Dataflow in Run 3

Fionn Bishop

University of Cambridge

15/03/22

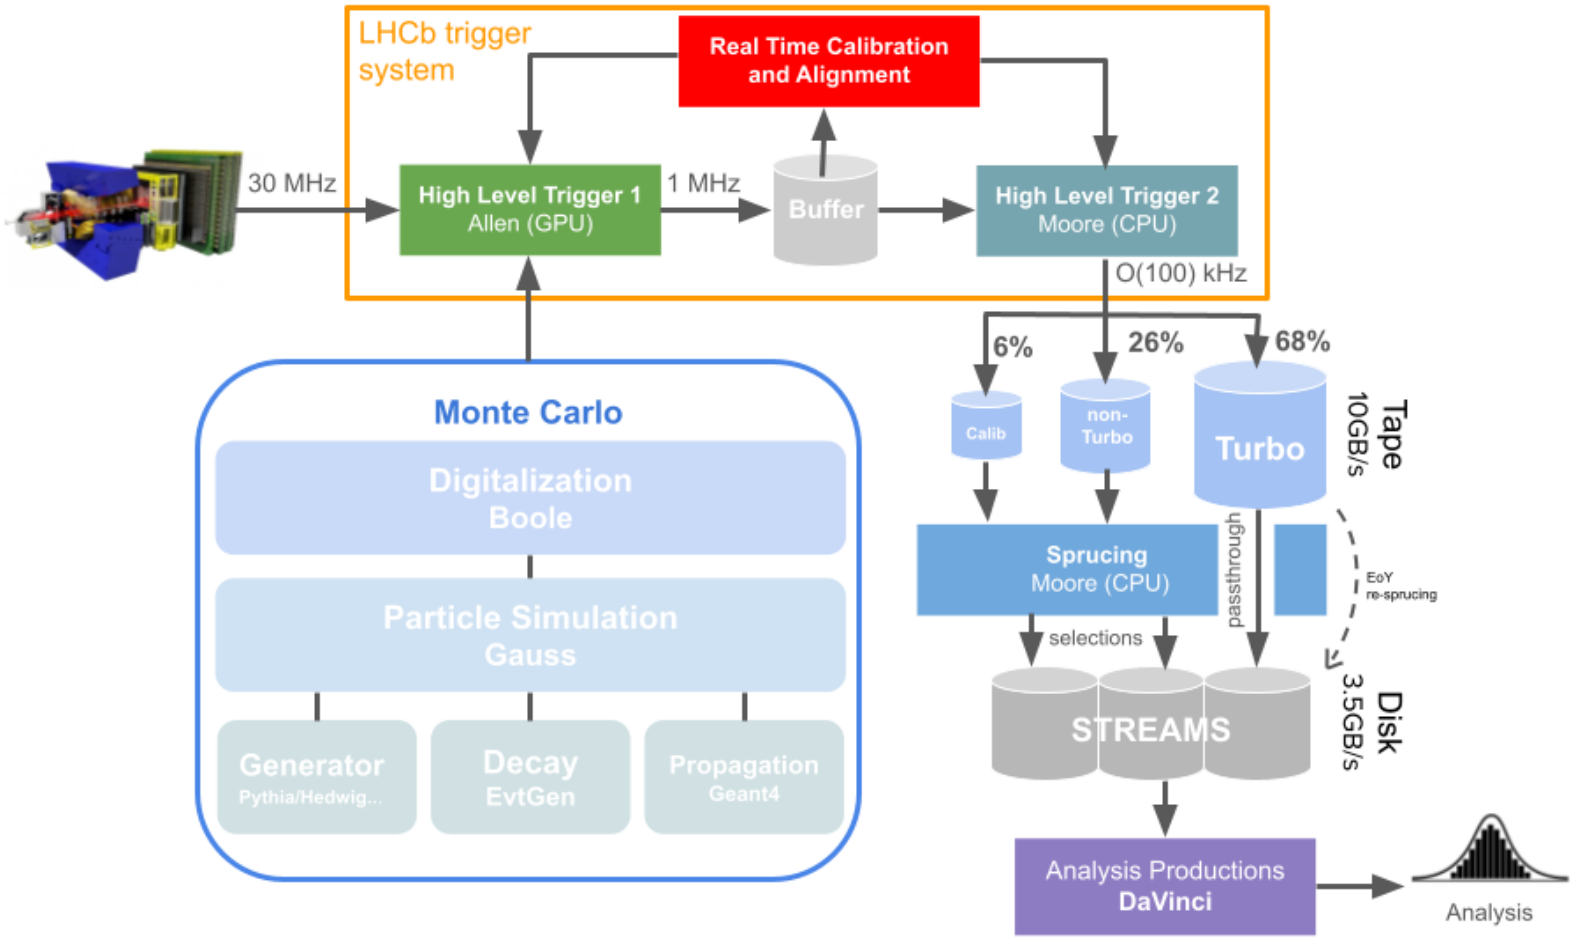
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)
O(100) kHz
6%
26%
68%
Calib
non-Turbo
Turbo
Tape
10GB/s
Monte Carlo
Digitalization
Boole
Particle Simulation
Gauss
Generator
Pythia/Hedwig...
Decay
EvtGen
Propagation
Geant4
Sprucing
Moore (CPU)
passthrough
EoY
re-sprucing
selections
STREAMS
Disk
3.5GB/s
Analysis Productions
DaVinci
Analysis

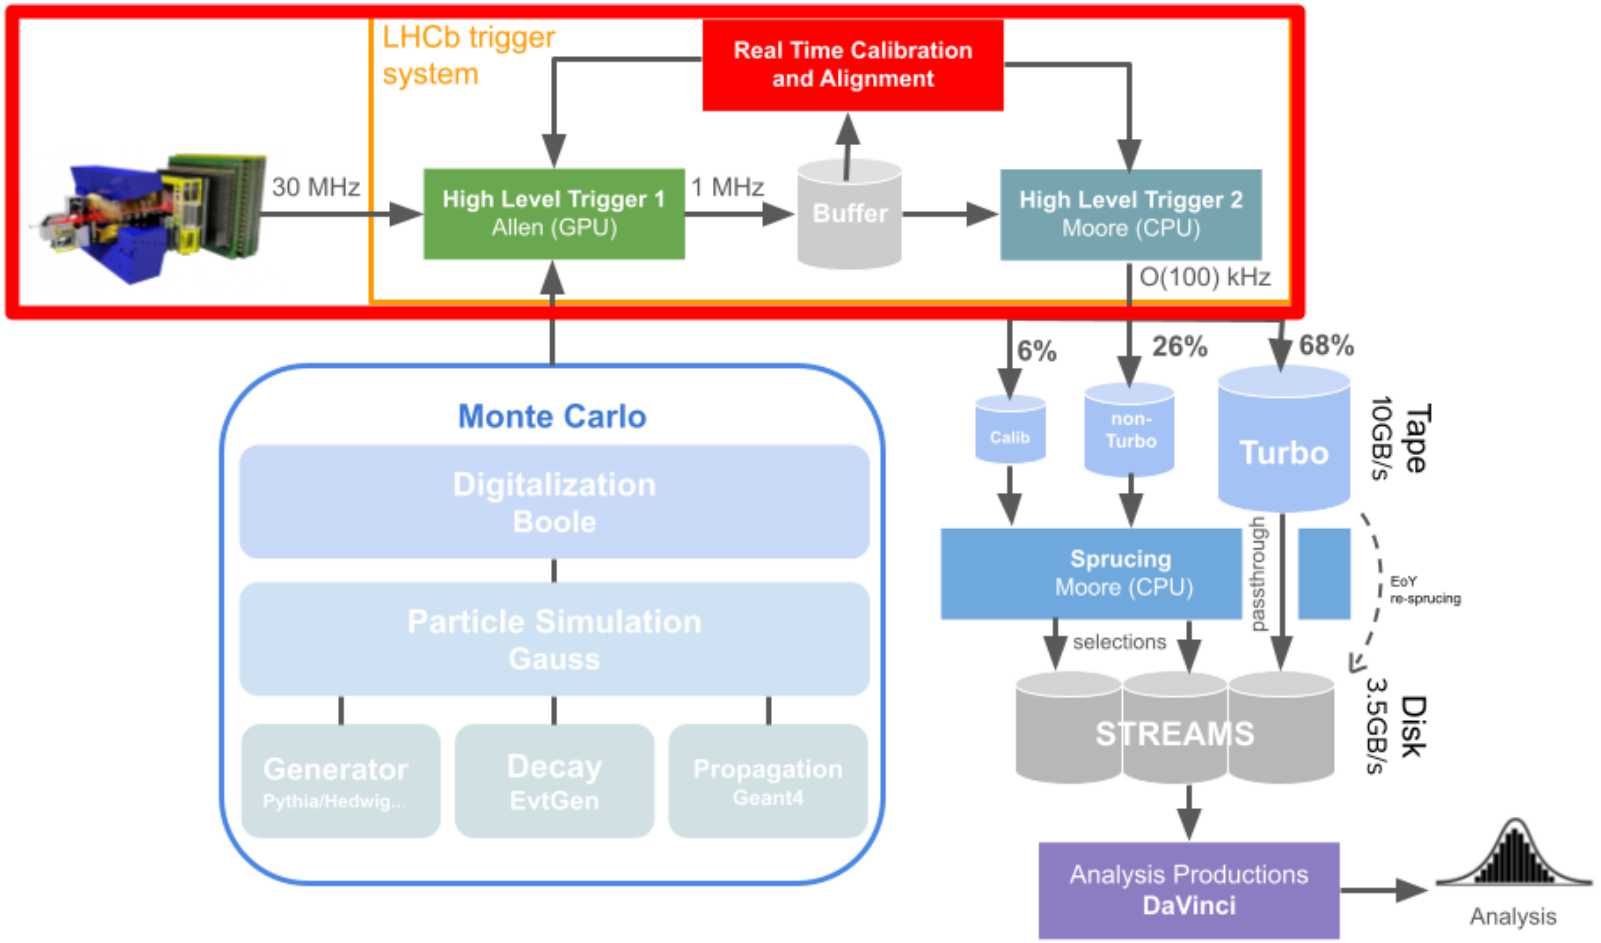
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)
O(100) kHz
6%
26%
68%
Monte Carlo
Digitalization
Boole
Particle Simulation
Gauss
Generator
Pythia/Hedwig...
Decay
EvtGen
Propagation
Geant4
Calib
non-Turbo
Turbo
Tape
10GB/s
Sprucing
Moore (CPU)
passthrough
EoY
re-sprucing
selections
STREAMS
Disk
3.5GB/s
Analysis Productions
DaVinci
Analysis

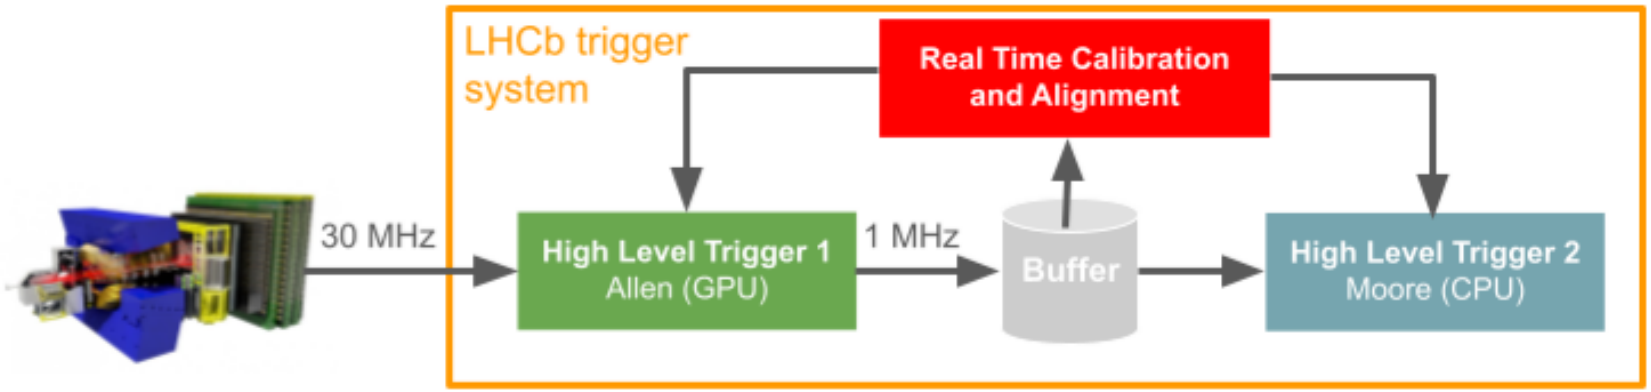

LHCb RTA

WP1 Data Structures	WP3 Selections	WP5 QA
WP2 Reconstruction	WP4 Align & Calib	WP6 R&D

Architecture: Event building servers

- Data acquisition from subdetectors: Field Programmable Gate Array (FPGA) cards
- FPGAs push data to main memory of server
- Receiver application in server interacts with FPGA

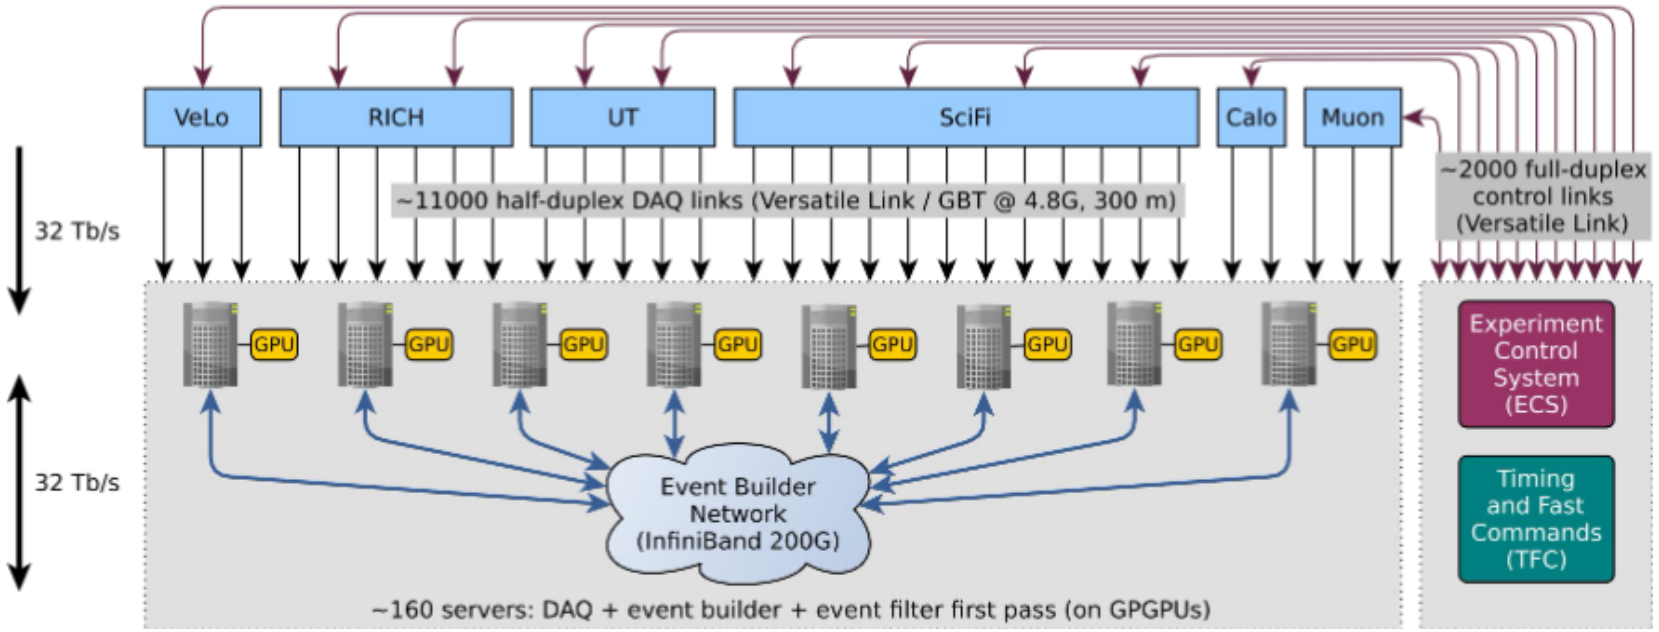

Architecture: Event building servers

- Receiver combines data from $\mathcal{O}(30K)$ bunch-crossings into packets
- Event building:
  - One server chosen as event-builder for each packet
  - All servers send raw data packets to event-builder server through receiver applications
  - Event-builder server combines raw data to form events
- Graphics Processing Units (GPUs): HLT1 filtering

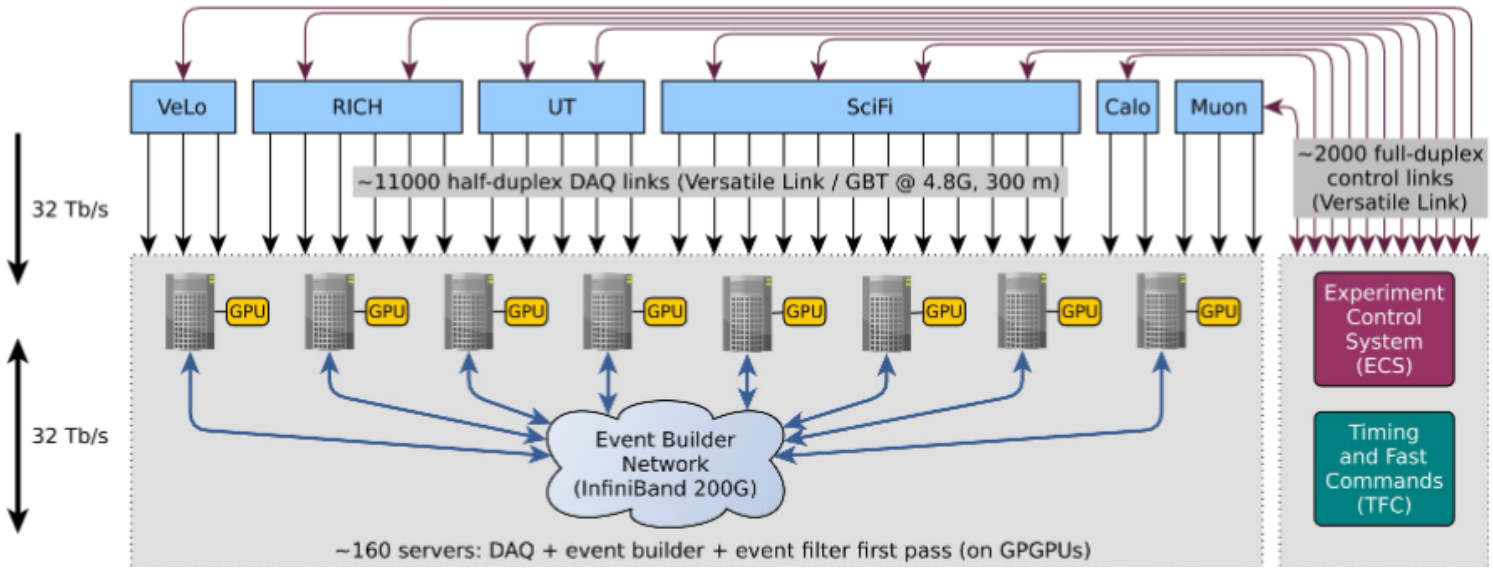

HLT1

- Fast and partial reconstruction
- Inclusive selections with one or two track signatures
- But must run at 30× higher input rate than Run 2!

→ Up to $\mathcal{O}(500)$ **GPUs**

  - Takes advantage of parallesisable nature of track reconstruction
  - Can be hosted in event building servers
    - → Reduce volume of data sent to event filter farm by at least factor 30
  - 200 NVIDIA RTXA5000 currently

- Implemented in Allen: `https://gitlab.cern.ch/lhcb/Allen`
  - Runs on CPU and GPU
- Starterkit: HLT1&Allen (Friday)

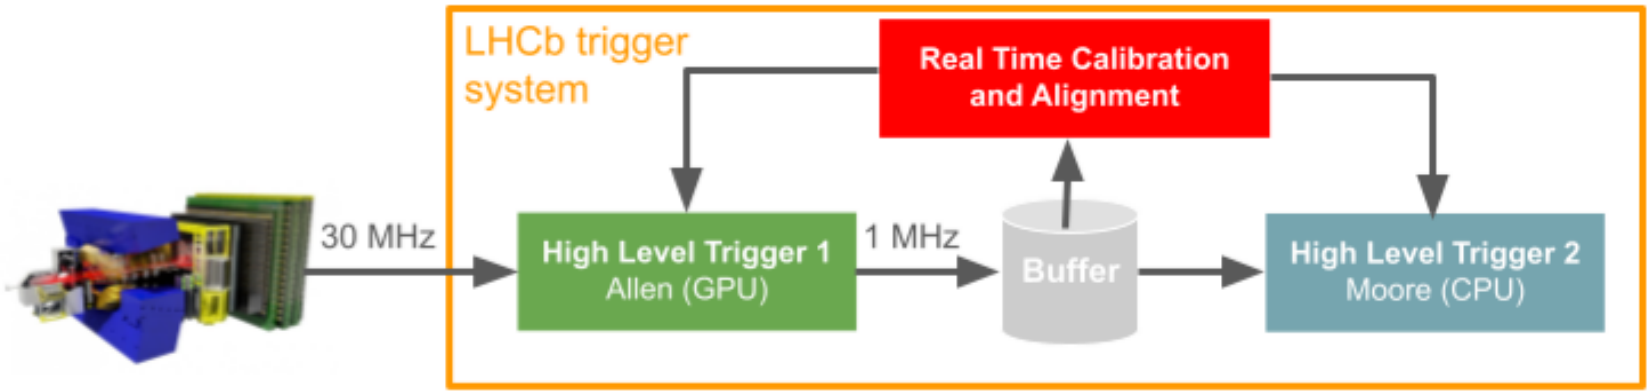
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)

Alignment and Calibration

- Uses dedicated calibration samples selected by HLT1
  - Stored in HLT1 output buffer
- Real-time calculation of alignment and calibration constants
  $\rightarrow$ Improved resolution and PID
- Input to HLT1 and HLT2
  - Offline-quality selection in HLT2

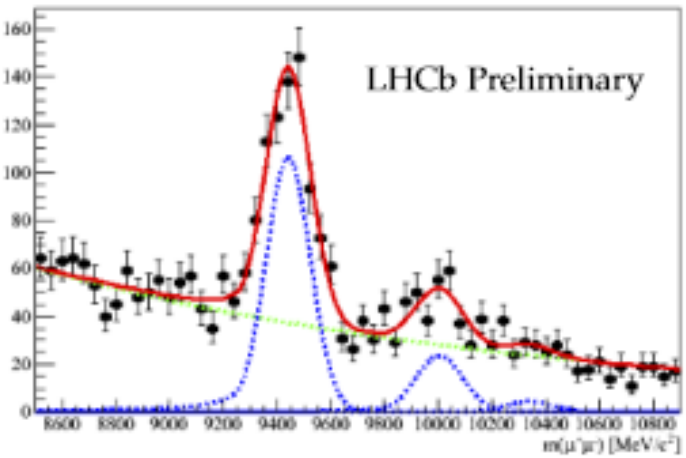

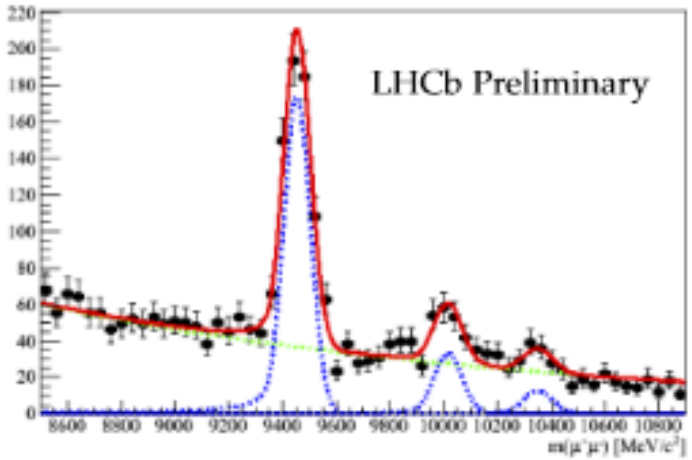

Alignment and Calibration

Alignment:

- Tracking system
- RICH mirrors
- Iterative procedure:
  - Event processing distributed over many computing nodes
  - One node collects input from processing nodes and calculates constants
  - Repeats until alignment has converged

Calibration:

- RICH
- CALO

HLT2

- Online on CPUs
- Full reconstruction
- Complex selection
  - Inclusive or exclusive HLT2 lines
- 10GB/s output to tape
- Implemented in Moore
- Documentation: 'Writing an HLT2 line'

Architecture

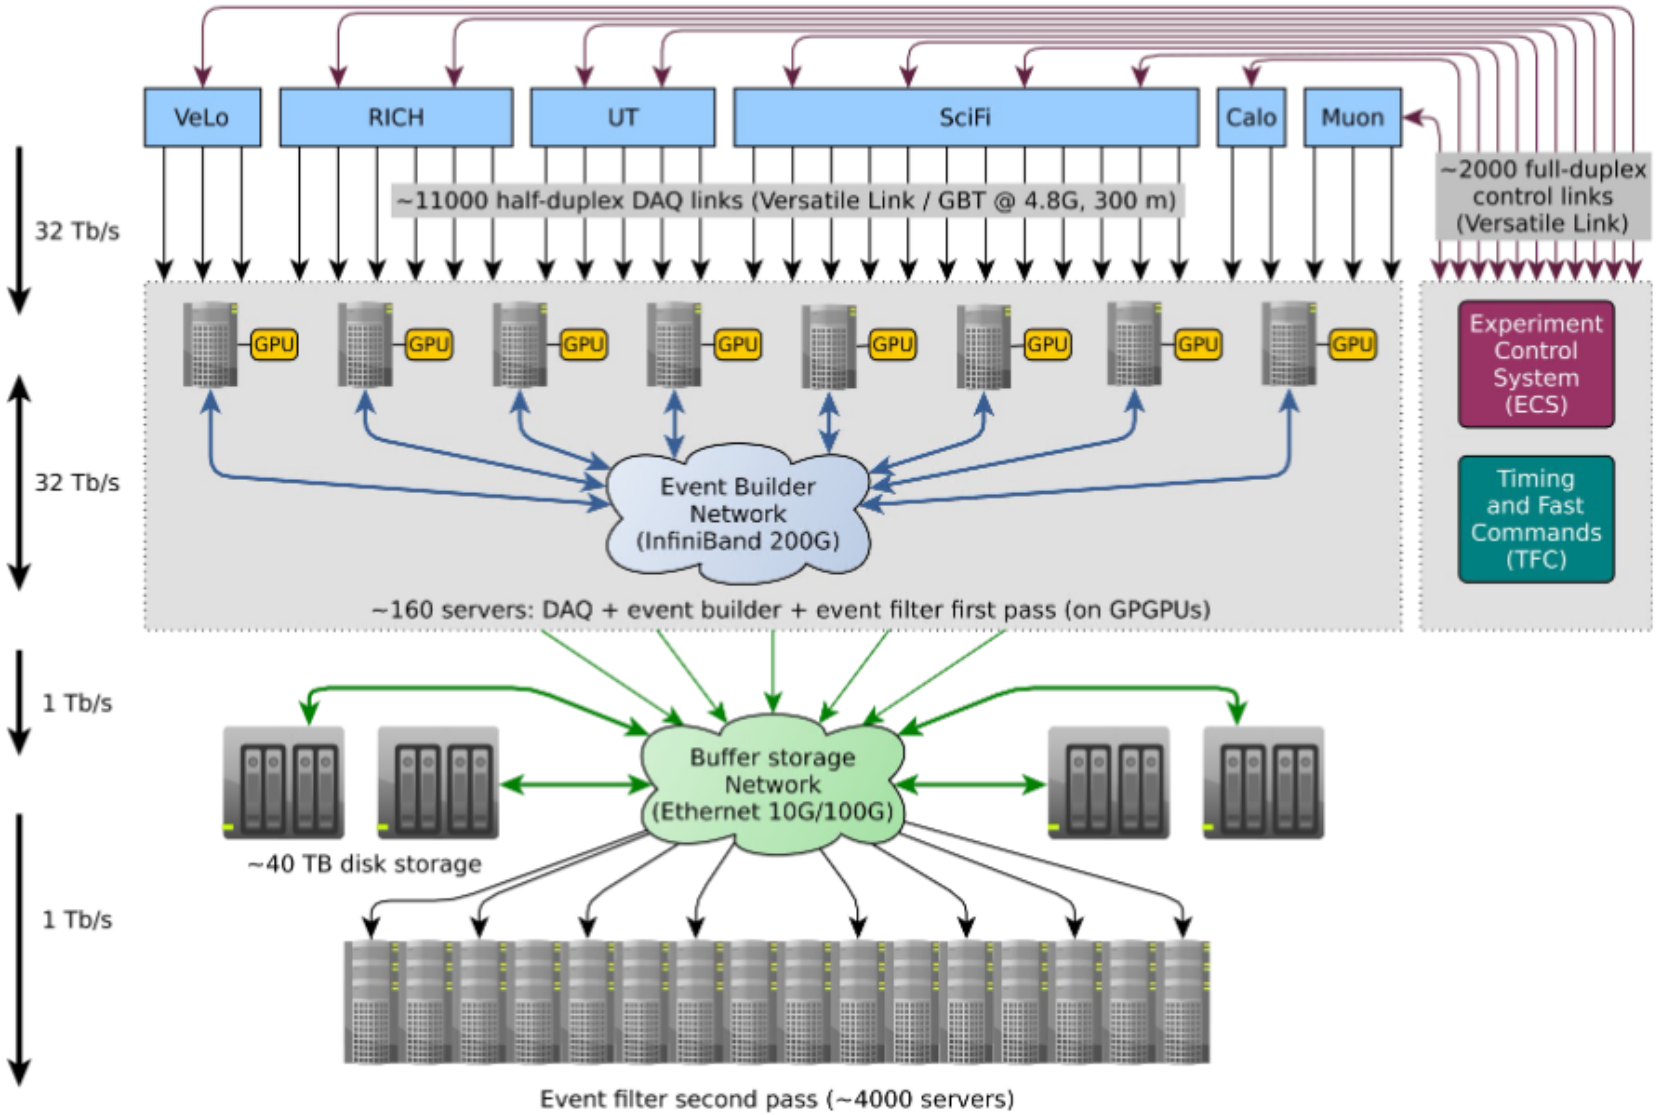

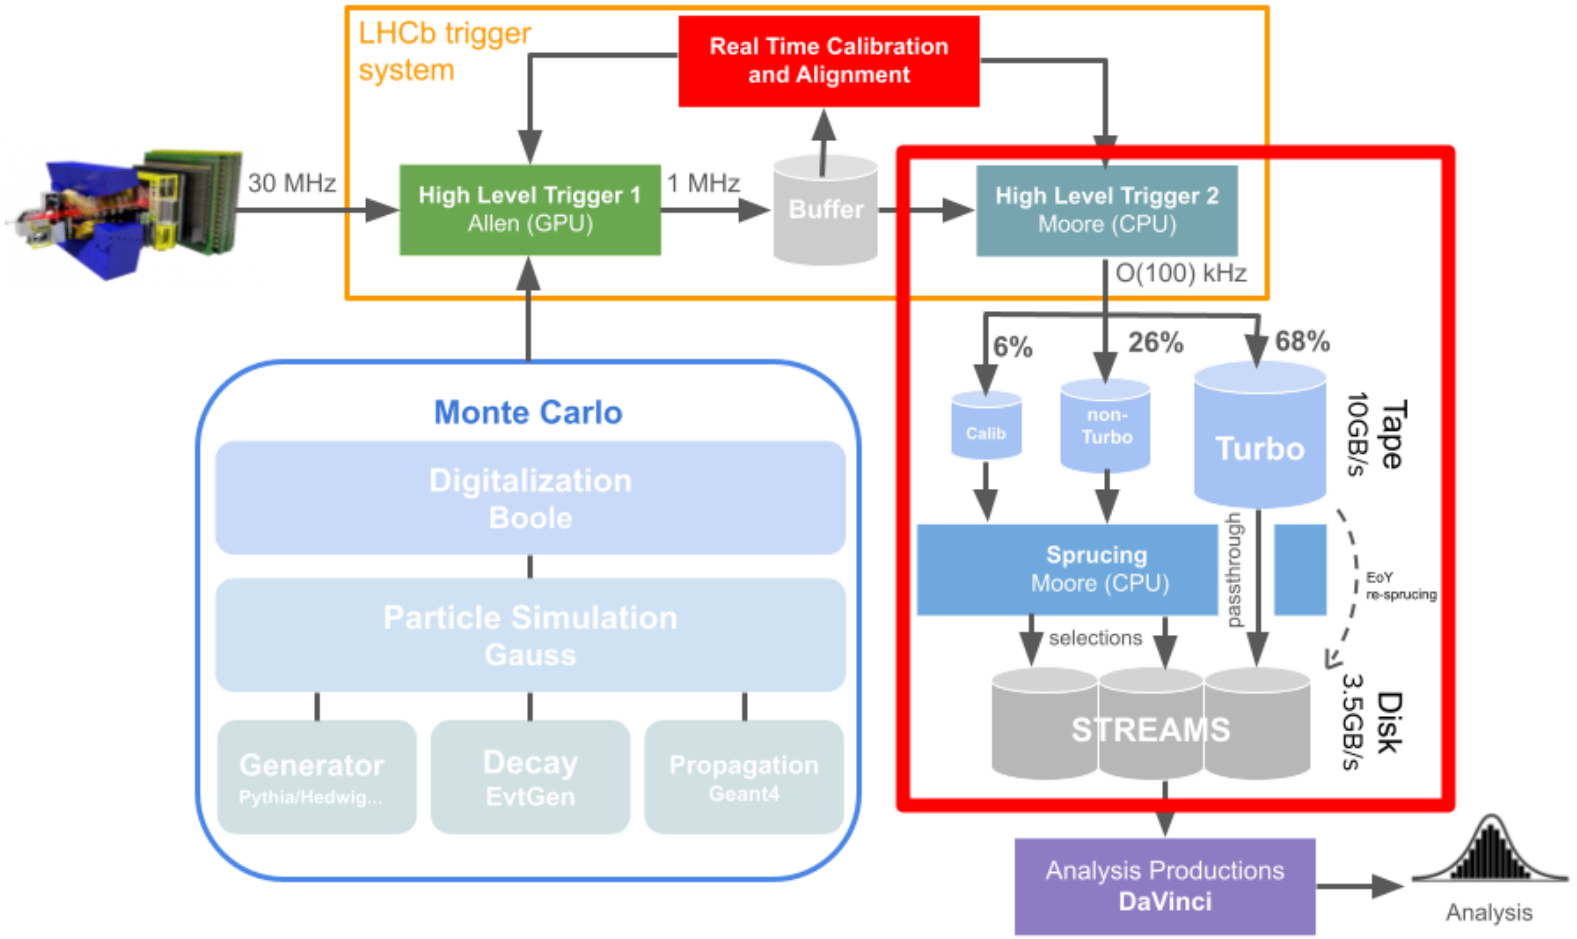
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)
O(100) kHz
6%
26%
68%
Calib
non-Turbo
Turbo
Tape
10GB/s
Sprucing
Moore (CPU)
passthrough
EoY re-sprucing
selections
STREAMS
Disk
3.5GB/s
Monte Carlo
Digitalization
Boole
Particle Simulation
Gauss
Generator
Pythia/Hedwig...
Decay
EvtGen
Propagation
Geant4
Analysis Productions
DaVinci
Analysis

HLT2 and Sprucing

- Sprucing:
  - 'Skimming'+'pruning'
  - Offline
  - Implemented in Moore with same base code as HLT2
  - Exclusive selection or passthrough mode
  - DPA project
- Data in streams on disk available for making ntuples
- Two routes: **Turbo** and **Sprucing** models

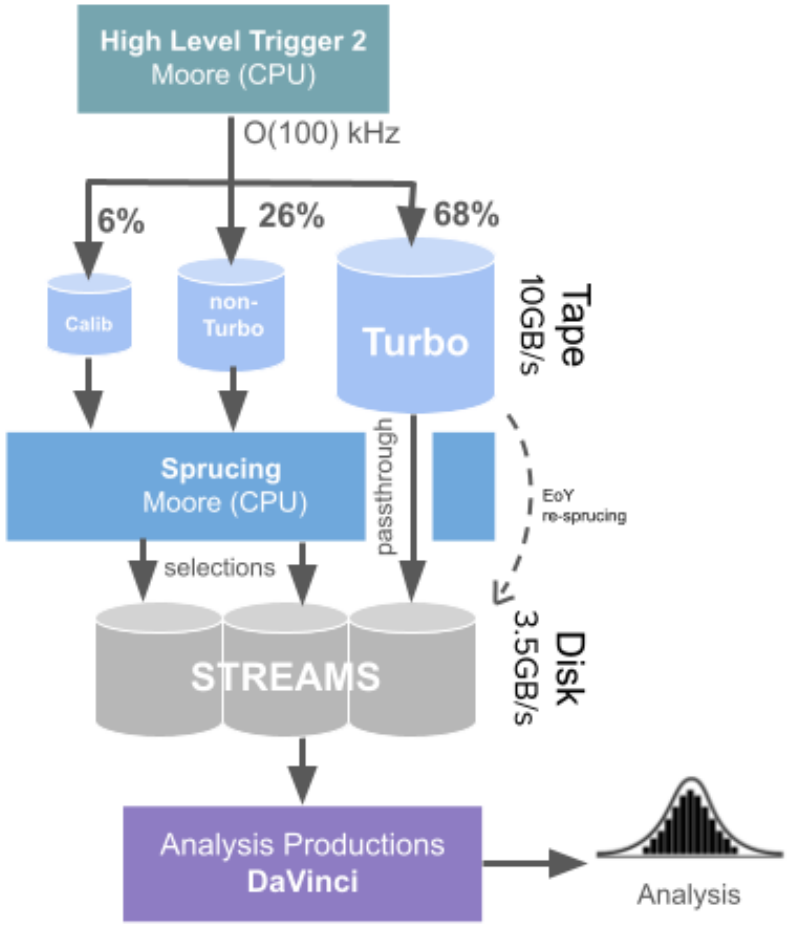

Turbo model

- Exclusive selection in HLT2
  - Offline-like quality due to real-time alignment and calibration
- Sprucing in passthrough mode $\rightarrow$ Save to disk
- Baseline in Run 3
- 68% events in 25% bandwidth $\rightarrow$ Most Physics per data on tape

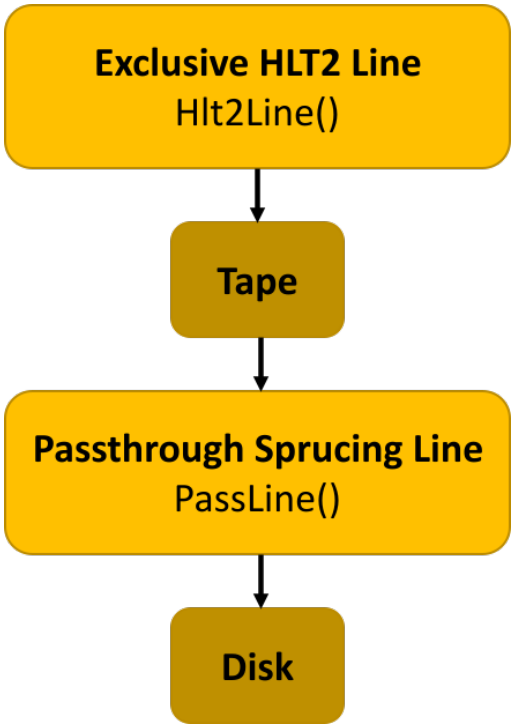

Sprucing model

- Tighter, exclusive selection over output of inclusive HLT2 lines
  - HLT2 line must have `persistreco=True`
- Similar to Run 1+2 stripping
- For selections with throughput too large to go straight to disk
- Input rate 5.9 GB/s*
- Output rate 0.8 GB/s*
  - Small decrease in event rate
  - Large decrease in event size

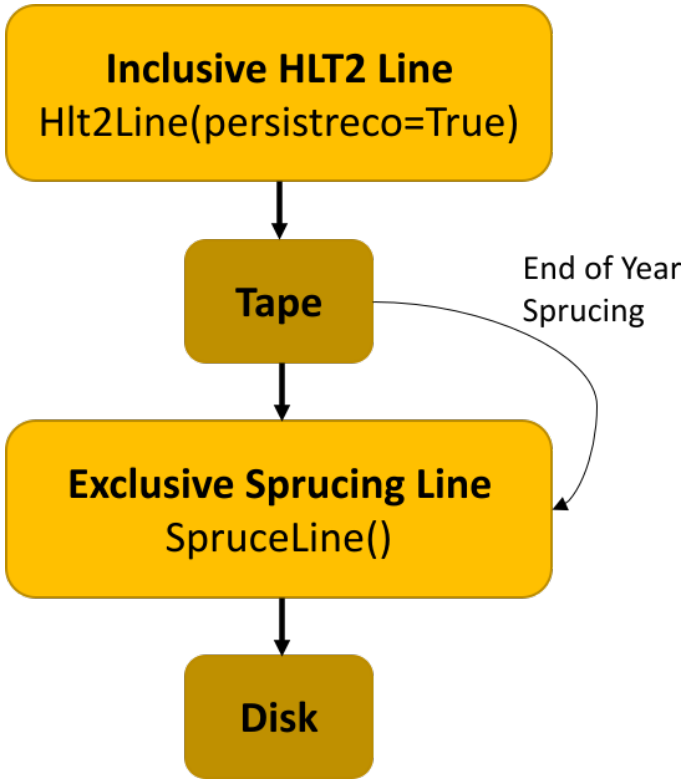

*LHCb-TDR-018

HLT2 and Sprucing

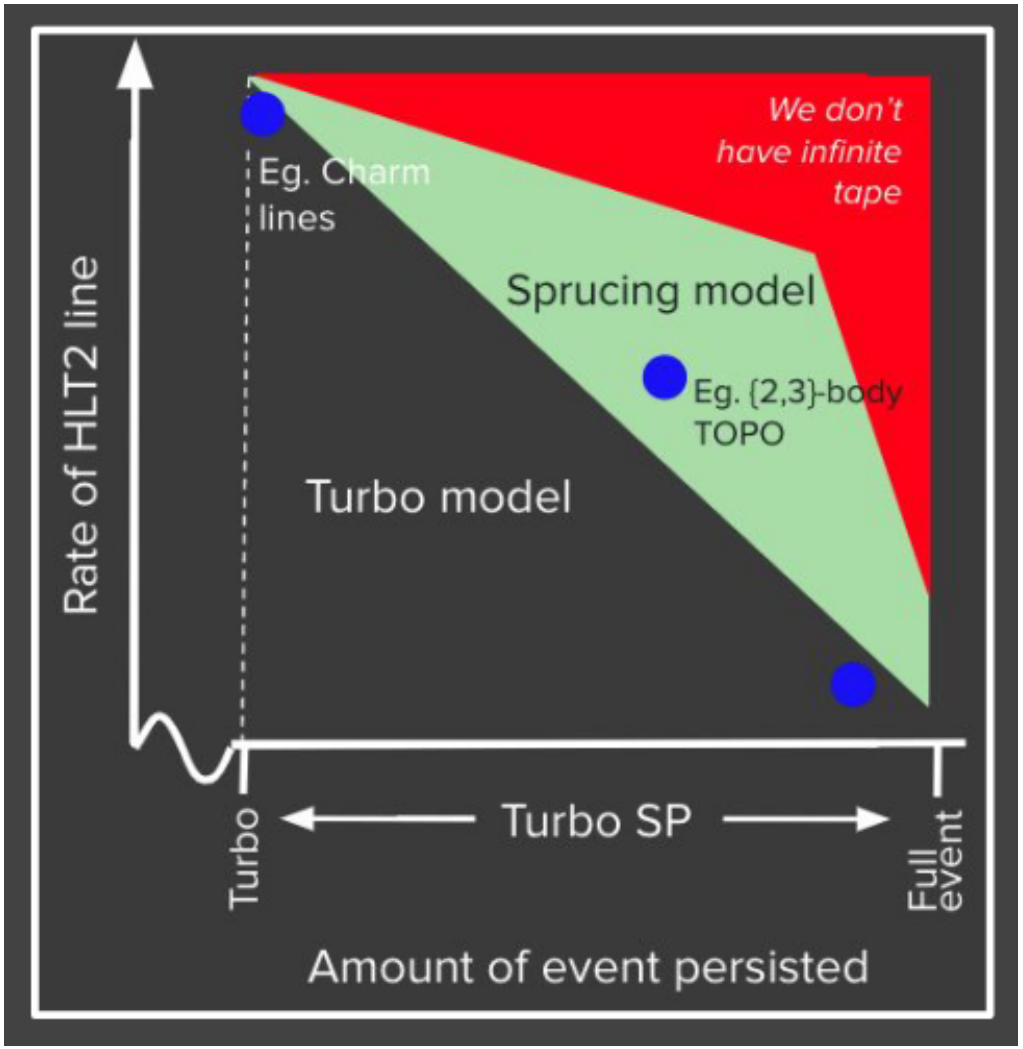

- Pure turbo only saves **signal candidate**
  - → Smaller event size → More events
- Can persist other information if specifically enabled
  - e.g. extra tracks, ECAL clusters

HLT2 and Sprucing: Selections

- HLT2 and sprucing use same code framework
- Use new **ThOr** functors
  - Python representation in configuration files
- ThOr functors reference

```
@configurable
def filter_particles(make_particles,
                     make_pvs=make_pvs,
                     trchi2todof_max=3,
                     mipchi2_min=4,
                     pt_min=250 * MeV,
                     p_min=2 * GeV):
    pvs = make_pvs()
    code = require_all(F.PT > pt_min, F.P > p_min, F.CHI2DOF < trchi2todof_max,
                       F.MINIPCHI2(pvs) > mipchi2_min)
    return ParticleFilter(make_particles(), F.FILTER(code))
```

BPVIPCHI2()	F.BPVIPCHI2(pvs)
BPVLTIME()	F.BPVLTIME(pvs)
CHI2VXNDOF	F.CHI2DOF
ETA	F.ETA

Starterkit:

- HLT2 line development (this morning)
- Sprucing (this afternoon)
- Testing HLT2 lines (Thursday)

HLT2 and Sprucing

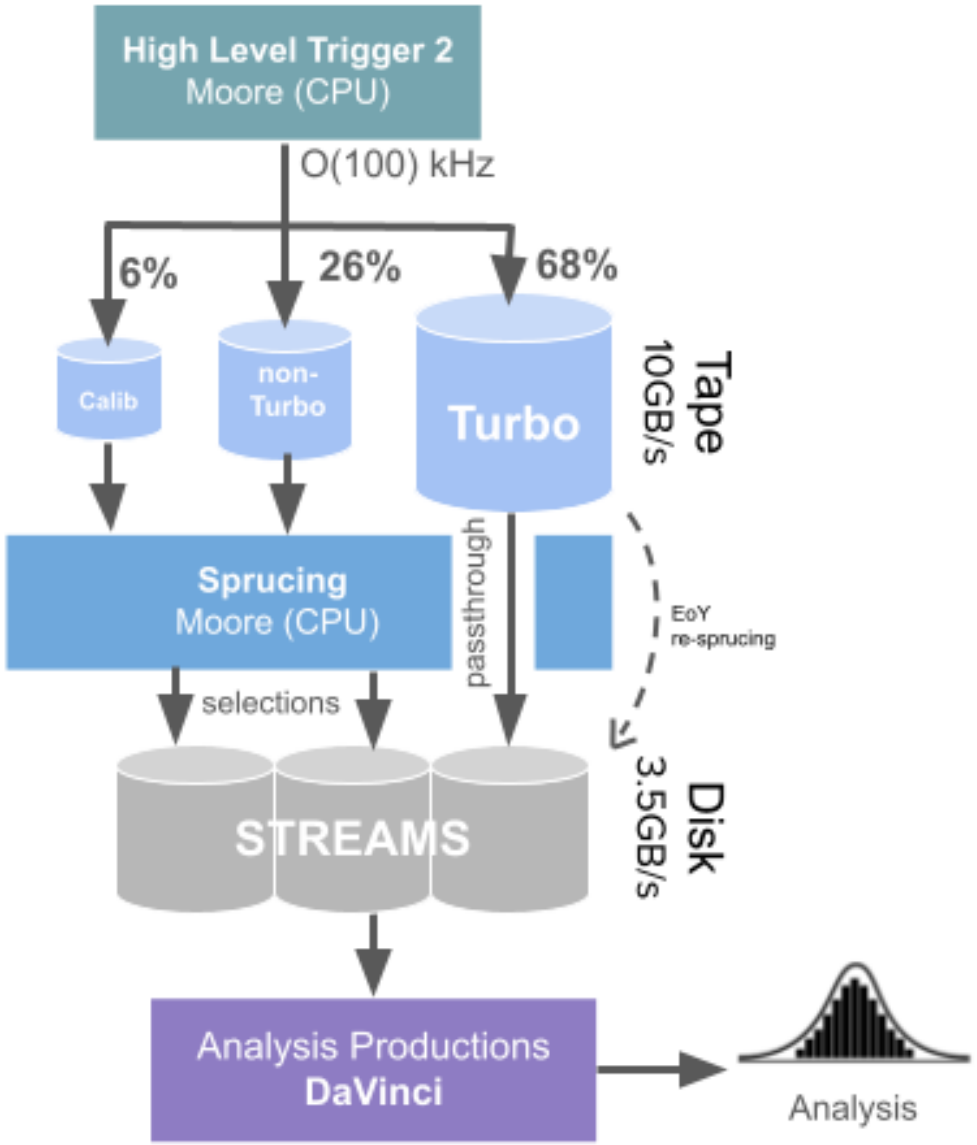

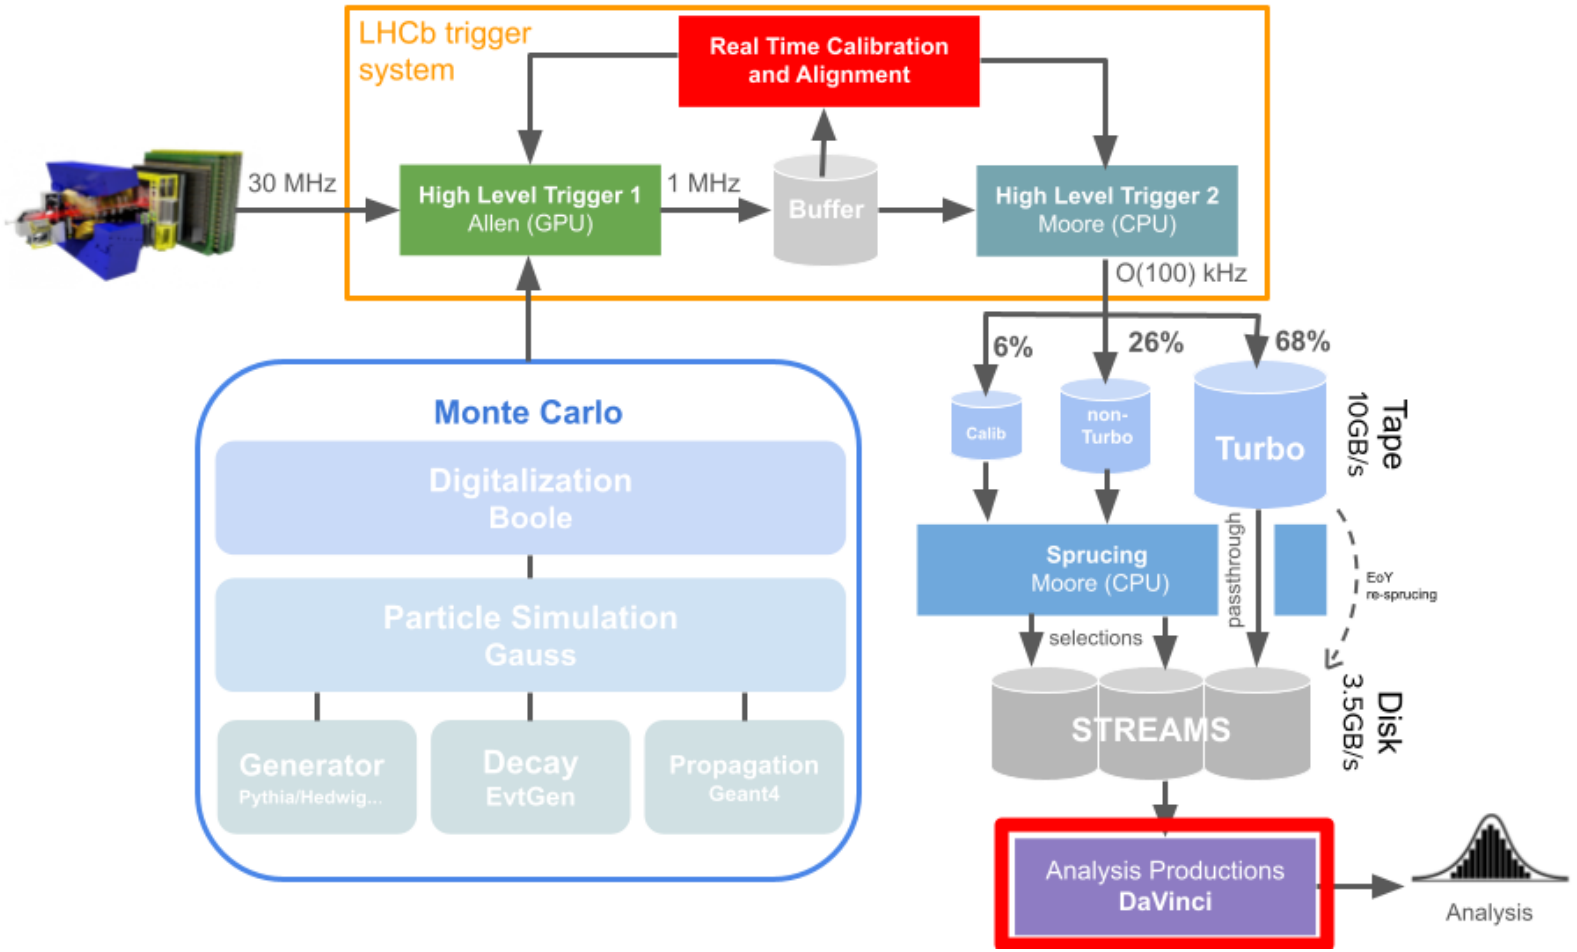
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)
O(100) kHz
6%
26%
68%
Calib
non-Turbo
Turbo
Tape
10GB/s
Monte Carlo
Digitalization
Boole
Particle Simulation
Gauss
Generator
Pythia/Hedwig...
Decay
EvtGen
Propagation
Geant4
Sprucing
Moore (CPU)
passthrough
EoY
re-sprucing
selections
STREAMS
Disk
3.5GB/s
Analysis Productions
DaVinci
Analysis

DaVinci and FunTuple

Major changes:

- ThOr functors
- FunTuple

FunTuple:

- Replaces TupleTools
- Greater flexibility for analyst in choice of variables
- Reduce storage and computing use

- Starterkit: FunTuple (tomorrow)
- Documentation: DPA WP3 Offline Analysis Tools

Analysis Productions

Before: Jobs submitted to grid using Ganga
Now: **Analysis productions** preferred

- Centralised and automated ntuple-making system
  - Runs DaVinci
- Runs on DIRAC
- Requires:
  - DaVinci options file
  - yaml configuration file
- Can already be used for Run 1+2 analyses

Analysis Productions

- Improved analysis preservation:
  - Ntuples well-preserved
  - Automatic preservation of configuration
- Improved sharing of ntuples between different analyses
- Less monitoring from user side
- Automatic testing of options files

Analysis Productions

- Access through straightforward web interface
- Starterkit: Analysis productions (Friday)
- Documentation: DPA WP2 Analysis Productions

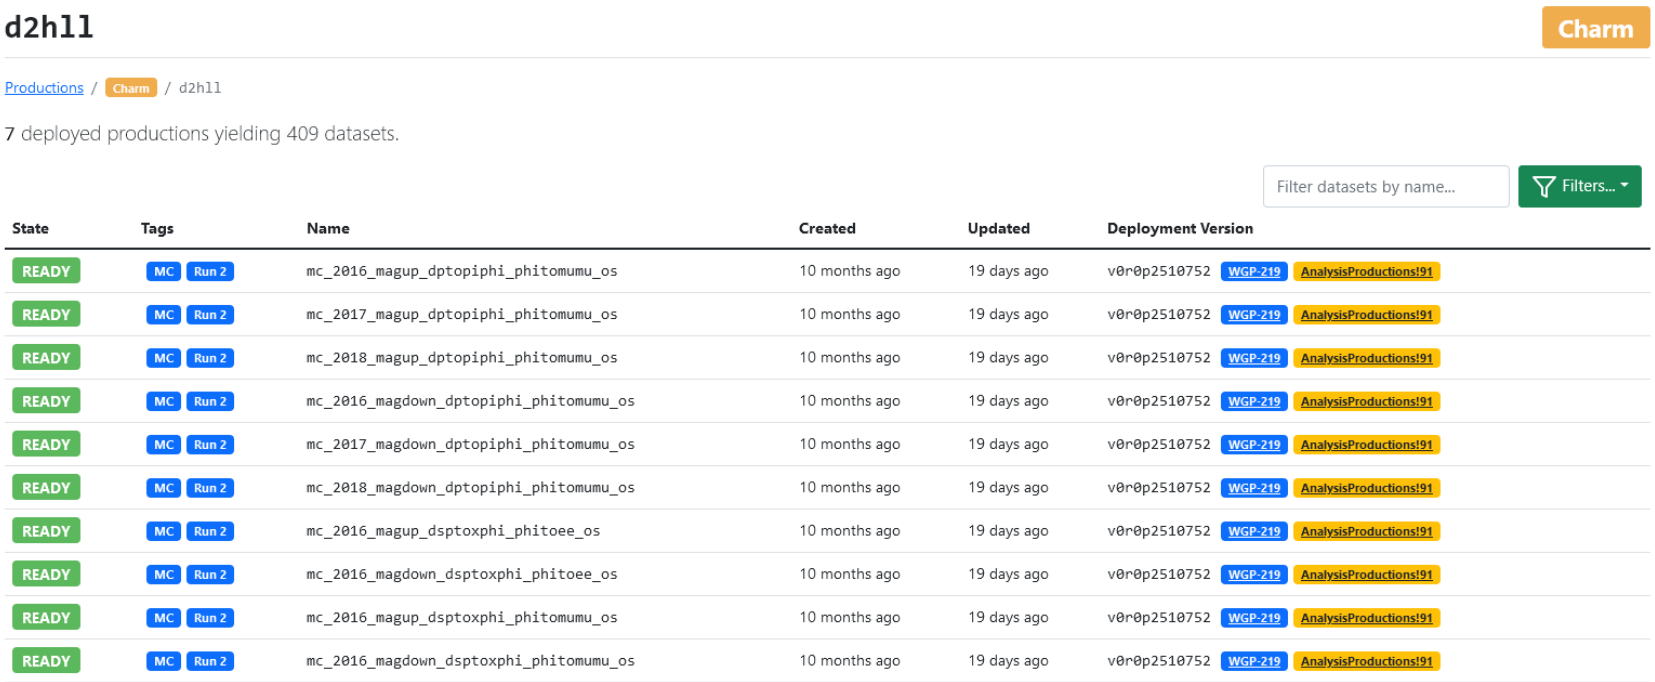

d2h11 Charm

Productions / Charm / d2h11

7 deployed productions yielding 409 datasets.

Filter datasets by name... Filters...

State	Tags	Name	Created	Updated	Deployment Version
READY	MC Run 2	mc_2016_magup_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2017_magup_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2018_magup_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2016_magdown_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2017_magdown_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2018_magdown_dptopiphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2016_magup_dsptoxphi_phitoee_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2016_magdown_dsptoxphi_phitoee_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2016_magup_dsptoxphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91
READY	MC Run 2	mc_2016_magdown_dsptoxphi_phitomumu_os	10 months ago	19 days ago	v0r0p2510752 WGP-219 AnalysisProductions!91

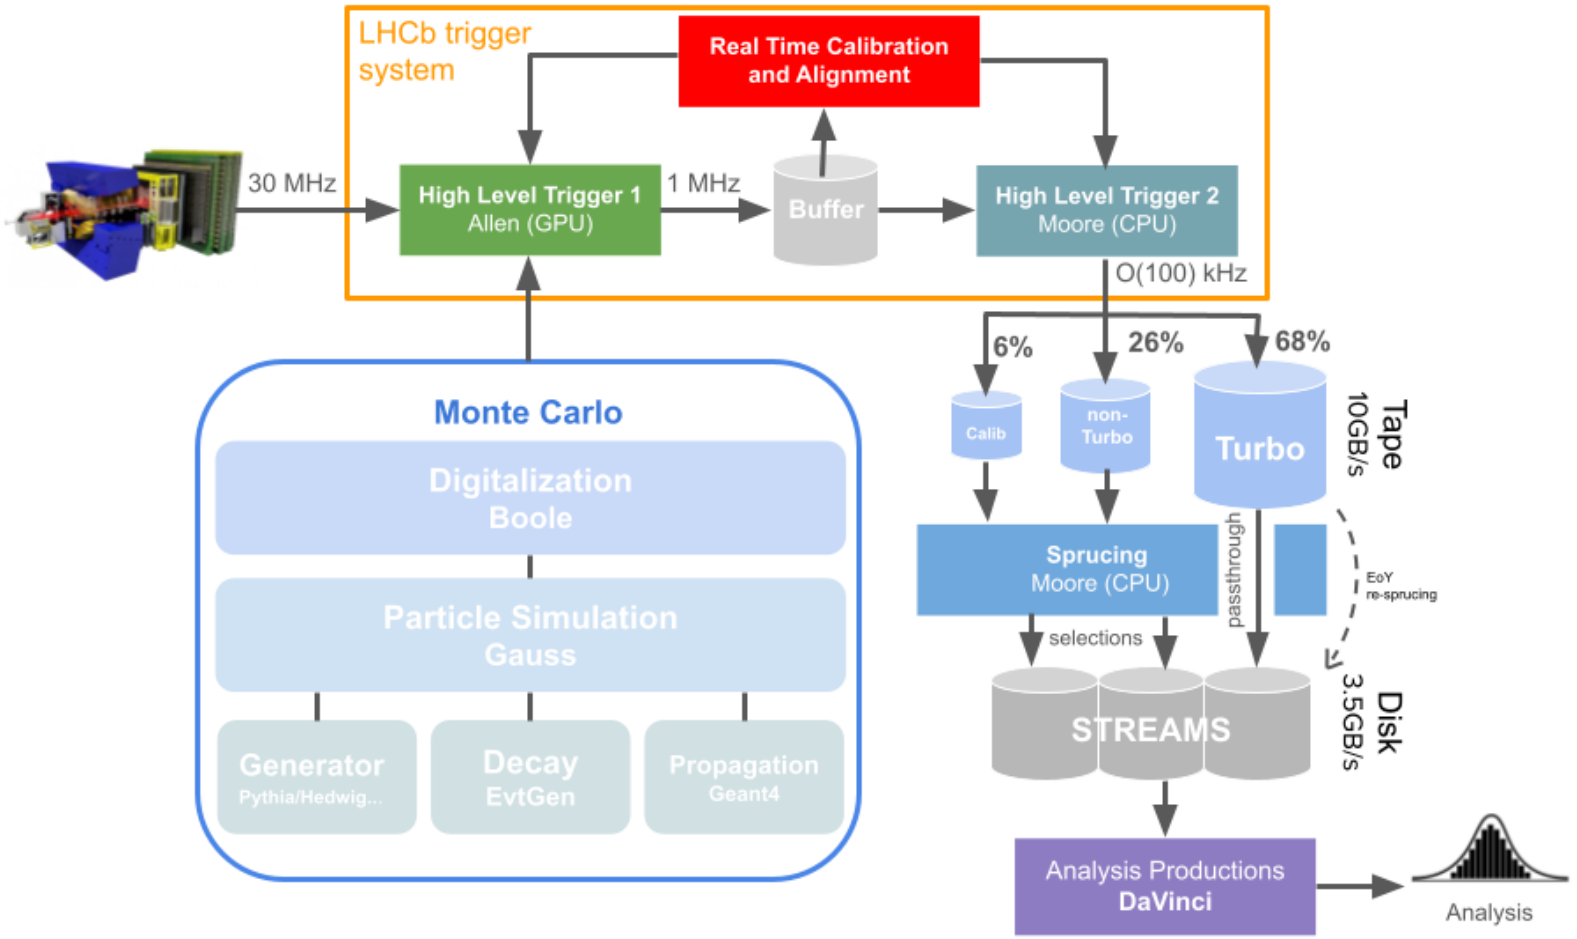
LHCb trigger system
Real Time Calibration and Alignment
30 MHz
High Level Trigger 1
Allen (GPU)
1 MHz
Buffer
High Level Trigger 2
Moore (CPU)
O(100) kHz
6%
26%
68%
Calib
non-Turbo
Turbo
Tape
10GB/s
Monte Carlo
Digitalization
Boole
Particle Simulation
Gauss
Generator
Pythia/Hedwig...
Decay
EvtGen
Propagation
Geant4
Sprucing
Moore (CPU)
passthrough
EoY
re-sprucing
selections
STREAMS
Disk
3.5GB/s
Analysis Productions
DaVinci
Analysis