



Computing Requirements for Geant4

Computing Performance Task

K. Genser, S. Y. Jun, J.Yarba

Fermilab

Jan. 26, 2021

Geant4 Computing Performance Task (I)

- Monitor Geant4 software (CPU & memory) throughout the development cycle to
 - Identify opportunities for code improvements
 - Provide feedback to the developers & working group leaders
- Why is it important:
 - Geant4 is used by every HEP experiment, and in other domains
 - One of the most CPU-heavy tasks in HEP experiments
 - Potential benefits
 - Performance improvement = \$\$\$
 - Detector design/optimization = \$\$\$...\$, reconstruction tools, physics background studies...
 - Promote best practices in developing efficient HEP software
- We have been carrying out this activity for about 10 years
- It is Fermilab institutional responsibility to Geant4
- **Many thanks to the experts for all the support**

Geant4 Computing Performance Task (II)

- Obligations
 - Provide results within 48 hours after the release notification
 - Close all open issues before the next release
- Workflow (as of decommissioned Wilson cluster)
 - Build Geant4 and applications – ~4 hours (total, wall clock), on tevnfsg4
 - Run CPU profiling jobs – batch (see following slides)
 - Open|SpeedShop + Geant4-based application
 - Run memory profiling jobs – ~24 hours (wall clock), on tevnfsg4
 - IgProf + Geant4-based application
 - Perform analysis on the results – ~12 hours (wall clock), on tevnfsg4
 - Reporting/publishing – on the head node and the central web server
 - <https://g4cpt.fnal.gov>
 - NOTE: Careful performance evaluation/optimization requires “quiet”, uniform, stable resources, i.e., a generic Grid node is not optimal

Geant4 Computing Performance Task (III)

- Number of profiling rounds per year: ~35

Release Type	Base	Associated	Total/Year
Development	10		10
Semi-annual	2 official	3 candidates/release	8
Patches	3 patches	2 candidates/patch	9
Special Requests	5-10		5-10
Total	20-25	12	32-37

The number of profiling rounds used to be larger in the past. However, we have been improving the workflow to efficiently use the available resources. For example, we have switched to using static libraries instead of shared ones which has allowed to:

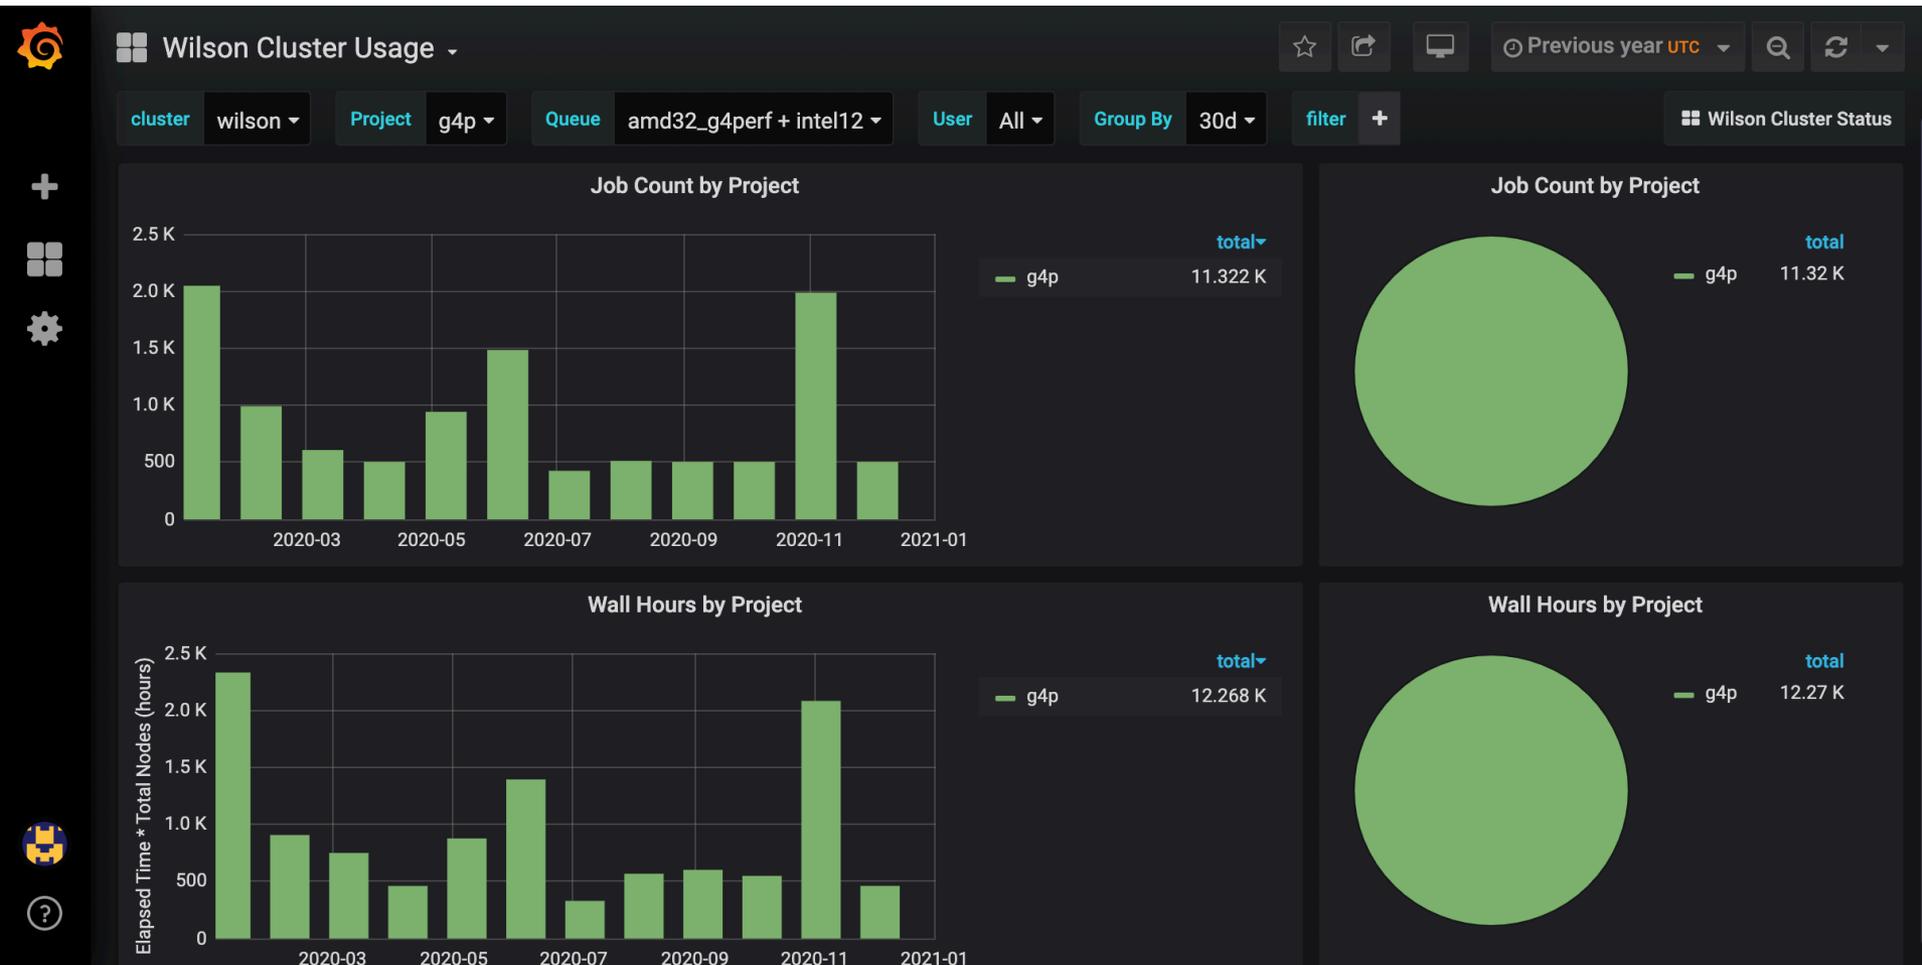
- Improve the speed by ~12%
- Improve the results stability, which in turn allowed to eliminate re-profiling of a previous release in the development cycle, and sometimes other releases, which we used to do in order to avoid systematic shifts occurring when using the shared libraries; this has cut the number of profiling rounds by at least 10 or more

Workload and Resources Used for Geant4 Computing Performance Task Until December 2020

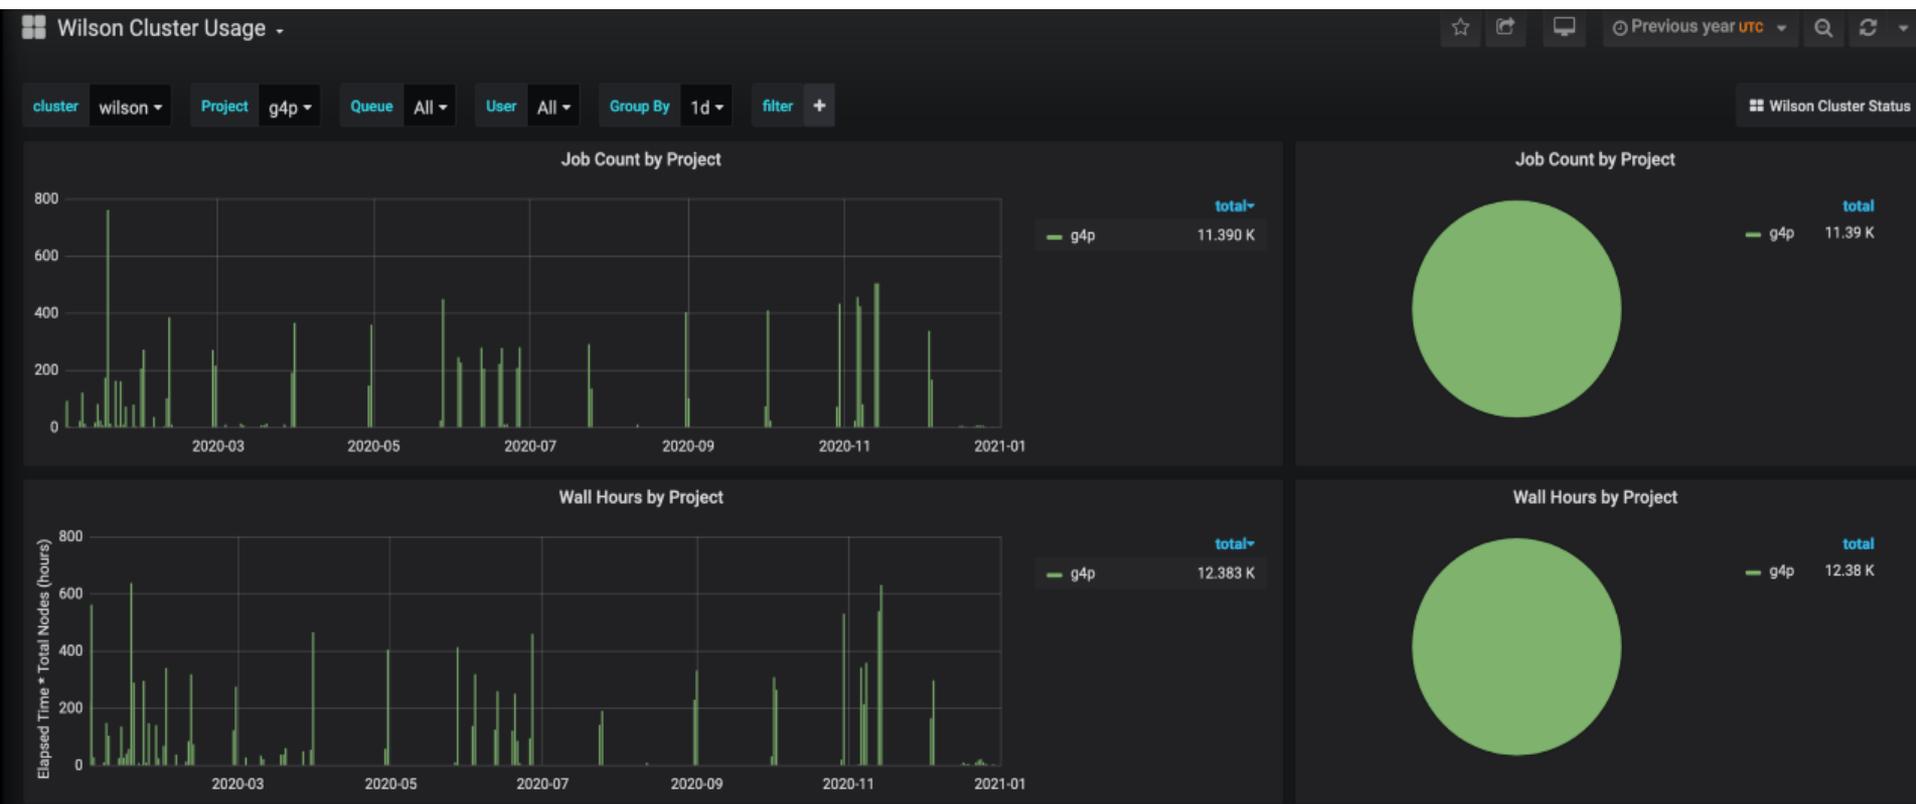
- Typically, we have ~170 profiling application configurations
- ~2/3 of them need to run 90 times (30cores x 3nodes, amd32)
- ~1/3 of them need to run 30 times (10cores x 3 nodes, intel12)
- Multiple runs are needed to minimize statistical fluctuations, as CPU time needs to be measured with an accuracy of ~1%
- Several applications may need to run for about 5-6 hours (per single run)
- Most of the jobs finish in between a few minutes and half an hour, giving way to subsequent jobs in the queue

Test Type(s)	Nodes	Cores
Sequential (amd32_g4perf)	21	32
Sequential (intel12)	25	12
Multi-threaded (amd32_g4perf)	2	up to 32
Multi-threaded (intel12)	2	up to 12
Memory (tevnfsg4)	1	1

Use of amd32_g4p+intel12 by g4p in 2020



Use of resources by g4p shown on the daily basis in 2020



Fraction of Resources Used by g4p in 2020



**g4p - 2.03% of the total use;
much less of the available time**

Summary

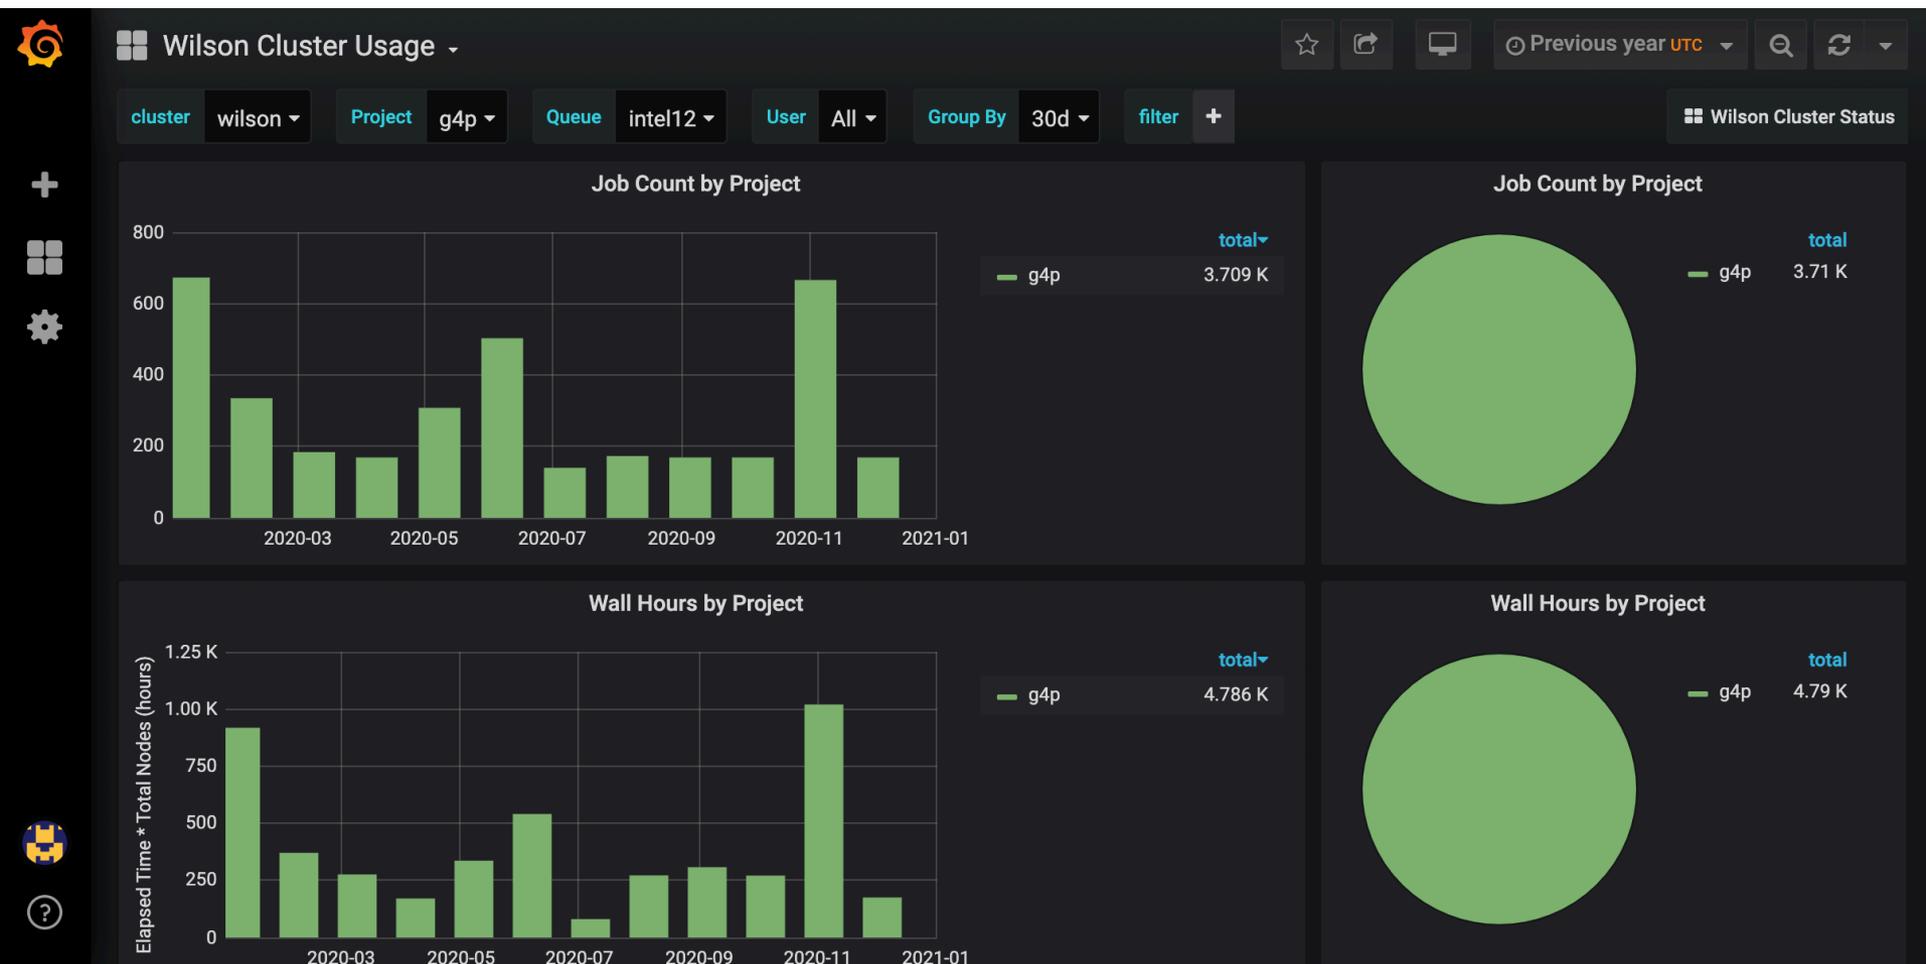
- Geant4 is a key software package in planning, optimizing, and running any HEP experiment
- Improving its performance translates into saving \$\$\$...\$
- Carrying on the Geant4 QA tasks requires resources
 - Certain groups of **results are expected to be delivered within 48 hours** from the notification of the new release which translates into **special priorities for using computing resources**
 - However, intense, **high-priority** use of computing resources is needed only **periodically and only for a limited time**
 - Overall load by the task is **~2% of the total time used by all the projects (and much less of the available cluster computing time)**
 - Having a dedicated queue allows to use resources only when needed, and release them to other users as soon as jobs complete
- **Many thanks to the Wilson cluster team for support and advice; we would not be able to carry the task without your help**

BACKUP SLIDES

Use of amd32_g4perf by g4p in 2020



Use of intel12 by g4p in 2020



Estimate of core hours potentially available per year on the decommissioned Wilson cluster:

$$365 \times 24 \times (32 \times 32 + 25 \times 12) = 11\,598\,240$$

Estimate of core hours potentially available per year on WC-IC:

$$365 \times 24 \times (90 \times 16) = 12\,614\,400$$