

Table of Contents

Usage of Tier1-BUFFER storage elements.....	1
Introduction.....	1
Tier1-BUFFER storage elements.....	1
Dataflow.....	1
Using Tier1-BUFFER for RAW reprocessing.....	1

Usage of Tier1-BUFFER storage elements

Introduction

In order to simplify the operations for processing productions that require access to data that is normally not resident on disk, we propose to use a temporary disk buffer space that needs to be cleaned up when the files are no longer needed.

Tier1-BUFFER storage elements

These SEs are located in the LHCb-Disk SRM space token, with a specific SPath obtained by appending /lhcb/buffer to the usual SPath. At CERN it is proposed to locate the CERN-BUFFER storage element on the EOS disk-only service rather than in Castor. Since this SE cannot be implemented and tested before several weeks, we have set up a temporary CERN-CASTORBUFFER that makes use of LHCb-Disk.

Dataflow

The reconstruction (or reprocessing) jobs will upload the resulting FULL.DST file (replacing the SDST) to their associated Tier1-BUFFER storage, making it available for any stripping production(s).

A standing DM transformation will replicate these datasets to the corresponding Tier1-RDST. This requires a dedicated plugin in the TransformationSystem (ReplicateToLocalSE that replicates to an SE at the same site where the files is, selected from a list of SEs). This plugin is tested and committed.

```
dirac-dms-add-replication --BK /LHCb/Collision12//RealData/Reco14//FULL.DST --Plugin ReplicateToL
```

A series of stripping productions will be launched on this (re)processing, each corresponding to a specific Processing-Pass (e.g. Stripping10, CaloFemtoDST,...). The corresponding jobs will be directed to sites associated to the Tier1-BUFFER SEs. If we want to include Tier2s for stripping, we should set the number of *Merge* jobs to 0 at those sites, as Tier1-BUFFER should also be used for uploading unmerged files.

Another DM transformation will be launched that will wait until the corresponding files have been processed by all processing passes. This is achieved by a new plugin DeleteReplicasWhenProcessed that accepts a series of parameters:

- SEs where to remove the replicas from (default Tier1BUFFER)
- List of processing passes to be watched
- A period at which the plugin should be executed in order to group files (default 6 hours).
- A cache life time for caching the list of productions corresponding to the required processing passes (default 24 hours)

```
dirac-dms-add-replication --BK /LHCb/Collision12//ReadData/Reco14//FULL.DST --Plugin DeleteReplic
--ProcessingPasses Stripping20,CaloFemtoDST --FromSEs Tier1-BUFFER -
```

As soon as a FULL.DST file will be found that has been completely processed, a removeReplica task will be created for removing the Tier1-BUFFER replica. There is no need to launch another removal transformation as long as the list of processing passes doesn't change, as the plugin will look every day for the list of associated productions.

Using Tier1-BUFFER for RAW reprocessing

UsingBufferSEs < LHCb < TWiki

With this DeleteReplicasWhenProcessed plugin, one can consider also using it for RAW data reprocessing. There is no need for prompt reconstruction as the files are already staged (although one could also use it in uploading RAW from the pit to CERN-BUFFER). The data flow could be the following:

- Once a file is staged (without need for pinning, or only for an hour), the stager sets a replication task to the relevant Tier1-BUFFER. The replica in Tier1BUFFER would then preferentially be used by the job when it starts. The cache handling is then much easier as we should only avoid to fill up the disk buffer. If a job starts before the file is replicated to the BUFFER, it will pick up the disk cache replica and the replication will proceed as well as the subsequent removal.
- A removal transformation waits for the RAW files to be processed for deleting them from the Tier1-BUFFER, using the DeleteReplicasWhenProcessed plugin, with a single (or more) processing passes (in case for example we run prompt reconstruction and reprocessing at the same time)

This data flow goes in the direction proposed by the DMS/SM TEG that jobs only access data from disk-only SEs and T1D0 SEs are only used for archival.

There is a modification required at the stager level though. In order to make it flexible, there could be a VO-dependent plugin called whenever files are staged, of which a DIRAC default would do nothing and the LHCb one would set this replication request.

-- PhilippeCharpentier - Last update 21-Jun-2012

This topic: LHCb > UsingBufferSEs

Topic revision: r5 - 2012-06-26 - PhilippeCharpentier



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback