

Title: Consistency between grid storage elements and file catalogue for the LHCb experiment

Overview:

DIRAC is a framework developed to provide a complete solution for using distributed computing resources. DIRAC has been developed in a very generic way that have made it suitable for serving many VOs. The LHCbDIRAC framework is the DIRAC extension specific to LHCb, one of the four experiments operating at the Large Hadron Collider of CERN, and is where the particular features requested by the LHCb community are implemented. LHCbDIRAC is organized in several systems providing all functionality for using a distributed computing infrastructure, including a data management system, which implements an interface to the underlying grid middleware and performs all operations of data transfer, registration and removal. In this paper we shall describe a new component of the LHCbDIRAC data management system which checks the consistency between the data stored on the the grid storage elements and the central file catalogue of the experiment.

Description:

In the distributed computing model of WLCG Grid Storage Elements (SE) are by construction completely decoupled from the File Catalogs (FC) where the experiment's files are registered. On the basis of the experience of managing large volumes of data in such an environment, inconsistencies have often happened either causing a waste of disk space, in case the data were deleted from the FC, but still physically on the SE, or serious operational problems in the opposite case, when some data registered in the FC was not found on the SE. Therefore, the LHCbDirac data management system has been equipped with a new dedicated service to ensure the consistency of the data stored on the SEs with the information reported in the FCs implementing systematic checks. The service relies on information provided by the sites who should make available to the experiment a full dump of their SEs on weekly or monthly basis. The objective is to spot any inconsistency above a certain threshold, that cannot only be due to the expected latency between the storage dump creation and the checks' execution, and in such case try and identify the problematic data.

In this talk we shall present the definition of a common format and procedure to produce the storage dumps that has been coordinated with the other LHC experiments in order to provide a solution as generic as possible that can suit all LHC experiments and will reduce the effort for the sites who are asked to provide such data. We will also present the LHCb specific implementation, which is in production since September 2011, for checking the consistency between SEs and FC and discuss the results.

Impact:

The benefit that the new component brings to the LHCb experiment is to reduce the amount of data that is physically stored on grid SEs, whereas it is not registered in the experiment central catalogue, thus avoiding any waste of disk space. Moreover, the analysis of the results obtained so far has shown a very limited level of inconsistency between grid SEs and the central FC, proving that LHCbDIRAC tools handle in a consistent way both the operation of uploading data to a SE and registering in the catalogue, and the subsequent removal of files from SE and catalogue. Finally, also some grid sites have benefit from this system, as it allowed to spot some internal inconsistency of their storage system.

Conclusions:

A new component for checking the consistency between grid SEs and the LHCb central FC has been implemented in LHCbDIRAC data management system. The checks are performed centrally by an LHCbDIRAC agent on the basis of storage dumps provided by sites, and produced with formats and tools common also to other LHC's experiments, in order to minimize the effort from the site's side. The main benefit of this system is to verify that all data stored on grid SEs are actually registered in the FC, thus avoiding any waste of disk space at sites. An additional result is to spot possible internal inconsistencies in the sites' storage systems.

Track: Operational services and infrastructure

Authors: Vincent Bernardoff, David Bouvet, Marco Cattaneo, Philippe Charpentier, Krzysztof Ciba, Peter Clarke, Joel Closier, Ricardo Graciani, Elisa Lanciotti, Zoltan Mathe, Victor Mendez, Raja Nandakumar, Daniela Remenska, Stefan Roiser, Vladimir Romanovski, Federico Stagni, Andrei Tsaregorodtsev, Mario Ubeda Garcia, Alexey Zhelezov, Natalia Ratnikova, Nicolo' Magini, Cedric Serfon

-- ElisaLanciotti - 05-Dec-2011

---

This topic: LCG > LHCbConsistencyChecksEGICF2012Abstract

Topic revision: r2 - 2011-12-06 - unknown



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors. Ideas, requests, problems regarding TWiki? Send feedback