



Validation tests of CNAF storage infrastructure

Luca dell'Agnello
INFN-CNAF

Reasons for the test

- LHC is coming ☺
 - Last chance to finalize the technical issues of our setup
- Need to validate and perform final tests of CNAF storage infrastructure
 - Validation of network interconnection with the farm
 - Validation the SAN setup
 - Validation of new hardware (EMC CX3-80)
 - Storage system tuning and performance measurement
 - CASTOR
 - GPFS/StoRM
 - Get use of Large file system with GPFS (>100 TB)
- Test of other storage systems (mainly for Tier2s)
 - Xrootd
 - dCache

Storage classes @ CNAF

- Implementation of 3 Storage Classes needed for LHC
- Disk0Tape1 (D0T1) → CASTOR
 - Space managed by system
 - Data migrated to tapes and deleted from when staging area full
- Disk1tape0 (D1T0) → GPFS/StoRM
 - Space managed by VO
- *Disk1tape1 (D1T1) → CASTOR*
 - *Space managed by VO (i.e. if disk is full, copy fails)*
 - *Large buffer of disk with tape back end and no gc*

Our setup for the test (1)

- 2 EMC CX3-80
 - ~ 130 TB raw each
 - 2 storage processors
 - 4 GB of RAM for SP
 - 500GB LCFC disks
 - 8x4 Gbps FC ports
- FC Director: Brocade 48000 (x1)
 - 128x4 Gbps FC ports
- 24 Disk Servers: Dell SE1950 (one used only for management)
 - Dual core bi-processor Xeon 1.6GHz
 - 4 GB RAM
 - Dual-port Qlogic 246x HBA
 - 1 gbps uplink to the switch
 - new OS used (SLC 4.4 64 bits)
 - Some kernel tunings to optimize network throughput
 - Kernel: 2.6.9-42.0.8.EL.cernsmp

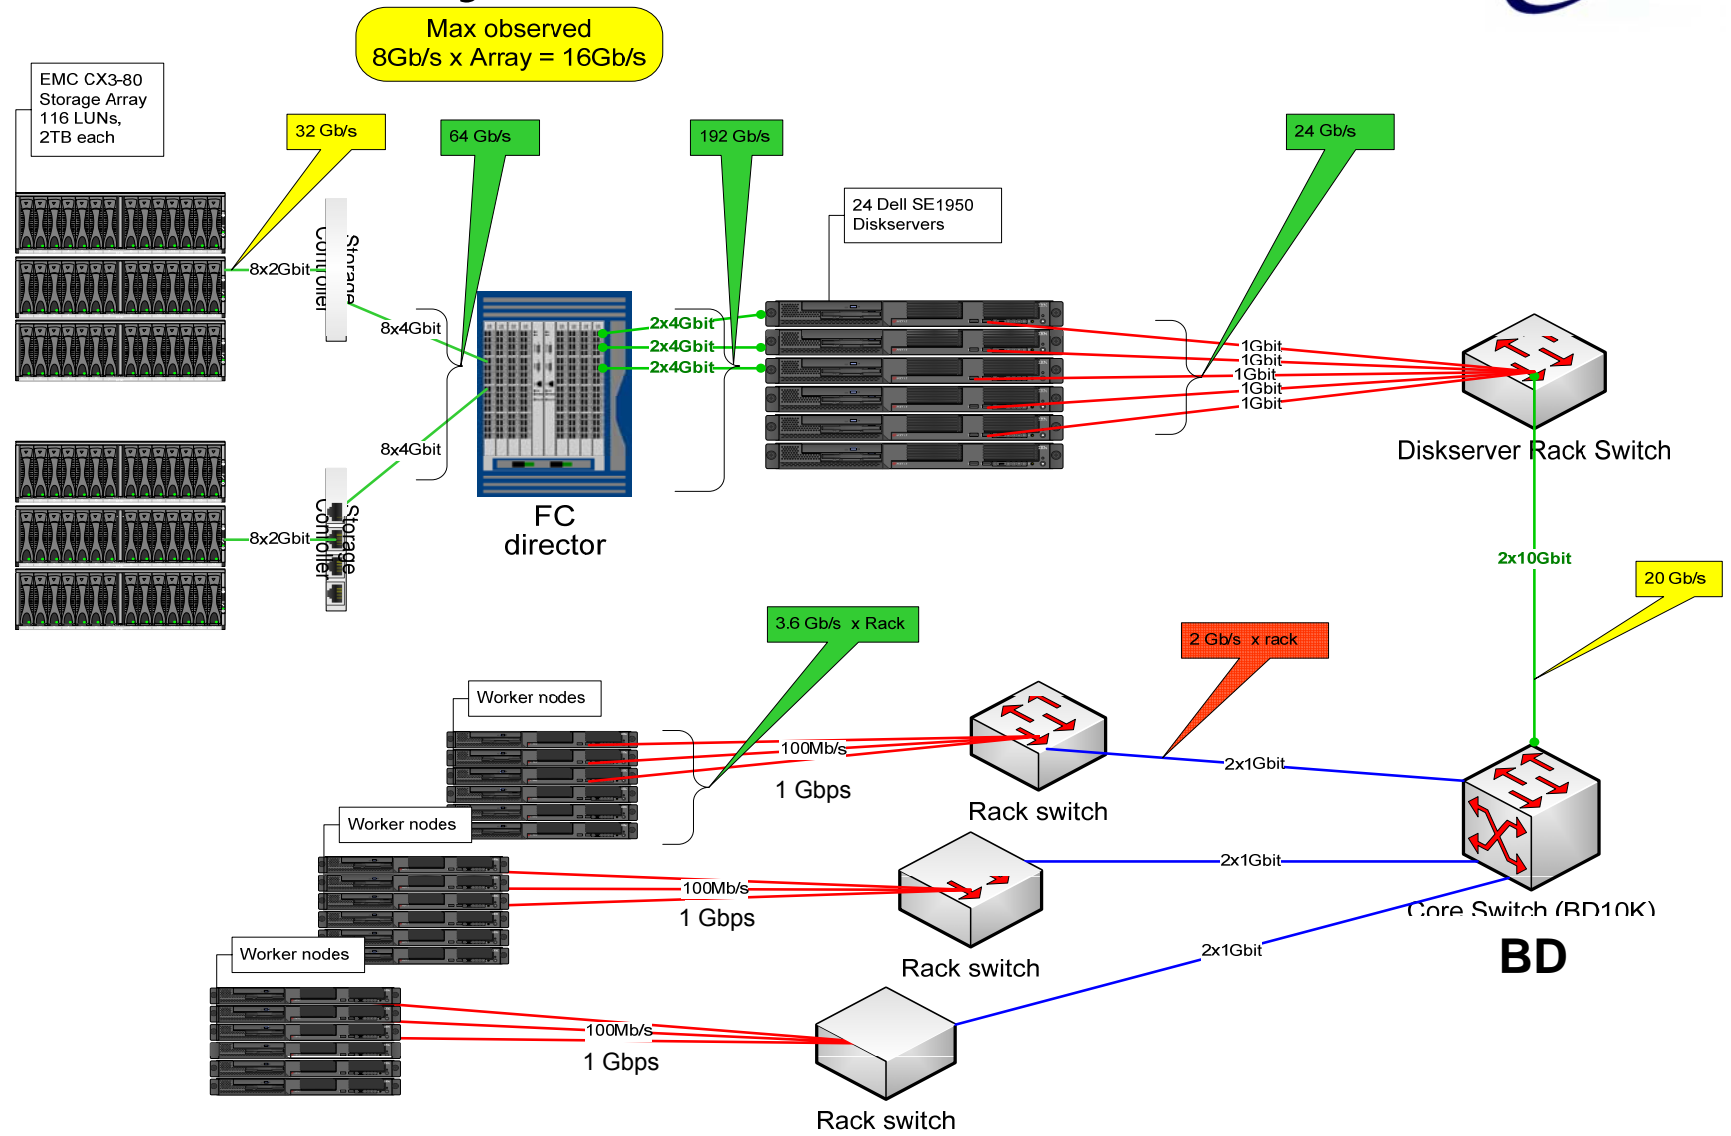
Our setup for the test (2)

- 8 wn racks (~ 1100 CPU slots)
 - 2 Gbps uplink from the rack switches
 - Standard OS (SLC 3.8 32 bits)
 - Standard mw (LCG 3.1)
 - Kernel 2.4.21-47.EL.cernsmp
- BD core switch interconnecting rack switches and disk servers switch
 - 2x10 Gbps links from BD to servers switch
 - 16 Gbps aggregated from BD to wn racks switches
 - ER16 switch not used due to interoperability issues
- Some minor problems with LDAP servers (used for local authz)

Kernel tuning on disk-servers

- Some kernel 2.6.9-42.0.8.EL.cernsmp parameters changed from their default values i.e.:
 - net.core.rmem max = 2097142 (default = 135168);
 - net.core.wmem max = 2097142 (default = 135168);
 - net.ipv4.tcp rmem = 4096 262144 2097142 (default = 4096 87380 174760);
 - net.ipv4.tcp wmem = 4096 262144 2097142 (default = 4096 16384 131072);
 - net.core.netdev max backlog = 2500 (default = 300);
 - vm.dirty ratio = 0 (default = 40);
 - vm.min free kbytes = 50000 (default = 1442).

Test-bed layout





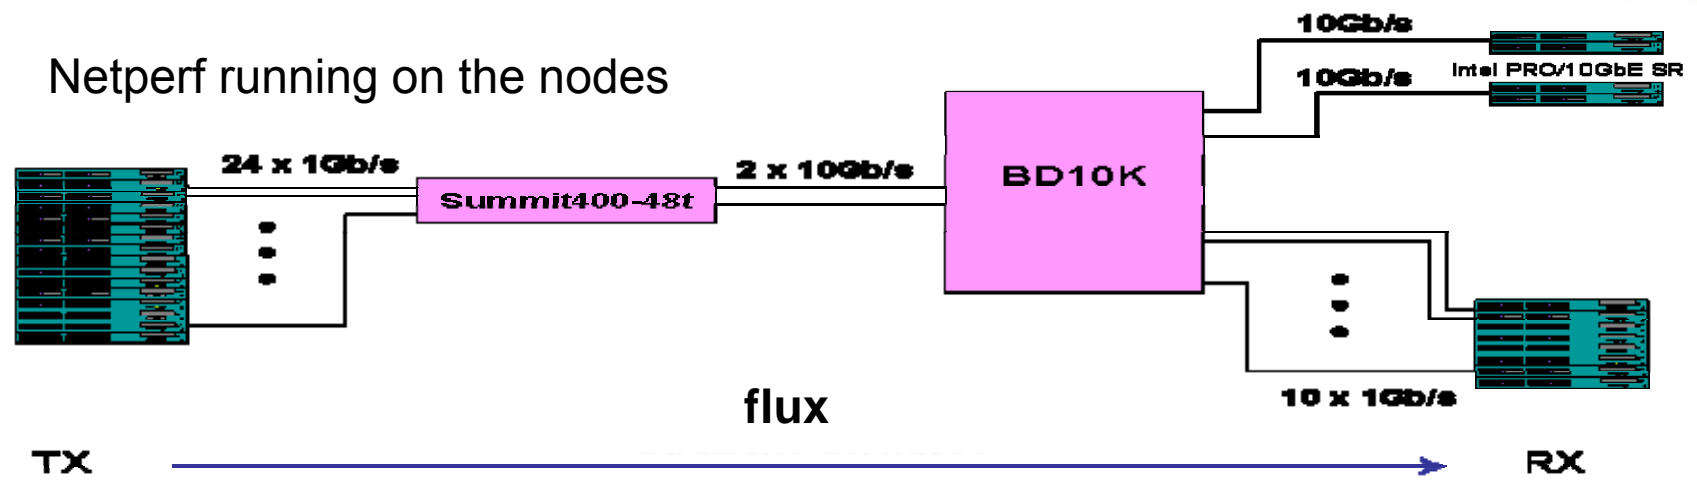
Test phases

- Network validation
- Local throughput measurement
- Remote throughput measurement
- Realistic LHC analysis
- SRM tests
 - First results just with StoRM
 - Tests with other SRMs in progress

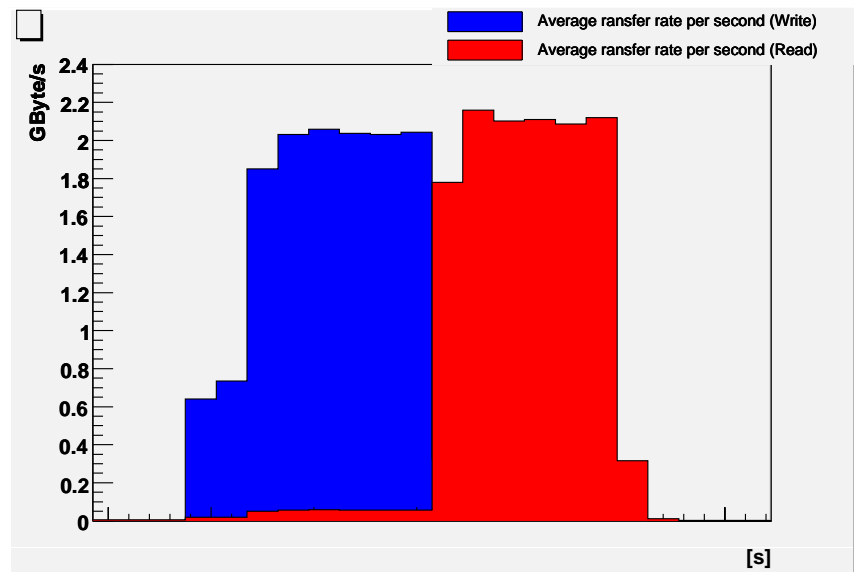


Network validation

Netperf running on the nodes




20 Gbps uplink validation



diskserver throughput \leftrightarrow wns

Storage systems configuration

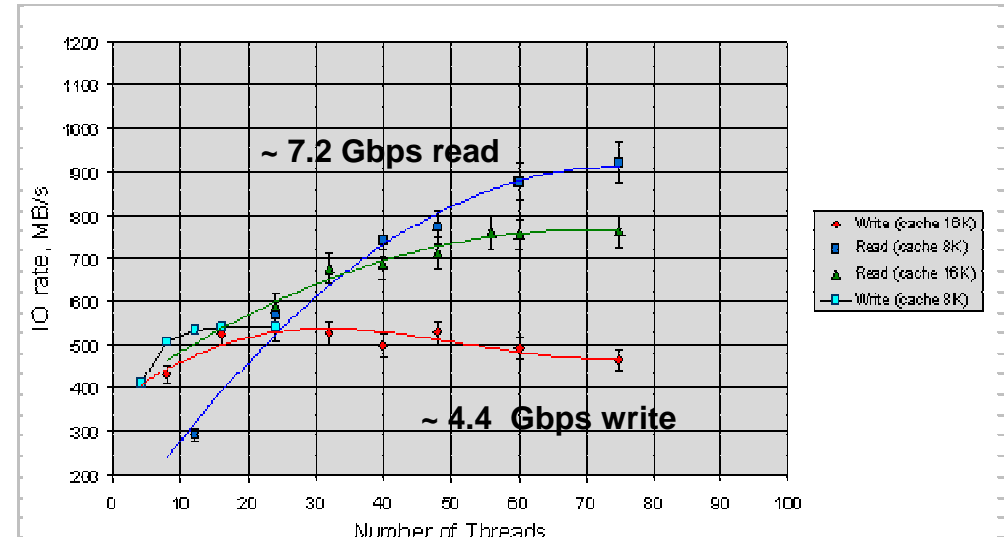
- **23 I/O servers**
 - CASTOR, GPFS, xrootd (xfs), dcache (xfs) installed on each server
- For CASTOR test instance of stager used
- For GPFS, production cluster used (but local tests on an ad hoc cluster)
 - **56 LUNs = 56 NSDs**
 - single file system of **110 TB**
 - **4 storage groups = 4 separated SAN zones**
 - Data and Metadata distributed across all NSDs
 - Each LUN has its primary and backup server



Local throughput measurement

single EMC CX3-80 with xfs
type of test:

```
dd if=/dev/zero of=/var/xxxx
bs=1024k
File size=12GB
```

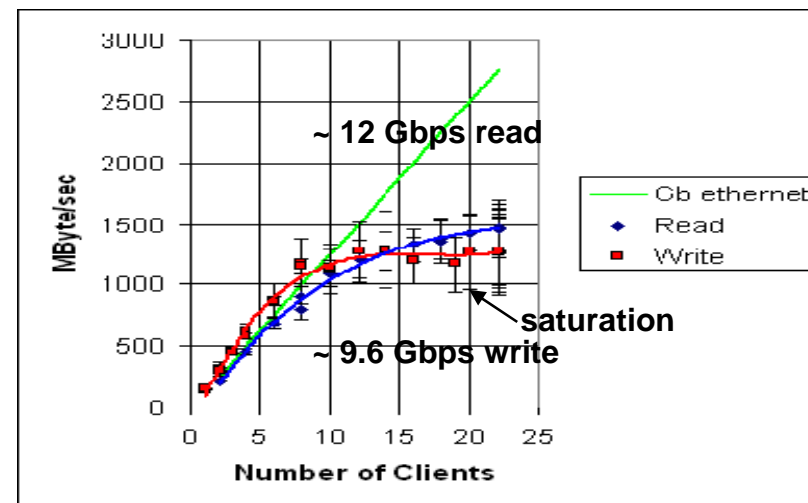


2 EMC CX3-80 with GPFS (v. 3.1.0-7)

24 server with Gbit and 2x4Gbit FC
connection

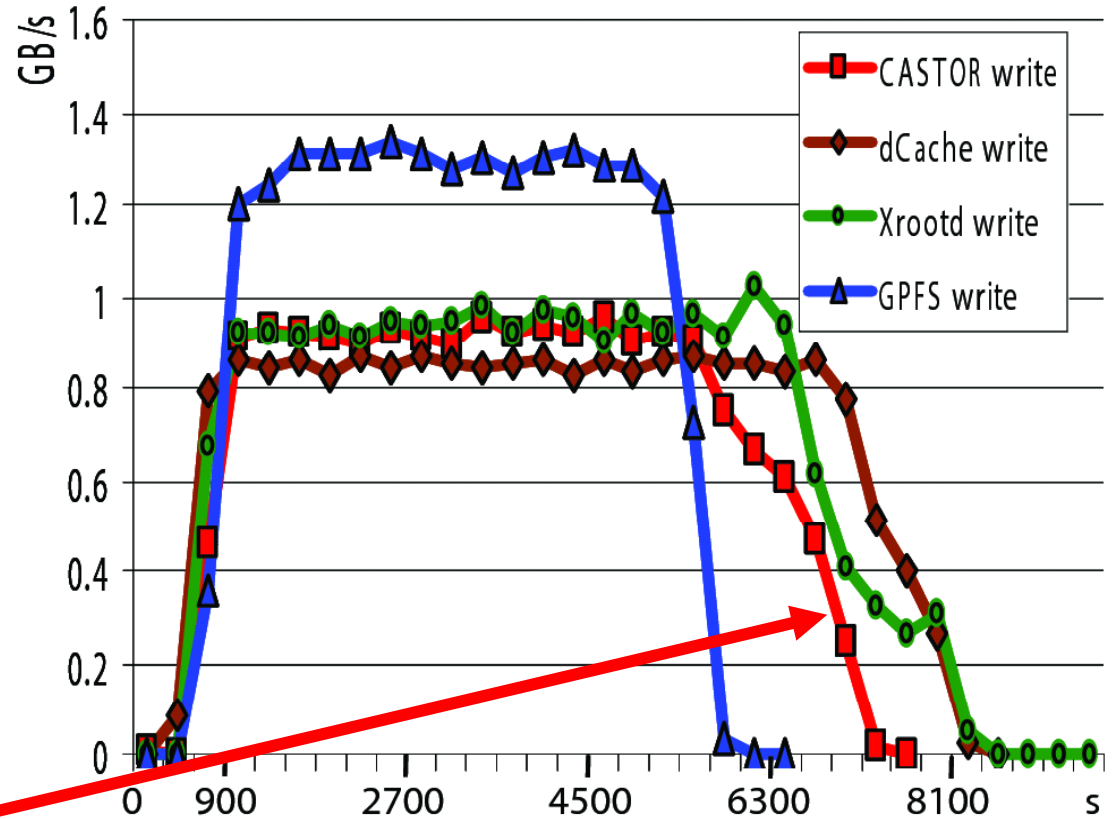
type of test:

```
dd if=/dev/zero of=/gpfs/file
bs=1024k
File size=12GB
```



Sequential Write throughput


- 5GB files (block size=64K)
- ~1100 CPU slots
 - 8 wn racks
 - 16 Gbps network available



Failures with Castor writes - integral of Castor curve less than the other ones

Sequential Read throughput

- Read back files previously written
 - 5GB files (block size=64K)
- ~1100 CPU slots
 - 8 wn racks
 - 16 Gbps network available

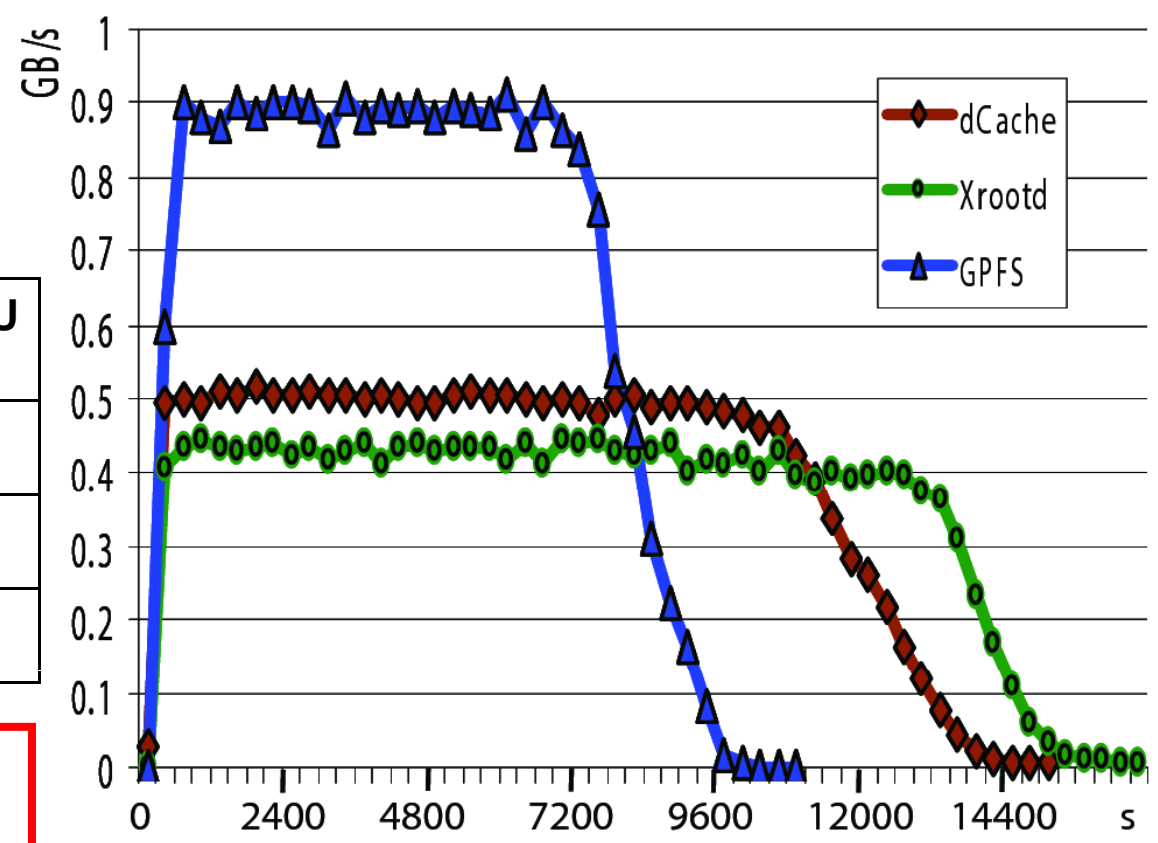


Failures with Castor reads plus failures in previous writes - integral of Castor curve less than the other ones

A real analysis case: LHCb DaVinci

- 1000 analysis jobs running on production farm and accessing same set of data via dCache, Xrootd and GPFS
- Network stats reported


Test	Total time [s]	Average CPU usage (%)
GPFS	9500	36.43
Dcache	14500	22.69
Xrootd	16000	21.17



CASTOR not considered in this test due to the large failure rate previously observed

Storm@CNAF for T0D1

- A (pre)production deployment for StoRM
 - 4 gridftp/GPFS servers
 - 38 TB of GPFS file systems
 - 3 load-balanced StoRM servers




Throughput test description

- Tests using low level tools (preliminary to FTS tests)
- Multithreaded script
 - Each thread keeps transferring the same source files (**LHCb simulated data files (~ 150MB)**) to always different destination files.
- Each thread does sequentially (for a configurable period of time):
 - PreparetoPut on StoRM
 - Storm polling until destination TURL ready
 - Globus-url-copy sourceturl-destturl
 - PutDone (on StoRM)
 - Is on StoRM for checking the size of the transferred file
- This process is iterated for a fixed time interval


Throughput tests

- From Castor (CERN) to GPFS/StoRM (CNAF)
- Total handled files ~100K
 - At least 400K interactions
- Failure rate in copying =0.2%
 - i.e. number of != 0 exit code from globus-url-copy
 - Failure rate due to Storm<0.1%
- Amount of data 17TB
- Bandwidth peak: 370MB/s




15 parallel streams/transfer
120 concurrent transfers

Removal Test




StoRM stress test (no throughput)


- Objective:
 - Test how many simultaneous requests the system can handle
 - What happens when the saturation is reached
- A main script launches a variable number of parallel processes.
- Each process performs 2 steps
 - List the content of the destination directory in StoRM and removes all the files in it.
 - Performs N subsequent ptp requests to the system polling it to get the TURL (No data transfer!).
- What we measured
 - Total time to perform the N requests
 - Percentage of failed requests



Stress test results

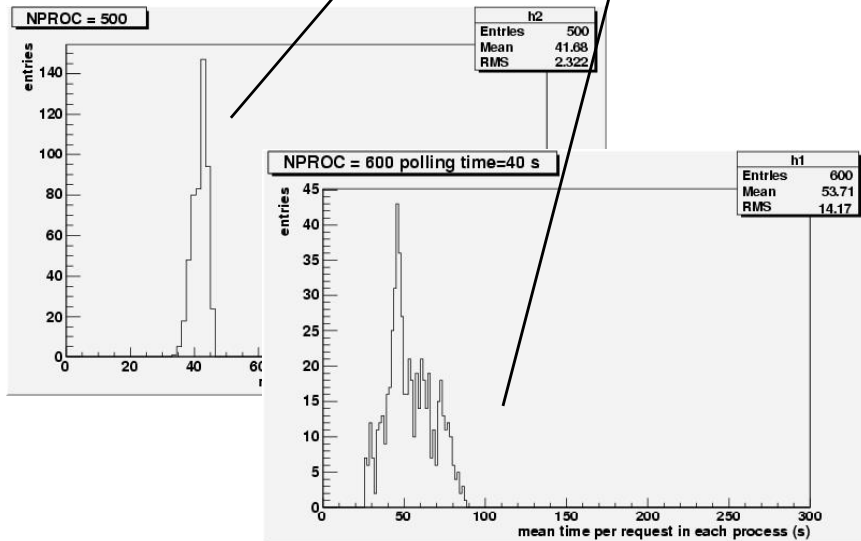


The mean time (delta time between ptp and successful polling) per request slightly increases with the number of parallel requests.



Failed requests vs number of parallel requests:

- Almost no failure up to 500
- For 600 parallel processes a noticeable rate of failures is observed



Conclusions

- The validation of GPFS setup has been successful
 - GPFS, dcache and xrootd all well performing
 - GPFS in particular demonstrated robustness and high performances
- GPFS more than suitable for usage in real world
 - Long experience at CNAF (> 3 years)
 - ~ 20 GPFS file systems in production at CNAF (~ 180 net TB) mounted on all farm wns
 - administration simple (no Oracle, no LSF, intuitive admin commands etc...)
 - Reliable: failure of a single file server does not affect data availability
 - SRM interface (StoRM) has been released
- CASTOR showed no surprises
 - Well known LFS plug-in problem
 - Upgrade end of this month ?
 - Issue to investigate: more stagers needed ? (CERN model)
- Confident to fit LHC and non LHC experiments with present model

Credits

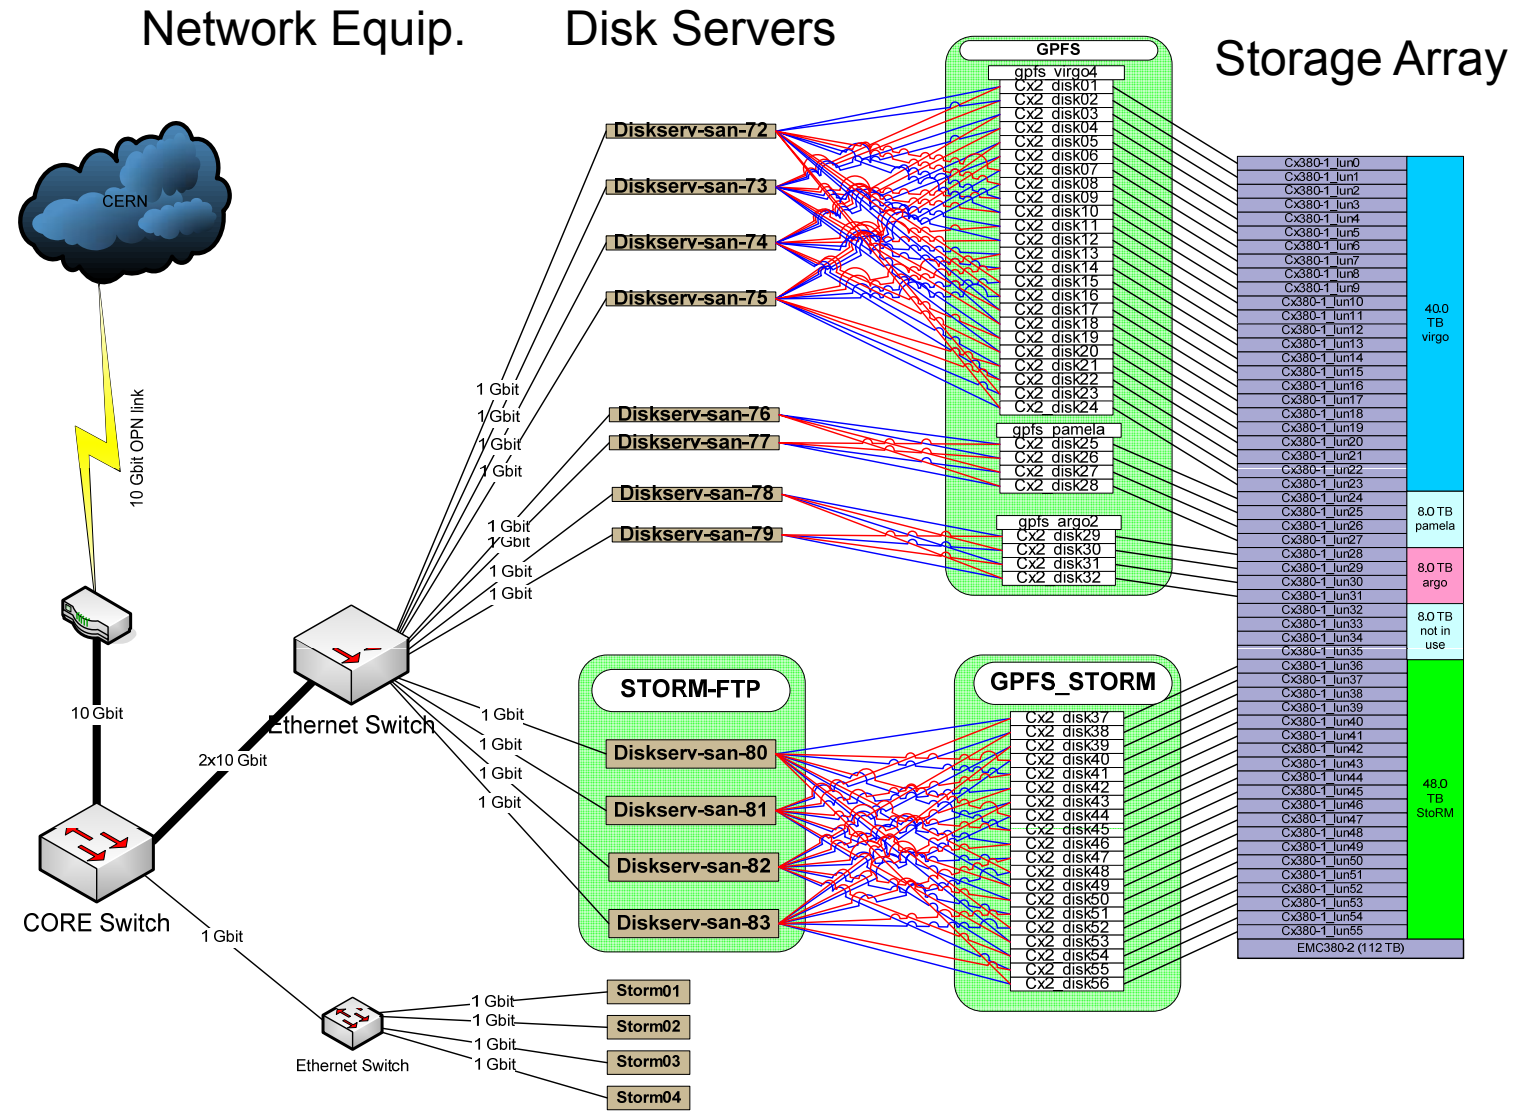
- Network setup: M. Bencivenni, D. Degirolamo, R. Veraldi, S. Zani
- Farm setup: M. Donatelli, A. Italiano, D. Salomoni
- Monitoring setup: F. Rosso, D. Vitlacil
- Storage hw setup: A. D'apice, PP. Ricci, V. Sapunenko
- Storage systems setup: G. Donvito, A. Fella, F. Furano, G. Lore, V. Sapunenko, V. Vagnoni, D. Vitlacil
- Storage systems tests: A. Carbone, L. dell'Agnello, G. Donvito, A. Fella, F. Furano, G. Lore, V. Sapunenko, V. Vagnoni


- Storm development team: A. Forti, L. Magnoni, R. Zappi
- Storm tests: E. Lanciotti, R. Santinelli, V. Sapunenko, V. Vagnoni



Infrastructure point of view:

File Systems





Please note
UI is a very old machine



storm-fe.cr.cnaf.infn



(PIII 1GHz, 512MB)

FE accepts request, authenticates
and queues data into DB



DNS balanced Front End

BE reads requests from DB and executes
them on GPFS (running clients)



06-Sep-07 CHEP 2007, Victoria
4 Gridftp servers running also GPFS servers

More details on the testbed

- Front-end: storm03-storm04
 - dual AMD Opteron 2.2 GHz, 4 GB ram
- Back-end: storm01
 - dual Intel Xeon 2.4 GHz, 2GB ram
 - also runs mysqld
- 4 GPFS disk servers
 - dual Intel Xeon 1.6 GHz, 4 GB ram
 - also running gridftpd

StoRM version 1.3-15



Throughput test description

- Tests using low level tools (preliminary to FTS tests, which will be finally used by LHCb)
- Multithreaded script . Each thread keeps transferring the same source files (**real LHCb DST and DIGI files O(100M)**) to always different destination files.
- Each thread does sequentially (for a configurable period of time):
 - PtP on Storm
 - Storm polling until destination TURL ready (up to 10 retries exponential varying time between 2 retries)
 - Globus-url-copy (or lcg-cp) sourceturl-destturl (source_surl-dest_turl) respectively
 - PutDone (on StoRM)
 - Ls on StoRM for computing the size transferred
 - Iterate the previous action until total time is reached

Sequential I/O throughput summary



write

Test	TB tbt	Transf. TB	TB Network	Disk Net Ratio	Efficiency (%)	Average Network Rate (GB/s)
GPFS	5.29	5.21	5.56	0.94	98.43	1.16
Castor	4.52	4.49	4.89	0.92	99.37	0.82
Dcache	5.24	5.22	5.43	0.96	99.51	0.77
Xrootd	5.24	5.24	5.59	0.94	100.00	0.85

read

Test	TB tbt	Transf. TB	TB Network	Disk Net Ratio	Efficiency (%)	Average Network Rate (GB/s)
GPFS	5.21	5.21	5.53	0.94	100.00	1.43
Castor	4.49	4.11	4.37	0.94	91.58	1.13
Dcache	10.44	10.18	10.71	0.95	97.52	1.29
Xrootd	10.49	10.44	11.10	0.94	99.54	1.00

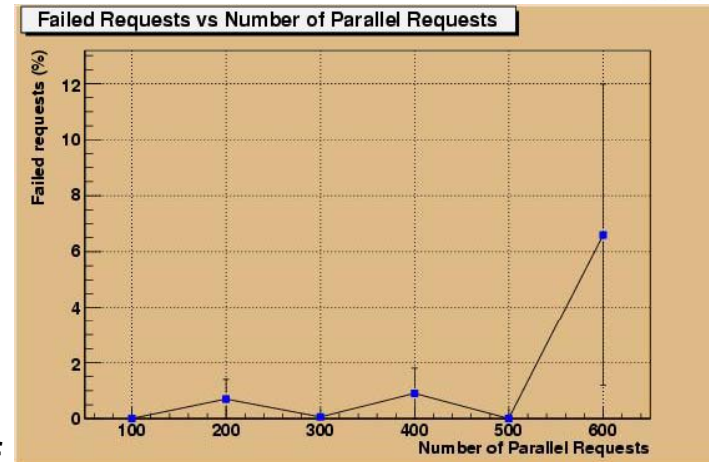
- **TB tbt**: foreseen TB based on the number of jobs effectively running
- **Transferred TB** : TB written on disk
- **Efficiency**: percentage of effectively transferred files
- **Average Network Rate (GB/s)**: mean of 5 mns. mean values

Stress test results (2)



Failed requests vs number of parallel requests:

- Almost no failure up to 500
- For 600 parallel processes a no negligible rate of failures is observed



Causes of the failed requests: mainly 3 types of error found.

1. “CGSI-gSoap: error reading token data: connection reset by peer” and “CGSI-gSoap: could not open connection! TCP connect failed in tcp_connect()”
2. Ls returns SRM_INTERNAL_ERROR “client transport failed to execute the RPC. HTTP response: 0.
3. Some client commands hung for hours (mainly statusptp)

•Almost 100% of failures for gSoap timeout occur in the first phase: when creating the destination directory or listing the content of the directories and deleting the files
▶ specific tests needed for **rm**, **ls**, **mkdir**

•Almost **no failure** in the ptp-statusptp phase