Sandia National Laboratories

Exceptional service in the national interest

KEEPING TRILINOS RUNNING PERFORMANTLY EVERY NIGHT, EVERYWHERE

Chris Siefert & Tpetra/Performance Team

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administrationunder contract DE-NA0003525.

SAND2024-14236PE

PROJECT GOALS

- Team: Chris Siefert, Jonathan Hu, Tim Fuller, Luc Berger-Vergiat.
- Extended Team: Brian Kelley, Carl Pearson, Curt Ober, Steve Kennon.

- Other Collaborators: James Elliott (APT), Sam Browne (Framework), Christian Glusa (Trilinos Leadership).

- Motivation: Taking ownership of Trilinos performance
  - Avoid (negative) surprises by the app teams, saving them time.
  - Reduce time for initial setup for apps / developers on new-to-them systems / updated libs.
  - Work with APT to stay at the most recent "good" version of libraries.
  - Understand behavior across time and across different systems.
  - Provide a way for app teams to know if it is safe to update Trilinos.

Note: We run performance tests, not the full test suite, so we do not guarantee that *everything* works.

WATCHR[1] PERFORMANCE SYSTEM

- Consists of 36 tests[2] run nightly on relevant systems covering Tpetra, Epetra, Intrepid2, MueLu, FROSch (new) and proxies for EMPIRE, SPARC, SIERRA-SD and SIERRA-TF.

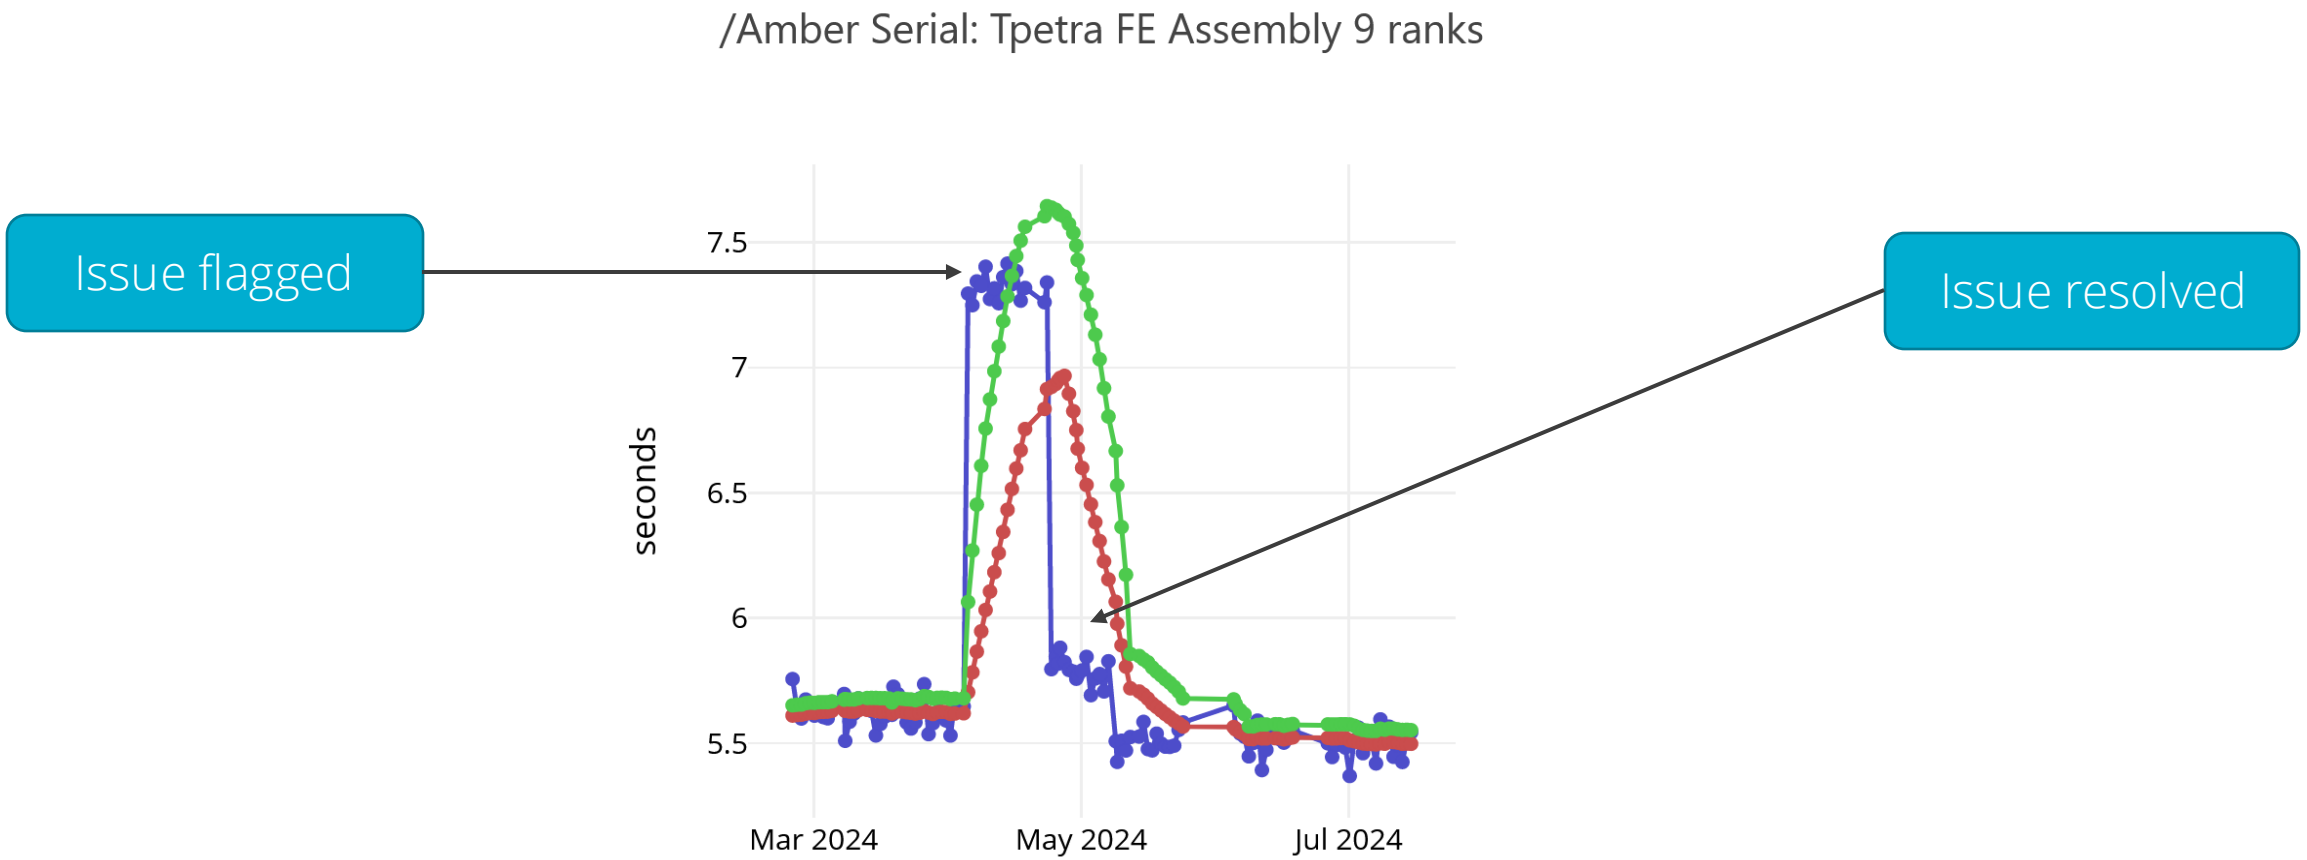

[1] https://github.com/sandialabs/watchr-core (developed by E. Ridgway)
[2] Not every test is run on each system --- Some are CPU or GPU only, others require certain TPLs, etc.

ZOOMING IN WITH WATCHR

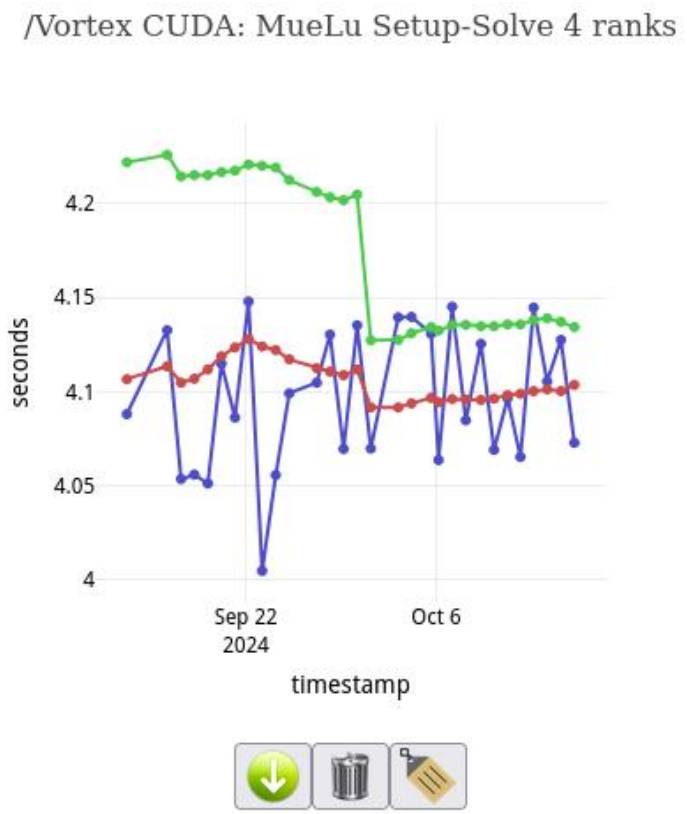

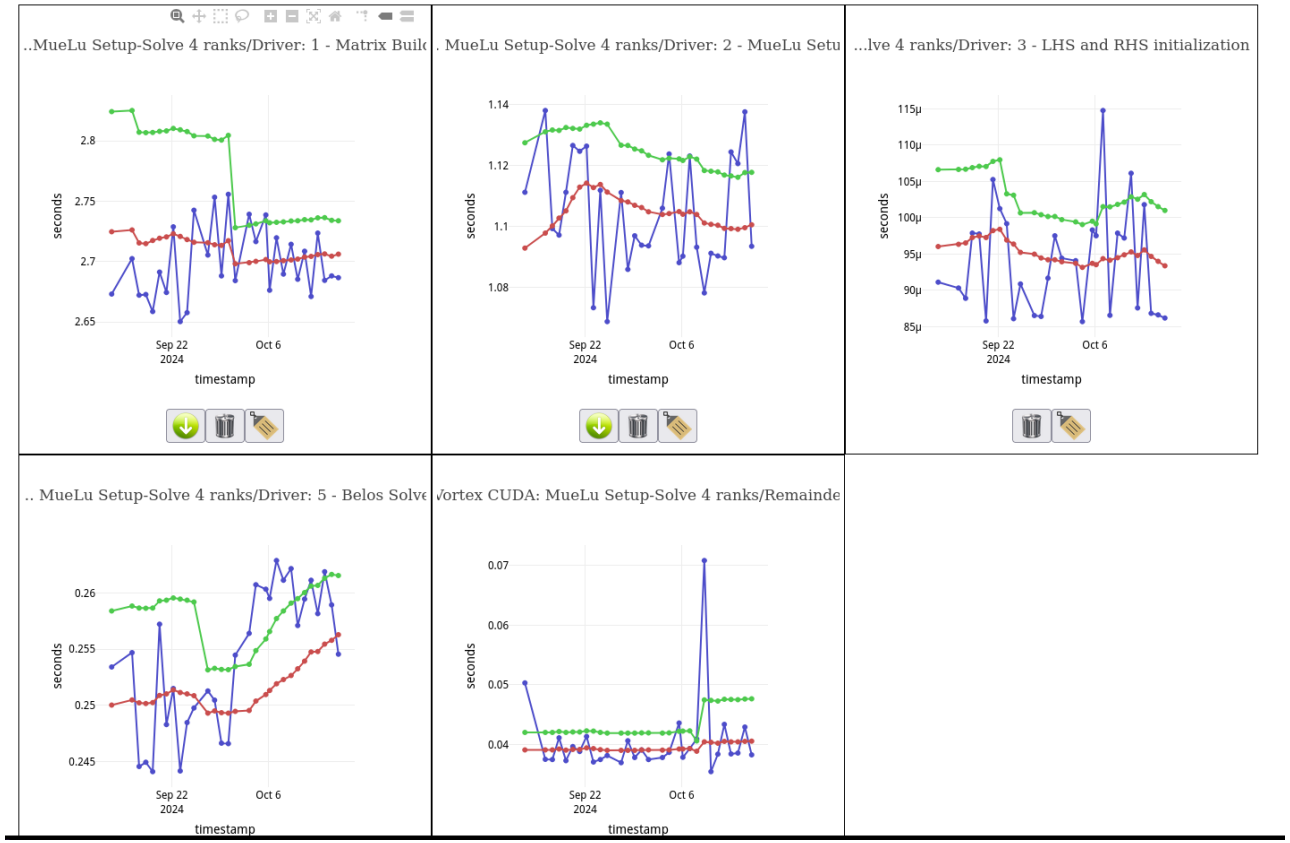

- We can zoom in to any Timer::StackedTimer in the proxy app.

SUCCESS STORY

- Kokkos 4.3 update switched the default behavior of operations in Kokkos::Serial.
  - They became atomic by default.
  - This makes Kokkos behavior correct if you use Kokkos::Serial within a threaded calculation.
  - Trilinos doesn't use Kokkos::Serial in this way, so performance degraded.

- After consultation with the Kokkos Team, we…
  - Changed the performance builds to use the -DKokkos_ENABLE_BYPASS_ATOMICS=ON
  - Emailed stakeholders recommending they do likewise.
  - Personal correspondence with stakeholders who had more questions.

Issue detected and resolved before apps even noticed.

Credit to Christian Trott for helping us diagnose the issue!

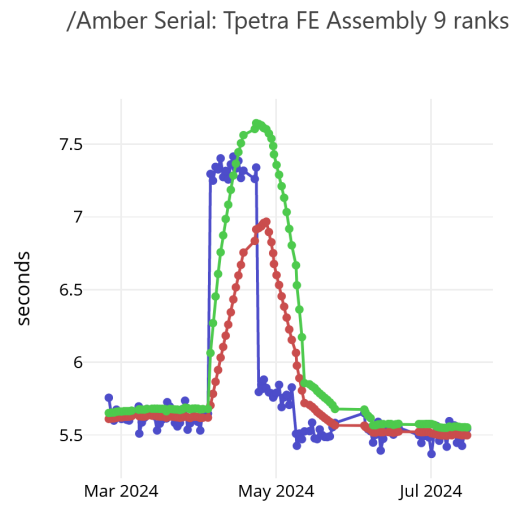

HPC CONFIGURATIONS REPO

- One-stop shop for users to get the nightly build script
  - https://gitlab-ex.sandia.gov/muelu/trilinos-hpc-configurations

- Worked with gitlab-ex admins to whitelist connections from LANL, LLNL, OLCF, NERSC & ALCF HPC resources.

- Sidesteps 2-factor issues.

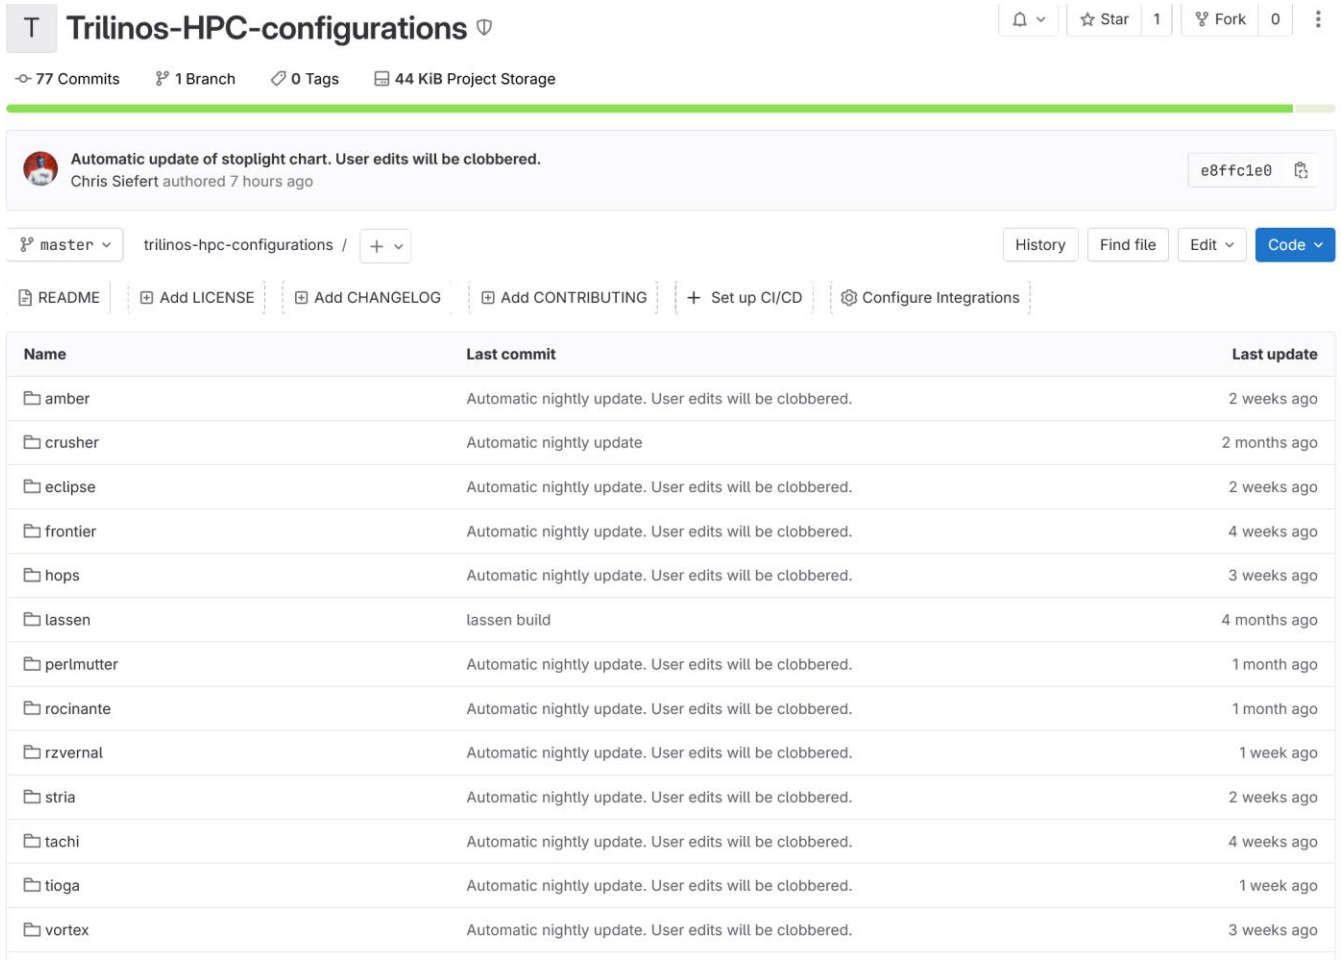

Credit to Christian Glusa for the idea!

WHAT DO WE MEAN BY “RUNS EVERYWHERE” ?

CPU
Amber (Sapphire Rapids/DDR)
Eclipse (Broadwell) [Also DevKokkos]
Rocinante (Sapphire Rapids/HBM)
Tachi (SPR/HBM) coming soon

NVIDIA GPU
Vortex (V100) UVM/NoUVM
Perlmutter (A100)
Hops (H100)
Venado (GH100) coming soon

AMD GPU
Frontier (MI250)
Tioga (MI250)
RZVernal (MI300) UM/NoUM
El Dorado (MI300) coming soon

Intel GPU
Sunspot (Intel Max)

DAILY STOPLIGHT CHART

- Did the tests run at all?
  - DST, Cron/Jenkins fails
- Did individual tests stop running?
  - Code crashes
- Has the performance degraded?
  - Still works, but worse.

Current builds managed by the Tpetra/Performance team

System	Days Run	Epetra	FROSch	Intrepid2	MiniEM	MueLu	SIERRA/SD	SIERRA/TF	SPARC	Tpetra
Amber	7/7	2.0 / 2 1.0	4.0 / 4 1.0	1.0 / 1 1.0	3.0 / 3 1.0	4.0 / 4 1.0	1.0 / 1 ---	7.9 / 8 1.0	---	13.0 / 13 1.0
Eclipse	7/7	2.0 / 2 1.0	4.0 / 4 1.0	1.0 / 1 1.0	3.0 / 3 1.0	4.0 / 4 1.0	1.0 / 1 ---	8.0 / 8 1.0	---	11.0 / 11 1.0
Rocinante	7/7	---	---	---	3.0 / 3 1.0	4.0 / 4 1.0	1.0 / 1 ---	---	---	13.0 / 13 1.0
Stria	7/7	2.0 / 2 1.0	---	1.0 / 1 1.0	3.0 / 3 1.0	4.0 / 4 1.0	1.0 / 1 ---	8.0 / 8 1.0	---	11.0 / 11 0.9
Vortex	4/7	---	---	---	3.0 / 3 1.0	4.0 / 4 1.0	1.0 / 1 ---	8.0 / 8 1.0	4.0 / 4 1.0	13.0 / 13 0.8
Frontier	7/7	---	---	---	1.0 / 1 1.0	4.0 / 4 1.0	1.0 / 1 ---	---	4.0 / 4 1.0	5.0 / 5 1.0
Perlmutter	7/7	---	---	---	---	3.0 / 3 1.0	1.0 / 1 ---	---	---	2.0 / 2 1.0
Hops	7/7	---	---	---	---	4.0 / 4 1.0	1.0 / 1 ---	---	---	5.0 / 5 1.0
RZVernal	7/7	---	---	---	---	1.9 / 2 1.0	1.0 / 1 ---	---	4.0 / 4 1.0	3.0 / 3 1.0
Tioga	7/7	---	---	---	1.0 / 1 1.0	4.0 / 4 1.0	1.0 / 1 ---	---	4.0 / 4 1.0	3.0 / 3 1.0

Credit to Michael Wolf for the idea! Last updated: 2024_07_09

CPU COMPARISONS

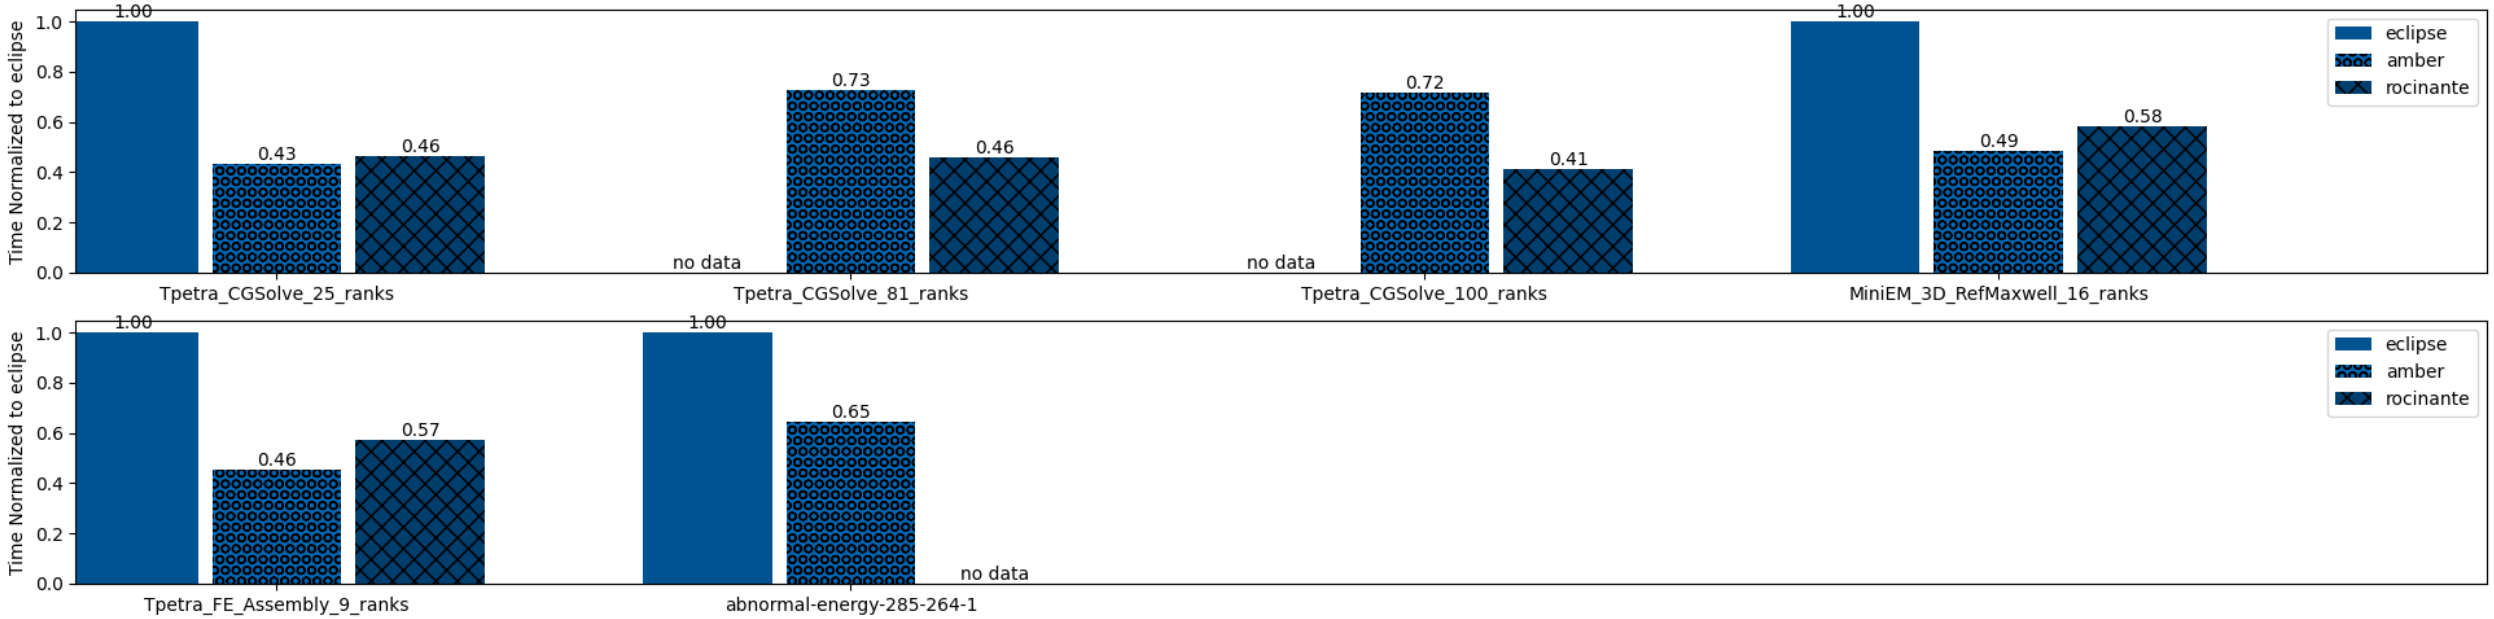

- If you're not using Amber for your CPU runs you should :)
- Once you start getting all the cores firing, you can see the advantage of the HBM on Rocinante.

GPU COMPARISONS

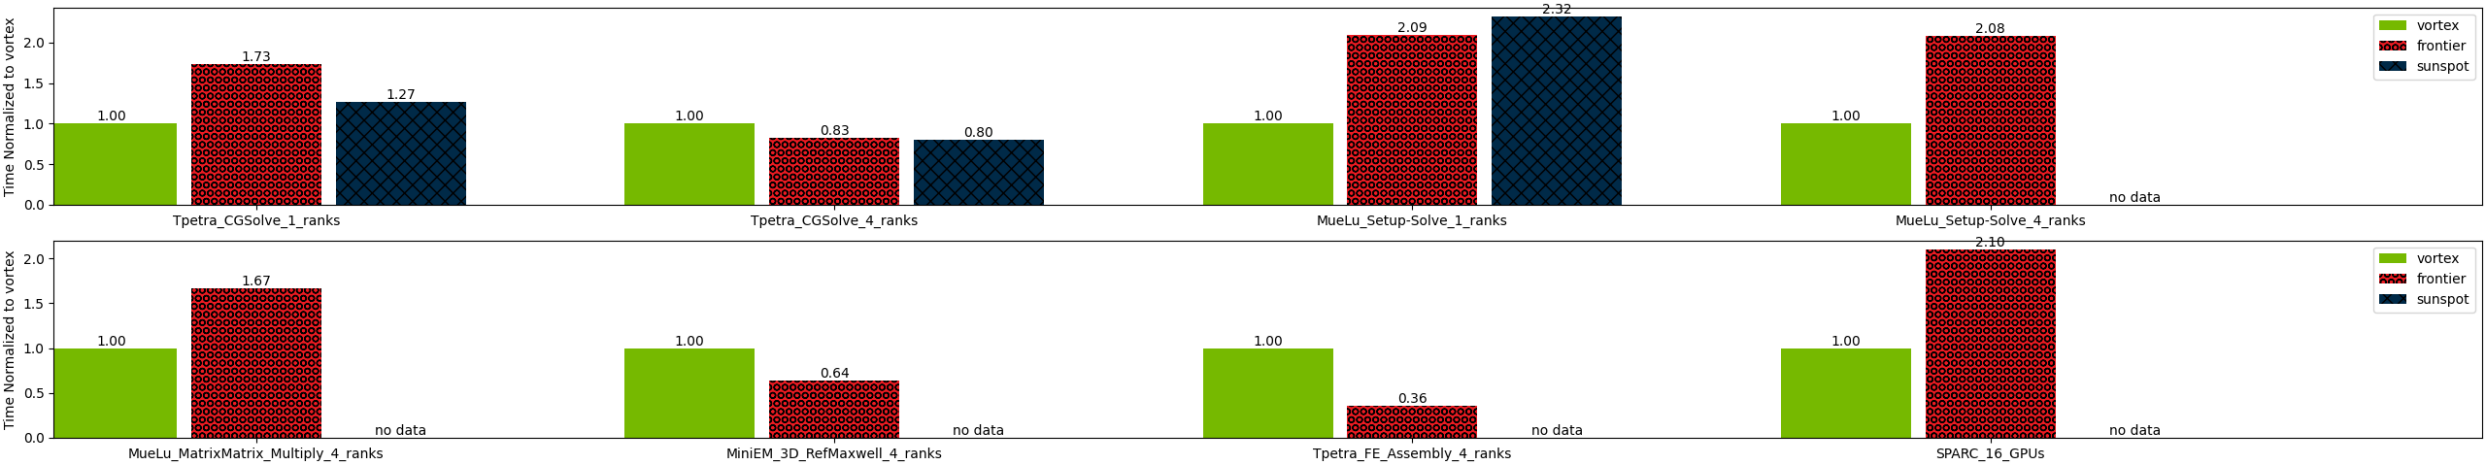

- MI300a comparisons not in this slide, but we have that data.
- Tests that aren't being run are obvious here.

UVM COMPARISONS

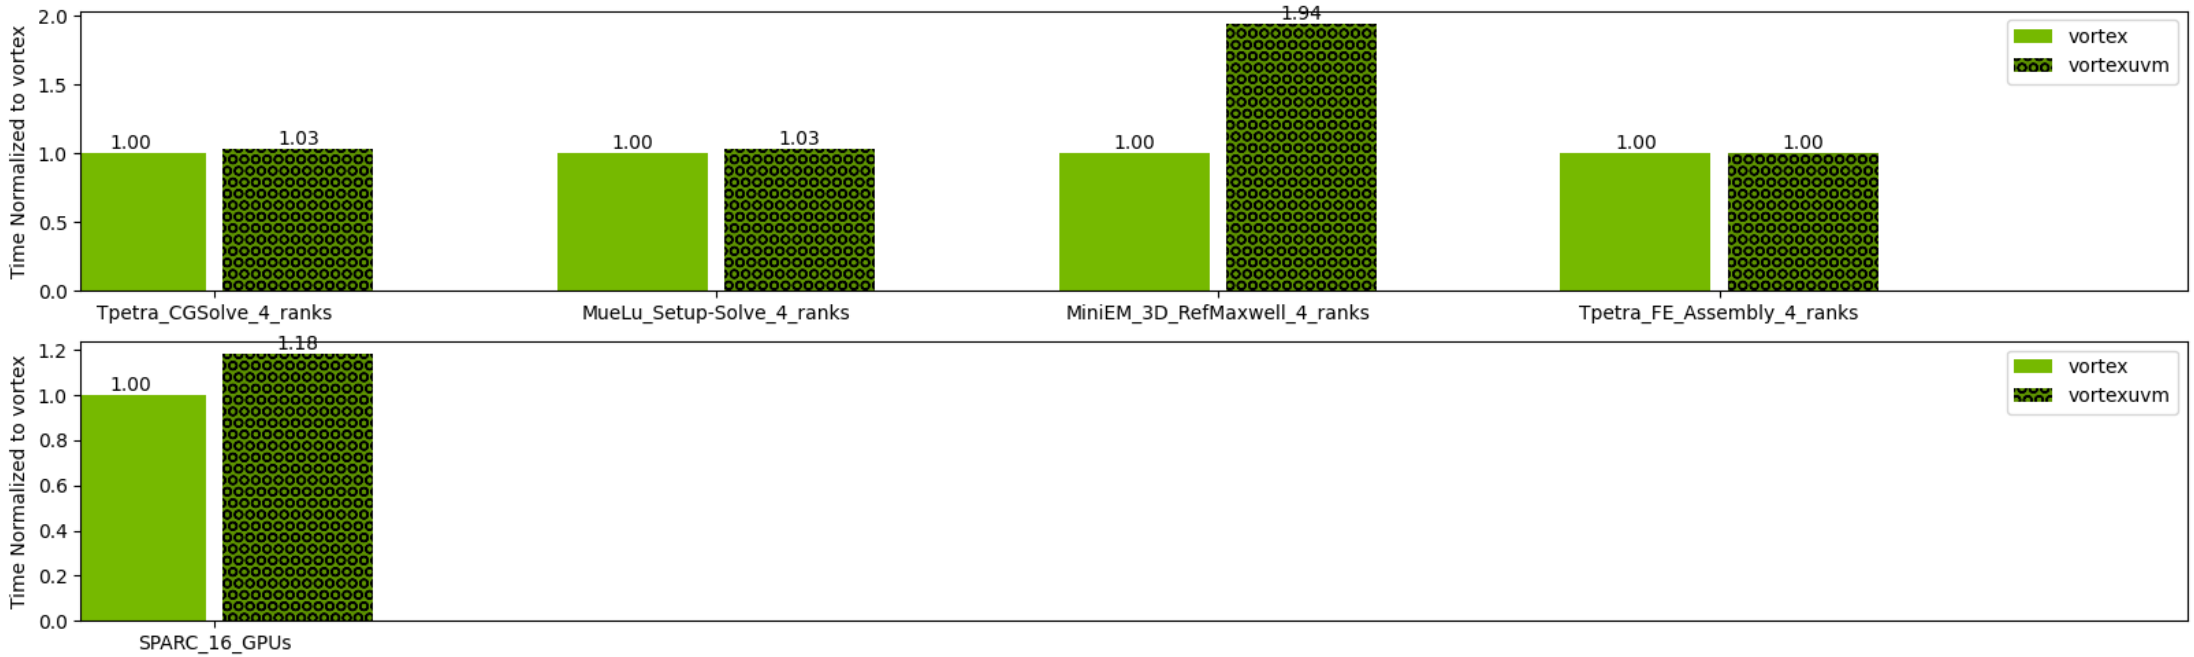

- If we have the data, we can make the comparison!

REMARKS & FUTURE DIRECTIONS

- Any Sandian can access the performance results.
  - https://gitlab-ex.sandia.gov/muelu/trilinos-hpc-configurations
  - Sierra/TF & SPARC team members already do this.
- If you need specific data extracted for your project, we can do that!
- The Tpetra/Performance team reviews the plots once a week.
- Coming soon: More performance testing w/ develop Kokkos / KokkosKernels.
- Question: Is *your* application on our dashboard? If not, why not?