Review of LHCb Computing Model

Discussion document

Version 130311

This document concerns the review of the LHCb computing model. It is a discussion document which collects together input and ideas arising from the consultation process.

In its present form it is only a collection of un-agreed ideas, and the reader should understand them in this spirit.

Section 1: The Past

Current Computing Model

A summary of the Computing Model as it had evolved to by the end of 2012 running is as follows:

- For the most part, production data processing (reconstruction, stripping) has been carried out at T0/T1 sites with the addition of selected T2 sites for re-processing.

- All user analysis work is carried out at T0/T1 sites if data access is required, and T0/T1/T2 sites if no data access is required.

- MC production work is carried out mostly at T2 sites, but also T0/T1 when available.

- The HLT farm is used whenever it is available.

- The data replication policy has evolved to the following in the face of disk and tape shortages:
    - All disk and tape data is at the T0 and the T1s
    - Raw data:
        - Tape: 1 copy at T0 + 1 copy at T1s
        - Disk: none
    - SDST/FULL.DST[1]:
        - Tape: 1 copy distributed between T0 and T1 at which created

[1] Note that FULL.DST is an evolution of the previous .SDST. It includes a copy of the RAW data in order to simplify access for stripping, but it therefore represents an additional RAW data copy as well.

        - Disk: none
    - Stripped DST:
        - Tape: 1 Archive
        - Disk: 4 copies of the latest processing and 2 copies of the previous processing. Distributed around the T0/T1 sites according to share.
    - MC: Archive copy + 3 replicas at T0/T1 sites

    - Physics group analysis: Physics groups access stripped DSTs or MDSTs on the Grid (at the T0/T1 sites holding the data). The output is dealt with in many different ways:
        - Often it is a subset of events – either a DST with a subset of events, a MDST or an Ntuple - which often ends up in a users private space, possibly on CASTOR or possibly at their home institute.
        - If a result of a WG production, then output datasets are left on the Grid and in the Book-keeping – these are then available to the group for further selections to be made to create Ntuples..etc.

Problems/Questions/Deficiencies – known and arising during the consultation

- **Disk space:** LHCb has had a perpetual disk space issue (partly this is due to success – i.e. LHCb can do so much more, including charm physics – so it is not a problem but a challenge).

- **Data replication and placement:** Our current data replication and placement policy has evolved from the original computing model out of necessity. It is largely a bespoke process which responds to disk space problems by reducing copies .etc. We therefore need to step back and review this afresh.

- **CPU power:** By and large LHCb has enough integrated CPU to match its needs, including the use of the HLT farm outside of LHC running. However we do suffer from peaks of use particularly during reprocessing periods which are in the run up to conferences. We need to ask whether there is any change we could make to our operational model which would smooth out the peaks and troughs.

- **MC planning:** We need to understand better our MonteCarlo needs, e.g. does volume scale with recorded luminosity, what fraction is trigger/stripping filtered, does it make sense to redo the whole simulation chain when there is a new reconstruction (rather than re-reconstructing existing MC), is the current way of dealing with spillover/pileup the best one?

- **Archive:** Review the concept of 'archive' of DST data: is one copy safe enough (probably because we can always recreate it) ?, what is use case for restoring an archived dataset (and how do we do it) ?, what is overlap with data preservation (and should we develop/use a DP infrastructure to take care of the archive). The archive is decoupled from online access – i.e. there is no simple recall from archive to disk.

- **Non Grid Working Group Space:** There is an important use case for shared disk space, which is not accessed via the Grid. Certainly some WG datasets can, and should, be produced and left on the Grid – typically major selections running over the entire processing, producing an output which will remain fairly static (this is what the WG production was developed for). However, WGs may then wish to create further derived data sets frequently (e.g. Ntuples, or small MDSTS used for mining) and make them accessible to a large number of people (e.g. for analysis). For these it is inappropriate to expect the users to "obtain" these regularly from the Grid as it is far to complex to do so. Often this leads to an individual keeping the "group dataset" in their private space. This is currently dealt with in an ad-hoc way by increasing that users quota. There are many drawbacks including the legacy issue when a user leaves.

- **User space on AFS:** A significant group of users (including CERN staff and other CERN based people) are constrained to use resources at CERN. This often leads then to having many problems associated with the very limited AFS space available. These people do not have access to large amounts of storage at home institutes. This brings particular difficulties when juggling with job outputs and often leads to significant loss of time.

- **Complexity and Knowledge:** The LHCb software is complex. There are a significant number of people who use it frequently and are relative experts and know what to do, or can find out from mailing lists, and can benefit from tutorials and Twikis. However there are also a significant number of people who do not use the software frequently enough and find it very complex to use and maintain a working job. Typically they are aware of a "few things" to do – but have no holistic knowledge. This leads to great barriers, particularly in using the Grid. Such lack of knowledge can lead to unintentional poor practice.

Section 2: THE FUTURE

Changing landscape

These are things happening "out there" in both WLCG and in the wider context, and which affect any New Computing Model.

- Storage Federations: There is work going on towards developing storage federations. A simplistic description is that the job requests a file and does

not care where it is. If the request cannot be satisfied locally it is redirected to somewhere within the federation which holds the data. The concept embodies a grand vision of anywhere-to-anywhere and much more use of the network. It is however not clear that this would in practice be a beneficial way to work. Much more understanding of how this could benefit us needs to be obtained.

- Parallelism and Multi-core: There is work going on to exploit multi-core devices efficiently and to see where parallelism may help. Its not clear that this affects the new computing model deeply but worth to keep in mind.

- Cloud resources at Tier-n sites: Whatever cloud resources means, we are probably going to have to use them – thus we will plan to first understand and then create the capability to use such resources.

Know Evolutions

These are things which a priori, we are pretty sure we will do.

- **Make more use of T2 sites**: we are almost obliged to make use of T2 resources, and in particular storage at T2s. It is very difficult to defend our storage problems unless we do. Ideas are presented below.

- **Better monitoring of performance of T1 and T2 sites** with clearly understandable metrics. If we are to expect a certain level of performance from an expanded set of sites, then it is only fair that we publish fair and clear metrics which indicate good and bad performance. This gives us a common currency with which to discuss performance problems with external sites. This monitoring can also include network transfer speed criteria.

- **Data migration and recall facility:** As part of the data management evolution we need to develop a data recall facility. This would allow data which is never touched to be migrated to nearline – such that it can be recalled to online disk if and when needed (in the way we are used to for Castor)

- **WG Space:** A significant fraction of users do not use the grid for analysis, instead using local facilities. We must consider again how to provide shared WG space which is accessible to them and which is not assigned to an individual user. Perhaps the move from CASTOR to EOS will play a role here. An alternative/addition would be to develop a simplified (standard) way to **"*go and get*" data from grid storage:** Many in practice have difficulty in the task of reliably transforming LFNs into a merged copy on their local machine. Even experts can struggle with this. If there were a simple, standard and robust way of gathering, merging, and transferring Ntuple and MDST datasets to local storage, then such space would more easily satisfy WG shared space requirements.

- **User space at CERN:** There is a sufficient problem for CERN/AFS users (i.e. those with no practical option of institute based storage) that this issue should be resolved. Again, the EOS user space may do this.

Possible new Computing Model ideas

- **Simple data parking:** In its simplest form this means taking more data than we can deal with immediately and feeding it off to offline tape. It would not be processed until well outside of proton collision time (possibly even LS2). In the simplest incarnation there is no differentiation – i.e. the parked data is the same as the processed data. This is unlikely to be an optimal strategy, but we include here merely as a baseline. The next bullet presents a better strategy.

- **Multiple trigger streams**: [motivated by Charm group] The parked data is not just more of the same and is not necessarily processed as a matter of course in the same way as the immediate data. In this model there would be several different types of RAW data containing different triggers. These would be fully independent, including lumi ..etc.. such that they can be analysed separately. The model would evolve as to what was in each stream, but the general principle would be that some was analysed "now" and some "later" and perhaps some "never unless a problem arises". As example:
    - The first type would be as we are used to, an "*immediate processing"* class of triggers. This would flow and be processed exactly as we do now. Its rate would be limited to that which the processing chain can deal with ~ 5 khz
    - The second type would contain a "*deferred processing"* class of triggers. This would simply flow to tape and not be processed until later. Physics groups would then make a case (to be approved by PPG perhaps?) to get on a list to "touch" these data. Such requests would be aggregated and the processing run at scheduled time in the future. If a development sample were needed then a request to process a small fraction could be made earlier.
    - The third type would be a *"mostly never process*" class of triggers. This would go only to tape. It would contain triggers for which it is expected that the full analysis could be done using HLT information (possibly augmented from what it is now – see next item) such that the raw data need only be accessed in case of unforeseen systematic uncertainties arising in some of the relevant analyses (and then possibly only a small fraction).

- **More processing in the HLT:** [motivated by Charm group] A general theme from the Charm group is to do more processing in the pit to reduce the data required as much as possible as close to source as possible. Perhaps an HLT3 to make an even more exclusive set of B triggers for the highest priority analyses (and then park the rest). For charm it could be useful in the context of making some high priority subsample of the inclusive D* trigger.

- **Prompt processing=end of year re-processing**: We have already evolved to a fairly efficient processing model and should optimise this further. The prompt processing will be automated such that it will be the final processing at the end of annual running (i.e. no further re-processing immediately after data taking). This will be achieved by introducing a delay to allow automated calibration and alignment to take place before the initial processing
  - A full restriping will take place shortly after the end of data taking. This will have a long life, and only incremental/bug fix strippings will happen.
  - The Ultimate reprocessing will then take place >~ 1 year later or possibly even only at the next LS.

- **Intelligent data replication and placement**. All data is archived, and so there is no risk of losing it. It might be argued that the current number of strippings (4) is too crude a policy, and DST data should only be replicated according to its likely use (the counter argument is availability in the face of T1 down-time). We will therefore consider developing a more intelligent automated system based upon data popularity. This means that some data sets may have more copies than others on the Grid. It is also noted that the C-RSG asked for all Tier centres to compile stats on data which is rarely accessed - so we should be in the position to know this for ourselves. A reduction in DST copied could be tied to more copied of an associated MDST – see next item.

- **More use of MDSTs stored on the T2s**: Perhaps we can use the opportunity of space available at T2s to store many MDSTs with additional information – and which would then be the natural source for working group Ntuple production. It has specifically been suggested by the Charm group that MDSTs (or similar) could be the normal way of doing analysis, meaning that it would only be necessary to maintain a single copy of the DST data itself on disk (not 4 copies). MDSTs would be produced by less frequent WG productions.

- **Is there a case for a (very-)compact general purpose data set ?** By and large LHCb production is still focused on the "rare-decay" paradigm, i.e. the analyses done by LHCb are based on few events out of a very large number of recorded events. The charm physics program already deviates from this, and also in the QCD sector there are some potential analyses which could benefit from access to a very large number of events. Therefore we can ask whether there is a case for a new very-compact data format containing only a small amount of information but for a much larger number of events. Of course study would be needed to verify that such a minimal information set would actually be useful and would not require going back to the DSTs

- **T2-D Centres (LAC Centres)**: This model would resurrect the old idea of designating some T2 sites being of sufficient size and performance as to allow (i) hosting of data and (ii) possible use for user analysis which requires data access. This is very similar to the original LAC definition (Local Analysis Centres). At this point it is not proven that this is a requirement, and there is

a good argument that if T2-Ds are used to offload much data processing, then all analysis can be carried out at the T0/T1 sites. However, if this were felt to be useful then we would select a few (~ 10 in the first instance) T2-D sites with a large amount of storage and a proven performance record. These sites would have to agree to a set of criteria (see appendix). Analysis jobs would then run in the normal way. They would be treated largely like the T0 and T1s are today, with explicit data management and work quotas.

Caveat: The parallel Distributed Computing Review has already noted that the way in which LHCb interacts with T1 sites is labour intensive for a small number of people. Unless this situation could be changed then it remains the strong feeling that operations team could not easily support analysis on data at T2-Ds.

- **Opportunistic computing:** Probably this isn't really a new CM item per se, but nevertheless interesting. Several in the computing group would be keen to develop opportunistic computing to run on desktops, home computers…etc. Depending upon take up this could provide a useful resource increase and might provide public outreach possibilities.

- **Use of private and/or commercial clouds:** In fact the current DIRAC can already send jobs to amazon. We will likely need to develop this to ensure we can take advantage of other cloud resources.

- **Documentation and examples:** This is not strictly CM, but no review in any context would be complete without noting the universal issue of having insufficient documentation. **It is emphasised that a lot of people have already put a lot of time into this sector, and their efforts are praised.** We also have tutorials, although there are not always well attended, and we have web based material, although it has been said that it is often hard to obtain a current example which goes from beginning to end and which will work (as examples go stale very quickly) . There is no easy solution to this – there are a few people able to provide documentation and they are fully occupied doing other essential tasks. Nevertheless, whilst the situation persists, there will be a significant fraction of LHCb who will not be able to perform computing tasks in the optimum way. This is an area where LHCb management may have to consider radical solutions, such as defining documentation on the tasks list, or crediting it at the same level as shifts (i.e. a group of people who provide a useful set of documentation are credited by equivalent shifts, and they obtain an ongoing shift credits for maintenance).

APPENDIX: T2-D(LAC) sites definition (Evolving Draft)

The concept of LAC sites was discussed in LHCb in the past (see https://twiki.cern.ch/twiki/bin/view/LHCb/LHCbNationalComputingBoard).

This appendix revisits this to define what we mean by a T2-D (Tier-2 with Data) and T2-A (Tier-2 for Analysis). The following criteria are an evolving draft.

A T2-D must

- Be open to the LHCb VO
- Be above a certain size which will be approximately > 5% of the aggregate T1 CPU and disk storage.
- Be able to offer LHCb at least 300 TB of disk storage to LHCb.
- Have performed reliably for LHCb in 2012 as a MC production or as an "attached T2" site (what is the metric here ??)
- Satisfy the requirements of automated LHCb site monitoring (to be developed) to be included as an active site.
- 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. However they must be in a position to progress any LHCb issues to the site authorities and to report back to LHCb
- Implements CERNVM-FS for software distribution.
- 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