

Progress Federating the Storage of USATLAS

R. Gardner
for the Federated ATLAS Xrootd group - FAX

21-Nov-11

FAX Working Group

- Led by Wei Yang, R.Gardner (US facilities integration)
- Doug Benjamin
- Andy Hanushevsky
- Hiro Ito
- Patrick McGuigan
- Shawn McKee
- Ofer Rind
- Horst Severini
- Sarah Williams
- Chicago Face-to-face workshop (Sep 12-13)
- bi-weekly meetings
- All notes available on twiki

ATLAS Goal

(the lofty one)

- Provide seamless access to every storage resource from every core
- direct access, copy-store, copy-cache, and sub-file copy cache, as warranted (depending on site, task requirements, etc)
- whether from an ANALY queue or off-grid (cloud, campus, or an HPC resource)

ATLAS Goal

(near term)


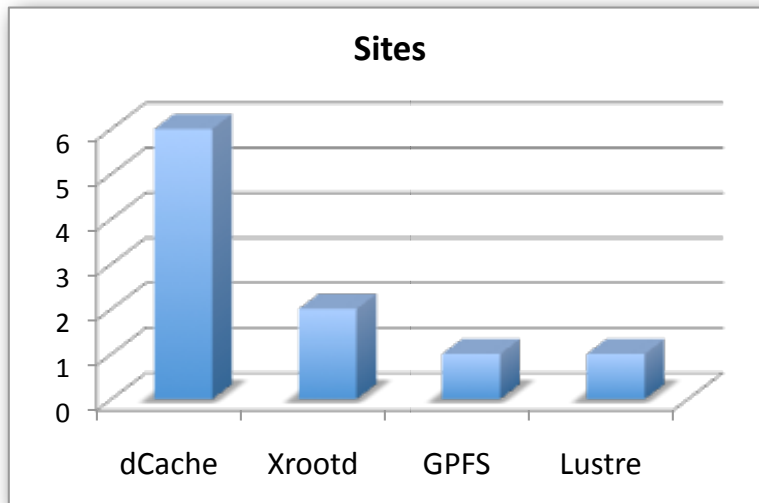
- T3 access to other T3s via global name
- T3 access to T2 via global name
- T3 access to T1 via global name
- T2 access to itself via global name
- T2 access to T1 via global name
- T2 access to other T2 via global name
- T2 access to T3 in the federation

Panda
ANALY
queue
or
T3g-like
tests

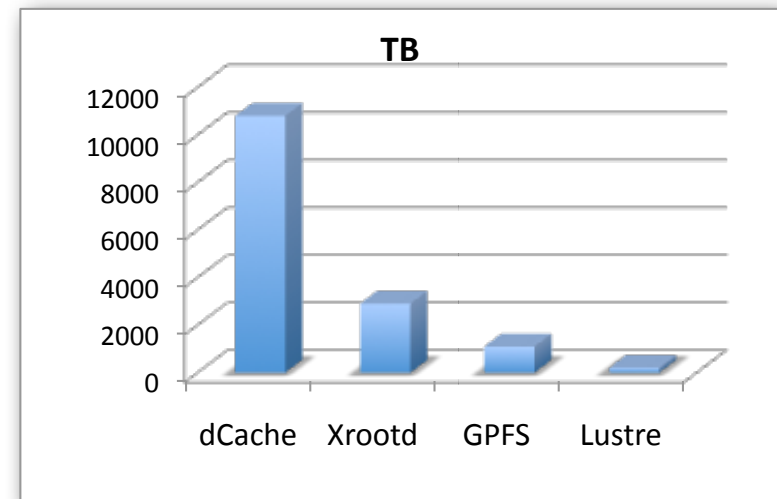
Data Sources in the US Cloud

- Storage resources at T1 and 5 Tier2 centers (10 sites) currently total 14.8 PB disk
- Three Tier2 centers are multi-site and share distributed storage resources across WAN (AGLT2, MWVT2, NET2)
- A large number of analysis T3g sites
- A small number of Tier3gs (grid) sites

open every channel possible for user access within reason

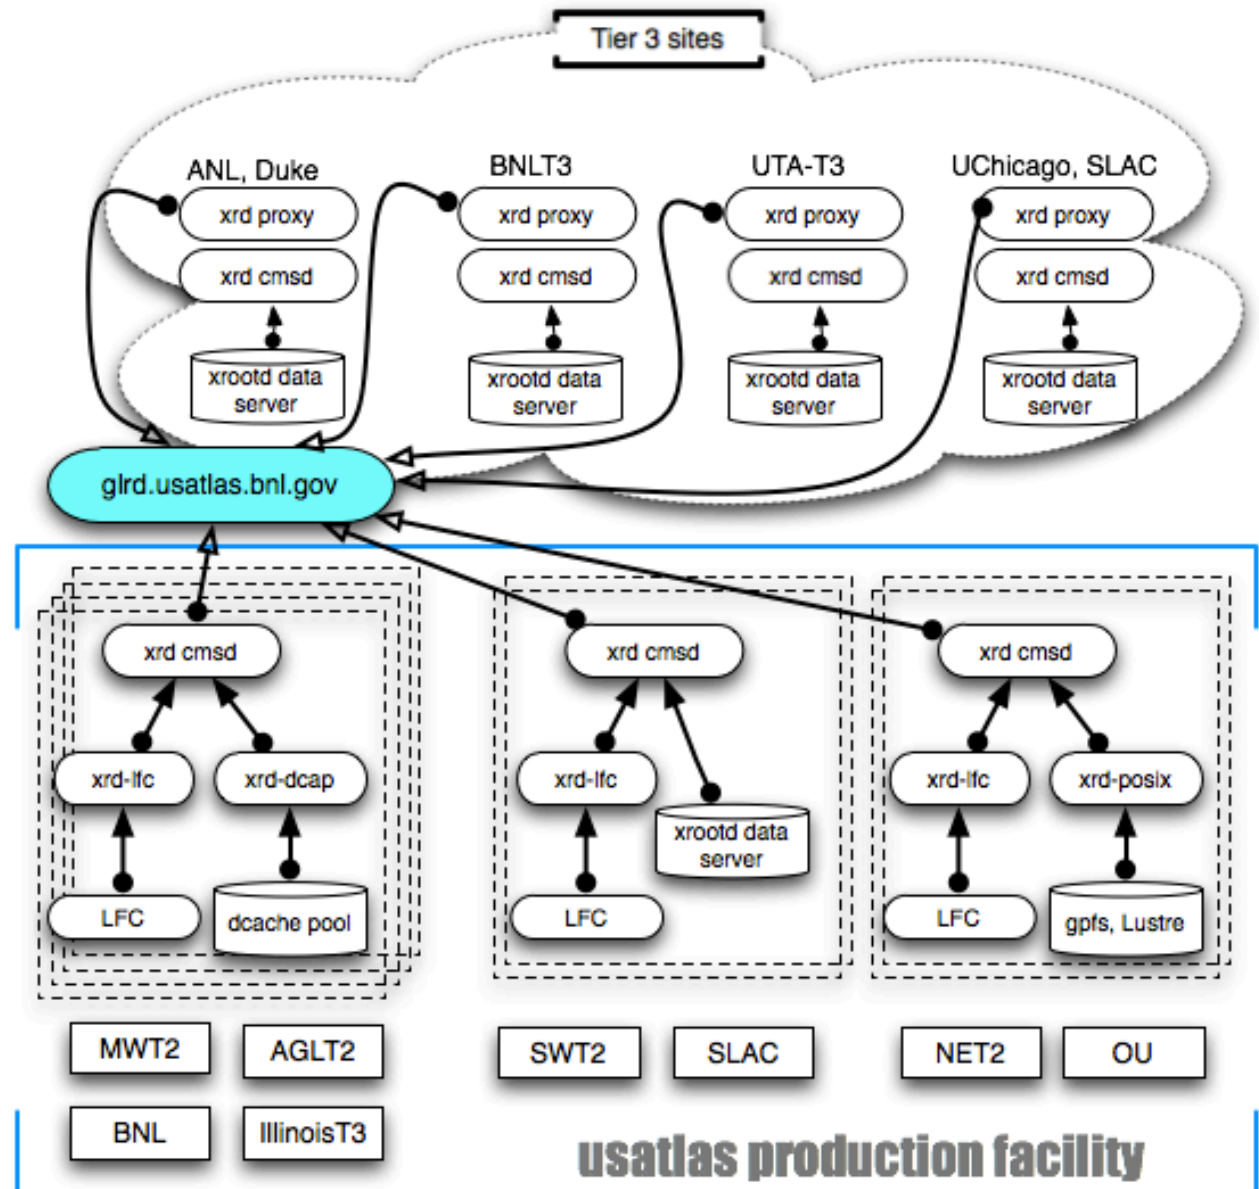


- Over 10 PB in dCache
- Over 2 PB in Xrootd
- All sites continue to add
- April '12 will have:
 - >2.2 PB at each T2
 - >6.3 PB at T1
 - >17.3 PB total



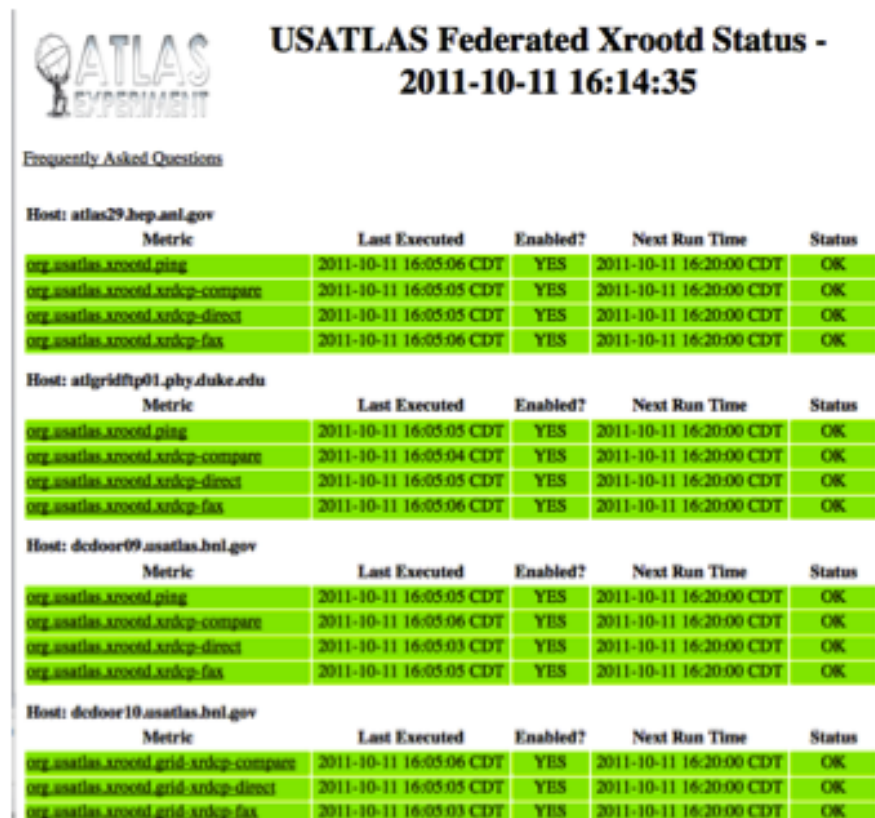
current deployment

- 6 tier3g
(4 co-located w/ t2 or tier1)
- all 5 tier2s
- BNL
- tier3gs
- OSG ITB site
- 15 endpoints to 4 backend storage types



FAX Status Monitor

- <https://uct3-xrdp.uchicago.edu:8443/rsv/>
- based on OSG RSV framework
- Probes sites every 15 minutes
- Tests direct and via global redirectory
- Also does simple ping and file comparison checks



USATLAS Federated Xrootd Status - 2011-10-11 16:14:35

Frequently Asked Questions

Host: atlas29.hep.anl.gov

Metric	Last Executed	Enabled?	Next Run Time	Status
org.asatlas.xrootd.ping	2011-10-11 16:05:06 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-compare	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-direct	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-fax	2011-10-11 16:05:06 CDT	YES	2011-10-11 16:20:00 CDT	OK

Host: atigrifftp01.phy.duke.edu

Metric	Last Executed	Enabled?	Next Run Time	Status
org.asatlas.xrootd.ping	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-compare	2011-10-11 16:05:04 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-direct	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-fax	2011-10-11 16:05:06 CDT	YES	2011-10-11 16:20:00 CDT	OK

Host: dcdoor09.asatlas.bnl.gov

Metric	Last Executed	Enabled?	Next Run Time	Status
org.asatlas.xrootd.ping	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-compare	2011-10-11 16:05:06 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-direct	2011-10-11 16:05:03 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.xrdcp-fax	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK

Host: dcdoor10.asatlas.bnl.gov

Metric	Last Executed	Enabled?	Next Run Time	Status
org.asatlas.xrootd.grid.xrdcp-compare	2011-10-11 16:05:06 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.grid.xrdcp-direct	2011-10-11 16:05:05 CDT	YES	2011-10-11 16:20:00 CDT	OK
org.asatlas.xrootd.grid.xrdcp-fax	2011-10-11 16:05:03 CDT	YES	2011-10-11 16:20:00 CDT	OK

Simple ROOT analysis job

- Higgs dilepton event skim over MC dataset:

```
mc10_7TeV.
```

```
116700.PowHegPythia_ggH110_WW2lep.merge.N  
TUP_SMWZ.e773_s933_s946_r2302_r2300_p591
```

- No TTreeCache buffering initially
- 20 input files, each ~ 550 MB (11 GB total)
- Placed this data at several FAX sites for testing

Simple 20 file test initial results


- Running jobs on a tier 3 node
- Takes 1.5 min if data were staged to local scratch
- Took 2.5 minutes to read the same dataset from MWT2's redirector
- Took 20 minutes when pointing to glrd.usatlas.org
- Welcome to the federation!

Redirector & endpoints

- The redirector has some configurability but as yet knows nothing of latency between clients and sources, or the performance of sites (default: first responder, round-robin thereafter) - load & timezone can be used
- Decide to first look at endpoints directly
- Choose Illinois (5 ms RTT): 6 min
- Choose AGLT2 (7 ms RTT): 5.5 min
- Choose BNL (25-42 ms RTT): 12.5 min

Latencies from mwt2.org

Host: atlas29.hep.anl.gov	4.0 ms
Host: atlgridftp01.phy.duke.edu	25.2 ms
Host: dcdoor09.usatlas.bnl.gov	24.6 ms
Host: dcdoor10.usatlas.bnl.gov	35-42 ms
Host: gk06.atlas-swt2.org	25-36 ms
Host: glrd.usatlas.org	35-42 ms
Host: griddev01.slac.stanford.edu	52 ms
Host: griddev02.slac.stanford.edu	52 ms
Host: gw01.tier3-atlas.uta.edu	22 ms
Host: itb3.uchicago.edu	< 1ms
Host: manage.aglt2.org	6.7 ms
Host: osgx3.hep.uiuc.edu	3.38 ms
Host: ouhep03.nhn.ou.edu	19.2 ms
Host: tier2-03.oceph.ou.edu	19.5 ms
Host: uct2-grid5.uchicago.edu	< 1ms
Host: uct3-xrdp.uchicago.edu	< 1ms
Host: lxplus.cern.ch	140 ms
Host: ccsqfatlasli02.in2p3.fr	133 ms



Scalability testing I: # clients

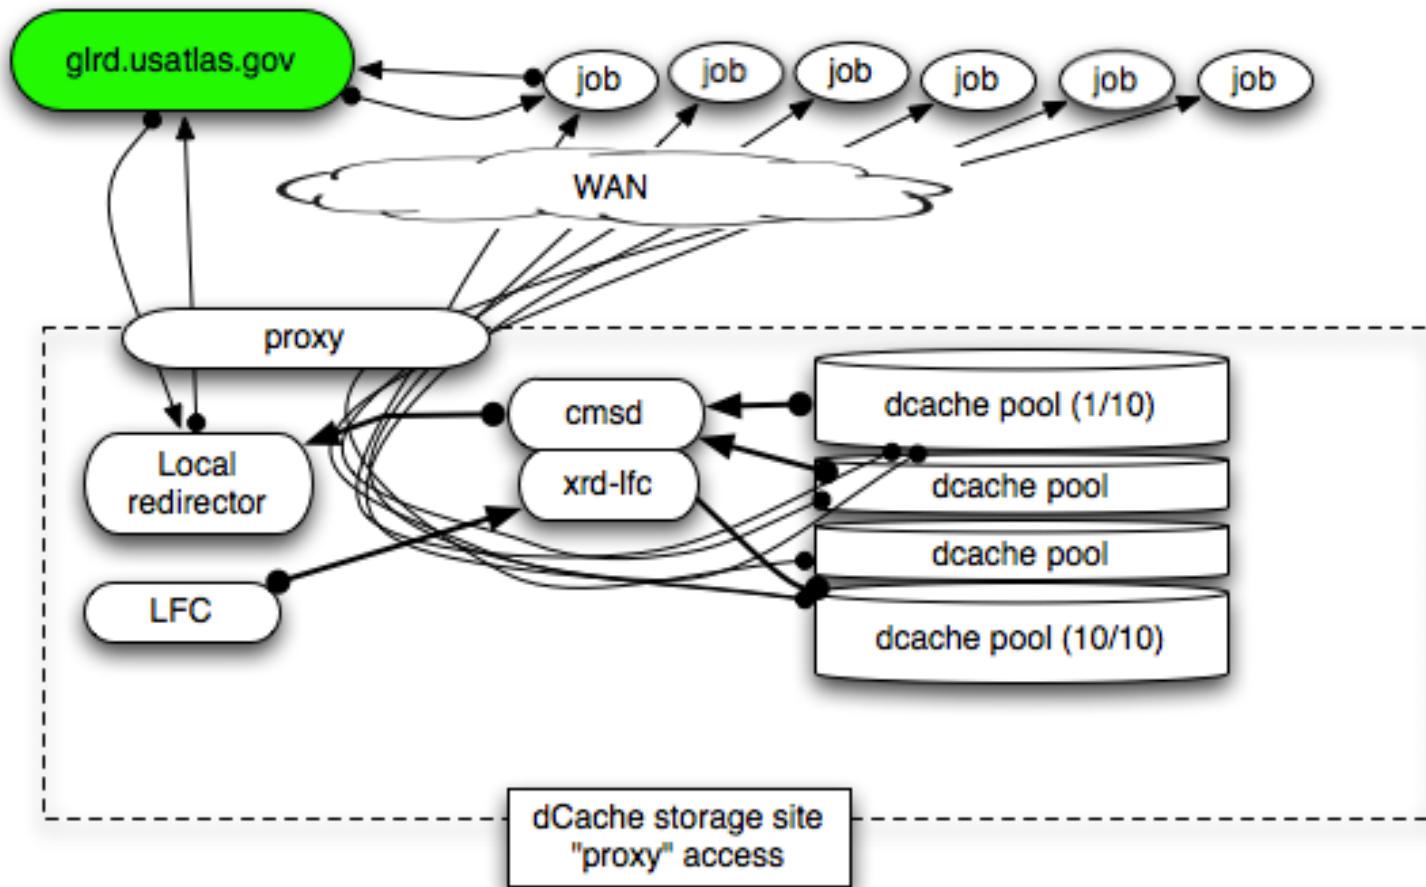
- Goal was to assess impact of potentially large numbers of clients
- Used occasion to tune up doors (proxy cluster at SLAC and BNL door)
- Submit simultaneously large numbers of analysis jobs over this data set

~200 job client test

- Ran jobs “under” a fully loaded MWT2 cluster - so not to infer too much about absolute performance
- 196 jobs, 2 jobs per node, 98 nodes

site	avg wall time (min)	failures	avg wall (min) post fix proxy work
mwt2	3.5	0	na
uct3	2.75	0	na
aglt2	30	0	na
slac	22	0	na
swt2	12	0	na
bnl	192	0	20
glrd	21	0	na
uct3-proxy	13	0	na (needs 10G)

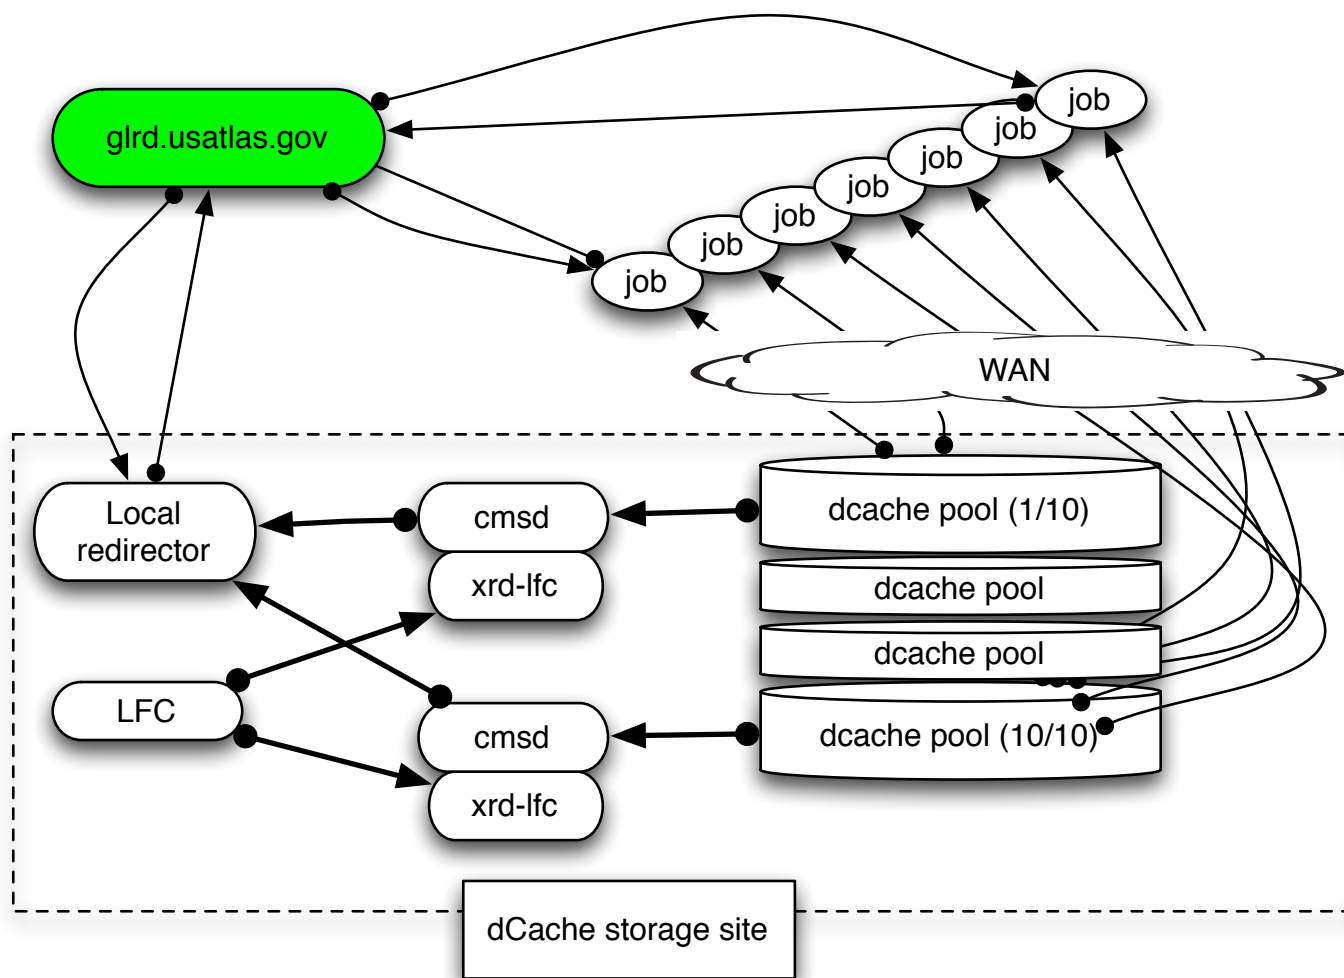
proxy access



access is
through a
dCache
“proxy”
server

bnl, aglt2

direct access to pools



run xrd
on every
pool

mwt2

Large scale job tests

- Same procedure as before (simultaneous launch)
- ~1000 jobs (1120 each)

site	avg wall time (min)	failures	avg wall (min) post fix proxy work
mwt2	6.3	0	na
IllinoisHEPT3	400	2	to do
swt2	35	0	na
slac	na	0	28

- ~5000 jobs (n.b. 24 per node! ... plus normal production)

site	# jobs submitted	failures	avg wall time (min)	avg wall (min) post fix proxy work
mwt2	4988	0	40	na
slac	2938	0	888	(5774 jobs) 70

Panda FAX testing

- Standard prun submission of same thing:

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1342632498 Robert W. Gardner Jr.	jobsetID= 2413 runGen-00-00-02	finished	2011-10-24 19:24	0:00:13	0:10:27	10-24 19:34	US/US.ANALY MWT2 , analysis-run	1000
In: mc10_7TeV.116700.PowHegPythia_ggH110_WW2lep.merge.NTUP_SMWZ.e773_s933_s946_r2302_r2300_p591/ Out: user.rwg.HWWlv.1319484215.0/								

- 10.5 minutes; longer due to stage-out (SRM)
- Now, submit using FAX-like submission directly to local redirector, using the global namespace (pfnList):

```
root://xrd.mwt2.org//atlas/dq2/mc10_7TeV/NTUP_SMWZ/
e773_s933_s946_r2302_r2300_p591/mc10_7TeV.
116700.PowHegPythia_ggH110_WW2lep.merge.NTUP_SMWZ.e773_s933_
s946_r2302_r2300_p591_tid408565_00/NTUP_SMWZ.
408565._000001.root.1
```

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1348883139 Robert W. Gardner Jr.	jobsetID= 2433 runGen-00-00-02	finished	2011-11-01 17:26	1:45:38	0:13:27	11-01 19:25	US/US.ANALY MWT2 , analysis-run	1000
Out: user.rwg.HWWlv.1320168377.0/								

- 13.5 minutes; increase could be attributed to N2N lookup - or unlucky node or other conditions, needs study

Panda-FAX (2)

- Now run the job at AGLT2, reading MWT2

<u>PandaID</u> , <u>Owner</u> , Working group	<u>Job</u>	<u>Status</u>	<u>Created</u>	<u>Time to start</u>	<u>Duration</u>	<u>Ended/ Modified</u>	<u>Cloud/Site, Type</u>	<u>Priority</u>
1348879908 Robert W. Gardner Jr.	jobsetID=2429 runGen-00-00-02	finished	2011-11-01 17:23	0:00:38	0:16:28	11-01 17:40	US/US.ANALY_AGLT2, analysis-run	1000
Out: user.rwg.HWW/vlv.1320168198.0/								

- 16.5 minutes (~20% hit) vs local
- Now use the global redirector glrd.usatlas.org, from AGLT2:

<u>PandaID</u> , <u>Owner</u> , Working group	<u>Job</u>	<u>Status</u>	<u>Created</u>	<u>Time to start</u>	<u>Duration</u>	<u>Ended/ Modified</u>	<u>Cloud/Site, Type</u>	<u>Priority</u>
1348882518 Robert W. Gardner Jr.	jobsetID=2431 runGen-00-00-02	finished	2011-11-01 17:25	0:00:10	0:21:35	11-01 17:47	US/US.ANALY_AGLT2, analysis-run	1000
Out: user.rwg.HWW/vlv.1320168321.0/								

- 21.5 minutes - but data could have come from anywhere

Panda-FAX (3)

- Tier 2 reading from a Tier3GS (IllinoisHEP):

<u>PandaID</u> , <u>Owner</u> , <u>Working group</u>	<u>Job</u>	<u>Status</u>	<u>Created</u>	<u>Time to start</u>	<u>Duration</u>	<u>Ended/ Modified</u>	<u>Cloud/Site, Type</u>	<u>Priority</u>
1348924098 Robert W. Gardner Jr.	jobsetID= 2437 runGen-00-00-02	finished	2011-11-01 18:28	0:00:40	0:18:32	11-01 18:47	US/US.ANALY IllinoisHEP , analysis-run	1000
Out: user.rwg.HWWvlv.1320172087.0/								

- 18 minutes (5 ms RTT) - similar to AGLT2

I 20 file test

- ANALY_MWT2 (data local) 19.5 mins

PandaID , Owner , Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1363242516 Robert W. Gardner Jr.	jobsetID= 2441 runGen-00-00-02	finished	2011-11-20 19:29	0:00:14	0:19:35	11-20 19:49	US/US.ANALY_MWT2 , analysis-run	1000
Out: user.rwg.HWWlv.1321817331.0/								

- ANALY_AGLT2 (data at MWT2) - 47 mins

PandaID , Owner , Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1363325024 Robert W. Gardner Jr.	jobsetID= 2447 runGen-00-00-02	finished	2011-11-20 22:12	0:01:22	0:47:09	11-20 23:00	US/US.ANALY_AGLT2 , analysis-run	1000
Out: user.rwg.HWWlv.1321827085.0/								

- ANALY_BNL_ATLAS (data at MWT2)
90 mins

PandaID , Owner , Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1363343536 Robert W. Gardner Jr.	jobsetID= 2451 runGen-00-00-02	finished	2011-11-20 22:36	0:00:13	1:29:54	11-21 00:06	US/US.ANALY_BNL_ATLAS_1 , analysis-run	1000
Out: user.rwg.HWWlv.1321828533.0/								

I20 file test (transatlantic)

- ANALY_CERN_XROOTD - 5h53m

PandaID , Owner , Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1363344462 Robert W. Gardner Jr.	jobsetID=2453 runGen-00-00-02 Out: user.rwg.HWWlv.1321828613.0/	finished	2011-11-20 22:37	0:04:15	5:53:57	11-21 04:35	CERN/CERN.ANALY_CERN_XROOTD , analysis-run	999

- ANALY_IN2P3-CC - 6h11m

PandaID , Owner , Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority
1363348433 Robert W. Gardner Jr.	jobsetID=2455 runGen-00-00-02 Out: user.rwg.HWWlv.1321828958.0/	finished	2011-11-20 22:43	0:02:05	6:11:51	11-21 04:57	FR/FR.ANALY_IN2P3-CC , analysis-run	999

summary of direct access federation testing

Mode	Tested	Type	relative performance
• T3 access to other T3s via global name		local script	Limited testing so far - T3gs tested
• T3 access to T2 via global name		local script	Good
• T3 access to T1 via global name		local script	Good
• T2 access to itself via global name		local script, Panda	Excellent
• T2 access to T1 via global name		local script, Panda	Good
• T2 access to other T2 via global name		Panda	Varies (region too large? not all proxy are tuned)
• T2 access to T3 in the federation		local script, Panda(!)	Can be good

- Keep in mind no TTreeCache'ing in these jobs
- Still a long way to go before anything like local performance

Conclusions

- Focus in this talk was on direct access tests directly to federated storage
- Proxy-fronted storage (esp dCache) need further work while direct pool access very good
- Demonstrated new functional modes of accessing data from within existing Panda ANALY queue framework
 - no changes needed; any ANALY queue can run FAX-like jobs
- Redirector scheduling to do yet & client buffer
- Other work (not discussed here but later) on FRM and internal site caching