

INFN CNAF TIER1 Castor Experience

CERN 8 June 2006

Ricci Pier Paolo pierpaolo.ricci@cnafe.infn.it

castor.support@cnafe.infn.it

TIER1 CNAF Experience



- Hardware and software status of our castor v.1 and castor v.2 installations and management tools
- Experience on castor v.2
- Planning for the migration
- Considerations
- Conclusion

Menpower



- At present there are 3 people at TIER1 CNAF working (at administrator level) for our CASTOR installation and front-ends:
- Ricci Pier Paolo Staff** (50% also activity in SAN/NAS HA disk storage management and test, Oracle adm) pierpaolo.ricci@cnaif.infn.it
- Lore Giuseppe Contract** (50% also activity in ALICE exp. as Tier1 reference, SAN HA disk storage management and test, managing Grid frontend to our resources)
giuseppe.lore@cnaif.infn.it
- Also we have 1 CNAF FTE contract working with the development team at CERN (started March 2005)
Lopresti Giuseppe giuseppe.lopresti@cern.ch

We are heavily outnumbered. We absolutely need the direct help of Lopresti from Cern in administering, configuring and third level support of our installation (Castor v.2)

Hardware & Software



At present our CASTOR (1.7.1.5 and 2.0.4.0) system is:



1 STK L5500 SILOS
partitioned with 2 form-factor
slots

About 2000 slots LTO-2 form
About 3500 slots 9940B form



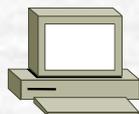
6 LTO-2 DRIVES with
2Gb/s FC interface

1300 LTO-2 TAPES
(200GByte) 250TByte



7 9940B DRIVES with
2Gb/s FC interface

1350 9940B TAPES
(200GByte)260TByte



Sun Blade v100 with 2
internal ide disks with
software raid-0 running
ACSL 7.0 (Single Point
of Failure)



13 Tapeservers, standard hardware is 1U Supermicro 3 GHz 2GB with 1 Qlogic 2300 F.C. HBA, STK CSC Development Toolkit rpm provided by CERN (with licence agreement with STK) ssi,tpdaemon and rtcpd.

All tapeserver have been re-installed with SL CERN 3.0.6, "quattorized" and all castor rpms upgraded to version 2.0.4.0

The 13 tapeservers are connected directly with the drive FC output:

DRIVE LTO-2 0,0,10,0 ->	tapesrv-0.cr.cnaf.infn.it
DRIVE LTO-2 0,0,10,1 ->	tapesrv-1.cr.cnaf.infn.it
DRIVE LTO-2 0,0,10,2 ->	tapesrv-2.cr.cnaf.infn.it
DRIVE LTO-2 0,0,10,3 ->	tapesrv-3.cr.cnaf.infn.it
DRIVE LTO-2 0,0,10,4 ->	tapesrv-4.cr.cnaf.infn.it
DRIVE LTO-2 0,0,10,5 ->	tapesrv-5.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,6 ->	tapesrv-6.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,7 ->	tapesrv-7.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,8 ->	tapesrv-8.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,9 ->	tapesrv-9.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,13->	tapesrv-10.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,14->	tapesrv-11.cr.cnaf.infn.it
DRIVE 9940B 0,0,10,15->	tapesrv-12.cr.cnaf.infn.it

In 2 years of activity we report that USING THE 9940B have drastically reduced the error rate (we report only 1-3% 9940 tape marked RDONLY due to SCSI error) and negligible hang problem

Hardware & Software



castor.cnaf.infn.it Central Machine 1 IBM x345 2U machine 2x3GHz Intel Xeon, raid1 with redundant power supply

O.S. Red Hat A.S. 3.0 Machine running all central CASTOR 1.7.1.5 services (Nsdaemon, vmgrdaemon, Cupvdaemon, vdqmdaemon, msgdaemon) and the ORACLE client for the central database

Installed by source, the central services will be migrated soon

castor-4.cnaf.infn.it ORACLE Machine 1 IBM x345

O.S. Red Hat A.S. 3.0 Machine running ORACLE DATABASE 9.i rel 2 for the Castor central daemons schemas (vmgr,ns,Cupv)

1 more x345 machine is in standby and is used for storing all the backup information of the ORACLE db (.exp .dbf) and can be used for replacing the above machines (castor and castor-4) if needed.

castor-1.cnaf.infn.it Monitoring Machine 1 DELL

1650 R.H 7.2 Machine running monitoring CASTOR service (Cmon daemon) NAGIOS central service for monitoring and notification.

Also contains the command rtstat e tpstat that are usually runned with the -S option over the tapeserver



Stagers with diskserver (v. 1.7.1.5): 1U

Supermicro 3 GHz 2GB with 1 Qlogic 2300 F.C. HBA accessing our SAN and running Cdbdaemon, stgdaemon and rfiod. 1 STAGER for each LHC Experiment and 2 GENERIC STAGERS installed by source.

disksrv-1.cnaf.infn.it ATLAS stager with 2TB direct connected

disksrv-2.cnaf.infn.it CMS stager with 3.2TB direct connected

disksrv-3.cnaf.infn.it LHCb stager with 3.2TB direct connected

disksrv-4.cnaf.infn.it ALICE stager with 3.2TB direct connected

disksrv-5.cnaf.infn.it TEST,PAMELA,ARGO stager

disksrv-6.cnaf.infn.it stager with 2TB locally (archive purpose LVD,alice TOF,CDF,VIRGO,AMS,BABAR, and other HEP experiment...)

Diskservers: 1U Supermicro 3 GHz 2GB with 1 Qlogic 2300 F.C. HBA accessing our SAN and running rfiod.



Storage Element front-end for CASTOR

castorgrid.cr.cnaf.infn.it (DNS alias load balaced over 4 machines for WAN gridftp)

sc.cr.cnaf.infn.it (DNS alias load balaced over 8 machines for SC WAN gridftp)

SRM1 is installed and in production

The access to the Castor system is

- 6) Grid using our SE frontends (from WAN)
- 7) Rfio using castor rpm and rfio commands installed on our WN and UI (from LAN)

Roughly **40%** (200TB / 500TB) of the total HSM space was effectively used by the experiments so far (3 years of official activity)

As TIER1 storage we offer “pure” disk as primary storage over SAN (preferred by the experiments) (Gridftp,nfs,xrootd,bbftp,GPFS)

Hardware & Software



CASTOR v.2 (2.0.4.0) servers (all the following servers runs SL CERN 3.0.6, "quattorized", castor installation using rpm)

☛ **castor-6** (1 IBM x345 2U server 2x3GHz Intel Xeon, raid1 system disks, with redundant power supply)

runs central stager services: STAGER + Request Handler + MigHunter + rtcplientd

☛ **castorlsf01** (HP Proliant DL360G4 1U server 2x3GHz Intel Xeon, raid1 system disks, with redundant power supply)

runs MASTER LSF server v. 6.1 (at present we run only a master instance of LSF for Castor v.2)

☛ **oracle01** (HP Proliant HA machine 2x3.6GHz raid 10 with redundant power supply)

runs the STAGER Database over Oracle 10.2.0.1.0

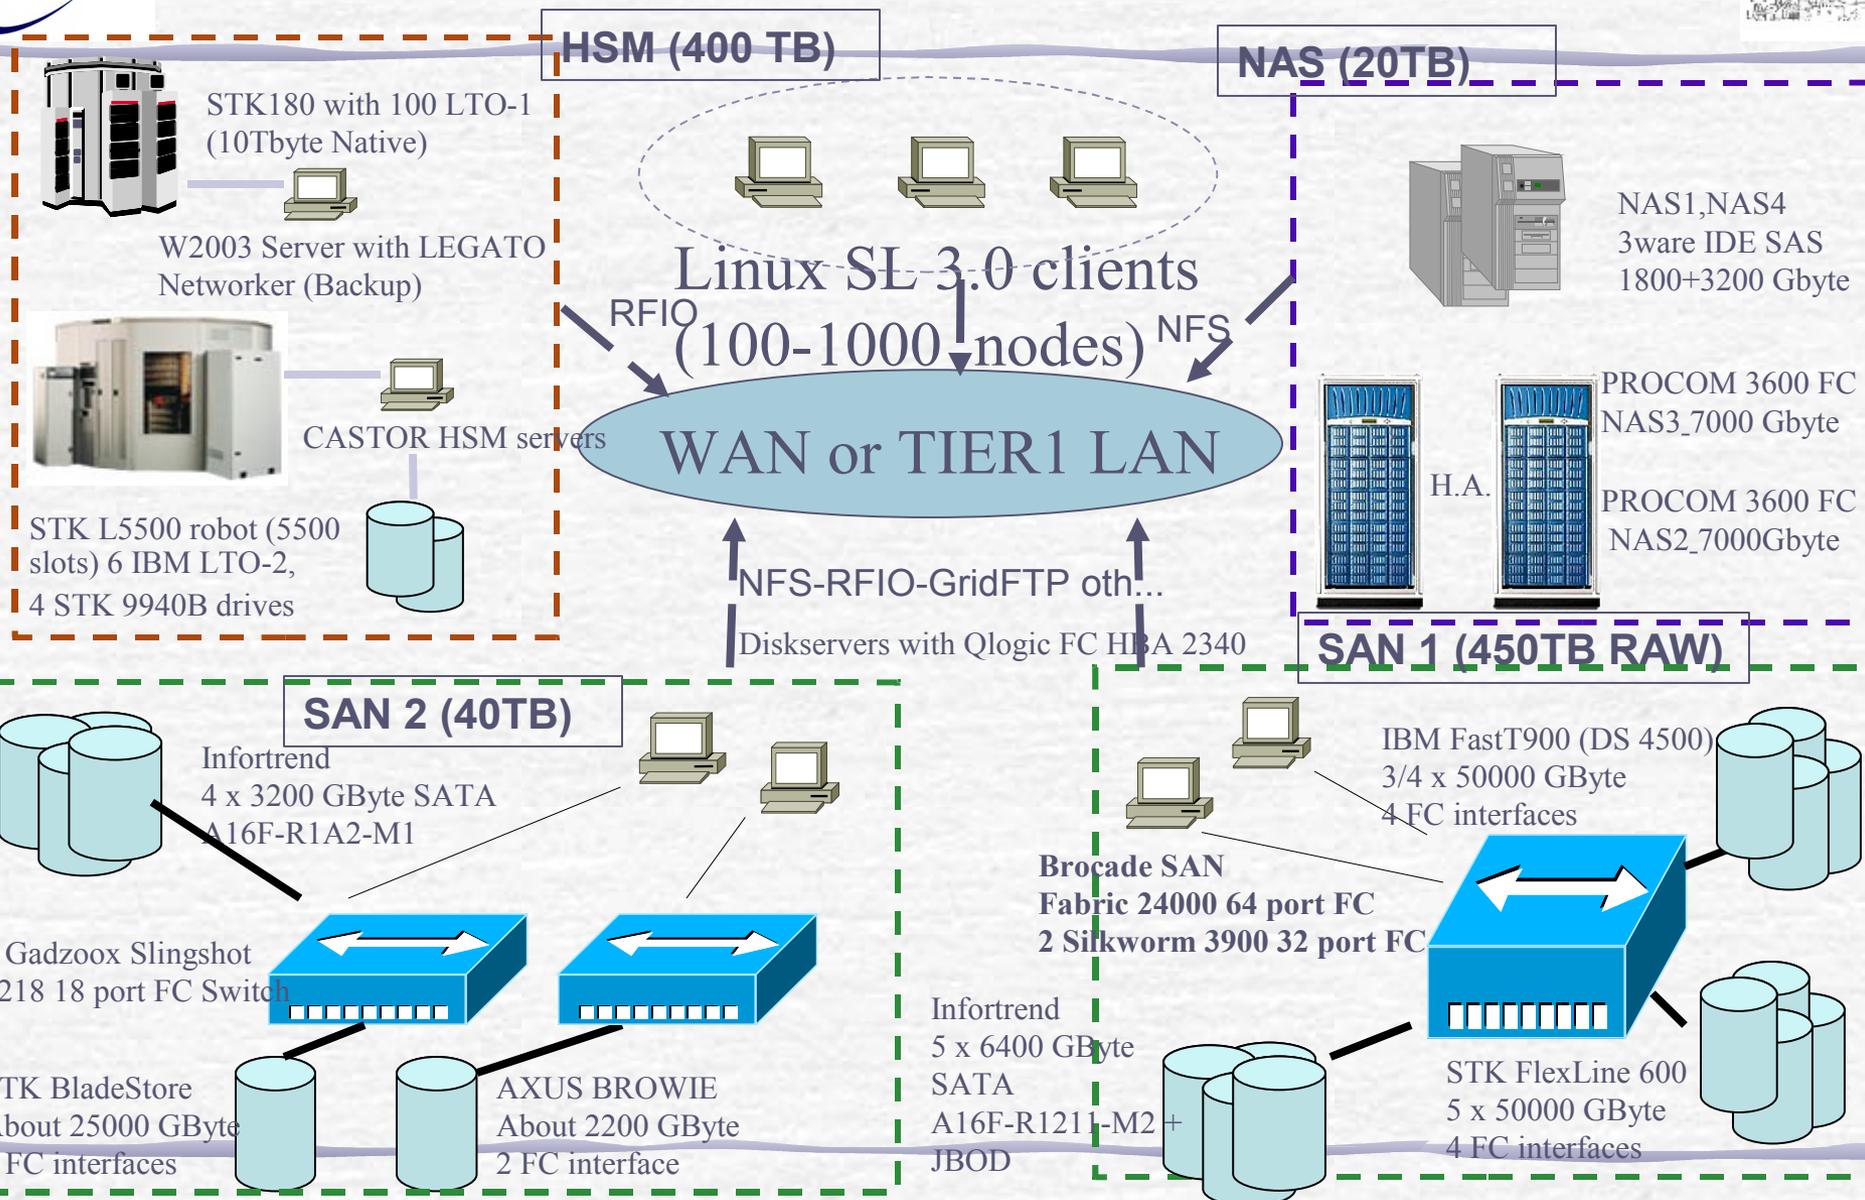


- ❏ **diskserv-san-13** (Supermicro 1U 3 GHz no hardware redundant) runs the services: DLF, DLF database, RMMASTER and EXPERT

- ❏ **castor-8** (1 IBM x346 2U server 2x3.6 GHz Intel Xeon, raid1 over 2 system disks and raid5 over 4 disks with redundant power supply) runs the new central services version 2.0.4-0 (nsdaemon,msgdaemon,vdqmserver,msgdaemon,Cupvdaemon)

The castor-8 machine will be the new central services machine. Some preliminary tests showed that the six castor v.1 stagers can use this machine without apparent major problems.

TIER1 CNAF Storage Overview





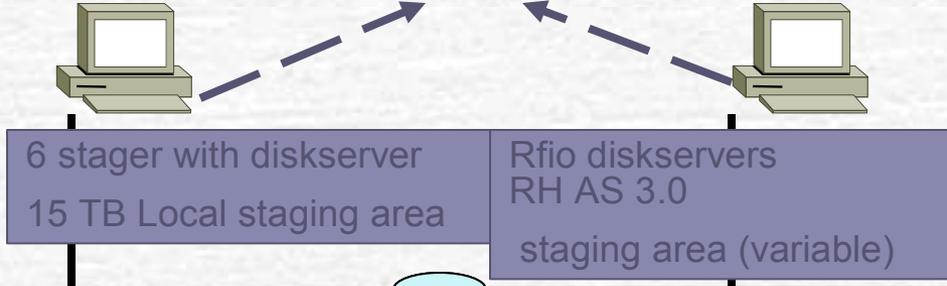
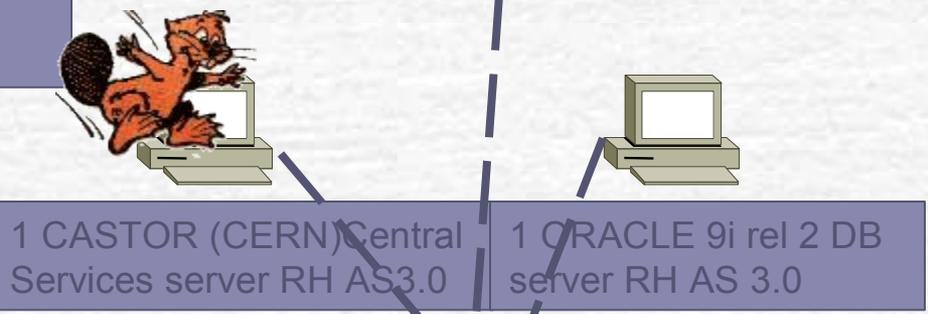
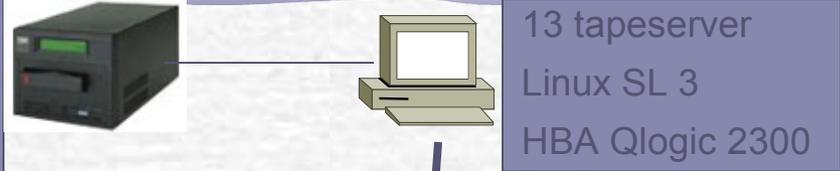
STK L5500 2000+3500
 6 drives LTO2 (20-30 MB/s)
 7 drives 9940B (25-30 MB/s)
 1300 LTO2 (200 GB native)
 1350 9940B (200 GB native)
 TOTAL CAPACITY with 200GB
250 TB LTO-2 (400TB)
260 TB 9940B (700TB)

Sun Blade v100 with 2 internal ide disks with software raid-1 running ACSLS 7.0 OS Solaris 9.0

CASTOR v.1

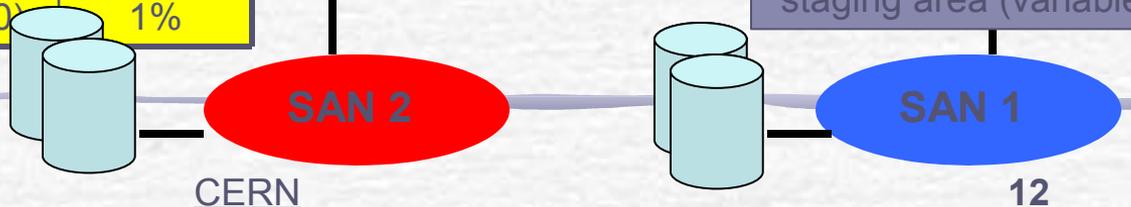


Point to Point FC 2Gb/s connections



EXPERIMENT	Stage (TB)	Tape (TB)	% ronly
ALICE	13	14(LTO-2)	8%
ATLAS	17	48(9940) 8 (LTO-2)	2% 30%
CMS	15	28(9940)	0%
LHCb	18	43(LTO-2)	10%
BABAR	4	20(LTO-2)	2%
(backup) VIRGO	1.5	5(LTO-2)	10%
CDF (backup)	1	9(LTO-2)	5%
AMS	3	5(9940)	0
ARGO+other	6	21(9940)	1%

Indicates Full redundancy FC 2Gb/s connections (dual controller HW and Qlogic SANsurfer Path Failover SW)



Monitoring and Notification



Nagios - Microsoft Internet Explorer

Indirizzo <http://castor-1.cnaf.infn.it/nagios/>

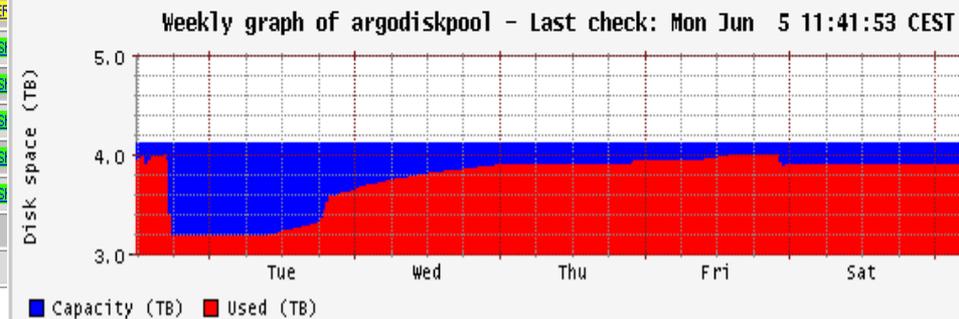
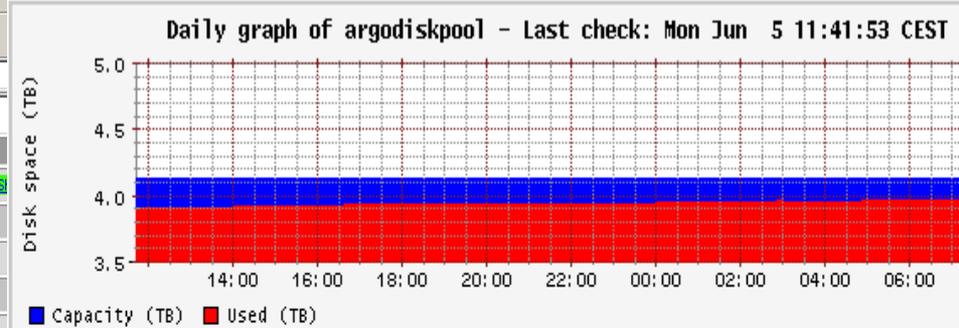
Nagios

- General
 - Home
 - Documentation
- Monitoring
 - Tactical Overview
 - Service Detail
 - Host Detail
 - Status Overview
 - Status Summary
 - Status Grid
 - Status Map
 - 3-D Status Map
 - Service Problems
 - Host Problems
 - Network Outages
 - Comments
 - Downtime
 - Process Info
 - Performance Info
 - Scheduling Queue
- Reporting
 - Trends
 - Availability
 - Alert Histogram
 - Alert History
 - Alert Summary
 - Notifications
 - Event Log

Host	Services
tapesrv-0	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-1	LOAD and NETWORK LOCAL DISK SPACE SSH TAPESERVER DAEMONS
tapesrv-10	LOAD and NETWORK LOCAL DISK SPACE SSH TAPESERVER DAEMONS
tapesrv-11	LOAD and NETWORK LOCAL DISK SPACE SSH TAPESERVER DAEMONS
tapesrv-12	LOAD and NETWORK LOCAL DISK SPACE SSH TAPESERVER DAEMONS
tapesrv-2	CPU HEALTH SM DISK HEALTH LOAD and NETWORK LOCAL DISK SPACE SSH TAPESER
tapesrv-3	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-4	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-5	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-6	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-7	CPU HEALTH SM CPU LOAD DISK HEALTH LOCAL DISK SPACE NETSTAT NETWORK SS
tapesrv-8	LOAD and NETWORK LOCAL DISK SPACE SSH
tapesrv-9	LOAD and NETWORK LOCAL DISK SPACE SSH

Central servers (central-servers)

Host	Services
castor	ALICE-LOG-TAPEPOOL AMS-TAPEPOOL ARGON-TAPEPOOL ATLAS-LOG-TAPEPOOL BABAR-TAPEPOOL CDC-TAPEPOOL CENTRAL DAEMONS CMS-LOG-TAPEPOOL CPU LOAD ITEAM-LOG-TAPEPOOL LHCb-LOG-TAPEPOOL LOCAL DISK SPACE MAGIC-TAPEPOOL NETSTAT NETWORK PAMELA-TAPEPOOL RAD EM xSerial SSH TAPE-DRIVES
castor-1	LOCAL DISK SPACE NETSTAT NETWORK SSH
castor-2	LOCAL DISK SPACE RAD EM xSerial SSH
castor-4	LOCAL DISK SPACE RAD EM xSerial SSH
castor-5	LOCAL DISK SPACE RAD EM xSerial SSH



MONITORING (Cmonitor)



Stage Pool

Real time
Disk to tape
streams
performance

Very useful for
tracing
real-time
bottleneck

Support for the
future
required!

Castor Monitoring

File Tape Activity Disk Status Servers

Monitoring Daemon Information

Monitoring daemon: castor-1.cnaf.infn.it:20001
Last update: 17/01/06 12:26:53

Drives currently down

Drive	Device Group Name
None	

Disk Volumes with more than 99% used:

Pool	Server	Name	Capacity
amsdiskpool@disksrv-6	disksrv-6.cnaf.infn.it	/area_staging/stage10	104861
amsdiskpool@disksrv-6	disksrv-6.cnaf.infn.it	/area_staging/stage11	104861
amsdiskpool@disksrv-6	disksrv-6.cnaf.infn.it	/area_staging/stage6	209714
argodiskpool@disksrv-5	disksrv-5.cnaf.infn.it	/mnt/disk2	9851
atlasdiskpool1@disksrv-1	disksrv-1.cnaf.infn.it	/bladestore1_sd1	980512
atlasdiskpool1@disksrv-1	disksrv-1.cnaf.infn.it	/bladestore1_sd2	980512

tapesrv-0 tapesrv-1 tapesrv-2 tapesrv-3 tapesrv-4 tapesrv-5 tapesrv-6 tapesrv-7 tapesrv-8 tapesrv-9

File

Server: tapesrv-6

EL0709.1 ACTIVE 99driv6@tapesrv-6 (994grp)
100% 1/1 00:00:07

Server: tapesrv-7

IDLE 99driv7@tapesrv-7 (994grp)
00:05:36

Server: tapesrv-8

EL0919 UNLOADING 99driv8@tapesrv-8 (994grp)
00:00:14

Server: tapesrv-9

IDLE 99driv9@tapesrv-9 (994grp)
00:15:01

disksrv-5.cnaf.infn.it/disksrv-5.cnaf.infn.it (castormn/te) → 40710 KB/s → EL0709.1

disksrv-1.cnaf.infn.it/disksrv-rfio-1.cnaf.infn.it (atlas-CHARF/zp) ← EL0919

Castor v.2 Experience



- ☞
December 2005: First servers installed with the direct help of Lopresti on-site. A single disk-only pool as test purpose
- ☞
January 2006: Problem with the two domains and the migration
- ☞
End of January 2006: Installed the first v. 2 tapeserver during the RAL External Operation Workshop, fixed some stager problem, first real tape migration
- ☞
February-March 2006: Installed new diskserver and experience over different file classes, problem with minor bugs (two domains) some fixes provided
- ☞
April 2006: SC4 failure due to LSF and nameserver compatibility and other minor bugs. Installation of a new machine (castor-8) with v.2 central services (nameserver) and upgrade of LSF, fix garbage collector problems
- ☞
May 2006: Re-run of SC4 over the new nameserver (OK)
- ☞
June 2006: Upgrade of all the tapeserver to latest version. Ready to migrate some of the LHC experiments to Castor v.2

Castor v.2 SC Experience



- r The castor v.2 stager and the necessary 2.0.4 nameserver on castor-8 has been used in preproduction during the Service Challenge rerun on May (after the problem during the official Service Challenge phase)
- r A relative good disk to disk bandwidth of 170MByte/s and a disk to tape bandwidth of 70MByte/s (with 5 dedicated drives) has been granted over a full week period.
- r We write an high quantity of data on tapes (about 6TByte/day) but we actually didn't test:
 - d) The access to the data in the staging area from our farming resource (test stress the staging area access)
 - e) The recalling system from tape with heavy requests of not-staged files in random order (tape stress of the stage-in procedure from tapes)

Migration Planning (stagers)



We have 3 hypothetical choices for migrating the six production stagers (and related staging area) to the castor v.2 stager

- 2) Smart Method: CERN could provide a script for directly converting the staging area from castor v.1 to castor v.2 renaming the directory and files hierarchy on the diskserver and adding the relative entry in the castor v.2 stager database. The diskserver are in such a way "converted" directly to castor v. 2.
- 3) Disk to disk Method: CERN could provide script for copying from the castor v.1 staging area to the castor v.2 stager without triggering a migration. We should provide new diskserver for castor v.2 during this phase with enough disk space for the "staging area" copy
- 4) Tape Method: The castor v.1 staging areas are dropped and new empty space is added in the castor v.2 stager. According to the experiment users usage a stage-in from tape of a large bunch of useful files is triggered to "populate" the castor v.2 stager

Since it will be very difficult for us to re-read a big number of file on our limited number of drive we cannot use the 3rd solution.

Migration Planning (services)



- ☛ The castor.cnaf.infn.it machine (running central services v. 1.7.1.5) will be dismissed
- ☛ A DNS alias "castor.cnaf.infn.it" will be created pointing to the castor-8.cr.cnaf.infn.it server
- ☛ The old castor.cnaf.infn.it machine will be reinstalled as castor-9.cr.cnaf.infn.it as a clone of castor-8 running in addition a vdqm replica
- ☛ All the native-replicable castor central services (nsdaemon,vmgrdaemon and Cupvdaemon) will run on both the machines
- ☛ We will obtain a DNS load balancing and High Availability installation of Castor for the central service (ns,vmgr,Cupv and vdqm using master and replica) (msgdaemon?)

Standard Diskservers Model



20 Diskservers with dual Qlogic FC HBA 2340
Sun Fire U20Z dual Opteron 2.6GHZ DDR
400MHz 4 x 1GB RAM SCSIU320 2 x 73 10K
10TB each diskserver

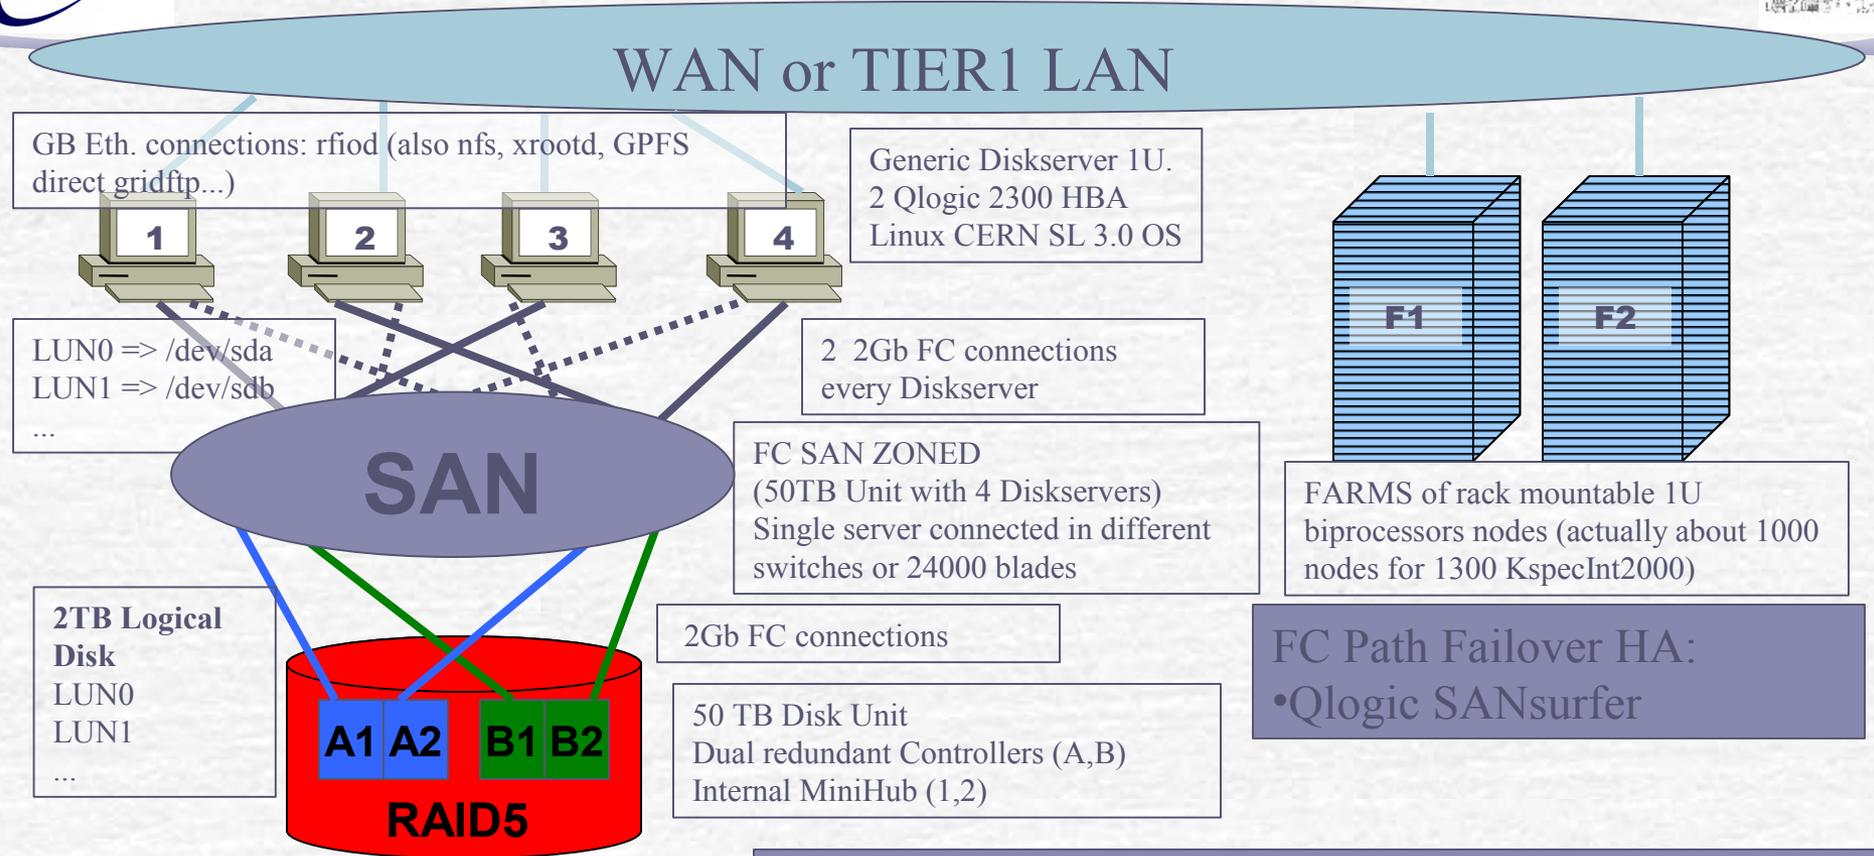


Performance on single volume not high (45w 35r) (but good 300-400MB/s aggregate)

Probably parallel I/O needed to optimize performance



Standard Diskservers Model



4 Diskservers every ~50TB
4 "High performance" servers
can reach the maximum
bandwidth of 200/400 MB/s
(peak of the storage subsystem)

Application HA:

- NFS server, with Red Hat Cluster AS 3.0
- GPFS with configuration NSD Primary Secondary
 /dev/sda Primary Diskserver 1; Secondary Diskserver2
 /dev/sdb Primary Diskserver 2; Secondary Diskserver3
- rfiid diskserver for Castor v.2 (to be implemented)**

Diskserver consideration



One of the major difference from the TIER0 model and our TIER1 is in the diskserver.

-  We have a small number of high performance diskserver with big quantities of disk storage connected (~12TByte). This will be the model also in the short-term period. A major failure in one diskserver could "cut down" the castor staging area of an essential disk area.
-  We have a SAN infrastructure that could provide everything needed for an High Availability System

The idea is that we would like to implement some sort of rfioid failover using Redhat Cluster (or equivalent) and registering virtual diskserver IP in the Stager and LSF catalog. We can do all the tests and the work on the cluster service but perheps some customization of Castor v.2 will be needed.

General considerations



- When we start 2 years ago the Castor External Collaboration (with our old Director Federico Ruggeri, PIC and CERN) the idea was that the Castor Development Team should take into account some specific customization needed in Tier1 sites (The original problem was the LTO-2 compatibility)
- To improve the Castor Development Team Tier1 CNAF agreed to provide manpower in terms of 1 FTE at CERN
- So far, the main activity requested from the Tier1 CNAF was support in term of current installation and help in upgrading to the Castor v.2
- After many years of production of Castor v.1 we become able to recover from mostly of the error conditions and contacts the Cern Support only as a "last resource"
- The situation is different in Castor v.2. From our point of view the software is more complicated and "centralized" and we lack in skills and tools to investigate and solve problems ourselves. Also the software itself is still in development...

General considerations (2)



- The Service Challenge 4 is just started and we still have the 4 LHC Experiments over Castor v.1 (only d-team is mapped on castor v.2). This is really a problem.

SO WHAT WE NEED TO MAKE CASTOR v.2 WORKING IN PRODUCTION? ("official requests")

- 3) The stager disk area migration should be concluded. Any solutions other than the "smart" method (diskserver direct conversion) could seriously influence the production and the current SC activity. We definitely ask a customization from CERN Development Team for this migration.
- 4) The castor central service migration and load balancing/high availability activity should be ended. This could be done by us probably with little support from CERN

General considerations (3)



- 1) After migrating the production over Castor v.2 the CERN support system should improve and grant a real-time direct remote support for the Tier1. Support team will have access to all the Castor, lfs and Oracle servers at Tier1 to speed up the support process. Due to the high "centralized" and complicated design of Castor v.2, and the lack of skills at the Cnaf Tier1 any problem could block the access to the whole Castor installation. If support is given only by e-mail and after many hours or days this could seriously affect the local production or analysis phases and translates in a very poor service.
- 2) The "direct" support could be done firstly by Lopresti since he can dedicate a fraction of time also to monitor and help administering our installation. But also all the other members of the development team should have the possibility to investigate and solve in real-time the high priority problems at the Tier1

General considerations (4)



The design of the Castor v.2 was supposed to overcome the old stager limits in such a way that a single stager instance could provide service for all the LHC experiments.

As Tier1 we don't want to find limits also in the Castor v.2 design that will force us to have multiple Lsf, stager, and Oracle instances to scale performance and capacity.

The idea is that, even when LHC and other “customers” will run at full service, the Tier1 expected data capacity and performance can be provided by a single LSF, stager and Oracle instance. The Development Team should take into account these considerations when optimizing the whole system evolution. We won't have the manpower and skill to manage a multiple installation of the castor v.2 services.

(one example: it seem that the new Oracle tablespace of the nameserver requires more space for the added ACL records. This translate in bigger datafile and perheps a single instance could not be enough in the next years. Is possible to prevent and optimize?)

Conclusion



- Our experience with Castor v.2 was overall good but we actually didn't test the heavy access to the staging area or tape recalling (perhaps the most critical parts?)
- The failure of the official phase of SC due to a LSF known bug suggests that the production installation needs an "expert" eye (at development level) for the administration, debugging and optimization of the system.
- Also the lack of a user-friendly command interfaces and documentation in general suggests that becoming a new administrator of Castor v. 2 won't be easy (Oracle query), and tracking/solving problem will be almost impossible without having a very good knowledge of the code itself and all the mechanisms involved.

Conclusion



- We agree to migrate all the stagers to Castor v.2 to help CERN with the support and for "standardizing" (quattor,rpms...) installations in the different Tiers.
- But, as part of the Castor External Collaboration, we ask that the Development Team should take into account all the needed and future CNAF Tier1 customizations.
- We ask that the scripts needed to optimize and speed up the migration process are developed by CERN.
- We ask also that when officially Castor v.2 at CNAF Tier1 will be in production a real-time first support will be granted at development level (with a contact in few hours in case of major blocking problems).
- Also the consideration about the peculiar CNAF diskservers model, the possibility of high availability rfioid and the scalability of Castor v.2 at Tier1 level should be taken into account (Castor must be designed to work also easily al Tier1 level!)