

June 15, 2006

Report of the CASTOR Review

“To assess the readiness of the *CASTOR-2* service for LHC production”

**held at CERN
June 6th - 9th 2006**

Review Committee:

Miguel Branco - CERN PH/ATC
John Harvey - CERN PH/SFT (chair)
Don Petravick - FNAL CD
Shaun de Witt - RAL e-Science Centre

Table of Contents

1. Executive Summary.....	3
2. Introduction.....	5
3. Functionality of the software and priorities for further development... 	6
4. Maturity of the software and the software development process.....	8
5. Status of the service infrastructure including database servers.....	10
6. Status of Operational Procedures and Personnel	11
7. Performance and scalability of overall service.....	13
8. Support of CASTOR at other sites	14
9. Planning and Future Directions	15
10. Appendix 1: Charge to the Review Committee	17
11. Appendix 2: Schedule of Review Meeting	18

1. Executive Summary

The committee was impressed by the consistency of the project's vision of CASTOR and by the enthusiasm, commitment and conviction being shown to make it work. Overall the committee believes that the project is on the right track, that the structure of the software and the software engineering practices meet many needs of the project.

With a vigorous development team and careful planning, we can see development and deployment of CASTOR keeping up to meet the complete requirements of the LHC experiments. Having said that, we can foresee many years of continuous work in the system - it is by no means mature. We can also foresee periods of operational distress associated with this process.

After deliberations the committee arrived at a number of recommendations for addressing issues of most concern for having *CASTOR-2* ready for LHC production work. These recommendations are highlighted in boxes in the body of the report and are repeated here for convenience:

- The *CASTOR-2* architecture should continue to allow a pluggable scheduler without enforcing any strong binding to LSF, unless and until clear positive results about LSF scaling are obtained.
- If *xrootd* is to be supported by *CASTOR-2* as an internal protocol, the development, integration and support model must be agreed with *xrootd* developers.
- Whilst it is clearly too early to claim that the software and the software process are mature the committee is happy to endorse the process and technologies adopted which clearly compare favourably to those in use in other software projects in the laboratory and elsewhere.
- Consideration should be given to use of new configuration management and build tools in order to better automate the release process and to facilitate support of different platforms.
- The new procedure of installing new versions of the CASTOR client library, *libshift*, under the LCG External Library Service is meeting a strong requirement from the LHC experiments and should continue in the future.
- Pursuing common source code management for *rfio* and the nameserver with the LCG/DPM project seems to be highly desirable.
- The committee recommends having adequate dedicated hardware, possibly using Oracle certified hardware, to ensure long-term support from Oracle.
- The project should seriously consider that ORACLE tuning be decoupled from storage system operations so that the ORACLE DBA's can begin their optimization work outside of production systems. This seems to require a

significant investment, namely the construction of a test harness and ORACLE test stand.

- The committee recommends that the development and operations teams review the list of workarounds, involving replacement of palliatives with features in the CASTOR core.
- The committee recommends that the procedures and tools used to register bugs/problems, and their follow-up, are reviewed and documented to ensure all problems are dealt with in a timely manner.
- Understanding of the operations support effort requires a project plan and effort tracking. The committee recommends this be instantiated and maintained.
- The committee recommends that the project organizes a meeting with the 4 LHC experiments in order to reach agreement on a set of requirements and priorities for Service Monitoring Tools and a plan for their implementation.
- CASTOR-2 will provide a critical service and will need the attention of a full-time DBA in order to tune and monitor performance, manage backups and consult for developers.
- The committee recommends that the project should aim to validate experimentally the scalability of the performance of *CASTOR-2* to establish a performance envelope for the LHC experiments. These tests may impact on the general availability of *CASTOR-2* and therefore should be planned in conjunction with WLCG Data and Service Challenges.
- The committee recommends that the project should aim to clarify the policy for the support of CASTOR with its collaborators with a view to specifying the minimal effort and expertise required to operate a reliable CASTOR service and the level of support CERN can offer. The project should not take on any deployments at additional sites until it can clearly see sustaining them without sacrificing maturity of the software.
- The committee suggests that FIO group management considers providing effort outside of the project to assist with the maintenance of a credible WBS and schedule covering both development and operation tasks.
- The committee recommends that the project management organize a 'delta review' in about 6 months to review the project plan and progress made.

Finally the committee would like to thank all the participants to the review for all their hard preparation work, the excellent quality of the talks and for all the help given to the committee members.

2. Introduction

The purpose of this review was to assess the readiness of the *CASTOR-2* service for LHC production. The review took place at CERN from Tuesday June 6th, starting at 14.00, until Friday June 9th, finishing at 13.00. The meeting took place in 32/1-A24 using VRVS.

The members of the review committee were:

- Miguel Branco - CERN PH/ATC
- John Harvey - CERN PH/SFT (chair)
- Don Petravick - FNAL CD-COMP & COMP FABRIC
- Shaun de Witt - RAL e-Science Centre

The Charge to the review committee can be found in Appendix 1 and the schedule of the review in Appendix 2. The scope of the review was the core CASTOR project and did not extend to the SRM interface. The CASTOR team gave presentations on the history of the project, on the architecture and design of the software, and on operational and performance issues. There were also several talks from those responsible for running Database Services and the Data Recording Challenges. On Thursday afternoon the review heard reports from the 4 LHC experiments on their *CASTOR-2* experiences and a report from the CNAF Tier 1 centre on the experience of deploying and using CASTOR in an outside computer centre. All material shown and used at the review may be found at the following URL:

<http://indico.cern.ch/conferenceDisplay.py?confId=2916>

The committee would like to compliment all the speakers, and in particular the CASTOR team, for the considerable efforts made to prepare this review. The talks were very well prepared, comprehensive in scope and detail and the openness of the team to present and discuss issues of potential concern was very refreshing.

3. Functionality of the software and priorities for further development

CASTOR-2 is an upgrade project to *CASTOR-1*. It reuses some *CASTOR-1* components, primarily those that touch tape and the namespace. It provides, among other things a new internal stager system. Activities in the new components are driven by interacting with a relational database and significant logic is implemented in the form of database-stored procedures. One goal for the project is to make the daemons stateless, so that they may be stopped and started, and also provide that the system scales by making the daemons replicable.

The project has based its performance requirements on heuristics from its own experience, which may or may not provide a solid indication of real access patterns. One of the risks in providing a storage system or storage system software is that clear statements of performance metrics are unavailable to developers.

Because of the nature of employment contracts of its staff, the project needs to cope with turnover of its developers. The project has taken steps to address this by selecting some high-level OTS components, a relational database (ORACLE) and a commercial batch scheduler (LSF).

There is a tension between portability and high performance. As we have noted, storage system development has not been given performance requirements in a top down fashion. The committee feels that the CERN deployment requires a development track that allows the system to obtain the highest levels of performance, resilience and reliability possible. This requires the support of ORACLE specific-features in *CASTOR* code and a high level of tuning of the database. Although we have heard no objections to ORACLE at the existing Tier one sites, we would hope that the evident sound software practices will localize the dependencies on ORACLE and allow the support of other databases at some future date should this become a requirement.

The adoption of a DB-centric architecture is ambitious, and provides an elegant solution providing it performs well. This approach has also been adopted by the HPSS community. The stager has been completely re-engineered to be stateless and thread-safe. The stateless architecture is commonly used in high availability systems within industry. However, the stateless architecture does require a great deal of effort in testing and development to ensure state integrity within the database in the case of failover. It seems there is still some work to do on this with the database occasionally ending up in unexpected states. Nonetheless, the architectural choice seems reasonable, and the design around this architecture is comprehensive, well documented and maintainable.

LSF, as currently used and configured, is a bottleneck on CASTOR-2, particularly the response time when opening a file. Bulk processing allows for up to ~10 requests/s but it appears difficult to go much above without changing the CASTOR LSF plug-in.

The CASTOR-2 architecture should continue to allow a pluggable scheduler without enforcing any strong binding to LSF, unless and until clear positive results about LSF scaling are obtained.

The CASTOR team has a process for incorporating non-native IO protocols into *CASTOR* servers and one such protocol is *xrootd*. *Xrootd* integration is currently being prototyped. Operational development to support an *xrootd* deployment has not begun. Evaluation of the experiments' interests and motivations for using *xrootd* must take place.

If *xrootd* is to be supported by CASTOR-2 as an internal protocol, the development, integration and support model must be agreed with *xrootd* developers.

The project faces stress and demanding requirements that the LHC experiments can help mitigate. Experiments should be encouraged to use large files, to be sensitive to tape layouts in their data recall, and provide performance metrics when it is possible to do so.

While out of scope, a satisfactory SRM deployment that serves the needs of the experiments will not be usable until SRM V2 is used in the WLCG. The project is not able to maintain its SRM V1 interface, as the developer has left, and no one is available to work on 'old paradigm' software modules. While the situation appears to have improved with the ongoing development of SRM V2, the integration of CASTOR-2 with high-level Grid interfaces is a critical requirement for the LHC experiments.

4. Maturity of the software and the software development process

Opportunities have been taken to re-engineer code in C++ and using an object-oriented approach documented in UML. The stager has been completely redesigned in order to address LHC requirements and incorporates new strategies for data migration and recall between disk and tape media. A new version of the VDQM will shortly be deployed which follows the same design approach and there is the intention to replace other components such as *rmmaster* with a new stateless implementation. However, there is still reliance in some parts on old, poorly understood code that will be difficult to maintain in the future.

The project has followed best software engineering practices and the committee was pleased to see a willingness to leverage on proven software framework design principles (GAUDI), benefiting from the component model and use of pure abstract interfaces that allow for easy adaptation to different implementations. In particular the core framework provides a strong, stable, well-document code base allowing the project to adapt to changing requirements and integrate new developments. Anticipating staff turnover, the project has brought changes to its software engineering infrastructure to include automatic code generation tools. The project models its designs and software using UML. The project is able to automatically generate abstract classes, basic data base interface code, and serialization code from its UML descriptions. Another significant by-product has been the production of diagrams for documenting the design. It is noted that the code generation and UML diagrams are based on the Umbrello product, which itself incurs some overhead since new releases are often not backwards compatible and require modifications to the code generation tool.

<p>Whilst it is clearly too early to claim that the software and the software process are mature the committee is happy to endorse the process and technologies adopted which clearly compare favourably to those in use in other software projects in the laboratory and elsewhere.</p>
--

The configuration files have evolved over a long period and have become difficult to maintain. This has been recognised by the development team and deserves to be given due attention. The project should anticipate requests for support of the client library on platforms other than Linux, in particular MacOS. The committee was pleased to see that porting to Windows is being done as a voluntary external contribution but also noted that there is no long-term commitment to continue this.

Consideration should be given to use of new configuration management and build tools in order to automate the release process and to facilitate support of different platforms.

The experiments have raised important issues concerning the deployment of new releases of the CASTOR client library, *libshift*. The software that is being built gets linked with the version of the library that happens to be in the node where the build is done and this software may not run in the nodes that have other versions installed. The experiments therefore request to have control on the version that is used for building the software. As such the wish is to treat the CASTOR library in the same way as any other external library. After discussion with the LCG Architect's Forum the decision was taken to install *libshift* in the External Software Service area managed by the LCG Applications Area along with other external libraries and this has now been done.

The new procedure of installing new versions of the CASTOR client library, *libshift*, under the LCG External Library Service is meeting a strong requirement from the LHC experiments and should continue in the future.

The project makes use of the DPM code base for development of the Nameserver although there is no common DPM/CASTOR-2 CVS repository yet. Two implementations of the *rfio* library also exist and this is causing some confusion for users.

Pursuing common source code management for *rfio* and the nameserver with the LCG/DPM project seems to be highly desirable.

5. Status of the service infrastructure including database servers

Appropriate database server platforms, with UPS or secure power supply, are required for reliable and scalable production CASTOR instantiation. Consideration should be given to platforms that are supported by ORACLE and the use of Oracle DataGuard should be considered. The committee understands the plan to deploy separate instances of CASTOR, which, given the maturity of the software, seems prudent.

The core of CASTOR-2 being stateless, and plans to reengineer some CASTOR-1 components to the same architecture, relies on a resilient database on resilient hardware. Corruptions to the database, or database failure without redundancy, would be an almost worst-case scenario. Good performance and reliability of the hardware (disks and tape servers) was reported. Plans for short-term improvements to the hardware infrastructure should minimize recent disturbances to the infrastructure. However load balancing and prioritized scheduling still needs to be worked on with high priority in order to ensure efficient and smooth operating conditions.

The committee recommends having adequate dedicated hardware, possibly using Oracle certificate hardware to ensure long-term support from Oracle.

The project should seriously consider that ORACLE tuning be decoupled from storage system operations so that the ORACLE DBA's can begin their optimization work outside of production systems. This seems to require a significant investment, namely the construction of a test harness and ORACLE test stand.

6. Status of Operational Procedures and Personnel

The load on the operations group has been particularly heavy and significant effort has gone into finding workarounds to problems whenever feasible rather than the core development group fixing bugs. The priority of the development team has been given to meeting requests of the LHC experiments in order to ensure their timely migration to using *CASTOR-2*. Now that this has been achieved it is expected that increased priority will now be given to attending to operational issues. The proposed 3 level operational support structure should ensure that appropriate levels of effort and expertise can be leveraged and an efficient use of staff achieved.

The committee recommends that the development and operations teams review the list of workarounds, involving replacement of palliatives with features in the *CASTOR* core.

The ability to detect report and manage the repair of operational problems is clearly of high importance. Several different mechanisms are in common usage (mailing lists, Remedy, Savannah) with the consequence that problems may not reach the appropriate expert and therefore be left untreated. The follow-up on progress of fixing bugs is equally important and consideration should be given to establishing procedures to ensure this happens.

The committee recommends that the procedures and tools used to register bugs/problems, and their follow-up, are reviewed and documented to ensure all problems are dealt with in a timely manner.

Much recent work has been done to improve the internal monitoring with the system. This is essential work, because deployed systems deal with diverse disk and tape technologies, and have a rich server structure. Much of the work seems recent and more work needs to occur to mature the system. While it seems *Quattor* and *Lemon* are an aid to *CASTOR* operators and maintainers, it seems there are still many tasks which must be done by hand, and which are non-trivial and possibly prone to human error. While concentrating on the implementation of robust core functionality should reduce the software maintenance burden, some effort will in the future need to go into codifying fault recovery scripts.

However, given the relative immaturity of *CASTOR-2* deployment, the support effort does not look too excessive, and will likely reduce in the next 6-18 months. This is a real testament to the skill of the development team and to the operations team, which has had to develop workarounds and fix faults to

keep the CASTOR-2 systems running smoothly.

Understanding of the operations support effort requires a project plan and effort tracking. The committee recommends this be instantiated and maintained.

Deployment of service-level monitoring for many aspects of the service has not begun, but is envisaged. Other aspects of monitoring where further development is required include display of the plots and presentation of data to experiments. All LHC experiments attach high importance to the provision of adequate Service Monitoring Tools with a level of functionality comparable to those used for CASTOR-1. Required information includes a more formal interface mechanism for ancillary information, for example tape maps that allow for optimized recall of data, and information that will allow problem resolution for the services relying on CASTOR. e.g. what transfers are pending, which are active, etc. At the same time the experiments indicated some differences in the information and interfaces they expect to use.

The committee recommends that the project organizes a meeting with the 4 LHC experiments in order to reach agreement on a set of requirements and priorities for Service Monitoring Tools and a plan for their implementation.

7. Performance and scalability of overall service

One risk in developing storage systems is the lack of firm performance requirements from the experiments. Heuristically, it would seem the rate of requests being served is a worry. Also, the ALICE-quoted figure of 8 seconds to open one file seems to not be a credible summary comment about CASTOR performance and should be further investigated for recently deployed versions of CASTOR-2.

Concerns about LSF were addressed earlier in this report, with the recommendation to treat it as a plug-in until its scalability is understood.

LSF aside, the intrinsic CASTOR scaling limit is the performance of the database. The work to scale CASTOR-2 has benefited from a professional ORACLE DBA. The project has not investigated the scalability of the *CASTOR-1* nameserver. The project knows of scaling limits in the ORACLE based logger, DLF. Most of the *CASTOR-1* components seem fairly robust, although because of their stateful nature when they do fail it makes it hard to recover fully. The *CASTOR-2* core is just now being used in production by the experiments, and much will be learned in the next few months.

CASTOR-2 will provide a critical service and will need the attention of a full-time DBA in order to tune and monitor performance, manage backups and consult for developers.

Evidence was shown that the fundamental limitations of the *CASTOR-1* implementation (e.g. stager unresponsive above 200k disk-resident files, lack of request scheduling) have been successfully overcome. Required performance has been demonstrated whilst utilizing relatively low level of resources (~20% CPU load). In addition the RDBMS is used for recording all transactions and allows for error recovery.

The committee recommends that the project should aim to validate experimentally the scalability of the performance of *CASTOR-2* to establish a performance envelope for the LHC experiments. These tests may impact on the general availability of *CASTOR-2* and therefore should be planned in conjunction with WLCG Data and Service Challenges.

8. Support of CASTOR at other sites

The project has been open to new collaborators from existing CASTOR sites with a view to CASTOR being used at Tier 1 centres. These people are encouraged to work with the CASTOR team for a period to gain experience and to contribute to the development effort. The model for sustaining this collaboration in the longer term envisages a community effort providing mutual support using an external operations mailing list. Such a scheme requires that each centre has sufficient operation support to sustain a reliable service that satisfies the centre's MoU commitments.

Several Tier-1 sites are now relying on CASTOR. It appears that these sites contribute development effort back into the project and in return the project has to support deployments at its collaborator's sites. The project is stressed by the deployments - for example there seems to be an expectation that the development team provide "level three" support to all sites.

The concept of external sites contributing to CASTOR development, enabling local expertise at remote sites, is a good model. However, it seems the external sites get out what they put in. The more effort they can contribute, the better they will understand *CASTOR* and become self-supporting.

<p>The committee recommends that the project should aim to clarify the policy for the support of CASTOR with its collaborators with a view to specifying the minimal effort and expertise required to operate a reliable CASTOR service and the level of support CERN can offer. The project should not take on any deployments at additional sites until it can clearly see sustaining them without sacrificing maturity of the software.</p>
--

9. Planning and Future Directions

The project has recently undergone a difficult transition with several of the most experienced people, both developers and operations staff, leaving to take up different responsibilities in the department or new appointments outside CERN. The loss of effort has been compensated by the recruitment of replacements who are currently gaining experience. This transition has coincided with the rollout of *CASTOR-2* and important, though non-vital, tasks have been sidelined as a result. However we were very impressed with the competence, commitment and team spirit of this newly formed team. Separation of development and operation roles has been good and ensures that healthy attention is given to providing a reliable and robust service.

Only recently has the project made a WBS and schedule for managing development tasks and therefore the schedule has not existed long enough to establish its predictive power. The committee feels that this work should be maintained and sustained and even extended to cover operational aspects of the work programme. The committee would have liked to see what impact the plan has on its tasks and high-level milestones, such as those visible to the WLCG. The project should be encouraged to develop accurate estimates for its tasks. Project management should evaluate these and provide feedback.

The committee suggests that FIO group management considers providing effort outside of the project to assist with the maintenance of a credible WBS and schedule covering both development and operation tasks.

It is the opinion of the committee that CERN's readiness is best handled by continuing and expanding the CASTOR project. The short-term project plan presented seemed reasonable. Providing for an ORACLE DBA seems important. Elements that were out of scope in this review, such as the SRM interface, also need to be completed and placed into production.

It is also noted that the effort profile looks marginal. Too many unexpected pressures on the development team, or loss of key staff, will lead to slippage, and therefore additional requirements should be reduced at this critical phase. If more effort could be leveraged out of external institutes, this might help and may also benefit the external sites themselves. The resources required for convergence are best assessed by planning and measuring the resources consumed over the next six months.

The project plan should be inspected to ensure adequate plans exist for the following areas:

(a) Complete all outstanding work on current *CASTOR-2* components paying particular attention to aspects related to robustness and fault tolerance.

Report of CASTOR Review 2006

(b) Identify and reduce current performance bottlenecks

(c) Look at re-engineering stateful CASTOR-1 components, particularly any which are unstable or do not seem to scale to the needs of LHC

The committee recommends that the project management organize a 'delta review' in about 6 months to review the project plan and progress made.

10. Appendix 1: Charge to the Review Committee

Assess the readiness of the *CASTOR-2* service for LHC production

Particular points to be addressed:

- Does the software have the required functionality? If not, is the functionality at least adequate? And what are the priorities for development?
- How mature/stable is the software? And the software development process?
- Is the service infrastructure, and particularly the database server hardware, appropriate?
- Are the necessary operational procedures in place (with adequate trained personnel) to ensure the MoU service reliability commitments can be met?
- Does the overall service have the required performance for the Tier0? For the CAF? What aspects of the software or system limit performance? Are these performance limits a concern?

Other issues

- Is the current policy for the support of *CASTOR* at other sites sustainable?

Future directions

If the *CASTOR* service is not demonstrably ready for LHC production, what are the areas to address in priority order? Or would the effort be better spent implementing an alternative solution and, if so, which?

11. Appendix 2: Schedule of Review Meeting

Tuesday 06 June 2006

Session 1 : Project Overview

- 14:00 Introduction and Overview (Olof Barring)
- 16:00 Architecture Overview (Giuseppe Lo Presti)
- 17:00 Planning & Resources (German Cancio Melia)

Wednesday 07 June 2006

Session 2: Description of Components

- 09:00 Functional description and status(Giuseppe Lo Presti/Olof Barring)
- 10:15 Technical design
 - Technology choices (DB, code generation, UML, .) (Sebastien Ponce)
 - CASTOR core framework (Giuseppe Lo Presti)
- 11:30 Protocols, clients and APIs (Sebastien Ponce)

Session 3: Software Process

- 14:00 Project Management Summmary (German Cancio Melia)
- 14:20 Requirements, intra-project communication & software repositories (Sebastien Ponce)
- 15:00 Configuration management, build & testing procedures (Giuseppe Lo Presti)
- 15:40 Dissemination & User Feedback (Jan van Eldik)

Session 4: Service Infrastructure

- 16:30 Hardware requirements for a *CASTOR* instance (Jan van Eldik)
- 17:30 Considerations for Database Servers (Eric Grancher)

- 19:00 Review Dinner

Report of CASTOR Review 2006

Thursday 08 June 2006

Session 5: Deployment and Operation

09:00 Operational Experiences (Jan van Eldik)

09:30 *CASTOR-1* Operations Infrastructure (Olof Baring for Tony Osborne)

10:00 *CASTOR-2* Operations Infrastructure & Integration (Miguel Marques Coelho dos Santos)

11:00 DB Server performance (Eric Grancher)

11:30 System Performance & Reliability (Bernd Panzer-Steindel)

Session 6: Reports from *CASTOR* Clients : CNAF and 4 LHC Experiments

14:00 CNAF experiences (Pier Paolo Ricci)

14:40 ALICE Experiences (Latchezar Betev)

15:20 ATLAS Experiences

16:30 CMS Experiences (Nick Sinanis)

17:10 LHCb Experiences (Andrew Smith)

Friday 09 June 2006

Session 7: Closeout

09:00 Meeting of Referees (closed)

Time for committee to prepare a statement of its findings

11:00 Findings of the Review Committee