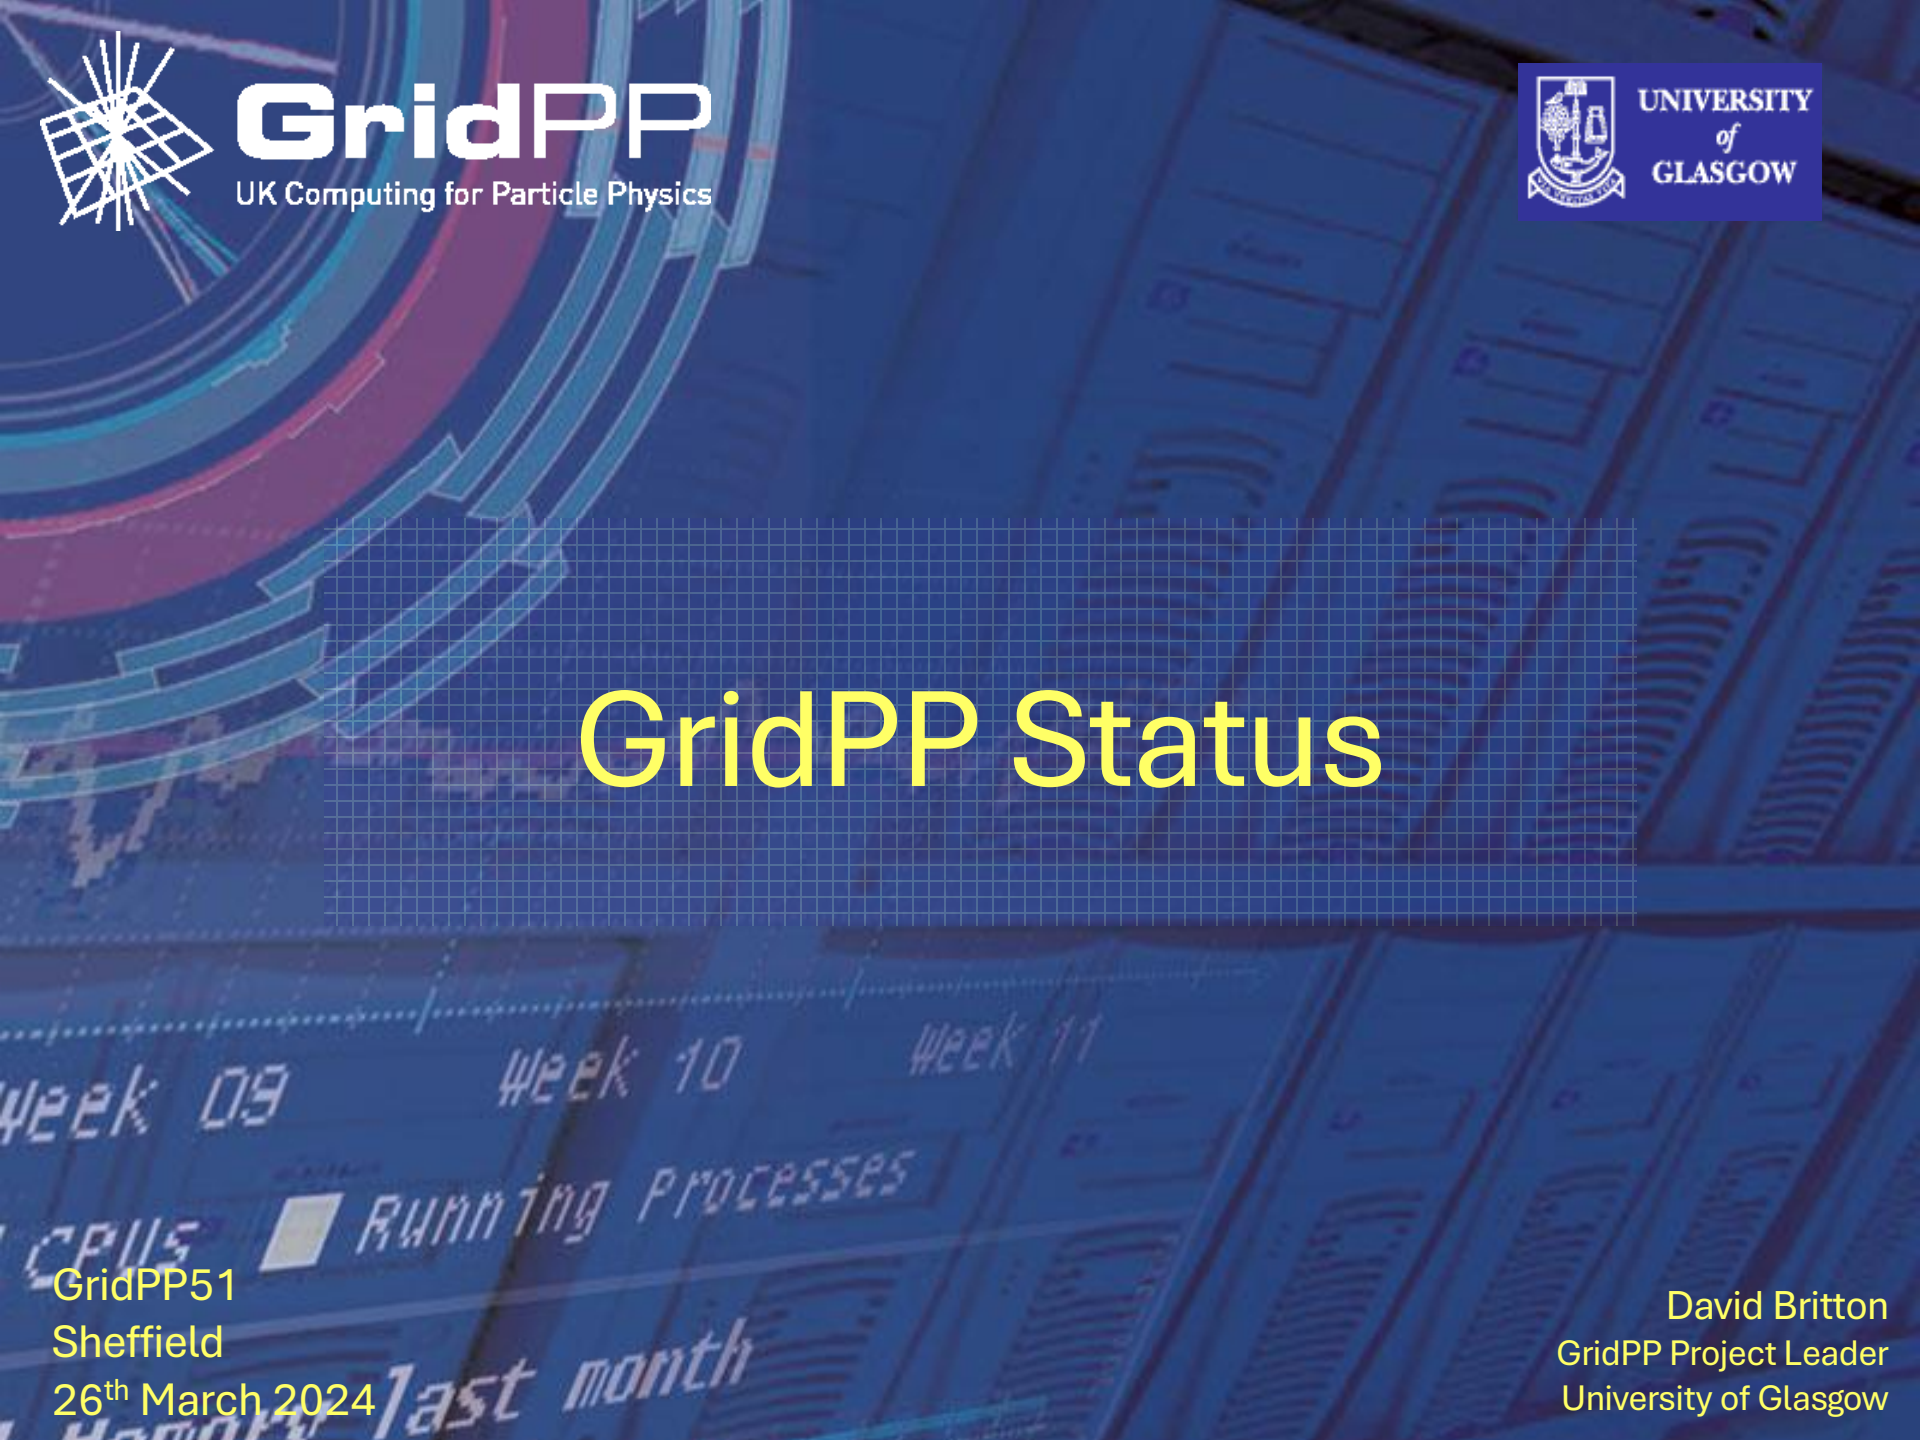
GridPP
UK Computing for Particle Physics
UNIVERSITY of GLASGOW
GridPP Status
Week 09
Week 10
Week 11
Running Processes
CPUs
last month
GridPP51
Sheffield
26th March 2024
David Britton
GridPP Project Leader
University of Glasgow

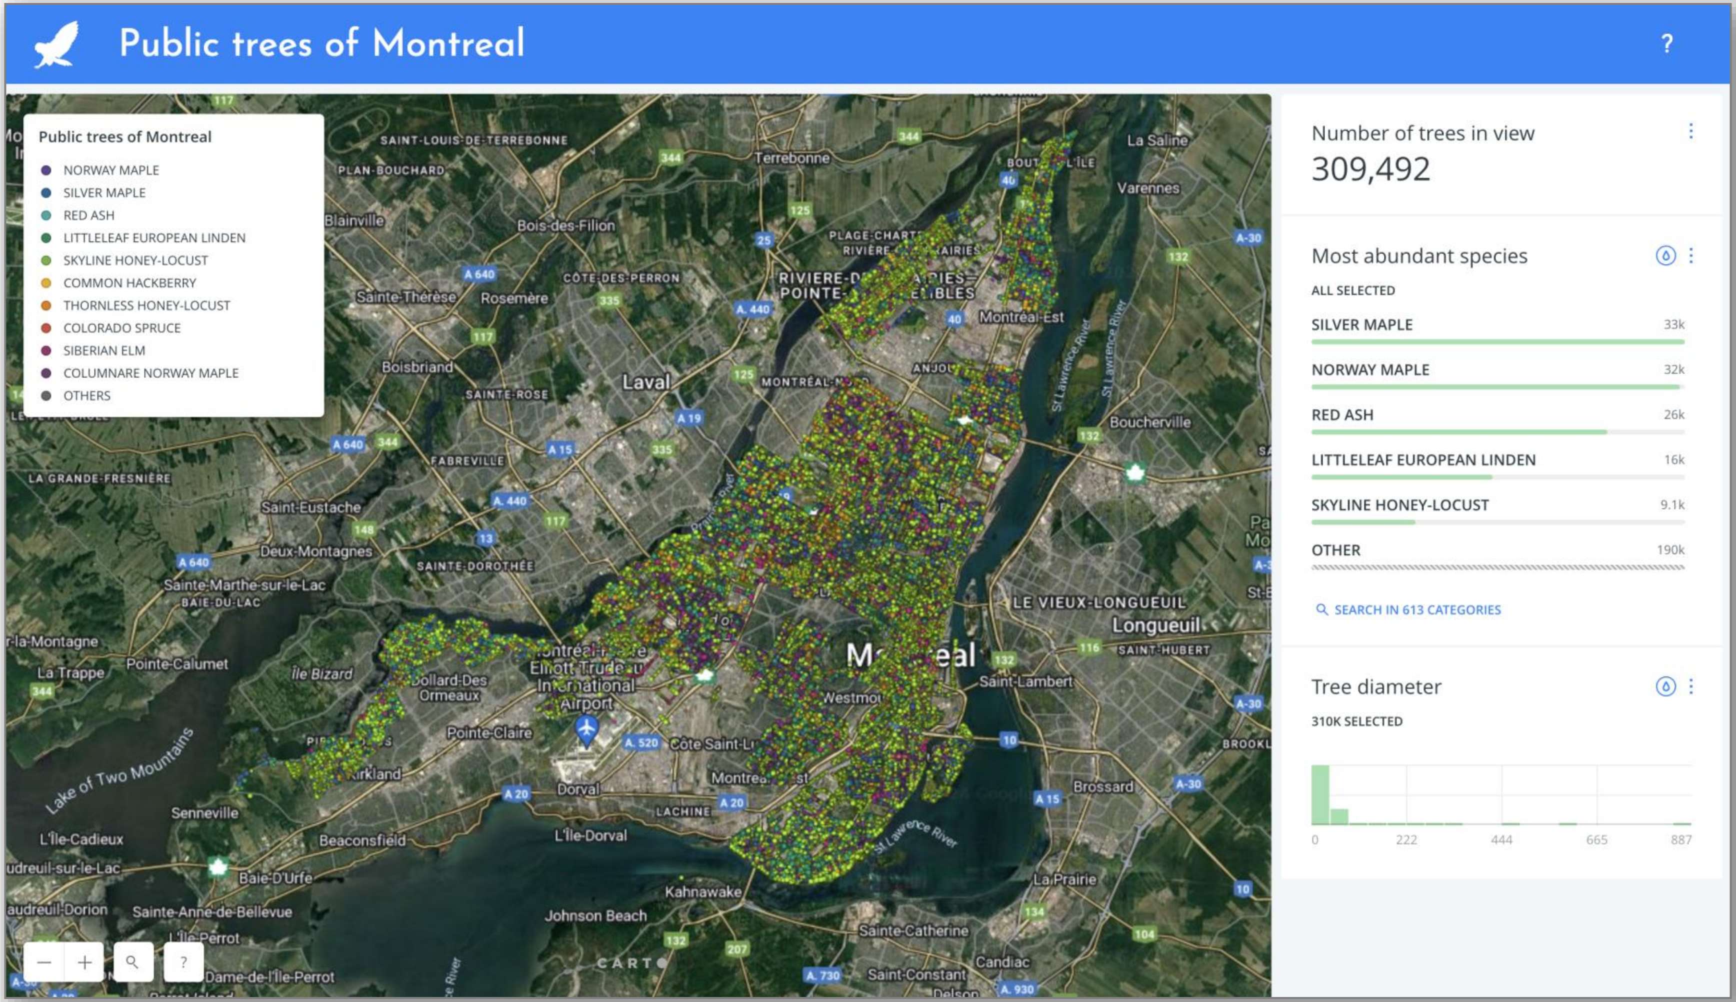
Public trees of Montreal
Public trees of Montreal
NORWAY MAPLE
SILVER MAPLE
RED ASH
LITTLELEAF EUROPEAN LINDEN
SKYLINE HONEY-LOCUST
COMMON HACKBERRY
THORNLESS HONEY-LOCUST
COLORADO SPRUCE
SIBERIAN ELM
COLUMNARE NORWAY MAPLE
OTHERS
Number of trees in view
309,492
Most abundant species
ALL SELECTED
SILVER MAPLE 33k
NORWAY MAPLE 32k
RED ASH 26k
LITTLELEAF EUROPEAN LINDEN 16k
SKYLINE HONEY-LOCUST 9.1k
OTHER 190k
SEARCH IN 613 CATEGORIES
Tree diameter
310K SELECTED
0 222 444 665 887

It's Back!

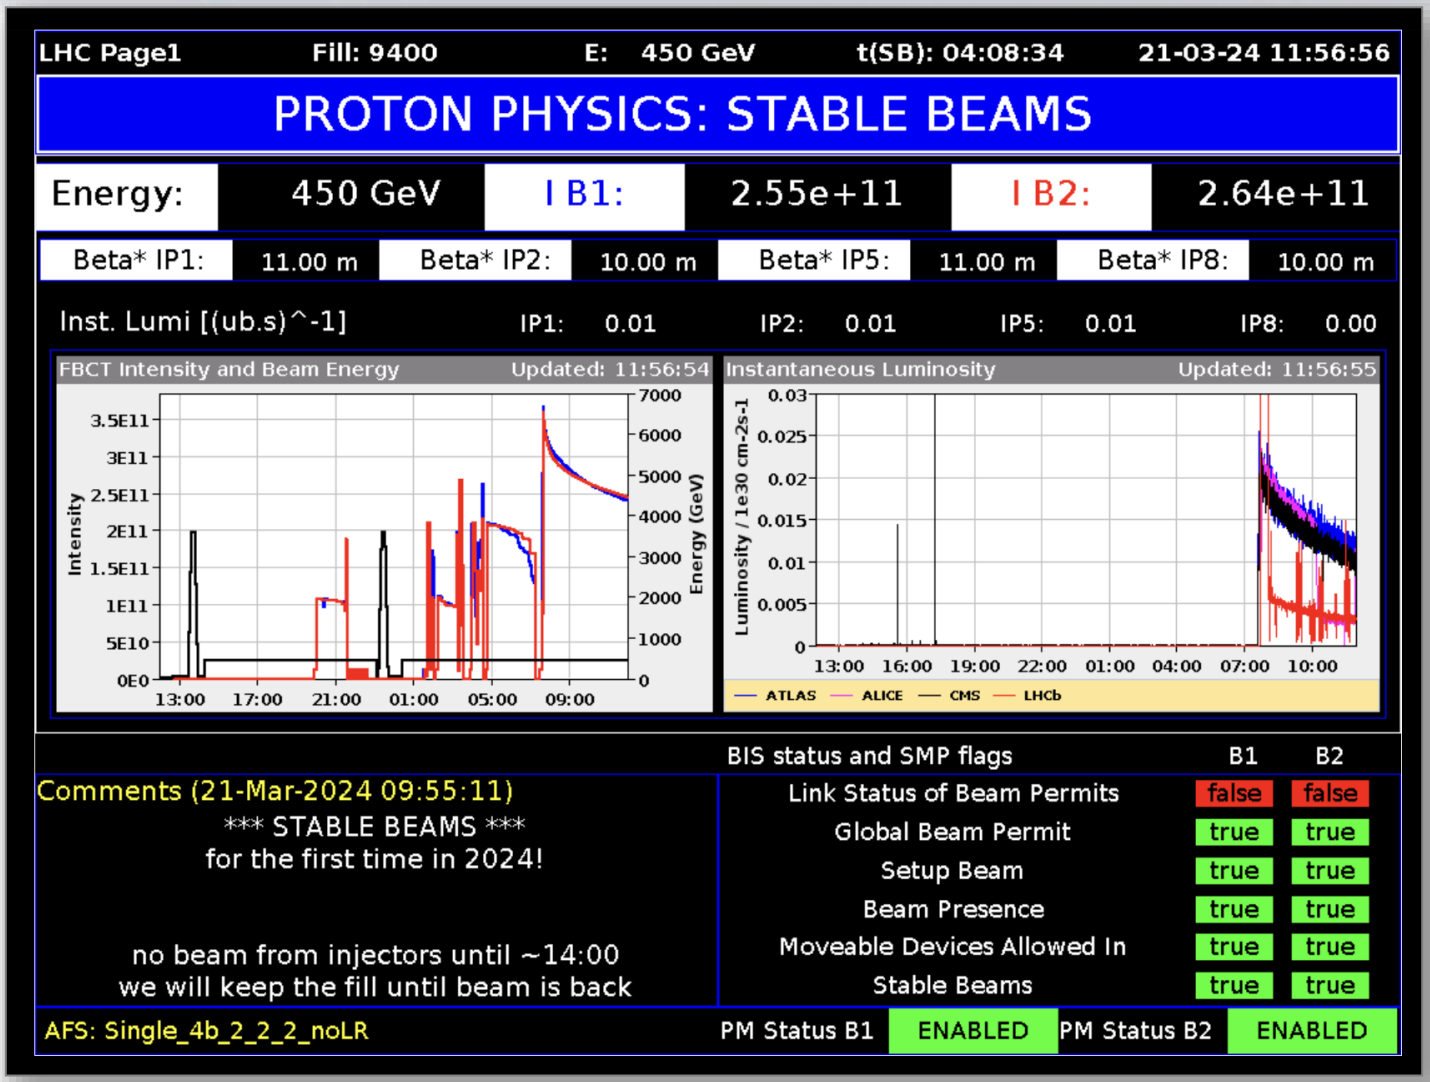
LHC Page1 Fill: 9400 E: 450 GeV t(SB): 04:08:34 21-03-24 11:56:56
PROTON PHYSICS: STABLE BEAMS
Energy: 450 GeV I B1: 2.55e+11 I B2: 2.64e+11
Beta* IP1: 11.00 m Beta* IP2: 10.00 m Beta* IP5: 11.00 m Beta* IP8: 10.00 m
Inst. Lumi [(ub.s)^-1] IP1: 0.01 IP2: 0.01 IP5: 0.01 IP8: 0.00
FBCT Intensity and Beam Energy Updated: 11:56:54
Intensity
Energy (GeV)
Instantaneous Luminosity Updated: 11:56:55
Luminosity / 1e30 cm-2s-1
ATLAS ALICE CMS LHCb
Comments (21-Mar-2024 09:55:11)
*** STABLE BEAMS ***
for the first time in 2024!
no beam from injectors until ~14:00
we will keep the fill until beam is back
BIS status and SMP flags B1 B2
Link Status of Beam Permits false false
Global Beam Permit true true
Setup Beam true true
Beam Presence true true
Moveable Devices Allowed In true true
Stable Beams true true
AFS: Single_4b_2_2_2_noLR
PM Status B1 ENABLED PM Status B2 ENABLED

Schedule (as of 1-March)

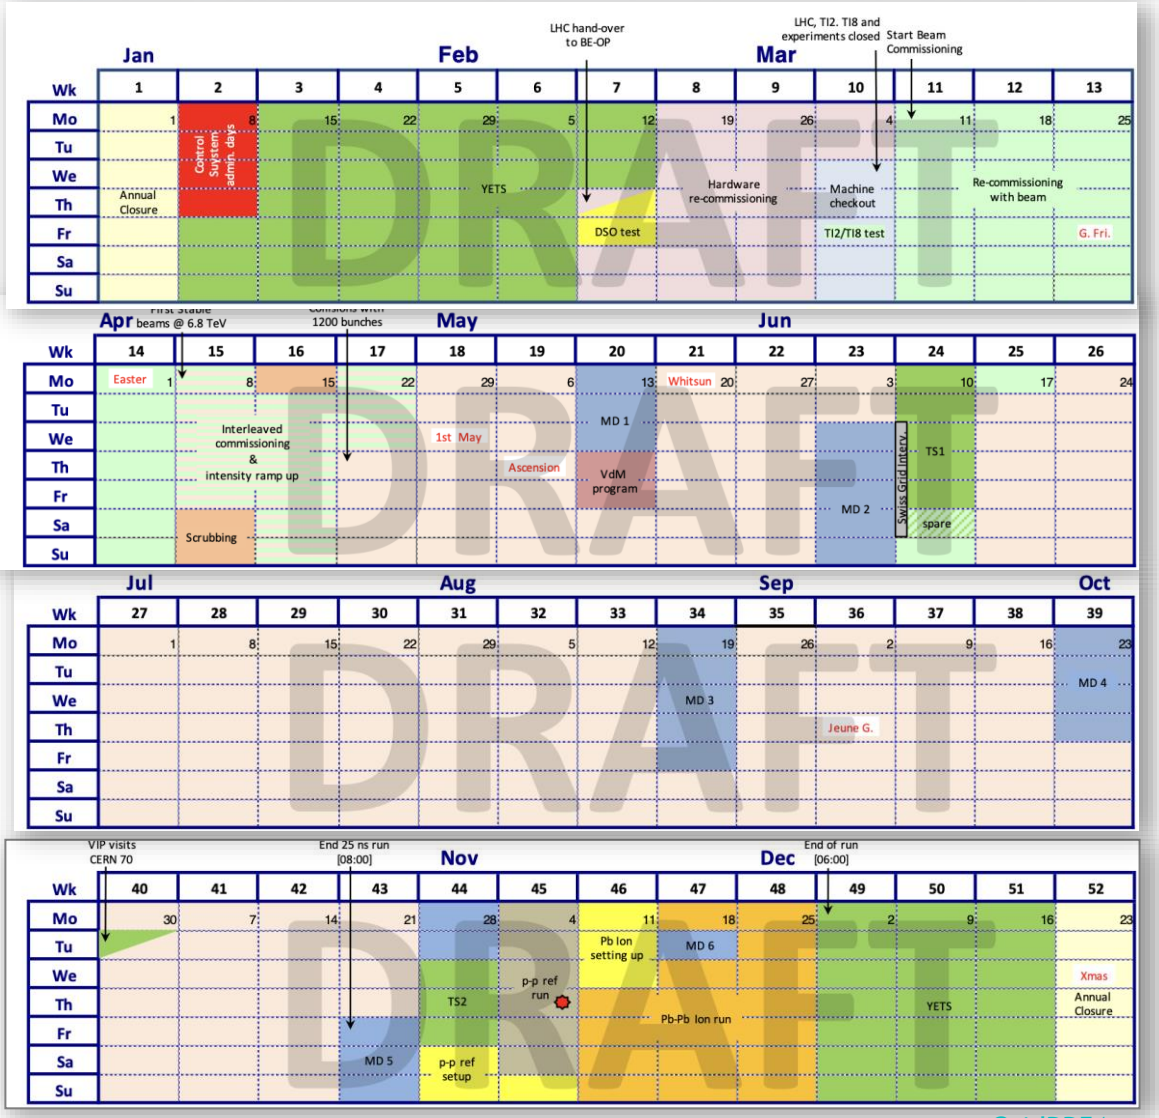

- First 13.5 TeV collisions expected March 26th!
- Then Stable Beams after Easter.
- Then ramp-up to full intensity by April 26th.

No decision yet on possible 5-week EYETS shift (decision currently expected at LMC 27 March): shift would gain 6 physics days (TS2 in 2025 removed)

Run -3

- 2024: Re-starting now; possible 5-week extension in October.
- 2025: The run is scheduled but will start 5-weeks later if 2024 is extended.
- 2026: Run-3 could be extended and LS3 delayed by 1 year.
- CERN will take a decision about LS3 in the week of 11-13 September 2024.

Since Last Time

- GridPP6 becoming GridPP7
  - GridPP7 grants issued.
  - No change in capital situation.
- DC24
  - Completed
- Exascale
  - Lots of noise…

DC24

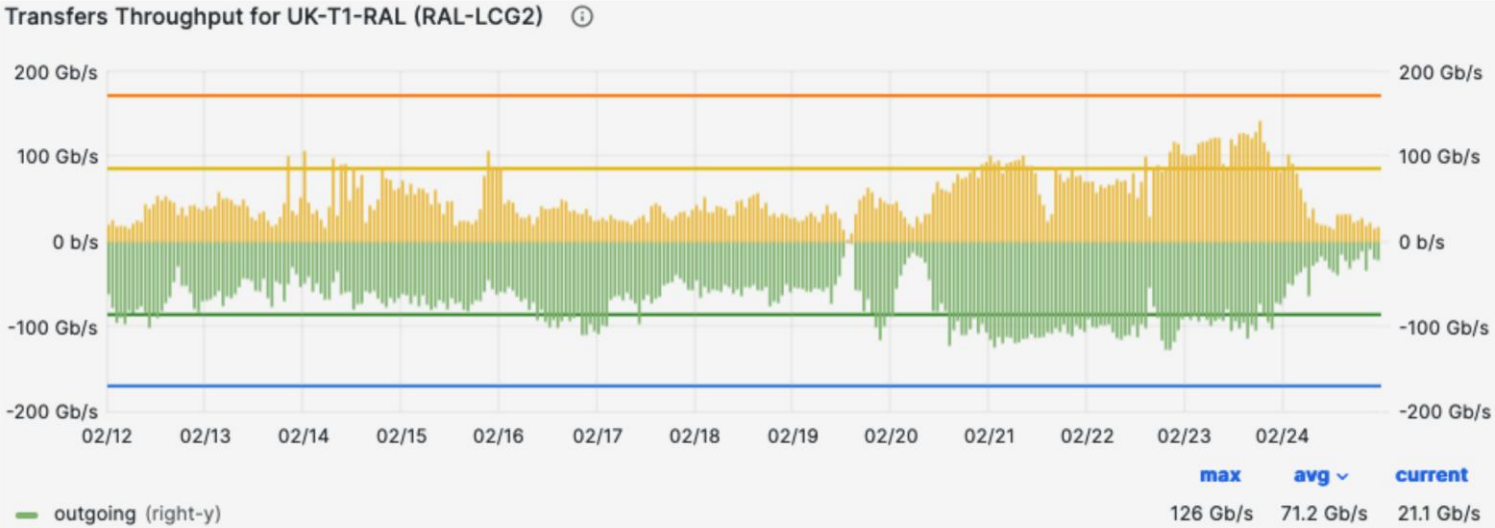

- What does success in DC24 look like? The exercise was designed to challenge the infrastructure and identify the weakest points.
  - If every site had worked flawlessly then that would not have been a success because either the challenge was not designed adequately, or the sites were over-engineered for what is currently needed. E.g., needing more Gateways to handle the DC24 load at RAL was not a failure.
- Success in DC24 is to identify the issues that are currently just below the surface so they can be addressed at the appropriate time.
  - RAL successfully identified and addressed issues around asymmetric routing, packet loss, and Explicit Congestion Notifications (ECN) tuning.

UK- Exascale

- Isambard-3 @Bristol to feature 384 Arm-based NVIDIA Grace (CPU) in early 2024. Nice of them to let us have 1 at Glasgow….
- Isambard-AI will comprise nearly 5,000 NVIDIA GH200 Grace Hopper (CPU+GPU) £225 million and 200+ petaflops… in 2024.
- DAWN @Cambridge: 1,000 Intel Data Centre GPU Max Series chips and 500+ Intel Xeon CPUs…. Up and running.
- Talk of a Phase 2 DAWN that will be 10x bigger… in 2024? Bids against £500m?
- Next Thing (few years): An ‘exascale machine’ – will probably be achieved with a GPU-only design. Likely to be a prototype phase first to ensure it scales.

- Bottom line is >>£1 billion being spent on the above.
- GridPP7 hardware shortfall is <<0.5% of this total…

- What about a WLCG allocation on UK HPC in the future?
- Experience from (at least) USA, Spain, Germany, Italy…

Can we use HPC?

- Of course, to first order the experiments (not us) determine whether HPC can be used... and agree with the CRSG whether they can be pledged.

- But.... HEP use cases have little technical need for High Performance Computing per se - there is little or no use of, or need for, MPI applications. It is not worth investing money on fast interconnects for HEP workloads.

- And... HEP workloads are gradually making more use of GPUs, but it's taken years of development. GPU-heavy HPC architectures are not really used efficiently by HEP work. Obviously want to avoid having idle GPUs for both economic and environmental reasons (GPUs idle at a much higher fraction of full-power than idle CPUs).

Can we use HPC?

- HPC machines have been described as 'snowflakes' because, historically, they have tended to be unique and melt away after 5 years. This has cost the HEP community a lot of work over the years and success has been variable:
  - VEGA HPC was designed with ATLAS in mind as a client – was very successfully used during the start-up phase before they found other users.
  - MareNostrum HPC @BSC cannot even run CVMFS as an edge service and after vast effort by HEP, can only be used for Sim workloads.
- WLCG attempting to create guidelines for 'minimum HEP requirements' that could be factored in at the design phase (eg EUROHPC).

Can we use HPC?

- So, technically yes, if designed correctly, and with significant (additional) investment of effort, we can use HPC if they are not just GPU…
- …but it is unlikely to be the economically or environmentally optimal choice… *Use of HPC in WLCG is largely a political decision.*
- WLCG work is *not* elective, it's deterministic from the science already approved; this means that it is not a good match to the normal competitive bid and allocation process typically used to share HPC resources.
- A more useful position would be to view HPC as potential opportunistic (above pledge) resource. The potential gain needs to be balanced with the effort needed to realise (though WLCG and the experiments hope to gradually reduce this).

Summary: GridPP Thoughts on Exascale

- GridPP supports the drive towards UK Exascale HPC resources to enable UK science.
- With possibly significant and sustained additional effort, HEP workloads could make use of such machines (if they are not too GPU-centric), though they may not be the optimal way of providing resources for HEP from an economic or environmental perspective.
- The most useful scenario is that HEP workloads back-fill unused cores (increasing the overall utilisation efficiency of the machine) and providing opportunistic resources above the pledge that helps the UK compete.

GridPP7 Strategic Objectives

- To deliver STFC's MoU commitment to the WLCG, meeting the challenges of steadily increasing data rates and data volumes throughout LHC Run 3.
  - Operations and Evolution
- To prepare for the 2029(?) start of HL-LHC (LHC Run 4) by influencing WLCG's future technical direction and contributing to developments.
  - Innovation and Development
- To provide broader benefit to STFC through partnership in IRIS and sharing the GridPP infrastructure and services with non-LHC and non-HEP communities.
  - Collaboration and Community

Innovation & Development

	WP-D Innovations summary		
1)	Tier-1 Data throughput improvements	**1.00**	Tier-1
2)	Tier-1 Container orchestration	**0.50** ~~**0.25**~~	Tier-1
3)	Tier-1 Token support	**0.25**	Tier-1
4)	Data Management	**0.60**	EDI(0.5); LAN(0.1)
5)	DOMA for analysis infrastructure for HL-LHC	**0.55** ~~**0.30**~~	MAN
6)	Energy and NetZero	**2.00** ~~**1.00**~~	GLA(1); QM+T1(0.5)
7)	GPU (Etc.)	**1.05** ~~**0.80**~~	T1(0.5); IC(0.25); GLA(0.2); LIV(0.1)
	Total	**5.95** ~~**4.20**~~	

1) Work focused on preparation for, and participation in, the WLCG Data Challenges.
2) Specific work proposed at RAL in conjunction with SwiftHEP (Kubernetes orchestration).
3) Deploy a unified authorisation service that will supports Tokens.
4) Effort required to do both planned and unplanned work:
   - (i) further XCache deployment, monitoring and support across UK for diskless sites;
   - (ii) support for the transition to token-based authorisation for SEs;
   - (iii) developing Rucio monitoring technologies;
   - (iv) development of services for remote XRootD monitoring;
   - (v) continued support for the GridPP StashCache instance;
   - (vi) development of data flow rate monitoring using FTS logs.

Innovation & Development

	WP-D Innovations summary		
1)	Tier-1 Data throughput improvements	**1.00**	Tier-1
2)	Tier-1 Container orchestration	**0.50** ~~**0.25**~~	Tier-1
3)	Tier-1 Token support	**0.25**	Tier-1
4)	Data Management	**0.60**	EDI(0.5); LAN(0.1)
5)	DOMA for analysis infrastructure for HL-LHC	**0.55** ~~**0.30**~~	MAN
6)	Energy and NetZero	**2.00** ~~**1.00**~~	GLA(1); QM+T1(0.5)
7)	GPU (Etc.)	**1.05** ~~**0.80**~~	T1(0.5); IC(0.25); GLA(0.2); LIV(0.1)
	Total	**5.95** ~~**4.20**~~	

5) DOMA and AFs:
   i. Delivery of data using Content Delivery techniques with caches (data challenge activities).
   ii. Integration of IAM and tokens.
   iii. Developing a GPU partition at a core Tier-2 as a proto-AF.
6) NetZero
   i. GLA: Energy efficiency of different hardware; Operation optimisation;
   ii. QMUL: 4M: Model/Measure/Monitor/Moderate
   iii. Tier-1: Storage energy efficiency.
7) GPU
   i. Tier-1: ALICE reconstruction and joint SwiftHEP work.
   ii. IC: CMS reconstruction.
   iii. Enabling GPUs on the Grid: GLA and LIV.

WLCG Strategy

- Despite our plans, we also need to remain agile in GridPP and do things that are most useful in the current context:
  - We probably won't do all the things we proposed.
  - We will probably do things we didn't propose.
  - We need to keep sight of the big picture.
- WLCG has effectively just started (1/Jan/24) a new 4-year phase and is developing a strategy document that will be one of the main discussions at the May WLCG workshop in Hamburg.
  - Strategy draft is based on interviews with many people (including Roger and me).
  - I have helped write and edit the draft.
  - Pete has provided comments.
- So GridPP is helping shape the WLCG strategy and GridPP7 development work should align with that as appropriate.

Summary+

- Run-3 is recommencing.
- GridPP7 Starts next week (with a day off!)
- We continue to have challenges with capital funding for hardware.
- We have some interesting development work to do but we should ensure we are aligned with the wider WLCG strategy, which we are helping to shape.

The GridPP52 collaboration meeting will take place in Ambleside 28th to 30th August 2024.