LHCb NCB MEMORANDUM

TITLE: **Introduction of T2-D sites to LHCb**

DATE: 2013-April-25

TO: NCB, Computing Management

FROM: NCB Chair (P.Clarke)

STATUS: DRAFT Version 1.0 (130425)

The idea of T2-D sites

This document discusses the introduction of T2-D sites to LHCb.

Currently LHCb only uses T2 sites for MC production and for reconstruction as sites attached to T1s.

The term T2-D refers to the extended use of T2 sites to hold Data (hence T2-D) and allow analysis. The model is that a T2-D will store stripped DST data according to the prevailing data management distribution policy. [Initially we may impose additional policies to ensure, for example, that only 1 of N copies is at a T2-D - details to be decided]. Analysis jobs requiring data will run at T2-D sites in the normal way.

The idea of T2-Ds has arisen for several reasons, including the immediate LHCb disk shortage problems, experience of the first years of running, the longer term position of LHCb compared to other experiments' use of the Grid, and the position we can take when arguing for more disk to the RRB.

The primary motivations can be summarised as:

- it is increasingly more difficult to argue externally that we cannot make use of disk resources at T2s when we argue for more disk
- or to argue that we alone amongst the LHC experiments cannot use T2s for data analysis.
- we have good experience with many T2s from the first LHC running phase

The primary concern (as identified in the Distributed Computing Review – see McNab/Templon report to be published) is that the operations team resources needed to interact with analysis sites scales with the number of sites, and so could overwhelm the team if many problems need chasing at T2-D sites.

To mitigate this:

1. In the first instance we will limit T2-D sites to a few (between 5 and 10) which have a track record with LHCb and/or are being proposed by non-T1 countries as a "T1 in lieu". Such T2-Ds will be selected in negotiation between the site and the operations coordinator.
2. We will use the experience gained in the pilot phase in 2013 to inform extension of the use of T2-Ds in the 2014 and beyond.

3. Each T2-D will be required to satisfy the set of requirements listed at the end of this document.
4. The performance of T2-Ds will be monitored by automated means (to be defined and developed). These metrics will give clarity to sites as to when they are performing adequately, or when they fall below a usable threshold.
5. We will appoint a T2-D coordinator from outside of the operations team, and with experience of running a T2 site. This person will take the main role in organizing the qualification of T2s, and be the primary liaison with them in case of performance issues.

It has also been suggested that we seek to augment the user support by using more experts from within LHCb. The DAST team already performs this function well and we will discuss with them how best to augment if necessary.

The aim is to be using a limited set of T2-Ds by the end of 2013.

Discussions so far with BR,CH and US (non-T1 countries).

Motivated by the LHCb "underpledge problem" (i.e. that if T1 countries pledge according to their national fraction of LHCb then, since T1 countries only account for 66% of LHCb, we guarantee an underpledge) I have spoken to the NCB reps of the three non-T1 countries who have >10 M&O authors and therefore make up a good fraction of the 33% non-T1 countries. These are Brazil, Switzerland and the US (I did not talk explicitly to Russia yet since we are expecting a T1 there soon anyway).

I asked whether it might be possible for them to provide a T2-D of a sufficient size to help LHCb. I have been very clear that this was an informal request to the NCB rep, and there was no implication of an obligation. The request was merely for "best efforts".

Responses have been positive from all in spirit, but in all cases the minimum disk requirement of 300 TB (see appendix) was seen to be a problem. We will therefore have to consider whether a "lead in time" might be acceptable, i.e. something like 100 TB now with the expectation to grow by the end of LS1.

Note that it still remains the strong wish of LHCb to find sites which satisfy the criteria, and could provide > 300 TB. This is relatively small compared to the T1s and is of the scale needed to solve the LHCB disk space problems in large part.

Next steps for NCB members

This document is the first comprehensive discussion document on where we are.

I would now ask NCB members to let me know if they think there are any candidate sites in their countries. In the first instance I will set up an informal discussion with the T2-D coordinator and the operations team to see if the site is likely to satisfy the criteria.

This matter will be further discussed at the next NCB.

APPENDIX: T2-D definition (evolving)

A T2-D must

- Be open to the LHCb VO
- Implements the necessary LHCb software framework, including CERN-VM for software distribution.
- Be able to offer LHCb at least 300 TB of disk storage to LHCb with commensurate CPU availability (in a similar ratio as T1 CPU/DISK)
- Have either
    - performed reliably for LHCb in 2012 as a MC production or as an "attached T2" site
    - or have equivalent assurances from the site and country concerned if this is being offered as a national contribution in lieu of T1.
- Satisfy the requirements of automated LHCb site performance monitoring which (to be developed).
- Be able to provide the name of a person who will act as the LHCb point of contact. It is not expected that this person is an LHCb member, nor that they attend LHCb operations meetings in general. However they must be in a position to progress any LHCb issues to the site authorities and to report back to LHCb. This person should join a special T2-D operations meeting once per month. This person should liaise with their national NCB rep and with the T2-D coordinator.
- Agree to adhere to prevailing WLCG uptime/scheduled downtime rules.
- Respond to WLCG problem escalation procedures.
- Have an adequate network connection which can support at least 1 Gbit/s of data transfer rate to CERN