

Status of the FAX federation

Rob Gardner
Computation and Enrico Fermi Institutes
University of Chicago

ATLAS Tier 1 / Tier 2 / Tier 3 Jamboree
December 10-11, 2012



Recall our goals



- Common ATLAS namespace across all storage sites, accessible from anywhere
- Easy to use, homogeneous access to data
- Identified initial use cases
 - Failover from stage-in problems with local SE
 - Now implemented, in production on several sites
 - Gain access to more CPUs using WAN direct read access
 - Allow brokering to Tier 2s with partial datasets
 - Opportunistic resources without local ATLAS storage
 - Use as caching mechanism at sites to reduce local data management tasks
 - Eliminate cataloging, consistency checking, deletion services
- WAN data access group formed in ATLAS to determine use cases & requirements on infrastructure

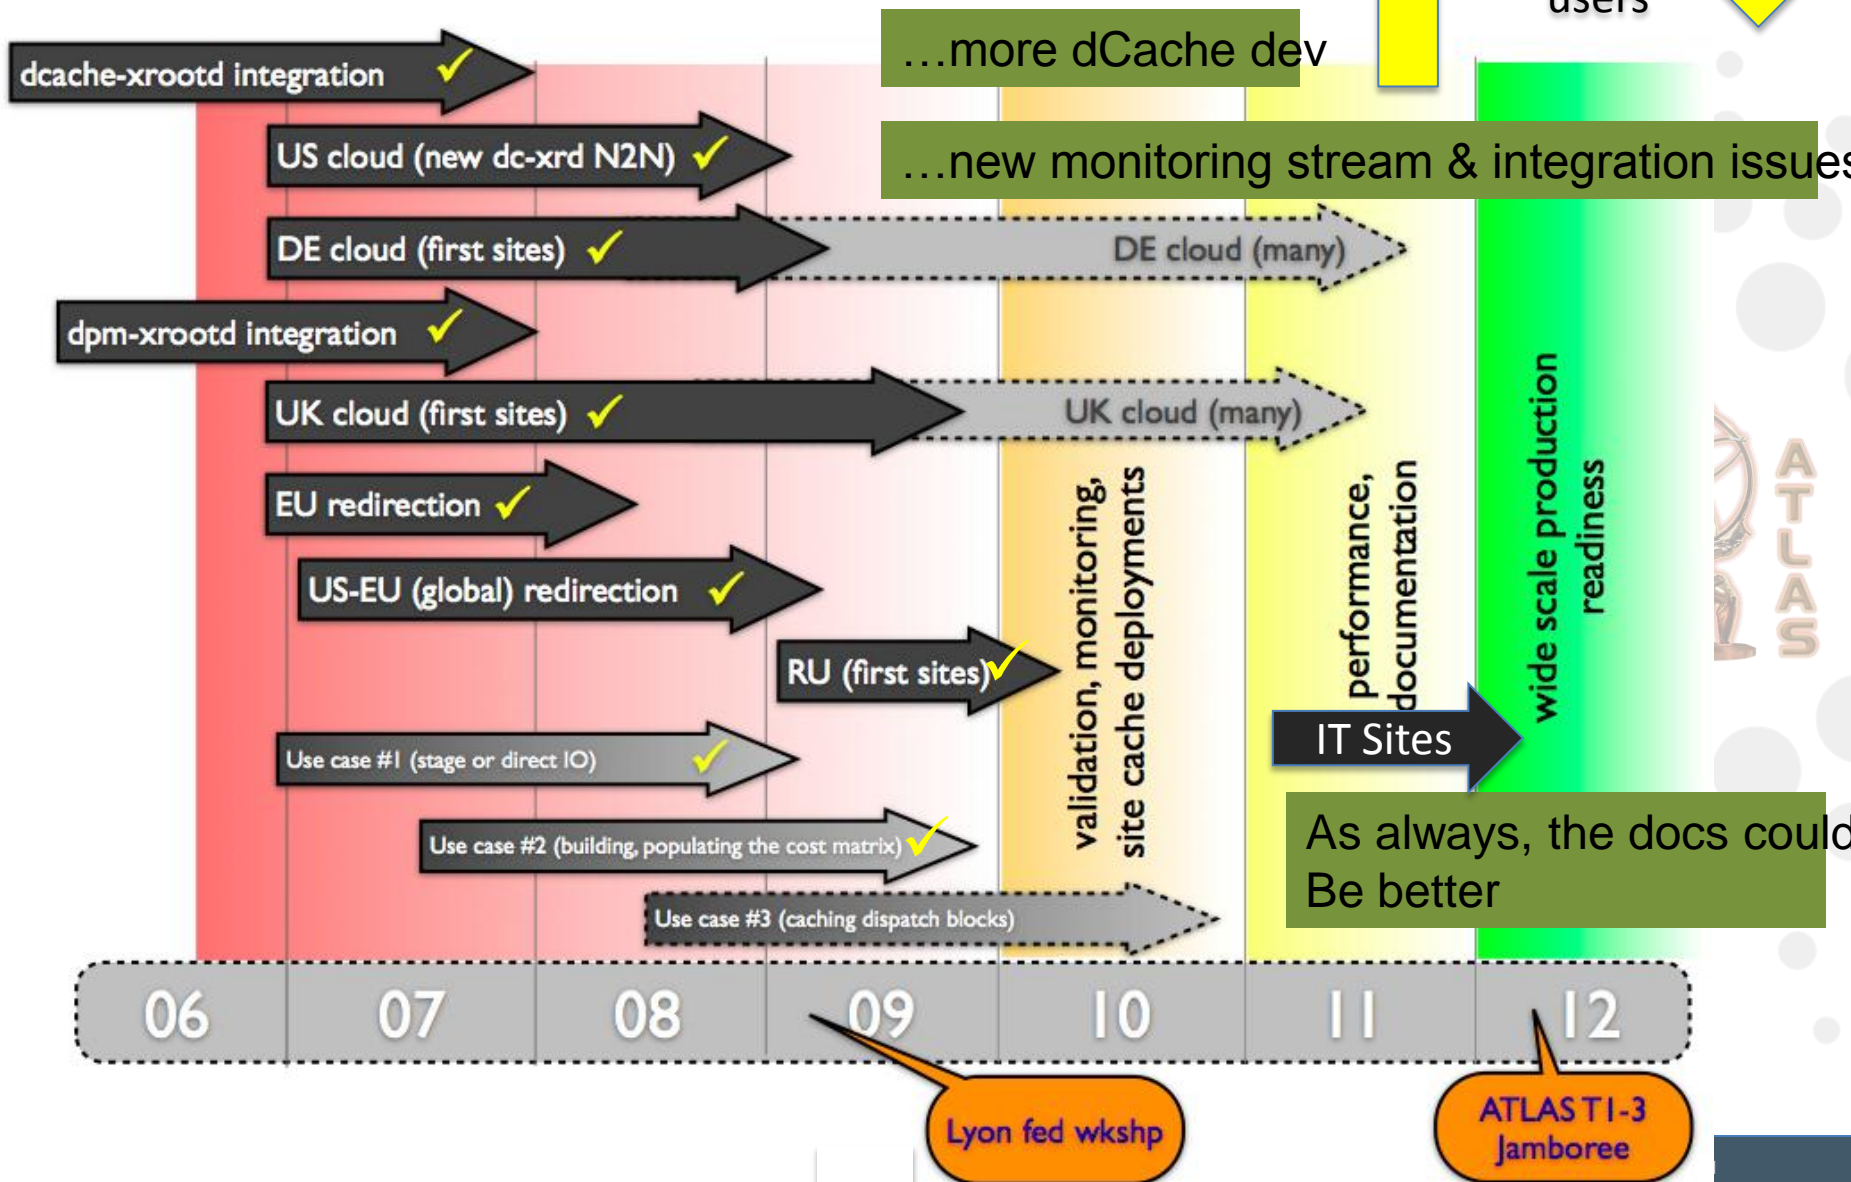




- Ongoing issues development & operations
- FAX in the pilot.. First use case working now
 - Later → Panda brokering
- Programmatic WAN testing
- FAX as supporting portable, lightweight applications on opportunistic resources (Tier 3, cloud, campus)



Roadmap



EU deployment status





RU cloud status.

• IHEP

- Site in production status.
- Current xroot version: 3.2.6-1
- No major problems observed
 - little concern because of the high memory consumption.

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
28123	xrootd	15	0	6213m	1.4g	2392	S	3.9	70.8	162:09.14	xrootd

- Automatic notification about test failures – are very useful. (Thanks to Ilija)

• JINR

- Site in production status.
- Current xroot version: 3.2.6-1
- No major problems observed, but memory...

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
5331	xrootd	20	0	3604m	1.5g	3204	S	4.0	38.1	2634:20	xrootd



UK Cloud FAX Status

- Same UK sites in FAX as at s/w week: ([Wiki page](#) of status)
 - ECDF; Glasgow; Oxford (T2Ds) (DPM) - Working
 - QMUL (T2D) (Lustre) – Working
 - RAL (T1) (Castor) – Works for CMS; “local” atlas path works
 - “Global” path doesn’t –issue contacting LFC?
- Nearly all other T2s in UK are DPM
 - Almost all have upgraded to EMI: DPM 1.8.4 (or later)
 - Allows easy Yaim config of FAX (confirmed this works)
 - Liverpool; Manchester; Lancaster willing to join
 - **but Xrootd 3.2.5 is broken in SL6**
 - David Smith confirmed this is fixed in 3.2.6 but also found a bug for SL5 – confirmed fixed now in devel.
 - Once 3.2.7 available from xrootd will be rebuilt for EPEL/EMI – probably released before Christmas... so these sites should be able to join next year





UK Testing and Other Issues (since last meeting)

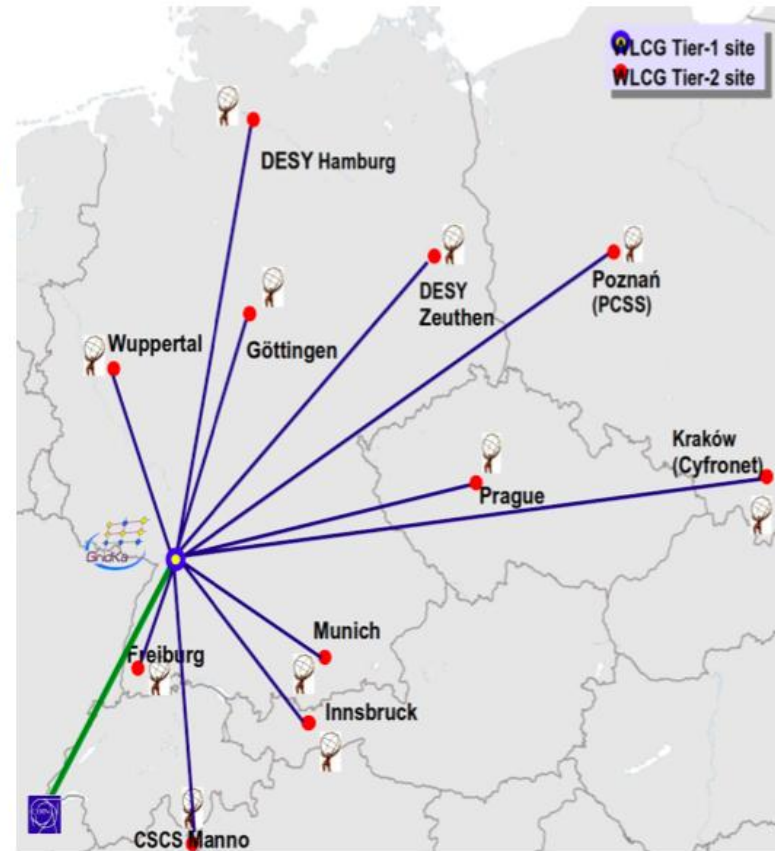
- All sites occasionally getting stuck in LFC lookup:
 - LFC host address changes, single threaded N2N sometimes (re)trying the “wrong” (cached) address.
 - **Fixed** by `setenv LFC_CONRETRY=0` (now in DPM YAIM config.)
- ANALY_GLASGOW_XROOTD queue
 - “local” access hammerclouds fine (as good as rfio at least)
 - For FAX test waiting to use new dq2-tools in cvmfs
 - Currently a python 2.6 issue
- **Real user (!)** at QMUL – all his files are local though
- **Ready for “dress-rehearsal”** mid-Jan with up-to 7 T2s and maybe T1
- Also planning WebDav deployment (tracked on same wiki)





DE cloud overview

- in DE:
 - 8 dCache sites (1 T1, 7 T2, 1 T3)
 - several T3 w/ Lustre/GPFS/Sonas
- outside DE:
 - DPM sites Prague, Cracow, Innsbruck
 - dCache in CSCS/Manno





DE FAX status

Sites:

- **LRZ-LMU:**
 - FAX node name (proxy): xrootd-lrz-lmu.grid.lrz.de:1094, tests ok
 - //atlas/dq2/user/Hironorilto/user.Hironorilto.xrootd.lrz-muenchen/user.Hironorilto.xrootd.lrz-muenchen-1M
- **MPPMU:**
 - FAX node name (proxy): xrootd-mppmu.grid.lrz.de:2094 , no tests
 - //atlas/dq2/user/Hironorilto/user.Hironorilto.xrootd.mpg/user.Hironorilto.xrootd.mpg-1M
- **Wuppertal**
 - FAX node name (proxy): up.pleiades.uni-wuppertal.de, local ok, redir fails
 - //atlas/dq2/user/Hironorilto/user.Hironorilto.xrootd.wuppertal/user.Hironorilto.xrootd.wuppertal-1M
 - needs reinstall of uptodate xrootd proxy ... in progress
- **Prague**
 - FAX node name (dpm) goliath100.farm.particle.cz, tests ok
 - //atlas/dq2/user/Hironorilto/user.Hironorilto.xrootd.prague/user.Hironorilto.xrootd.prague-1M
- **Desy-GridLab:** tests of dCache plugin, → Ilija





DE FAX prospects

- Documented FAX-Proxy setup for dCache in DE cloud Wiki:
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/FAXDECloud>
- We are waiting for more sites to join
 - we had expected GridKa to join
 - ongoing operational issues with tapes and storage have priority
 - currently very little manpower&resources available at sites.
 - some sites have strict policies how to operate services
 - require dedicated node for xrootd proxy installation, well connected to dcache nodes
 - not straightforward to fulfill
 - Plan B is deployment of dCache plugin for xrootd door (Ilija), needs no extra node
 - requires upgrade to recent dCache version
 - extended downtime & service interruption
 - rather wait until after spring confs in March 2013
- Dress rehearsal January:
 - Currently active sites should be ok, maybe GridKa as well.





- All the DPM (v1.8.4) sites in IT have FAX enabled since the last week of November 2012
 - INFN-FRASCATI
 - INFN-NAPOLI-ATLAS
 - INFN-ROMA1
- All FAX-enabled sites are pointing to the IT redirector at CERN
 - atlas-xrd-it.cern.ch
- Storm sites not ready yet
 - Test facility (XROOTD test machine in front of Storm) requested to CNAF
 - Still to be implemented and tested
 - Major Storm sites in Italy are INFN-T1 and INFN-MILANO-ATLASC



Alessandro de Salvo



■ Installation issues

- Not so much, the most complex task was to identify the list of ports to be opened on the routers!
- Update to the new N2N plugin transparent, no evident problem seen in DPM sites

■ Operational problems

- Some test dataset missing in Frascati and Napoli, needed for the transfer tests, should we ask manually for a dataset transfer in such cases?



■ Monitoring

- All the DPM sites correctly showing in the monitoring
- ...but not transferring so much data!
- We would like to have a real test to see if our infrastructure is really working properly



CERN IT
Department
CH-1211
Genève 23
Switzerland

DSS FAX: Federating EOS

CERN IT
Department

- Issue: (non-trivial) mapping of “global” name to EOS-local, need lookup in LFC - see [EOS Federation plans A.Peters \(IN2P3 workshop\)](#)
- N2N translation done via new service – see [LFC plug-in for EOS, E.Sindrilaru, A.Peters](#)
 - Currently will look into a single location only (`/eos/atlas/atlasdatadisk`), prototype/config issue
- Deployed on 2 VMs (`atlas-xrd-eos-n2n.cern.ch`) – works.

EOS N2N already looks into several locations -
`match=atlasdatadisk,atlasgroupdisk,atlasscratchdisk`

FAX phone conf – CERN redirector status - 1





CERN IT
Department
CH-1211
Genève 23
Switzerland

Other FAX infrastructure at CERN

CERN IT
Department

- Redirectors for EU, UK, DE, FR, RU (and ASIA)
- Common setup:
 - 2x VMs (on different hypervisors), same OS+xrootd versions
 - behind common DNS alias,

```
all.manager meta any atlas-xrd-eu.cern.ch.+ 1098  
xrootd.redirect atlas-xrd-eu.cern.ch:1094 ? /atlas
```

- Minor differences:
 - Site access control (to prevent mistakes, != security)
 - User access control (debug)
 - Topology – who subscribes where
- Minimal effort to run
 - Can keep or can give away to “real” regional redirectors
 - No real advantage to have them at CERN, some penalty for faraway sites



4



Development & integration .. operations





- SSB and WLCG transfer dashboard with cost matrix decision algorithm
- Xrootd instabilities seen in the UK cloud – perhaps related to N2N blocking at LFC
- FAX extensions to ATLAS information system AGIS
- Need new monitoring f-stream at all sites
- Stand-alone cmsd for dcache sites
- xrootd.org repository & EPEL policy (site guidance, esp. DPM)
- Several dCache specific issues, and many releases under test (1.9.12-22+, 2.2.4,...); f-stream, proper stat response and checksum support from dcache-xrootd doors
- Moved US sites to r.o. LFC
- SLC6 issues – X509 and VOMS attribute checking
- UDP collector service with f-stream format under test
- Functional testing probes & publishing into ActiveMQ and dashboards
- Monitoring will have to be validated at all stages
- FAX-enabled pilot site mover in production at two Tier 2s
- Documentation for users & site admins





Need a xrootd release with fixed f-stream (3.3.0-rc2)

Move all xrootd sites to new version

Test new xrootd monitoring settings

Have all xrootd servers auto-restart and auto report problems (request to Xrootd team)

A standalone version of cmsd to be used in tandem with dCache xrootd doors (request to Xrootd team)

Issue of xrootd repository: Xroot.org/CERN or EPEL (WLCG operations task force)





Version 2.2.4 in production at AGLT2

Version 1.9.12-23 in production at MWT2

xrootd4j plugin working

N2N plugin working

Monitor plugin working

Have other 1.9.12-X dCache sites need to upgrade to get monitoring



Have dCache properly respond to xrootd stat call (request to dCache team)

Implement kxr_query checksum request (request to dCache)



Italian DPM sites sensitive to version of xrdcp used to get file (ongoing investigation)

N2N for RAL (Castor) probably nothing fundamental but needs a bit of sustained effort

Shaun deWitt

Will setup a call to discuss (Wei, Ilija, Shaun)





German cloud

(more from Guenter's slides)

Wuppertal not working

DESY-HH works on a test node with dCache 2.4. Production on 1.9.12-12

Deployment in France? Where next?



EOS not redirecting properly

From Jan: Issue identified (yesterday), new version+ configuration not yet deployed but expect coming week



The problem on RHEL6 is fixed in 3.2.6 but the same fix also broke GSI security on RHEL5 (due to openssl). So 3.2.6 can't be deployed on RHEL5.

X509 on SLC6 is broken

VOMS attribute verification

Grid jobs running at some sites (apparently CERN and Glasgow at least) can't access FAX . Possible problem with limited proxy.






Test and certify f-stream

Finalize development of dCache Monitor plugin

Have the dashboard results "certified"

Things dashboard needs (proper host/user site name, VO, application)





Make RSV publish to ActiveMQ

Move FAX topology viewer to AGIS/AMQ



Add to WAN tests AGIS/AMQ



Improve on all parts of TWiki

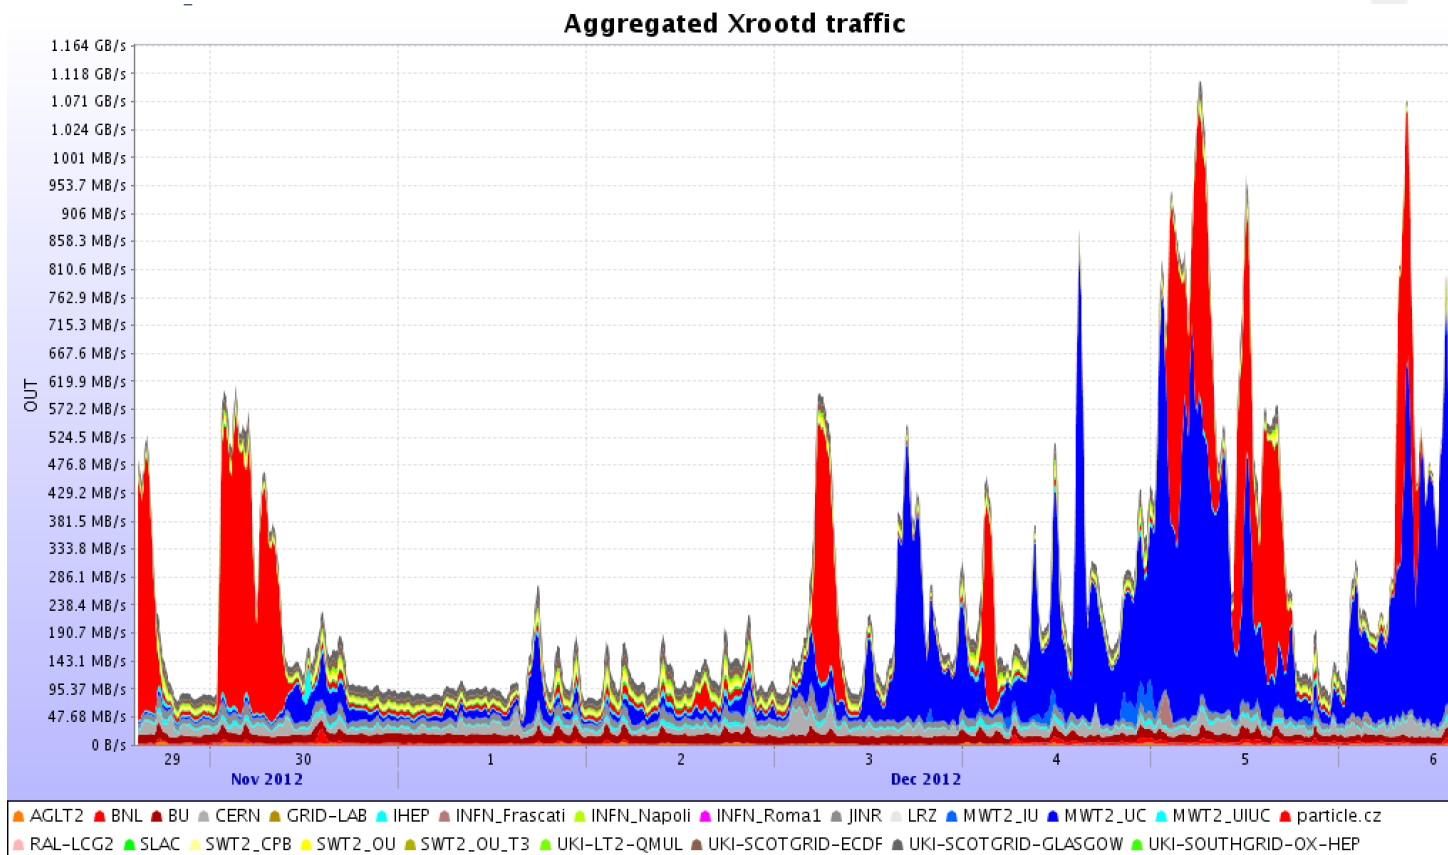
FAX site certification matrix update



Make our own PAT tutorial



- Growing but could be faster.
 - US ATLAS Hadronic Final State Forum 2012 users



Functional status & cost performance



Show 200 entries Copy Print Save view: Network

Site Name	Site Info		Network Measurements
	Source	Destination	
ANALY_MWT2_to_ANALY_CERN_XROOTD	MWT2	CERN-PROD	2.074
ANALY_SWT2_CPB_to_ANALY_CERN_XROOTD	SWT2_CPB	CERN-PROD	2.079
ANALY_OU_OCCHP_SWT2_to_ANALY_CERN_XROOTD	OU_OCCHP_SWT2	CERN-PROD	2.083
ANALY_CERN_XROOTD_to_ANALY_AGLT2	CERN-PROD	AGLT2	6.143
ANALY_MWT2_to_ANALY_NET2	MWT2	BU_ATLAS_Tier2	6.6
ANALY_CERN_XROOTD_to_ANALY_NET2	CERN-PROD	BU_ATLAS_Tier2	6.794
ANALY_AGLT2_to_ANALY_CERN_XROOTD	AGLT2	CERN-PROD	7.769
ANALY_NET2_to_ANALY_NET2	BU_ATLAS_Tier2	BU_ATLAS_Tier2	7.982
ANALY_AGLT2_to_ANALY_NET2	AGLT2	BU_ATLAS_Tier2	8
ANALY_SWT2_CPB_to_ANALY_NET2	SWT2_CPB	BU_ATLAS_Tier2	8
ANALY_OU_OCCHP_SWT2_to_ANALY_NET2	OU_OCCHP_SWT2	BU_ATLAS_Tier2	8.25
ANALY_CERN_XROOTD_to_ANALY_CERN_XROOTD	CERN-PROD	CERN-PROD	10.513
ANALY_OU_OCCHP_SWT2_to_ANALY_OU_OCCHP_SWT2	OU_OCCHP_SWT2	OU_OCCHP_SWT2	13.631
ANALY_SWT2_CPB_to_ANALY_SWT2_CPB	SWT2_CPB	SWT2_CPB	14.206

There are many more components as discussed at the Lyon storage federations workshop in September



ATLAS Federated Xrootd Status - 2012-10-08 07:15:51

Frequently Asked Questions

- ANALY_QMUL_to_ANALY_CERN_XROOTD
- ANALY_ILLINOISHEP_to_ANALY_AGLT2
- ANALY_ILLINOISHEP_to_ANALY_MWT2
- ANALY_ILLINOISHEP_to_ANALY_SWT2_CPB
- ANALY_wuppertalprod_to_ANALY_CERN_XROO
- ANALY_QMUL_to_ANALY_NET2
- ANALY_ILLINOISHEP_to_ANALY_NET2
- ANALY_QMUL_to_ANALY_AGLT2
- ANALY_wuppertalprod_to_ANALY_AGLT2
- ANALY_wuppertalprod_to_ANALY_SWT2_CPB
- ANALY_wuppertalprod_to_ANALY_MWT2
- ANALY_ILLINOISHEP_to_ANALY_CERN_XROO
- ANALY_wuppertalprod_to_ANALY_NET2

Host: atl-prod09.slac.stanford.edu (atl-prod09.slac.stanford.edu)

Metric	Last Executed	Enabled?	Next Run Time	Status
org.usatlas.xrootd.grid.xrscp.compare	2012-10-08 07:05:00 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.direct	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.fax	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.ping	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK

Host: atlas-cm4.bu.edu (atlas-cm4.bu.edu)

Metric	Last Executed	Enabled?	Next Run Time	Status
org.usatlas.xrootd.grid.xrscp.compare	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.direct	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.fax	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.ping	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK

FAX dashboard – sites transfer matrix



FAX Monitoring

TRANSFER MATRIX (2012-10-25 00:00 to 2012-11-08 00:00 UTC SLIDING)



▼ Summary

Interval
Last 14 days

Access type
Local access
Remote access

Transfer VS Reading
Reading
Copy

Servers

Sources
Countries:
Sites:
Host:
Grouping: COUNTRY

Destinations
Countries:
Sites:
Host:
Grouping: COUNTRY SITE

▶ Interval

▶ Access type

▶ Transfer VS Reading

▶ Servers

▶ Sources

Matrix | Transfer Plots | Access Plots | Site Statistics | Custom Ranking Plots

Efficiency
 Throughput
 Successes
 Errors

0 % 100 %

100 %

0 %

SOURCES

DESTINATIONS

	TOTAL-	Germany+	Russian-Federation+	UK+	USA+
TOTAL-	98 %	98 %	100 %	99 %	94 %
Germany LRZ-LMU+	86 %	86 %			
Russian-Federation JINR-LCG2+	96 %		96 %		
Russian-Federation RU-Protvino-IHEP+	84 %		84 %		
Switzerland CERN-CMSTEST+	100 %	100 %	100 %	100 %	100 %
UK UKI-LT2-QMUL+	100 %			100 %	
UK UKI-SCOTGRID-ECDF+	100 %			100 %	
UK UKI-SCOTGRID-GLASGOW+	76 %		100 %	74 %	100 %
USA AGLT2+	100 %		100 %		100 %
USA BNL-ATLAS+	55 %				55 %
USA BU_ATLAS_Tier2+	100 %		100 %	100 %	100 %
USA IllinoisHEP+	11 %				11 %
USA MWT2+	100 %		100 %	100 %	100 %
USA MWT2_ATLAS_IU+	96 %	100 %	100 %	100 %	95 %
USA OU_OCHEP_SWT2+	100 %		100 %	100 %	100 %
USA SWT2_CPB+	100 %		100 %	100 %	100 %
USA WT2+	98 %			100 %	98 %





- Alessandro, Artem, Danila and Alexey added FAX support to AGIS
 - Database side
 - Web UI
 - Export to JSON
 - Tutorial (on the way)
- Current FAX topology data inserted (Ilija)
- Need to decide on who will take care data are up-to-date.
 - Since it should not change often we could charge one person with doing it
 - But it is rather simple to maybe sites could do it too
- Keep in mind: currently no e-mail field associated to FAX SE or redirector in AGIS
 - Currently FAX topology monitor has handpicked people to send mail to when FAX SE or redirector is down
 - Move to AGIS will mean that people responsible for site will be getting mails



- We have an on-going set of meetings with two WLCG working groups
 - WLCG Federated Data Storage (F. Furano Chair)
 - Explores issues generally and assesses approach taken by each of the LHC experiments
 - WLCG Xrootd task force (D. Giordano)
 - New group forming to coordinate deployment and operations
 - Will seek to engage Grid infrastructure providing groups (EGI/EMI, OSG) for support
- Both of these will help bring FAX into normal production operations



In the pilot, in Panda





FAX in the Pilot

- Available since pilot version 54a (August 16)
- Update for user analysis jobs in v 54b (September 26)
- Tested “manually” (i.e. with faked stage-in failures) using dev pilots and test jobs on US test sites; UTA_PAUL_TEST, ANALY_MWT2_TEST
- Now used with user jobs on ANALY_MWT2
- Currently running tests at ANALY_SWT2_CPB and ANALY_AGLT2
- More tests on more sites under way
- Schedconfig activation: allowfax = 'true', faxredirector = '..'



FAX Activation Algorithm

- When the pilot encounters a stage-in error, it will retry the transfer (after a 5 min delay).
- If the second stage-in attempt fails as well, switch to FAX
- A special FAX site mover has been developed (can also be used as primary site mover, tested – works fine)
- If FAX fails, report original failure reason, the entire job will be retried
- Requires CVMFS site with an updated DQ2 client (2.3.0)
 - Performs a local root setup, uses xrd-cp from cvmfs, needs new option `-p` for dq2-list-files



“Diskless” Use-cases

- FAX mechanism offers some interesting use-cases
- A stage-in problem will trigger the FAX algorithm as mentioned in the previous slide, bypassing the local SE for stage-in
- Planned SE shutdown; FAX as primary copy tool
 - Set `schedconfig.copytoolin = "fax"` (switch off any direct access)
 - But what about stage-out?
 - Stage-out to alternative SE (from T-2 to T-1) in development
 - Similar mechanism; after repeated stage-out attempts, the pilot switches to attempting stage-out to an alternative SE (belonging to the SE of the corresponding T-1 [within the same cloud as the job] using `lcg-cp`)
 - Local SE bypassed



- Sites with FAX enabled: ANALY_MWT2, ANALY_SWT2_CPB
- Currently testing FAX at: ANALY_OX, ANALY_GLASGOW_XROOTD, ANALY_CERN_XROOTD, ANALY_LANCS, UKI-LT2-QMUL, ANALY_ECDF and ANALY_LRZ using updated pilot wrapper
- Pilot factory wrappers need to setup python 2.6 + latest DQ2 client (2.3.0)
 - Done for pyfactory wrapper by Peter Love
- Introduced FAX usage info to job metrics
 - If at least one successful FAX transfer, pilot adds “filesWithFAX= N_1 filesWithoutFAX= N_2 ” to job metrics





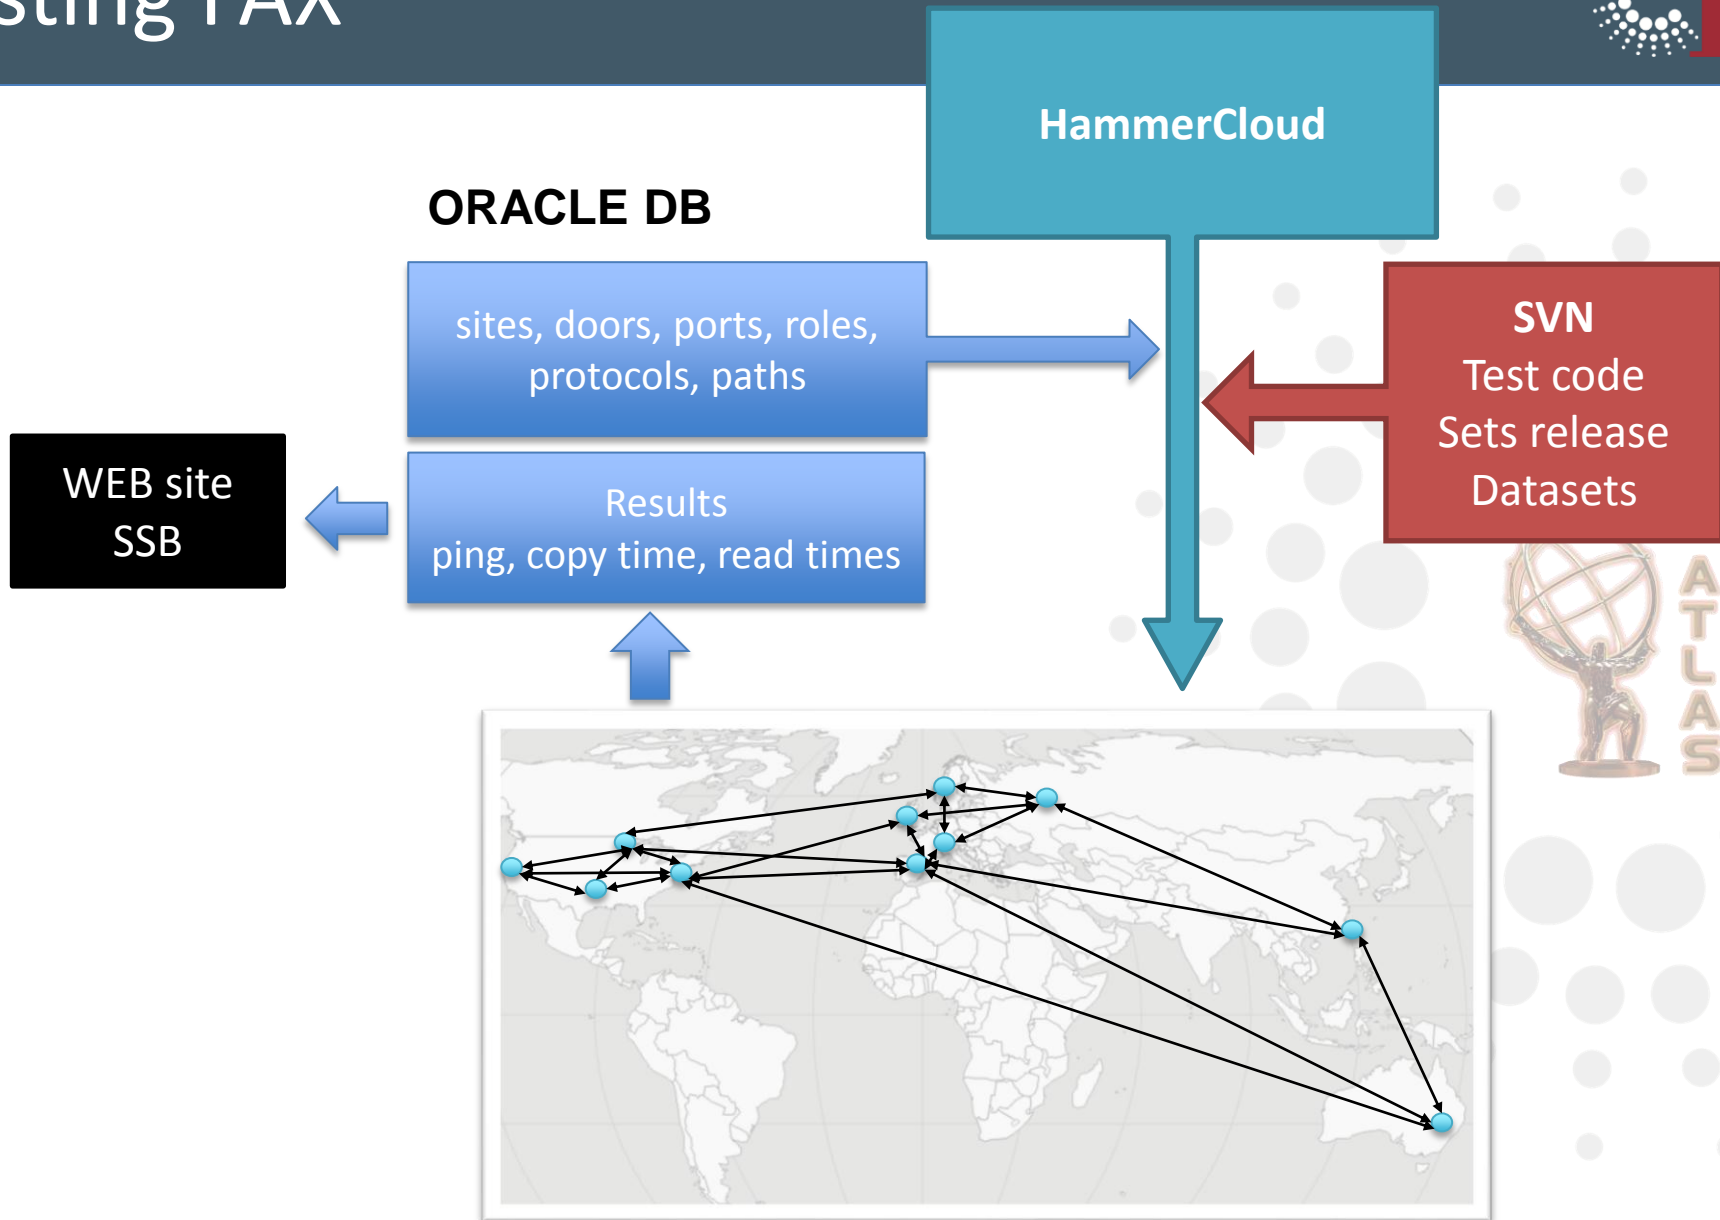
- Problems with lcg-cp seen on some sites (ANALY_GLASGOW, ANALY_LANCS, UKI-LT2-QMUL):
 - lcg-cp: error while loading shared libraries:
libcgsi_plugin.so.1: cannot open shared object file: No such file or directory
 - Bad PYTHONPATH
 - lcg-cp works fine e.g. at RAL with the new wrapper; however RAL is not yet part of the federation
- Intermittent transfer problems at ANALY_CERN_XROOTD
 - Reads for a while and then gets interrupted
[xrootd] Total 1320.18 MB |=====>..... | 52.72 % [1.1 MB/s]
Last server error 3010 ('Unable to open /atlas/dq2/./AOD.638838._000004.pool.root.1; Permission denied')
- Server error 10000 at ANALY_OX
 - Transfers (all) fail with errors like
Last server error 10000 (") Error accessing path/file for root://atlas-xrd-uk.cern.ch//atlas/dq2/mc11_7TeV/AOD/e971_s1310_s1300_r3043_r2993/mc11_7TeV.127627.simple_Gt_G750_T1200_L150_herwigpp_susy.merge.AOD.e971_s1310_s1300_r3043_r2993_tid637585_00/AOD.637585._000001.pool.root.2



WAN performance



Testing FAX





- HC submits 1 job/day to all of the “client” nodes. Client node is the one using the data. It is an ANALY queue
 - All the “server” sites have one and the same dataset. Server sites are the ones delivering data.
- Each job, each half an hour, in parallel:
 - Pings of all of the “server” sites.
 - Copies a file from a site (xrdcp/dccp)
 - Reads the file from a root script
 - Uploads all the results to Oracle DB at CERN
- Result are shown at:
<http://ivukotic.web.cern.ch/ivukotic/WAN/index.asp>
- Results are also given in JSON format to SSB: <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#currentView=Network+Measurements&highlight=false>




- Test file
 - standard ATLAS 760 MB D3PD with 5k branches and 13kevents
- Measurements
 - “direct copy”: time (in seconds) for xrdcp to site
 - “read time” time required to read 10% randomly selected consecutive events using default TTreeCache of 30 MB




For jobs at MWT2 (client location)



ping [ms] for 10% default cache



read time for 10% default cache



by Ilija Vukotic

FAX 'Dress Rehearsal' ("FDR")



- Steps towards usability:
 - Define an FDR sufficient to cover most anticipated user workloads, including user docs
 - Define a rehearsal period: ~ week and metrics (and needed monitoring)
 - Poll for site volunteers and an ad-hoc FAX OPS team
 - Execute the FDR; gather monitoring statistics and accounting data
- Propose spend December preparing
 - Identifying five exemplar use cases that can be run by FAX OPS team
 - Preparing a clean set of tutorial-like documents
 - Preplace example datasets
 - Load test redirectors and sites against with examples
 - Solve the main problem of federated access to datasets
- Week of January 21 going live with real users





- FAX usability inches forward – but growing pains due to:
 - Standardizing metrics
 - dCache components
 - Many more sites and SEs
- First pilots bits are in production at a couple sites
- Co-located Tier3 users using FAX doors for LOCALGROUPDISK analysis
- Offers storage access from opportunistic or cloud
- Offers “diskless” use case which would be very attractive to sites for storage admin purposes
- Dress rehearsal being planned for week of January 21



Extra: FAX usability



How to use it? Part - I



- Datasets should be registered
 - All the grid produced datasets are automatically registered independently if these are part of official production or simply result of a user's job.
 - If files are not registered it is trivial to do so. Very detailed description how to do this is given <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/DQ2ClientsHowTo>.

- Have your ATLAS grid certificate
 - Make a proxy

```
source /afs/cern.ch/project/gd/LCG-share/current_3.2/etc/profile.d/grid_env.sh
voms-proxy-init -voms atlas
```

- setup DQ2

```
source /afs/cern.ch/atlas/offline/external/GRID/ddm/DQ2Clients/setup.zsh
```

- Make sure your code uses TTreeCache!

CVMFS version 





- Setup environment

```
export
ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase
alias setupATLAS='source
${ATLAS_LOCAL_ROOT_BASE}/user/atlasLocalSetup.sh'
export ALRB_localConfigDir=$HOME/localConfig
setupATLAS
localSetupGLite
```

- Make a proxy

```
voms-proxy-init -voms atlas
```

- setup DQ2

```
localSetupDQ2Client
```





- Check that datasets exist at one of the federated sites

```
dq2-ls -r myDataSetName
```

- Find gLFN's of input datasets

- Find closest redirector to compute site. List is here:

<https://twiki.cern.ch/twiki/bin/view/Atlas/FaxRedirectors>

- Do



```
export STORAGEPREFIX=root://closestRedirector:port/
```

- make a file with the list of all the gLFN's

```
dq2-list-files -p  
data12_8TeV.00201556.physics_Muons.recon.DESD_ZMUMU.f437_m716_f437 >  
my_list_of_gLFNS.txt
```



- From ROOT

```
TFile *f =  
TFile::Open("root://myRedirector:port//atlas/dq2/user/ilijav/HCTest/user.i  
lijav.HCTest.1/group.test.hc.NTUP_SMWZ.root");
```

- From prun

Instead of giving --inDS myDataset option, provide it with --pfnList my_list_of_gLFNS.txt



- copy files locally

```
xrdcp  
root://xrddc.mwt2.org:1096//atlas/dq2/user/ilijav/HCTest/user.ilijav.HCtes  
t.1/group.test.hc.NTUP_SMWZ.root /tmp/myLocalCopy.root
```