



Status of LHCb SRM2.2 Testing

Andrew C. Smith



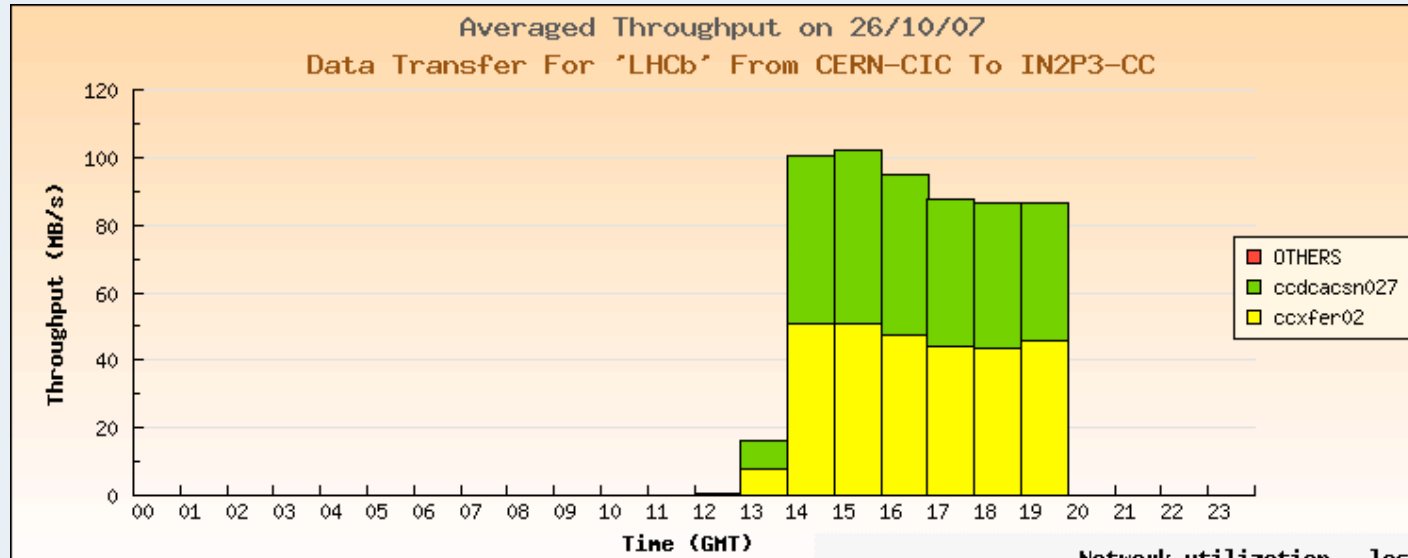
- Pre-staged 12k file on CERN Castor
 - ~2.2TB on dedicated disk pools
 - Ensure file availability

- FTS Channels configured on production FTS(2)
 - Source
 - srm.cern.ch production SRM1 endpoint
 - Destinations
 - ccsrmtstv2.in2p3.fr (files: 50, streams: 5)
 - ant2.grid.sara.nl (files: 40, streams: 5)
 - srm-v2.cr.cnaf.infn.it (files: 40, streams: 3)
 - storm-fe.cr.cnaf.infn.it (as above)



- Submitted replication requests to DIRAC DMS
 - Overloads the FTS channels with 3 FTS jobs
 - Maintain load
 - Each FTS job contains 100 files
 - Retain control



On the surface...



- IN2P3
- 26th October


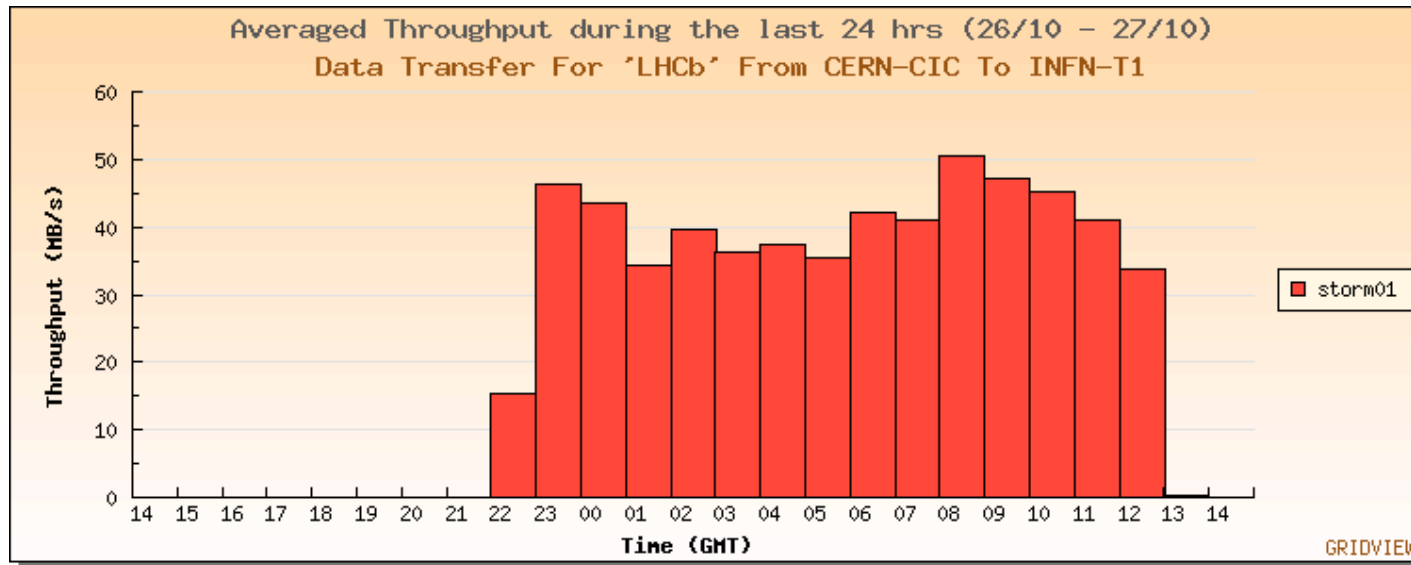


- All 12k transfer successfully
- No transient errors
- ~ 90-100MB/s





...On the surface...

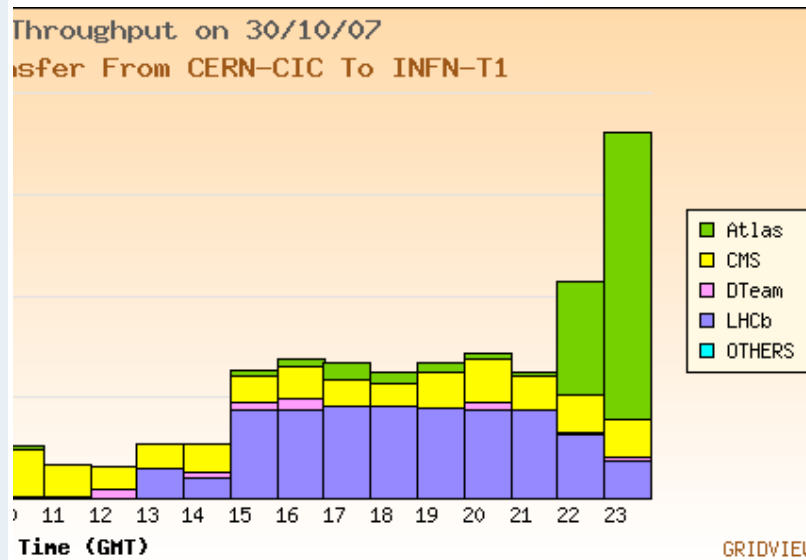
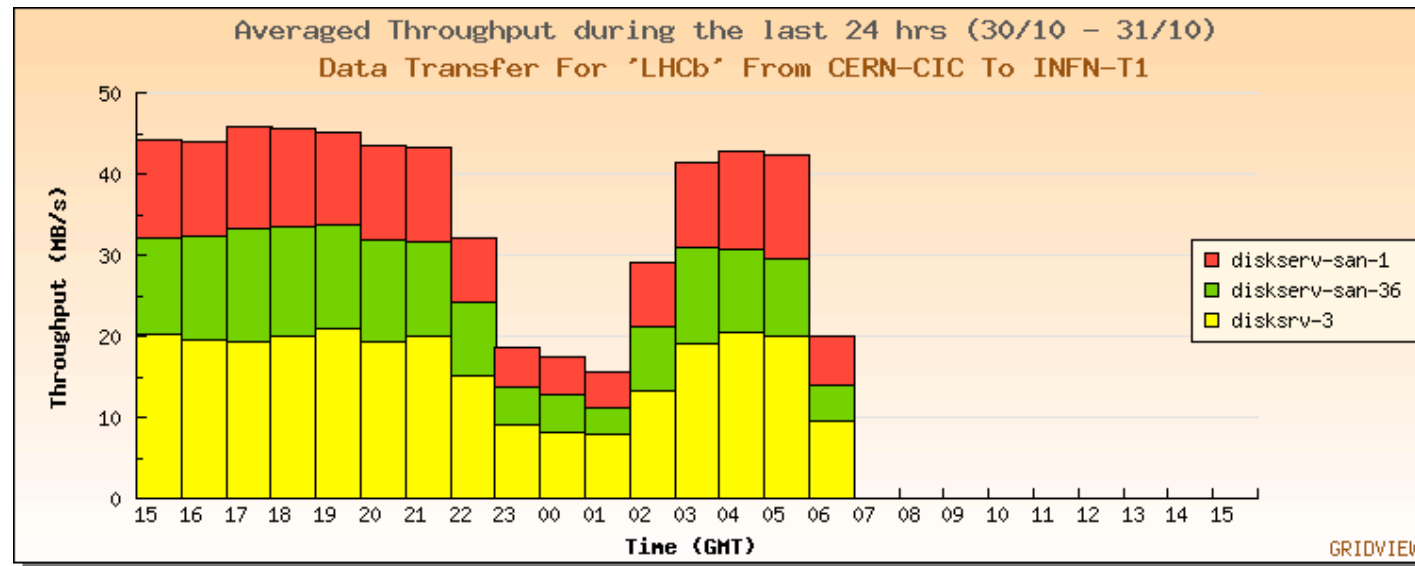


■ CNAF Storm

- 26-27th October
- Successful: 11749
- Failed: 267
- ~40-50MB/s



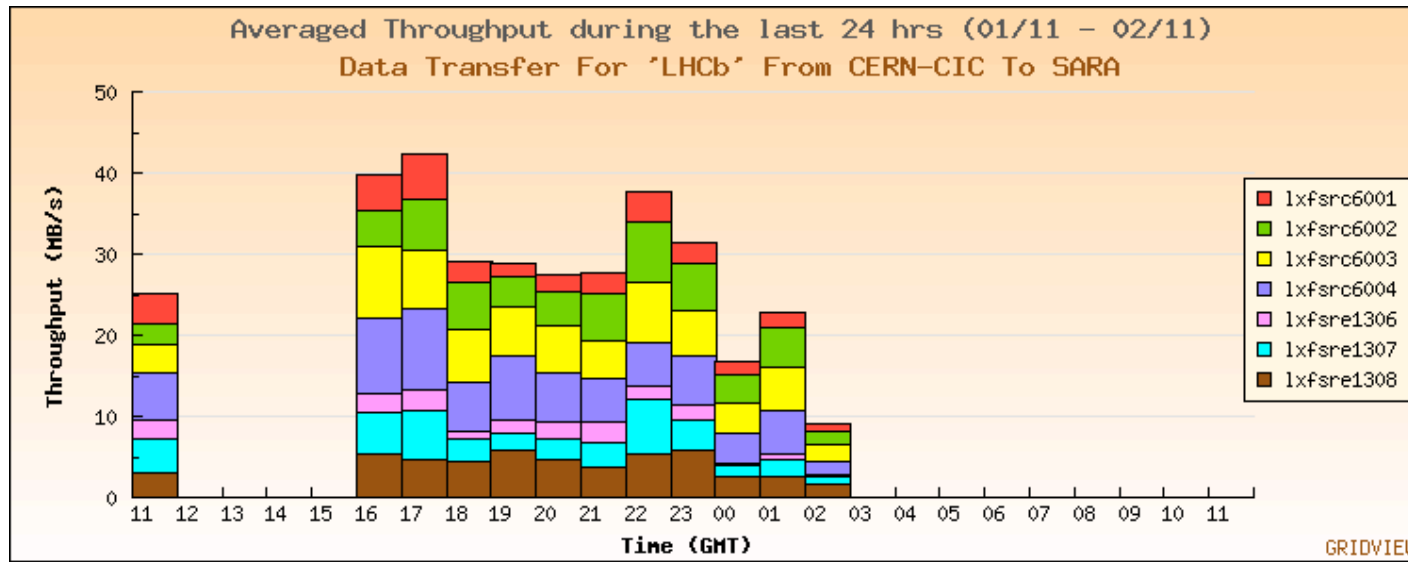
...On the surface...



- CNAF Castor
 - 30-31st October
 - Successful: 11878
 - Failed: 138
 - ~ 40MB/s
 - Dip in throughput due to Atlas activity.



...On the surface.



- SARA dCache
 - 1st October
- All 12k transfer successfully
 - Despite problems with MSS backend
 - Some transient errors
- ~30-40MB/s



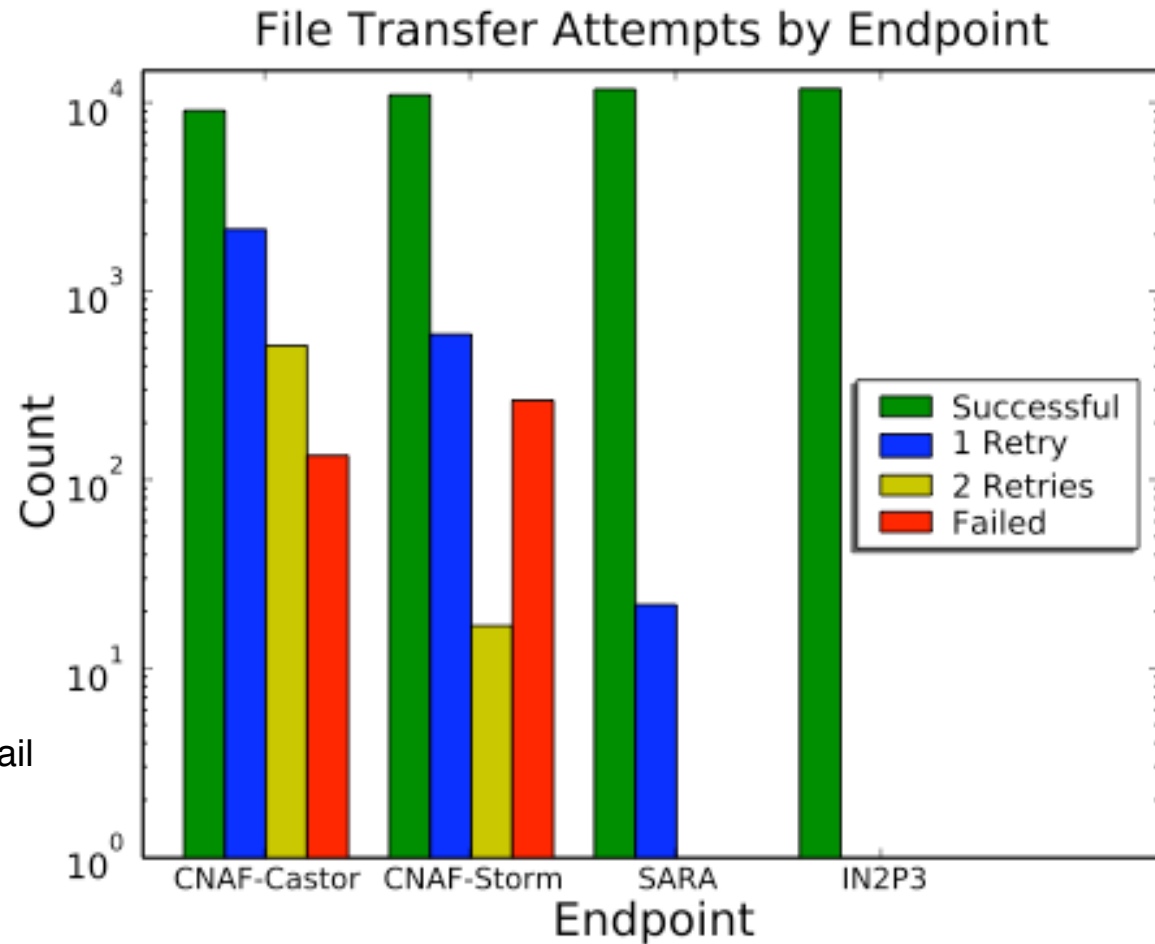
Transient vs. Terminal Errors

- IN2P3
 - No errors!!

- SARA
 - No retries > 1

- CNAF Castor
 - Many errors
 - Decreasing with retries

- CNAF Storm
 - Many errors
 - # Failed > # 2 Retries
 - Files terminally fail after previous error?
 - See next slide





- CNAF Storm

- For all errors

Source error?

Source error

DESTINATION error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare Destination file in 180 seconds	612
DESTINATION error during PREPARATION phase: [FILE_EXISTS] Cannot srmPut file because it already exists!	251
TRANSFER error during TRANSFER phase: [GRIDFTP] a system call failed (Connection refused)	1
SOURCE error during TRANSFER phase: [GRIDFTP] a system call failed (Connection refused)	1

- For failed files

DESTINATION error during PREPARATION phase: [FILE_EXISTS] Cannot srmPut file because it already exists!	250
DESTINATION error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare Destination file in 180 seconds	17

- Looks like the '[REQUEST_TIMEOUT]' error precipitates the '[FILE_EXISTS]' error.

- Operation timeouts not cleaned up properly internally by Storm/FTS?



...Errors Observed...

■ CNAF Castor

Resolved during
exercise

DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] rfio_open error: No such file or directory (error 2 on diskserv-san-1.cr.cnaf.infn.it)	2352
DESTINATION error during PREPARATION phase: [REQUEST_TIMEOUT] failed to prepare Destination file in 180 seconds	394
DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] stage_put error: Unknown internal error	40
DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] Not able to find the version of castor in the database Original error was ORA-00904: "SCHEMAVERSION": invalid identifier	13
DESTINATION error during TRANSFER phase: [INVALID_SIZE] empty file size returned	3
DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] Internal error	2
DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] stage_put error: Could not find suitable filesystem for diskserv-san-36.cr.cnaf.infn.it:/invalidPath/	2
DESTINATION error during PREPARATION phase: [FILE_EXISTS]	2
DESTINATION error during PREPARATION phase: [INTERNAL_ERROR] No SRM method factory found for the requested version []	1
DESTINATION error during PREPARATION phase: [GENERAL_FAILURE] stage_put error: Could not find suitable filesystem for diskserv-san-1.cr.cnaf.infn.it:/invalidPath/	1
SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStagerInterface.c:2145 Address already in use (errno=9, serrno=98)	1

Source error



...Errors Observed...

- SARA dCache

Source error?

Source error

Source error?

Source error

Source error

Source error


DESTINATION error during TRANSFER phase: [PERMISSION] the server sent an error response: 553 553 Permission denied, reason: CacheException(rc=10001;msg=/pnfs/grid.sara.nl/data/lhcb/production/DC06/v1-lumi2/00001368/DIGI/0001/filename.digi not found)	5
TRANSFER error during TRANSFER phase: [GRIDFTP] a system call failed (Connection refused)	5
TRANSFER error during TRANSFER phase: [GRIDFTP] the server sent an error response: 451 451 rfiio read failure: Connection closed by remote end.	4
TRANSFER error during TRANSFER phase: [GRIDFTP] the server sent an error response: 425 425 Cant open data connection.	3
SOURCE error during PREPARATION phase: [REQUEST_FAILURE] [SRM_FAILURE] [SrmGet] failed: SOAP-ENV:Client - Operation now in progress	2
SOURCE error during PREPARATION phase: [GENERAL_FAILURE] CastorStagerInterface.c:2145 Address already in use (errno=9, sermo=98)	1
SOURCE error during PREPARATION phase: [REQUEST_FAILURE] [SRM_INVALID_REQUEST] [SrmGet] failed: SOAP-ENV:Client - RequestRegister.c:182 request record not found in spool, open(srm102.cern.ch:/tmp/0x21940498.xml) and open(srm102.cern.ch:/tmp/0x21940498) failed: Timed out	1

- Very few errors attributable to SARA dCache



...Errors Observed.

- IN2P3 dCache

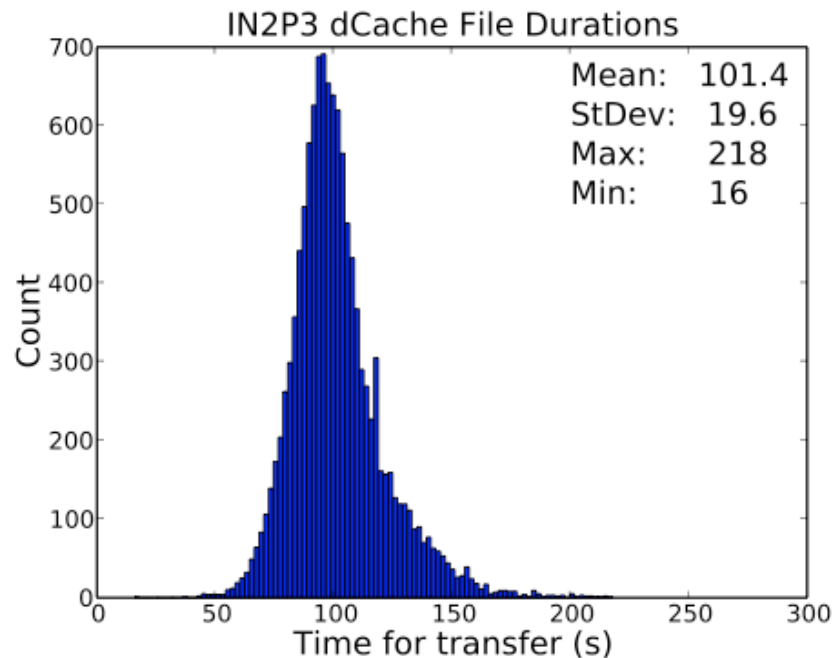
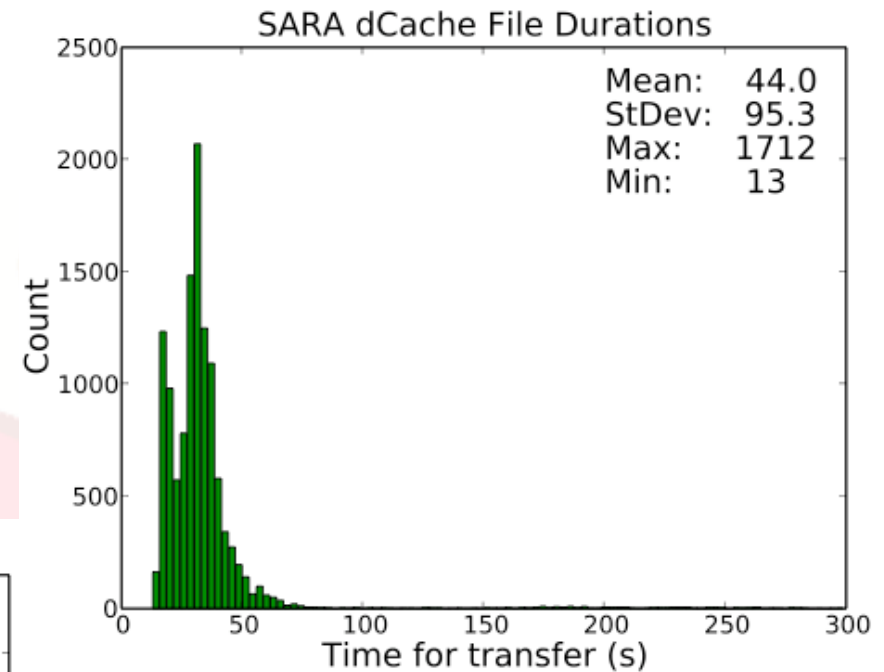




File Transfer Times...

■ SARA

- Exceptionally quick
 - ~92% < 50s
- But StdDev large
 - 1% > 500s
 - 0.3% > 1000s
 - Small percentages but significant with FTS



■ IN2P3

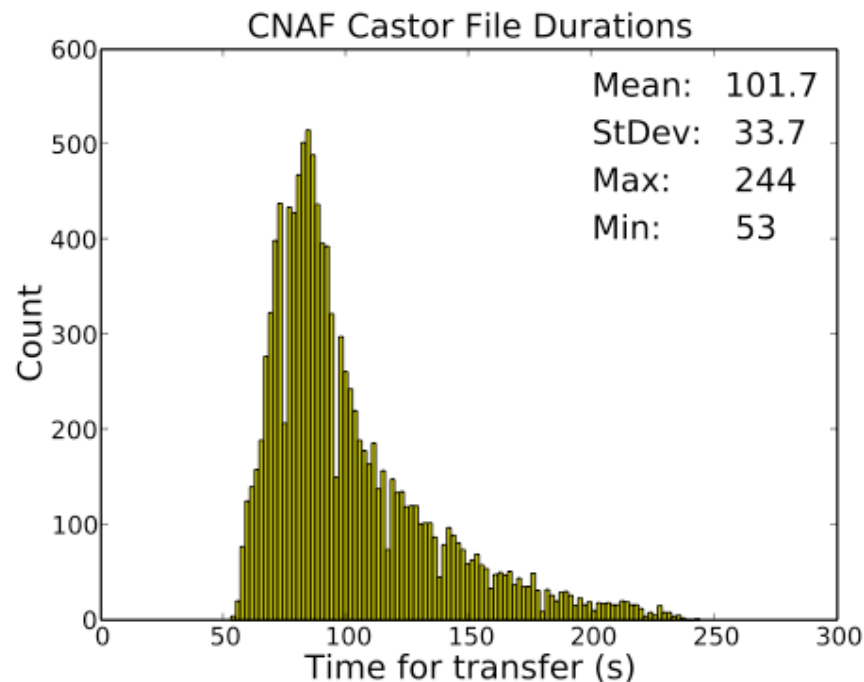
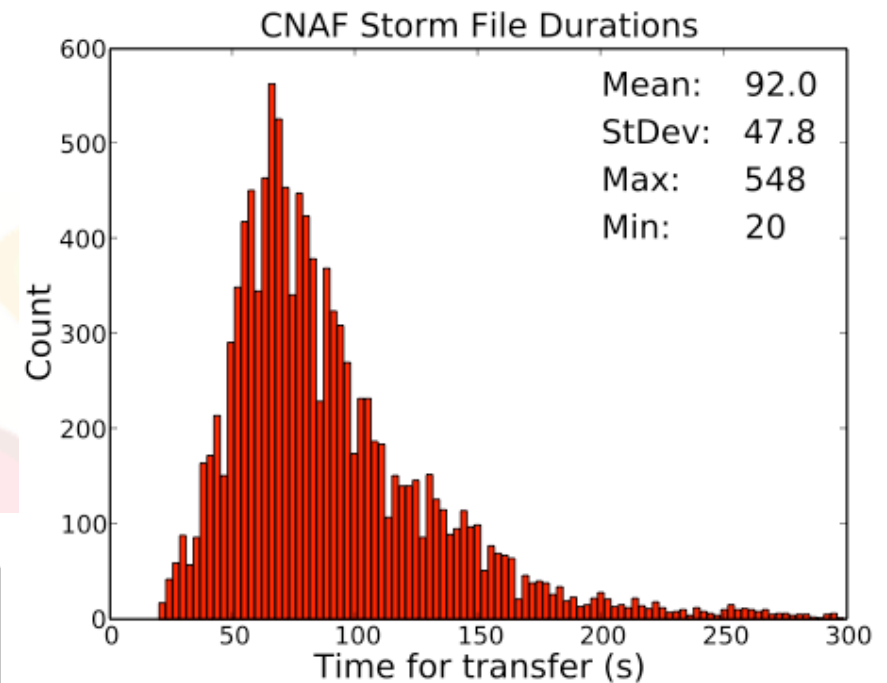
- Consistent in performance
 - Lowest observed StDev
 - 97% 50s < t < 150s
- Average times comparable to other tested endpoints...



...File Transfer Times...

■ CNAF Storm

- Slightly better mean performance than IN2P3
- Higher StdDev
 - 78% $50s < t < 150s$
- Also suffers from long tail
 - 3.3% $> 200s$

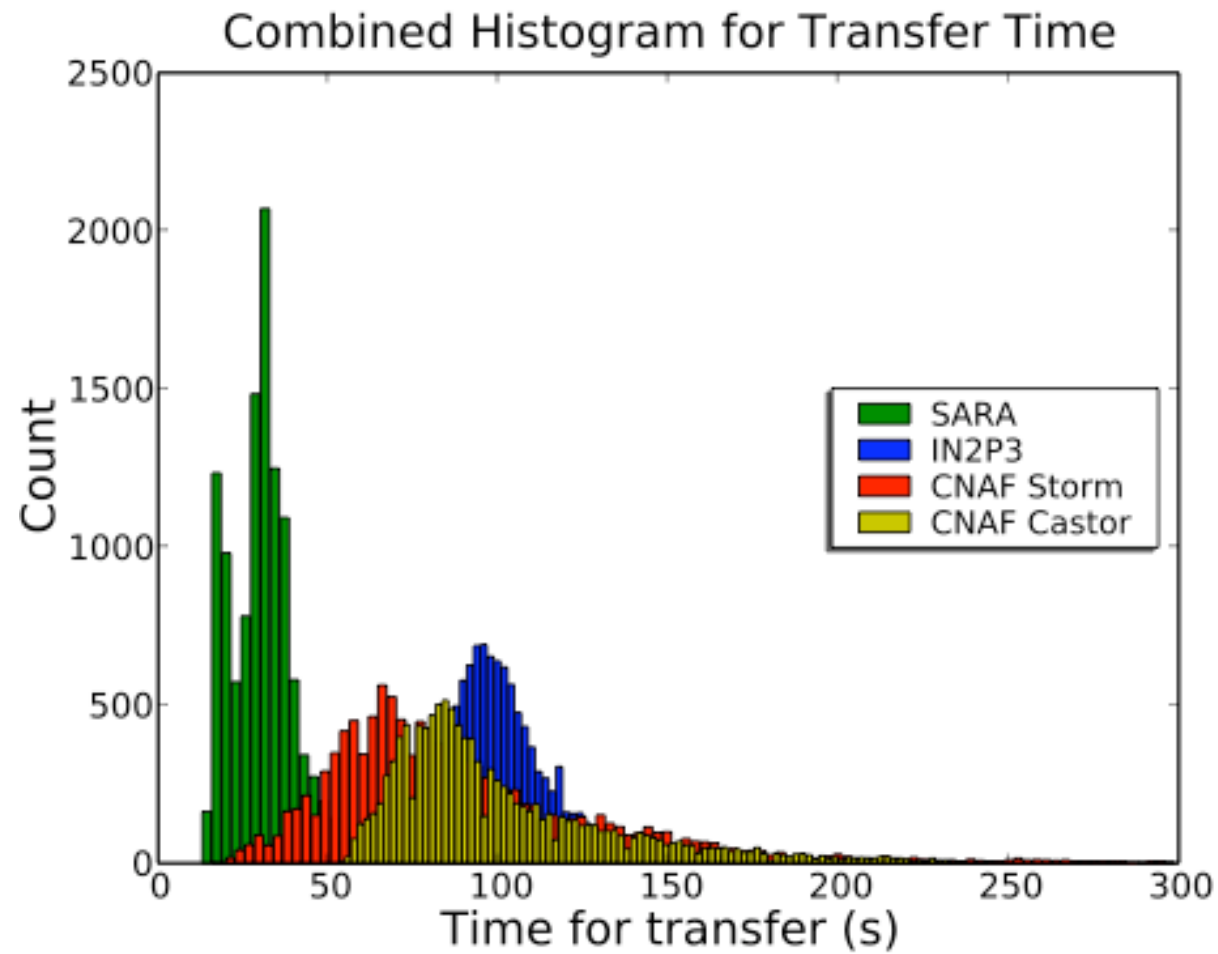


■ CNAF Castor

- Similar profile to Storm
 - Slightly higher mean
 - 2.0% $> 200s$



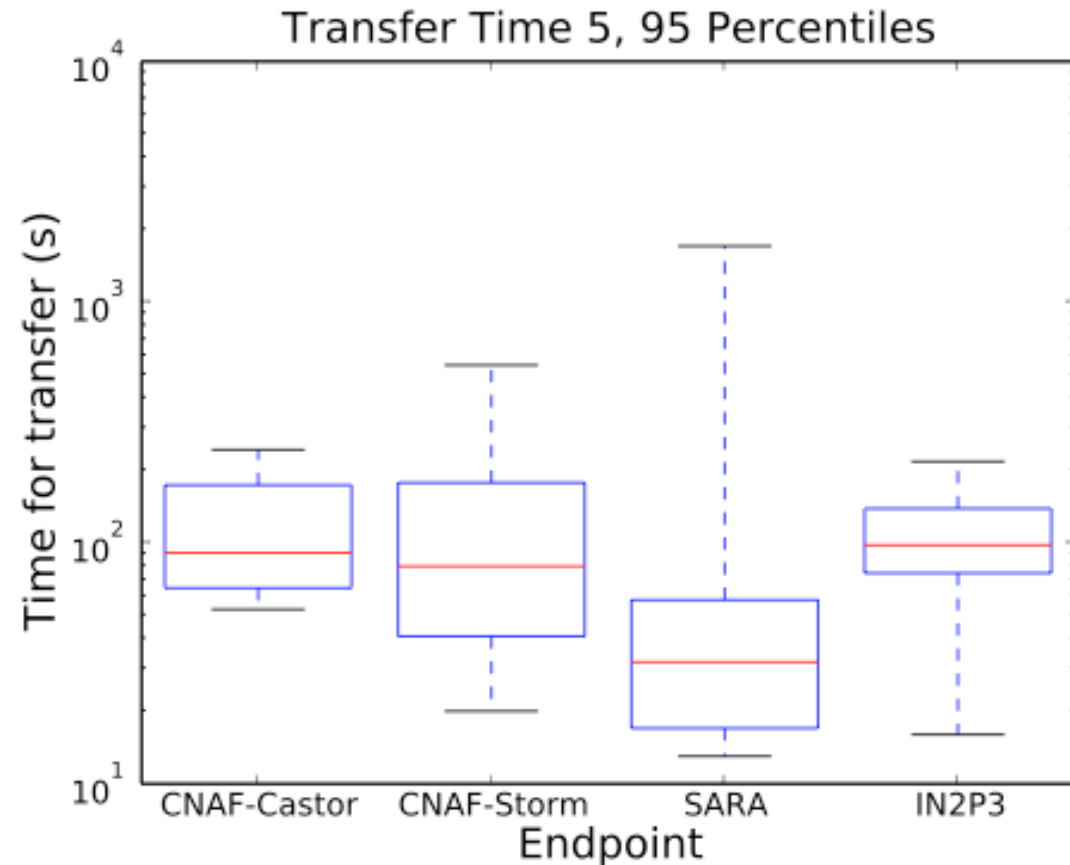
...File Transfer Times...





...File Transfer Times.

- StdDev misleading in representing the long tail data
- Blue box represents the 5,95 percentiles
 - More defined
- Red line
 - Mean
- Top/bottom ticks
 - Max/Min values





- SARA
 - lcg-gt followed by dccp
 - OK
- IN2P3
 - lcg-gt followed by dccp
 - OK
- GRIDKA
 - WN running old gLite version
 - lcg-utils incompatible
- CNAF Castor
 - lcg-gt followed by rfcv
 - OK
- CNAF Storm
 - Still being tested
 - Read from ROOT application





- Perform the same test with GridKA

- Experience gained with the destination SRMs
 - Repeat the same test using CERN SRM2 as source
 - Waiting for some configuration to be resolved

- Tools ready to perform bulk removal
 - Using the gfal libraries
 - Plan to remove data immediately after transfer
 - Determine the performance of this operation
 - From initial tests should be able to remove faster than transferring
 - Continual transfer-removal over sustained period
 - How long? 1 week? 1 month?



- **Stress tests**
 - LHCb throughput requirements (40MB/s aggregated) already met
 - Certainly not the other VOs
 - Space shouldn't be a problem
 - FTS Transfer slots are
 - LHCb share on production channels modest
 - Testing activities shouldn't interfere with production activities

- **Comparative studies**
 - Hard to draw comparisons between different instances
 - Configuration different for each channel
 - Dedicated testing channels?
 - Testing different parameters for individual SRMs
 - Nice to have testing channels
 - Modify channel parameters and gauge performance

- **Are we aiming to break the SRMs?**



Last slide.

