IT Service Level Description

Service Name: **CASTOR Service**

Service Description

CASTOR ("CERN Advanced STORage manager") is a hierarchical storage management (HSM) system developed at CERN used to store physics production files and user files. Files can be stored, listed, retrieved and accessed in CASTOR using command line tools or applications built on top of the different data transfer protocols like RFIO (Remote File IO), ROOT libraries, GridFTP and XROOTD. CASTOR manages disk cache(s) and the data on tertiary storage or tapes. There are some hundreds of millions of files and 10s of petabytes (10^15 bytes) of data in CASTOR.

CASTOR provides a UNIX like directory hierarchy of file names, but not fileystem-like access - the CASTOR name space can viewed and manipulated only through CASTOR client commands and library calls.

CASTOR usage is covered by the CERN Computing rules (OC5 - as any other IT service), and in particular the File Service subsidiary rules. In particular, CASTOR is **not** suitable for confidential data.

The CASTOR service is split into several instances - one for each LHC experiment, and CASTORPUBLIC for everybody else. Space allocation and access control is handled by the experiment for their instances.

Group and Section Responsible

IT-FIO-FS

Clients

The clients are CERN or Grid users that need to persistently store large volumes of physics' data. The majority of users are from the large LHC experiments, although the service is also used by non-LHC experiments and individual users from other departments.

Criticality

The CASTOR service enables data taking, physics production (reconstruction, calibration, simulation) as well as physics analysis, which are core missions for CERN. If the service goes down, data recording will eventually stop, and physics production and analysis will be delayed.

Dependencies

The service relies on:

- computer centre infrastructure (e.g. power, cooling)
- network and DNS availability
- ORACLE (used to store all metadata)
- license servers administering the LSF license key (used for scheduling)
- management tools: Quattor/CDB, CDBSQL, Lemon

Service Support Level

Level 0: automatic: RAID and auto-migration from gradually failing machines

Level 1: The computers on which the service runs are under operator surveillance in the Computer Centre 24*365. They will be restarted on malfunction and on simple problems.

Level 2: A failure of the service which requires the intervention of an expert will be dealt with during normal working hours as soon as the alarm is raised. Outside working hours serious service problems are handled by the data services piquet – only certain classes of issues will trigger a callout of the piquet. Questions sent to helpdesk@cern.ch will be answered during working hours.

Service Reliability

The reliability of the service is measured by the following monitoring tools:

- Availability of public CASTOR resources: . http://cern.ch/sls/service.php?id=CASTOR
- Values of KPIs (Key Performance Indicators): http://cern.ch/sls/kpi.php.

Files availability: individual CASTOR files on disk can be temporary inaccessible because of hardware problems on the diskserver or filesystem they reside on. If there is a copy of the file on tape, the file will normally be recalled from there. If there is no copy on tape, the file can be unavailable for long periods (days) in case of sudden hardware failure. “Slow” or gradual hardware failures should be detected and files will be moved away from the failing machine.