

CASTOR and StoRM in May CCRC08: Configuration, experience and questions for developers

Elisabetta Ronchieri

INFN CNAF

On behalf of INFN T1 and RAL T1

Overview

- INFN CNAF
 - atlas, cms, lhcb, and alice
- RAL
 - atlas,cms and lhcb

ATLAS

- StoRM (T0D1)
 - 1 end-point storm-fe.cr.cnaf.infn.it
 - 2 FE, 1 BE, 4 gridftp, 150 TB
- Castor (T1D0)
 - 1 end-point srm-v2.cr.cnaf.infn.it
 - 2 FE, 1 BE, 1 srm DB Oracle single instance
 - 6 disk-servers, 36 TB

ATLAS Issues and Solutions

- ATLAS adopted both castor srm v2 and StoRM endpoints
 - T1D0 with CASTOR, T0D1 with GPFS/StoRM
- Main issues with gridftp layout:
 - Multiple retransmissions of the same data between gridftp and GPFS, overloading the servers:
 - GPFS configured to use 1 MB block-size
 - SLC3 gridftp servers use 64 KB block-size for writing on disk (i.e., on GPFS) while SLC4 use 256 KB
 - Problem raises if no. of streams/file > no. GPFS disk-servers
 - Adopted Solution:** lower no. of streams/file on CNAF FTS server
 1. Not fully effective
 2. **No retransmissions but low throughput**

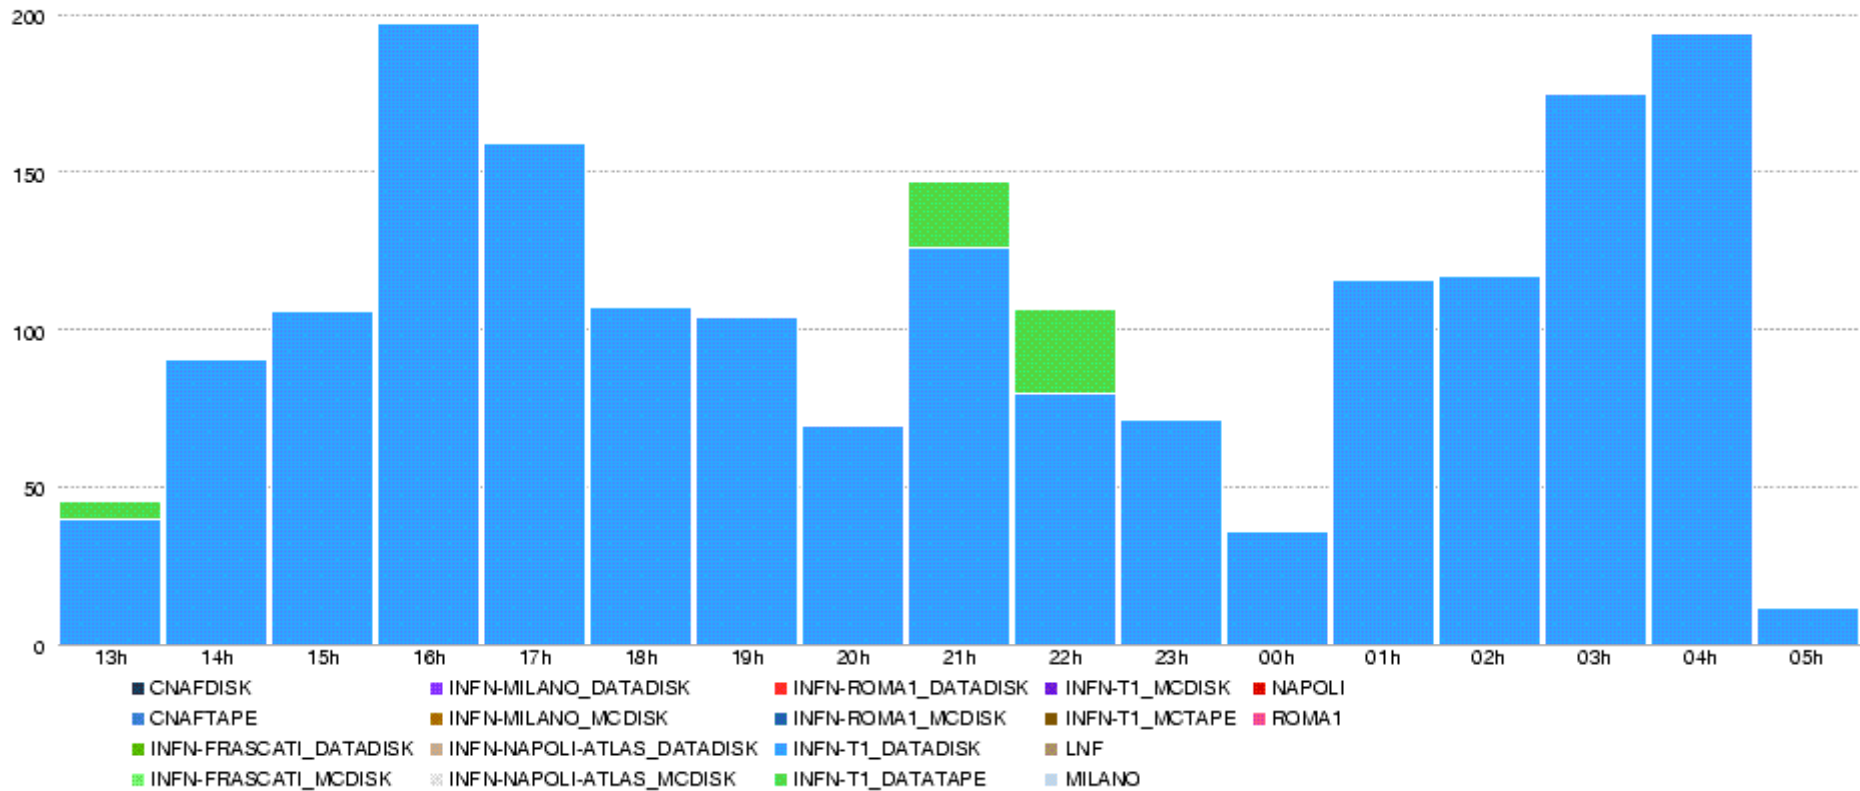
ATLAS Issues and Solutions

- Main issues with gridftp layout:
 - gridftp servers → Slow (in and out) data transfers and backlog
 - Solution:** at first 3 SLC3 gridftp servers then moved to 4 SLC4 gridftp servers
 1. double transfer rate compared with SLC3
 2. improved retransmission issues due to block size of 256KB, therefore FTS streams increased from 3 to 5
- Added new 20 TB mounted on one only channel causing slow transfers for data that wrote on ATLADATADISK
 - Adopted solution:** moved these TB on a different GPFS pool and removed in writing for 'in' data
 - Next solution:** mount these TB on two channels on the old GPFS pool (at the end of CCRC)

ATLAS Issues and Solutions

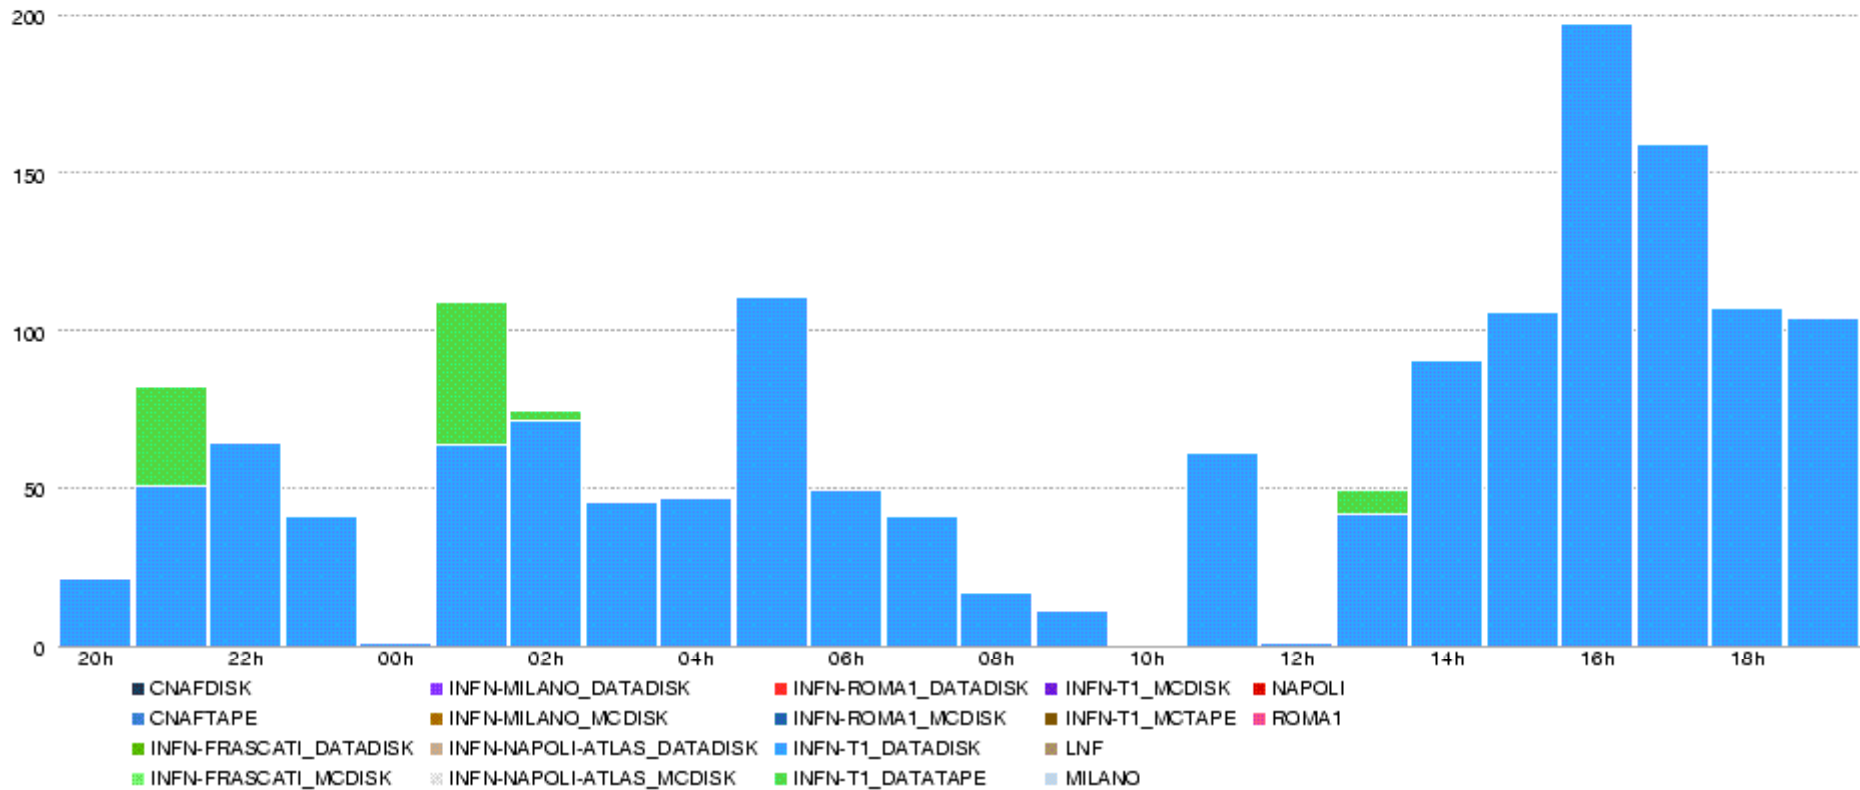
- StoRM issue:
 - Disappearance of space token ATLASDATADISK from the SRM configuration observed during a failure into yaim/quattor configuration
 - No loss of data but only transfers failures
 - Under investigation from the StoRM support team** (i.e., problem reproduced)
- (minor) CASTOR issue:
 - Some HW failures on a disk-server
 - Solution:** disk-server disabled promptly
 - Configuration issue on a disk-server
 - Solution:** fix it promptly
- General issue:
 - srm get errors due to FTS timeout of 3600 sec
 - Solution:** reduced FTS get timeout from 3600 sec to 3000 sec for all sites

ATLAS



- Throughput from 13:00 of the 29th of May to 17:00 of the 30th of May.
- Growth from the 13:00 to 17:00 as soon as gridftp server was added
- Average throughput of 120 MB/s

ATLAS



- Throughput from 20:00 of the 28th of May to 20:00 of the 29th of May
- Growth from the 14:00 to the 17:00

CMS

- StoRM (T0D1)
 - 1 end-point `storm-fe-cms.cr.cnaf.infn.it`
 - 2 FE, 1 BE, 6 gridftp, up to 40TB
- Castor (T1D0)
 - 1 end-point `srm-v2-cms.cr.cnaf.infn.it`
 - 2 FE, 1 BE, 1 srm DB Oracle single instance
 - 17 disk-servers, 134 TB

CMS

- CMS adopted both castor srm v2 and StoRM end-points
 - T1D0 with CASTOR, T0D1 with StoRM/GPFS
- Main issues with gridftp layout
 - gridftp server late deployment → server overload
 - Started with 1 dedicated gridftp server (slc4) - 1st week
 - **Solution:** Ramped up to 6 gridftp servers (slc4) – 4th week
 - Limitation of the HW/GPFS infrastructure deployed
 - but exploited at nominal rate (useful feedback for further deployment)
 - Rate of 100/140 MB/s in writing
 - 200MB/s when writing and reading concurrently (import ~ 80-100 MB/s export ~ 100-120 MB/s)
 - gridftp block-size
 - More problematic for ATLAS (common issue)

CMS

- Main issues with CASTOR:
 - Farm activity peaks caused saturation of available slots on some disk-servers
 - Number of slots is limited in order not to have too gridftp connections
 - **Possible solution (under study):** subdivide the slots for:
 - Rfio connections (up to 100)
 - Gridftp connections (up to 10-30)
 - Tape families handled manually
 - Safe solution for the first test but not efficient
 - Functionally a success 😊 but cumbersome to manage
 - **Possible solution:** develop migration policies in order to automatise the migration of more tape families

CMS

- GPFS (common) issue:
 - Random access latency to software area
 - difficulties with CMSSW release installation
 - random high rate of CMS specific SAM test failures
 - 60% of availability at CNAF (weeks 1 and 2)
 - 100% of availability (weeks 3 and 4)
 - Again problems this week
- Solution:** currently migrated part of the VOs sw on another FS/HW


LHCB

- StoRM (D1T0)
 - 1 end-point srm-fe-lhcb.cr.cnaf.infn.it
 - 2 FE, 1 BE, 2 gridftp
- Castor (D0T1)
 - 1 end-point castorgrid.cr.cnaf.infn.it
 - 4 FE, 1 BE
 - 1 end-point srm-v2.cr.cnaf.infn.it
 - 2 FE, 1 BE, 1 srm-db oracle single instance
 - svcClass lhcb, 1 disk-server, 2.7TB
 - svcClass lhcb_raw, 2 disk-servers, 3.6TB

LHCb tasks

- Receive and reconstruct raw data 9% of 1 month of data taking
 - assuming a 50% machine cycle efficiency
 - Production of rDST data - T1D0
 - Use of Castor (SRM 2.2)
- Stripping of data
 - Input data: RAW & rDST - T1D0
 - Use of Castor (SRM 2.2)
 - Output data: DST - T1D1 and T0D1
 - Use of GPFS (Storm) and GPFS+TSM
- Distribution of DST data to all other T1s
 - Use of FTS - T0D1
 - Use of GPFS (Storm)

LHCb: T0→CNAF transfer




- FTS transfer on SRM 2.2
- The expected rate was achieved without any problems
 - Peak rate excess the requirements
- Low instability of SRM 2.2
 - Sometimes srm-v2 did not worked properly due to a known BE bug (occasional crashes, solved into the incoming release 1.3-27)

LHCb: reconstruction

- Submitted what is expected: 9% of the whole raw data
 - 77% is the success/submitted rate of reconstruction jobs
 - Main problem → GPFS: slow software area access, sometimes blocking
 - i.e., some jobs went in stall
 - broken switch one of the main cause
 - FIXED
- Robust solution:** currently migrated part of the VOs sw on another VO-independent filesystem
- The problem for LHCb appeared the 29th of May, before that the success/submitted efficiency was 93%

LHCb: stripping

- 4/6 Tier-1s run stripping
 - CNAF, PIC, RAL and GridKa



DST data produced at CNAF and distributed to other T1

- T1D1 was used to store the DST master copy
 - Successfully used GPFS+TSM to migrate data on tape
 - 4k files stored on tape, ~1TB
 - 0% failure rate

ALICE

- Castor (D0T1)
 - 1 end-point `castorgrid.cr.cnaf.infn.it`
 - 4 FE, 1 BE
 - 4 disk-server, 25.5TB

ALICE

- *alice* upgraded software on vobox, migrating from sl3 to sl4 and blocking its production activity of at least two weeks.
- Current problems are two:
 1. A thick latency period accessing to software area
 - e.g., *ls* operation needs about ten seconds
 2. Initial problems with a *srm v2* misconfiguration by the experiment

Questions for Developers

- StoRM
 - Monitoring Service
 - Administrative Interface
 - Improvements of FAQs
 - How to recognize StoRM problems from its logs
 - Defining a kind of map between errors and actions

CASTOR Experiences @ RAL

Shaun de Witt

CCRC08 Post Mortem 12 July
2008

ATLAS

- CASTOR slow down
 - Rm and put done requests became very slow
 - SRM reported 'too many threads busy with CASTOR'
- Needed database intervention
 - Would not have been solved quickly if it had not been seen at CERN week before.

LHCb

- Reconstruction jobs failed due to rfio problems
 - Caused core dumps in processing software
 - Reasons still unclear
- Worked around by copying data to job rather than opening file in CASTOR.

CMS

- User tape mounts brought production jobs to a halt
 - Jobs were replicating skimming
 - Now have limited user jobs on farm
 - No longer use 'fair share'
- Would have been desirable to be able to prioritise tape recalls.

General

- Tape servers were 'flaky'
 - They would stop in the middle of recalls and require manually stopping and restarting daemons
 - Seen by CMS and LHCb
- Reasons still not well understood
 - - we are running very old castor version (2.1.3) on tape servers
 - - hope upgrade will help
 - Problems seen elsewhere?