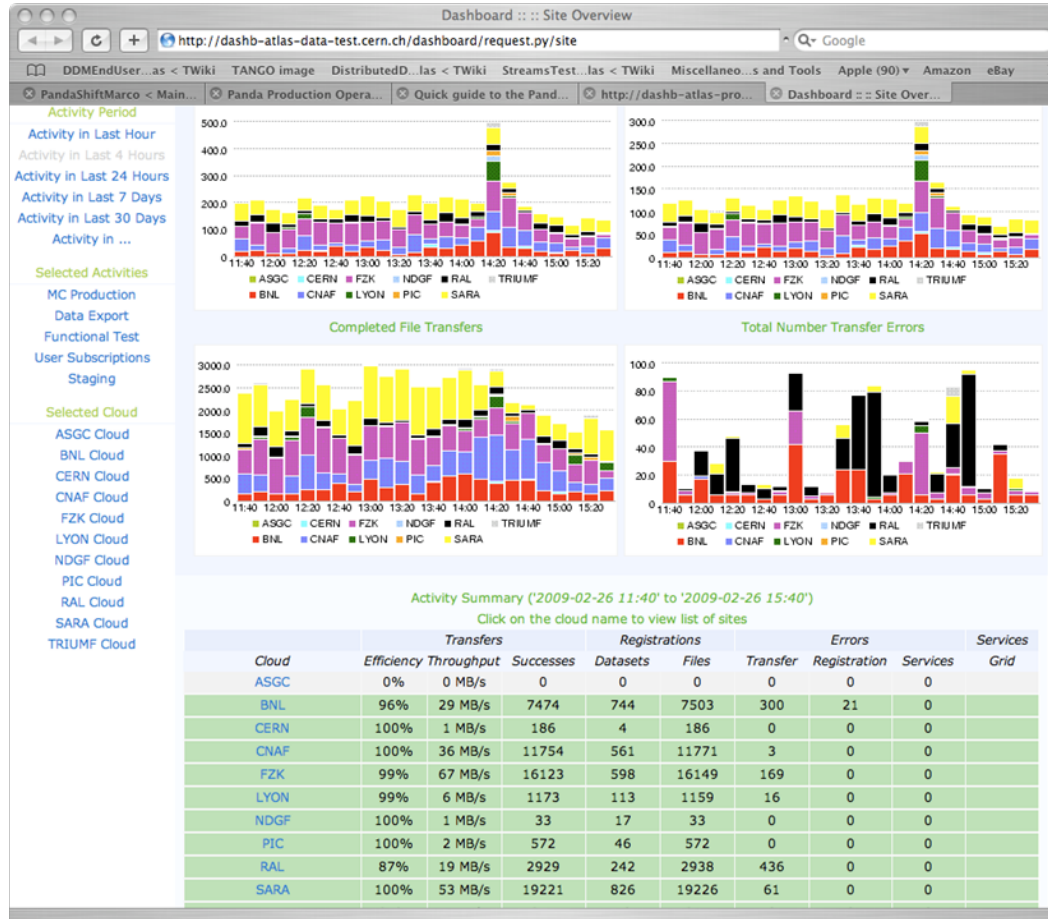
Form Fields:

- Request Parameters:** Dataset Pattern (text), Data Format (None), Software Version (None), Destination Tier (cloud, tier), Req Type (1beam)
- Control Parameters:** Data Management Mode (DataTransfer), Validity (transfer only) (OneTimeCopy), Priority (Immediate), Transfer Volume (100 %), Comments (text)

[Continue](#)
For data submitting certificate is needed. Please follow this [link](#) to register certificate.

<http://panda.cern.ch:25980/server/pandamon/query?mode=reqsubs0>

DDM Dashboard



Can be used to check the status of your (and other) subscriptions
Use drill-down menus

DDM Dashboard detail

The screenshot shows the DDM Dashboard interface for MWT2_DATADISK. The browser address bar indicates the URL: http://dashb-atlas-data-test.cern.ch/dashboard/request.py/dataset?site=MWT2_DATADISK. The dashboard includes navigation tabs for Overview, Dataset Info, Page Help, User Guide, and Feedback. A summary section shows the following counts: QUEUED (0), COMPLETE (351), BROKEN (0), INCOMPLETE (0), STAGED (0), and CANCELED (0). A table below lists dataset entries with columns for Last Update, Dataset Name, and State. The table shows 351 datasets, all in a COMPLETE state, with various dataset names and last update timestamps.

Last Update	Dataset Name	State
2009-02-26 14:26:10	ccrc08_run2.1020091000000L.physics_Minbias.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:26:08	ccrc08_run2.1020091000000L.physics_Jet.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:26:06	ccrc08_run2.1020091000000L.physics_Bphys.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:26:04	ccrc08_run2.1020091000000L.physics_Muon.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:26:01	ccrc08_run2.1020091000000L.physics_Muon.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:59	ccrc08_run2.1020091000000L.physics_Jet.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:57	ccrc08_run2.1020091000000L.physics_Bphys.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:55	ccrc08_run2.1020091000000L.physics_Egamma.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:53	ccrc08_run2.1020091000000L.physics_Jet.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:05	ccrc08_run2.1020091000000L.physics_Minbias.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:03	ccrc08_run2.1020091000000L.physics_Muon.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:25:00	ccrc08_run2.1020091000000L.physics_Egamma.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:24:58	ccrc08_run2.1020091000000L.physics_Minbias.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:24:56	ccrc08_run2.1020091000000L.physics_Egamma.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 14:24:53	ccrc08_run2.1020091000000L.physics_Bphys.merge.AOD.o0_r0_t0	COMPLETE
2009-02-26 09:01:11	user.HironoriIto.Dq2Ping_4_03	COMPLETE
2009-02-26 08:01:20	user.HironoriIto.Dq2Ping_4_02	COMPLETE
2009-02-26 07:01:07	user.HironoriIto.Dq2Ping_4_01	COMPLETE
2009-02-26 06:00:57	user.HironoriIto.Dq2Ping_4_00	COMPLETE

Here you see the subscriptions to Chicago (BNL cloud)

Some observations

- ▶ Works generally for $O(\text{TB})$ datasets
- ▶ Large number of small files are inefficient (avoid $O(100)$ files datasets)
- ▶ DDM operations for support



DQ2 Clients

- ▶ Use the how-to page as reference (open in browser)
 - ▶ Live demo (terminal)
 - ▶ setup
 - ▶ voms proxy
 - ▶ `dq2-ls, -L ROAMING, name*`; hat's in the output
 - ▶ containers (= physics datasets)
 - ▶ `_tidXXXX` datasets
 - ▶ non `_tid` datasets – **DO NOT USE THIS!**
 - ▶ `dq2-list-dataset-replicas-container`, explain output
 - ▶ which site (SE) has all your files?
 - ▶ # of complete (datasets) = Total datasets
 - ▶ verify that the files are there: `dq2-ls -f` (verify path `-p`)
 - ▶ use `-L site_that_you_want_to_check`
-



DQ2 Clients (2)

▶ (cont)

- ▶ Create PoolFileCatalog.xml dq2-ls -P
 - ▶ --replace-io=REPLACEIO ('regularexpression^newvalue')
 - ▶ e.g. srm://uct2-dcl.uchicago.edu(:[0-9]+)(/srm/managerv2SFN=)/pnfs/^dcache:/pnfs/
- ▶ Errors in dq2-get
 - ▶ no replica available (check dq2-ls -r)
 - ▶ files not copied: try different 'protocol': ng,lcg,lcgl,srm,dcap,rftio,...
 - ▶ checksum error: report to DDM operations

▶ Help

- ▶ DDM operation savannah portal or [hn-atlas-DDMOperations](#)
- ▶ hn-atlas-dist-analysis-help hypernews



Demo 1

```
[uct3-edge5] /ecache/marco/test_dq2-get > ./share/wlwg-client/setup.sh
> ./share/dq2-client/setup.sh
> voms-proxy-init -voms atlas
Enter GRID pass phrase:
Your identity: /DC=org/DC=doegrids/OU=People/CN=Marco Mambelli 325802
Creating temporary proxy ..... Done
Contacting lcg-voms.cern.ch:15001
[/DC=ch/DC=cern/OU=computers/CN=lcg-voms.cern.ch] "atlas" Done
Creating proxy ..... Done
Your proxy is valid until Thu Feb 26 22:13:02 2009
> dq2-ls -L ROAMING mc08.107003.singlepart_e_Et25.recon.AOD.*
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r604_tid032272_sub03552716
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r604_tid032272_sub03551456
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563
...
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r604/
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r604_tid032271
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r604_tid032272
```



Demo 2

```
> dq2-list-dataset-replicas-container \  
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563/  
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027926:  
  INCOMPLETE:WISC_MCDISK  
  COMPLETE:AGLT2_MCDISK,BNL-OSG2_MCDISK,BNLPANDA,...  
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027925:  
  INCOMPLETE:RRC-KI_MCDISK  
  COMPLETE:AGLT2_MCDISK,BNL-OSG2_MCDISK,BNLPANDA,...  
Container name: mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563/  
Total datasets: 2
```

Summary:

SITE	/ # COMPLETE / # INCOMPLETE / TOTAL		
UKI-SCOTGRID-GLASGOW_MCDISK	1	0	1
WISC_MCDISK	1	1	2
CERN-PROD_MCDISK	2	0	2
...			
SARA-MATRIX_MCDISK	1	0	1
MWT2_UC_MCDISK	2	0	2

Demo 3

```
> dq2-ls -L MWT2_UC_MCDISK -f mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027926
mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027926
[X] AOD.027926._00507.pool.root.3 48B4D116-38AE-DD11-B876-00A0D1E7FB98
md5:ba76b8c771a11c03311272462855d0fb 14325205
[X] AOD.027926._00196.pool.root.1 5E39D1C3-31AE-DD11-9C54-00A0D1E7BC5E
md5:6a2fd3ad0464c4da30faaf05ca55b86a 13767732
[X] AOD.027926._00249.pool.root.1 0A947510-35AE-DD11-9E81-00A0D1E70C54
md5:797dfea0f1b20983231e18f96c6e060e 14279200
...
[X] AOD.027926._00209.pool.root.1 3C2D1AC7-31AE-DD11-9687-00A0D1E7B60E
md5:2b85fcf06529024c8ec3d5b7725e5f58 14399761
[X] AOD.027926._00765.pool.root.1 B0773A0C-A8AF-DD11-ABD6-00A0D1E50051
md5:099b1c154adf4e4376def5afd93c3382 14354310
total files: 720
local files: 720
total size: 10300924545
date: 2008-11-11 04:53:23
> dq2-ls -L MWT2_UC_MCDISK -f mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027925 | \
tail -n 4
total files: 80
local files: 80
total size: 1134078206
date: 2008-11-11 04:38:26
> dq2-ls -L MWT2_UC_MCDISK -fp mc08.107003.singlepart_e_Et25.recon.AOD.e342_s462_r563_tid027926
```

