

# Upgrade computing TDR, a roadmap

Many people<sup>1</sup>.

<sup>1</sup>*Many places*

## **Abstract**

There is no place like the future.



# 1 Introduction

This document defines the roadmap towards a Technical Design Review (TDR) document for the computing model of the upgrade LHCb [1] experiment.

## 1.1 Wider Context

There are many factors which are at play and which will affect the nature of the distributed computing infrastructure which will exist at the time of Run-III

The projections of the combined data volumes of all experiments at the time of Run-III compared to the computing resources that are likely to be fundable, lead us to conclude that by just carrying on as normal, there is a significant risk that the required computing resources will not be sufficient. This assumption is driving all parties concerned, including funding agencies, the WLCG and experiments to devise alternative paradigms.

Relevant factors include:

- It is expected that all funding agencies will face continued tight constraints upon the funding available for computing. It is common to assume that only flat cash will be available (i.e. a constant number of currency units), but it is worth bearing in mind that even this may be optimistic.
- Whereas the LHC programme has been essentially the only data intensive scientific programme in the past, this will not be the case by 2020. Amongst other endeavours the SKA will potentially be set to produce even greater data rates, and other major astronomical instruments such as LSST and EUCLID will produce significant amount of data. Therefore we will not be the only activity seeking large computing resources.
- Recognising the previous reality, all of the major funding agencies within Europe are converging upon a policy of sharing infrastructure across many of the activities for which they are responsible. This means that the major data centres will need to evolve to a more activity-neutral offering, which is more easily used by different disciplines. High on the radar in this respect is deployment of a virtualised infrastructure and a simple single sign on framework.
- It is not axiomatic that provision of resources exclusively in home grown data centres in science laboratories will remain the most cost effective way to provision. Accordingly projects to pilot a hybrid cloud infrastructure are going ahead. These will seek to set up cloud resources provided both by traditional laboratory data centres and commercial providers. We would be prudent to assume that such a model will be prevalent in 2020.
- If EU funds are to be available (and this is a big if) they will only be available for highly shared infrastructures. From the funding agency point of view it is mandatory that particle physics engages in such cost sharing activities. This will probably take

38 place at a level above the experiments, but the fact that it is happening will have  
39 some effects upon HEP provisioning.

- 40 • In some quarters there is an expectation to make use of more diverse resources.  
41 ATLAS already make use of spare cycles on high end HPC machines for Monte Carlo  
42 simulation.

43 Naturally some of the directions indicated above are speculative, and individuals will  
44 have their own guesses as to which are likely to become reality. However it is incumbent  
45 upon us to bear all of these points in mind when building our upgrade computing model.

## 46 2 Collaborative Tools

47 The purpose of a collaborative tool is to support a group to accomplish a common  
48 goal. Within our collaboration, work is usually organised into groups at different levels  
49 (subdetector, analysis, institutions, ...). Our main goal is performing a physics measurement:  
50 this is achieved via several steps, from detector maintenance to data taking, processing  
51 and analysis, each step generally requiring the interaction of different people and the  
52 sharing of information. The goal of the Collaborative Tools session was first to understand  
53 the status of collaborative practices within our collaboration, and then start thinking  
54 how we can improve how we work together, exploiting tools and techniques currently  
55 available. Improvement not only in terms of scientific outcome but also in terms of working  
56 conditions.

### 57 2.1 Collaborative tools in LHCb: current status

58 A questionnaire was submitted to LHCb collaboration to understand the the level of  
59 collaborative practices and collect feedback. About 100 people answered the question-  
60 naire. The main outcome of the questionnaire was that collaborative tools for archival of  
61 documentation and software are widely used within our collaboration. People involved  
62 in core software development tend to use standard tools, e.g. svn and gitlab repositories  
63 for software archival. Instead people working only on data analysis tend to use many  
64 different tools, also commercial ones, which makes sharing documentation and software  
65 more difficult. Moving to work organization, the usage of tools such as issue trackers  
66 is very common among software developers, contrary to people working only on data  
67 analysis. Code review techniques are not popular in both groups. The detailed results of  
68 the questionnaire can be found at this link.

### 69 2.2 Main topics of discussion

70 During the session we had discussion about the following items:

- 71 • **Automated analysis workflows.** Our analysis are usually made of several steps.  
72 Creating a detailed workflow of the analysis can have several advantages, e.g. replicate

73 the analysis on a different sample, share the analysis with other people, keep track  
74 (archive) of all the details of the analysis. Among the different tools to develop  
75 workflows, *gitlab* is already used with success within our collaboration. This tool  
76 makes code sharing and testing easier; as an example, it features the possibility of  
77 producing a private branch where algorithm modifications can be reviewed and tested  
78 before being merged with the main branch. It is now also used in a few analyses  
79 and has proven to be useful in this environment as well.

- 80 • **Reproducible software.** The *everware* project aims at giving users an infrastruc-  
81 ture to make easy to run someone else's code. Its main ingredients are the code,  
82 archived in a software repository, a dockerfile which describes the environment and  
83 an executable README, e.g. a jupiter notebook to describe the analysis flow. Such  
84 a tool is also useful to make the usage of code easier. As an example, new machine  
85 learning algorithms could be promoted thanks to simple analysis examples hosted in  
86 a *everware*-like container.
- 87 • **Analysis documentation and preservation.** A fundamental ingredient of an  
88 analysis is documentation. People in LHCb use many different media for this (slides,  
89 minutes, twiki, jira, ....) and different archival tools. Can we aim at a higher level of  
90 uniformity for analysis documentation? The ideal analysis logbook should allow to  
91 record different types of information, from text to plots, intermediate results and  
92 code. It needs to record any changes to a topic, and allow to share entries with other  
93 people for commenting and reviewing. A tool which is intended to help preserving  
94 all analysis ingredients, from the very first step to the final publication is the Cern  
95 Analysis Preservation framework LHCb is developing together with Cern-IT and  
96 other LHC experiments.
- 97 • **Software tests.** Software testing is very important and users should collaborate  
98 with core software developers in writing tests. To enforce code testing we should  
99 agree on a set of guidelines - e.g. new code or fixes must come with a test - and on a  
100 procedure to make sure this guidelines are met - e.g code reviewing.
- 101 • **Code review.** Code review is a common practice in industry. It has several advan-  
102 tages, as enforcing a common style of coding, improving overall the documentation  
103 of the code, identifying bugs or problematic parts in the code as early as possible,  
104 improving coding skills. The goal is to achieve the best possible code and minimise  
105 bugs, not judging the author's ability to write code. Code should be commented and  
106 merge requests should be as small as possible. Different tools exists but *github* and  
107 *gitlab* have easy to use GUIs to comment on specific lines in the code for comments  
108 and discussion.
- 109 • **Guidelines for contributing.** In order to increase the number of people contribut-  
110 ing to LHCb software development we should agree on a set of guidelines. The idea  
111 is to create a handbook on how to contribute to LHCb software. This handbook  
112 should tell who can contribute, on which subjects, which rules should be followed,

113 where people can find more information about the software to be developed. This  
114 handbook can become one of the Starterkit lessons.

## 115 **2.3 Conclusions and next steps**

116 From the questionnaire and the discussions during the session we can conclude that it is  
117 becoming clear that people want to change towards a more effective collaboration. For  
118 almost all technical challenges we already have the tools and solutions. For many there  
119 are even working prototypes tested in LHCb at the moment. We have started to identify  
120 psychological and organisational obstacles on the way to collaborative working and in the  
121 next future we will work to overcome them. We have started to collect documentation  
122 about the collaborative tools discussed during the session in the repository at this link.

## 123 **3 Data processing/access and analysis models**

124 Physics WGs would be OK with having centralised ntuple production, at least for  
125 stripping lines going to DST. This would allow a reduced number of DST stream replicas  
126 and possibly higher retention. A dedicated WG liaison would take care of organising  
127 the “ntupling” that could run at fixed intervals (e.g. two weeks). Such analysis train  
128 would allow for a more uniform and efficient data management. It would however require  
129 that the WG release their code several days ahead of the train departure. It is also  
130 envisaged for some analysis a more direct usage of  $\mu$ DST, without having to pass for  
131 ntuple production. This would require a better documentation and standardisation of the  
132 muDST format and tools to access it.

133  
134 In the upgrade the current workflow with a Stripping will break both because of  
135 the size of the raw and of the cpu requirements needed to perform periodic strippings.  
136 Moreover, with a fully software trigger, the stripping retention would also be very likely to  
137 increase to values at which the dataset is not significantly reduced anymore (at least in  
138 terms of number of events). Turbo stream should probably become the “default” in the  
139 upgrade, many existing stripping lines could in principle make perfectly valid turbo lines  
140 with no impact on the analysis feasibility. However some WG would still need information  
141 from the entire event, or at least a significant part of it. Therefore, although in a reduced  
142 bandwidth fraction with respect to Turbo, the complete event should be kept in some  
143 form. Developments in the triggers together with the agreement between online and  
144 offline reconstruction would potentially lead to DST production directly as output of the  
145 trigger. These DST would contain all the combinatorics performed during the trigger that  
146 should not to be re-performed offline (as it is currently done in the stripping).

147  
148 The possibility of having streaming at the trigger level would allow different level of  
149 information propagation. For instance, one could imagine having streams with no raw  
150 banks or only some of them, streams in between a muDST and DST where only part of

151 the event is saved (e.g. saving only tracks within a cone around the trigger candidate).  
152 In addition, all triggered events could be “flagged” into an event index (which could be  
153 produced by the successor of the stripping) so that users could quickly read single event  
154 via random access, this would remove the need to loop over large files, having to store  
155 related info such as flavour tagging and isolation variables and reduce the number of  
156 replicas. The result of the event indexing could be stored either in files as it has been  
157 proposed in the past (ETC files) or via a centralized service. Such a workflow however  
158 requires careful evaluation in terms of performance.

159

160 Propositions were made in order to make the Event Indexer service more useful. They  
161 include an HTTP/REST API, indexing of the MC and Turbo data, a bandwidth division  
162 feature or a functionality to search for events similar to a sample of events given as  
163 input. Interest was shown in these features, but it requires more thoughts before deciding  
164 whether they could show helpful.

165

166 The middleware to support user analysis and productions is and will stay DIRAC.  
167 An evaluation of its scalability has been made, taking into account an increase of traffic,  
168 data and maintenance. The conclusion is that new technologies must be integrated, and  
169 some core part of DIRAC must be replaced, but overall only evolutions are needed, not a  
170 revolution.

171

172 In order to cope with the increased amount of data, a more dynamic and flexible  
173 placement of data could be achieved using the data popularity and taking advantage of  
174 the analysis train paradigm.

175

176 Efforts from the institutes should be made in order to enhance computing skills of  
177 students and PhDs and each working group should keep an up-to-date analysis to be used  
178 as a best practice example. In the same spirit as the the starter kit, intermediate tutorials  
179 focusing on best practices for the various tools would be appreciated.

## 180 4 Event Model Summary

181 The current event model has been very successful in some aspects and less so in others.  
182 It has enabled developers to write code that performs various reconstruction tasks and  
183 analysts to create the selections they need by representing reconstructed objects and  
184 decay trees in an intuitive way. This ease of use of event model objects should be kept.

185 At the same time, as a result of the Array of Structures (AoS) design of the event  
186 model classes, in its current form exploitation of modern CPU features such as SIMD  
187 are more difficult. In addition, current event model objects are not composable, so if  
188 information needs to be added to objects without modifying them, they need to be copied.  
189 This costs significant memory, which in turn leads to suboptimal usage of resources.

190 While current event model objects were intended to be read-only after they had been

191 added to the Transient Event Store (TES), this was not obvious in the way they are  
192 retrieved from the TES by algorithms. As a result, parallelisation of algorithms is more  
193 difficult and would often require cloning of objects.

194 Since computing resources will be more and more limited compared the rate of events  
195 that needs to be processed, speed and memory efficiency should be core guidelines for the  
196 design of the new event model. New event model objects are therefore recommended to:

- 197 • be composable,
- 198 • have Structure of Arrays (SoA) memory layout,
- 199 • become read-only after their initial creation and “filling”,
- 200 • use `float` precision where possible.

201 Transformation of event model objects between their on-disk and in-memory repre-  
202 sentation will remain needed, for example to be able to efficiently store objects that need  
203 double representation during calculations, but fewer bits on disk.

204 Information such as luminosity and simulation statistics are best stored at the file  
205 level. The storage of this type of data should be reviewed to become more extensible  
206 and light-weight to allow for example production information and simulation statistics to  
207 be added. If this data can be read from the file in a lightweight way, it can be used to  
208 configure jobs or to check their configuration.

209 If the code required to read files is available as a single library, or a small set of  
210 libraries, with minimal dependencies, this would facilitate analyses and improve the forward  
211 compatibility of files.

212 To ensure data can be read in the future, data preservation should be an important  
213 factor in the way the event data is stored.

## 214 5 Hardware and Dataflow

215 All discussions of hardware must be tightly coupled to discussions of software. While the  
216 focus of discussions was what hardware plus software would be sufficiently performant  
217 for the online system, several people emphasized the need to process at least as much  
218 simulated data as real data. This may require that similar hardware resources be available  
219 offline and online. In addition, the life-cycle cost of the hardware and software, including  
220 capital costs, operating costs, and development plus maintenance costs for software were  
221 identified as critical issues.

222 We discussed three families of CPUs:

- 223 • “Ordinary” Intel and AMD x86 processors;
- 224 • ARM64 processors;
- 225 • OpenPOWER processors.

226 It is expected that all three can run our software. *Building the LHCb software stack*  
227 *regularly on an instance of each platform (continuous integration) was identified as an*  
228 *action item.* In addition to planning for Run 3, this will allow us to identify bugs in  
229 existing software that are not caught when building on just one architecture.

230 We discussed two general classes of accelerators in detail:

- 231 • GPUs, including both NVIDIA and AMD boards;
- 232 • Intel Xeon Phi boards with circa 60 CPUs each.

233 The GPUs may be more performant in terms of bang per buck and bang per watt. Daniel  
234 Campora compared a vectorized CPU velopix algorithm with an OpenCL velopix algorithm  
235 implemented on a traditional Intel Xeon node, an NVIDIA GTX 980, an AMD HD 7970,  
236 and an Intel Xeon Phi. The cost for a standard unit of performance was judged to be  
237 roughly a factor of three lower for the AMD GPU than for a traditional CPU node, and  
238 that for the NVIDIA GPU only somewhat worse than that of the AMD GPU. Alexey  
239 Badalov discussed a co-processor manager in our session, and Illya Shapoval presented a  
240 description of GaudiHive in a parallel session. The common conclusion is that low latency  
241 interfaces for Gaudi-like algorithms to use accelerators seamlessly should work, be it on  
242 the same system or in an offloaded execution. The question of *how many events in parallel*  
243 *are required to benefit from an offloaded execution* is dependent on the specifics of the  
244 algorithms in consideration, host-device transmissions and efficiency of the scheduling,  
245 and should be addressed in demonstrators to come.

246 Marco Corvo described integrating GPU hardware and an interface to Gaudi into the  
247 EFF. To decide whether to adopt GPUs, Xeon Phi’s, or other accelerators as part of  
248 the vision for Run 3, *we should rapidly develop an inventory of up to six key algorithms*  
249 *used in the HLT, and build demonstrators to see whether these technologies can provide*  
250 *enough additional performance to justify the additional development effort required to build*  
251 *and maintain the requisite software.* It was noted that much of the design work to make  
252 algorithms parallel and vectorize them will also improve performance on “ordinary” Intel  
253 CPUs. Rainer Schwemmer reported that the existing HLT software is currently using  
254 about half the CPU cycles available on the farm due to cache misses, and one might  
255 reasonably hope for 90%. Manuel Schiller discussed how to design/re-design software to be  
256 cache-friendly and pointed out that the advocated approaches also promote parallelization  
257 and vectorization. As it is possible that writing high-quality, properly vectorized code  
258 for Intel CPUs will increase performance significantly, *we need to develop and evaluate*  
259 *optimized CPU versions of the demonstrator algorithms in parallel with those for the*  
260 *accelerators to make useful comparisons* (excuse the pun).

261 Utilizing modern vector and parallel architectures effectively requires a high level of  
262 software design and coding expertise. Therefore, *we should provide extensive training and*  
263 *mentoring in this area.* This should include

- 264 • hands-on workshops to teach modern, high-quality C++ and parallel programming  
265 languages such as CUDA and OpenCL;
- 266 • formal code reviews, both to provide feedback to the original coders and to provide  
267 opportunities for the reviewers to learn from others' examples;
- 268 • clear documentation of design choices so that later developers understand what was  
269 done and why; this might allow good designs to serve as templates for later code  
270 development.

271 Collaborating with other experiments has the potential to advance our understanding  
272 of what is required for Run 3, and should be pursued as an integral part of the software  
273 R&D plan. For example, grid resources are often “common”. If we want to use particular  
274 architectures offline (CPU or accelerator), they are more likely to be supported if other  
275 experiments also adopt them. As lawyers like to say, *time is of the essence*; we need to  
276 implement the action items ASAP. Repeating what was said before:

- 277 • We should build the LHCb software stack regularly on ARM64 and OpenPOWER  
278 platforms in addition to x86 platforms.
- 279 • We should rapidly develop an inventory of up to six key algorithms used in the HLT,  
280 and build demonstrators using accelerators to see how much additional performance  
281 these technologies might provide. In parallel, we need to develop and evaluate  
282 optimized CPU versions of the same algorithms to make useful comparisons.
- 283 • We should provide extensive training and mentoring in developing state-of-the art  
284 software.

285 As usual, all is easier said than done. Identifying and deploying resources for these activities  
286 is necessary to move forward with the road-map and TDR.

## 287 **6 Framework and Scheduling**

288 Discussions of software must be coupled to discussions of hardware – it is the evolution of  
289 the available computing hardware over the last decade that is forcing us to rethink the  
290 basics of our framework.

291 Even though Moore’s law has continued to evolve over the last decade, and will for the  
292 next few years, the trend of ever faster single-threaded CPU performance has been broken  
293 about ten years ago. Instead, the additional transistors have been invested into multiple  
294 CPU cores on a die, and, within these cores, into wider execution units. Furthermore,  
295 the gap between the processing units (PUs) and memory has gotten wider, and feeding  
296 the processing units with data is becoming more than ever a bottleneck.

297 As a result of these trends, decisions made (more than) ten years ago are no longer  
298 appropriate. For example, leveraging the 'embarrassingly parallel' nature of distinct  
299 collisions, by running multiple independent processes is actively contributing to increasing  
300 cache-miss, which subsequently stall these processes. Hence it seems necessary to bite  
301 the bullet, and deal with multi-threaded processing of events instead. This comes with its  
302 own complications and overheads, but will help in a more efficient utilization of the cache,  
303 as multiple cores will work on a coherent dataset.

304 However, one has instead to face the problems of enabling multiple cores to work  
305 on closely related data without trampling on each others. The solution to this problem  
306 which has been accepted in industry (eg. Microsoft, Apple, Google, Intel, ... ) is what is  
307 known as a 'task-based' system. A demonstrator of such a task-based system has been  
308 implemented on top (or to be more correct, beneath) the current Gaudi framework.

309 From this demonstrator it has become clear that, even though it is possible to 'transform'  
310 Gaudi into a task-based system, the domain-specific code which is written on top of Gaudi  
311 has to be changed for the system to be able to scale properly – both in performance, as  
312 well in understanding its configuration and operation.

313 The main changes pertain to making algorithms and tools 'stateless' so that their main  
314 methods of execution can be wrapped inside independent tasks, and (thus) be run in  
315 parallel. Also, their input and output data must be explicitly declared in a manner which  
316 is accessible to the framework, so that the framework is able to properly schedule these  
317 tasks dynamically, insuring that all cores remain busy.

318 By performing these changes, the domain specific code is not cluttered with the details  
319 of the concurrency framework. This is not only needed to allow us to follow the evolution  
320 of the C++ standard in this area but also to avoid people making assumptions on the  
321 specific implementation which will still need to evolve. As a side-effect, these changes  
322 typically make the individual algorithms more focussed on a single task, making them  
323 easier to comprehend, and also more 'regular', i.e. some of the current 'boiler plate' is  
324 abstracted away, and dealt with in dedicated, shared framework code.

325 To be able to write algorithms this way, several requirements have to be made on the  
326 event model. First and foremost, once data is globally made available, it must remain  
327 immutable. This is due to the fact that the order of execution becomes more dynamic,  
328 only bound by control and data flow dependencies. Changes 'in situ' complicate this  
329 tremendously, not only for the required scheduling, but also for people to comprehend the  
330 behaviour of the system.

## 331 **7 Conclusion**

332 We gonna do what they said can't be done.

## 333 Acknowledgements

334 We express our gratitude to our colleagues in the CERN accelerator departments for the  
335 excellent performance of the LHC. We thank the technical and administrative staff at the  
336 LHCb institutes. We acknowledge support from CERN and from the national agencies:  
337 CAPES, CNPq, FAPERJ and FINEP (Brazil); NSFC (China); CNRS/IN2P3 (France);  
338 BMBF, DFG and MPG (Germany); INFN (Italy); FOM and NWO (The Netherlands);  
339 MNiSW and NCN (Poland); MEN/IFA (Romania); MinES and FANO (Russia); MinECo  
340 (Spain); SNSF and SER (Switzerland); NASU (Ukraine); STFC (United Kingdom); NSF  
341 (USA). We acknowledge the computing resources that are provided by CERN, IN2P3  
342 (France), KIT and DESY (Germany), INFN (Italy), SURF (The Netherlands), PIC (Spain),  
343 GridPP (United Kingdom), RRCKI (Russia), CSCS (Switzerland), IFIN-HH (Romania),  
344 CBPF (Brazil), PL-GRID (Poland) and OSC (USA). We are indebted to the communities  
345 behind the multiple open source software packages on which we depend. We are also  
346 thankful for the computing resources and the access to software R&D tools provided  
347 by Yandex LLC (Russia). Individual groups or members have received support from  
348 AvH Foundation (Germany), EPLANET, Marie Skłodowska-Curie Actions and ERC  
349 (European Union), Conseil Général de Haute-Savoie, Labex ENIGMASS and OCEVU,  
350 Région Auvergne (France), RFBR (Russia), GVA, XuntaGal and GENCAT (Spain), The  
351 Royal Society and Royal Commission for the Exhibition of 1851 (United Kingdom).

## 352 References

- 353 [1] LHCb collaboration, R. Aaij *et al.*, *LHCb detector performance*, Int. J. Mod. Phys.  
354 **A30** (2015) 1530022, [arXiv:1412.6352](#).