

Proposed setup for CASTOR 2 fresh installations

Sebastien Ponce

October 21, 2005

This document aims at proposing a default setup for a fresh CASTOR 2 installation. The goal is to try to homogenize the different sites in order to ease debugging.

This installation is meant to be a small scale setup (although hints are given on how to extend it). It should anyway be enough for TIER 1 activities.

The description of the tape part (including tapeservers, robotic, etc...) is not given since it is common with CASTOR 1. However, take care that the tape server software has to be upgraded to CASTOR 2 on all machines. This CASTOR 2 version is still fully compatible with old CASTOR 1 clients but is needed for the new CASTOR 2 clients.

1 Hardware setup

The machines needed for running CASTOR 2 are :

- Disk servers 2 to 4 nodes
- Central daemons 3 nodes
- Databases 2 nodes

The diskserver nodes are identical to the CASTOR 1 setup. They do not require anything specific.

Concerning the central servers (both databases and central daemons), the accent should be put on high available nodes rather than powerful ones. The nodes used at CERN have hot swappable redundant powersupplies and hot swappable, RAID 1 mirrored disks. They have 2GB of memory and for the ones running databases a swap partition of size equal to 3 times the memory size.

Otherwise, CPU and disk space are not really a problem so single CPU machines with 200GB of disk are perfectly fine.

All central daemons and database machines should be 32 bits architectures. Disk servers can be 32 or 64 bits.

2 Software installations outside CASTOR

2.1 OS

All central deamon machines should run a standard SLC3 linux distribution. The database machines should run a RHES 3 (RedHat Enterprise Server 3) distribution since ORACLE requires it.

2.2 ORACLE

CASTOR 2 was tested succesfully with both ORACLE 10.1.0.4 and ORACLE 10.2.0.1. However, we discovered that you should never use the client of one version with the server of another one. Although ORACLE says there is no problem, we experienced segmentation faults under high load in such cases.

We suggest that the stager database is called `castorstager` and has as main user `castor_stager`. Along the same line, we suggest that the DLF database is called `castordlf` and has a main user called `castor_dlf`.

2.3 LSF

LSF version should be greater or equal to 6.1. You will need to create a single queue called `castor`. You will also need to define the machine features correctly. Each machine should have as feature all the protocols it supports (e.g. `rftio` and `rotd`) as well as the diskpool it belongs too.

A separate guide concerning LSF for CASTOR 2 can be found at :

http://castor.web.cern.ch/castor/DOCUMENTATION/CASTOR2_DEPLOYMENT/FILES/LSF_%40_CASTOR2_guide.html

It gives a lot of examples on how to administrate LSF in CASTOR 2, including the way to define LSF queues and features. It also provides a description of LSF configuration files and how to edit them.

3 CASTOR 2

3.1 Components and their deployment

There are 8 main daemons in the CASTOR 2 installation : `scheduler`, `rtcpclientd`, `MigHunter`, `stager`, `request handler`, `dlf`, `rmmaster` and `expert daemon`. We grouped them in 6 logical units in the following way :

- `scheduler`
- `rtcpclientd` and `MigHunter`
- `stager`

- request handler
- dlf
- rmmaster and expert daemon

We recommend that each of this unit appears in the dns as if it would be a separate machine, with the following aliases : castorscheduler, castorrtcpcl, castorstager, castorrh, castordlf and castorrmmaster.

The mapping of the units/aliases to the real machines should then be done as follow :

machine 1 rtcpcl, MigHunter, stager and request handler

machine 2 dlf and rmmaster

machine 3 scheduler

Figure 1 summarizes the different components and their locations.

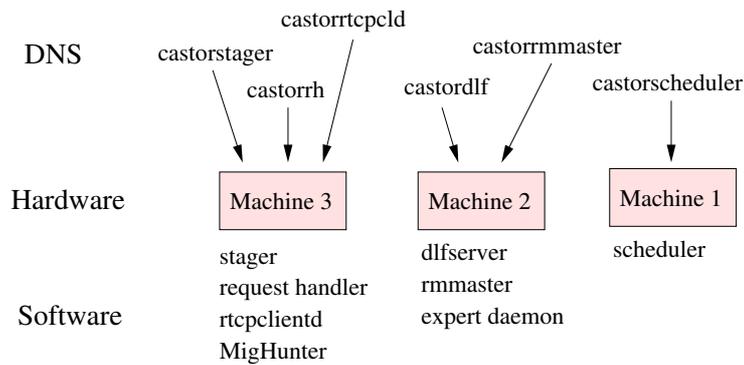


Figure 1: Setup of the CASTOR 2 machines

3.2 Software repository

You can find the CASTOR 2 RPMs on the `castor` web site :

http://savannah.cern.ch/files/index.php?group=castor&thread_max=1

RPMs are distributed for both i386 and ia64 architectures. However, the ia64 distribution only contains diskserver related packages.

The RPMS you should install on the machines are the following :

machine 1

castor-commands
castor-dbtools
castor-dlf-client
castor-hsmtools
castor-lib
castor-lib-oracle
castor-msg-client
castor-ns-client
castor-rfio-client
castor-rh-server
castor-rmmaster-client
castor-rtcopy-client
castor-rtcopy-clientserver
castor-rtcopy-messages
castor-scriptlets
castor-stager-client
castor-stager-server
castor-tape-client
castor-upv-client
castor-vdqm-client
castor-vmgr-client

machine 2

castor-dlf-client
castor-dlf-server
castor-expert-server
castor-hsmtools
castor-lib
castor-lib-oracle
castor-msg-client
castor-ns-client
castor-policies
castor-rfio-client
castor-rmmaster-client
castor-rmmaster-server
castor-rtcopy-client
castor-rtcopy-messages
castor-scriptlets
castor-stager-client
castor-tape-client
castor-upv-client
castor-vdqm-client
castor-vmgr-client

machine 3

castor-commands
castor-dbtools
castor-devel
castor-dlf-client
castor-hsmtools
castor-job
castor-lib
castor-lib-oracle
castor-lsf-plugin
castor-msg-client
castor-ns-client
castor-rfio-client
castor-rmmaster-client
castor-rtcopy-client
castor-rtcopy-messages
castor-scriptlets
castor-stager-client
castor-tape-client
castor-upv-client
castor-vdqm-client
castor-vmgr-client

On top of these packages, castor-sysconfig should be installed on all nodes and the files contained in it revisited for the local configuration.