

LHCbDIRAC limits on the number of streams

9th computing workshop

C. HAEN

16th May 2017

Limits for the streams out of the HLT

Limits

- More streams \Rightarrow more writers
 - merging overhead (maybe solved ?)
 - Online storage: IO limitation to some extent
- Bandwidth to CASTOR is limited
- Per file overhead for transfers
- Operationally, if different processing \Rightarrow more productions & transformations

Status

- Not OK to divide streams: operational issues offline
- Not OK to add streams: sorry, we're full

Limits for the streams out of the stripping

Limits

- More streams \Rightarrow more output files \Rightarrow higher likelihood that one fails to upload \Rightarrow operationally (very) annoying.
- Large operational overhead for follow up (runs flushing, consistency checks, etc)
- More stream \Rightarrow more merging jobs \Rightarrow more operational work
- More streams \Rightarrow more input for merging jobs \Rightarrow input data (MINIBIAS stream in Stripping28)

Status

- Technically, adding 1 or 2 more streams would not kill LHCbDIRAC by itself. 50 would...
- Operationally, difficult to add any more streams.

Few more notes...

- In terms of resources and operations, the less overlap the better
- In Run 3, CPU will be an issue: multiplying streams and merging would mean more cpu intensive jobs which we cannot afford
- Ideally, all streams out of the stripping should be balanced