

TotalView 2025.1 Release Notes

Updated: March 10, 2025



Table of Contents

1 Additions and Updates	4
AMD GPU Agent and Wave Thread Debugging.....	4
TotalView Parallel Configuration Definitions Updates	5
Multi-Dimensional Array Debugging Enhancements	5
Process and Group Hold.....	9
TotalView Documentation Changes.....	10
Bug fixes, security updates, and performance improvements	10
2 Platform Updates	11
3 Bug Fixes for TotalView 2025.1.....	12
4 Deprecation Notices.....	13
Solaris	13
TotalView Individual.....	13
5 Known Issues	14
Microsoft Windows Subsystem for Linux (WSL/WSL2)	14
Python Debugging	15
Licensing	15
macOS	16
Linux	18
Linux - Ubuntu	23
CUDA.....	23
ROCm / AMD GPU.....	25
Cray	25
AIX.....	31

These release notes contain a summary of new features and enhancements, late-breaking product issues, migration from earlier releases, and bug fixes.

Supported product versions and end-of-life information

To take advantage of new product features and improvements, and to ensure that your product is supported, install a current version. For a list of supported product versions, along with maintenance and end-of-life (EOL) information, see:
<https://portal.perforce.com/s/article/TotalView-Maintenance-Lifecycle>

PLEASE NOTE: The version of this document in the product distribution is a snapshot at the time the product distribution was created. Additional information may be added after that time because of issues found during distribution testing or after the product is released. To be sure you have the most up-to-date information, see the version of this document at:
<https://help.totalview.io/>.

1 Additions and Updates

AMD GPU Agent and Wave Thread Debugging

TotalView 2025.1 adds the ability to debug in two different modes while debugging ROCm code on AMD GPUs, agent-thread mode and wave-thread mode. Prior to this release TotalView only provided agent-thread mode debugging. Developers can toggle a process between the two modes during the TotalView debugging session. By default, TotalView uses agent-thread mode since it is known to perform better in the debugger because it is handling far fewer debugger threads per process. However, wave-thread mode enables more granular execution control at the AMD GPU wavefront level by leveraging TotalView's asynchronous thread control features.

Asynchronous thread control refers to a set of TotalView features that allow the debugger to individually control the execution of "threads" in a process. This granular thread execution control enables you to run and halt individual threads, single-step a group of threads in lockstep, hold and release the execution of individual threads, and create stop-thread and thread barrier breakpoints. When a ROCm process is in wave-thread mode, many of the asynchronous thread control features apply to the AMD GPU wave threads.

Benefits of using asynchronous wave thread debugging include the following:

- **Wave-Level Execution Control**
Increase execution-control precision by managing individual wavefronts. By controlling threads at the wave level, you can specify the execution of individual wavefronts for a finer level of control. You can set and clear stop-thread breakpoints, hold and release threads, and perform single-steps for specific waves, giving you more precise execution control than in agent-thread mode.
- **Enhanced GPU Breakpoint Functionality**
Set breakpoints at the wave level to stop only the triggering wave threads while allowing other threads in the process to continue running. Stop-thread and thread-specific conditional breakpoints allow all waves to reach the same location.
- **Improved Data Visibility**
Examining variables and data across wave threads within a process allows you to view variable values across the waves in a process.

For more information about using the new AMD GPU wave-thread debugging mode, see the AMD ROCm Wave Thread Debugging section in the TotalView Help at Using TotalView > GPU Debugging > Debugging AMD ROCm Programs.

TotalView Parallel Configuration Definitions Updates

When launching a parallel job from the TotalView user interface, developers select a parallel configuration from the Debug a Parallel Program session dialog. For the TotalView 2025.1 release, the available types of parallel configurations have been updated. Below is a summary of the changes:

- Various Open MPI definitions for versions up to Open MPI are combined into Open MPI. A separate Open MPI 5 parallel configuration has been created to optimize how TotalView acquires processes launched under Open MPI 5.
- Parallel configurations for different versions of MPICH, SLURM, MVAPICH and others have been cleaned up.
- Obsolete and non-supported parallel configuration definitions have been removed.

See the Loading Programs from the Session Editor section in the TotalView Help for more information about launching parallel programs from TotalView's user interface.

For information on customizing existing or adding your own parallel configuration definitions, see the MPI Startup Customizations section in the TotalView Help.

Multi-Dimensional Array Debugging Enhancements

The ability to debug multi-dimensional arrays has been enhanced in TotalView 2025.1 with the addition of the new Data Table panel in the Array View.

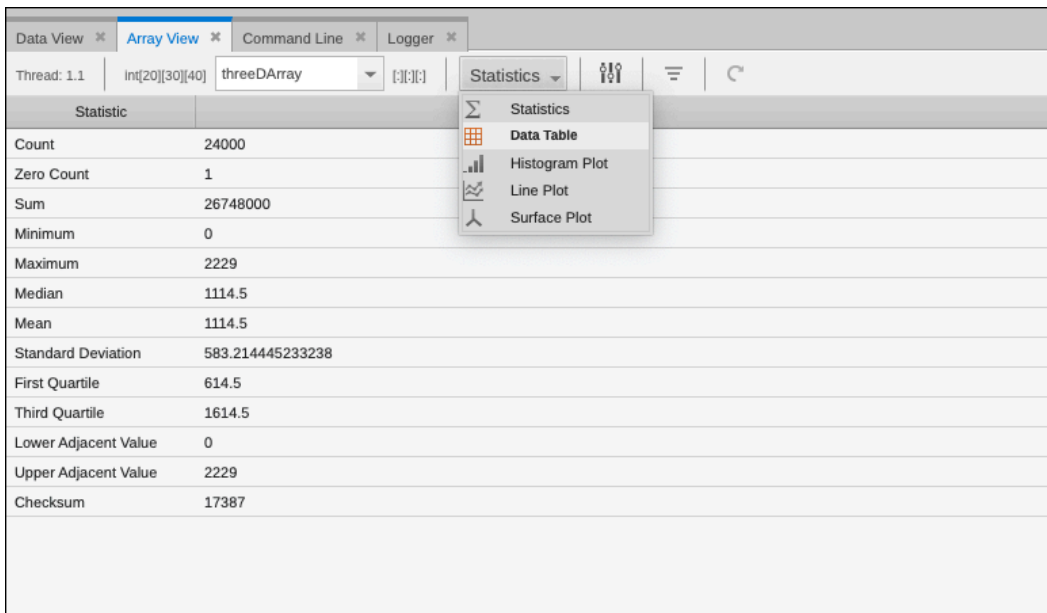


Figure 1 - New Array View Data Table Option

Once selected, the Data Table panel displays two dimensions of array data for the array expression added to the Array View.

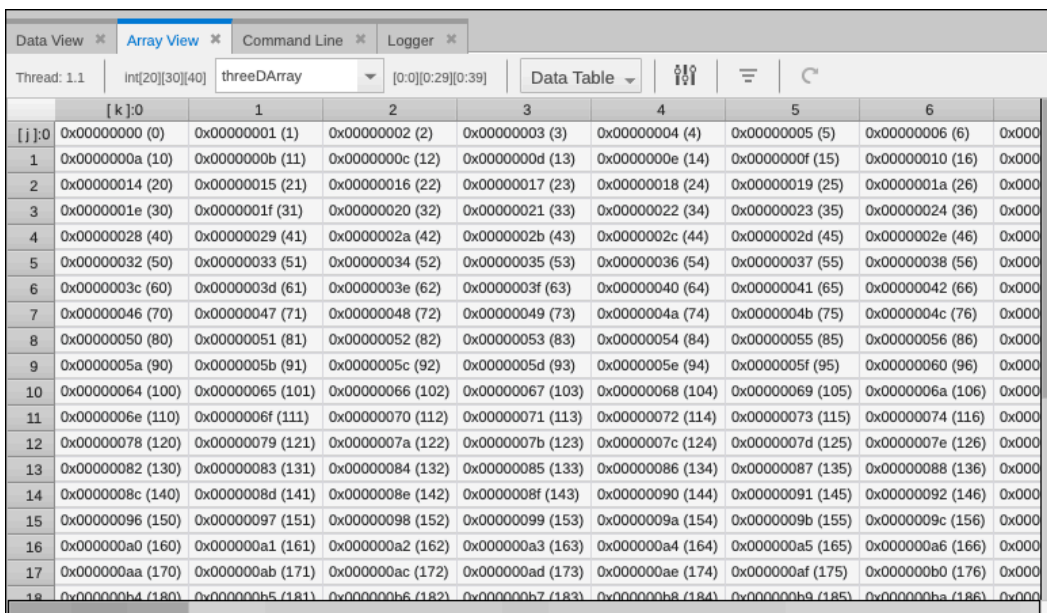


Figure 2 - Data Table Display of 3D Array Data

Which two dimensions of array data that are displayed is configured through the Array View Configuration Options dialog, selectable from the options icon in the Array View toolbar.



Figure 3 - Array View Configuration Options

In figure 3, dimensions *j* and *k* of a three-dimensional array are selected. Use the Slice area of the dialog and the drop-down options in the Dimension column to define which slice of the array to display in the Data Table. Slices begin with index *i*. The following screenshot shows data after selecting slices *i* and *k* to be displayed and the Start and End dimensions are adjusted for the dimensions of the array.

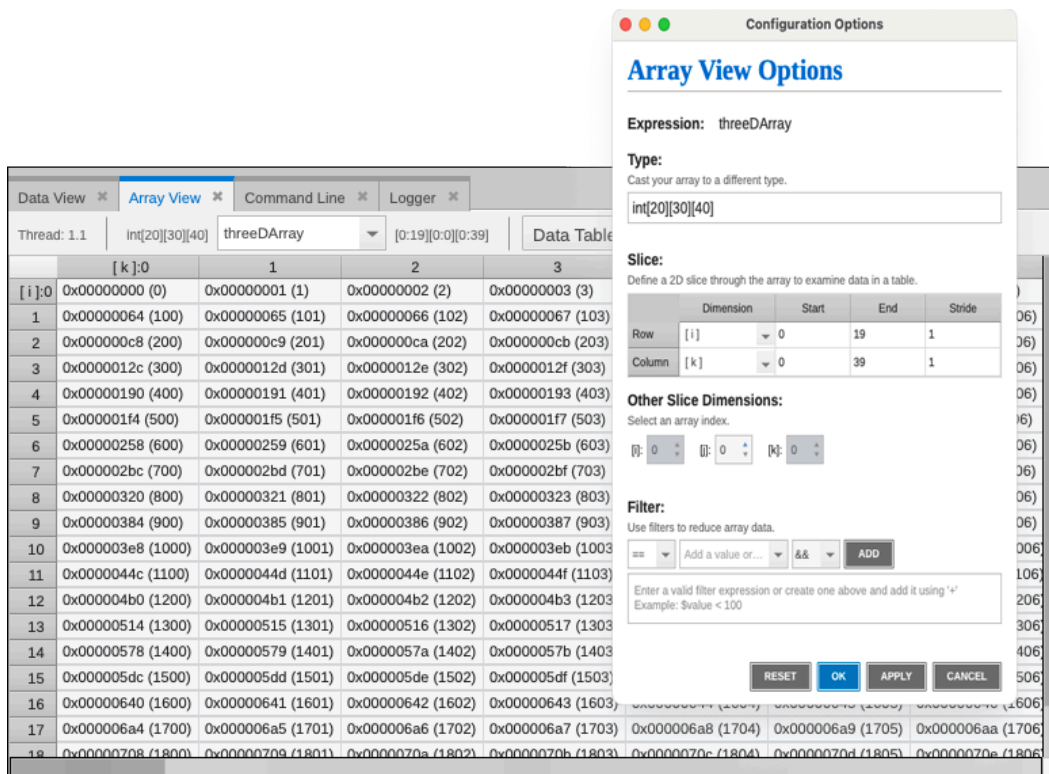


Figure 4 - Display Slices i and k

You can refine the displayed array data by adjusting the slice’s start and end indexes and setting a stride for the elements. Use the *Other Slice Dimensions* controls on the dialog to instruct TotalView which data indexes to traverse for the other slice dimensions not being displayed in the Data Table.

When the Array View display is changed to another view panel, such as the Line graph, the Configuration Options dialog provides a checkbox option to “Use Data Table Slice”. This will bring in the slice defined for the Data Table display and use the array data for the selected display panel. The follow diagram shows this option selected for the Line panel display.

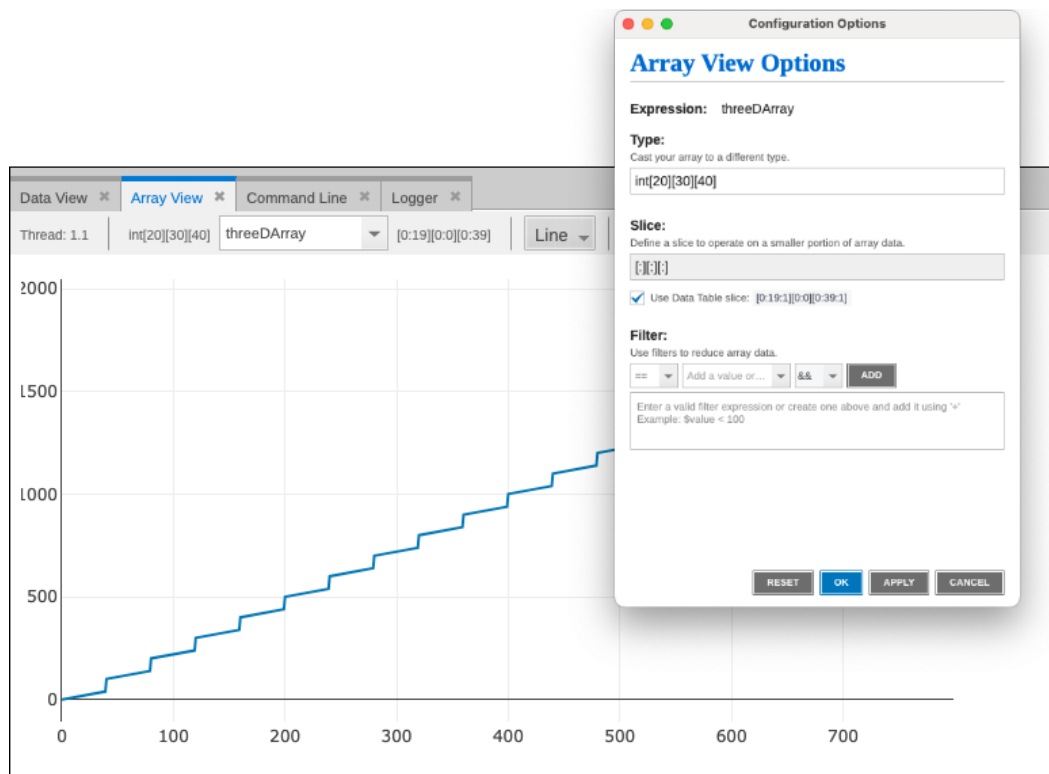


Figure 5 - Multi-Dimensional Array Data in Line Panel

Process and Group Hold

TotalView 2025.1 adds the ability to hold a process and the threads within it so they are not run when a TotalView command such as Group > Go is issued. This is useful when trying to synchronize the execution of a multi-process application. The following new menu items have been added to the Process menu:

- Process > Hold – Toggle menu item that tells TotalView to hold or release the current process. A held process will ignore all commands that would cause a thread to run.
- Process > Hold Threads – Tells TotalView to hold all threads in the current process. Compared to Process > Hold, when you hold a process none of the threads will ever run; that is, a go has no effect. Whereas Process > Hold Threads will hold user threads and maintain their state, but manager threads will be run when a process is run with a Go command.
- Process > Release Threads – Releases all threads that were held in the process.

Similarly, the ability to hold and release processes within a Control Group can be controlled with new menu items added to the Group menu.

- Group > Hold – Hold all the processes that are part of the same Control Group associated with the current thread in focus. Any commands to run processes within the group will be ignored.
- Group > Release – Release all processes that were held in the Control Group.

The Group > Hold/Release menu items are a convenience. If some Processes were held and need to be released, it is more convenient to use the Group/Release rather than visiting each held process to release it from its held state. If a Control Group contains many processes and you want to execute only a small number of those processes, it is quicker to do Group/Hold and then release only those processes that you want to execute.

TotalView Documentation Changes

As part of the TotalView 2025.1 release, the system used to generate TotalView's documentation was replaced. The new HTML documentation produced has improved formatting, navigation, and search capabilities. HTML documentation is distributed with TotalView and easily accessible through in product help through the Help menu. It is also available online at <https://help.totalview.io>. PDF documentation will no longer be produced or distributed with the product. Documentation for TotalView's legacy user interface will not be converted to the new documentation system and will be minimally maintained. HTML and PDF versions of the Classic UI documentation will still be shipped with the product and available online.

Bug fixes, security updates, and performance improvements

Numerous bug fixes, security fixes, and performance improvements have been addressed, including update of open-source third-party software packages used by TotalView.

2 Platform Updates

TotalView 2025.1 adds support for the following platforms:

Operation Systems:

- Fedora 40 and 41

Compilers:

- Intel oneAPI 2025

CUDA:

- CUDA 12.6

ROCm:

- ROCm 6.3

3 Bug Fixes for TotalView 2025.1

- TVT-42033 Watchpoint isn't halting execution in MPI program when variable changes
- TVT-42347 Stack backtrace of one thread can take over a minute when debugging GPU code
- TVT-42395 MacOS Sequoia - Private Wi-Fi address feature changes hostid periodically causing users TV licenses to fail
- TVT-42462 TotalView CLI fails to prompt yes/no question under parallel launch
- TVT-42463 TV library tracking fails if statically-linked binary exec's dynamically-linked binary
- TVT-42542 Runtime compiled GPU code that includes the DWARF embedded-source extension causes crash

4 Deprecation Notices

Solaris

We dropped support for Solaris starting with the first TotalView release in 2025. The final 2024 release with Solaris support will continue to be made available for download.

TotalView Individual

TotalView Individual will no longer be eligible for product renewals starting in 2025. Existing perpetual licenses will continue to operate but will not support newer versions of TotalView. Customers wishing to utilize newer versions of TotalView will need to move to the TotalView Developer for HPC or regular TotalView for HPC license. Contact your TotalView sales representative for more information.

5 Known Issues

High DPI Display Known Limitations and Recommendations

Based on our testing, we identified that DPI settings between 100-150, particularly scenarios where the scale and layout is set to 125% on Windows, do not always trigger the application's auto-scaling features as expected. This can result in display issues, such as some widgets appearing off-screen or UI elements not scaling correctly.

We are actively working on a solution to this specific scaling range and expect to address it in an upcoming release. In the meantime, we recommend the following workarounds for users experiencing display at these settings:

Adjust Scale and Layout

If possible, set your display's scale and layout factor to 100%. This adjustment has shown to mitigate the scaling issues effectively.

Additional Configuration Options

For users encountering widgets off-screen or requiring further adjustments to the application's scale factor, manually setting the **QT_SCALE_FACTOR** environment variable can offer a temporary solution. Adjusting this value between 0.8 and 0.95 has been beneficial in our testing.

For more detailed information on managing high DPI settings within a Qt application, visit <https://doc.qt.io/qt-5/highdpi.html>.

Microsoft Windows Subsystem for Linux (WSL/WSL2)

TotalView does not yet support debugging in a Microsoft WSL/WSL2 environment. Some preliminary testing has been done and there are networking issues that need to be resolved. The TotalView UI does work correctly in WSL and allows for easy remote debugging using the Remote UI feature. If you are interested in TotalView supporting WSL, please send a request to support-TotalView@perforce.com.

Python Debugging

Anaconda

TotalView supports debugging of the python interpreter in release 4 of Anaconda but this functionality does not work with the recent release of Anaconda 5. A change in the way it builds the python interpreter has broken the ability to debug python in TotalView for the python Anaconda 5 distribution.

Licensing

TotalView's FlexNet Embedded (FNE) version 2022.12 based license server implementation requires Java version 11 or higher

In TotalView 2024.1 the FNE based license server implementation was upgraded to FNE version 2022.12. The FNE 2022.12 License Server Administration Guide lists Java versions 8 and 11 as prerequisites. TotalView's testing has found however that OpenJDK 8 did not work. "Unrecognized option: --add-opens=java.base/java.lang=ALL-UNNAMED" appeared in the console output, and the license server could not complete initialization. Upgrading to OpenJDK version 11 eliminated the error and allowed the FNE server to initialize. Accordingly, TotalView's FNE Configure_License script now checks for Java 11+.

Model definition files cannot be loaded with FlexNet Embedded (FNE) version 2022.12

Model definitions were used in past FNE based license servers for TotalView to manage license access among users. An error is seen when attempting to load a model definition file in TotalView's 2022.12.0-0 FNE based license server. License access control with model definitions is therefore unsupported.

TotalView releases built with FlexNet Publisher 11.16.6 must have licenses served by a license server at 11.16.6 or higher

FlexNet Publisher client library version 11.16.6 is built into our recent releases. This means, according to FlexNet Publisher's component version compatibility rules (near the end of FNP's License Administration Guide PDF), the license server *must be* at v11.16.6 or higher. Although these rules have long been in place with no problems, we've recently been receiving reports of license checkout failures when using the license server v11.13.1 from previous TotalView releases. In this case the vendor daemon's debug log file shows "(toolworks) Request denied: Client (11.16) newer than Vendor Daemon (11.13). (Version of vendor daemon is too old. (-83,21049))". As noted in FNP's License Administration Guide

PDF, this issue can be avoided by making sure our latest license server components are in place.

TotalView sometimes cannot acquire license due to FlexNet bug

If you are using Linux Power or AIX then you are still affected by the bug in the FlexNet Publisher software that results in TotalView's inability to acquire a license when your license file contains multiple licenses with different maintenance expiration dates (i.e. the 4th field on the INCREMENT line). The licensing software skips some of the licenses in this case. If you know that your license file is being read and is correct, and you think you might be running into this bug, we recommend that you add the "sort" keyword and value (such as sort=1, sort=2, sort=3) to each INCREMENT line in the license file in any order. This bug has been reported to Flexera and is identified as SIOC-000145042.

Here is an example of adding the "sort" keyword:

```
SERVER linux-power 0050569b402c
VENDOR toolworks
INCREMENT TotalView_Enterprise toolworks 2014.1231 permanent 1 \
    sort=1 VENDOR_STRING="processors=16 platform=linux-power" \
    SIGN=3350A26C395A
INCREMENT TotalView_Enterprise toolworks 2015.1231 permanent 1 \
    sort=2 VENDOR_STRING="processors=16 platform=linux-ia64" \
    SIGN=C7D11FB667C8
INCREMENT TotalView_Enterprise toolworks 2016.1231 permanent 1 \
    sort=3 VENDOR_STRING="processors=16 platform=linux-x86_64" \
    SIGN=BEC7534A248A
```

FlexNet Embedded Licensing Technology

The Linux ARM64 and Linux PowerLE platforms require FlexNet Embedded for their licensing technology while the remaining TotalView platforms use FlexNet Publisher. In order to share tokens across all of the TotalView platforms we have also been porting them to support FlexNet Embedded as well. If you need to share tokens across Linux ARM64, Linux PowerLE and the other TotalView platforms, customers should contact license@perforce.com to set up a new license file for the FlexNet Embedded license server.

macOS

OpenGL Support

Beginning with TotalView 2023.2 TotalView's user interface on macOS is compiled for OpenGL support to enable drawing of 3D surface plots. At least version 2.1 of OpenGL is required. If an older version is detected OpenGL support will be disabled. Updating your graphics drivers may be required. If TotalView fails to start due to an OpenGL compatibility or driver issue support can be disabled by setting the environment variable QMLSCENE_DEVICE to softwarecontext:

```
env QMLSCENE_DEVICE=softwarecontext totalview
```

Trace Error when Debugging MPI Applications on Monterey

When debugging an MPI application on Monterey TotalView may fail with a fatal "trace_attach" error as it tries to attach to the MPI processes.

Big Sur Dynamic Linker Cache

MacOS Big Sur ships with a built-in dynamic linker cache for all system-provided libraries. This is a change from previous versions of macOS. As part of this change, copies of dynamic libraries are no longer present on the filesystem. As the result, TotalView debugger will fail to construct full backtraces that extend into the dynamic system libraries.

Physical console access needed when running TotalView on macOS High Sierra

Due to new security changes in macOS High Sierra and higher, TotalView will only run from the console and cannot be run through a remote desktop technology such as VNC.

With multiple displays attached to a macOS machine, some TotalView windows may not be noticeable

When a user attaches an external monitor to a running Macbook Pro, or adds multiple displays to a Mac, window rearrangement may move some windows off-screen. This may result in a TotalView modal window not being found until you use the Mac command to display all the windows (Mission Control). This appears to be an interaction between XQuartz and Darwin. It has been seen in Mavericks, but it's possible it will show up in other releases. There may be a workaround in System Preferences->Mission Control by disabling "Displays have separate Spaces."

Physical console access needed when starting TotalView

Starting in Mountain Lion, OS X security policies require that users meet a password challenge in order to use TotalView, and the challenge can be issued only to the console

(the OS X desktop). After the password challenge is met once, you can run TotalView repeatedly from the same login session without further challenges.

It is possible to work around this need for physical access with the following steps. The first few of these are likely already set in order to allow TotalView to run.

- Install XQuartz and TotalView
- Ensure every user needing debugging is in the `_developer` group
- Allow X11 forwarding in the `sshd_config` file (disabled by default)
- In a terminal window enter the following two commands:
 - `DevToolsSecurity -enable` (this step is optional if this was already enabled)
 - `sudo security authorizationdb write system.privilege.taskport allow`

Visualizer fails under macOS Sierra and Later using XQuartz

Attempts to use the visualizer tool fails with a message 'Error: attempt to add non-widget child "dsm" to parent "vismain"' which supports only widgets.

Windows

TotalView Remote Client

TotalView remote client on Windows supports the drawing of 3D surface plots. However, some users with specific GPU configurations might encounter issues due to system-level GPU blacklisting by the underlying Qt WebEngine (based on Chromium).

Users experiencing issues with the 3D Surface Plot are advised to set the environment variable `QTWEBENGINE_CHROMIUM_FLAGS` with the value --ignore-gpu-blacklist` before launching the application. This bypasses the GPU blacklist, potentially enhancing the feature's performance.`

Bypassing the GPU blacklist can lead to unstable behavior in rare cases, depending on your hardware and drivers. If you encounter any problems after applying this workaround, please revert the changes and contact our support team for further assistance.

Linux

OpenGL Support

Beginning with TotalView 2023.2 TotalView's user interface on Linux x86-64 and Linux ARM64 is compiled for OpenGL support to enable drawing of 3D surface plots. At least version 2.1 of OpenGL is required. If an older version is detected OpenGL support will be

disabled. Updating your graphics drivers may be required. If TotalView fails to start due to an OpenGL compatibility or driver issue support can be disabled by setting the environment variable `QMLSCENE_DEVICE` to `softwarecontext`:

```
env QMLSCENE_DEVICE=softwarecontext totalview
```

Reverse Debugging Crash on Ubuntu 22.04 / CentOS 9

Newer Linux distributions such as Ubuntu 22.04 and CentOS 9 changed restartable sequences structures and could cause a crash when reverse debugging. The `rseq` (restartable sequence) system call is not supported. This includes attaching to running programs that previously issued this system call. `glibc`, from version 2.35 onwards, attempts to register a restartable sequence with the kernel, preventing the reverse debugging engine from attaching to running processes using those versions of `glibc`. If you have any trouble using reverse debugging on these platforms the `glibc` behavior can be disabled by adding `glibc.pthread.rseq=0` to the `GLIBC_TUNABLES` environment variable of the target application. To avoid the crash simply set the environment variable to 0:

```
export GLIBC_TUNABLES=glibc.pthread.rseq=0
```

Intel 17 and 18 compilers not generating debug information for a declared integer

The Intel 17 and 18 version compilers are not generating proper debug information for declared integers in Fortran applications. As a result, TotalView is not able to properly evaluate the variable expression and display the variable values. This is a compiler, bug but a workaround is available by simply adding the `-debug extended` compilation flag option which adds symbols for local scalar variables and parameters.

Split-DWARF and `.gdb_index` support, and related options

State variable **`TV::dwarf_global_index`** is a boolean flag that controls whether or not TotalView considers using the DWARF global index sections (`.debug_pubname`, `.debug_pubtypes`, `.debug_typenames`, etc.) in executable and shared library image files. It defaults to **`true`**. It may be useful to set this flag to **`false`** if you have an image file that has incomplete global index sections, and you want to force TotalView to skim the DWARF instead, which may cause TotalView to slow down when indexing symbol tables. Command option **`-dwarf_global_index`** sets the flag to **`true`**, and **`-no_dwarf_global_index`** sets the flag to **`false`**.

State variable **TV::gdb_index** is a boolean flag that controls whether or not TotalView considers using the `.gdb_index` section in executable and shared library image files. It defaults to **true**. It may be useful to set this to **false** if you have an image file that has an incomplete `.gdb_index` section and you want to force TotalView to skim the DWARF instead. Command option **-gdb_index** sets the flag to **true**, and **-no_gdb_index** sets the flag to **false**.

ReplayEngine On-Demand Records Can Show Invalid Stack Trace

In some circumstances in which a ReplayEngine debugging session is driven to the beginning of recorded history, the debugger will display an invalid stack trace and stack frame. We have observed this when debugging a ReplayEngine recording file that was created during a live debugging session in which ReplayEngine was enabled on-demand.

To recover a valid stack, simply step or continue the session - that is, move forward in history. If the beginning of history is specifically of interest, it can be reached directly by opening a CLI window and issuing the command `"dhistory -go_time 1"`.

OpenMPI 1.8.4 with ReplayEngine enabled on older Linux releases.

We have observed a problem with the combination of Replay Engine, Open MPI 1.8.4, and older Linux releases such as RHEL5 (Red Hat Enterprise Linux). When the MPI runtime system closes shared libraries during its startup, a `munmap(2)` system call may attempt to unmap memory which is in use by Replay Engine. The error message "Unsupported memory access with syscall (11). Conflict with replay private internal memory." is displayed. This error is unrecoverable, but it may be possible to use Replay Engine on the same application by enabling it after the application has completed its `MPI_Init` call. We have not seen this problem with newer Linux releases such as RHEL6.

TotalView Message Queue and Intel MPI 5.0

By default, the TotalView Message Queue will not work with Intel MPI 5.0 without setting correctly `LD_LIBRARY_PATH` to the Intel MPI debug libraries. This can be done by sourcing one of the `"mpivars.sh/csh"` scripts provided by Intel with an added `"debug"` argument. For example, issue the command `"source PATH/impi/5.0.3.048/bin64/mpivars.sh debug"`, making sure to replace `PATH` with the path to your Intel MPI compiler installation. TotalView will then properly pick up the MPI message queue information and display it in its Message Queue window.

Memory Debugging and Intel MPI 5.0+

If a user wants to do memory debugging and they statically link their MPI program with the Intel MPI 5.0+ libraries, MemoryScape will detect a Double Allocation error. This is because, starting with Intel MPI 5.0, the MPI libraries redefine `free()` and MemoryScape depends on the system `free()` to see the deallocations. To work around this problem, the user will need to link dynamically or fall back to the Intel MPI 4.0+ libraries.

Debugging IBM Platform MPI Jobs in TotalView

Users have seen some issues when trying to use TotalView on an IBM Platform MPI job. If you try to launch the job from the Session Manager, or the Parallel Tab of the Startup Parameters window, TotalView may show an error that the target program has crashed while trying to load shared libraries. When run under `mpirun`, this error does not show since `mpirun` sets up the environment correctly. One can avoid the problem by setting the environment variable `LD_LIBRARY_PATH` to add the path to the library directory containing the missing libraries. The libraries should be in the 'lib' directory that is the same level as the 'bin' directory containing `mpirun`.

While testing the above issue it was noted that the classic launch method of

```
totalview mpirun -a -np 4 ./foo
```

did not appear to work correctly. TotalView would attach to all the processes, but only rank 0 was held at the point where the job went parallel. The other processes would run to a point where they were waiting on rank 0. To work around this, launch through the GUI as described above, or use the `-tv` switch for `mpirun`

```
mpirun -tv -np 4 ./foo
```

Newer Linux kernels that prohibit non-root access to /proc/self/pagemap and ReplayEngine

If non-root access to `/proc/self/pagemap` is prohibited, the ReplayEngine will emit an ignored assertion warning when the program being debugged enters record mode for the first time. Furthermore, any unknown syscalls will subsequently be handled a little more slowly.

The change affects Ubuntu 15.04 and likely other new distribution releases.

While using ReplayEngine, attaching to 32-bit application from 64-bit hosts sometimes fails

On some 64-bit hosts, attaching to a 32-bit target fails and results in a crash. The underlying technology behind ReplayEngine assumes that in a 64-bit environment, the target is also a 64-bit application and was not explicitly designed to support a mixed environment.

Benign Warning Messages Displayed when ReplayEngine is run on SuSE Linux Enterprise Server 11 with Service Pack 1. (SLES 11 SP1)

As of the 2016.06 release, ReplayEngine no longer fails when attempting to do replay mode operations (move the target backward into history) on Linux x86-64 platforms running SuSE Linux Enterprise Server 11 with Service Pack 1 but some warning messages such as the following are displayed:

```
127083 client/set_current_child.c:456:set_current_child_fn
[78347:78347]: Failed to restore process name for pid 78357: -5
127084 client/set_current_child.c:179:set_current_child_fn
[78347:78347]: Failed to set process name for pid 78357: -5
```

These messages are benign, and your reverse debugging session will work normally. We have only observed this problem on SLES 11 SP1. It is possible however, that other platforms running the same Linux kernel version (2.6.32.12-0.7) will also run into this problem. The only available workarounds are to update the OS to a more recent release (for example, installing Service Pack 2).

std::string shows as opaque value compiled with Clang 3.5

When trying to debug a program compiled with Clang 3.5 that uses a `std::string` variable, the variable is listed as type `std::string:64` with a value of

```
Opaque std::basic_string<char,std::char_traits<char>,std::allocator<char>>
```

When the same program is compiled with the Clang 3.3 compiler, the `std` string is seen as a simple STL container, and the string value is seen as `'abcd'`.

Clang supports several optimizations to reduce the size of debug information in the binary. These optimizations work based on the assumption that the debug type information can be spread over multiple compilation units. For instance, Clang does not emit type definitions for types that are not needed by a module and could be replaced with a forward declaration. Further, Clang only emits type info for a dynamic C++ class in the module that contains the vtable for the class.

It has been found that using the **-fstandalone-debug** option which turns off these optimizations works around the problem with the opaque value above.

The **-fstandalone-debug** option is useful when working with 3rd-party libraries that don't come with debug information.

Note that Clang never emits type information for types that are not referenced at all by the program.

-fstandalone-debug is the default on macOS.

-fno-standalone-debug is the default on Linux-x86-64. To work around the opaque value problem above, use the **-fstandalone-debug** option.

Linux - Ubuntu

Memory debugging by linking against the TotalView libraries may not work

There are several cases in which it is recommended to link the Heap Interposition Agent (HIA) into the target program to allow memory debugging without having to enable it in the GUI each time. Starting in Ubuntu 12, the linker does not link in libraries that are not directly used by the program. This means that the link line for the agent, with TVLIB pointing to the TotalView library,

```
gcc -g -o memprog memprog.c -L$TVLIB -ltvheap -Wl,-rpath,$TVLIB
```

may not pick up the HIA. Instead, add the option "-Wl,--no-as-needed" before the inclusion of the tvheap library. The new compile/link line will look like

```
gcc -g -o memprog memprog.c -Wl,--no-as-needed -L$TVLIB -ltvheap -Wl,-rpath,$TVLIB
```

CUDA

CUDA 11.8, 12.0, and 12.1 and later New Debug Implementation Issues

CUDA 11.8 introduced a new unified CUDA debugger backend, but several serious problems were identified while testing TotalView against the new implementation. The only viable workaround is to force CUDA to fall back to the legacy CUDA debug implementation by setting the CUDBG_USE_LEGACY_DEBUGGER environment variable to 1, e.g., `export CUDBG_USE_LEGACY_DEBUGGER=1`. The CUDA 12.0 r525 driver release should address the issues such that the new unified CUDA debugger backend works correctly with TotalView.

CUDA 12.0 *might* require setting `CUDBG_USE_LEGACY_DEBUGGER=1`, depending on the driver patch level. NVIDIA fixed the CUDA unified debugger backend problems after some CUDA 12.0 driver packages had *already* been released, so older versions of CUDA 12.0 drivers might have the same problems as CUDA 11.8. NVIDIA's advice for TotalView users who are

using the CUDA 12.0 driver is to upgrade to the latest version of the CUDA 12.0 driver, or upgrade to CUDA 12.1 driver or later.

CUDA 12.1 and later requires *not* setting CUDBG_USE_LEGACY_DEBUGGER. Starting with CUDA 12.1, NVIDIA does not support legacy CUDA debug implementation and setting the environment variable is known to cause debugger failures.

Note that support for Hopper requires using the unified debugger backend, so for systems with Hopper cards, do *not* set CUDBG_USE_LEGACY_DEBUGGER.

Note that the environment variable must be set (or unset) the same way in all the debugger processes (TotalView client and servers) and application processes. The easiest way to accomplish this, especially for MPI+CUDA jobs, is to set the environment variable in your shell start-up script (`~/cshrc`, `~/bashrc`, etc.). Below, are examples of setting CUDBG_USE_LEGACY_DEBUGGER only if CUDA Version 11.8 is installed on the node:

Example for csh:

```
test -f /bin/nvidia-smi && \  
  /bin/nvidia-smi | /usr/bin/grep 'CUDA Version:.*11\.8' >/dev/null && \  
  setenv CUDBG_USE_LEGACY_DEBUGGER 1
```

Example for bash:

```
test -f /bin/nvidia-smi && \  
  /bin/nvidia-smi | /usr/bin/grep 'CUDA Version:.*11\.8' >/dev/null && \  
  export CUDBG_USE_LEGACY_DEBUGGER=1
```

NVIDIA Pascal Unified Memory Debugger Internal Error

CUDA applications running on Pascal under the debugger may cause a debugger internal error (for example, a SEGV) when the application process exits. Known driver versions that have this problem are r361 and r375, however the problem may exist in other driver versions. The debugger internal error typically involves debugging a CUDA application that exits after using unified memory. NVIDIA has confirmed the driver error, and plans to release a fixed driver version. A release date is not yet available. If the problem occurs, TotalView will exit with an internal error.

Dynamic parallelism not fully supported

With CUDA, we have limited support for dynamic parallelism. We plan improvements to our functionality for displaying the relationships between dynamically launched kernels and navigating the various running kernels.

Layered textures not supported

TotalView does not yet support CUDA layered textures. If you try to examine a layered texture in the TotalView Data Pane, a “Bad address” message will be displayed and you will see “ERROR: Reading Texture memory not currently supported” displayed on the console. If you require layered textures support, please contact TotalView support at support-TotalView@perforce.com and let us know how you are using textures so we can develop the best solution to support you.

ROCm / AMD GPU

AMD GPU Core File Debugging

Currently debugging GPU core files generated from code running on AMD GPUs is not supported yet.

Reverse Debugging and AMD GPU Debugging

Currently reverse debugging during an AMD GPU debugging session is not supported due to issues with the reverse debugging engine dealing with GPU kernel launches.

Cray

Debugging your program within supercomputer environments can often be challenging. Reference the sections below to learn pointers on how to successfully run TotalView on Cray EX (Shasta) systems, enable memory debugging and perform reverse debugging on your program within a Cray environment.

Running TotalView on a Cray EX (Shasta) system

TotalView has been tested on several Cray EX systems with different configurations. The following are recommendations for running TotalView on several specific system configurations. If your system does not match these configurations, please contact the TotalView support team at support-TotalView@perforce.com to have them help guide you to the proper configuration to use. If you are using NVIDIA GPUs on your Cray EX system, please see the “Cray EX Configured with NVIDIA GPUs” below.

Cray EX Configured with: Cray Shasta / PALS

The Cray Shasta environment with the PALS workload manager contains several issues that have been worked around in TotalView, primarily by disabling the use of the Cray Common Tools Interface (CTI) library. When CTI is disabled, TotalView will use its traditional MRNet/SSH based tree instantiation instead of an MRNet/CTI tree.

More seriously is a critical problem with PALS where it has been observed that "mpiexec.pals" fails when trying to debug multiple processes per node. In this case, the "mpiexec.pals" program currently fails and prints the error message: "Error: RPC timeout after 15s, received 0 of 1 responses". This bug has been recorded by HPE as PE-37769 and preliminary investigation shows that there is a bug in the PMI layer. Only the first process on a node calls the barrier. PMI needs a change to make all processes on a node wait at the barrier.

The following are the recommended requirements for successfully launching TotalView on a Cray EX system configured with Cray Shasta / PALS. Example launch commands follow the requirements.

1. The Cray Common Tools Interface (CTI) library and Parallel Application Launching Service (PALS) packages that are installed on the system have several bugs that prevent the TotalView MRNet/CTI support from working. Perforce has modified TotalView to work around these bugs by avoiding the use of MRNet/CTI, and instead forcing the use of a traditional MRNet/SSH launch approach. Until the CTI bugs are resolved, set the TVD_DISABLE_CRAY=1 environment variable before starting TotalView.
2. TotalView's MRNet/SSH tree instantiation requires the following:
 - a. SSH from the front-end (login) node to the compute nodes, and from the compute nodes to other compute nodes. If SSH is not allow or not configured correctly, MRNet/SSH tree instantiation will fail.
 - b. Each user must configure SSH to allow password-less SSH. That is, it must be possible for the user to SSH from one node to another without a password. There are numerous articles on the Internet that describe how to do this (for example, <https://linuxize.com/post/how-to-setup-passwordless-ssh-login/>).
 - c. SSH must be configured to disable strict host key checking. TotalView is configured to do this automatically when run in the MRNet/SSH mode by setting the TVD_DISABLE_CRAY=1 environment variable. The "totalview"

startup script does this by passing the following configuration options to the TotalView executables:

```
-xplat_rsh ssh
-xplat_rsh_args '-x -o StrictHostKeyChecking=no'
```

3. TotalView requires the MPI starter program to implement the MPIR Process Acquisition Interface, which allows it to attach to a parallel program. On Cray Shasta/PALS system, the "mpiexec.pals" command must be used to launch jobs, because the "cray mpiexec" command does not support MPIR. The options for the "mpiexec.pals" program are similar to "cray mpiexec", so generally it is possible to use "mpiexec.pals" in place of "cray mpiexec". Use "mpiexec.pals --help" for more information on "mpiexec.pals".
4. The "--no-transfer" option should be provided to "mpiexec.pals", which disables transferring application binaries to the compute nodes. Using this option avoids two issues when debugging:
 - a. Caching the application binaries. The debugger will use the original copies of the application binaries rather than copying them from a compute node into the TotalView library cache, which saves time and space.
 - b. Complains from the debugger regarding the system deleting the transferred copies of the application binaries: "We can no longer read the executable '/run/palsd/...'".
5. The "mpiexec.pals" program has a bug where it will print the following message:

```
Usage: cray pals apps signal create [OPTIONS] APID
Try 'cray pals apps signal create --help' for help.
```

```
Error: Error received from server: 503 Service Unavailable
```

There appears to be no consequences to the error message. The command was trying to signal the ranks with a SIGCONT but, it is unnecessary since when running under the debugger it just stops the process again.

6. The "mpiexec.pals" program currently fails when trying to debug multiple processes per node. Use the "-ppn 1" option to restrict the job to use one process per node. When attempting to debug multiple processes per node, mpiexec,pals will print the error message: "**Error: RPC timeout after 15s, received 0 of 1 responses**".

Here is an example of some commands to allocate 4 nodes, run an MPI application with one process per node on 4 nodes, and run that application under TotalView:

```
# Set TVD_DISABLE_CRAY=1 to disable use of Cray CTI launch
export TVD_DISABLE_CRAY=1

# Test that a normal launch on these nodes works
mpiexec.pals -n 4 -ppn 1 ./mpi_hello_world

# Launch the job under TotalView, disable transferring binaries
# to compute nodes, and run one process per node
totalview -args mpiexec.pals --no-transfer --verbose \
  -n 4 -ppn 1 ./mpi_hello_world
```

Cray EX Configured with: Cray Shasta / SLURM

The Cray Shasta environment with the SLURM workload manager contains an issue that can be worked around in TotalView by disabling the use of the Cray Common Tools Interface (CTI) library.

The following are the recommended requirements for successfully launching TotalView on a Cray EX system configured with Cray Shasta / SLURM. Example launch commands follow the requirements.

1. The Cray Common Tools Interface (CTI) library package that is installed on the system has bugs that prevent the TotalView MRNet/CTI support from working. Perforce has modified TotalView to work around these bugs by avoiding the use of MRNet/CTI, and instead forcing the use of a traditional MRNet/SSH launch approach. Until the CTI bugs are resolved, set the `TVD_DISABLE_CRAY=1` environment variable before starting TotalView.
2. TotalView's MRNet/SSH tree instantiation requires the following:
 - a. SSH from the front-end (login) node to the compute nodes, and from the compute nodes to other compute nodes. If SSH is not allow or not configured correctly, MRNet/SSH tree instantiation will fail.
 - b. Each user must configure SSH to allow password-less SSH. That is, it must be possible for the user to SSH from one node to another without a password. There are numerous articles on the Internet that describe how to do this (for example, <https://linuxize.com/post/how-to-setup-passwordless-ssh-login/>).
 - c. SSH must be configured to disable strict host key checking. TotalView is configured to do this automatically when run in the MRNet/SSH mode by

setting the `TVD_DISABLE_CRAY=1` environment variable. It does this by setting the following internal configuration options through the TotalView CLI commands:

```
-xplat_rsh ssh
-xplat_rsh_args '-x -o StrictHostKeyChecking=no'
```

Here is an example of some commands to allocate 4 nodes, run an MPI application with one process per node on 4 nodes, and run that application under TotalView:

```
# Set TVD_DISABLE_CRAY=1 to disable use of Cray CTI launch
export TVD_DISABLE_CRAY=1

# Allocate a node for the job
salloc -C gpu -N 1 -G 4 -t 30 -q regular -A nvendor_g

# Test that a normal launch on these nodes works
srun -n 4 ./mpi_hello_world

# Launch the job under TotalView
totalview -disable_cray -mrnet_super_bushy \
  -args srun -n 4 ./mpi_hello_world

# If running a job that uses CUDA, launch the job under |
# TotalView without CUDA.
totalview -disable_cray -no_cuda \
  -args srun -n 4 ./mpi_hello_world
```

[Cray EX Configured with NVIDIA GPUs](#)

We have not fully tested debugging NVIDIA GPUs on Cray EX systems, therefore it is not yet supported and may not work. We have seen application hangs with one of our CUDA test programs, so we have reason to believe that there is a bug in TotalView, the CUDA debug API, or the NVIDIA kernel driver on these systems. Further investigation is needed. We'd appreciate feedback on your experiences debugging GPU code on Cray EX systems. Note the following:

- a. Due to limitations in the CUDA debug API, it is not possible for a single debugger process to debug more than one CUDA process per node. It is possible to work around this limitation in MPI applications by using the TotalView `"-mrnet_super_bushy"` option. This option will force TotalView to create an MRNet tree that consists of one debugger server process per MPI process.

- b. To disable CUDA debugging in TotalView, use the “-no_cuda” option. This option might avoid CUDA-related application or debugger hangs; however you will not be able to debug code running on the GPUs.

Memory debugging on Cray systems

Use the pointers below to help achieve a successful memory debugging session within the Cray environment:

- **Install TotalView on a shared file system visible to the Cray compute nodes.**
In order for the required memory debugging shared libraries to be located when your program is running on a compute node it is best to install TotalView on a shared file system.
In order for the required memory debugging shared libraries to be located when your program is running on a compute node it is best to install TotalView on a shared file system.
- **Statically linking your program against the tvheap_cnl library.**
One of the most foolproof ways of enabling memory debugging is to statically link against the tvheap_cnl library that is shipped with TotalView. The static library fully supports multithreaded applications. Statically linking your application with the tvheap_cnl library will automatically enable memory debugging in your program when debugged under TotalView and MemoryScape. See the “Linking Your Application with the Agent” discussion in the User Guide for more information on how to statically link your applications with the library.
One of the most foolproof ways of enabling memory debugging is to statically link against the tvheap_cnl library that is shipped with TotalView. The static library fully supports multithreaded applications. Statically linking your application with the tvheap_cnl library will automatically enable memory debugging in your program when debugged under TotalView and MemoryScape. See the “Linking Your Application with the Agent” discussion in the User Guide for more information on how to statically link your applications with the library.
- **Dynamically linking your program against the tvheap_cnl library.**
It is possible to dynamically link your application against the dynamic version of the tvheap_cnl library. In this scenario the tvheap_cnl library must be visible to the Cray Compute Node systems, either through a shared file system or by the Cray environment automatically staging the applications shared library dependencies on the compute node.
- **Do not enable memory debugging on the aprun starter process.**
Turning on memory debugging for the aprun starter process will cause it to fail and

prevent the job from starting. The proper way to enable memory debugging is to use the static or dynamic linking options described above.

Reverse debugging on Cray systems

With the 2021.3 release of TotalView, when trying to enable reverse debugging on a process the error message “Couldn’t attach to at least one group member” is raised. This issue will be addressed in the next version of TotalView.

Use the pointers below to help achieve a successful reverse debugging session within the Cray environment:

- **Do not enable reverse debugging on the aprun starter process.**
Turning on reverse debugging for the aprun starter process will cause it to fail and prevent the job from starting. The proper way to turn on reverse debugging is to launch your parallel job and reach a breakpoint in the code and then dynamically turn on the Replay Engine reverse debugging option. It will begin recording the execution of the program from that point forward.
- **Reverse debugging not working on inter-node jobs on Cray systems**
When debugging processes within an inter-node job, the reverse debugging engine will crash, causing the debugging session to become unusable.

AIX

Debugger Target Hangs After Fork/Exec on AIX 7.2 and 7.3

Targets launched through a fork/exec on AIX 7.2 and 7.3 occasionally hang in low-level ptrace() system calls as TotalView tries to write to the target processes address space. This has been traced to a kernel bug in the underlying AIX operating system. Default 1163968 has been filed with IBM.

TotalView Help Does Not Display with Firefox on AIX

Firefox on AIX is unable to display the TotalView help due to an error in the processing of the help set. A workaround is to view the TotalView help set online at <https://help.totalview.io> with another browser.