CASTOR 2.1.8 Upgrade – RAL Tier1 Strategy

Matthew Viljoen 19/3/09

Introduction

CASTOR 2.1.8 has now been released by CERN, and CERN intend to upgrade their production instance to 2.1.8 as soon as possible before data taking.

2.1.8 contains many bug fixes to 2.1.7 (the version at RAL) and new functionality, notably:

- xrootd integration (lower latency and better scalability, suitable for data analysis)
- dedicated tool for draining disk servers
- improved user accounting, security, better repacking etc.

Further feature enhancements to version 2.1.8 are frozen with new CASTOR development work being channelled to version 2.1.9. This is likely to be released in 2010 (during data taking) when 2.1.7 will be rendered unsupported unless CERN change their support model.

This document briefly summarises the reasons for and against RAL upgrading to this version, including a timeline of preparations needed for such an upgrade.

Pros

- The additional functionality listed above (especially if we will be required to support data analysis during data taking).
- Being on the same version as CERN during data taking, facilitating problem solving between CERN and RAL.
- CERN will not have to change their software support model in supporting the version that RAL is running in production when 2.1.9 is released. There is some uncertainty about this, caused by contradictory information from CERN. Although CERN have said they will not leave us unsupported if we stay with 2.1.7, they have also said that they will not support more than two versions of CASTOR.

Cons

- At present the CASTOR service at RAL is stable compared to the past. 2.1.8 includes major changes from 2.1.7 and is (at the time of writing) untested in a production environment. However, we believe that none of the changes are critical to successful data taking at RAL.
- We have not yet been given clear signals that the major experiments wish to perform intensive data analysis at RAL, and recent performance (including the ATLAS 10m test) indicates that the CASTOR service at RAL is currently meeting most experiments' needs.
- No other Tier1 gave any firm indication at the CASTOR F2F that they would upgrade before data taking.

We will only consider upgrading to 2.1.8 if CERN's experience of running it in production indicates it is stable and it has passed our certification and stress testing satisfactorily, and there is enough time to adequately perform this certification and testing.

Due to the time constraints of when we need to be stable prior to data taking, it is prudent for us to plan an upgrade to 2.1.8 by starting the certification procedure as soon as the production release of 2.1.8 is available. This will include comprehensive testing of all configuration changes, the database schema upgrade procedure and stress testing.

In the event of multiple subsequent releases to address problems uncovered by running 2.1.8 in production at CERN (this was the case during the 2.1.7 release) it is very unlikely that we will have time to perform multiple certification cycles and upgrades before data taking.

Best-case Upgrade Timeline

10/3/09 - 2.1.8-6 released by CERN

+1wk

16/3/09 - CERN starts upgrade of their production CASTOR instances

(+6wks per subsequent release, if critical bugs are uncovered at CERN leading to further development and a new release). This includes time necessary to detect bugs, fix the bugs, certify the new release and upgrade, as well as time to check that the new release performs as expected in production.

16/3/09 – RAL starts testing discreet changes that can form part of an upgrade, e.g. (VDQM2, changes to LSF, updated tape servers). These changes may be installed even if we stay on 2.1.7. At the same time, we can prepare the certification and pre-production instances and databases for 2.1.8 certification and stress testing. This includes making the pre-production database more like the production database by consisting of multiple nodes.

+ 7wks

4/5/09 - RAL starts full certification on certification instance and stress testing on pre-production. This will be done on a pre-created snapshot copy of each production VO on the preproduction system. This copied database can then be discarded once we are confident that there will be no problems with the database side of the upgrade.

+3wks, based on previous certifications of CASTOR. This includes working with CERN and dealing with any RAL-specific issues with our configuration.

25/5/09 - RAL are in a position to upgrade pending successful certification

The above timeline implies the following:

- If we move CASTOR to R89, the latest date by when we can upgrade by is 15/6/09, Andrew Sansum's date by when the service must be stable. The above plan only gives us an upgrade

window of 3 weeks before this date. Should there be any unforeseen delays to the above schedule, we will not be able to upgrade. Such a short upgrade window will also prevent us staggering the upgrade over different CASTOR instances with breaks (e.g. by 2 weeks) between upgrades to monitor performance and, if necessary, resolve problems with CERN. It is for this reason that we do not believe an upgrade to 2.1.8 is possible if we move to R89.

- If we are *not* moving CASTOR to R89, we need to have a stable CASTOR service by 1/08/09. The above plan would then give us approx. 2 months of flexibility before this date - this should be sufficient time to perform a 2.1.8 upgrade, barring unforeseen problems.

Note that if we choose *not* to upgrade to 2.1.8 before data taking, it will be desirable to postpone this upgrade until data taking finishes in 2010 due to the complexity of the upgrade. If we are forced to perform such an upgrade during data taking, the likely downtime per instance will only be estimated once we have completed certification.

Conclusions

Regardless of when we upgrade to 2.1.8 it makes sense to start the preparation work listed above now and prepare our certification and pre-production instances for 2.1.8. This will put us in a good position for future upgrades.