



The logo for LRZ, consisting of the lowercase letters "lrz" in a bold, white, sans-serif font, centered within a light blue square.

TotalView Training for LRZ

Dean Stewart – Principal Sales Engineer Perforce

Bill Burns – Vice President, Software Engineering
Perforce

Susanne Horn – SMB GmbH

Confidentiality Statement

The information contained in this document is strictly confidential, privileged, and only for the information of the intended recipient. The information contained in this document may not be otherwise used, disclosed, copied, altered, or distributed without the prior written consent of Perforce Software, Inc.

Agenda

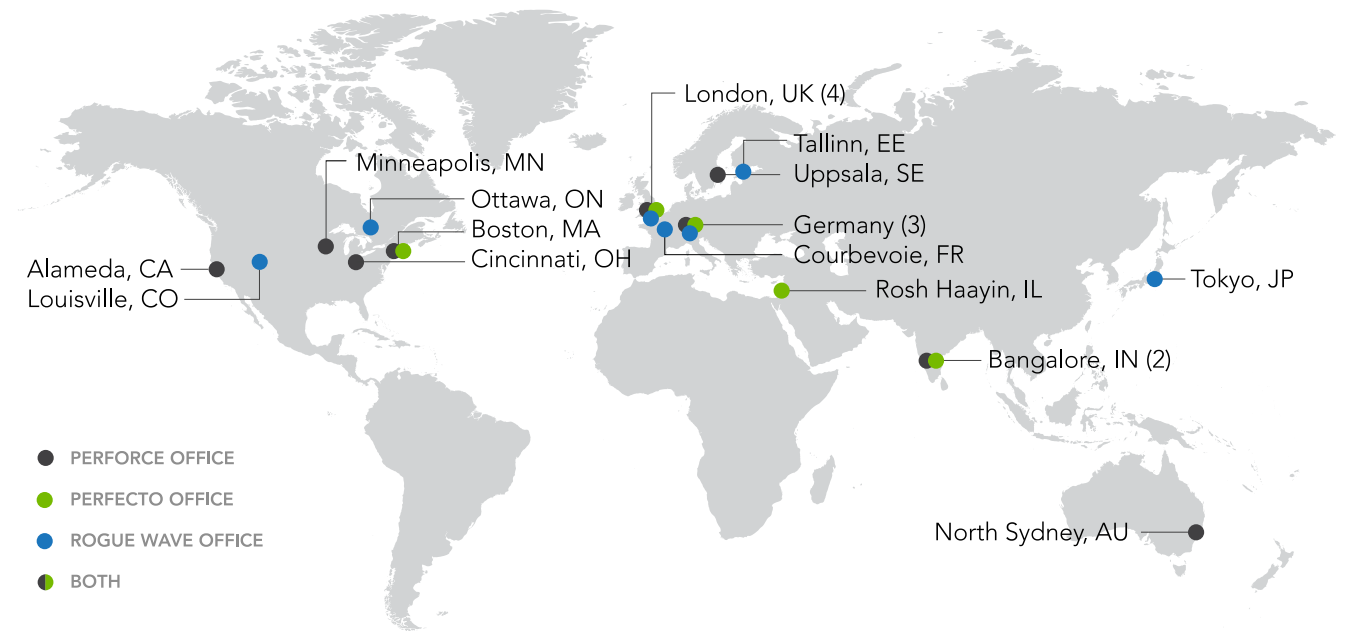
Agenda

- Introduction to Perforce
- HPC Debugging with TotalView
- Multithreaded debugging (OpenMP and pthreads)
- Remote debugging
- Multiprocess debugging
- Understanding complex data structures (advanced C++ and data debugging)
- Debugging CUDA applications
- Mixed Language C/C++ and Python debugging
- Reverse debugging with ReplayEngine
- Memory debugging (memory leaks, buffer overflows, dangling pointers)
- Batch debugging with tvscript
- Debugging I/O bottlenecks
- Best practices for debugging with TotalView
- Q&A

Introduction to Perforce

Global Footprint

- Customers in 80 countries
- 9,000 customers worldwide
- More than 250 of the Fortune 500
- Customers deploying multiple products
- 16 offices and 4 data centers which give us global reach
- 1,700 employees



Perforce Portfolio



Digital Creation

P4 Server

P4 DAM

P4 Plan

ALM

IPLM

P4 ONE



Developer Tools

Akana

OpenLogic

Zend

JRebel

Gliffy

Components

TotalView



Testing Automation

Perfecto

BlazeMeter

Klocwork

Helix QAC



Platform Automation

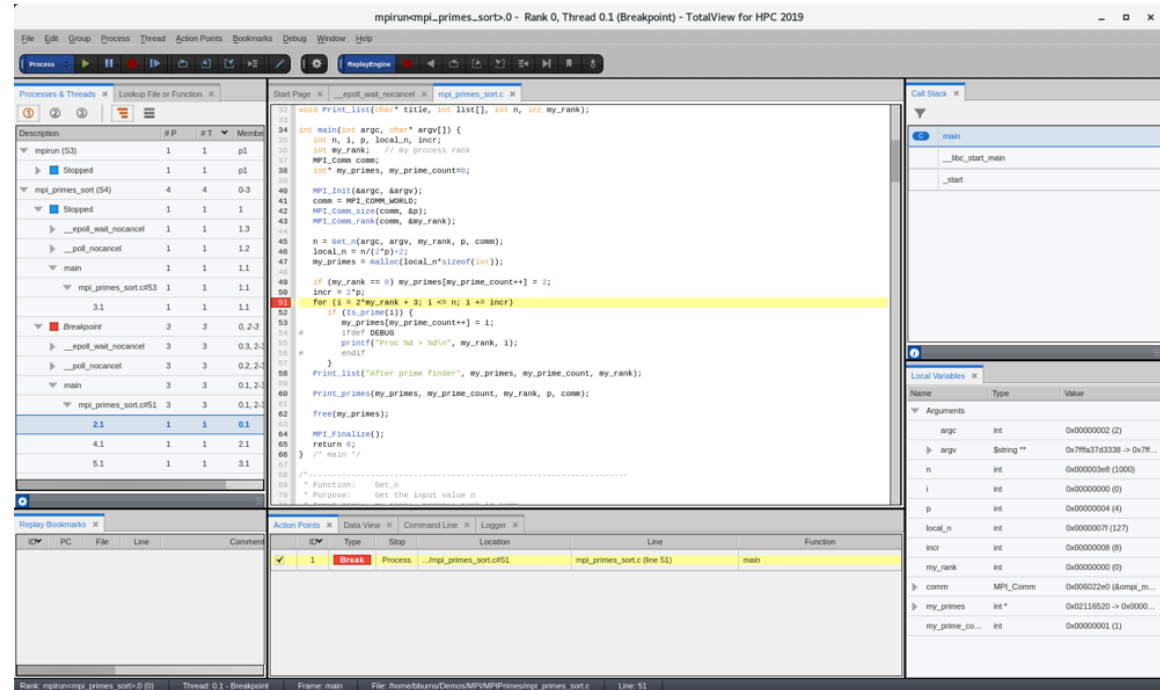
Puppet

Delphix

HPC Debugging with TotalView

HPC Debugging with TotalView

- Comprehensive multi-process/thread dynamic analysis and debugging
- Debug hybrid MPI/OpenMP applications
- Advanced C, C++ and Fortran support
- NVIDIA CUDA GPU debugging support
- AMD / ROCm GPU debugging
- Integrated reverse debugging
- Mixed language C/C++ and Python debugging
- Memory debugging and leak detection
- Batch/unattended debugging



LANGUAGES



OPERATING SYSTEMS



APPLICATIONS



PLATFORMS



Multithreaded Debugging

Multithreaded Debugging with TotalView

- Built from the beginning to support debugging parallel programs
- Supports parallel technologies
 - Processes: fork/exec, MPI, etc
 - Threads: pthreads, TBB, OpenMP, GPU threads
- Provides parallel constructs across debugger features
 - Easily understand overall program state with an aggregate view of all processes and threads
 - Control program behavior with group, process, and thread breakpoints and stepping operations
 - Easily view program data across processes and threads
- Scalable to thousands of threads and processes
- Multi-GPU support

Description	# P	# T	Members
tx_fork_loop (S3)	4	4	p1-4
Breakpoint	4	4	p1-4
Breakpoint	4	4	p2.1, p4.1, p3.2,...
snore	4	4	p2.1, p4.1, p3.2,...
tx_fork_loop.cxx#682	4	4	p2.1, p4.1, p3.2,...
1.3	1	1	p1.3
2.1	1	1	p2.1
3.2	1	1	p3.2
4.1	1	1	p4.1
Stopped	4	8	p1.1, p3.1, p1-2,...
__select_nocancel	2	3	p1-2.2, p2.3
<unknown line>	2	3	p1-2.2, p2.3
1.2	1	1	p1.2
2.2	1	1	p2.2
2.3	1	1	p2.3
snore	3	5	p1.1, p3.1, p4.2,...
tx_fork_loop.cxx#682	3	5	p1.1, p3.1, p4.2,...
1.1	1	1	p1.1
3.1	1	1	p3.1
3.3	1	1	p3.3

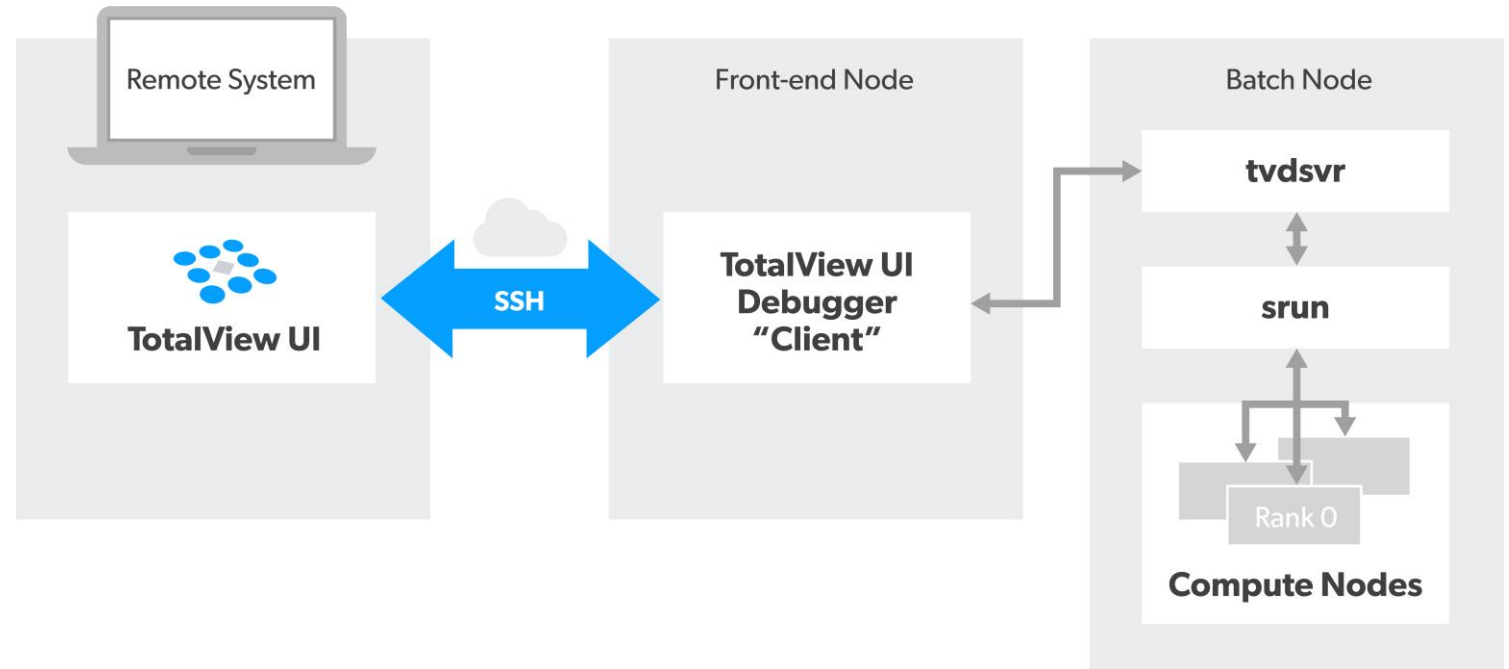
OpenMP/OMPD Support

- Extends and productizes existing TotalView OpenMP/OMPD debugging support
 - Focuses on supporting the Clang, AMD Clang/Flang, and HPE CCE C, C++, and Fortran compilers / OMPD libraries
- Single-step into and out of parallel regions automatically
- Display OpenMP runtime state of parallel/task regions, threads, control variables, and ICVs
- Enhanced stack displays using OMPD-provided parallel/task region information
 - Filter-out OpenMP runtime frame
 - Annotate parallel/task region frames with “#pragma omp parallel” / “#pragma omp task”
 - Insert “parent back-link” frames to focus on a parallel region’s encountering thread/frame
- Support for OMPD interface defined by the OpenMP 5.0 - 5.2 APIs
- Note: OMPD-based features are CPU only
 - An OMPD library implementation for GPUs does not exist yet

Remote Debugging

Remote Debugging with TotalView

- Combine the convenience of establishing a remote connection to a cluster and the ability to run the TotalView GUI locally
- Front-end GUI architecture does not need to match back-end target architecture (macOS front-end -> Linux back-end)
- Secure communications
- Convenient saved sessions
- Once connected, debug as normal with access to all TotalView features
- Windows, macOS, Linux Support



Multiprocess Debugging

Starting a Parallel Program Session from the UI

From New Parallel
Session page select:

Parallel
Environment

Launch Style

Number of tasks

Number of nodes

Starter arguments

Click Start Session
to save and launch

The screenshot shows the 'Session Editor' window with the following sections:

- Session Details:** Session Name (dropdown menu with placeholder text: [Enter or select a session name, e.g. myprogram with ReplayEngine]).
- Parallel Details:** Parallel Environment (dropdown menu with 'Slurm srun' selected and a 'REQUIRED' tag).
- Parallel Launch Style:** Radio buttons for 'Native' (selected, with subtext 'Use the MPI starter process to launch the job (recommended)') and 'Bootstrap' (with subtext 'Use the bootstrap server to indirectly launch the job (if Native fails)').
- Tasks (-n):** Input field with placeholder text [Enter the number of tasks].
- Nodes (-N):** Input field with placeholder text [Enter the number of nodes].
- Additional Starter Arguments:** Input field with placeholder text [Enter starter arguments as needed].
- Program Details:** File Name (dropdown menu with placeholder text [Enter program path and name, e.g. /home/smith/myprogr...], a 'REQUIRED' tag, and a 'BROWSE...' button).
- Arguments:** Text area with placeholder text [Enter any program arguments. Ex. -option foo].
- Debug Options:** Python Debugging (checkbox for 'Enable call stack filtering for Python') and Memory Debugging (checkbox for 'Enable memory debugging').
- Program Environment:** Working Directory (input field).

At the bottom right, there are three buttons: 'RESET', 'LOAD SESSION', and 'CANCEL'.

Starting a Parallel Program Session from the Command Line

General Command Line: `totalview --args <starter> -n ## <partition> <program>`

MPI	Startup Command
Linux under SLURM	<code>totalview --args srun -n 16 -p pdebug <program></code>
Open MPI / MPICH / Intel MPI	<code>totalview --args mpirun -np 16 <program></code>

SLURM

- Use `salloc` or `sbatch` to grab an allocation
- Use `srun` to start the job
`totalview -args srun -N 4 -n 16 <program> <program args>`
- `sbatch tv.sbatch` to submit a batch job
 - The batch script may need to define a `DISPLAY` variable to forward TotalView's display
- Use of `tvconnect` can also simplify a parallel debugging session launch

TotalView Reverse Connections

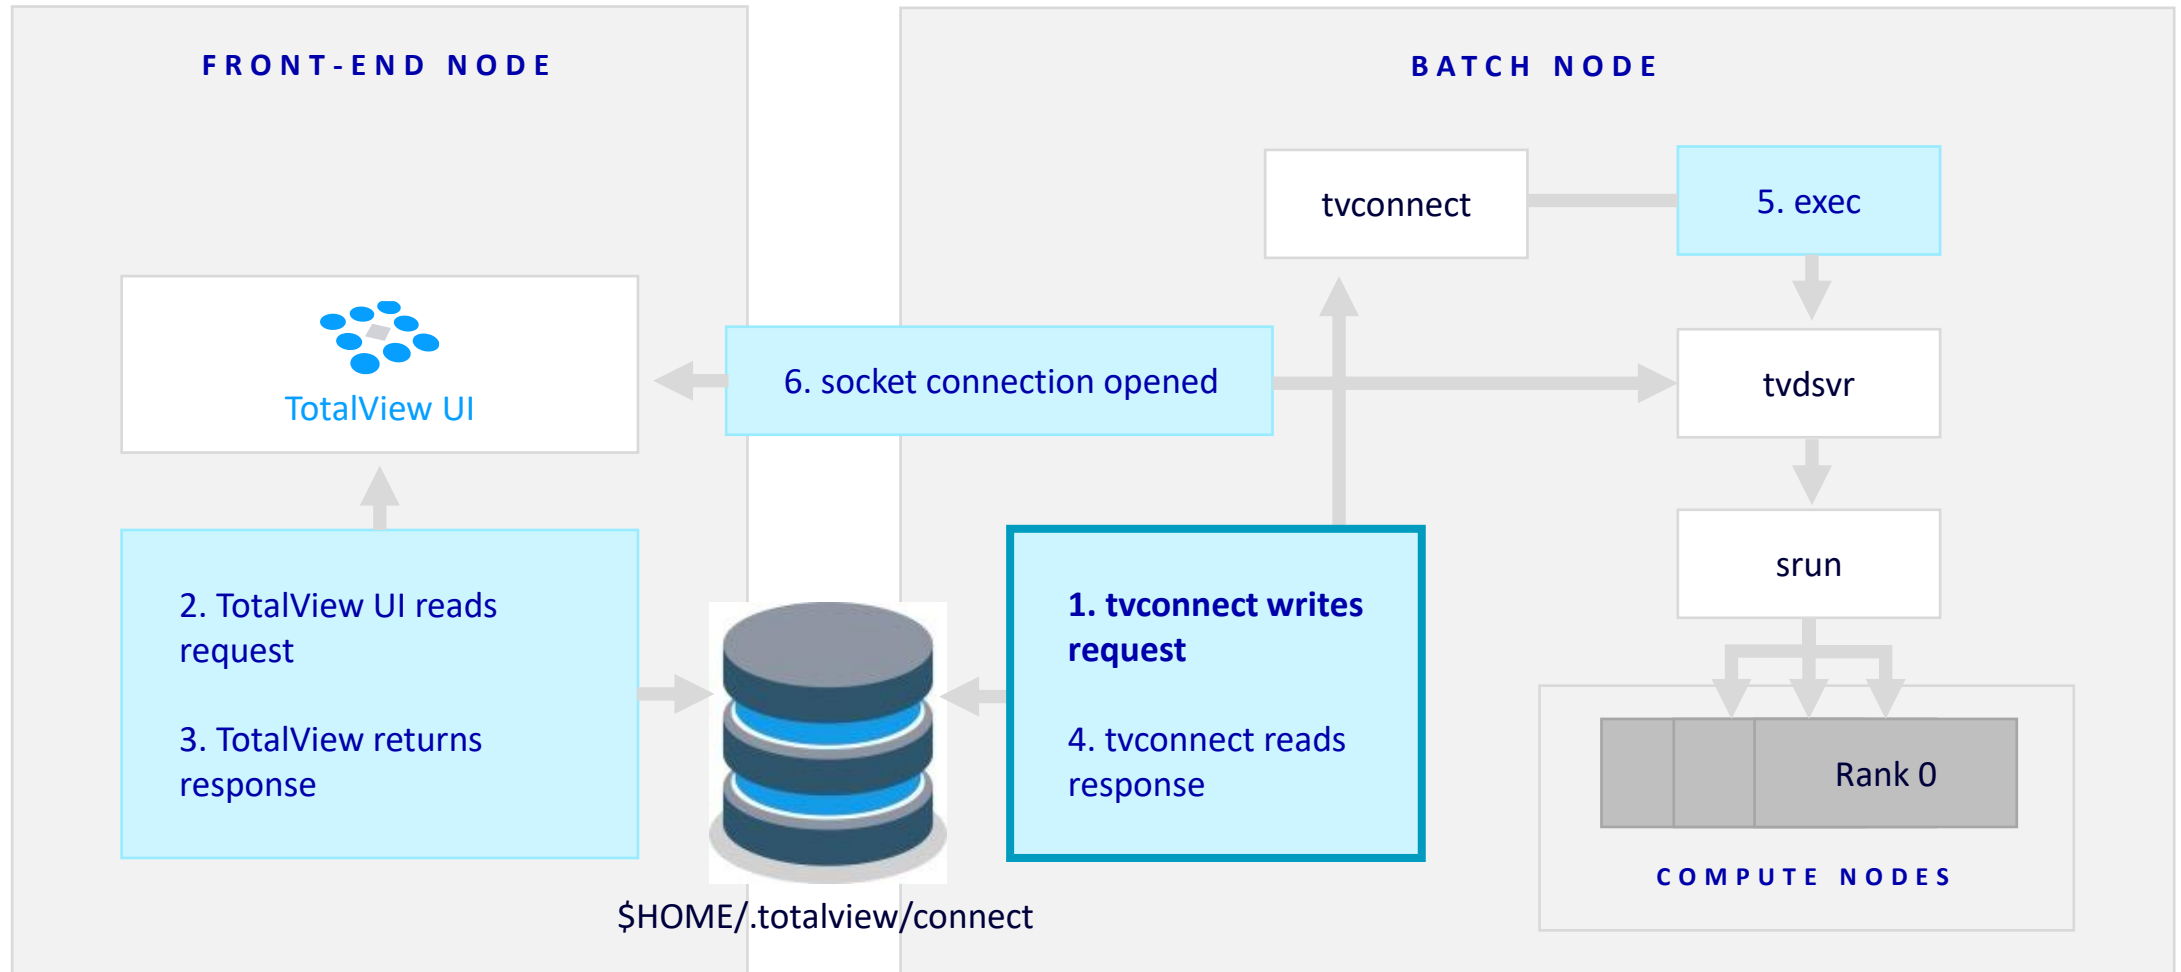
The Problem:

- Establishing an interactive debugging session in a cluster environment can be difficult
 - Timing issues when submitting through a job manager and when the job runs
 - The organization of modern HPC systems often makes it difficult to deploy tools such as TotalView
 - The compute nodes in a cluster may not have access to any X libraries or X forwarding
 - Launching a GUI on a compute node may not be possible

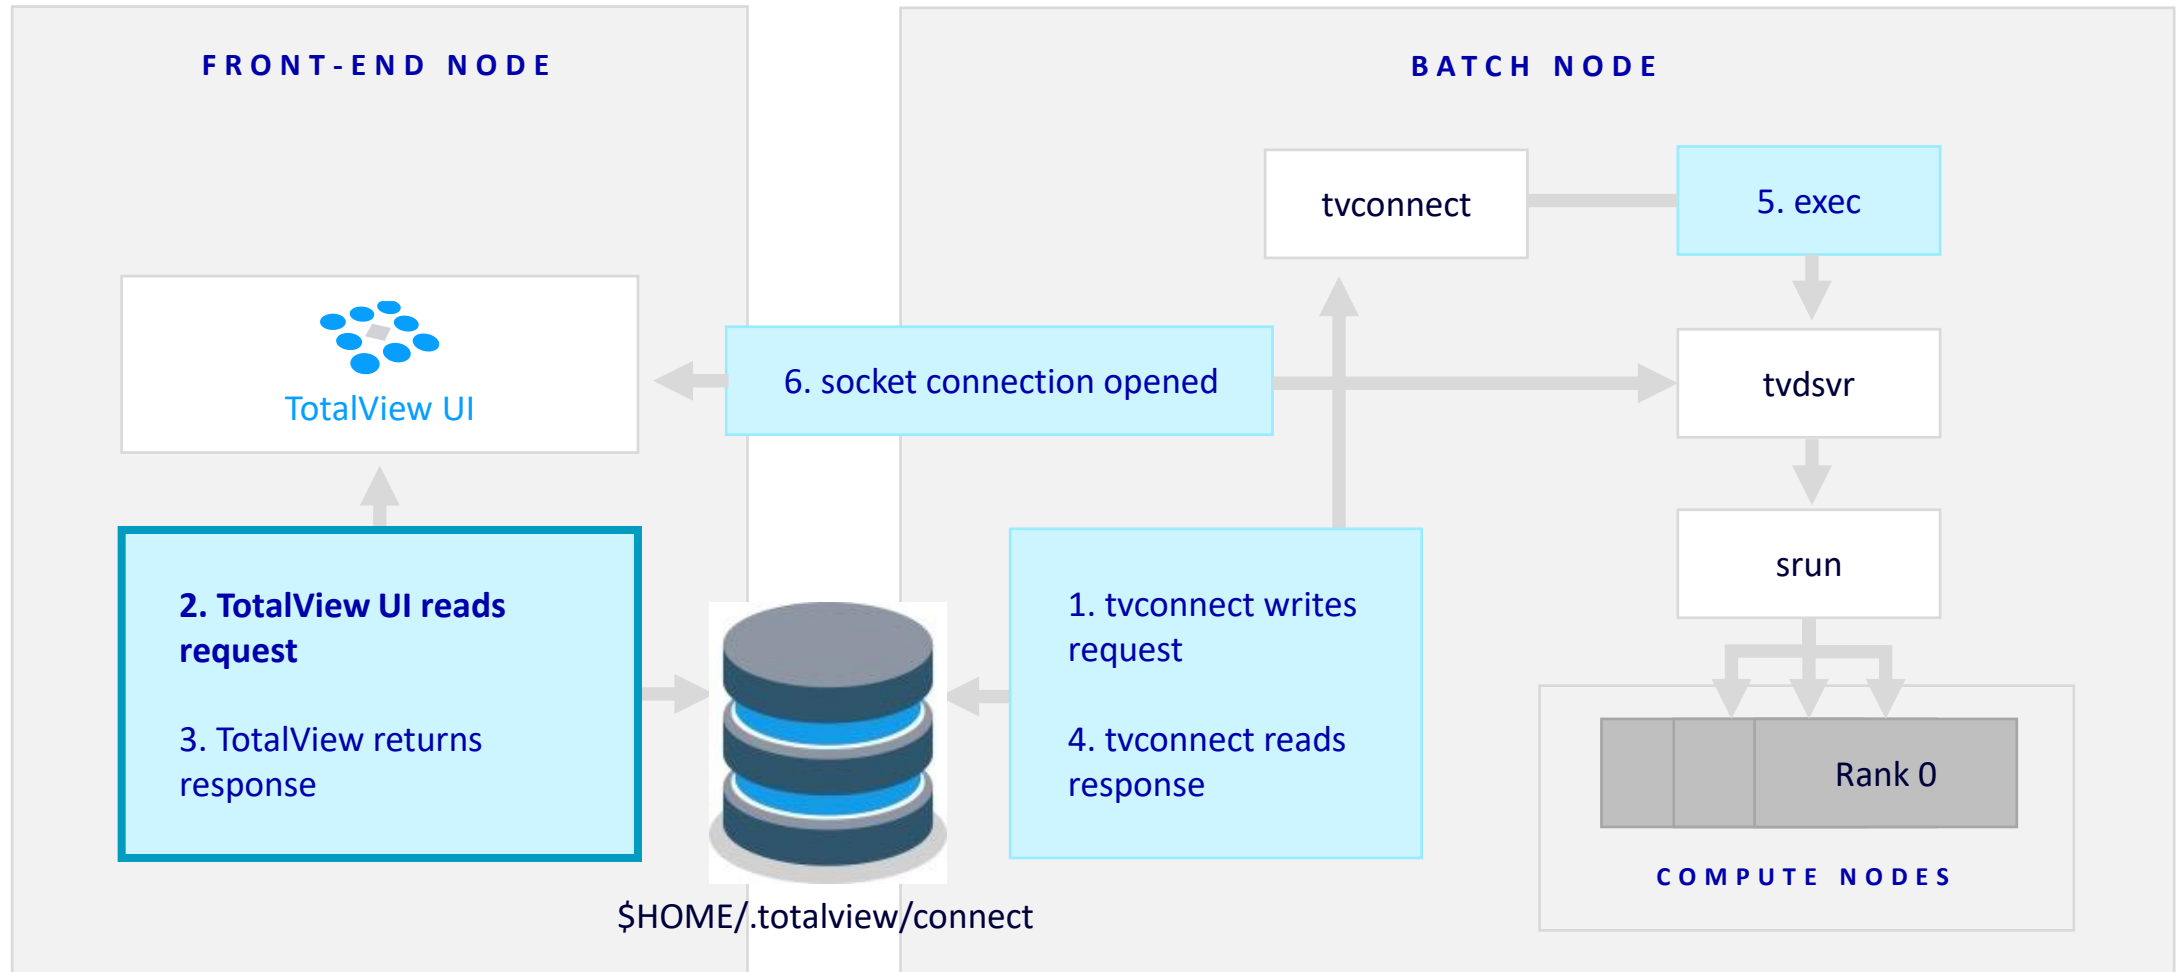
The Solution:

- Disconnect starting debugger UI from the backend job launch and debug session acquisition
- TotalView Reverse Connect workflow enables developers to start the TotalView UI on a front-end node and, when a job is run in the cluster, a remote TotalView reverse connect agent connects it back to the waiting UI

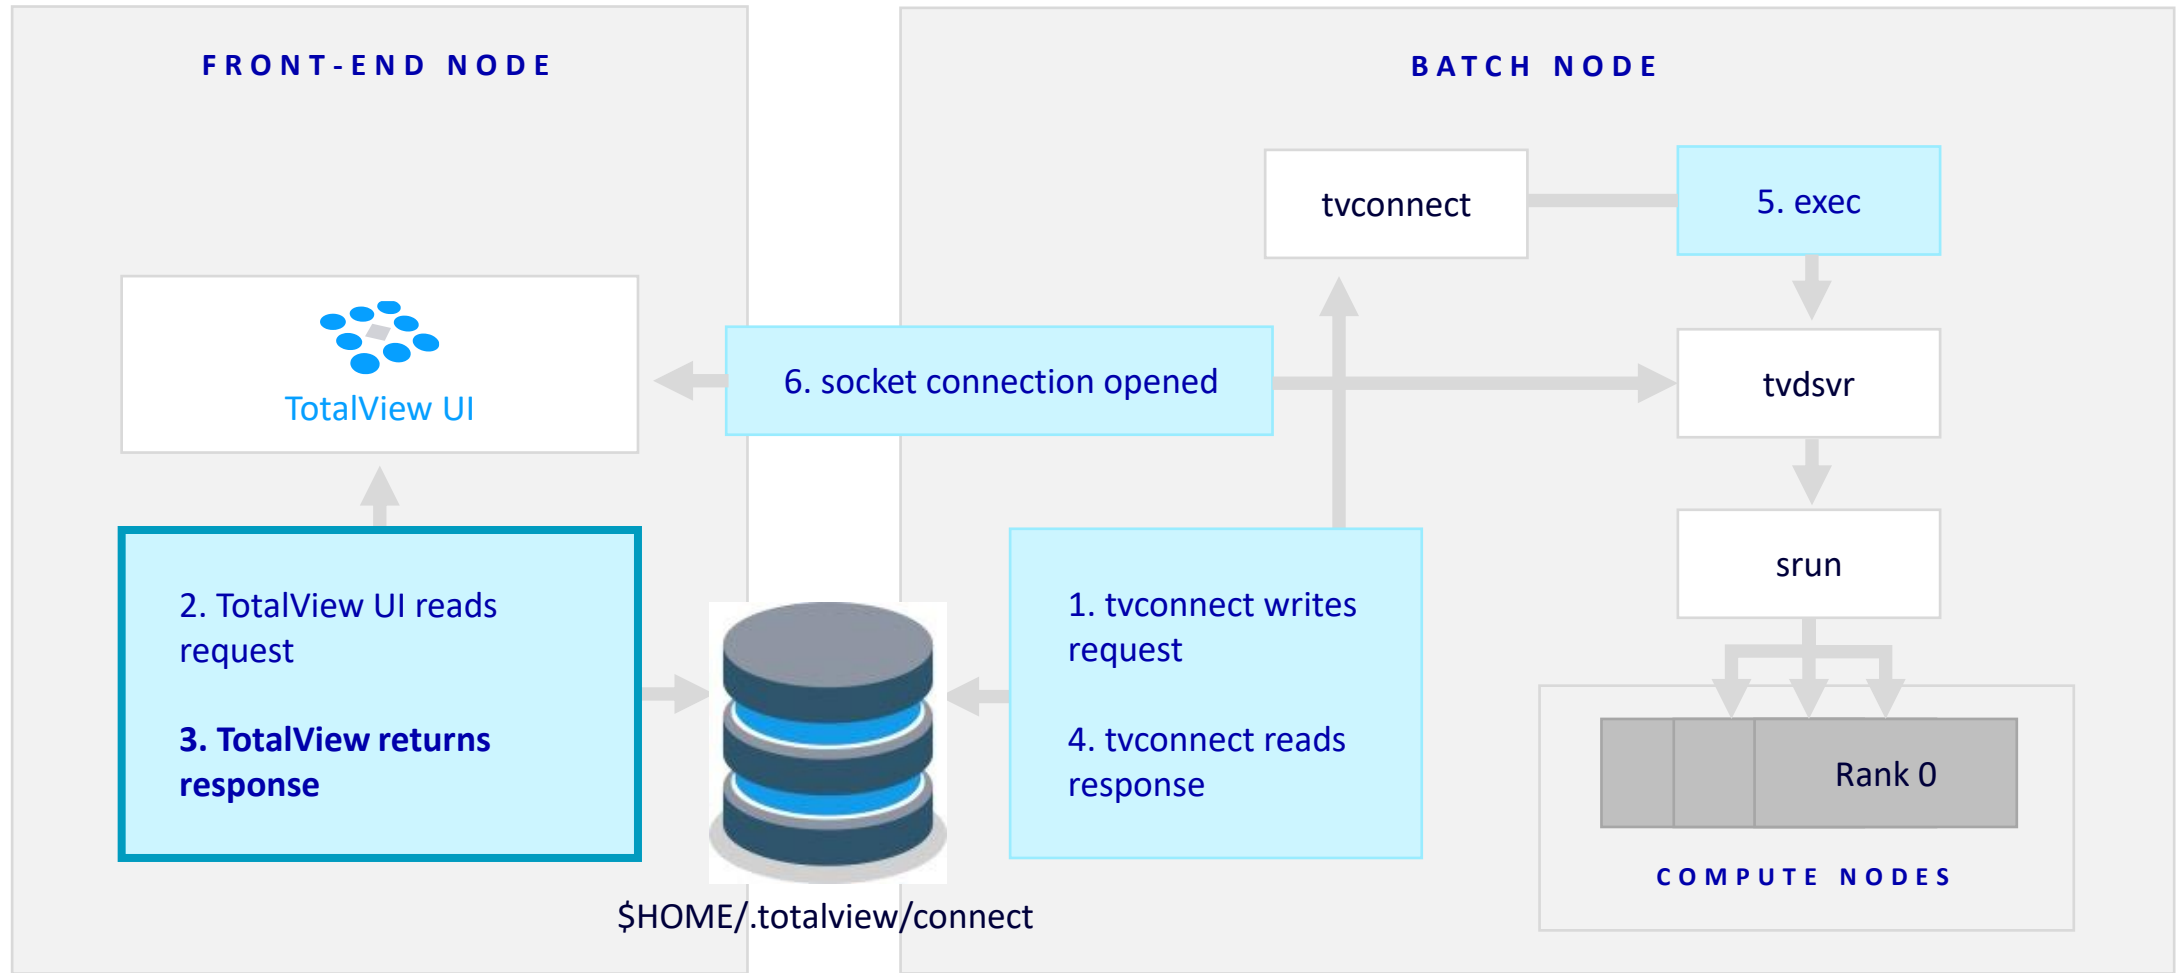
Reverse Connection Flow



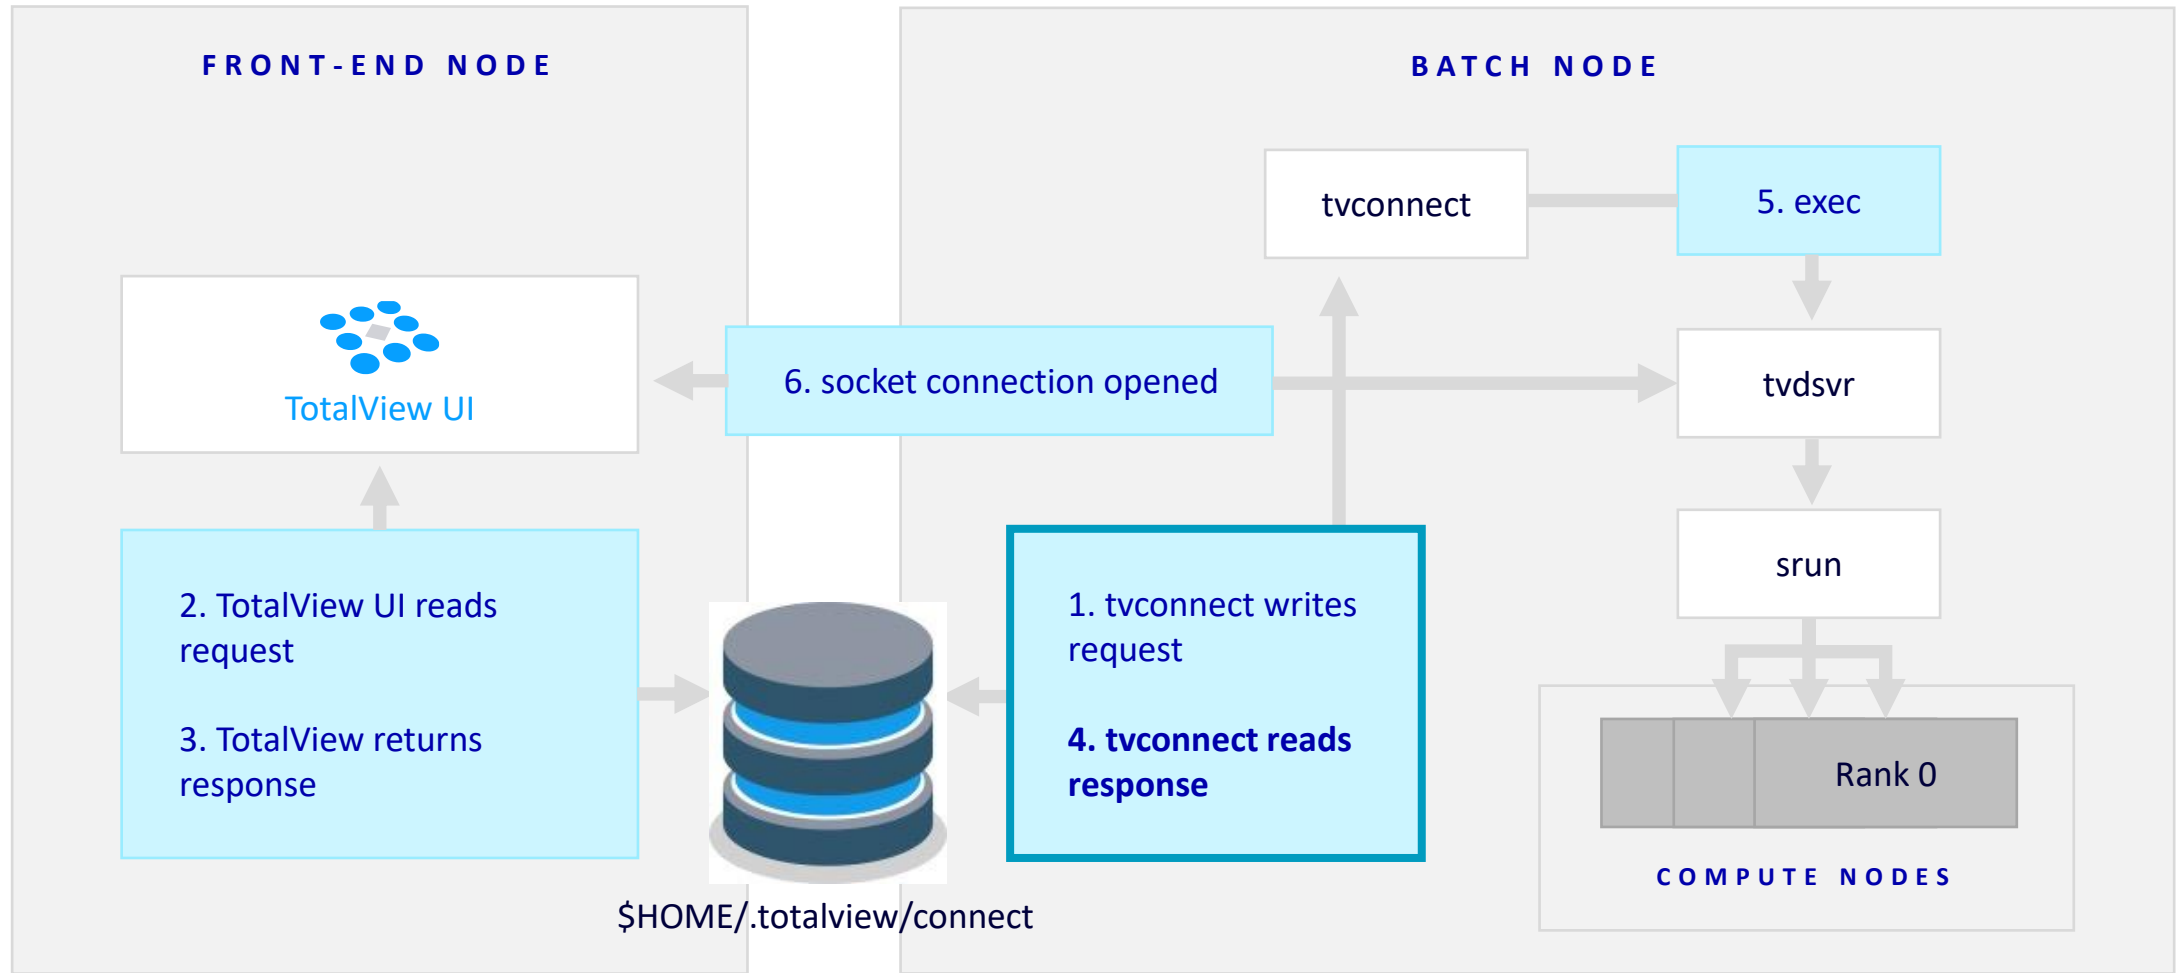
Reverse Connection Flow

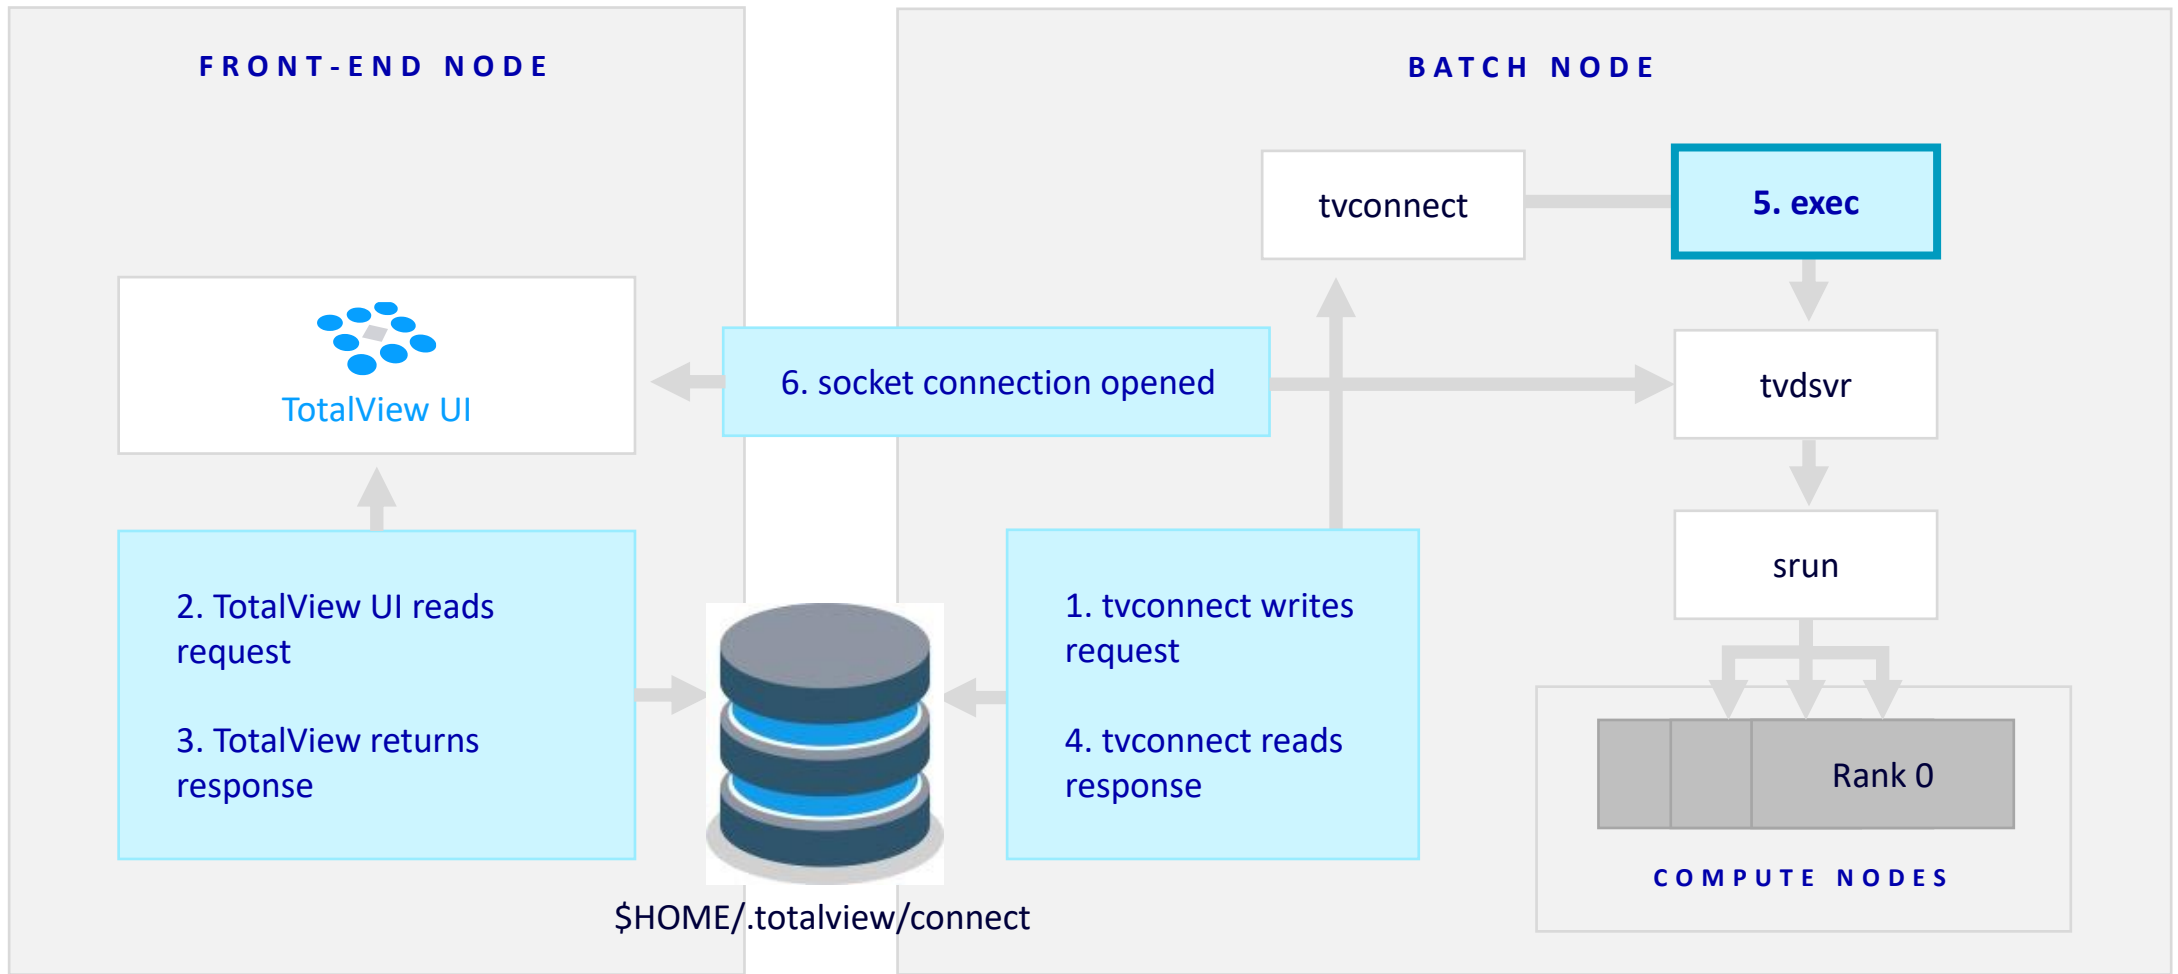


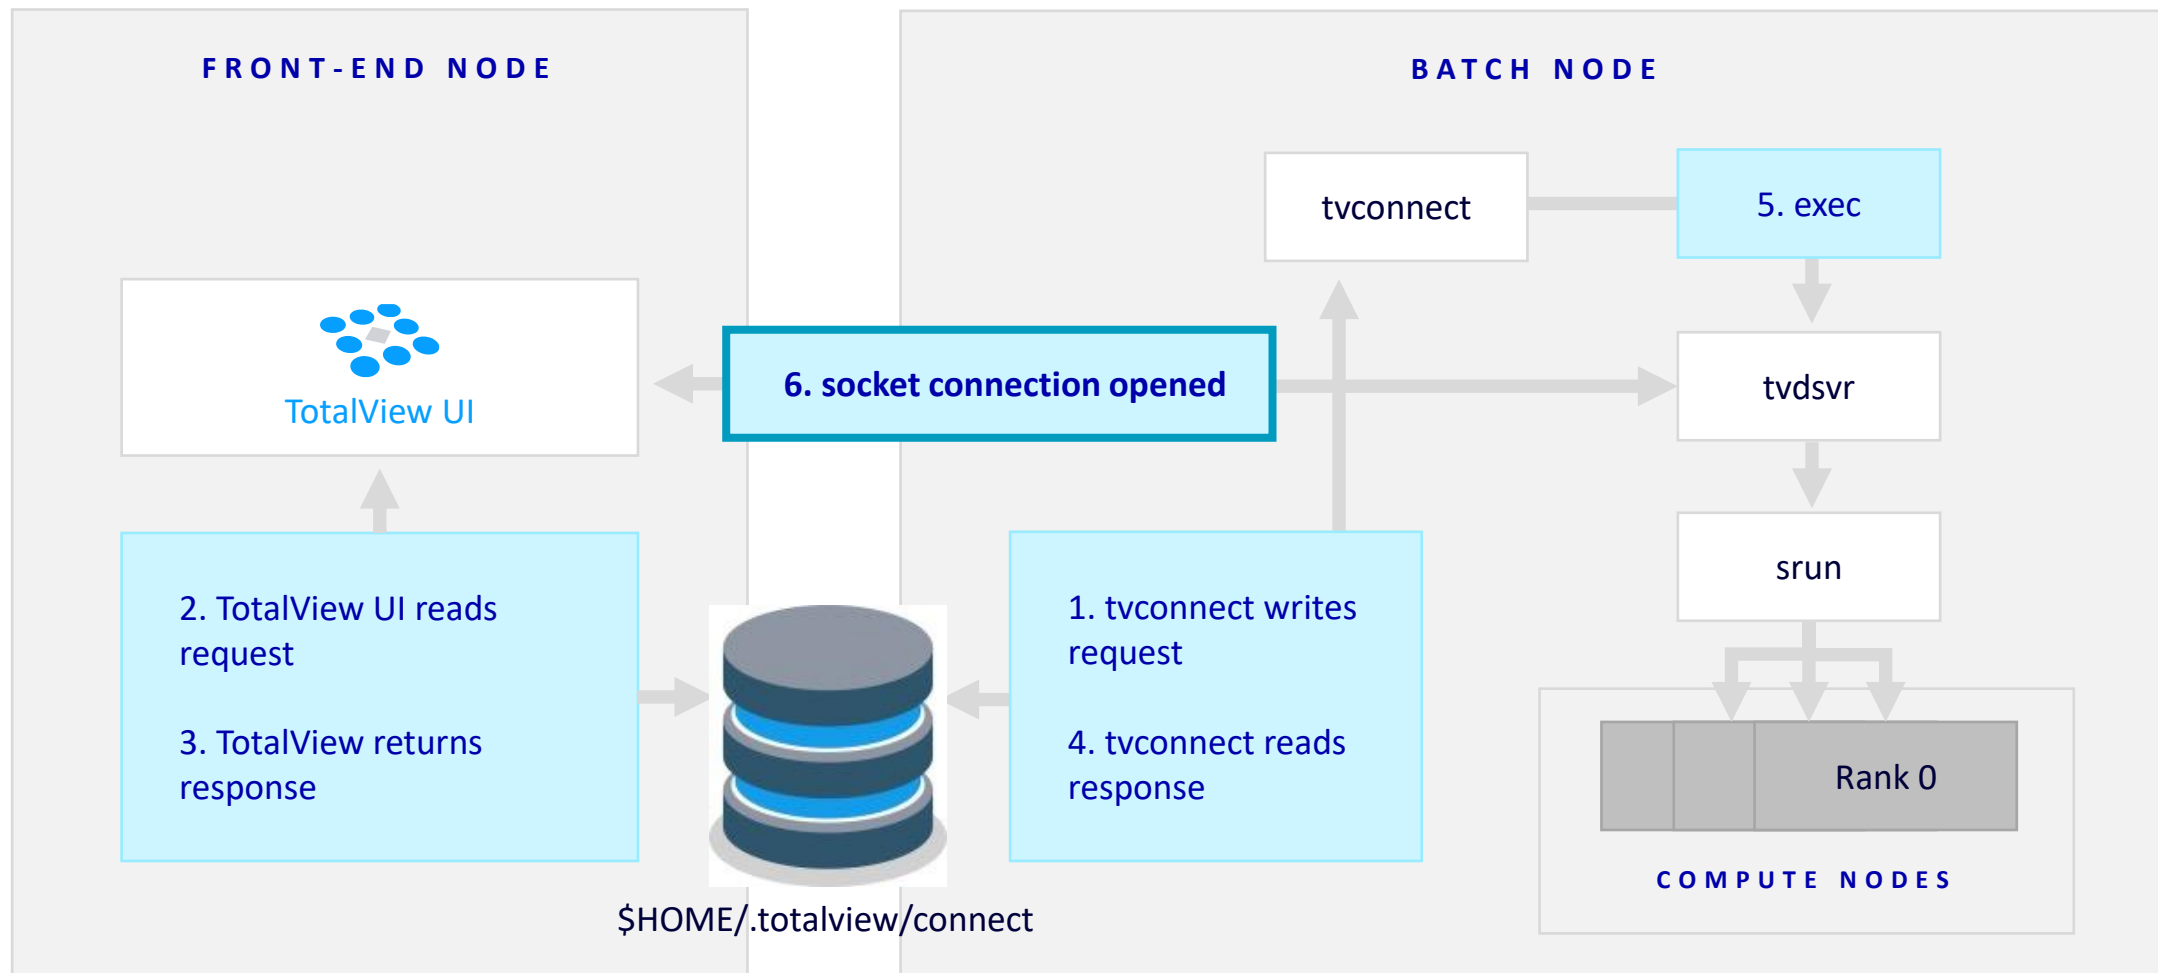
Reverse Connection Flow



Reverse Connection Flow



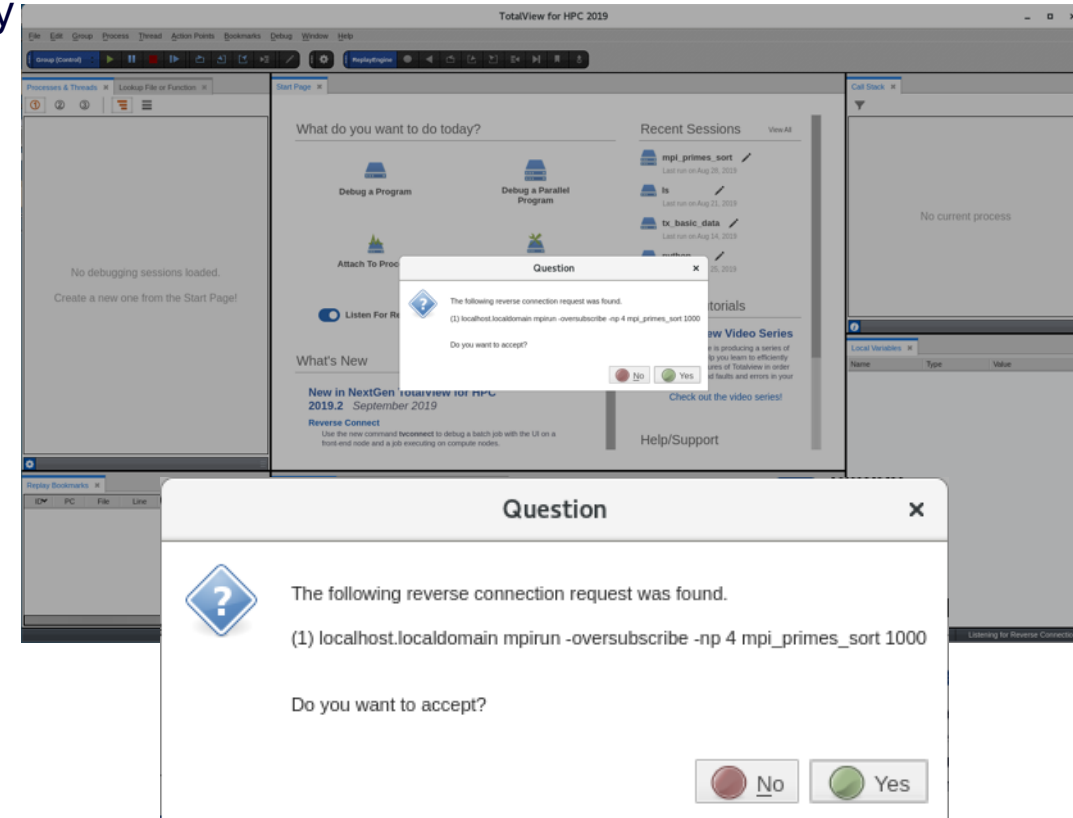




Batch Script Submission with Reverse Connect

- Start a debugging session using TotalView Reverse Connect.
- Reverse Connect enables the debugger to be submitted to a cluster and connected to the GUI once run.
- Enables running TotalView UI on the front-end node and remotely debug jobs executing on the compute nodes.
- Very easy to utilize, simply prefix job launch or application start with “tvconnect” command.

```
#!/bin/bash
#SBATCH -J hybrid_fib
...
#SBATCH -n 2
#SBATCH -c 4
#SBATCH --mem-per-cpu=4000
export OMP_NUM_THREADS=4
tvconnect srun -n 2 --cpus-per-task=4 --mpi=pmix ./hybrid_fib
```



Parallel Debugging Group, Process and Thread Control

Select either

- Group (Control)
- Group (Share)
- Process
- Thread

The screenshot displays the TotalView debugger interface. On the left, a tree view shows the hierarchy: Group (Control), Group (Share), Process, and Thread. Below this is a table with columns for Description, # P, # T, and Members. A blue arrow points to the 'Group (Control)' option in the tree. The main window shows the source code of 'exec1.cpp' with a breakpoint set at line 15. The call stack on the right shows the current frame 'main' and its callers. The bottom status bar indicates the current rank and thread information.

Description	# P	# T	Members
Running	1	1	p1
<unknown address>	1	4	p1.1-4
1	1	4	p1
Breakpoint	4	4	0-3
exec1.cpp#15	2	2	0-1.1
6	1	1	0
7	1	1	1
exec2.cpp#15	2	2	2-3.1
8	1	1	2
9	1	1	3
<unknown line>	4	8	0-3.2-3

```
1 #include <mpi.h>
2 #include <iostream>
3 #include <cstring>
4
5 int main(int argc, char** argv) {
6     MPI_Init(&argc, &argv);
7     int rank, size;
8     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
9     MPI_Comm_size(MPI_COMM_WORLD, &size);
10    if (size != 4) {
11        std::cerr << "This program should be run with 4 MPI ranks." << std::endl;
12        MPI_Finalize();
13        return 1;
14    }
15    if (rank == 0) {
16        std::string message = "Hello from Executable 1!";
17        MPI_Send(message.c_str(), message.size() + 1, MPI_CHAR, 2, 0, MPI_COMM_WORLD);
18    }
19    MPI_Finalize();
20    return 0;
21 }
22
```

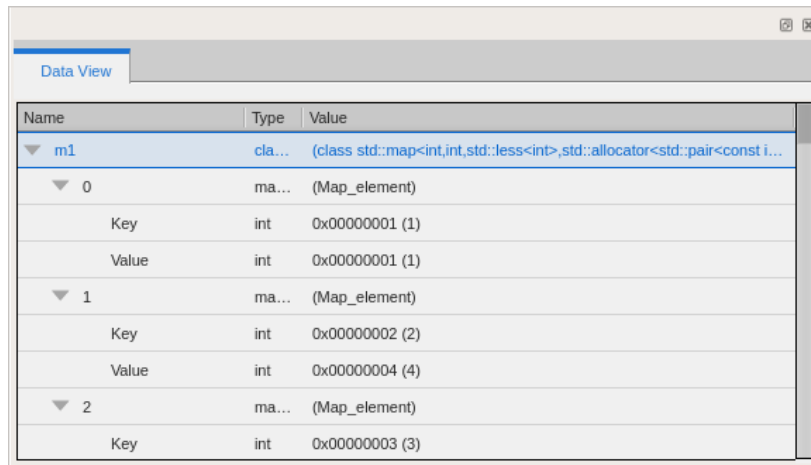
Name	Type	Value
Arguments		
argc	int	0x00000001 (1)
argv	Sstring **	0x7ffc9b78e9d8 -> 0x...
rank	int	0x00000000 (0)
size	int	0x00000004 (4)

Understanding Complex Data Structures

C++ Container Transformations

TotalView transforms many of the C++ and STL containers including:

array, vector, deque, forward_list, list, set, multiset, map, multimap, stack, queue, priority_queue, pair, tuple, unique_ptr, shared_ptr, weak_ptr and others.

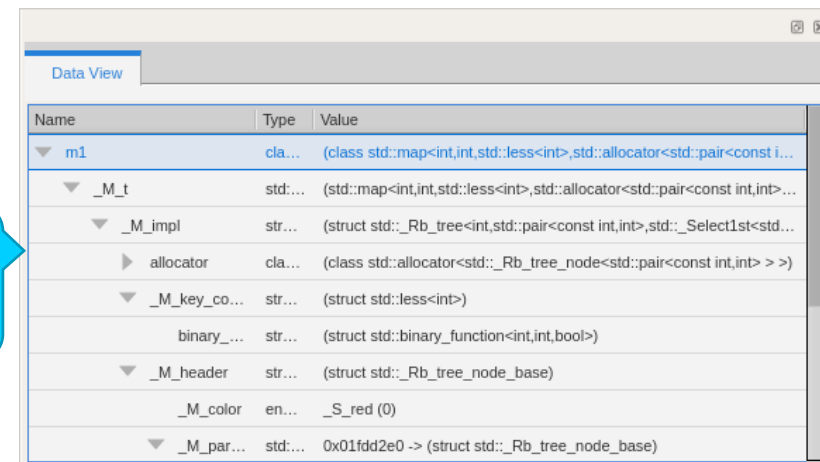


A screenshot of a 'Data View' window showing a transformed map container. The table has three columns: Name, Type, and Value. The root node is 'm1' with type 'cla...' and value '(class std::map<int,int,std::less<int>,std::allocator<std::pair<const i...'. It is expanded to show three elements (0, 1, 2). Each element has a 'Key' and 'Value' sub-property, both of type 'int'.

Name	Type	Value
▼ m1	cla...	(class std::map<int,int,std::less<int>,std::allocator<std::pair<const i...
▼ 0	ma...	(Map_element)
Key	int	0x00000001 (1)
Value	int	0x00000001 (1)
▼ 1	ma...	(Map_element)
Key	int	0x00000002 (2)
Value	int	0x00000004 (4)
▼ 2	ma...	(Map_element)
Key	int	0x00000003 (3)

See This!

Instead of This



A screenshot of a 'Data View' window showing the internal structure of a map container. The table has three columns: Name, Type, and Value. The root node is 'm1' with type 'cla...' and value '(class std::map<int,int,std::less<int>,std::allocator<std::pair<const i...'. It is expanded to show internal members like '_M_t', '_M_impl', 'allocator', '_M_key_co...', 'binary_...', '_M_header', '_M_color', and '_M_par...'. Each member has its corresponding type and value.

Name	Type	Value
▼ m1	cla...	(class std::map<int,int,std::less<int>,std::allocator<std::pair<const i...
▼ _M_t	std:...	(std::map<int,int,std::less<int>,std::allocator<std::pair<const int,int>...
▼ _M_impl	str...	(struct std::_Rb_tree<int,std::pair<const int,int>,std::_Select1st<std...
▶ allocator	cla...	(class std::allocator<std::_Rb_tree_node<std::pair<const int,int> > >)
▼ _M_key_co...	str...	(struct std::less<int>)
binary_...	str...	(struct std::binary_function<int,int,bool>)
▼ _M_header	str...	(struct std::_Rb_tree_node_base)
_M_color	en...	_S_red (0)
▼ _M_par...	std:...	0x01fdd2e0 -> (struct std::_Rb_tree_node_base)

Array Statistics

Easily display a set of statistics for the filtered portion of your array

Start Page × combined.cxx × vol ×

Array: vol (Thread: 1.1) UPDATE

Slice: Type: double[20][20]

Statistic	Value
Count	400
Zero Count	0
Sum	597909.794
Minimum	6.28304
Maximum	7273.40418
Median	915.75308
Mean	1494.774485
Standard Deviation	1566.23066267959
First Quartile	283.91487
Third Quartile	2259.14557
Lower Adjacent Value	6.28304
Upper Adjacent Value	5195.2887
Nan Count	0
Infinity Count	0
Denormalized Count	0
Checksum	45976

Viewing Array Data

- The Array View provides convenient array debugging and visualization features.
- The Data Table displays a 2D slice of a multi-dimensional array.
- Use the Array View Options to control the slice and stride of the array data.

The screenshot displays the 'Array View' window with the following components:

- Statistics Panel:** A table showing summary statistics for the array.
- Data Table:** A 2D slice of the array data, showing memory addresses and values.
- Configuration Options Dialog:** A window titled 'Array View Options' for customizing the view.

Statistic	Value
Count	24000
Zero Count	1
Sum	26748000
Minimum	0
Maximum	2229
Median	
Mean	
Standard Deviation	
First Quartile	
Third Quartile	
Lower Adjacent Value	
Upper Adjacent Value	
Checksum	

[j]:0	1	2	3	4
0x00000000 (0)	0x00000001 (1)	0x00000002 (2)	0x00000003 (3)	0x00000004 (4)
0x0000000a (10)	0x0000000b (11)	0x0000000c (12)	0x0000000d (13)	0x0000000e (14)
0x00000014 (20)	0x00000015 (21)	0x00000016 (22)	0x00000017 (23)	0x00000018 (24)
0x0000001e (30)	0x0000001f (31)	0x00000020 (32)	0x00000021 (33)	0x00000022 (34)
0x00000028 (40)	0x00000029 (41)	0x0000002a (42)	0x0000002b (43)	0x0000002c (44)
0x00000032 (50)	0x00000033 (51)	0x00000034 (52)	0x00000035 (53)	0x00000036 (54)
0x0000003c (60)	0x0000003d (61)	0x0000003e (62)	0x0000003f (63)	0x00000040 (64)
0x00000046 (70)	0x00000047 (71)	0x00000048 (72)	0x00000049 (73)	0x0000004a (74)
0x00000050 (80)	0x00000051 (81)	0x00000052 (82)	0x00000053 (83)	0x00000054 (84)
0x0000005a (90)	0x0000005b (91)	0x0000005c (92)	0x0000005d (93)	0x0000005e (94)
0x00000064 (100)	0x00000065 (101)	0x00000066 (102)	0x00000067 (103)	0x00000068 (104)
0x0000006e (110)	0x0000006f (111)	0x00000070 (112)	0x00000071 (113)	0x00000072 (114)
0x00000078 (120)	0x00000079 (121)	0x0000007a (122)	0x0000007b (123)	0x0000007c (124)
0x00000082 (130)	0x00000083 (131)	0x00000084 (132)	0x00000085 (133)	0x00000086 (134)
0x0000008c (140)	0x0000008d (141)	0x0000008e (142)	0x0000008f (143)	0x00000090 (144)
0x00000096 (150)	0x00000097 (151)	0x00000098 (152)	0x00000099 (153)	0x0000009a (154)
0x000000a0 (160)	0x000000a1 (161)	0x000000a2 (162)	0x000000a3 (163)	0x000000a4 (164)
0x000000aa (170)	0x000000ab (171)	0x000000ac (172)	0x000000ad (173)	0x000000ae (174)
0x000000b4 (180)	0x000000b5 (181)	0x000000b6 (182)	0x000000b7 (183)	0x000000b8 (184)

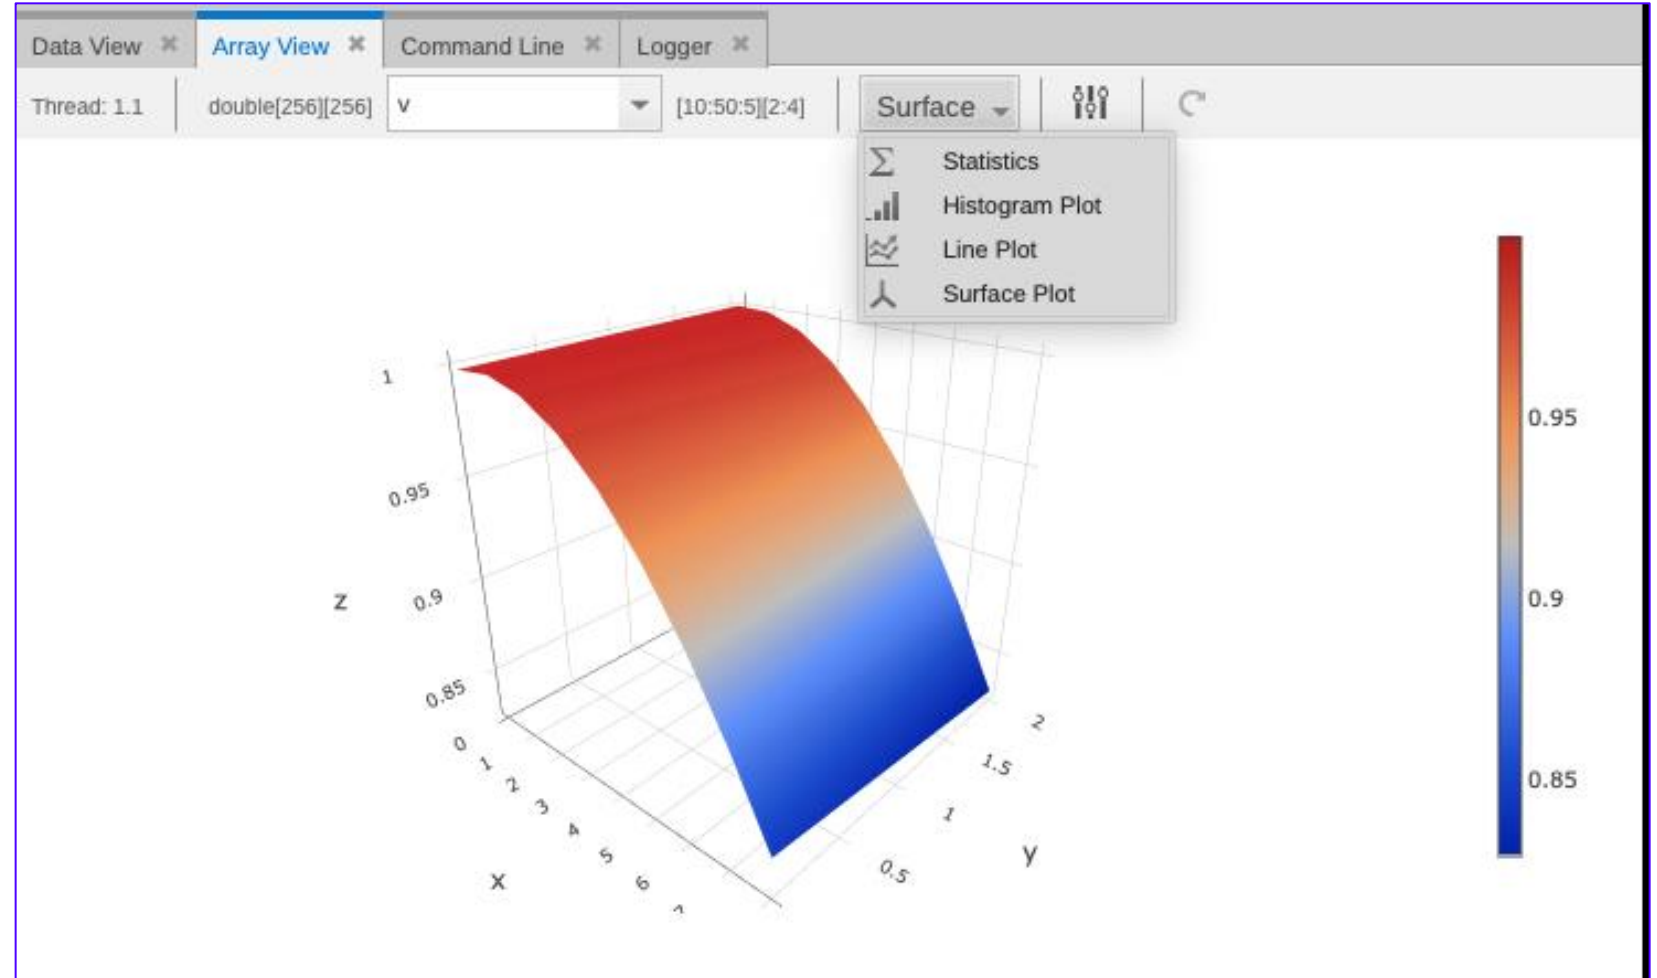
Array View Options Configuration:

- Expression:** threeDArray
- Type:** Cast your array to a different type. `int[20][30][40]`
- Slice:** Define a 2D slice through the array to examine data in a table.

Dimension	Start	End	Stride
Row [j]	0	29	1
Column [k]	0	39	1
- Other Slice Dimensions:** Select an array index. `[i]: 0`, `[l]: 0`, `[m]: 0`
- Filter:** Use filters to reduce array data. `==` Add a value or... `&&` ADD
Example: \$value < 100

Visualizing Array Data

- Select the array in the source code pane or Data View
- Right click and select "Add to Array View"
- Select visualization from the Array View drop down
 - Statistics
 - Histogram Plot
 - Line Plot
 - Surface Plot



Debugging CUDA Applications

Debugging CUDA Applications

- NVIDIA Tesla, Fermi, Kepler, Pascal, Volta, Turing, Ampere and Hopper
- NVIDIA CUDA 9, 10, 11 and 12
 - With support for Unified Memory
- Debug 64-bit CUDA programs
- Features and capabilities include
 - Support for dynamic parallelism
 - Support for MPI based clusters and multi-card configurations
 - Flexible Display and Navigation on the CUDA device
 - Physical (device, SM, Warp, Lane)
 - Logical (Grid, Block) tuples
 - CUDA device window reveals what is running where
 - Support for types and separate memory address spaces
 - Leverages CUDA memcheck



Mixed Language C/C++ and Python Debugging

Python without Filtering

The screenshot shows a debugger interface with the following components:

- Process and Thread List:** Shows a process named 'python3.7-...' with a thread 'Br...' and a function 'fact'.
- Code Editor:** Displays C++ code for a 'fact' function. Line 5 is highlighted in yellow.
- Call Stack:** Lists the call stack entries, including C++ and C entries. A blue circle highlights the entries from 'fact' down to 'call_function'. A green box labeled 'Glue code' points to the C++ entries.
- Local Variables:** Shows the argument 'n' with a value of '0x00000003 (3)'.
- Action Points:** Shows two breakpoints: one at 'dot' and one at 'fact'.

```
1 /* File: example.c */
2
3 #include "tv_example.h"
4
5 int fact(int n) {
6     if (n < 0){ /* This should probably return an error, but this is simpler */
7         return 0;
8     }
9     if (n == 0) {
10        return 1;
11    }
12    else {
13        /* testing for overflow would be a good idea here */
14        return n * fact(n-1);
15    }
16 }
17
18 int getSquare(int n) {
19     return n * n;
20 }
```

Name	Type	Thread ID	Value
[Add New Expression]			

Name	Type	Value
Arguments		
n	int	0x00000003 (3)

ID	Type	Stop	Location
1	Break	Group	dot
2	Break	Group	fact

Process: 1 (2943) python3.7-dbg Thread: 1.1 (2943) - Breakpoint Frame: fact File: ..jects/Python/Python Examples/tv_python_example.cpp Line: 5

Python with filtering

The screenshot shows a debugger interface with the following components:

- Top Bar:** File, Edit, Group, Process, Thread, Action Points, Bookmarks, Debug, Window, Help. A toolbar with a 'ReplayEngine' button is also present.
- Left Panel (Processes & T...):** A table showing process information.

Description	# P	# T	Members
python3.7-...	1	1	p1
Br...	1	1	p1
fact	1	1	p1.1
	1	1	p1.1
- Code Editor:** Displays C++ code for a factorial function. Line 5 is highlighted.

```
1 /* File: example.c */
2
3 #include "tv_example.h"
4
5 int fact(int n) {
6     if (n < 0){ /* This should probably return an error, but this is simpler */
7         return 0;
8     }
9     if (n == 0) {
10        return 1;
11    }
12    else {
13        /* testing for overflow would be a good idea here */
14        return n * fact(n-1);
15    }
16 }
17
18 int getSqaure(int n) {
19     return n * n;
20 }
```
- Call Stack:** Shows the current call stack.

Language	Function
C++	fact
Py	callFact
Py	pySupportedTypes
Py	<module>
C	main
	__libc_start_main
- Data View:** A table for monitoring data.

Name	Type	Thread ID	Value
[Add New Expression]			
- Action Points:** A table of breakpoints.

ID	Type	Stop	Location
1	Break	Group	dot
2	Break	Group	fact
- Local Variables:** Shows the current frame's variables.

Name	Type	Value
Arguments		
n	int	0x00000003 (3)
- Status Bar:** Process: 1 (2943) python3.7-dbg | Thread: 1.1 (2943) - Breakpoint | Frame: fact | File: ...jects/Python/Python Examples/tv_python_example.cpp | Line: 5



Reverse Debugging with ReplayEngine

Reverse Debugging with TotalView

- Reverse debugging provides the ability for developers to go back in execution history
- Activated either before program starts running or at some point after execution begins
- Capturing and deterministically replay execution
- Enables stepping backwards and forward by function, line, or instruction
- Run backwards to breakpoints
- Run backwards and stop when a variable changes value
- Saving recording files for later analysis or collaboration
- Reverse debugging is an add-on capability to TotalView

```
Start Page * common-main.c *
15 static void restore_sigpipe_to_default(void)
16 {
17     sigset_t unblock;
18
19     sigemptyset(&unblock);
20     sigaddset(&unblock, SIGPIPE);
21     sigprocmask(SIG_UNBLOCK, &unblock, NULL);
22     signal(SIGPIPE, SIG_DFL);
23 }
24
25 int main(int argc, const char **argv)
26 {
27     /*
28      * Always open file descriptors 0/1/2 to avoid clobbering files
29      * in die(). It also avoids messing up when the pipes are dup'ed
30      * onto stdin/stdout/stderr in the child processes we spawn.
31      */
32     sanitize_std fds();
33
34     git_setup_gettext();
35
36     git_extract_argv0_path(argv[0]);
37
38     restore_sigpipe_to_default();
39
40     return cmd_main(argc, argv);
41 }
42
```

Memory Debugging

Memory Debugging with TotalView

Memory Debugging Features

- Leak detection
- Dangling pointer detection
- Heap status
- Automatically detect allocation problems
- Memory Corruption Detection
- Memory Block Painting
- Memory Hoarding
- Lightweight memory block tracking

The screenshot displays the TotalView IDE interface with several panels open for memory debugging:

- Processes & Threads:** A table showing process details. The selected process is 'Process 1 (5381): filterapp-leaks'.
- Leak Report <p1>:** A table showing memory usage for the selected process.

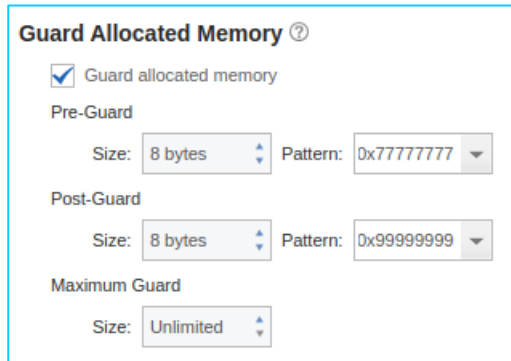
Process	Bytes	Count
Process 1 (5381): filterapp-leaks	2.69 MB	13680
myClassB.cxx	2.60 MB	13272
myClassA.cxx	69.00 KB	138
filterapp-leaks.cxx	16.36 KB	270
- Call Stack:** A list of stack frames, including 'open64', 'TV_thread_lock_recursive_file_lock', 'TV_thread_is_multi_threaded', 'TV_pthread_mutex_unlock', and 'malloc'.
- Code Editor:** Shows C++ source code for 'filterapp-leaks.cxx' with line numbers 222 to 249. The code includes a loop and a conditional statement.
- Action Points:** A table showing configured breakpoints.

ID	Type	Stop	Location	Line
1	Break	Process	...lterapp-leaks.cxx#344	... (line 344)
2	Break	Process	...lterapp-leaks.cxx#333	... (line 333)
- Standard Input:** A text input field with a 'SEND' button.

Guard Blocks

Guard allocated memory

When selected, the Memory Debugger writes guard blocks before and after a memory block that your program allocates



The screenshot shows a dialog box titled "Guard Allocated Memory" with a help icon. It contains a checked checkbox labeled "Guard allocated memory". Below this, there are three sections: "Pre-Guard" with "Size: 8 bytes" and "Pattern: 0x77777777"; "Post-Guard" with "Size: 8 bytes" and "Pattern: 0x99999999"; and "Maximum Guard" with "Size: Unlimited".

Pre-Guard and Post-Guard Size:

Sets the size in bytes of the block that the Memory Debugger places immediately before and after the memory block that your program allocates

Pattern:

Indicates the pattern that the Memory Debugger writes into guard blocks. The default values are `0x77777777` and `0x99999999`

Memory Painting

Memory Block Painting ⓘ
 Paint allocations Pattern:
 Paint deallocations Pattern:

- Useful to track down when you read uninitialized or undefined memory
 - Before it has been initialized
 - After it has been freed
- Memory Painting allows you to inject known values into memory upon allocation or deallocation
 - Good values don't correspond to any valid address
 - Have a distinctive look
 - When cast to different types

Memory Hoarding

Hoard Deallocated Memory ?

Hoard deallocated memory

Maximum KB to hoard:

Maximum blocks to hoard:

Automatically release blocks when memory gets low

Warn when hoard size drops below:

Hoard Status:

- **Memory Hoarding**
 - Stops the memory manager from reusing memory blocks
 - Can detect certain memory errors
- **Hoard Low Memory Controls**
 - Automatically release hoarded memory when available memory gets low, allowing your program to run longer
- **Hoard Low Memory events**
 - MemoryScape can stop execution as notification that the hoard dropped below a particular threshold. This provides an indication that the program is getting close to running out of memory.

Batch Debugging with TVScript

Batch Debugging with tvscript

- A straightforward language for unattended and/or batch debugging with TotalView and/or MemoryScape
- Usable whenever jobs need to be submitted or batched
- Can be used for automation
- A more powerful version of printf, no recompilation necessary between runs
- Schedule automated debug runs with *cron* jobs
- Use in continuous integration (CI)
- Expand its capabilities using Tcl

tvscript examples

Simple example

```
tvscript \  
-create_actionpoint "method1=>display_backtrace -show_arguments" \  
-create_actionpoint "method2#37=>display_backtrace \  
    -show_locals -level 1" \  
-event_action "error=>display_backtrace -show_arguments \  
    -show_locals" \  
-display_specifiers "noshow_pid,noshow_tid" \  
-maxruntime "00:00:30" \  
~/work/filterapp /filterapp -a 20
```

MPI example

```
tvscript -mpi "Open MPI" -tasks 4 \  
-create_actionpoint \  
"hello.c#14=>display_backtrace" \  
~/tests/MPI_hello
```

Debugging I/O Bottlenecks

Debugging I/O Bottlenecks

1. Identifying Stalled Processes (Hangs)

If your application is waiting on I/O, you can use TotalView to identify which processes are affected.

Attach to a Running Job: Use TotalView to attach to an already running, slow, or hung job. Once attached, click "Halt" to freeze the processes.

Examine Call Stacks: Look at the call stacks for various processes in the Process Window. Processes stuck in I/O will often have stacks showing functions related to file system calls (e.g., read, write, fsync) or MPI I/O functions.

Focus on Process Group: Use the "Processes and Threads" window to identify if all nodes are waiting for I/O (a potential collective I/O issue) or if only a few are (a potential load imbalance).

2. Using Memory Debugging for I/O-Related Memory Issues

I/O bottlenecks are sometimes caused by inefficient memory management, such as excessive buffering.

Detect Memory Leaks/Excessive Allocation: Use the Leak Report to see if memory is filling up, which can cause excessive paging to disk.

Debugging I/O Bottlenecks

3. ReplayEngine for Post-Mortem Analysis

If the I/O issue causes a crash or is hard to reproduce, you can use Reverse Debugging (ReplayEngine).

Record Execution: Start recording with ReplayEngine.

Step Backward: If a crash or hang occurs, step backward to identify the exact point where a large file was opened, written to, or where a communication stall initiated.

4. Combining with Performance Tools

TotalView is often used alongside performance tools like TAU profiler from ParaTools which provide detailed I/O statistics.

TotalView can help you connect directly to a job that has been flagged as slow by a profiler to see *exactly* what the code is doing at that moment.

Best Practices for HPC Debugging with TotalView

Best Practices for HPC Debugging with TotalView

- TotalView can't find the program source
 - Did you compile with -g ?
 - How to adjust the TotalView search paths? Preferences -> Search Path
- Python Debugging
 - Making sure proper system debug packages are installed for Python
- Understanding different ways to stop program execution with TotalView Action Points
 - Using a watchpoint on a local variable
- Focus
 - Diving on a variable that is no longer in scope. Check the Local Variables window for in scope variables
 - Totalview doesn't change focus to the thread hitting a breakpoint? Set Action Point Preferences to "Automatically focus on threads/processes at breakpoint"

Best Practices for HPC Debugging with TotalView

- MPI Debugging
 - Differences in launching MPI job from within the TotalView UI vs the command line.
 - TotalView runs an MPI program without stopping? Set the Parallel Preferences to “Ask What To Do” in After Attach Behavior
 - Using wrong attributes in processes and threads view
- Reverse Debugging
 - Running out of memory by not setting the maximum memory allocated to ReplayEngine
 - Defer turning on reverse debugging until later in program execution to avoid slow initialization phases
 - Adjust reverse debugging circular buffer size to reduce resources
- Memory Debugging
 - Starting with All memory debugging options enabled rather than Low
 - Not setting a size restriction for Red Zones
 - Issues with getting memory debugging turned on in an MPI job? May have to set LD_PRELOAD environment variable or worst case, prelink HIA

Best Practices for HPC Debugging with TotalView

- Reverse Connect with tvconnect
 - When I use Reverse Connect I get the following obscure message: *myProgram is an invalid or incompatible executable file format for the target platform*
 - The message indicates an incompatible file format but most often this occurs if the program provided to tvconnect for TotalView to debug cannot be found. The easiest way to resolve problem is to provide the full path to the target application, e.g., `tvconnect /home/usr/myProgram`
- How do I get help?
 - How to submit a support ticket? support-TotalView@perforce.com
 - Where is TV documentation (locally and on the internet). <https://help.totalview.io/>
 - Are there videos I can watch to learn how to use TotalView? <https://perforce.com/resources/videos>

Resources and Documentation

Resources and Documentation

PERFORCE

TotalView PRODUCTS & SOLUTIONS RESOURCES CUSTOMERS SUPPORT TRY FREE

The Most Advanced Debugger for HPC

Debug HPC Applications Written in C, C++, Fortran, and Python

TRY FREE

Why Do Top HPC Developers Use TotalView for Debugging Code?

You need special tools for multithreaded, multiprocess, and GPU-specific applications. TotalView is a powerful debugging solution that meets the unique and demanding requirements of HPC developers.

See why industry leaders use TotalView to get unprecedented HPC code visibility and control.

SEE CASE STUDIES

SEND FEEDBACK

Visit totalview.io for more information

DOCUMENTATION
<https://help.totalview.io/>

VIDEO TUTORIALS
<https://www.perforce.com/resources/videos>

BLOG
<https://www.perforce.com/blog>

Thank you

Name: Dean Stewart

Email: dstewart@perforce.com

Name: Bill Burns

Email: bburns@perforce.com

Name: Suzanne Horn

Email: Susanne.Horn@SMB-Net.de