

ATLAS Experience with SRM2

Birger Koblitz, CERN-IT

for the ATLAS DDM team

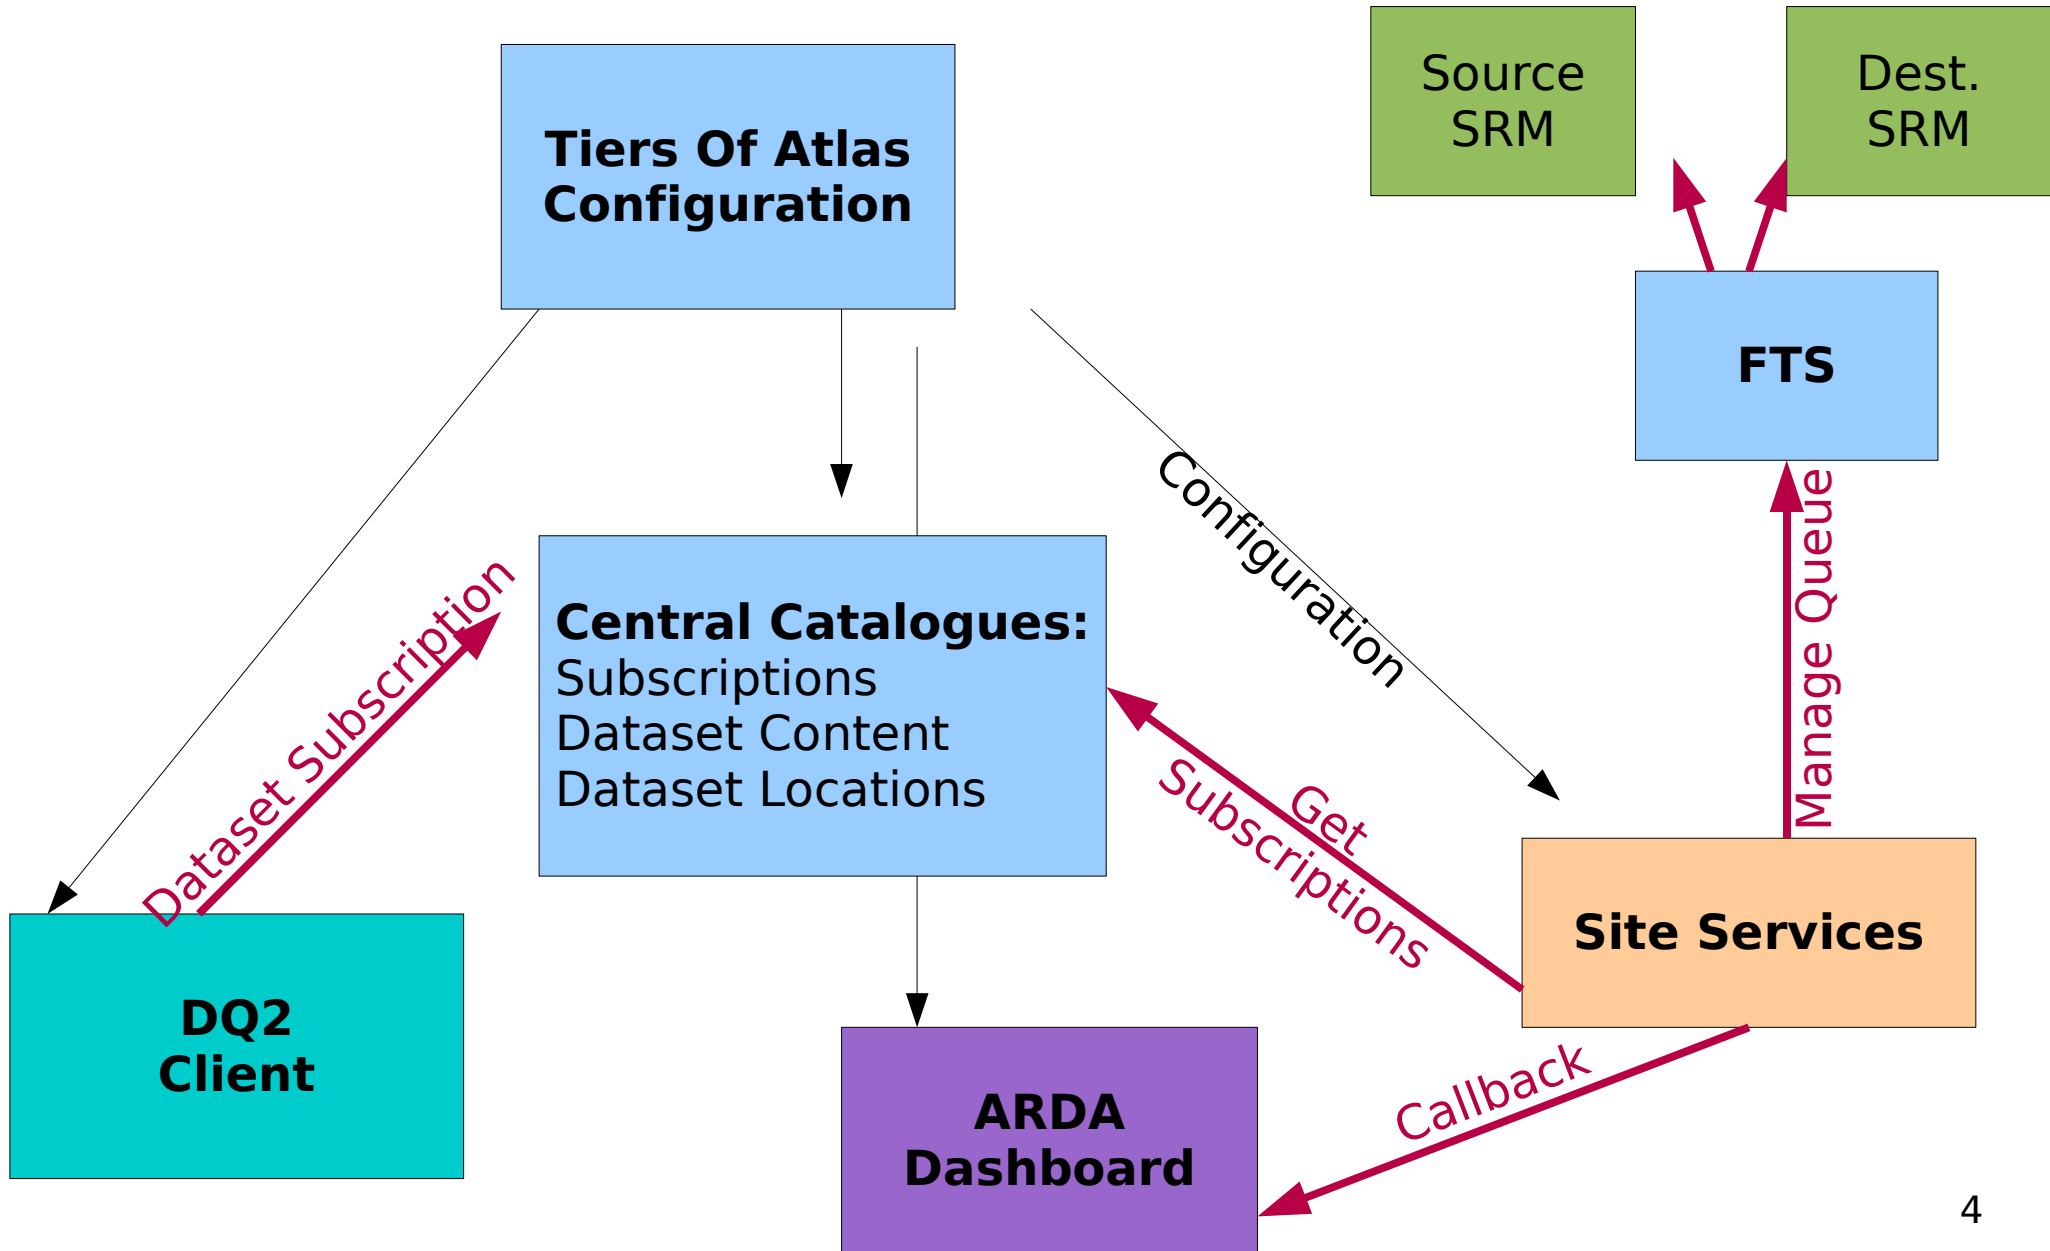
Overview

- Objectives
- SRM2 test setup: ATLAS-DDM
- Results for many small files at NDGF, CNAF
- Results with large T0-export files
- First impression of the srm-pin/stage tools
- ATLAS requirements on SRM tools

Objectives

- Test new SRM2 implementations with the existing ATLAS DDM system
 - Do usability, not functional tests!
 - Understand what we have to change in DDM
 - Allows using the standard queues on FTS, standard Certificates
 - Understand how fast Atlas DDM can be adjusted to work with Test-Sites
- Understand new features:
 - Pinning/Staging
 - Space Tokens

Test setup with ATLAS DDM



Tests with many Small Files

- Sent 70k small files of 100kB size to sites
- Small files should stress the SRM most...
- Picked first sites that were read: NDGF & CNAF
- No modifications in DDM:
 - No Space Tokens used
 - Not making use of e.g. Pre-staging functionality
- Writing tested only...

OVERVIEW

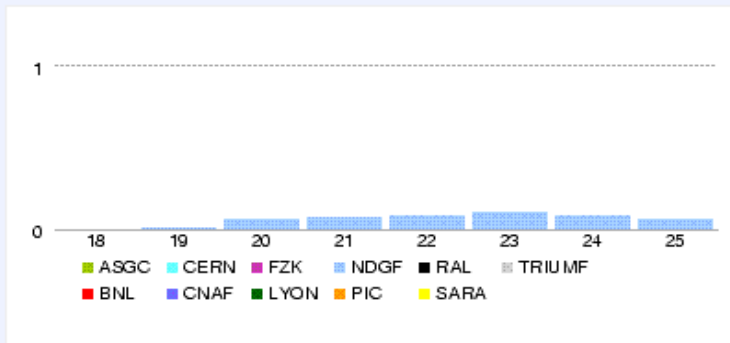
OVERVIEW Activity

- [Activity in Last Hour](#)
- [Activity in Last 4 Hours](#)
- [Activity in Last 24 Hours](#)
- [Activity in Last 7 Days](#)
- [Activity in Last 30 Days](#)
- [Activity in ...](#)

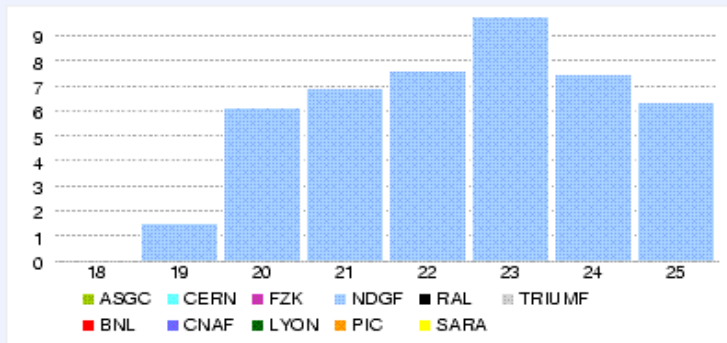
Cloud Activity

- [ASGC Cloud](#)
- [BNL Cloud](#)
- [CERN Cloud](#)
- [CNAF Cloud](#)
- [FZK Cloud](#)
- [LYON Cloud](#)
- [NDGF Cloud](#)
- [PIC Cloud](#)
- [RAL Cloud](#)
- [SARA Cloud](#)
- [TRIUMF Cloud](#)

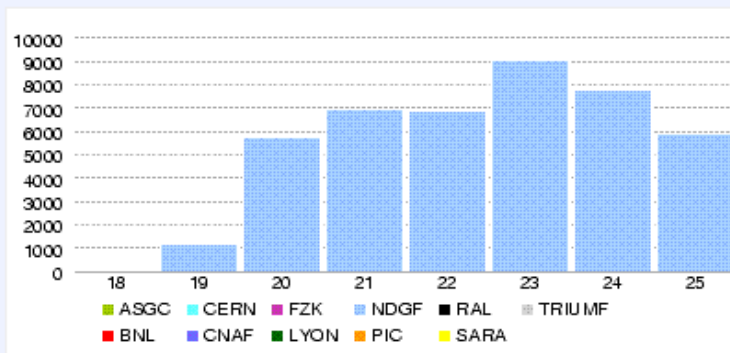
Throughput (MB/s)



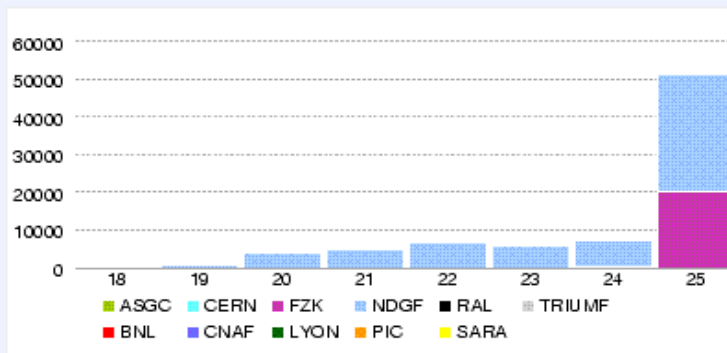
Data Transferred (GBytes)



Completed File Transfers



Total Number Errors



Activity Summary (Last 168 Hours)

Click on the cloud name to view list of sites

Cloud	Transfers			Services		Errors				
	Efficiency	Throughput	Files Done	Datasets Done	DQ	Grid	Transfer	Local	Remote	Central
ASGC	0%	0 MB/s	0	0			0	0		
BNL	0%	0 MB/s	0	0			0	0		
CERN	0%	0 MB/s	0	0			0	0		
CNAF	0%	0 MB/s	0	0			0	0		
FZK	0%	0 MB/s	0	20			874502	0		
LYON	0%	0 MB/s	0	0			0	0		
NDGF	75%	0 MB/s	41220	0			13562	88406		
PIC	0%	0 MB/s	0	0			0	0		
RAL	0%	0 MB/s	0	0			0	0		
SARA	0%	0 MB/s	0	0			0	0		
TRIUMF	0%	0 MB/s	0	0			0	0		

All Errors from CERN Castor...

NDGFTST	75%	0 MB/s	41220	0	OK	13562	88406		
State from FTS: Failed; Retries: 3; Reason: SOURCE error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare source file in 60 seconds									8168
State from FTS: Failed; Retries: 3; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStageInterface.c:2162 Required tape segments are not all accessible (errno=0, serrno=0)									4204
State from FTS: Failed; Retries: 3; Reason: TRANSFER error during TRANSFER phase: [TRANSFER_TIMEOUT] gridftp_copy_wait: Connection timed out									509
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [INVALID_PATH] CastorStageInterface.c:2162 No such file or directory (errno=0, serrno=0)									184
State from FTS: Failed; Retries: 2; Reason: Failed to get proxy certificate from myproxy-fts.cern.ch . Reason is Error in bind()									147
State from FTS: Failed; Retries: 3; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] cannot continue since no size has been returned after PrepareToGet or SrmStat									90
State from FTS: Failed; Retries: 3; Reason: SOURCE error during TRANSFER phase: [GRIDFTP] a system call failed (Connection refused)									71
State from FTS: Failed; Retries: 3; Reason: SOURCE error during PREPARATION phase: [INVALID_PATH] specified file(s) does not exist									39
State from FTS: Failed; Retries: 3; Reason: SOURCE error during TRANSFER phase: [GRIDFTP] the server sent an error response: 550 550 lcg00118.grid.sinica.edu.tw:/flatfiles/SE00/at/stage/alpgen.008162.AlpgenJimmyZtautauNp2VBFCut._00504.tar.gz.104633: not a plain file.									21
State from FTS: Failed; Retries: 3; Reason: SOURCE error during TRANSFER phase: [TRANSFER_TIMEOUT] globus_ftp_client_size: Connection timed out									18
State from FTS: Failed; Retries: 3; Reason: SOURCE error during TRANSFER phase: [GRIDFTP] the server sent an									

CNAF-StoRM works well, too...

OVERVIEW


OVERVIEW Activity

- [Activity in Last Hour](#)
- [Activity in Last 4 Hours](#)
- [Activity in Last 24 Hours](#)
- [Activity in Last 7 Days](#)
- [Activity in Last 30 Days](#)
- [Activity in ...](#)

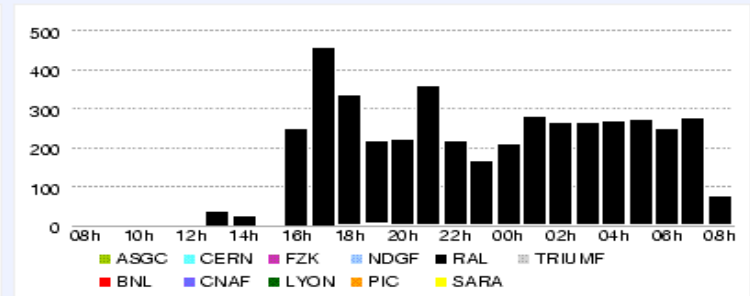
Cloud Activity

- [ASGC Cloud](#)
- [BNL Cloud](#)
- [CERN Cloud](#)
- [CNAF Cloud](#)
- [FZK Cloud](#)
- [LYON Cloud](#)
- [NDGF Cloud](#)
- [PIC Cloud](#)
- [RAL Cloud](#)
- [SARA Cloud](#)
- [TRIUMF Cloud](#)

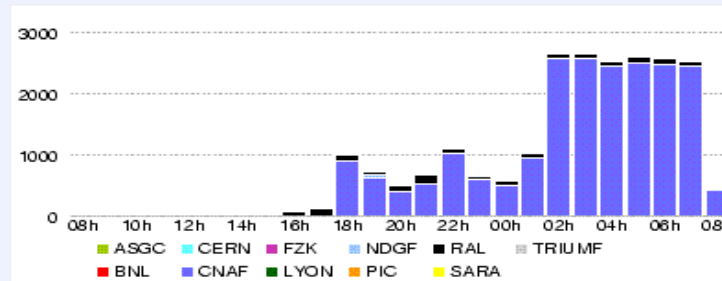
Throughput (MB/s)



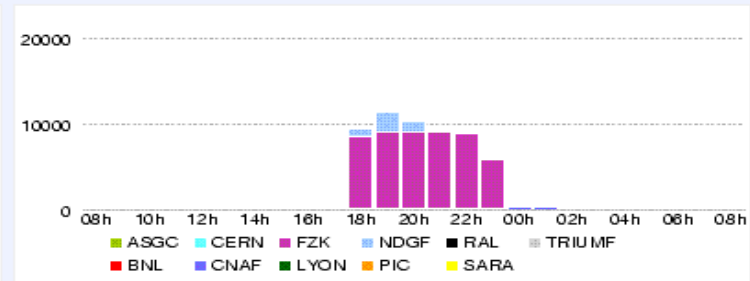
Data Transferred (GBytes)



Completed File Transfers



Total Number Errors



Activity Summary (Last 24 Hours)

Click on the cloud name to view list of sites

Cloud	Transfers			Services		Errors				
	Efficiency	Throughput	Files Done	Datasets Done	DQ	Grid	Transfer	Local	Remote	Central
ASGC	0%	0 MB/s	0	0			0	0		
BNL	0%	0 MB/s	0	0			0	0		
CERN	0%	0 MB/s	0	0			0	0		
CNAF	93%	0 MB/s	22587	0			1667	0		
FZK	0%	0 MB/s	0	0			49707	0		
LYON	0%	0 MB/s	0	0			0	0		
NDGF	1%	0 MB/s	45	0			5055	0		
PIC	0%	0 MB/s	0	0			0	0		
RAL	84%	54 MB/s	1489	840			290	0		
SARA	0%	0 MB/s	0	0			0	0		
TRIUMF	0%	0 MB/s	0	0			0	0		

Errors again from Castor...

CNAF	93%	0 MB/s	22587	0		1667	0		
Click on the site name to go to the site page									
CNAFDISK	0%	0 MB/s	0	0	OK	0	0		
CNAFTAPE	0%	0 MB/s	0	0	OK	0	0		
CNAFTEST	93%	0 MB/s	22587	0	OK	1667	0		
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStagerInterface.c:2162 Required tape segments are not all accessible (errno=0, serrno=0)									1266
State from FTS: Failed; Retries: 1; Reason: DESTINATION error during PREPARATION phase: [FILE_EXISTS] Cannot srmPut file because it already exists!									235
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [INVALID_PATH] CastorStage rInterface.c:2162 No such file or directory (errno=0, serrno=0)									81
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare source file in 180 seconds									32
State from FTS: Failed; Retries: 1; Reason: DESTINATION error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare Destination file in 180 seconds									25
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [INVALID_PATH] specified file(s) does not exist									15
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] cannot continue since no size has been returned after PrepareToGet or SrmStat									7
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStagerInterface.c:2145 Internal error (errno=0, serrno=1015)									3
State from FTS: Failed; Retries: 1; Reason: SOURCE error during TRANSFER phase: [GRIDFTP] a system call failed (Connection refused)									1
State from FTS: Failed; Retries: 1; Reason: DESTINATION error during PREPARATION phase: [INVALID_PATH] Invalid path									1
State from FTS: Failed; Retries: 1; Reason: SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStagerInterface.c:2162 Is a directory (errno=0, serrno=0)									1

Conclusions from First Phase

- SRM2 works very well as a drop-in replacement for SRM1, at least for the dCache and StoRM implementations
- Pre-staging should be able to help a lot in efficiency
- Unfortunately: Files in SRM-Test instances also picked up for transfers to other sites
 - Reading tested in uncontrolled way, too!
- Next phase: Need to integrate better with ongoing tests...

Writing large files

- Now integrated with ongoing T0 exercise:
 - Export tests from CERN with large files to T1 sites
 - Tests mainly throughput: 3GB files
- Add new site: LYON
 - BNL had trouble with dCache
 - FZK not in production BDII
 - LAL not considered (T2)

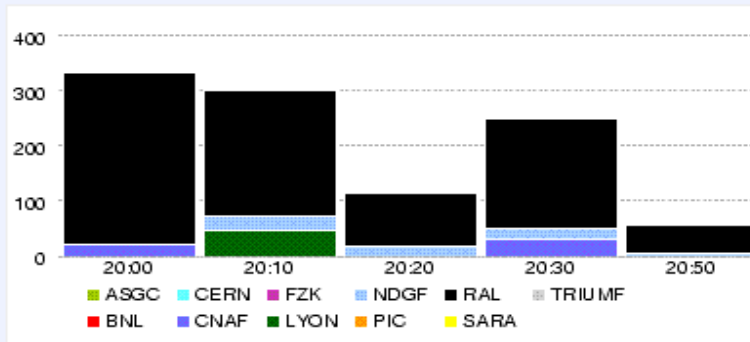
OVERVIEW Activity

- Activity in Last Hour
- Activity in Last 4 Hours
- Activity in Last 24 Hours
- Activity in Last 7 Days
- Activity in Last 30 Days
- Activity in ...

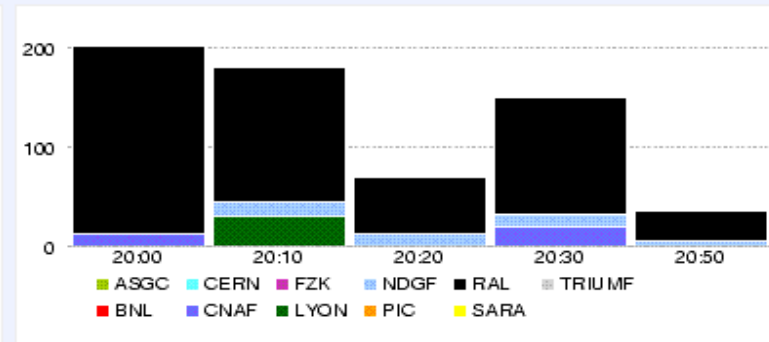
Cloud Activity

- ASGC Cloud
- BNL Cloud
- CERN Cloud
- CNAF Cloud
- FZK Cloud
- LYON Cloud
- NDGF Cloud
- PIC Cloud
- RAL Cloud
- SARA Cloud
- TRIUMF Cloud

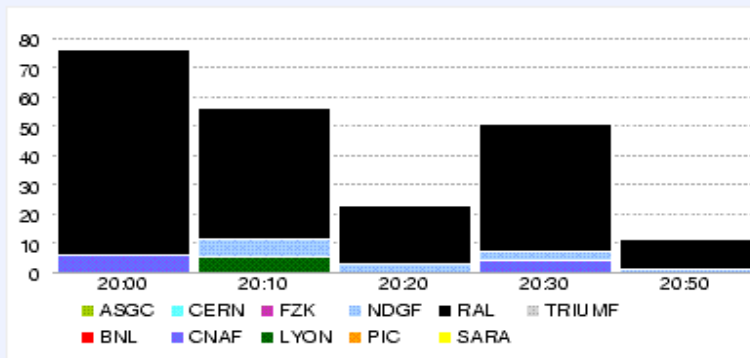
Throughput (MB/s)




Data Transferred (GBytes)



Completed File Transfers



Total Number Errors



Activity Summary (Last 1 Hours)

Click on the cloud name to view list of sites

Cloud	Transfers				Services		Errors			
	Efficiency	Throughput	Files Done	Datasets Done	DQ	Grid	Transfer	Local	Remote	Central
ASGC	0%	0 MB/s	0	0			0	0		
BNL	0%	0 MB/s	0	0			0	0		
CERN	0%	0 MB/s	0	0			0	0		
CNAF	12%	9 MB/s	10	2			76	0		
FZK	0%	0 MB/s	0	0			0	0		
LYON	56%	8 MB/s	5	1			4	0		
NDGF	100%	11 MB/s	13	3			0	0		
PIC	0%	0 MB/s	0	0			0	0		
RAL	51%	147 MB/s	189	55			184	0		
SARA	0%	0 MB/s	0	0			0	0		
TRIUMF	0%	0 MB/s	0	0			0	0		

StoRM

- CNAF problems result from StoRM dynamic space allocation feature:

CNAF	0%	0 MB/s	0	0			700	0		
Click on the site name to go to the site page										
CNAFDISK	0%	0 MB/s	0	0	OK		0	0		
CNAFTAPE	0%	0 MB/s	0	0	OK		0	0		
CNAFTEST	0%	0 MB/s	0	0	OK		700	0		
State from FTS: Failed; Retries: 1; Reason: DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] Space Management step in srmPrepareToPut failed!										697
State from FTS: Failed; Retries: 1; Reason: TRANSFER error during TRANSFER phase: [GRIDFTP] an end-of-file was reached (possibly the destination disk is full)										3
LNF	0%	0 MB/s	0	0	OK		0	0		

- Currently under investigation...

SRM-pin tool preview

- Understood that we will use BringOnline
- But to see whether this is useful for us, we need to change our data management software
 - Could be interesting to do pre-staging before scheduling FTS transfer

Conclusions

- SRM2 implementations can be used as drop-in replacements for SRM1 instances
- Problems observed seem to be due to backend configuration or CASTOR at CERN
- StoRM dynamic space allocation under investigation
- Atlas DDM will now regularly use the SRM2 endpoints for its export tests
- Need to include T2s (test DPM)

Future Plans

- Need to include T2s (test SRM2 of DPM)
- Test SRM2 at CERN
 - Read from SRM2
 - Test usage of Space Tokens
(needs updates to DDM, first)
- Need to test new SRM2 features
 - Bring-Online can allow pre-staging
 - srm-ls?

ATLAS Requirements

- lcg-xx tools:
 - Used intensively by Production System, Analysis
 - Need to work with SRM2
 - Support for Space Tokens
- gfal:
 - Used on Worker Node via Root to access files
 - Sometimes used in Ganga instead of lcg-tools on WN
 - Updates needed to work with SRM2
 - Python wrapper needed

Tools Needed for Direct SRM Access

- Need direct access to SRM2 on CLI:
 - High interest in Distributed Data Management
 - srm-ls: needed to check consistency with LFC
 - srm-stage: intended to pre-stage files before transfer with DDM
- Tools should be scriptable (i.e. no Java)
- Tools should not look into BDII
- Tools should be OSG interoperable

SRM access library

- Need direct SRM access library in Python
 - No detour via BDII, please
 - Should be thread-safe
 - Should work against multiple catalogues concurrently
 - Either swigged C or Python client generated from SRM 2.2 WSDL
 - OSG interoperability important

Additionally...

- For consistency of catalogues on Sites:
 - List LFC content by SURL date, needed to clean up catalogues