HPC at COMPASS: The Blue-Waters project

Riccardo Longo
On behalf of COMPASS

CERN PRACE workshop
CERN
22 October 2018

The COMPASS experiment

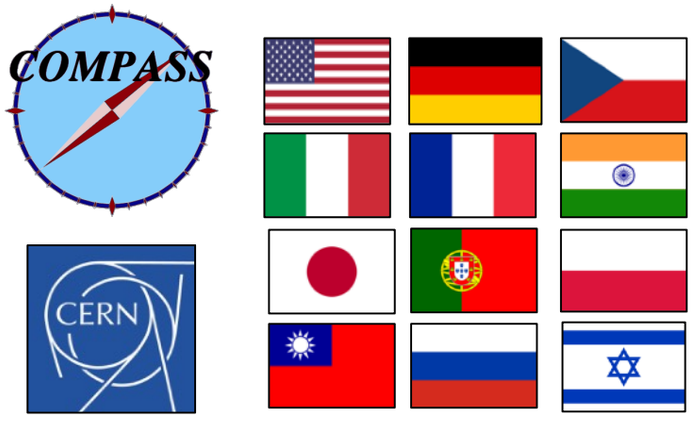

- High tracking and PID power;
- Polarized targets (NH_3, ^6LiD);
- μ, π, K, p beams;
- Wide polar acceptance (± 180 mrad);

- Fixed target experiment;
- SPS North Area;
- M2 beam-line;
- Data taking started in 2002;
- Wide physics programme

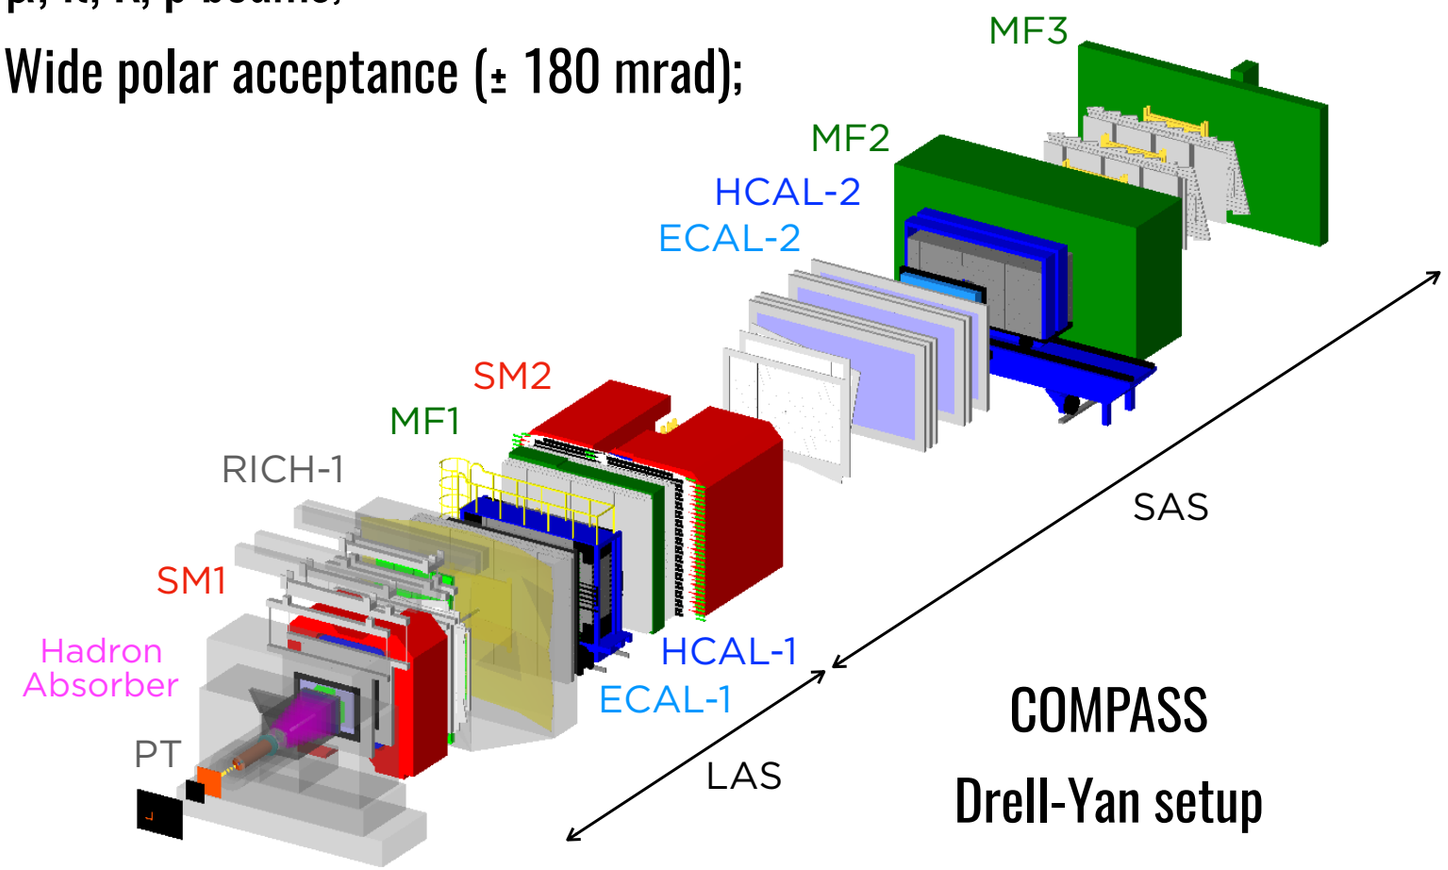

COMPASS

Drell-Yan setup

Phase I (2002 - 2011)

- Hadron spectroscopy
- Nucleon spin structure via polarized DIS and SIDIS

Phase II (2012 - 2021)

- Primakoff + DVCS pilot run (2012)
- Drell-Yan (2015, 2018)
- DVCS + Unpolarized SIDIS (2016-2017)
- T-polarized SIDIS (D target) (2021 [new])

The COMPASS experiment

- Fixed target experiment;
- SPS North Area;
- M2 beam-line;
- Data taking started in 2002;
- Wide physics programme

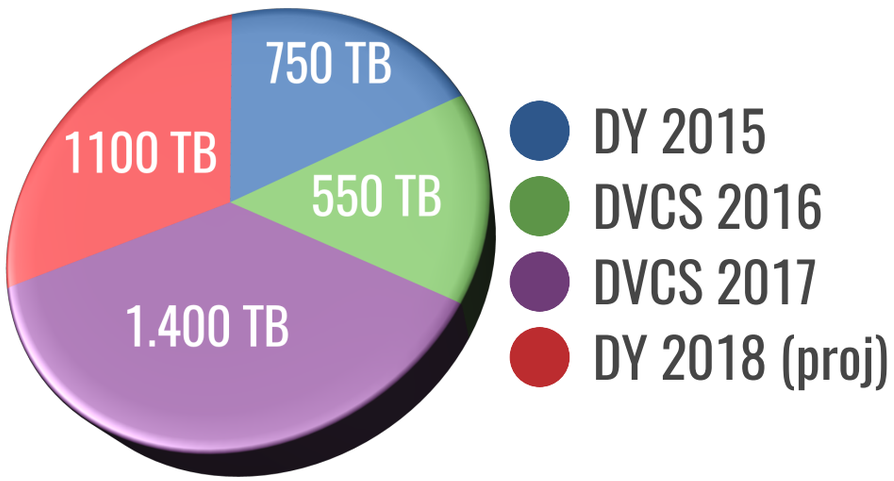

- ~4 PB of raw data to be processed;
- Real Data production campaign are carried out at CERN;

Phase I (2002 - 2011)
- Hadron spectroscopy
- Nucleon spin structure via polarized DIS and SIDIS

Phase II (2012 - 2021)
- Primakoff + DVCS pilot run (2012)
- Drell-Yan (2015, 2018)
- DVCS + Unpolarized SIDIS (2016-2017)
- T-polarized SIDIS (D target) (2021 [new])

- Many precision measurements:
  - Massive MC samples generation (~1e9 events);
  - Precise knowledge of the response of the apparatus;
  - Several analysis to be carried out;

Heavily demanding tasks from computing and storage power point of view!

COMPASS MC-chain

TGEANT

- **MC generation:**
  - Event generation carried out using event generators (Pythia[6-8], HepGen++, LEPTO, Django…);
  - Setup response simulated via Geant4;
  - Precision of the simulation, time window chosen to simulate the pile-up, number of events, etc… determine the processing time;

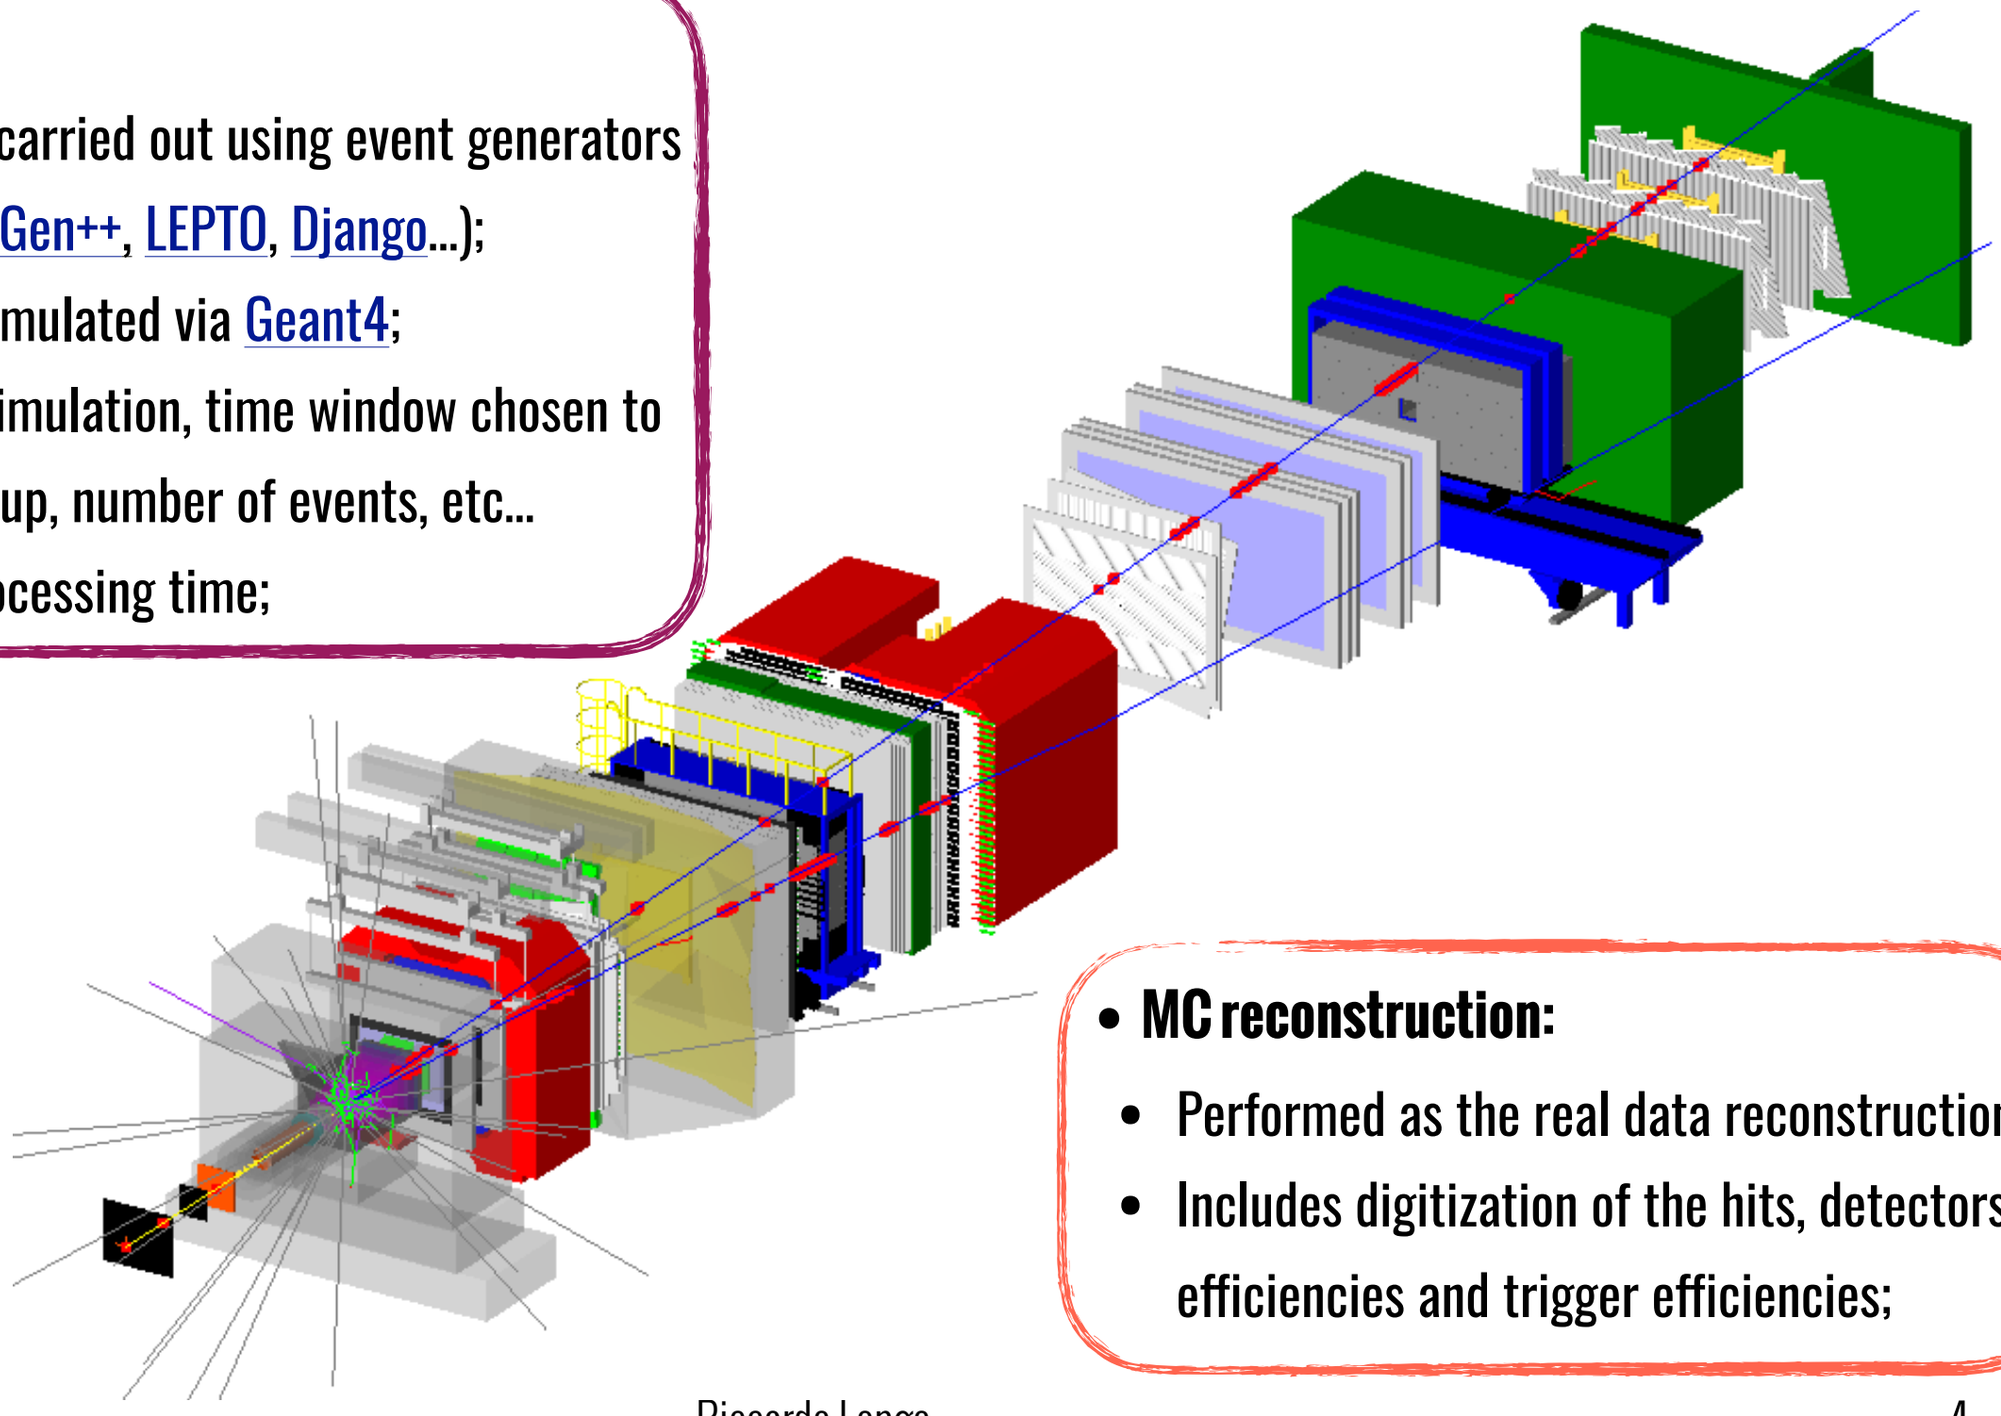

- **MC reconstruction:**
  - Performed as the real data reconstruction;
  - Includes digitization of the hits, detectors efficiencies and trigger efficiencies;

The Blue Waters project

- Petascale supercomputer at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign.
- The NSF awarded $208 million for the Blue Waters (BW) project. Installation done in 2012.
- Currently, BW is granted until 31th March 2019.

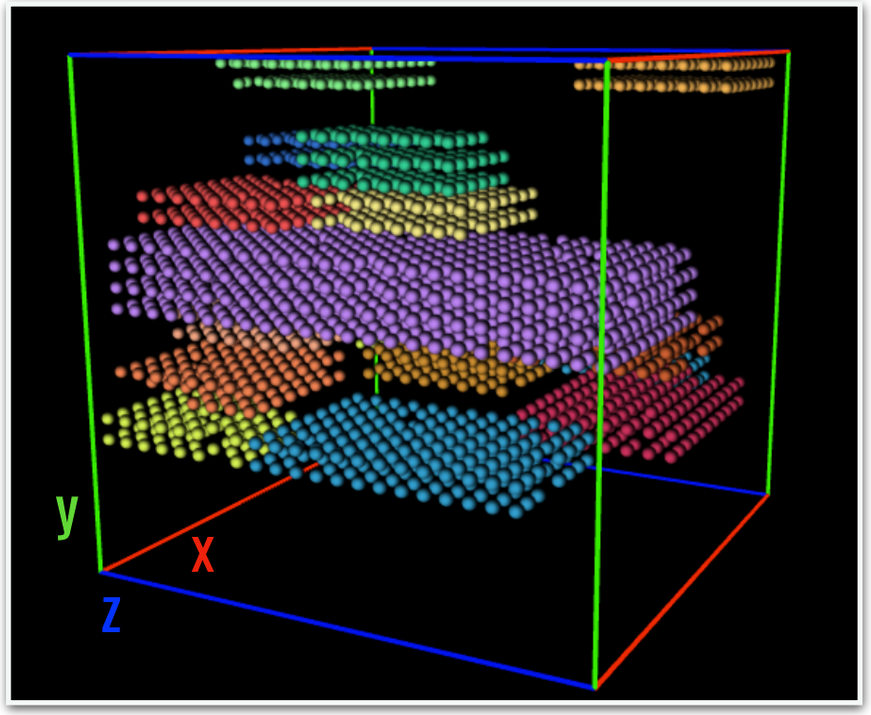

Memory	Disk storage	Tape storage	Speed
1.5 petabytes	26 petabytes	380 petabytes	13.3 petaFLOPS

node type	remark	CPU's/node	memory/node	memory per CPU	# of nodes
XK	heterogeneous x86 CPU & ("accelerated") GPU	16	32 GB	2 GB	4.200
XE	dual CPU x86 processor	32	64 GB	2 GB	22.500

NSF allocation for COMPASS @ BW

- Computing power at BW is distributed in allocations of node-hours;
- Feasibility of the proposals for allocations are tested in “exploratory” phases;
- Currently COMPASS is one of the biggest allocations granted at BW;

	Jan	Feb	Mar	Apr	May	June	July	Aug	Sep	Oct	Nov	Dec
2016			✉						✉		✉	
2017												
2018										NOW		
2019												

node hours (~~used up~~) ~~40,000~~ ~~200,000~~ ~~9,440,000~~ 9,440,000

✉ proposal submission

Exploratory **Campus** **PRAC-A** **PRAC-B** Allocations

	disk	tape	purging?
home/user	1 TB	5 TB	no
project	500 TB	6,000 TB	no
scratch	1,000 TB	—	after 30 days

- Extensive storage availability to cope with the petascale resources at the disposal.

NSF allocation for COMPASS @ BW

- Constant deployment of resources on a daily-basis, allowing to push forward COMPASS MC production and analyses.
- ~ **160 M Core-hours** used in the last 6 months.

1 node hour
=
32 Core hours

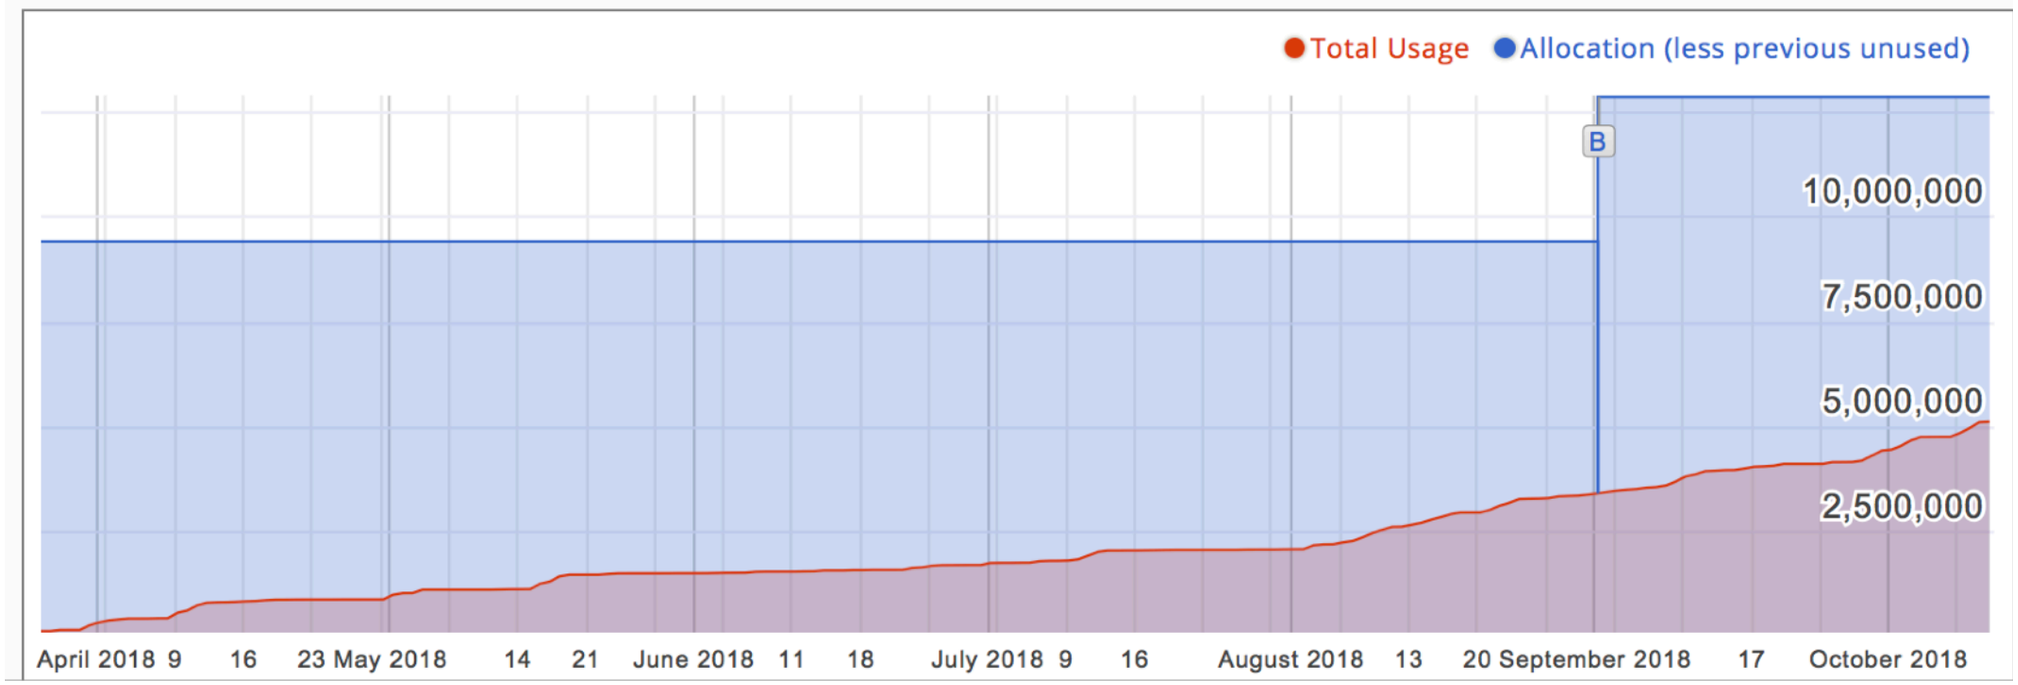

	disk	tape	purging?
home/user	1 TB	5 TB	no
project	500 TB	6,000 TB	no
scratch	1,000 TB	—	after 30 days

- Extensive storage availability to cope with the petascale resources at the disposal.

COMPASS MC @ Blue-Waters

- **Determination of Multi-Dimensional acceptance**
  - MC data encodes both reconstructed and generated info on each event;
  - Experimental acceptance is determined using MC, to correct data for experimental effects (geometrical acceptance, reconstruction effects, detector efficiency …);
  - Multi-dimensional analysis in different kinematical variables allows to strongly reduce the model dependence of the corrections;

BW-resources give a possibility to carry out extensive MC-simulations and a number of other CPU-consuming studies.

COMPASS data transfer @ Blue-Waters

- **Data Transfer and processing:**
- Raw data are moved from CERN to BW disk using FTS3 and tarred and stored on BW tape;
- Typically, we process ~ 160.000 jobs in parallel on BW (~25% of the XE nodes)
- Data are used for various test and technical-productions and studies (efficiency, user analyses)

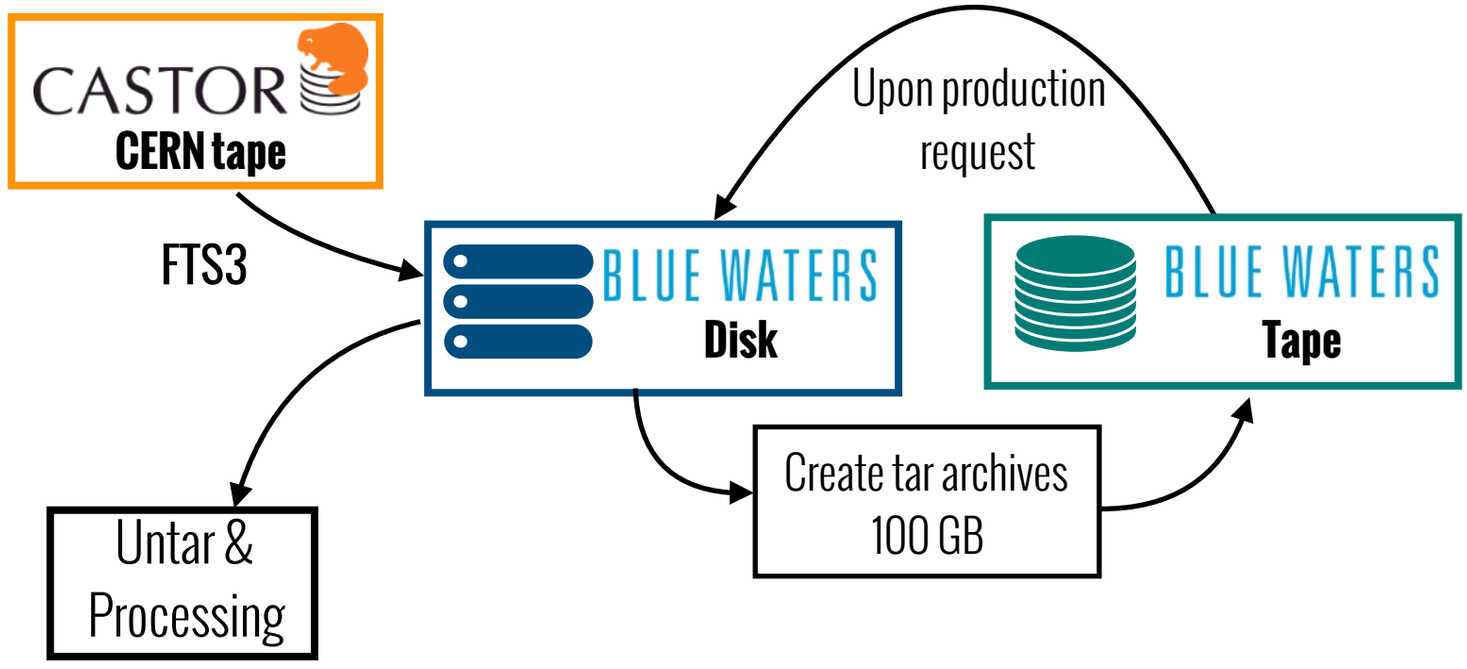

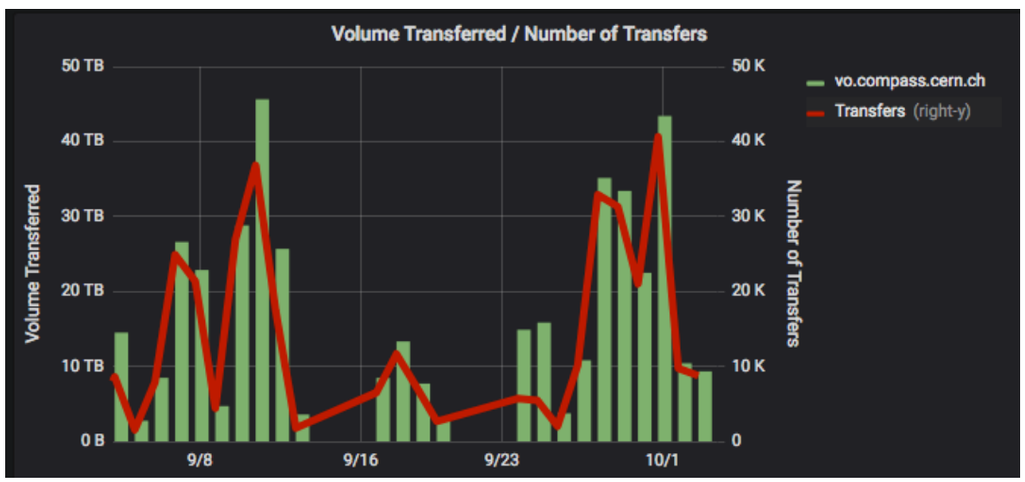

Example: 2018-data transfer from CERN to BW in the last month;

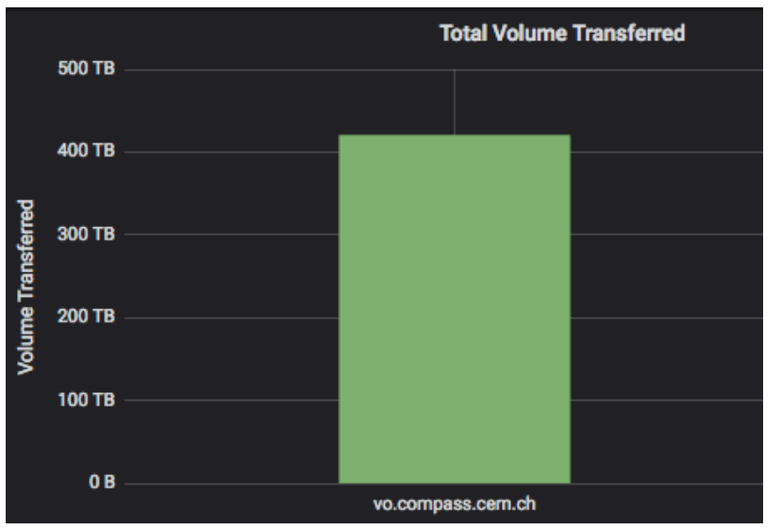

- **Massive data transfer using FTS3**
- Allows for fast copy of TB of data from CERN to BW;

COMPASS efficiency extraction@ Blue-Waters

- **Extraction of 2D efficiency maps**
  - Efficiency are sampled over each period, randomly selecting a portion of stable data;
  - COMPASS has ~ 220 detector planes;
  - E.g. 2015 DY run:

9 data taking periods × ~220 detector planes = ~ 10 Mjobs !!!!

- Extremely time consuming task;
- Allows for COMPASS an unprecedented estimation of the detector performance over the year;
- Significant improvement in the simulation of the detector response;

COMPASS efficiency extraction@ Blue-Waters

- **Extraction of 2D efficiency maps**
  - Efficiency are sampled over each period, randomly selecting a portion of stable data;
  - COMPASS has ~ 220 detector planes;
  - E.g. 2015 DY run:

9 data taking periods **x** ~220 detector planes **=** ~ 10 Mjobs !!!! **x** 4 years of COMPASS-II data taking **!!!**

- Extremely time consuming task;
- Allows for COMPASS an unprecedented estimation of the detector performance over the year;
- Significant improvement in the simulation of the detector response;
- Can be accomplished in reasonable time only making use of HPC resources like BW;

Summary

- BW quickly became important for many COMPASS tasks (MC, 2D maps extraction, user analysis).
- News from last week: BW is going to be extended through December 2019 (new proposal for April - December 2019).
- Frontera system (recently funded by NSF, more info here) at the University of Texas is expected to go into service in the second half of 2019.
- COMPASS is exploring the possibility to submit a proposal, once the call has been opened;

Thank you!

The Blue Waters sustained-petascale computing project is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. BW allocation for COMPASS is also part of the "Mapping Proton Quark Structure using Petabytes of COMPASS Data" PRAC allocation supported by the National Science Foundation (award number OCI 1713684).