Computing at the LHC

and its transformation towards the HL-LHC

Sebastian Wozniewski

DPG Frühjahrstagung 2025 – Göttingen

Georg-August-Universität Göttingen
In publica commoda
seit 1737

Worldwide LHC Computing Grid (WLCG)

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

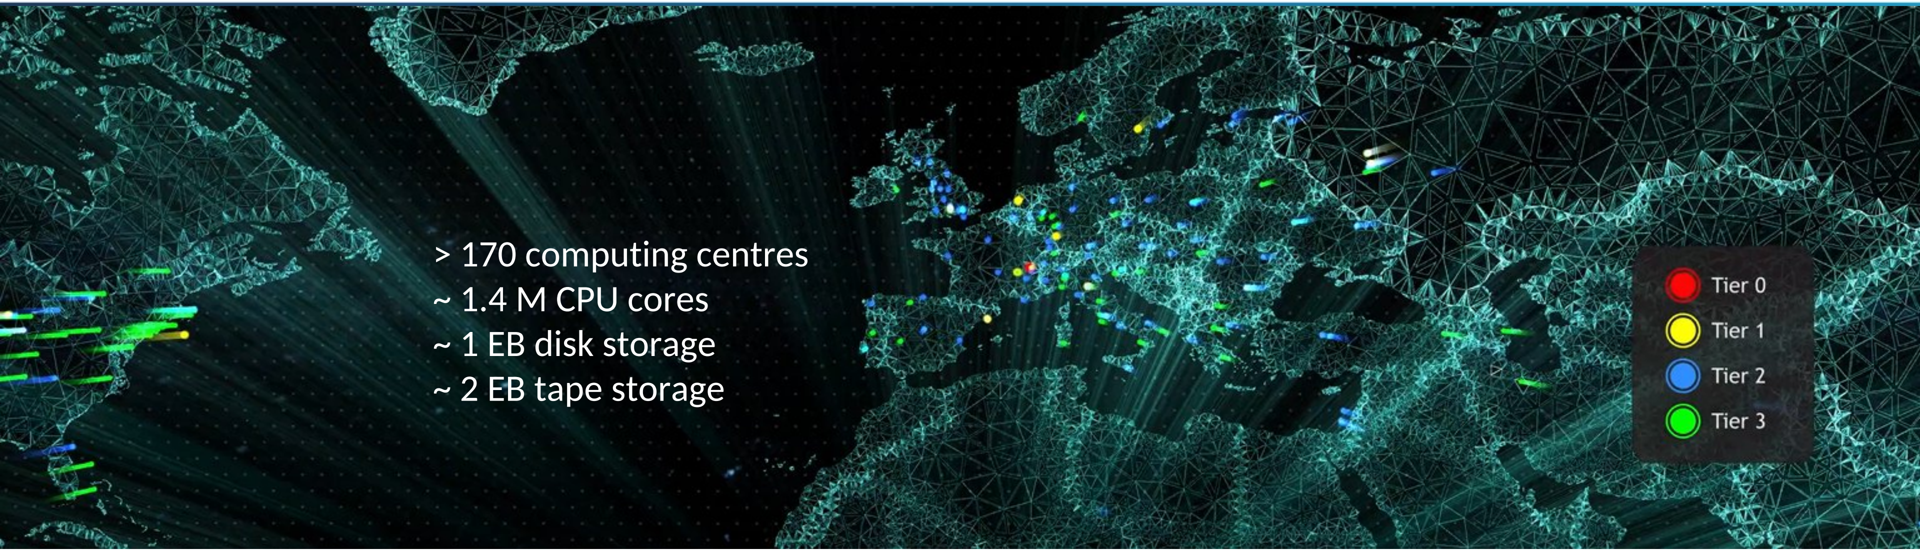

> 170 computing centres
~ 1.4 M CPU cores
~ 1 EB disk storage
~ 2 EB tape storage

Worldwide LHC Computing Grid (WLCG)

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

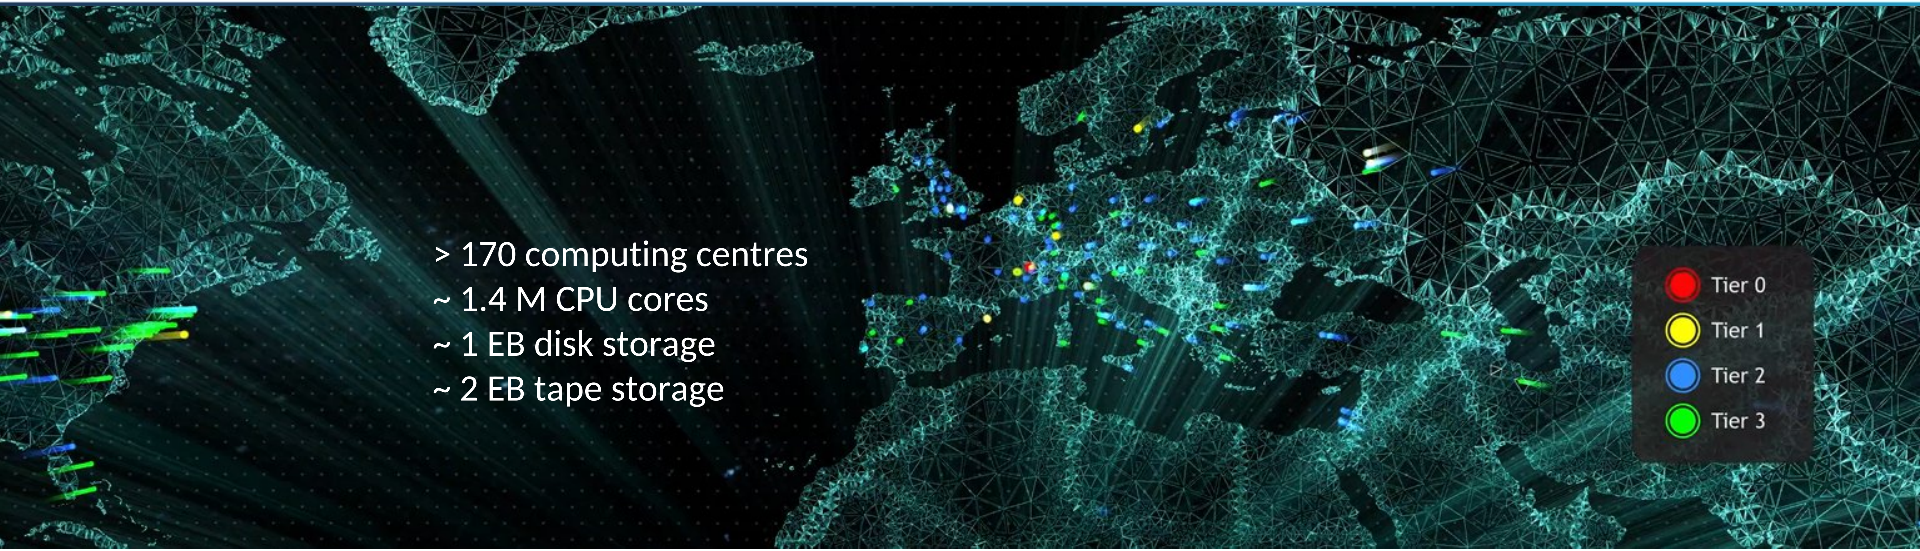

> 170 computing centres
~ 1.4 M CPU cores
~ 1 EB disk storage
~ 2 EB tape storage

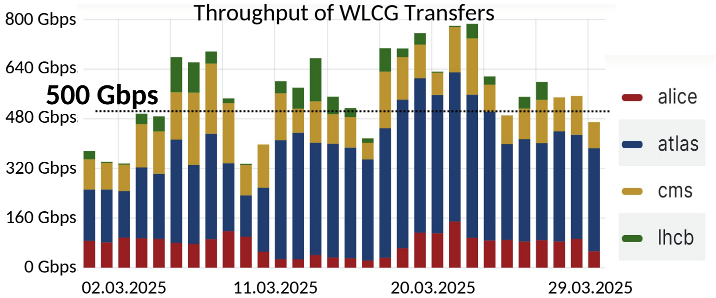

Worldwide LHC Computing Grid (WLCG)

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

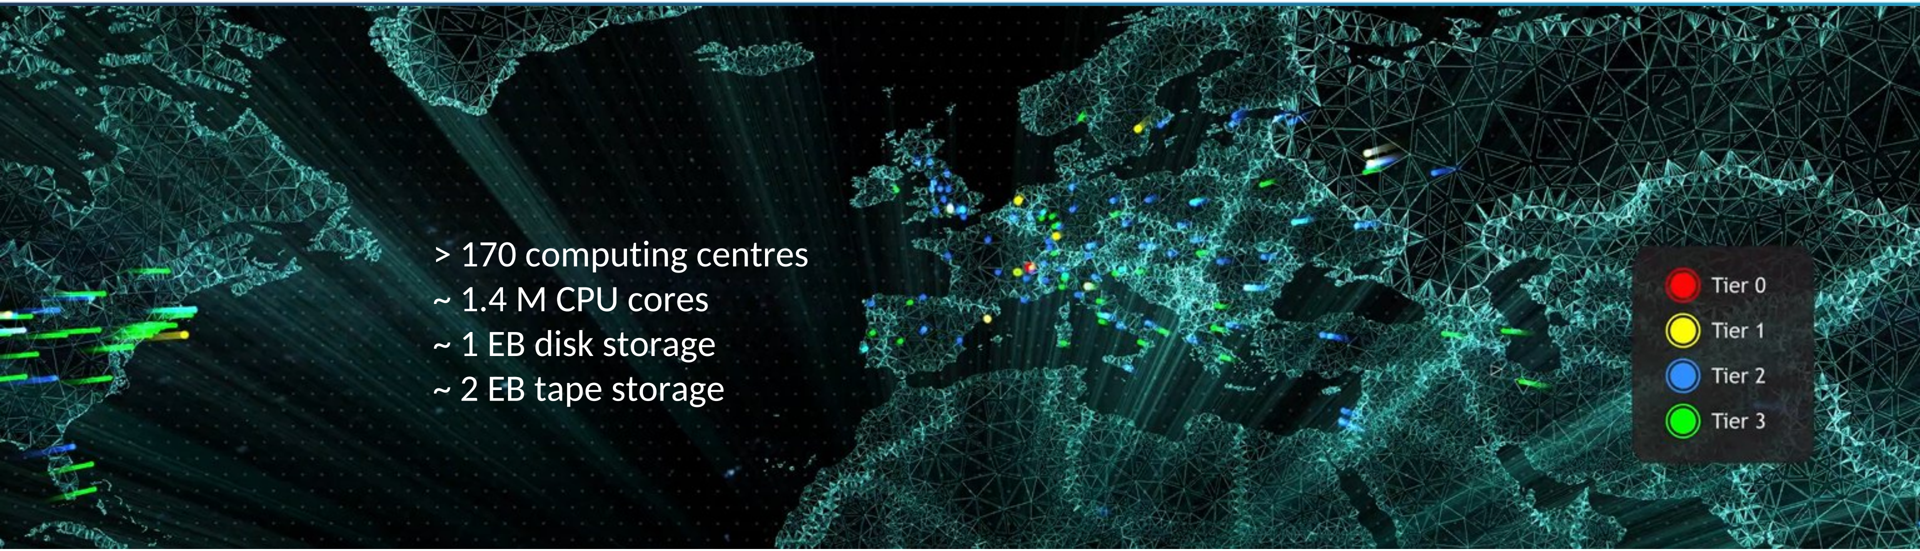

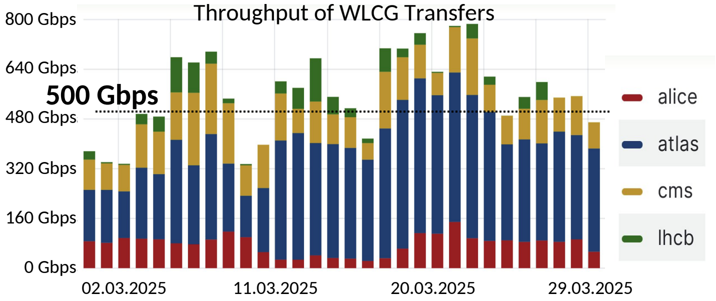

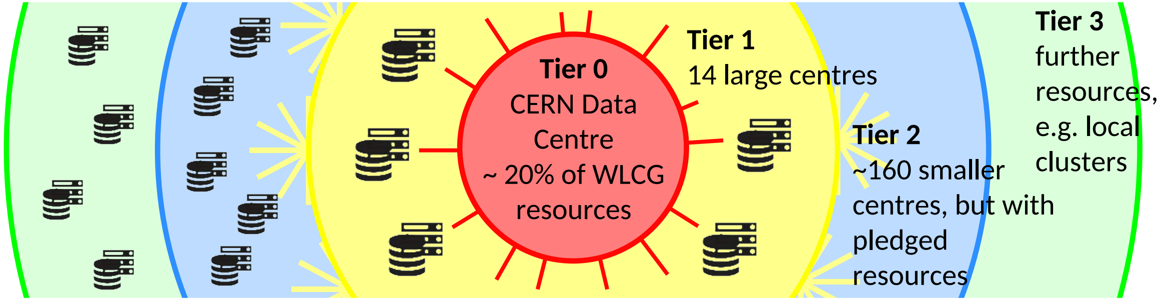

WLCG in Germany

German Tier 1 centre „GridKa“ at KIT

Tier 2 centres at DESY, GSI, MPP and universities

WLCG resources by country within last year:
(excluding Tier-0)

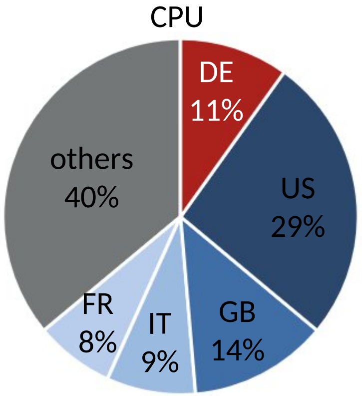

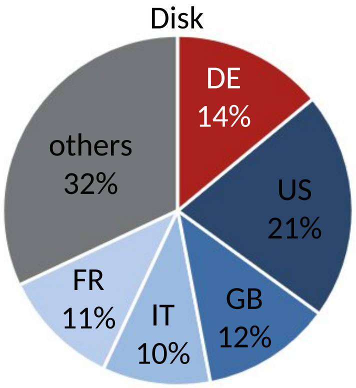

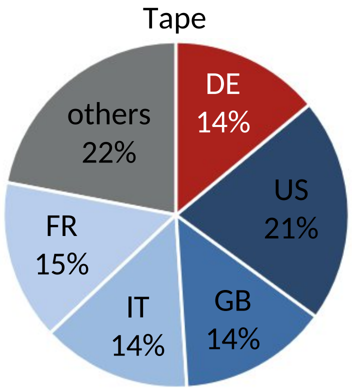

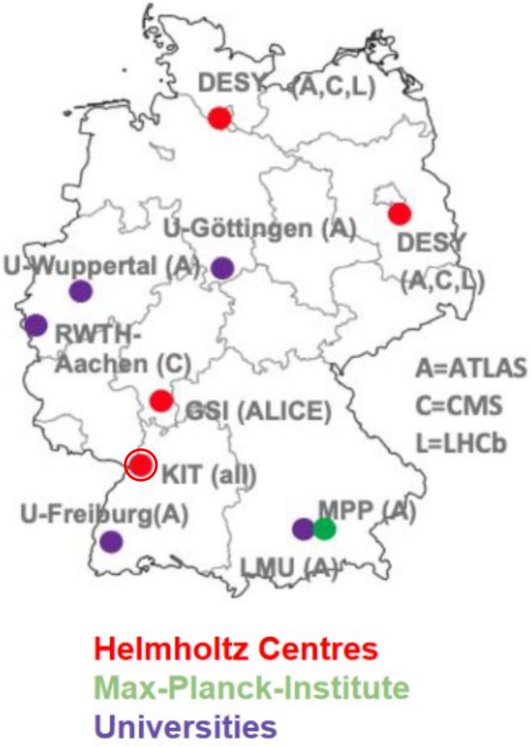

LHC / HL-LHC Plan

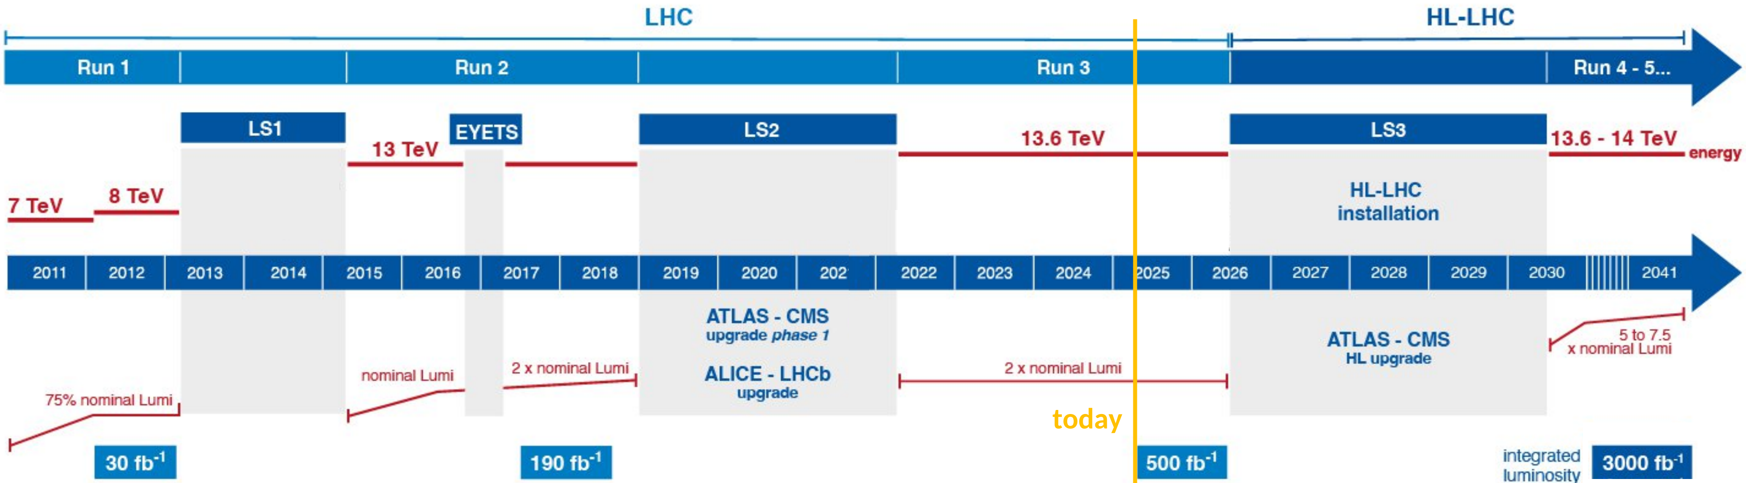

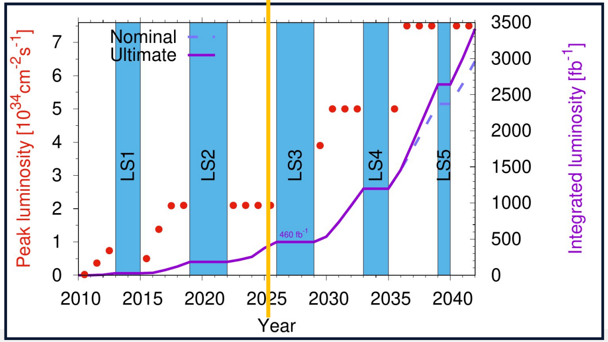

Challenging trigger rate and event complexity!

90% of integrated luminosity yet to come!

Resource prospects for the HL-LHC

CPU

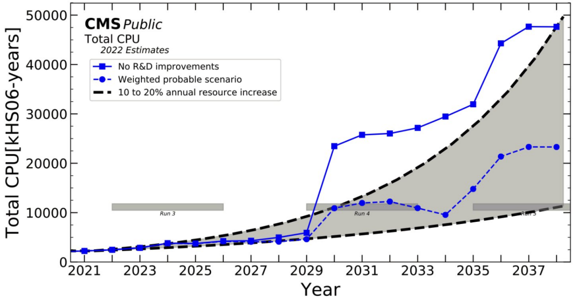

Disk

Tape

Resource prospects for the HL-LHC

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

CPU

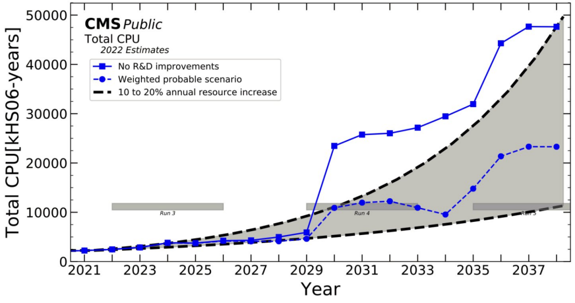

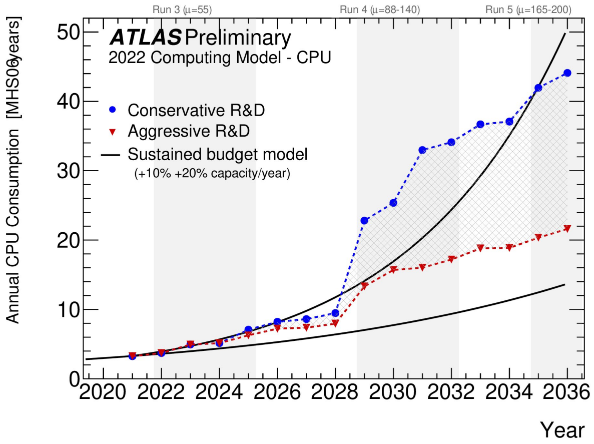

Disk

Tape

Resource prospects for the HL-LHC

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

CPU

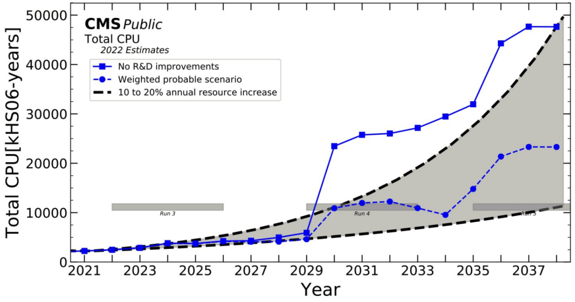

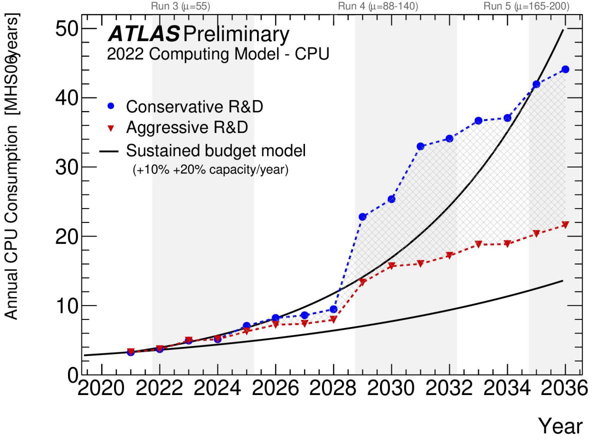

Disk

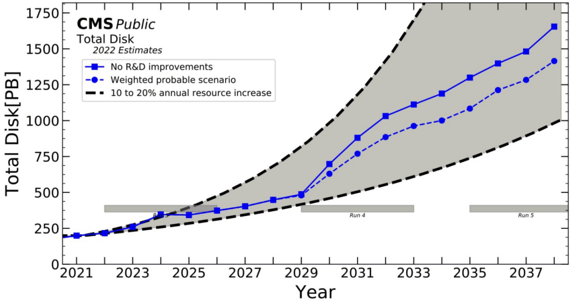

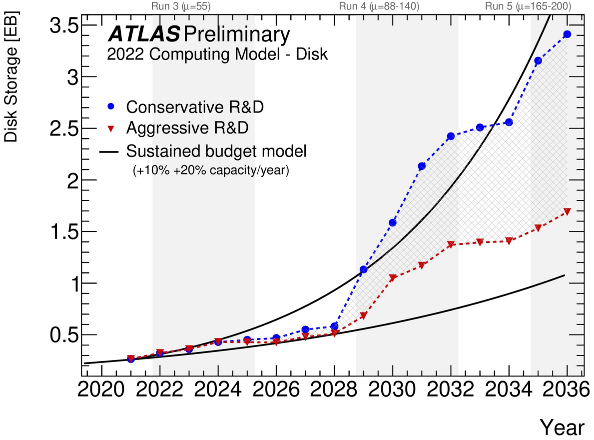

Tape

Resource prospects for the HL-LHC

CPU | **Disk** | **Tape**

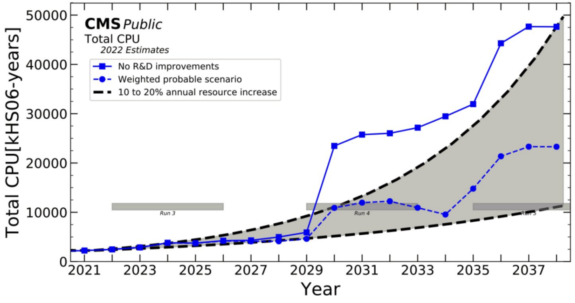

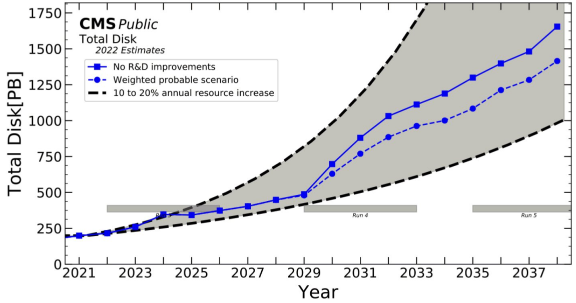

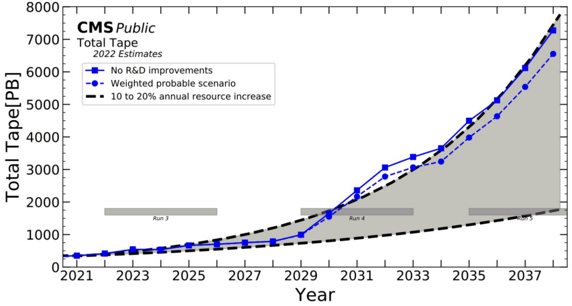

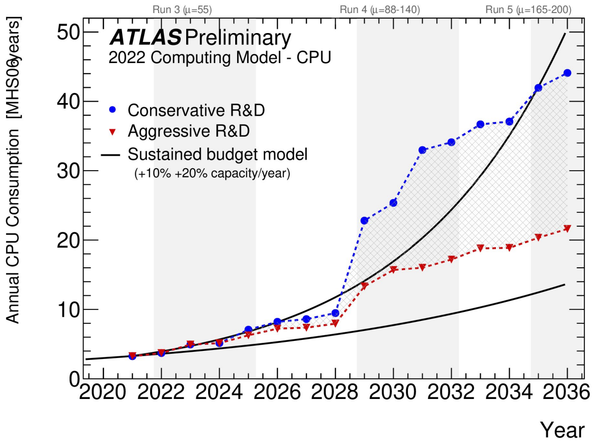

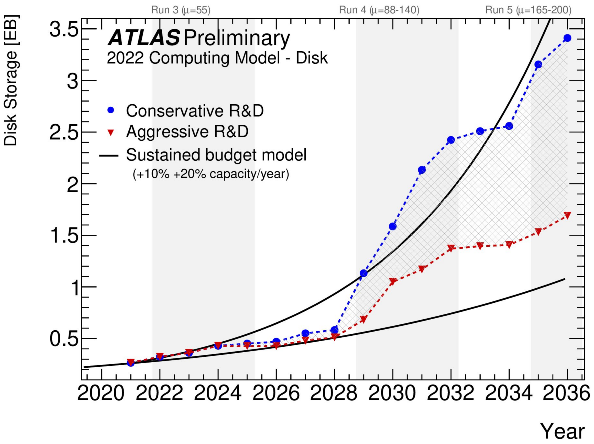

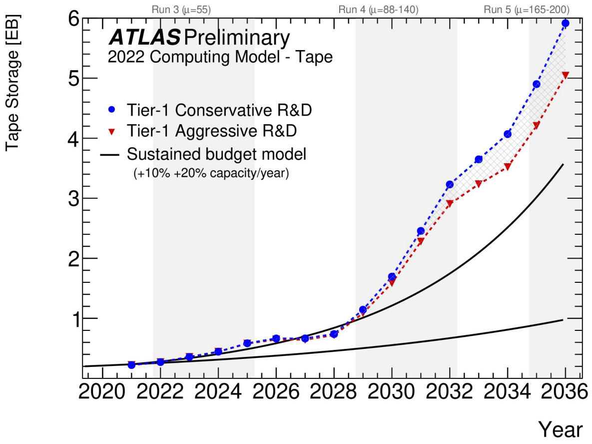

Resource prospects for the HL-LHC

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

CPU

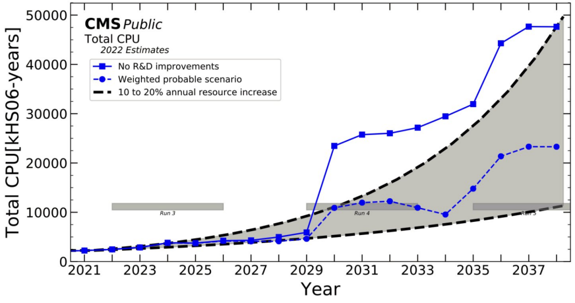

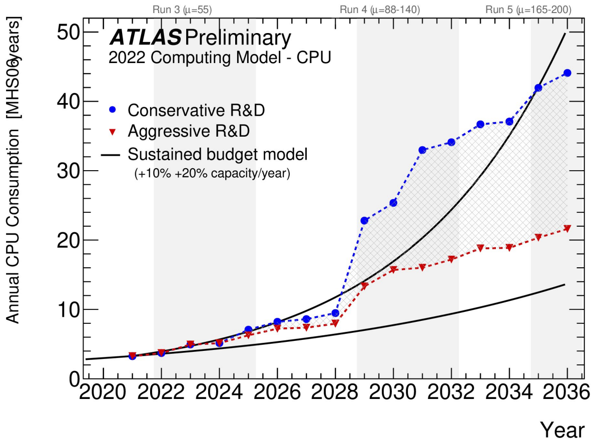

Disk

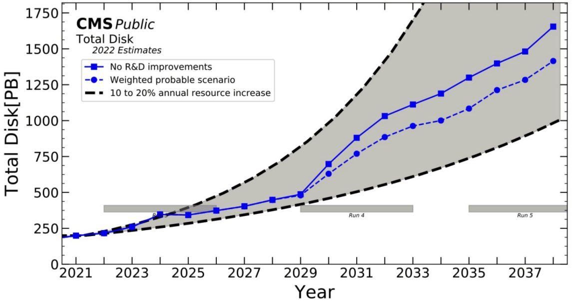

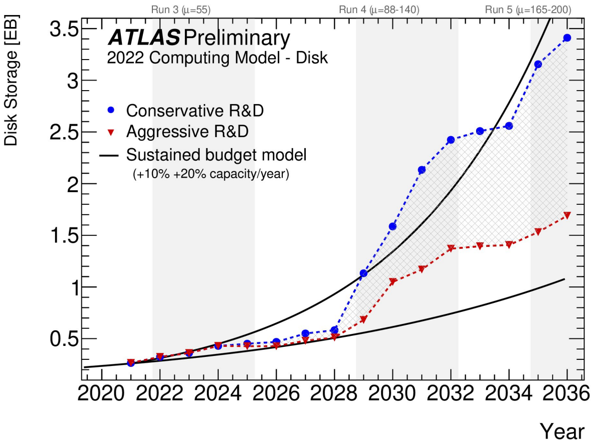

Tape

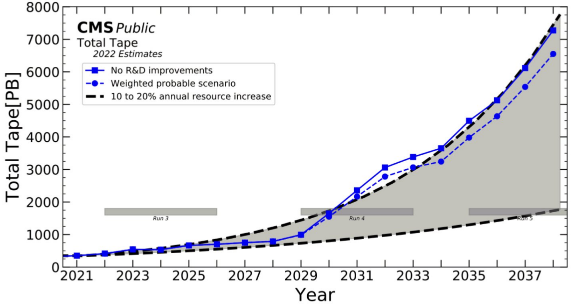

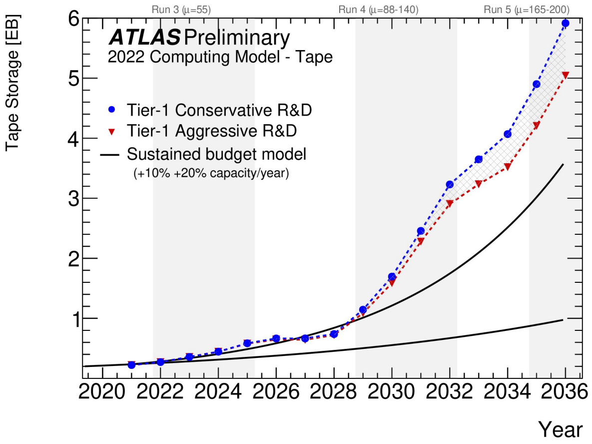

R&D and efficient budget usage necessary to meet budget limits!

Efficiency and Sustainability

Money not the only limiting factor: CO_2 emissions, raw material, environmental impact…
Efficiency more important than ever!

Efficiency and Sustainability

Money not the only limiting factor: CO_2 emissions, raw material, environmental impact…
Efficiency more important than ever!

Usage:
Development cycles, algorithms,
data processing snapshots, data formats,
user awareness, machine learning

scientific value / flop or stored byte

Resource provision:
Power consumption of the servers,
Power Usage Effectiveness (PUE*) of the entire data centre, hardware replacments

flops or stored byte / power

Power supply: Type, cost, availability
(<-> flexibility of resources)

power / money or emission

*PUE = power consumption data center / power consumption hardware

Efficiency and Sustainability

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Money not the only limiting factor: CO_2 emissions, raw material, environmental impact…
Efficiency more important than ever!

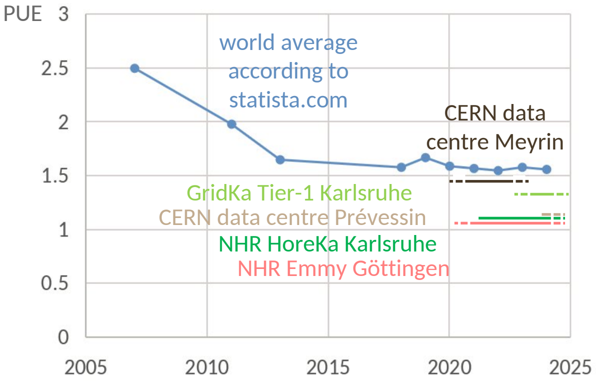

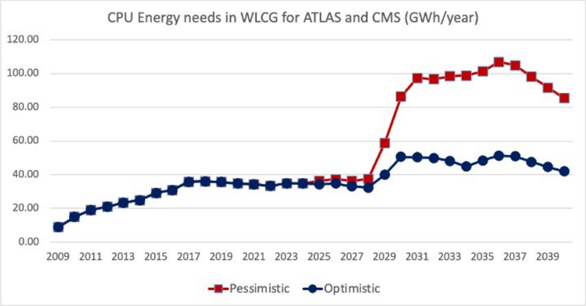

CHEP23, Simone Campana, WLCG

Usage:
Development cycles, algorithms,
data processing snapshots, data formats,
user awareness, machine learning

scientific value / flop or stored byte

Resource provision:
Power consumption of the servers,
Power Usage Effectiveness (PUE*) of the entire data centre, hardware replacments

flops or stored byte / power

Power supply: Type, cost, availability
(<-> flexibility of resources)

power / money or emission

*PUE = power consumption data center / power consumption hardware

Efficiency and Sustainability

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Money not the only limiting factor: CO_2 emissions, raw material, environmental impact…
Efficiency more important than ever!

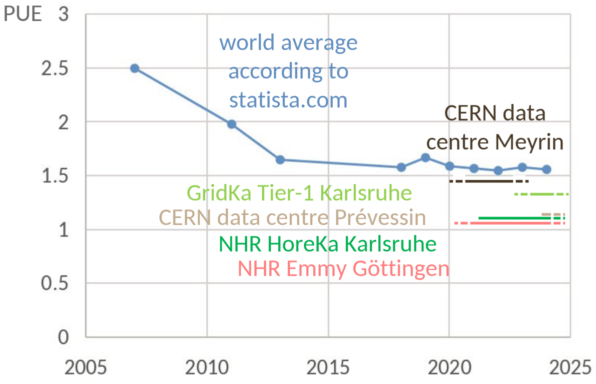

Usage:
Development cycles, algorithms, data processing snapshots, data formats, user awareness, machine learning

scientific value / flop or stored byte

Resource provision:
Power consumption of the servers, Power Usage Effectiveness (PUE*) of the entire data centre, hardware replacments

flops or stored byte / power

Power supply: Type, cost, availability (<-> flexibility of resources)

power / money or emission

WLCG & LHC Collaborations pursuing extensive R&D roadmaps towards HL-HLC + many local efforts of sites or national communities

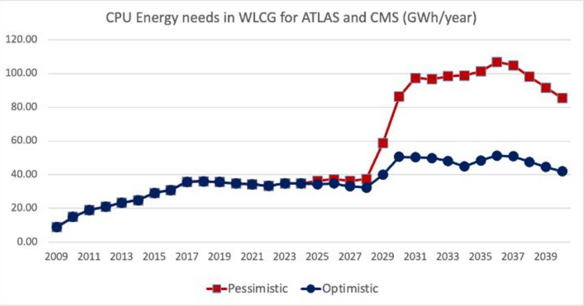

CHEP23, Simone Campana, WLCG

*PUE = power consumption data center / power consumption hardware

Data Challenges

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Expected data rates T0 -> T1 for HL-LHC:

ATLAS & CMS:	1000 Gbps	
ALICE & LHCb:	200 Gbps	
	1200 Gbps	x2 (same for T1 -> T2)
	2400 Gbps	x2 for bursts
=> DC nominal rate:	**4.8 Tbps**	

Data Challenges

Expected data rates T0 -> T1 for HL-LHC:

ATLAS & CMS:	1000 Gbps	
ALICE & LHCb:	200 Gbps	
	1200 Gbps	x2 (same for T1 -> T2)
	2400 Gbps	x2 for bursts
=> DC nominal rate:	**4.8 Tbps**	

Series of data challenges carried out on WLCG:

Year	% of HL-LHC
2021	10
2024	25
~2027	50
~2029	100

Data Challenges

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Expected data rates T0 -> T1 for HL-LHC:

ATLAS & CMS:	1000 Gbps	
ALICE & LHCb:	200 Gbps	
	1200 Gbps	x2 (same for T1 -> T2)
	2400 Gbps	x2 for bursts
=> DC nominal rate:	**4.8 Tbps**	

Series of data challenges carried out on WLCG:

Year	% of HL-LHC
2021	10
2024	25
~2027	50
~2029	100

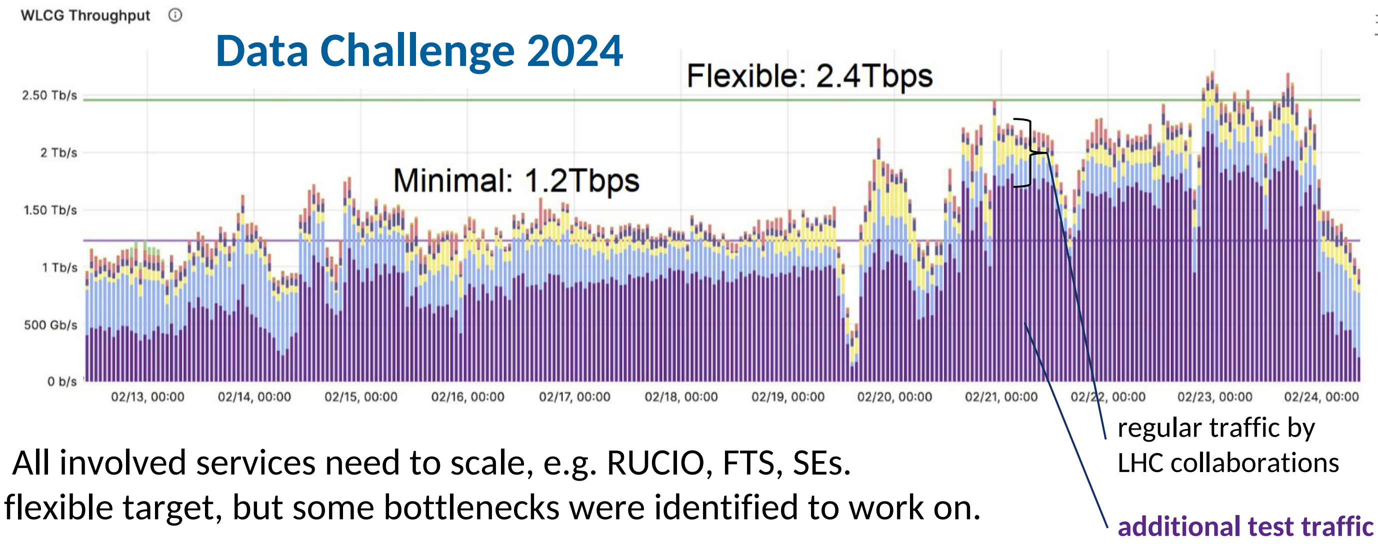

https://zenodo.org/records/5532452

WLCG/HSF Workshop 2024, Katy Ellis, RAL

It‘s not just about bandwidth! All involved services need to scale, e.g. RUCIO, FTS, SEs.
Data Challenge 2024 reached flexible target, but some bottlenecks were identified to work on.

Data Formats

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data

Event content:

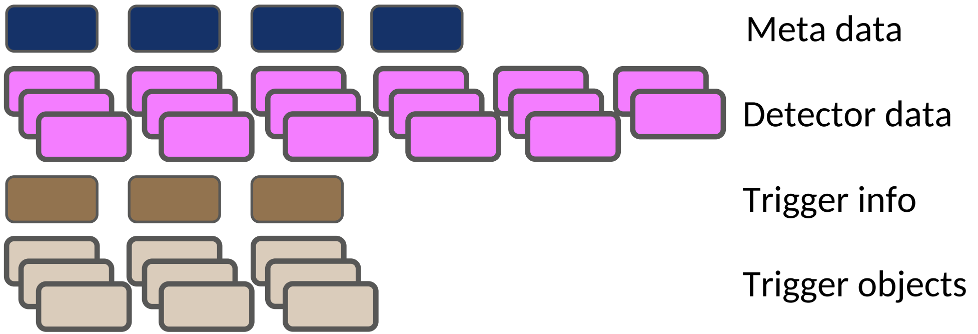

Data Formats

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects

Event content:

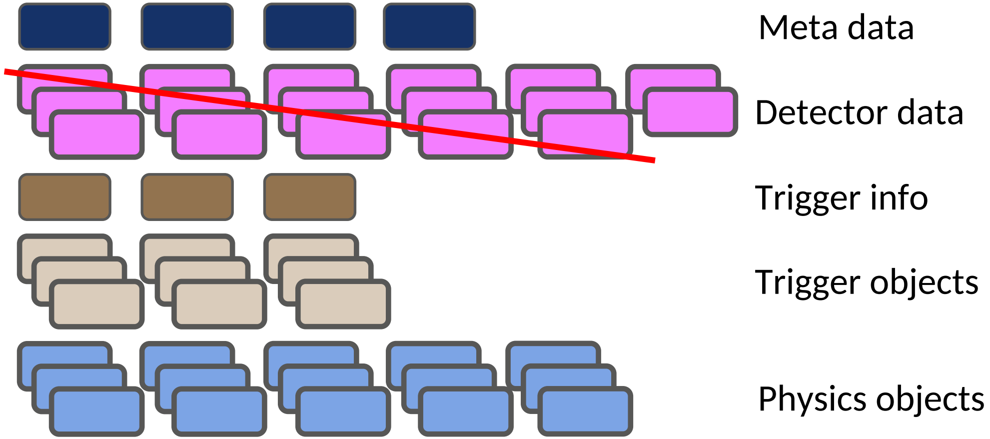

Data Formats

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects
- slimmed physics objects

Event content:

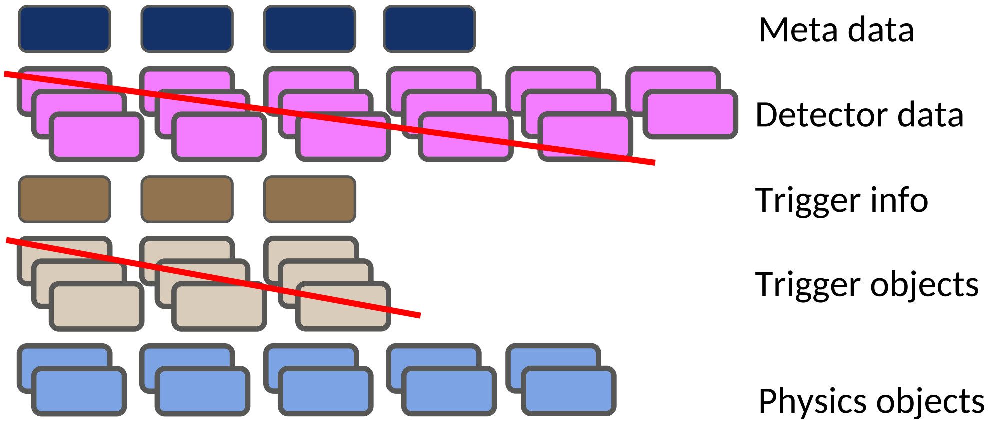

Data Formats

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects
- slimmed physics objects
- analysis ntuples

Event content:

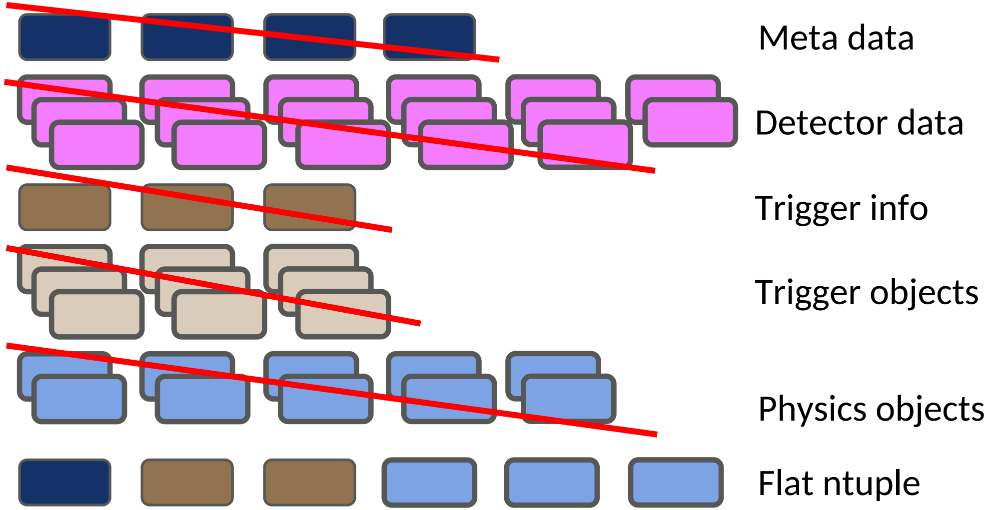

Data Formats

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects
- slimmed physics objects
- analysis ntuples

=> more standardised and light-weight analysis data formats in favour of storage and CPU consumption

Event content:

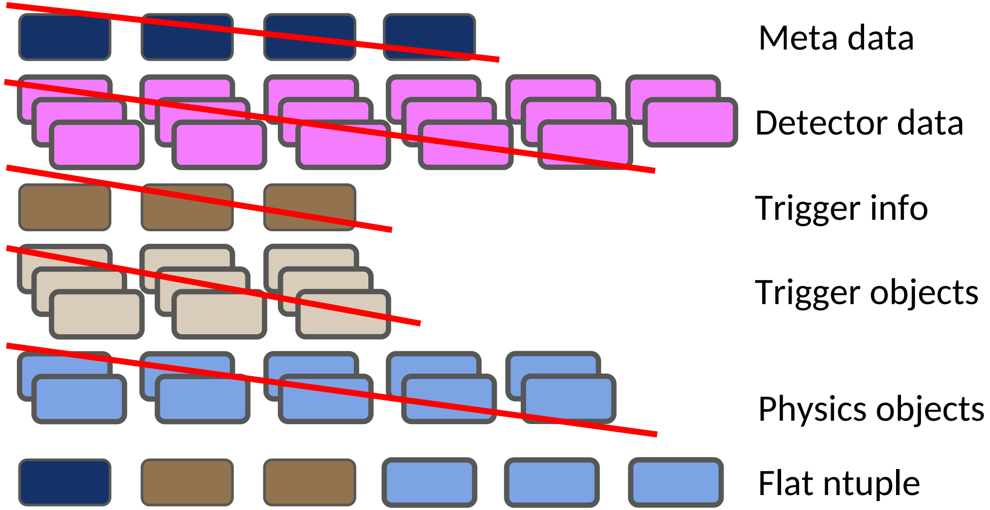

Data Formats | RNTuple

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects
- slimmed physics objects
- analysis ntuples

=> more standardised and light-weight analysis data formats in favour of storage and CPU consumption

Event content:

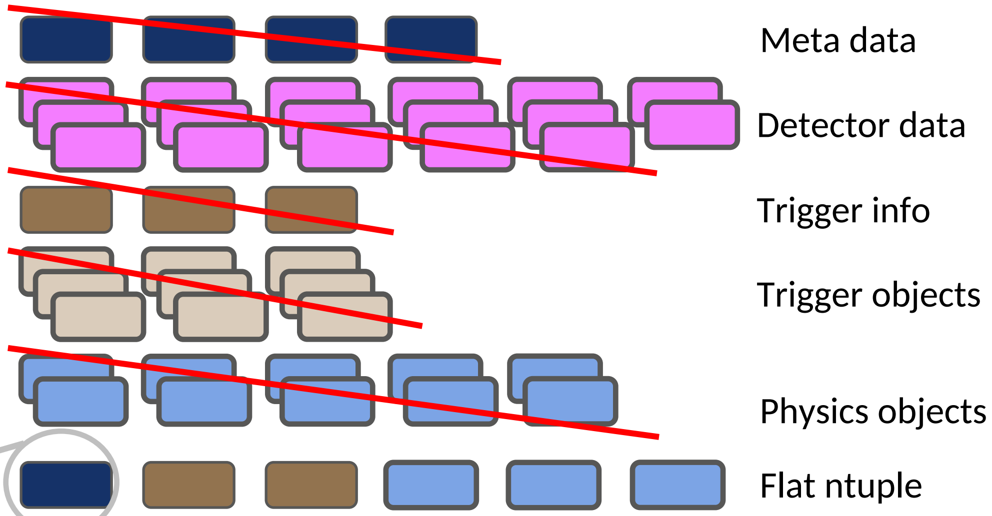

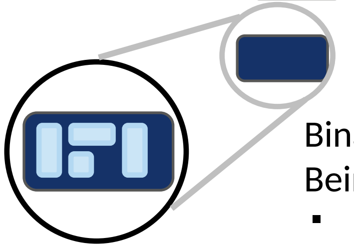

Binary format **TTree** for >25 years!
Being replaced by new optimized **RNTuple**:

- adapted to modern hardware
- asynchronous & parallel I/O
- new basic types, e.g. f16
- more efficient storage of booleans and collections

First numbers reported by ATLAS & CMS for the disk size reduction TTree->RNtuple:
6% - 40%
depending on the object data format

CHEP24

Data Formats | RNTuple

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Consolidation of **object data formats** ongoing issue over LHC-Runs:

- raw data
- reconstructed physics objects
- slimmed physics objects
- analysis ntuples

=> more standardised and light-weight analysis data formats in favour of storage and CPU consumption

Event content:

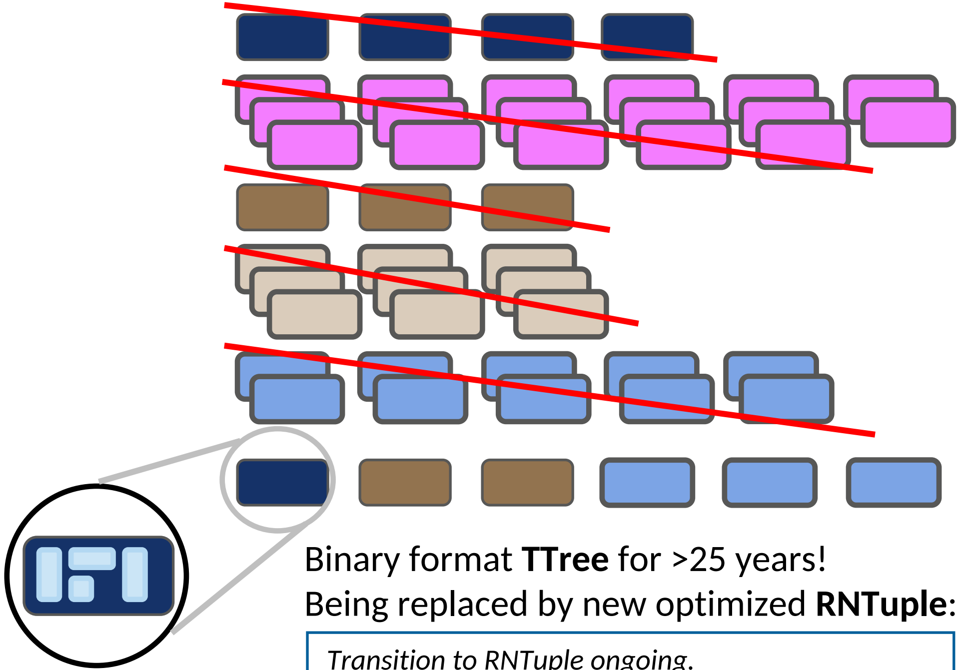

Binary format **TTree** for >25 years!
Being replaced by new optimized **RNTuple**:

Transition to RNTuple ongoing.
RNTuple ready but work on adaptation and validation of code necessary!

First numbers reported by ATLAS & CMS for the disk size reduction TTree->RNtuple:
6% - 40%
depending on the object data format

CHEP24

CPU architectures and accelerators

Other hardware solutions than **x86 CPUs** can improve speed and energy efficiency for specific tasks:

- **GPUs** used in many places already, e.g. HLT/online data processing, ML, …

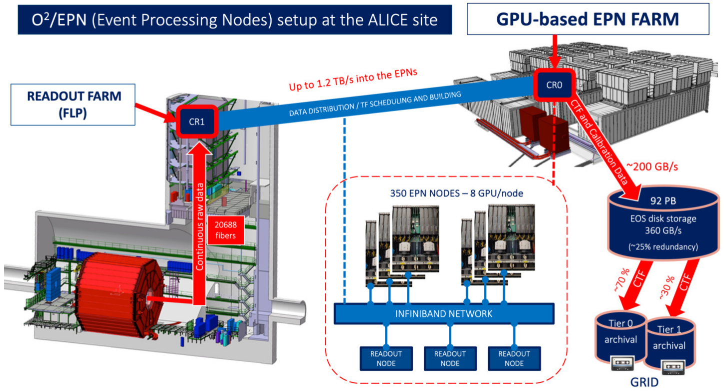

CPU architectures and accelerators

Other hardware solutions than **x86 CPUs** can improve speed and energy efficiency for specific tasks:

- **GPUs** used in many places already, e.g. HLT/online data processing, ML, …
- **ARM CPUs** energy efficient, but depends on clock frequency and strong competition with new x86 models.
  LHC collaborations are working on making workflows ARM-compatible => flexibility for future decisions.
  ATLAS even accepting ARM pledges since 2025.

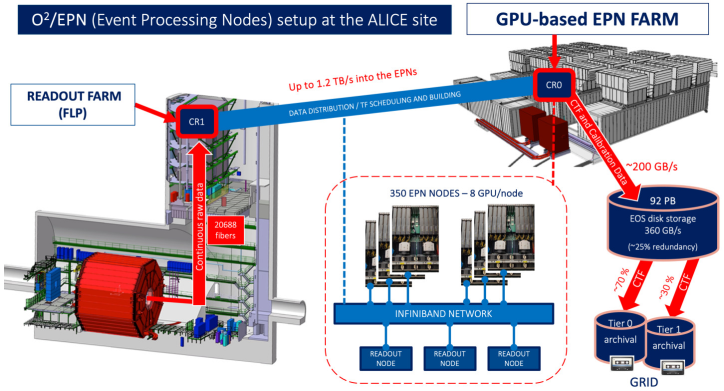

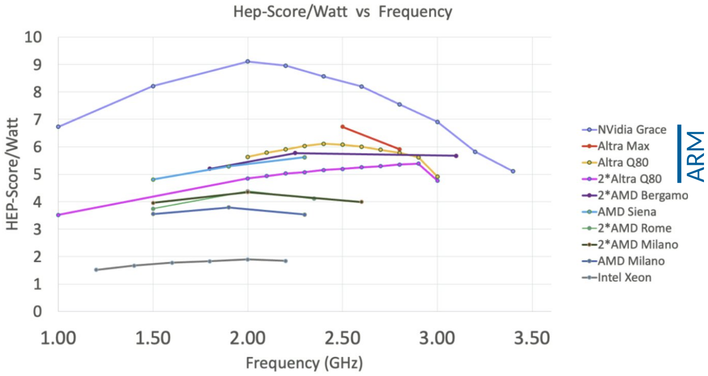

WLCG Grid Deployment Board, D. Britton

CPU architectures and accelerators

Other hardware solutions than **x86 CPUs** can improve speed and energy efficiency for specific tasks:

- **GPUs** used in many places already, e.g. HLT/online data processing, ML, …
- **ARM CPUs** energy efficient, but depends on clock frequency and strong competition with new x86 models.
  LHC collaborations are working on making workflows ARM-compatible => flexibility for future decisions.
  ATLAS even accepting ARM pledges since 2025.

Market very dynamic and hard to predict!
(new models; other customers and applications, e.g. AI)

=> CERN, LHC Collaborations, Sites and dedicated boards keep watching the market and increase flexibility.

Commercial cloud providers taken into account for hardware testing.

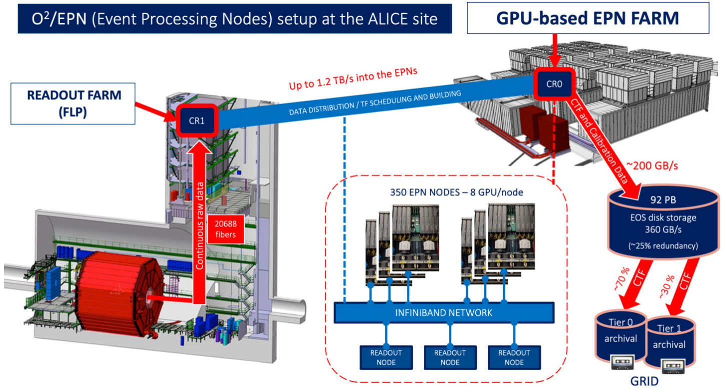

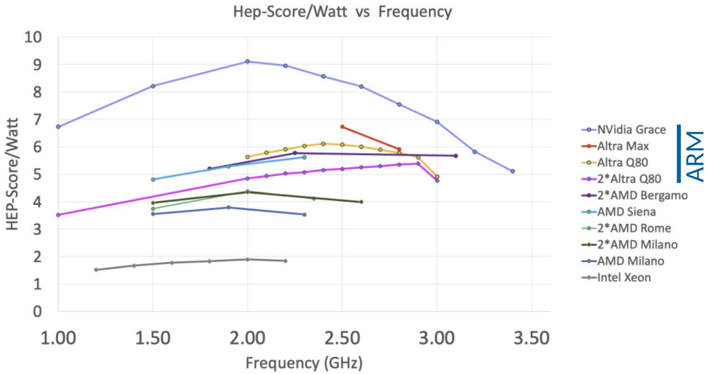

WLCG Grid Deployment Board, D. Britton

CPU architectures and accelerators

Other hardware solutions than **x86 CPUs** can improve speed and energy efficiency for specific tasks:

- **GPUs** used in many places already, e.g. HLT/online data processing, ML, …
- **ARM CPUs** energy efficient, but depends on clock frequency and strong competition with new x86 models.
  LHC collaborations are working on making workflows ARM-compatible => flexibility for future decisions.
  ATLAS even accepting ARM pledges since 2025.

Market very dynamic and hard to predict!
(new models; other customers and applications, e.g. AI)

=> CERN, LHC Collaborations, Sites and dedicated boards keep watching the market and increase flexibility.

Commercial cloud providers taken into account for hardware testing.

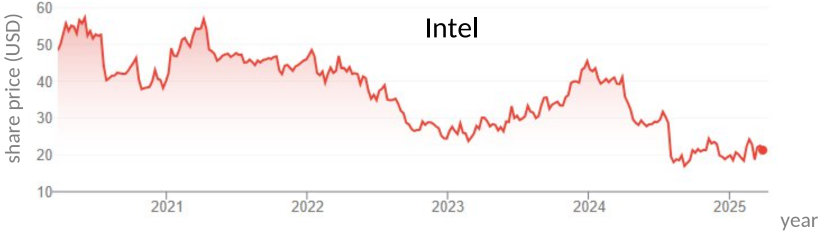

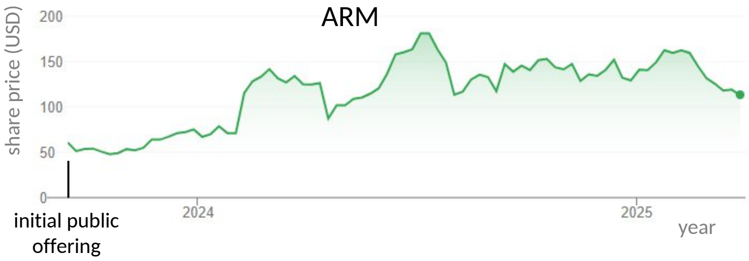

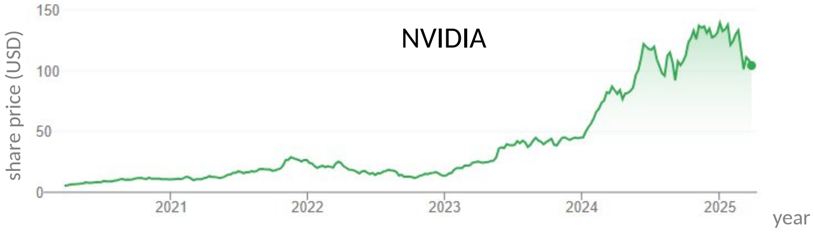

Fast(er) simulation and reconstruction

some examples

Potential improvements in all simulation (generator & detector) and reconstruction steps, in particular AI driven!

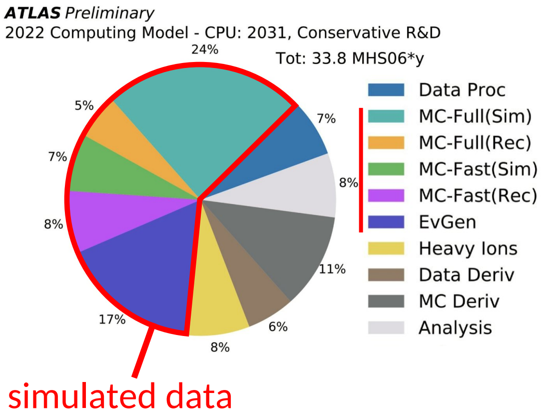

Fast(er) simulation and reconstruction

some examples

Potential improvements in all simulation (generator & detector) and reconstruction steps, in particular AI driven!

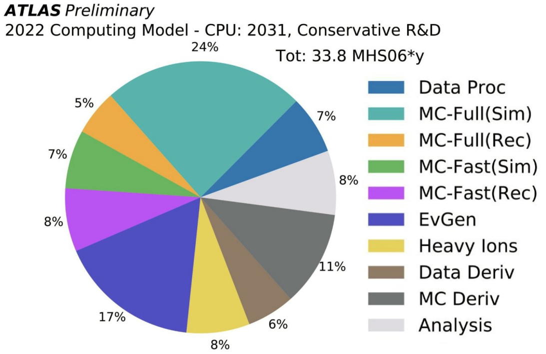

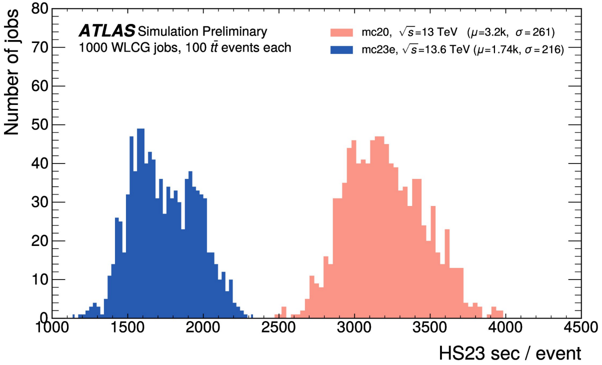

Fast(er) simulation and reconstruction

some examples

Potential improvements in all simulation (generator & detector) and reconstruction steps, in particular AI driven!

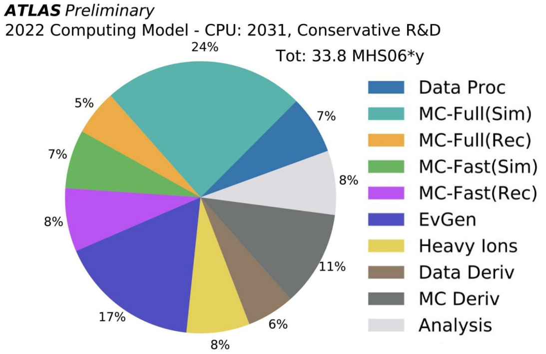

NN driven improvement in SHERPA importance sampling:

unweighting eff. ϵ_{uw} process/sampling		LO QCD $n=0$	$n=1$	$n=2$	$n=3$	$n=4$	NLO QCD (RS) $n=0$	$n=1$
$W^+ + n$ jets	SHERPA	$2.8\cdot10^{-1}$	$3.8\cdot10^{-2}$	$7.5\cdot10^{-3}$	$1.5\cdot10^{-3}$	$8.3\cdot10^{-4}$	$9.5\cdot10^{-2}$	$4.5\cdot10^{-3}$
	NN	$6.1\cdot10^{-1}$	$1.2\cdot10^{-1}$	$1.0\cdot10^{-3}$	$1.8\cdot10^{-3}$	$8.9\cdot10^{-4}$	$1.6\cdot10^{-1}$	$4.1\cdot10^{-3}$
	Gain	2.2	3.3	1.4	1.2	1.1	1.6	0.91
$Z/\gamma^* + n$ jets	SHERPA	$3.1\cdot10^{-1}$	$3.6\cdot10^{-2}$	$1.5\cdot10^{-2}$	$4.7\cdot10^{-3}$		$1.2\cdot10^{-1}$	$5.3\cdot10^{-3}$
	NN	$3.8\cdot10^{-1}$	$1.0\cdot10^{-1}$	$1.4\cdot10^{-2}$	$2.4\cdot10^{-3}$		$1.8\cdot10^{-3}$	$5.7\cdot10^{-3}$
	Gain	1.2	2.9	0.91	0.51		1.5	1.1

[SciPost Phys. 14 (2023) 4, 079]

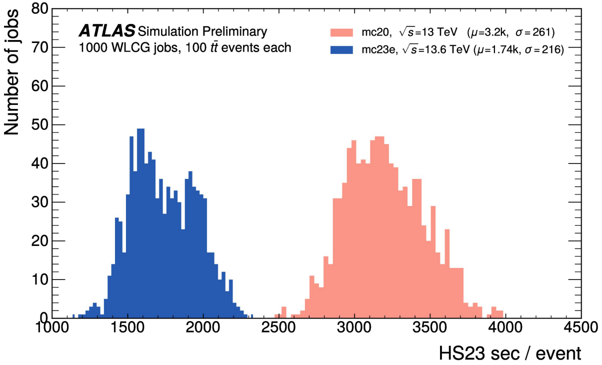

Fast(er) simulation and reconstruction

some examples

Potential improvements in all simulation (generator & detector) and reconstruction steps, in particular AI driven!

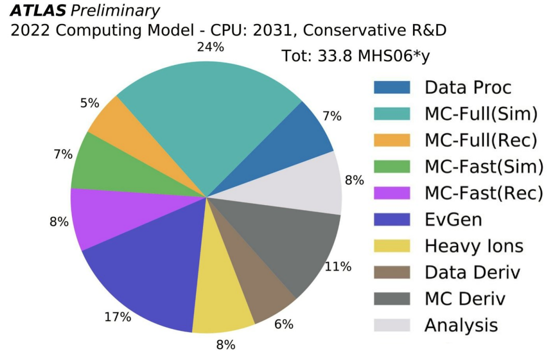

NN driven improvement in SHERPA importance sampling:

unweighting eff. ϵ_{uw}		LO QCD					NLO QCD (RS)	
process/sampling		$n=0$	$n=1$	$n=2$	$n=3$	$n=4$	$n=0$	$n=1$
$W^+ + n$ jets	SHERPA	$2.8\cdot10^{-1}$	$3.8\cdot10^{-2}$	$7.5\cdot10^{-3}$	$1.5\cdot10^{-3}$	$8.3\cdot10^{-4}$	$9.5\cdot10^{-2}$	$4.5\cdot10^{-3}$
	NN	$6.1\cdot10^{-1}$	$1.2\cdot10^{-1}$	$1.0\cdot10^{-3}$	$1.8\cdot10^{-3}$	$8.9\cdot10^{-4}$	$1.6\cdot10^{-1}$	$4.1\cdot10^{-3}$
	Gain	2.2	3.3	1.4	1.2	1.1	1.6	0.91
$Z/\gamma^* + n$ jets	SHERPA	$3.1\cdot10^{-1}$	$3.6\cdot10^{-2}$	$1.5\cdot10^{-2}$	$4.7\cdot10^{-3}$		$1.2\cdot10^{-1}$	$5.3\cdot10^{-3}$
	NN	$3.8\cdot10^{-1}$	$1.0\cdot10^{-1}$	$1.4\cdot10^{-2}$	$2.4\cdot10^{-3}$		$1.8\cdot10^{-3}$	$5.7\cdot10^{-3}$
	Gain	1.2	2.9	0.91	0.51		1.5	1.1

[SciPost Phys. 14 (2023) 4, 079]

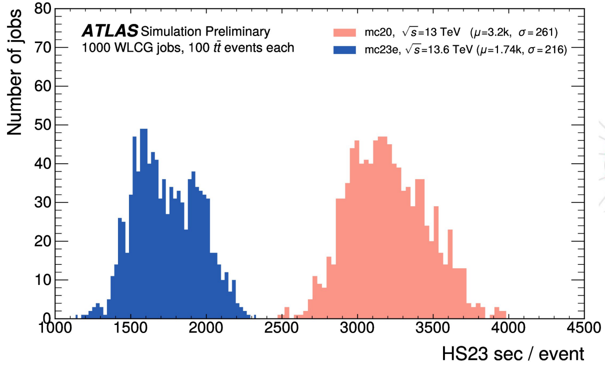

GNN driven track reconstruction „GNN4ITk“:

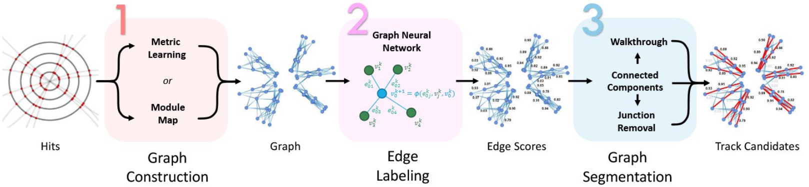

Track reco default ~5s, fast (HLT) ~1s, GNN4ITk ~0.5s
Tuning and thourough evaluation ongoing.

ATL-PHYS-PUB-2024-018

Fast(er) simulation and reconstruction

some examples

Potential improvements in all simulation (generator & detector) and reconstruction steps, in particular AI driven!

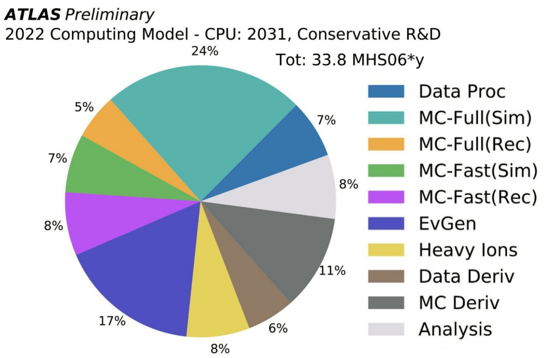

NN driven improvement in SHERPA importance sampling:

unweighting eff. ϵ_{uw} process/sampling		LO QCD $n=0$	LO QCD $n=1$	LO QCD $n=2$	LO QCD $n=3$	LO QCD $n=4$	NLO QCD (RS) $n=0$	NLO QCD (RS) $n=1$
$W^+ + n$ jets	SHERPA	$2.8 \cdot 10^{-1}$	$3.8 \cdot 10^{-2}$	$7.5 \cdot 10^{-3}$	$1.5 \cdot 10^{-3}$	$8.3 \cdot 10^{-4}$	$9.5 \cdot 10^{-2}$	$4.5 \cdot 10^{-3}$
	NN	$6.1 \cdot 10^{-1}$	$1.2 \cdot 10^{-1}$	$1.0 \cdot 10^{-3}$	$1.8 \cdot 10^{-3}$	$8.9 \cdot 10^{-4}$	$1.6 \cdot 10^{-1}$	$4.1 \cdot 10^{-3}$
	Gain	2.2	3.3	1.4	1.2	1.1	1.6	0.91
$Z/\gamma^* + n$ jets	SHERPA	$3.1 \cdot 10^{-1}$	$3.6 \cdot 10^{-2}$	$1.5 \cdot 10^{-2}$	$4.7 \cdot 10^{-3}$		$1.2 \cdot 10^{-1}$	$5.3 \cdot 10^{-3}$
	NN	$3.8 \cdot 10^{-1}$	$1.0 \cdot 10^{-1}$	$1.4 \cdot 10^{-2}$	$2.4 \cdot 10^{-3}$		$1.8 \cdot 10^{-3}$	$5.7 \cdot 10^{-3}$
	Gain	1.2	2.9	0.91	0.51		1.5	1.1

[SciPost Phys. 14 (2023) 4, 079]

End-to-end „FlashSim“:

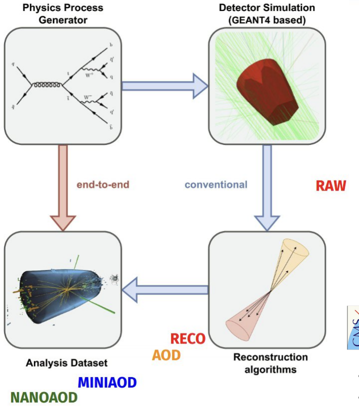

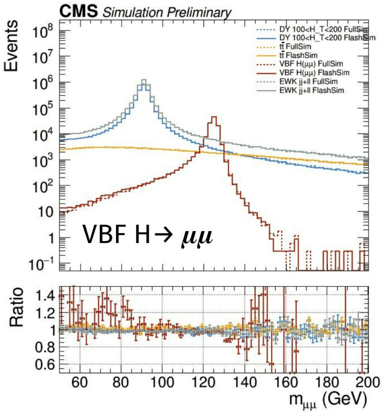

CHEP2024, Andrea Rizzi

Runs at O(100) Hz.
Development ongoing.

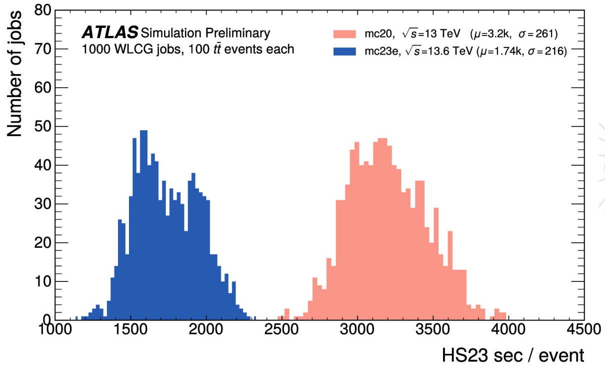

GNN driven track reconstruction „GNN4ITk“:

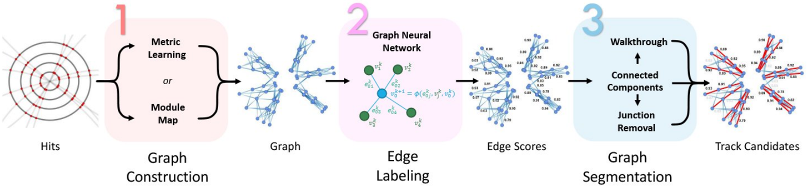

Track reco default ~5s, fast (HLT) ~1s, GNN4ITk ~0.5s
Tuning and thourough evaluation ongoing.

ATL-PHYS-PUB-2024-018

German R&D activities – Overview

- **BMBF**-funded **FC-AC** project: Provision of Tier-2 resources for ATLAS and CMS and R&D for related computing operations

BMBF Action Plan: **ErUM-Data** = dedicated funding line within the ErUM program („**Er**forschung von **U**niversum und **M**aterie“) for the development of the data infrastructure and digital competences. Covers:

- **FIDIUM** project: Additional R&D for the transformation of computing infrastructures in ErUM for coping with the demanding future needs
- **KISS** project: AI for faster simulation of scientific data
- **ErUM-Data-Hub**: networking, knowledge transfer, workshops

NFDI (**N**ationale **F**oruschungs**daten**i**nfrastruktur) = umbrella organisation of 26 **DFG**-funded consortia of various research fields, in particular:

- **PUNCH4NFDI**: **P**articles **U**niverse **N**u**c**lei **H**adrons
  Prime goal: set up a federated and "FAIR" science data platform

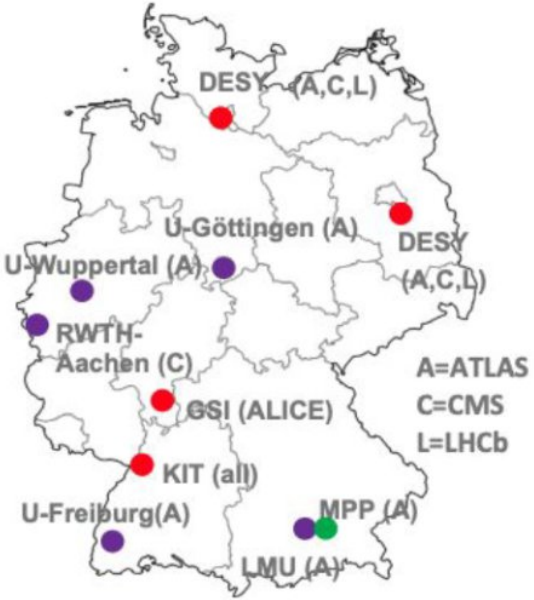

Helmholtz Centres
Max-Planck-Institute
Universities

FIDIUM & Infrastructure transformation

R&D for the transformation of computing infrastructures in ErUM for coping with the demanding future needs.

- Tools for the integration of heterogenous resources, e.g. HPC centres:
  - Resource management: COBalD/TARDIS
  - Accounting: AUDITOR
  - …
- Tools for distributed data storage:
  - Caching
  - Improved authorization mechanisms
  - Monitoring
  - …
- Testing and optimization under realistic conditions

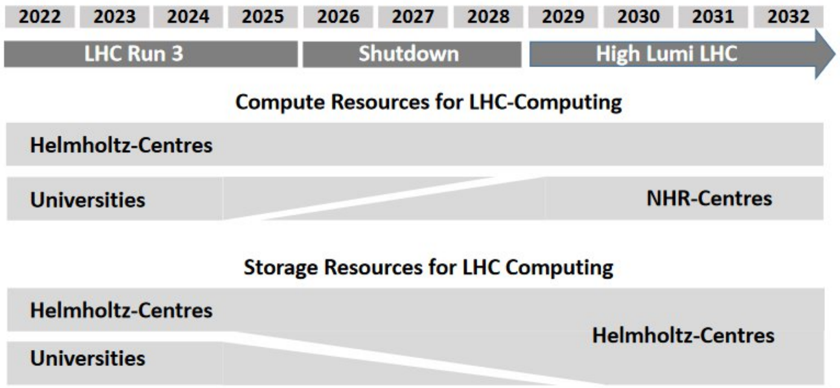

FIDIUM & Infrastructure transformation

R&D for the transformation of computing infrastructures in ErUM for coping with the demanding future needs.

- Tools for the integration of heterogenous resources, e.g. HPC centres:
  - Resource management: COBalD/TARDIS
  - Accounting: AUDITOR
  - …
- Tools for distributed data storage:
  - Caching
  - Monitoring
  - Improved authorization mechanisms
  - …
- Testing and optimization under realistic conditions

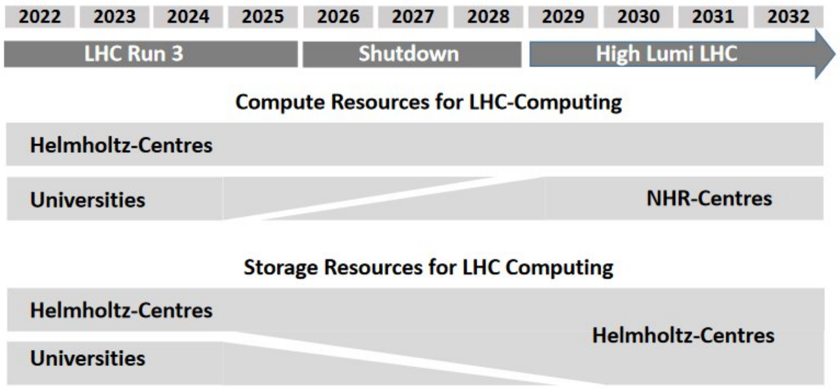

<u>3 successor projects planned to start in Oct 2025</u>:

SUSFECIT (federated computing):
- dynamic resource provision
- aware of power supply
- forecasting
- extended accounting

FUSE (federated storage):
- federated dCache
- caching solutions
- cloud storage

FAIRUM (analysis facilities):
- overarching access portal
- userfriendly interfaces (interactive+batch)
- flexible usage of multiple AFs
- support infrastructure

PUNCH4NFDI

NFDI consortium of **particle, astro-, astroparticle, hadron and nuclear physics** (~9000 scientists in Germany)

Prime goal: a federated and "FAIR" science data platform.
FAIR: Findable, Accessible, Interoperable, Reproducible

Furthermore: knowledge transfer, offering services for efficient research data analysis.

No own resources or 'parallel' resources, mainly interfaces between existing storage and compute resources!

Successor project „PUNCH2" in preparation

Visit booth outside in the lobby!

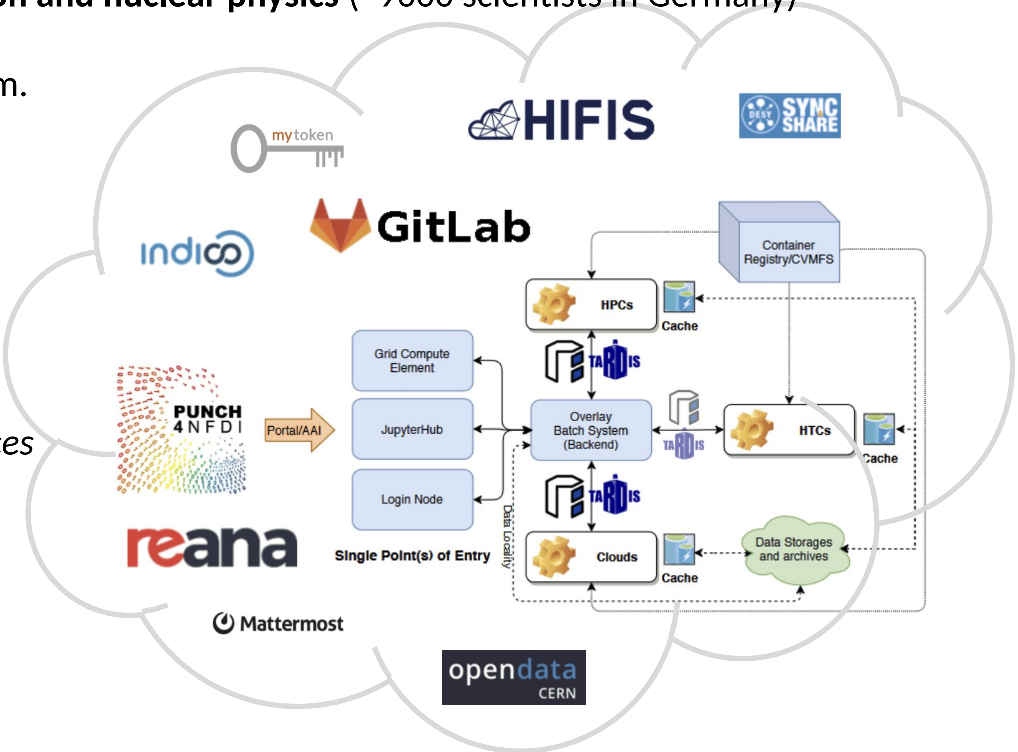

HEP in PUNCH4NFDI

„Do HEP people need PUNCH4NFDI?“

PUNCH compute resources -> no, can‘t compete with HEP infrastructure

BUT: Technology transfer important! HEP provides much knowledge about big federated computing infrastructures. Many services will look familiar to you.
Other communities‘ infrastructures growing to similar scales and challenges!

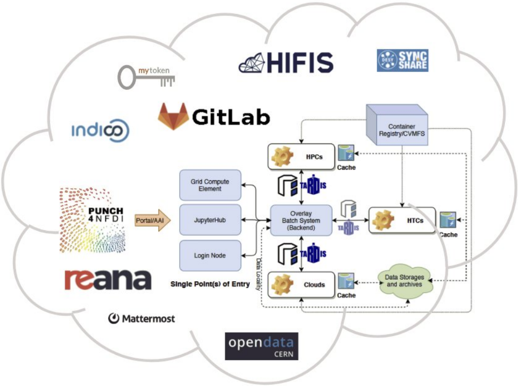

HEP in PUNCH4NFDI

„Do HEP people need PUNCH4NFDI?“

PUNCH compute resources -> no, can‘t compete with HEP infrastructure

BUT: Technology transfer important! HEP provides much knowledge about big federated computing infrastructures. Many services will look familiar to you.
Other communities‘ infrastructures growing to similar scales and challenges!

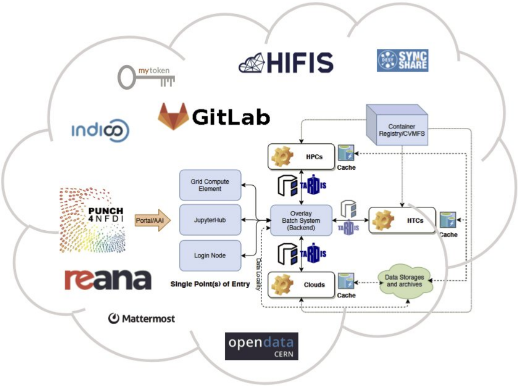

<u>Some thoughts on data FAIRness</u>:

- LHC collaborations putting growing effort in open data provision
- Yet, other communities ask: Why does it take years to publish raw data?
- FAIRness sounds good, but it‘s not trivial:
  - How achieve it with EBs of data?
  - Not just storage, but also accessibility. Not just permissions, but also download rate. Not just data, but also resources to analyse it…
  - How sustainable is FAIRness? What is the value of data (in 100 years)?

Summary

Scientific computing at the Exabyte-Scale

Efficiency at all levels highly important to meet resource limitations and output goals

External developments of hardware and power supply + own R&D for usage + interplay

Collaboration between research fields more and more important:

- knowledge exchange
- common infrastructures & synergies

Thank you for your attention!

Acknowledgements: *Z. Marshall, A. Quadt, A. Schmidt, M. Klute, G. Duckeck, M. Giffels, S. Schumann, Th. Kuhr, S. Korn, M. Horzela, I. Lakomiec*