

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

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 on the Rogue Wave web site:

<http://www.roguewave.com/support/product-documentation/totalview.aspx>

Additions and Updates

TotalView “What’s New” Splash Screen

In order to better promote new features in TotalView, a What’s New splash screen will be displayed after a new version of the product is installed. The dialog provides convenient information about the latest new features and links to learn more about them. The splash screen can easily be dismissed and not shown again until the next new release is installed. To view the What’s New dialog again, simply select the “What’s New in TotalView” menu item from the Help menu.

TotalView’s New “NextGen” UI Officially Supported

Starting with 2018.2, TotalView’s new user interface is now officially supported and no longer an Early Access feature. For existing TotalView users, the default UI is still the TotalView Classic UI, but the default can easily be changed by changing the preference on the Display tab in the Preferences dialog. For brand new TotalView users, the new UI will be displayed by default. New features continue to be added to the new UI. Your feedback helps prioritize what features from the Classic UI need to be added to the new UI. Please send email to tv-beta@roguewave.com with your thoughts about the new UI and any feature requests.

CUDA Debugging Model and Unified Display Improvements

With the 2018.2 release, TotalView improves on the ability to easily set action points within CUDA applications and applications that dynamically load shared libraries with **dlopen**. In either case, until the CUDA or shared library code is loaded, the information

required for setting a breakpoint is not available to the debugger. To address this issue, TotalView now allows setting a breakpoint on any line in the Source View, whether or not it can identify executable code for that line. The breakpoint becomes either a pending breakpoint or a sliding breakpoint until the CUDA or shared library code is loaded at runtime. For more information, see Chapter 27, CUDA Debugging Model and Unified Display, in the TotalView for HPC User Guide.

New User Interface Improvements

TotalView's new user interface, activated through the Display Preferences panel or using the `-newUI` command line option, continues to deliver new enhancements to make debugging your applications even easier. If you have any feedback about the new user interface, requests for new or missing features or any problems please send email to tv-beta@roguewave.com.

- Welcome to TotalView's New UI page
A "Welcome to TotalView's New UI" page has been added to the TotalView's new user interface. When you start the new UI this page is displayed with resources about the new UI and how to learn to use it.
- Set PC
The ability to set the PC to a new line location has been added. To change the PC, simply select a valid line in the source code and then choose the Thread | Set PC menu item or simply hit the "p" key.
- Data View improvements
A number of enhancements to better display changing data and data collections during program execution.
- Action Point improvements
Changes to better work with CUDA Unified Breakpoints and showing the accurate state of action points as dynamic code from CUDA kernels or shared libraries have been made to the new UI.
- Multi-process debugging improvements
Various aspects of the new user interface now use the share group of the process in focus for operations and data displayed. For example, only source files and action points related to the current share group in focus are shown. When focus to another process in a different share group is made, the source view and action point view will update to show the related files and action points.
- Bug fixes and improvements
Numerous bug fixes and minor improvements have been made to the new UI.

Display `std::string` Value Without Diving

New with 2018.2, TotalView now displays the resulting string value of a `std::string` without having to dive on the string and view it in the Data window! The string is now displayed in

tooltips, in the local variable pane and other locations where you would expect the string value to be displayed.

QString Type Transformation

TotalView 2018.2 automatically transforms instances of type QString in Qt4 and Qt5 applications. No longer will the user be required to locate and manipulate the underlying character data to a human-friendly format. TotalView now does it for you.

Manage Single-Stepper Skip Rules

TotalView now provides the ability to define single-stepper “skip” rules that modify the way source-level single stepping works. These rules identify functions that you are not interested in debugging. Skip rules can be defined to skip *over* a function or *through* a function. In skip over cases, the debugger does not step into the function, but rather over it. These are useful for skipping over library functions such as C++ STL code. Skip through rules tell the debugger to ignore any source-line information for the function, so that single stepping does not stop at source lines within the function. If the function being skipped through calls another function, that call is handled according to the original single-stepping operation. Skip through is most useful for callback or thunk functions. For more information and examples, see documentation for the **dskip** command in the TotalView Reference Guide.

Platform Updates

TotalView 2018.2 introduces support for the following platforms:

Compilers:

- Absoft 18

Bug Fixes for 2018.2

- TVT-26312 Application built using Split DWARF gets Fatal Error on startup
- TVT-26209 No breakpoint boxes shown for lines in a lambda function
- TVT-26091, TVT-18806 Add "skip" functionality when stepping.
- TVT-26006 Library dependency problem running TotalView on ARM64 RHEL 7.
- TVT-25805 Put user id in FNE log file.
- TVT-25701 Fix documentation URL link
- TVT-25648 Qt5 QString transformation
- TVT-25643 Internal error when selecting a Stack Trace entry
- TVT-25407 -genv in launch string causing MPI application to not run under PBS when memory debugging is enabled.
- TVT-25081 TotalView with memory debugging under SLURM gets -geometry invalid argument.
- TVLH-6713 Support displaying std::string without having to dive.
- TVLH-6671 Program arguments dropped in new UI with -load_session
- TVLH-1562 Allow a user to set a new location for the PC in the new UI.

Deprecation Notices

Mac OS X El Capitan

TotalView release 2018.1 was the last release to support Mac OS X El Capitan.

IBM Blue Gene/L and Blue Gene/P

TotalView release 2018.1 was the last release to support the IBM Blue Gene/L and Blue Gene/P systems.

RedHat Enterprise Linux 5

Starting with release, 2018, TotalView no longer supports RHEL 5.

Sun Solaris on x86-64 (Opteron)

Starting with release, 2018, TotalView will no longer support Sun Solaris on x86-64 (Opteron).

Intel IA-64 Linux

TotalView no longer supports Intel IA-64 Linux.

32-bit macOS Application Debugging Support on High Sierra and later OS versions

Apple is phasing out support for 32-bit applications beginning with their latest High Sierra release and because of this TotalView will no longer support debugging 32-bit applications on High Sierra and later versions of macOS.

Known Issues

Python Debugging

Anaconda

TotalView supports debugging of the python interpreter in release 4 of Anaconda but is not working with the recent release of Anaconda 5. Something in the way they build the python interpreter has broken the ability to debug python.

Ubuntu 16.04

When debugging python on Ubuntu 16.04 TotalView is not detecting the python interpreter automatically and does not turn on python filtering. Filtering can be turned on by clicking the “filter” icon in the toolbar of the Call Stack view.

Licensing

TotalView releases built with FlexNet Publisher 11.13.1 must have licenses served by a license server at 11.13.1 or higher

FlexNet Publisher client library version 11.13.1 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 v11.13.1 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 v11.12.1 from previous TotalView releases. In this case the vendor daemon’s debug log file shows “(toolworks) Request denied: Client (11.13) newer than Vendor Daemon (11.12). (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

Linux ARM64 and Linux PowerLE use FlexNet Embedded Licensing Technology

The Linux ARM64 and Linux PowerLE platforms use FlexNet Embedded for their licensing technology while the remaining TotalView platforms use FlexNet Publisher. As a result, customers using Team Plus tokens are unable to share their tokens with the Linux PowerLE and Linux ARM64 platforms. In this case, customers should contact license@roguewave.com to set up a new license file for the FlexNet Embedded license server that includes a portion of the Team Plus tokens. TotalView 2018.1 will provide full Team Plus support across the platforms.

macOS

Physical console access needed when running TotalView on macOS High Sierra

Due to new security changes in macOS High Sierra, TotalView will only run from the console and cannot be run through a remote desktop technology such as VNC. We are still assessing what changes need to be made to TotalView so that it will run remotely on macOS High Sierra systems.

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 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.

Linux

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 a number of 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 a number of 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 8.0 Debugger API Internal Error

Certain NVIDIA driver versions associated with CUDA 8.0 may provoke "CUDA Debugger API internal error" messages from TotalView or CUDA-GDB. Known driver versions that have this problem are r361, r367, and r375, however the problem may exist in other driver versions. The internal error typically involves debugging a multi-thread application, when multiple host threads in the process are launching CUDA kernels. 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 may print a message and the program may appear to hang.

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@roguewave.com and let us know how you are using textures so we can develop the best solution to support you.

Solaris

[Oracle Studio 12u4 – TotalView unable to evaluate virtual function calls](#)

When debugging applications compiled with the latest version of the Oracle Studio 12u4 compiler, TotalView is unable to call virtual functions through its expression system. This appears to be a shortcoming in debug information from the compiler and should be addressed in cooperation with the Oracle compiler team.

SGI

[Memory debugging MPI programs on SGI systems needs special linking](#)

The TotalView and MemoryScape memory debugging Heap Interposition Agent (HIA) technology conflicts with the SGI memory manager when used in MPI programs. The easiest way to get around this problem is to disable the SGI memory manager by unsetting the `MPI_MEM_ALIGN` environment variable. Without this variable set, the SGI memory manager will not be loaded and the HIA will work correctly, enabling memory debugging to take place.

Cray

Debugging your program within supercomputer environments can often be challenging. Reference the sections below to learn pointers on how to successfully enable memory debugging and perform reverse debugging on your program within a Cray environment.

[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.
- **Cray TotalView Support Module is not required for TotalView 8.15.0.**
As of TotalView 8.15.0 and its use of the MRNet tree framework TotalView no longer requires the Cray TotalView Support Library to be installed in order to run. If the

MRNet tree framework is turned off then the Cray TotalView Support Module will be required.

- **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.
- **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

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 and TotalView exits.