



Exceptional service in the national interest

WARDEN: A NEW SOFTWARE FOR VISUALIZING NESTED TIMER DATA

Performance Monitoring for Trilinos

Presenter: Dane Camacho

NNSA-IMPACT Intern

HPSF Conference 2026, Chicago, IL



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

WHAT IS WARDEN?



- WARDEN is a way to visualize an array of plots based on nested XML data
- Each graph is a different test's time series data with the date on the x-axis and time (in seconds) on the y-axis
- Turns data that looks like this....

```
1 <?xml version="1.0"?>
2 <performance-report date="2025-09-11T08:57:00" name="nightly_run_2025_09_11" time-units="seconds">
3   <metadata key="Trilinos Version" value="cc5ec2b8bd"/>
4   <timing name="PERLMUTTER: Belos Block GMRES 1 ranks" value="0.990551">
5     <timing name="Galeri: Laplace 3D generation" value="0.0121108"/>
6     <timing name="Belos: Operation Op*x" value="0.00070374"/>
7     <timing name="Belos: BlockGmresSolMgr total solve time" value="0.386192">
8       <timing name="Belos: ICGS[2]: Orthogonalization" value="0.342478">
9         <timing name="Belos: ICGS[2]: Ortho (Norm)" value="0.00449125"/>
10        <timing name="Belos: ICGS[2]: Ortho (Inner Product)" value="0.1549"/>
11        <timing name="Belos: ICGS[2]: Ortho (Update)" value="0.176669"/>
12        <timing name="Remainder" value="0.00641672"/>
13      </timing>
14      <timing name="Belos: Operation Op*x" value="0.0385307"/>
15      <timing name="Remainder" value="0.00518355"/>
16    </timing>
17    <timing name="Remainder" value="0.591544"/>
18  </timing>
19 </performance-report>
```

WHAT IS WARDEN?


- WARDEN is a way to visualize an array of plots based on nested XML data
- Each graph is a different test's time series data with the date on the x-axis and time (in seconds) on the y-axis
- Into data that looks like this

WARDEN Performance Dashboard

Which Dataset?

FrontierData

How Many Days?



Which Test?

^top ^

Current Subtest: /



HOW DOES SANDIA USE IT?



- Trilinos performance team reviews the WARDEN dashboard at weekly meetings to examine testing for the last week and across architectures
- Team flags any changes in performance
 - Root-cause analysis (code, environment, something else?)
 - In case of regression, fix directly or inform responsible code teams
- We currently monitor
 - Trilinos develop branch
 - Trilinos with Kokkos VOTD

THE TEAM'S WORKFLOW PRIOR TO WARDEN



- Team reviewed graphs on previous software which spanned multiple pages and was time consuming.
- Manual git diff to see which change caused the performance
- Performance speedups and regressions were noted in a separate text file
- Easy to miss when a machine stopped reporting

WHAT DOES THE INPUT LOOK LIKE?



```
1 <?xml version="1.0"?>
2 <performance-report date="2025-09-11T08:57:00" name="nightly_run_2025_09_11" time-units="seconds">
3   <metadata key="Trilinos Version" value="cc5ec2b8bd"/>
4   <timing name="PERLMUTTER: Belos Block GMRES 1 ranks" value="0.990551">
5     <timing name="Galeri: Laplace 3D generation" value="0.0121108"/>
6     <timing name="Belos: Operation Op*x" value="0.00070374"/>
7     <timing name="Belos: BlockGmresSolMgr total solve time" value="0.386192">
8       <timing name="Belos: ICGS[2]: Orthogonalization" value="0.342478">
9         <timing name="Belos: ICGS[2]: Ortho (Norm)" value="0.00449125"/>
10        <timing name="Belos: ICGS[2]: Ortho (Inner Product)" value="0.1549"/>
11        <timing name="Belos: ICGS[2]: Ortho (Update)" value="0.176669"/>
12        <timing name="Remainder" value="0.00641672"/>
13      </timing>
14      <timing name="Belos: Operation Op*x" value="0.0385307"/>
15      <timing name="Remainder" value="0.00518355"/>
16    </timing>
17    <timing name="Remainder" value="0.591544"/>
18  </timing>
19 </performance-report>
```

OVERVIEW OF WARDEN'S FEATURES



- Easy swapping between machines
- Dive-in System
 - The user can click on graphs to see the graphs for subtimers within that timer
- System stopped reporting notice
- History and Performance modes
- "Go To Trillinos Commits" button
- Annotation System
 - The user can click on specific dates to make annotations across any machine

MACHINE SWAPPING

- Performance Data from each machine can be swapped between



WARDEN Performance Dashboard

Which Dataset?

How

Which

60 days

^top

Current Subtest: /

A screenshot of a dropdown menu for dataset selection. The menu is open, showing a list of dataset names. The second item, 'EIDoradoData', is highlighted with a light blue background. The list includes: 'PerlmutterData', 'EIDoradoData', 'EIDoradoDevKokkosData', 'EIDoradoUnifiedMemoryData', 'FrontierData', 'HopsDevKokkosData', and 'PerlmutterData'. To the right of the menu is a horizontal slider with a circular handle, labeled '60 days'. Below the menu is a button labeled '^top' and the text 'Current Subtest: /'.

PerlmutterData
EIDoradoData
EIDoradoDevKokkosData
EIDoradoUnifiedMemoryData
FrontierData
HopsDevKokkosData
PerlmutterData

DIVING-IN



- When diving into a graph, you can see the subtimers of that graph





Which Test?

days

[^top](#) [^](#)

Current Subtest: /ROCINANTE: Tpetra FE Assembly 1 ranks



DIVING-IN



- When diving into a graph, you can see the subtimers of that graph

[^top](#) [^](#)

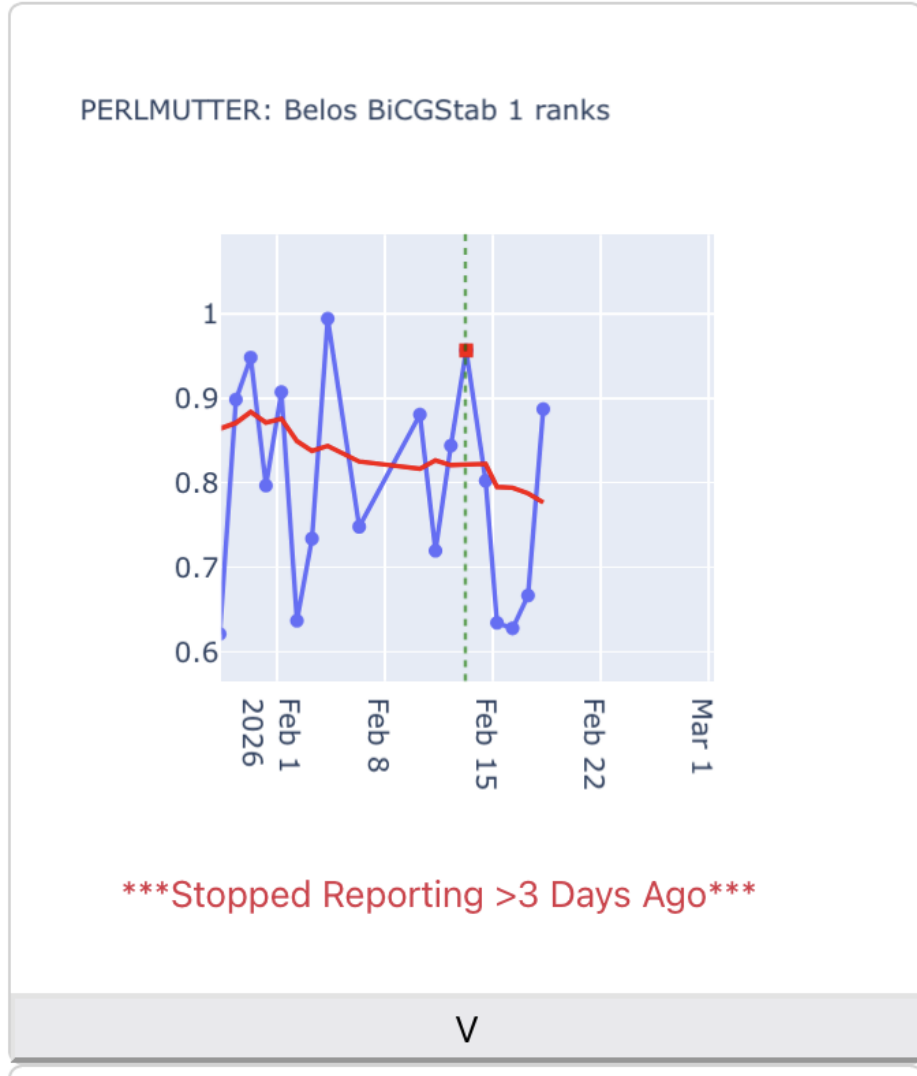
Current Subtest: /ROCINANTE: Tpetra FE Assembly 1 ranks[2] FillComplete (Graph)



STOPPED REPORTING NOTICE



- One feature the team wanted was the ability to quickly see whether a test has stopped running
- When a test stops receiving data for more than three days, there will be a notice on the graph



HISTORY AND PERFORMANCE MODES

- Performance mode – can only see the last 30 days, good for quick review
- History mode – can see all data points for all tests, takes longer to load

History Performance

WARDEN Performance Dashboard

Which Dataset?

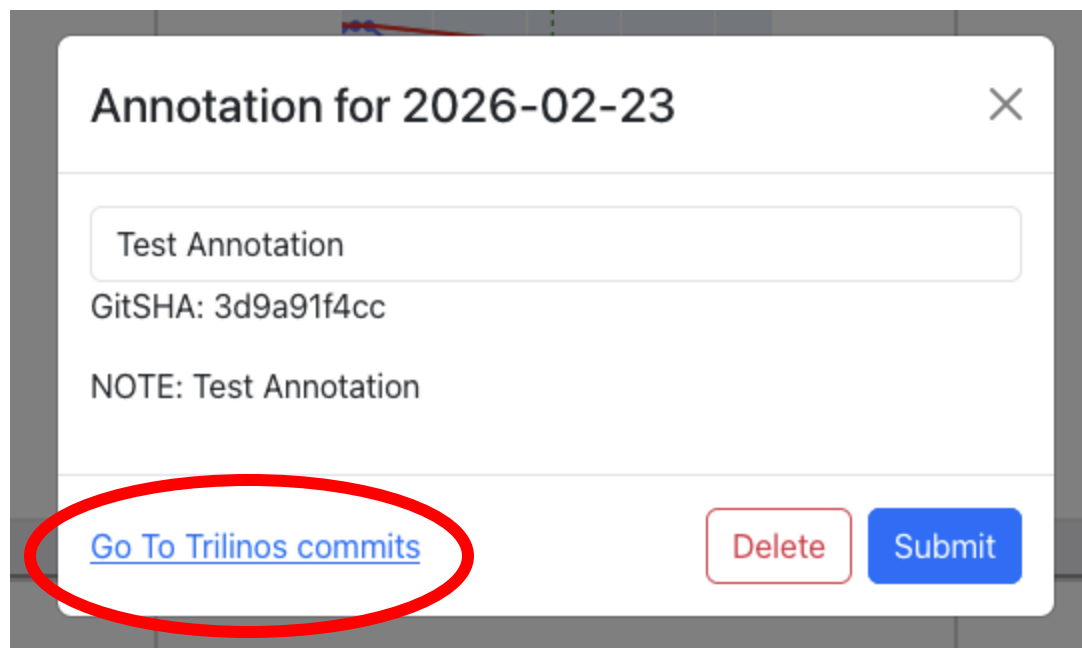
PerlmutterData





"GO TO TRILINOS COMMITS" BUTTON

- The user can press "Go To Trilinos commits" to see all commits from the previous day's GitSHA to the next day's on the Trilinos Github



"GO TO TRILINOS COMMITS" BUTTON



trilinos / Trilinos Public

Notifications Fork 609 Star 1.4k

Code Issues 274 Pull requests 69 Discussions Actions Projects Wiki Security Insights

Comparing changes

Choose two branches to see what's changed or to start a new pull request. If you need to, you can also [compare across forks](#) or [learn more about diff comparisons](#).

base: b29cfd9a2c ← compare: 3d9a91f4cc

Commits 122 Files changed 126 17 contributors

Commits on Jan 19, 2026

- build(deps): bump styfle/cancel-workflow-action from 0.12.1 to 0.13.0** ... Verified 8163cd3 <>
dependabot[bot] authored on Jan 19 ✖

Commits on Feb 2, 2026

- Tpetra makeColMap: Move to device** ... 3eeab7f <>
cgcgcg committed on Feb 2

Commits on Feb 3, 2026

- build(deps): bump step-security/harden-runner from 2.14.0 to 2.14.1** ... Verified 018e86b <>
dependabot[bot] authored on Feb 2 ✖
- Fix CodeQL Action download to use container certs** ... bf2a460 <>
achauphan committed on Feb 3

Commits on Feb 4, 2026

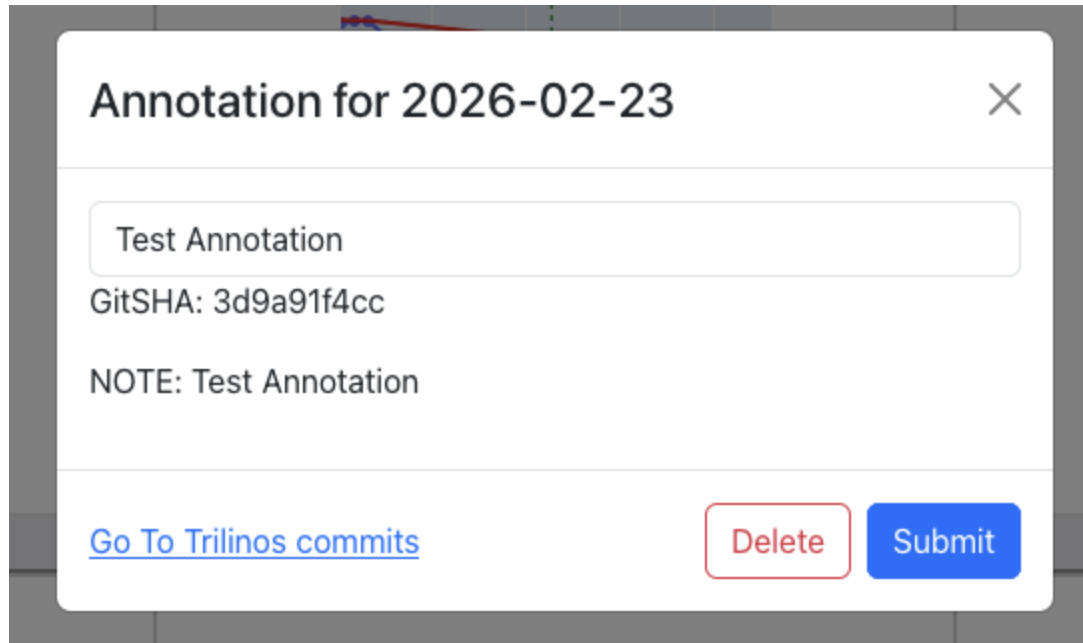
- Add googletest to CodeQL environment instance** ... de2d5da <>
achauphan committed on Feb 3 ✖
- Remove manually added comment for action version** ... c68d5c7 <>
achauphan committed on Feb 3 ✖

Commits on Feb 5, 2026

- Zoltan2 Tpetra adapters: Move constructors to device** ... deef550 <>
cgcgcg committed on Feb 5

ANNOTATION SYSTEM

- The user can click on a data point and add an annotation for that date on all tests for a machine
- This is useful for sudden spikes or dips in performance



Annotation for 2026-02-23

Test Annotation

GitSHA: 3d9a91f4cc

NOTE: Test Annotation

[Go To Trilinos commits](#) Delete Submit



FUTURE PLANS



- Generalize WARDEN to be used outside of Sandia
 - WARDEN is currently not open source, but is in the process of being open sourced
 - Any kind of nested XML data can be used with WARDEN
- Automated "stopped reporting" notifications
- Ability to add annotation across all machines at one time
- Changepoint Detection to see when performance variation becomes statistically significant
- Automated Changepoint Detection (sending an email when there is a significant change)
 - The team currently reviews the data once per week
- Auto-Annotation for software updates and other significant reasons for a dip or spike