Reducing Muon EDM size

Edward Moyse obo Muon Software

Current Sizes (ESD)

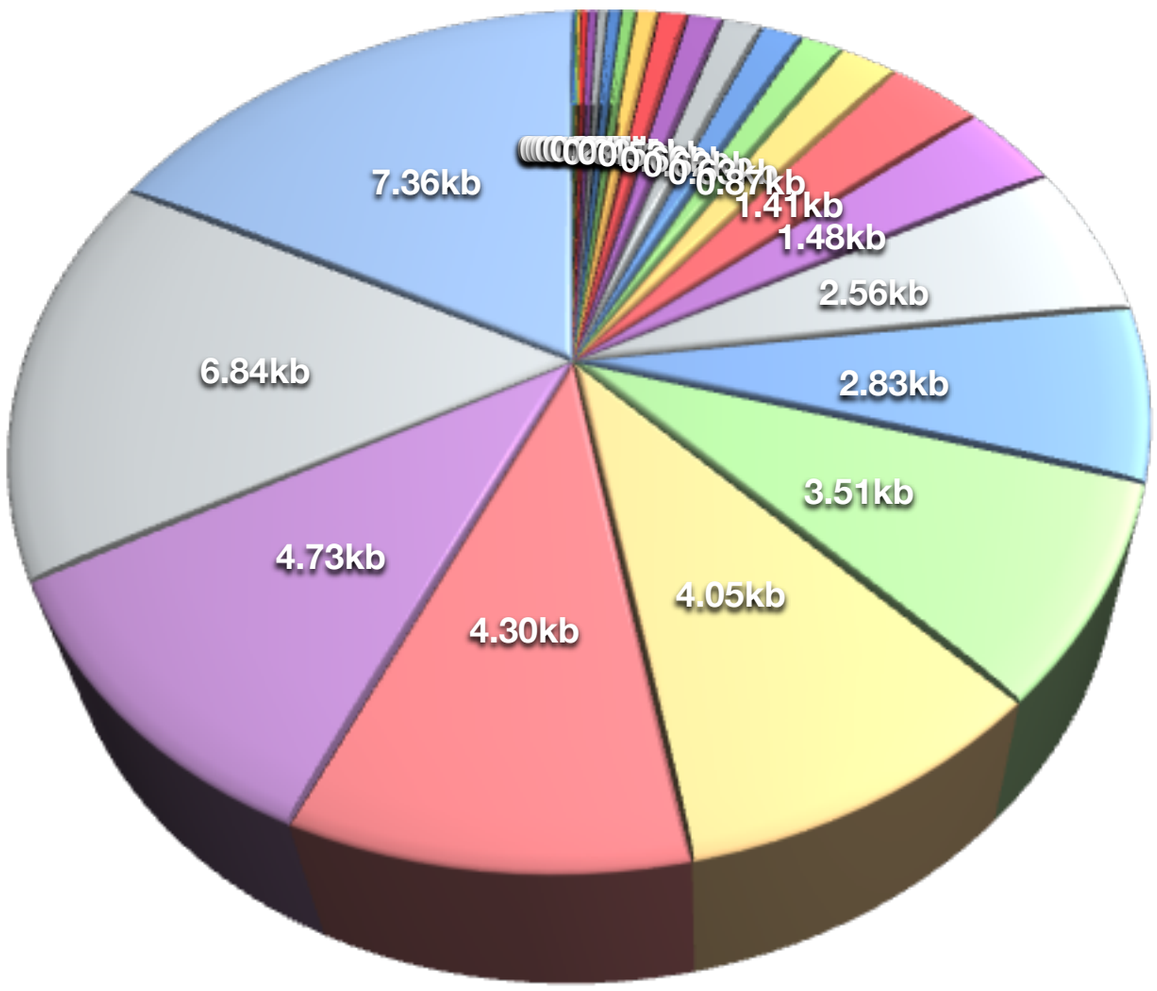

- Muon::ChamberT0s_p1_MooreMuonChamberT0s
- Muon::ChamberT0s_p1_MboyMuonChamberT0s
- Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollectionNextBC
- Muon::MuonPRD_Container_p2<Muon::RpcCoinData_p1>_RPC_triggerHits
- Muon::TgcPrepDataContainer_tlp1_TGC_MeasurementsNextBC
- Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollectionPriorBC
- Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollection
- Muon::CscPrepDataContainer_tlp1_CSC_Clusters
- Muon::TgcPrepDataContainer_tlp1_TGC_MeasurementsPriorBC
- Muon::TgcPrepDataContainer_tlp1_TGC_Measurements
- CscRawDataContainer_p3_CSCRDO
- Muon::RpcPrepDataContainer_tlp1_RPC_Measurements
- Rec::TrackParticleContainer_tlp1_MuonboyTrackParticles
- Rec::TrackParticleContainer_tlp1_MuonboyMuonSpectroOnlyTrackParticles
- Rec::TrackParticleContainer_tlp1_MooreTrackParticles
- Trk::SegmentCollection_tlp1_ConvertedMBoySegments
- Muon::CscStripPrepDataContainer_tlp1_CSC_Measurements
- Trk::SegmentCollection_tlp1_MooreSegments
- Trk::TrackCollection_tlp3_ConvertedMBoyMuonSpectroOnlyTracks
- MuonMeasurements_tlp1
- Muon::MuonPRD_Container_p2<Muon::MdtPrepData_p2>_MDT_DriftCircles
- RpcPadContainer_p2_RPCPAD
- Trk::TrackCollection_tlp3_ConvertedMBoyTracks
- Trk::TrackCollection_tlp3_MooreTracks
- MdtCsmContainer_p1_MDTCSM

/castor/cern.ch/atlas/atlascerngroupdisk/det-muon/Bytestream2010/
MuonBytestreamRun155697_selectedCbMu10Events_rec15.6.10.ESD.pool.root

Current sizes (2)

Size/Evt kb	Container Name
0.030	Muon::ChamberT0s_p1_MooreMuonChamberT0s
0.031	Muon::ChamberT0s_p1_MboyMuonChamberT0s
0.042	Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollectionNextBC
0.101	Muon::MuonPRD_Container_p2<Muon::RpcCoinData_p1>_RPC_triggerHits
0.141	Muon::TgcPrepDataContainer_tlp1_TGC_MeasurementsNextBC
0.152	Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollectionPriorBC
0.202	Muon::TgcCoinDataContainer_tlp2_TrigT1CoinDataCollection
0.211	Muon::CscPrepDataContainer_tlp1_CSC_Clusters
0.330	Muon::TgcPrepDataContainer_tlp1_TGC_MeasurementsPriorBC
0.438	Muon::TgcPrepDataContainer_tlp1_TGC_Measurements
0.527	CscRawDataContainer_p3_CSCRDO
0.599	Muon::RpcPrepDataContainer_tlp1_RPC_Measurements
0.619	Rec::TrackParticleContainer_tlp1_MuonboyTrackParticles
0.631	Rec::TrackParticleContainer_tlp1_MuonboyMuonSpectroOnlyTrackParticles
0.871	Rec::TrackParticleContainer_tlp1_MooreTrackParticles
1.408	Trk::SegmentCollection_tlp1_ConvertedMBoySegments
1.477	Muon::CscStripPrepDataContainer_tlp1_CSC_Measurements
2.562	Trk::SegmentCollection_tlp1_MooreSegments
2.828	Trk::TrackCollection_tlp3_ConvertedMBoyMuonSpectroOnlyTracks
3.507	MuonMeasurements_tlp1
4.047	Muon::MuonPRD_Container_p2<Muon::MdtPrepData_p2>_MDT_DriftCircles
4.304	RpcPadContainer_p2_RPCPAD
4.728	Trk::TrackCollection_tlp3_ConvertedMBoyTracks
6.841	Trk::TrackCollection_tlp3_MooreTracks
7.356	MdtCsmContainer_p1_MDTCSM

Current sizes, what next?

- For ESDs the bulk of the size is:
  - MDT, RPC RDOs, Tracks and Segments
  - Tracks/Segments consist of Tracking parts (Trk::SegmentCollection and Trk::TrackCollection) & Muon parts (MuonMeasurements_tlp1).
- What can we do?
  - Consider whether we need RDOs still? And if so, what timescale can we drop them on.
    - **Feedback**: need RPCs, but can probably drop everything ~immediately. Needs testing!
  - Compress/'flatten' data - already done for MDT PRDs and gained us about 30%
    - Compression based on e.g. storing Identifier offsets rather than full Identifier. Flattening is copying base-class information into concrete class i.e. no inheritance in persistent class
      - Talk on MDT compression: http://indico.cern.ch/getFile.py/access?contribId=2&resId=0&materialId=0&confId=39834
  - So consider compressing/flattening:
    - Tracking parameters
    - Remaining PRDs
    - ROTs

TGC PRDs

- We currently store PRDs for 3 bunch crossings.
  - Current, prior, post
  - Very small - sum to ~0.5kb/event
- Adding a bcid bitset and merging the containers was discussed in the last Muon SW meeting
- Conclusions
  - Would be more convenient for some clients & would give (tiny) reduction in diskspace....
  - ... gut serious technical challenges (mainly to do with handling schema evolution) & clients will need updating
  - For now, add bcid bit and think again later about merger.

Susumu Oda

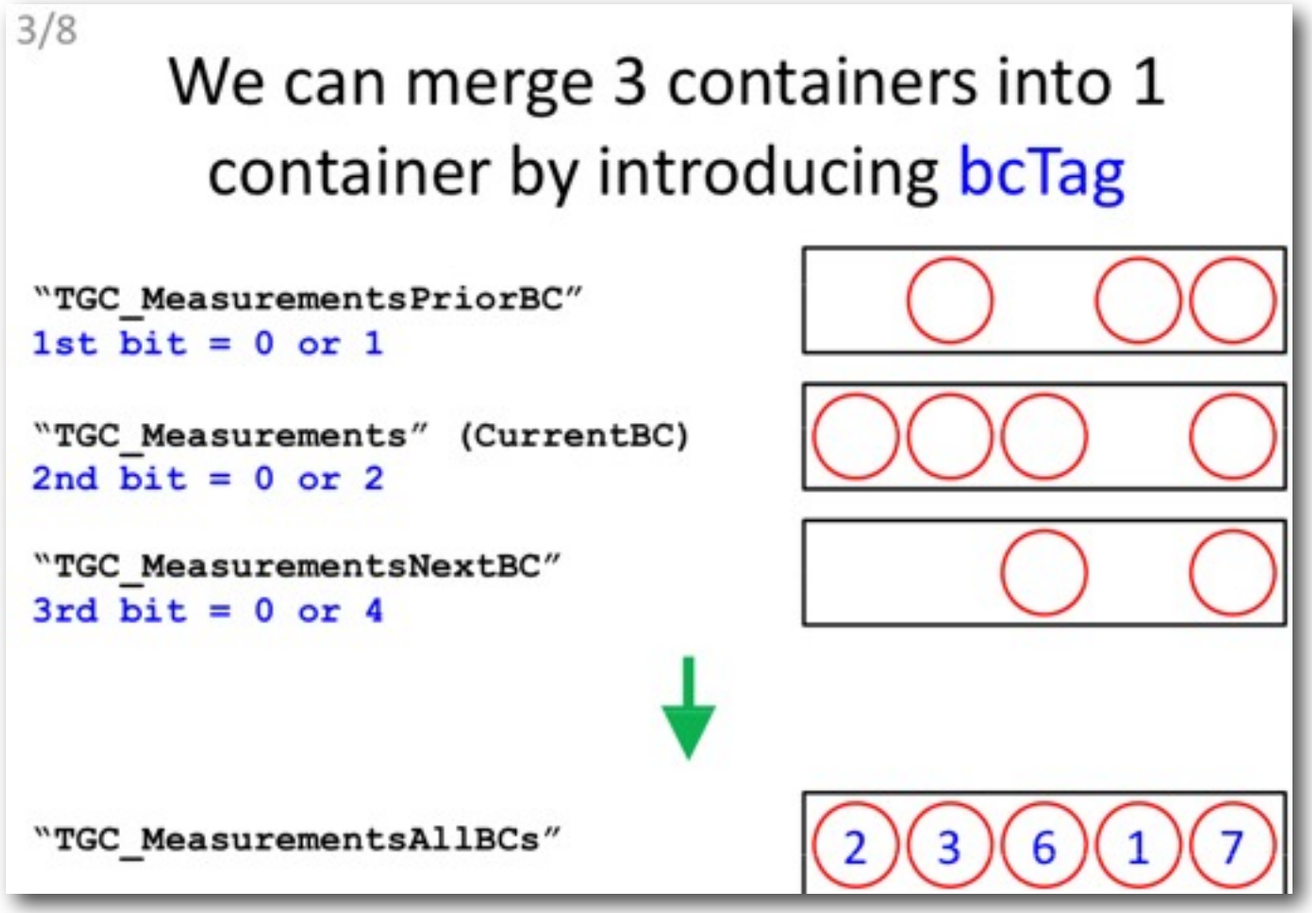

- http://indico.cern.ch/getFile.py/access?subContId=1&contribId=5&resId=0&materialId=slides&confId=114680

Conclusions

- Dropping non-essential RDOs could gain us ~8kB / event, based on current datasets (will obviously scale with cavern background).
  - Feasibility / desirability of dropping e.g. CSC strip PRDs needs to be discussed still.
- Compressing data will probably give us much smaller gains (but will also improve read speed).
  - Worth doing anyway - only problem: manpower.
- Not much to gain by merging TGC containers, and significant technical issues:
  - Deferred for now.
- Aside: muons should also think about dropping support for older releases
  - Not really a ESD reduction gain, but might save some memory by reducing size of convertor libraries (which are fairly large).