GlueX-doc-4356-v1

Offline Software Status Winter 2020

Mark M. Ito
Jefferson Lab
GlueX Collaboration Meeting
February 15, 2020

Outline

- MCwrapper: Present and Future
- CCDB Stabilization
- HDGeant4: Event-Level Efficiency Calculations
- Miscellaneous Topics
- Outlook and Conclusions
- Communication Channels

MCwrapper

Thomas Britton

Growth

- Continued growth and use (perhaps more now that I am no longer Hall-D)
  - About 40% of GlueX has used MCwrapper-bot
  - 883 submitted projects
  - Ever more use by individuals
- This growth has come with some pains…
  - File locking issue
  - Scosg16 self ddos
  - Excessive work disk usage

Pains

- Starting in late December MCwrapper-bot projects started to stall, forever stuck at 99+%
  - This is still a mystery
    - What is known is mcsmear seems to play a part in it
- Other projects stall much much sooner
  - Smattering of problems many related to use of SWIF
    - File locking issues, file system congestion
    - Fun fact: MCwrapper broke SWIF and set up a terrible feedback loop
      - Currently Mcwrapper-bot does not submit to SWIF but plans are to re-enable JLab submissions soon via SLURM directly
- Many more safe guards have been put in place to combat potential missing data

CCDB saga

- Started using mysql a long while ago
  - Hallddb under-powered causing storms
- Moved to sqlite when on the farm
  - Brought the group disk down
- Switched to using one of Mark I's 100 copies on the work disk
  - File locking issue
- Switched to hallddb-farm
  - Storms
- Switched back to 100 copies
  - Work disk load issues
- Took a copy from dtn1902
  - Current state
- We are investigating moving back to hallddb-farm....

Future

- settle on this whole ccdb thing...
- Redo JLab submits for MCwrapper-bot
  - Use SLURM directly
  - SWIF(2) will still be supported
- handle reading in an input hddm file containing more than one run number
  - this is less trivial than I had hoped
- add in generator post processing
  - evt_gen
  - other custom fiducials

CCDB Stabilization

- Multiple copies on the work disk
  - Maintaining this system with 100 copies
  - Work disk locking problems solved by server reboot
- New MySQL Server
  - hallddb-farm.jlab.org
  - DNS alias for four VMs: hallddb-a, b, c, and d
  - Jobs on farm are pointed here for CCDB constants
- CCDB 2.0
  - New intermediate table should speed up queries
  - Needs benchmarking
  - Look for call for beta testers
- SQLite files with small subset of database table rows (run-specific, time-specific)

Comparisons of Geant3 and Geant4

Nilanga: $\gamma p \rightarrow K^+\Sigma^0,\ \Sigma^0 \rightarrow \Lambda\gamma,\ \Lambda \rightarrow p\pi^-$

- Distributions for K^+
- Strange dips in G3 distributions

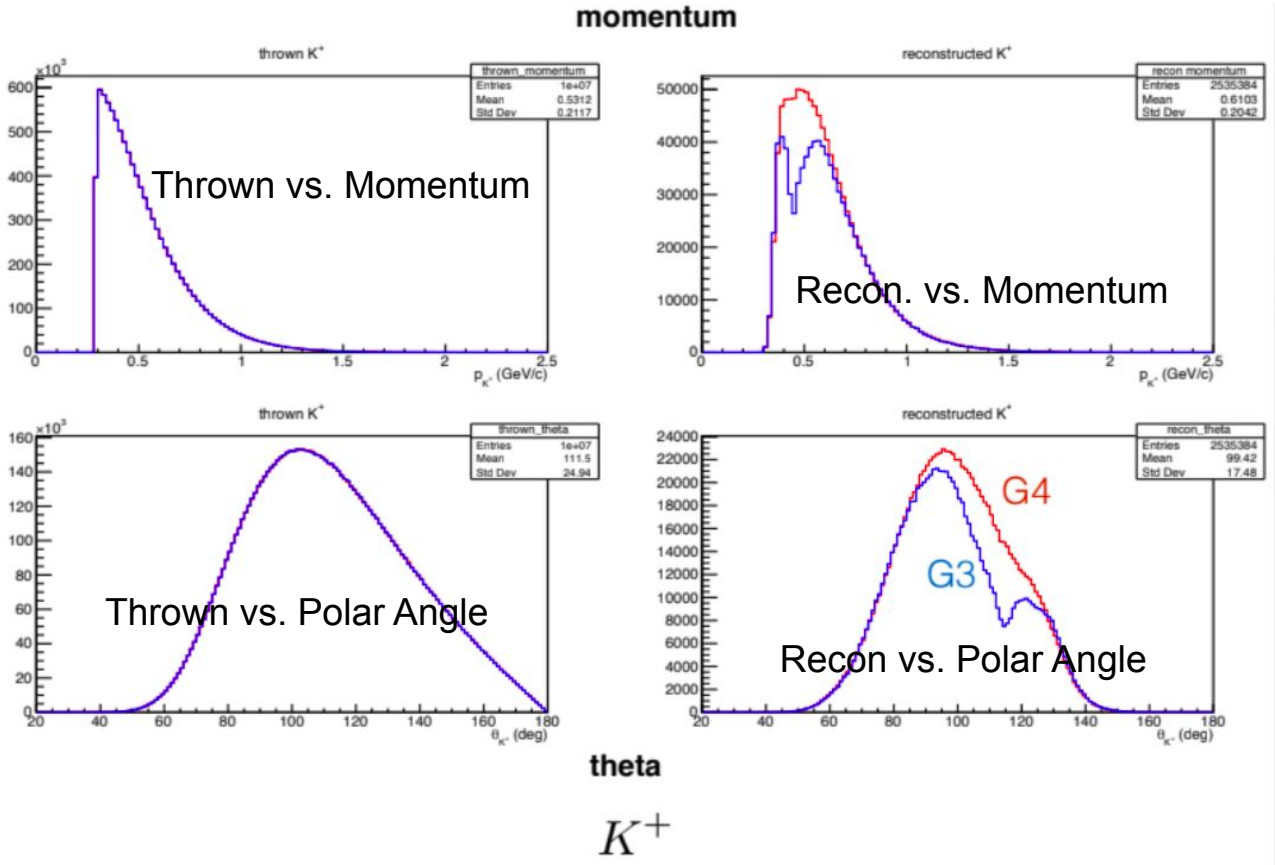

Colin: $\gamma p \to \eta \pi^- \Delta^{++},\ \eta \to \gamma\gamma,\ \Delta^{++} \to p\pi^+$

- Overall event acceptance different as function of polar angle of eta-pi system in the Gottfried-Jackson frame.

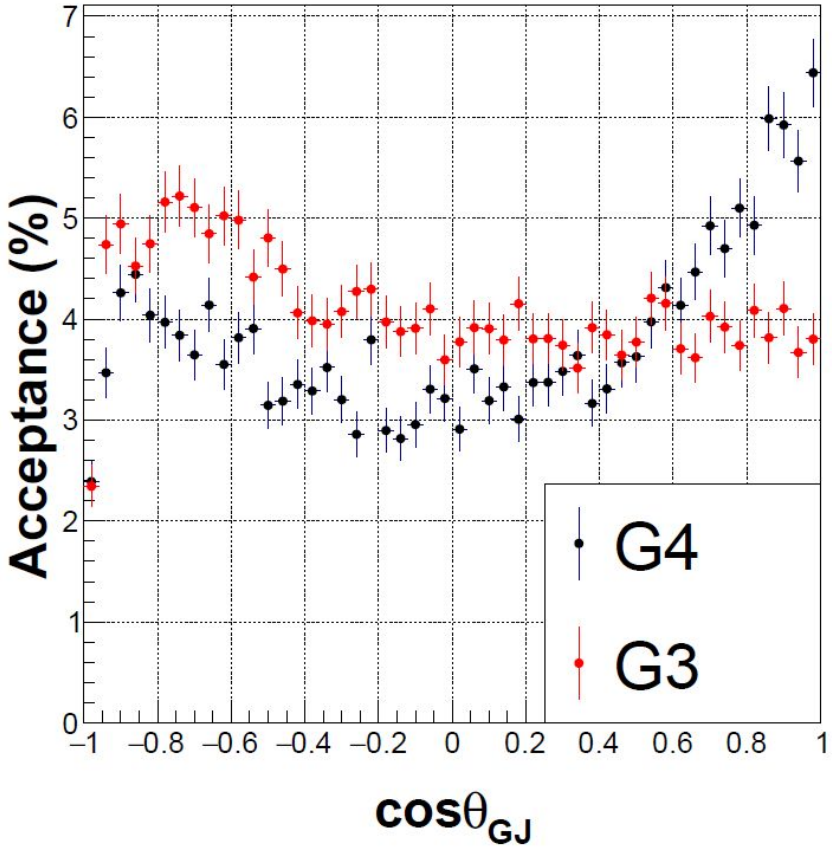

Alex: $\gamma p \rightarrow \rho^0 p,\ \rho^0 \rightarrow \pi^+ \pi^-$

- Overall event acceptance as a function of beam photon energy shows disagreement between Geant3 and Geant4

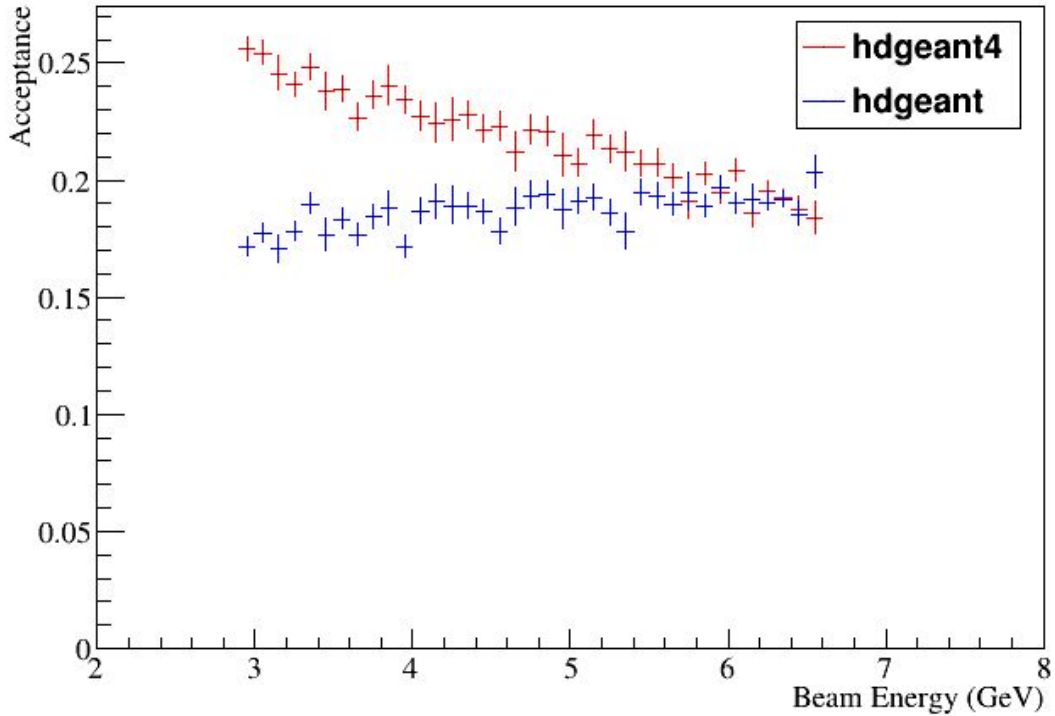

Alex: $\gamma p \to \rho^0 p,\ \rho^0 \to \pi^+ \pi^-$

- BCAL timing for pions shows high side tail for Geant3
- Analysis library has default cut at ±1 ns

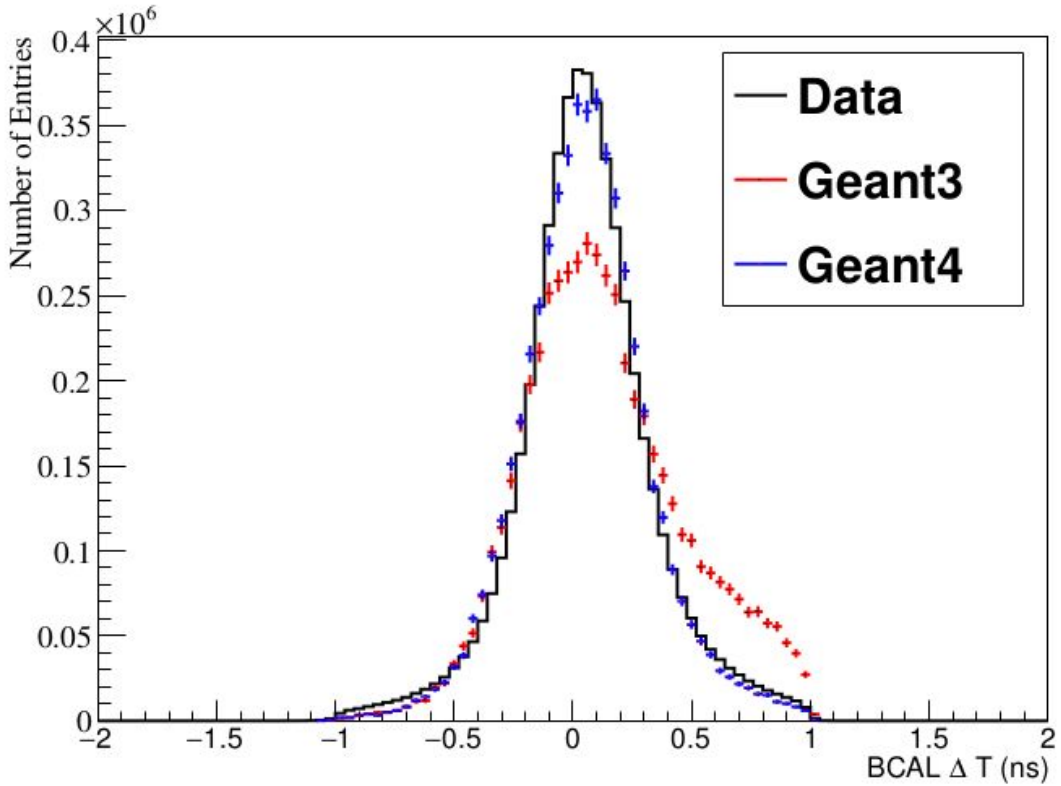

Peter: $\gamma p \to K^+ \Lambda(1520),\ \Lambda(1520) \to pK^-$

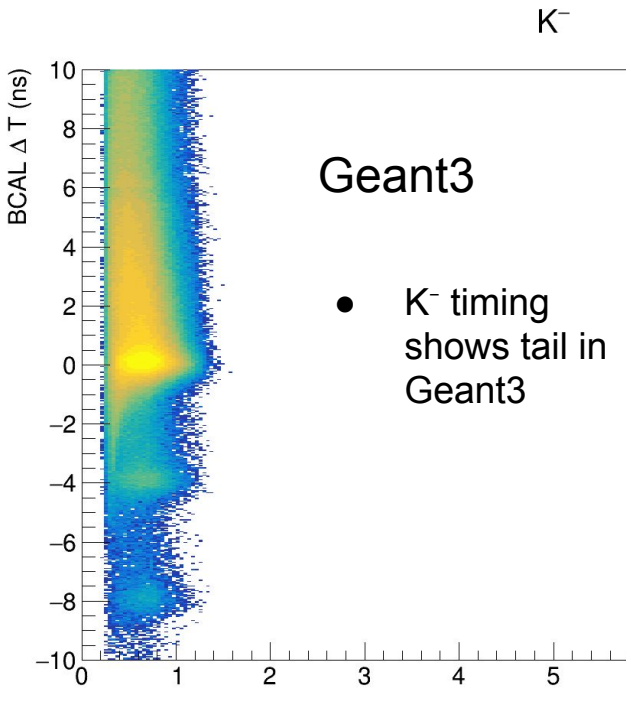

- K^- timing shows tail in Geant3

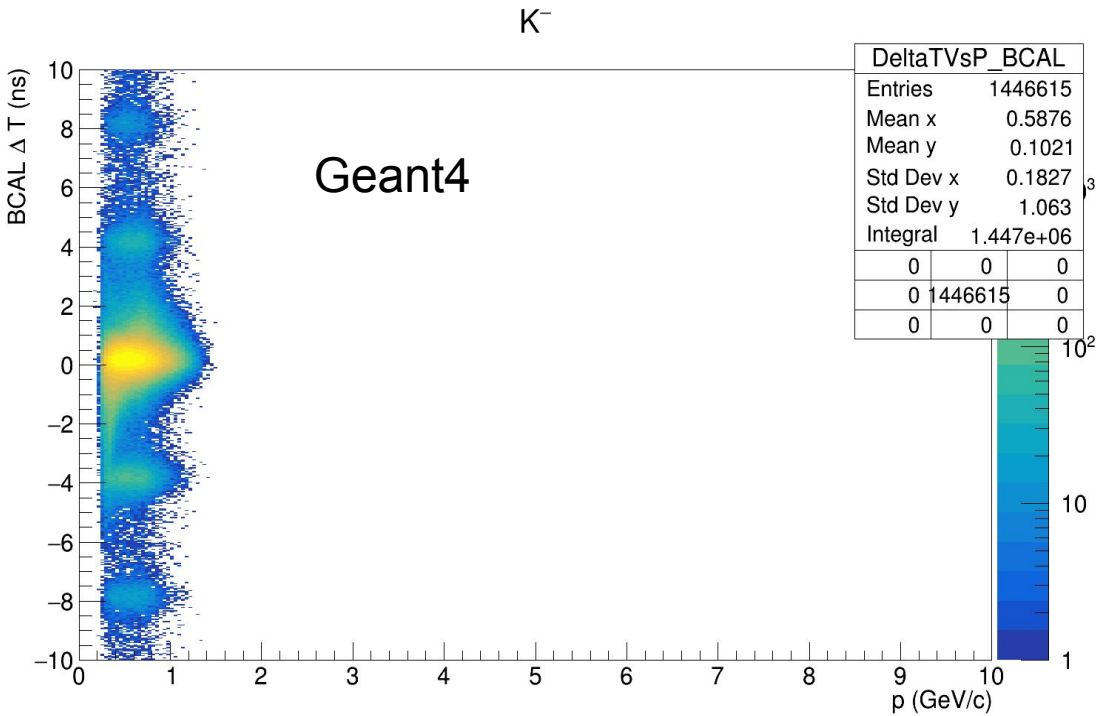

Alex: $\gamma p \rightarrow \rho^0 p,\ \rho^0 \rightarrow \pi^+ \pi^-$

- Widen timing cut to ±5 ns
- Much better agreement

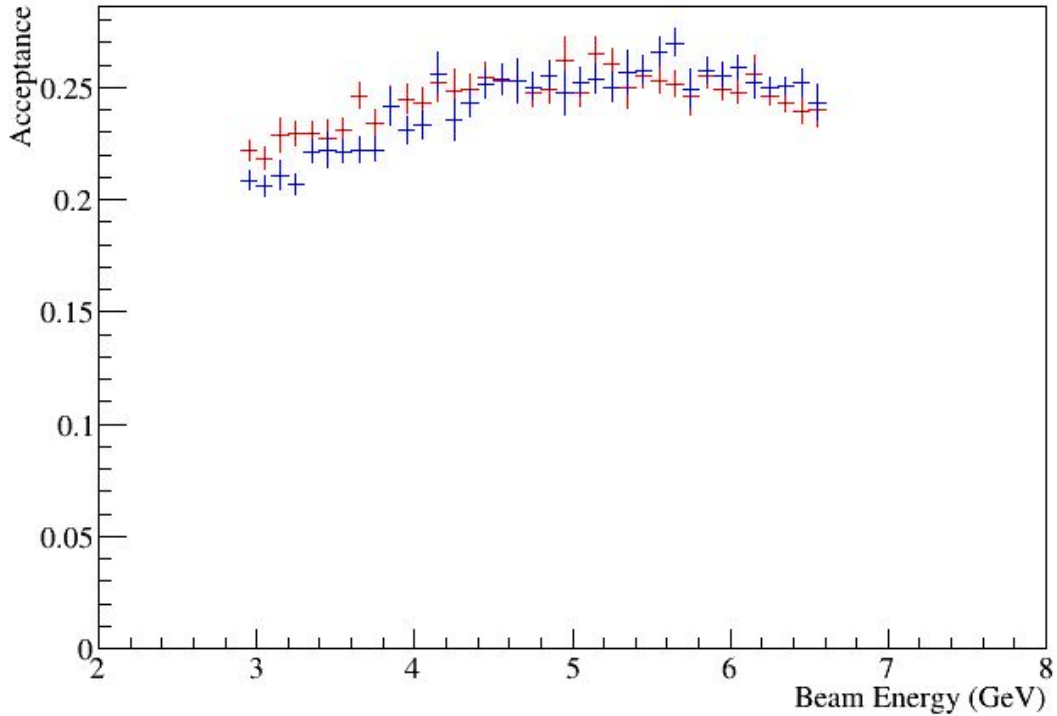

Conversion from Python 2 to Python 3

- Support for Python 2 officially ended January 1, 2020

> *Being the last of the 2.x series, 2.7 will receive bugfix support until 2020. Support officially stops January 1 2020, but the final release will occur after that date.*

- Most of our scripts are in Python 2
- Most of our scripts will not work with Python 3
  - Notably all of our SConstructs and SConscripts
- Most of the changes are minor, e.g., change “print x” to “print(x)”
  - ...but not all of them…
- Eventually the conversion will have to be done
  - Volunteers?

JLab Scientific Computing Resources

- More farm nodes, Linux distribution = CentOS 7.7 (see plot)
- More Lustre disk space with new Lustre software, now at 3.0 PB from 1.5 PB, will to go 5.0 PB soon
- Support on the CentOS 7.7 farm for
  - XROOT
  - CVMFS
  - Singularity
- https://jupyterhub.jlab.org is available for JupyterLab work using farm nodes
- New-to-us Slurm resource-manager/job-scheduler opens up useful options (see figure)

New Farm Nodes

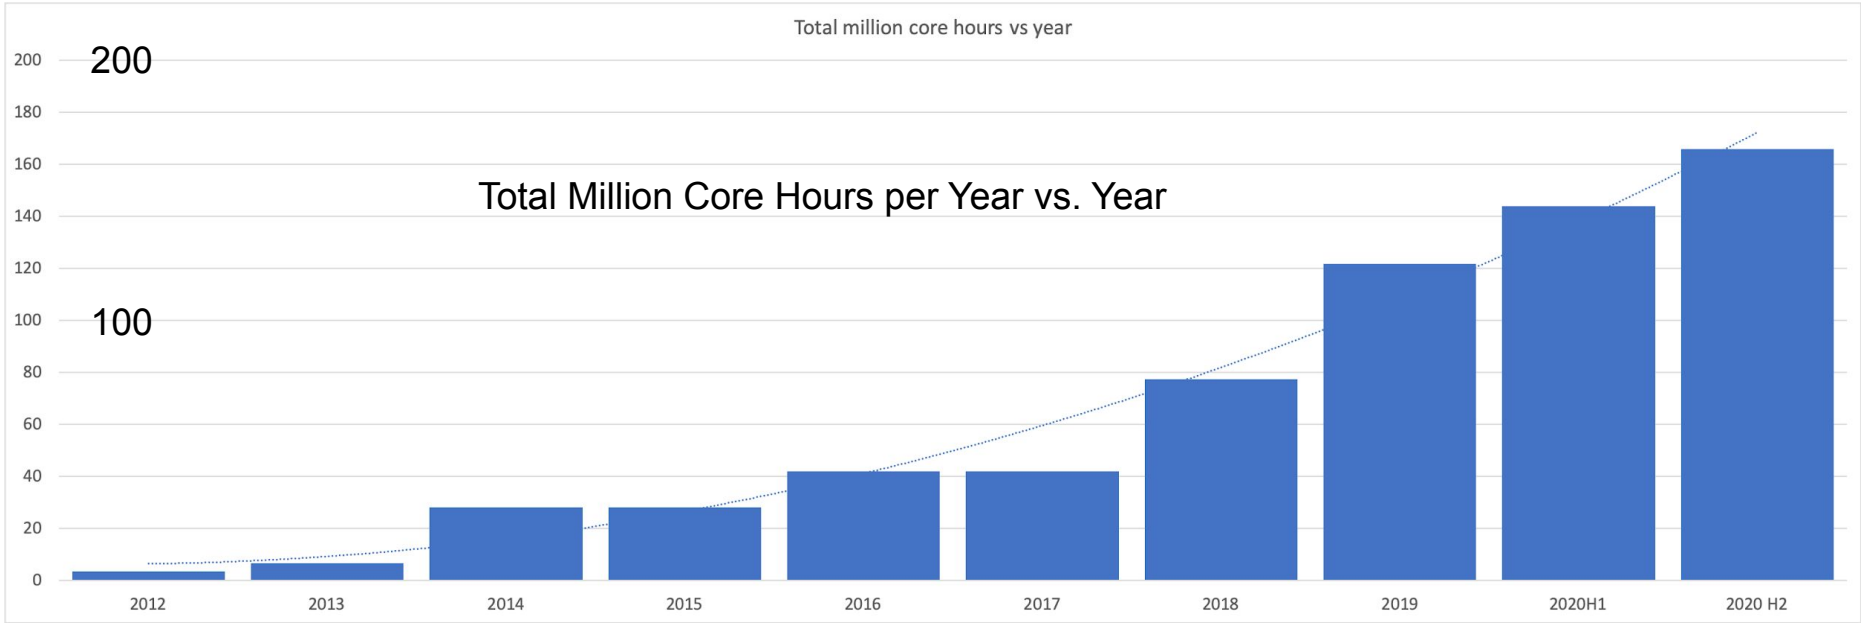

Total Million Core Hours per Year vs. Year

200

100

2012 2013 2014 2015 2016 2017 2018 2019 2020H1 2020H2

New Farm Priority Scheme

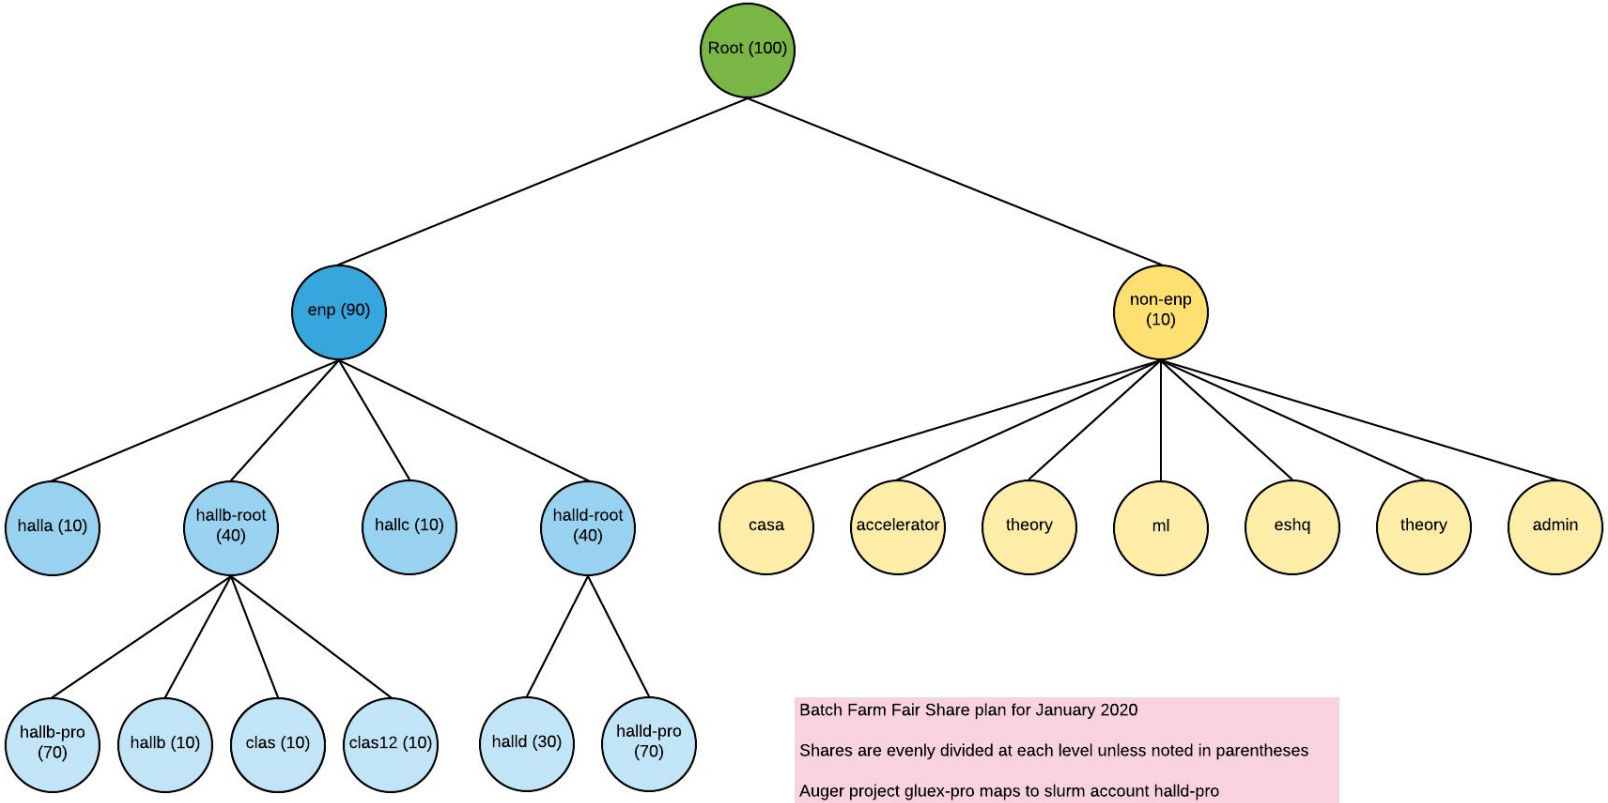

Batch Farm Fair Share plan for January 2020

Shares are evenly divided at each level unless noted in parentheses

Auger project gluex-pro maps to slurm account halld-pro
Auger project hallb-pro maps to slurm account hallb-pro
Auger user accounts see no changes

Summary and Outlook

- Summary
  - MCwrapper
    - User base continues to grow
    - System continues to adapt to changing conditions
  - HDGeant4
    - Dedicated meetings continue
    - Mysteries arise and are worked on
  - JLab scientific resources showing healthy growth
- Outlook
  - Support package modernization needed: Python, ROOT, Geant4,...
  - Standard interface for off-site computing
  - Software quality assurance procedures need upgrade