INTEL® PARALLEL STUDIO XE 2016 CLUSTER EDITION

For Distributed Performance

Intel® Parallel Studio XE 2016 development suite - Empowering Faster Code Faster

Delivering HPC Development Solutions

- Over 20 years
- Industry Collaboration on Standards
- Developed with Performance & Scaling with Intel hardware

Meeting the Challenges

- Boosting Performance
- Increasing Scalability
- Increasing Productivity

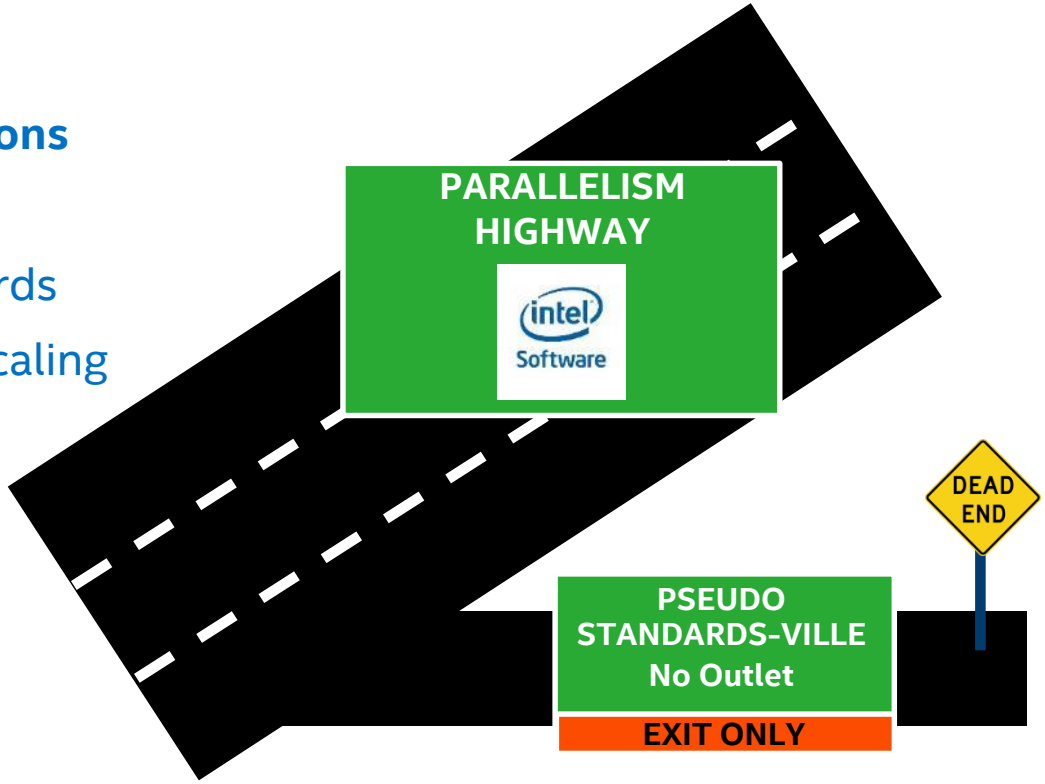

How Intel® Parallel Studio XE 2016 helps make ***Faster Code Faster*** for HPC

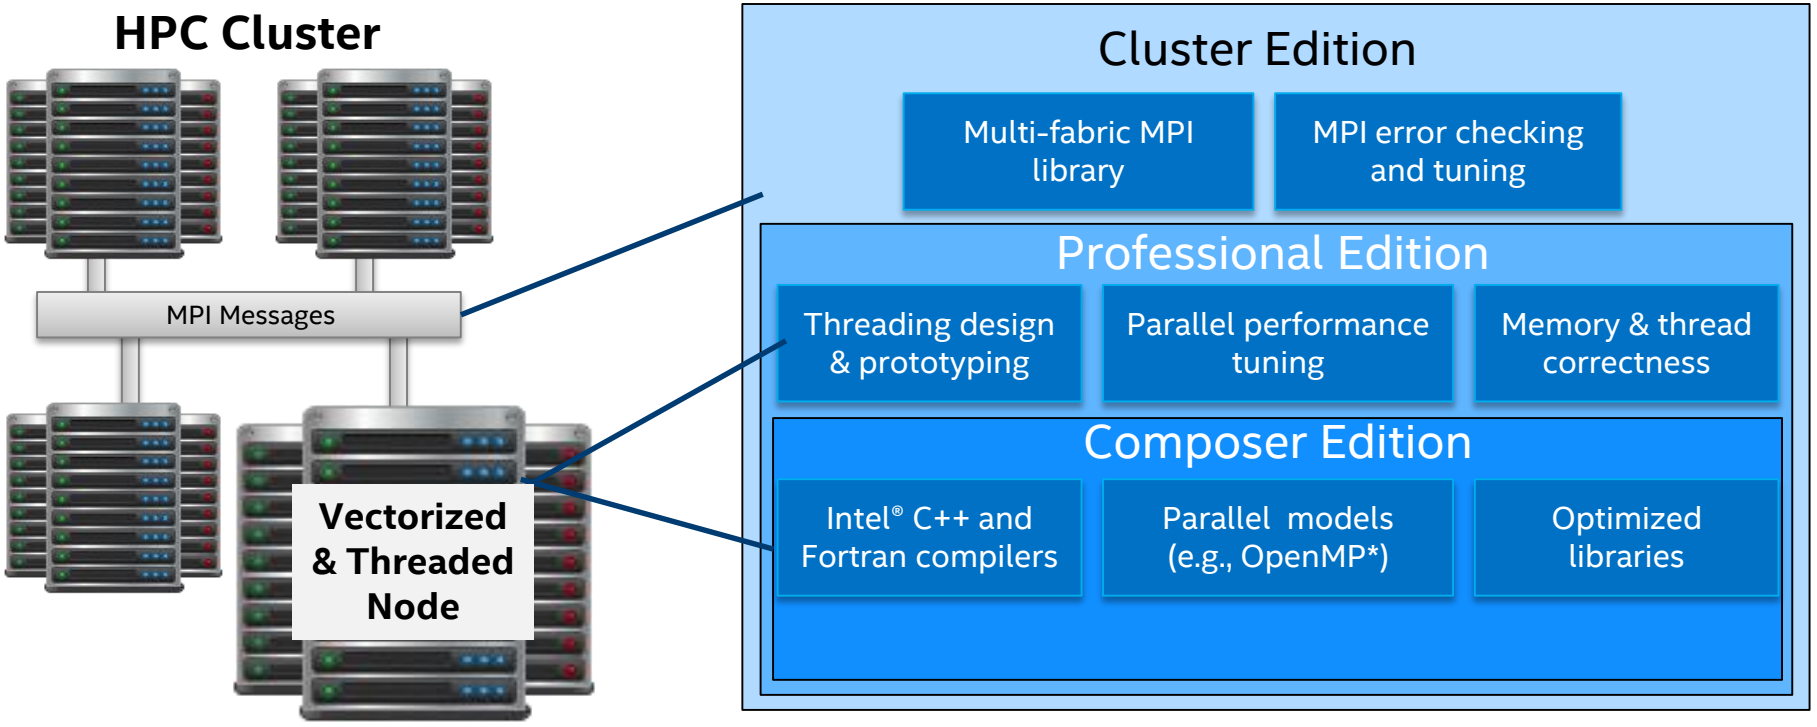

Intel® MPI Library

Value Proposition

What	• **Intel's High Performance MPI Library**
Why	• **Scale Performance –** Tuned for Latest Intel Architectures • **Scale Forward –** Multicore and Manycore Ready • **Scale Efficiently –** Flexible Fabric Selection & Compatibility
How	• **Standards Based –** Built on Open Source MPICH Implementation • **Sustained Scalability –** Tuning for Low Latencies, High Bandwidth & Increased Processes • **Multi Fabric Support –** Supports Popular High Performance Networking Fabrics

Intel® MPI Library Overview

Optimized MPI application performance
- Application-specific tuning
- Automatic tuning

Lower latency and multi-vendor interoperability
- Industry leading latency
- Performance optimized support for the latest OFED capabilities through DAPL 2.x

Faster MPI communication
- Optimized collectives

Sustainable scalability beyond 262K cores
- Native InfiniBand* interface support allows for lower latencies, higher bandwidth, and reduced memory requirements

More robust MPI applications
- Seamless interoperability with Intel® Trace Analyzer and Collector

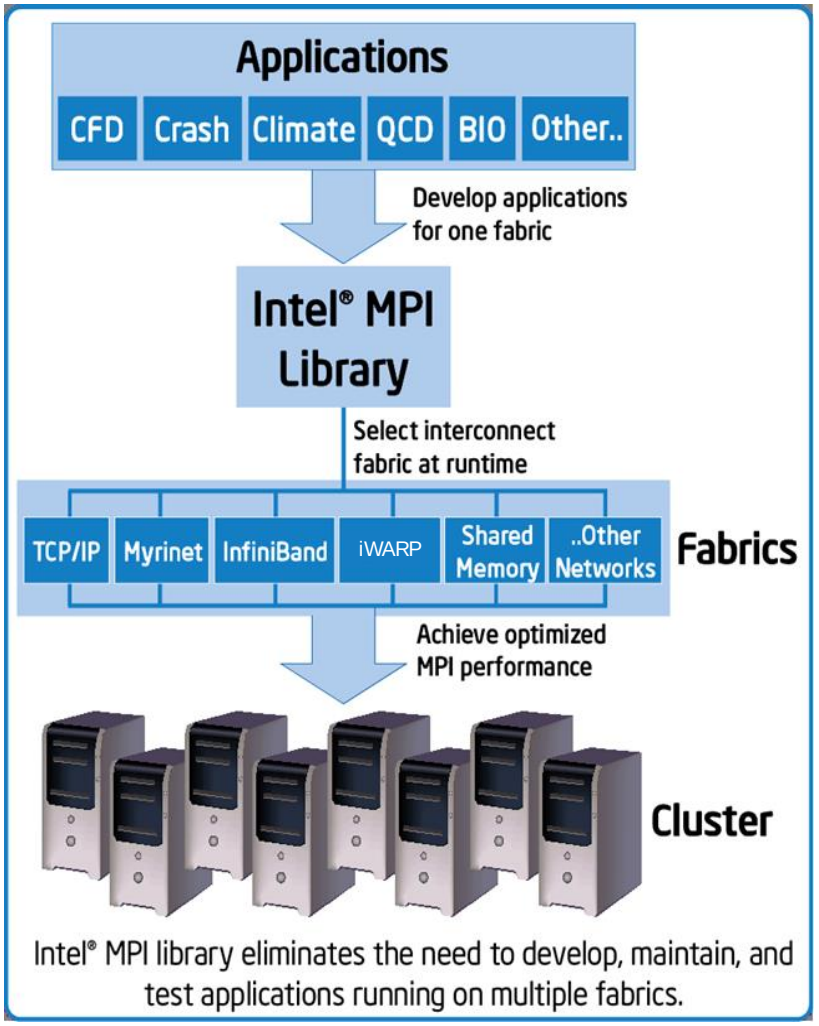

Superior MPI performance on Linux* 64

288 Processes, 8 nodes (InfiniBand + shared memory)

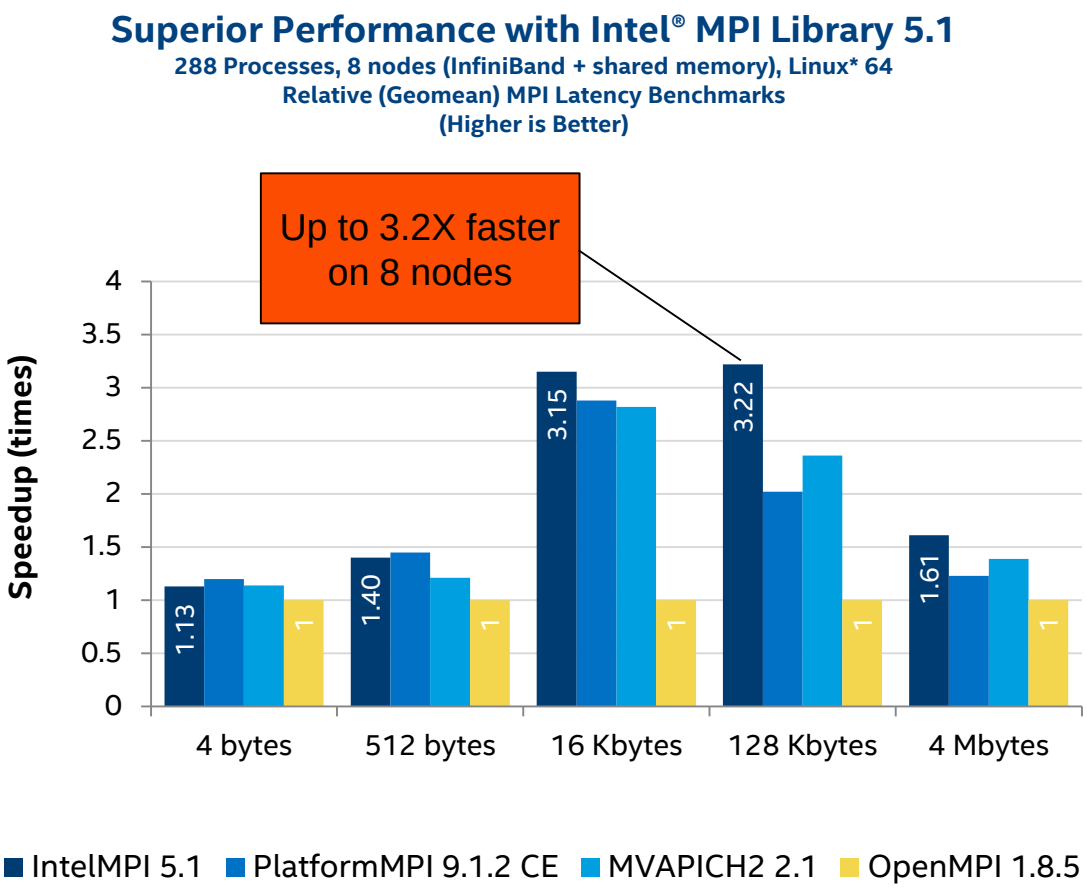

Configuration Info:

Hardware: CPU: Dual Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz; 64 GB RAM. Interconnect: Mellanox Technologies MT27600 [Connect-IB].

Software: RedHat* RHEL 6.5; OFED 3.12-1; Intel® MPI Library 5.1; Intel® MPI Benchmarks 4.1 (built with Intel® C++ Compiler XE 15.0.3 for Linux*);

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

Superior MPI performance on Linux* 64

1792 Processes, 64 nodes (InfiniBand + shared memory)

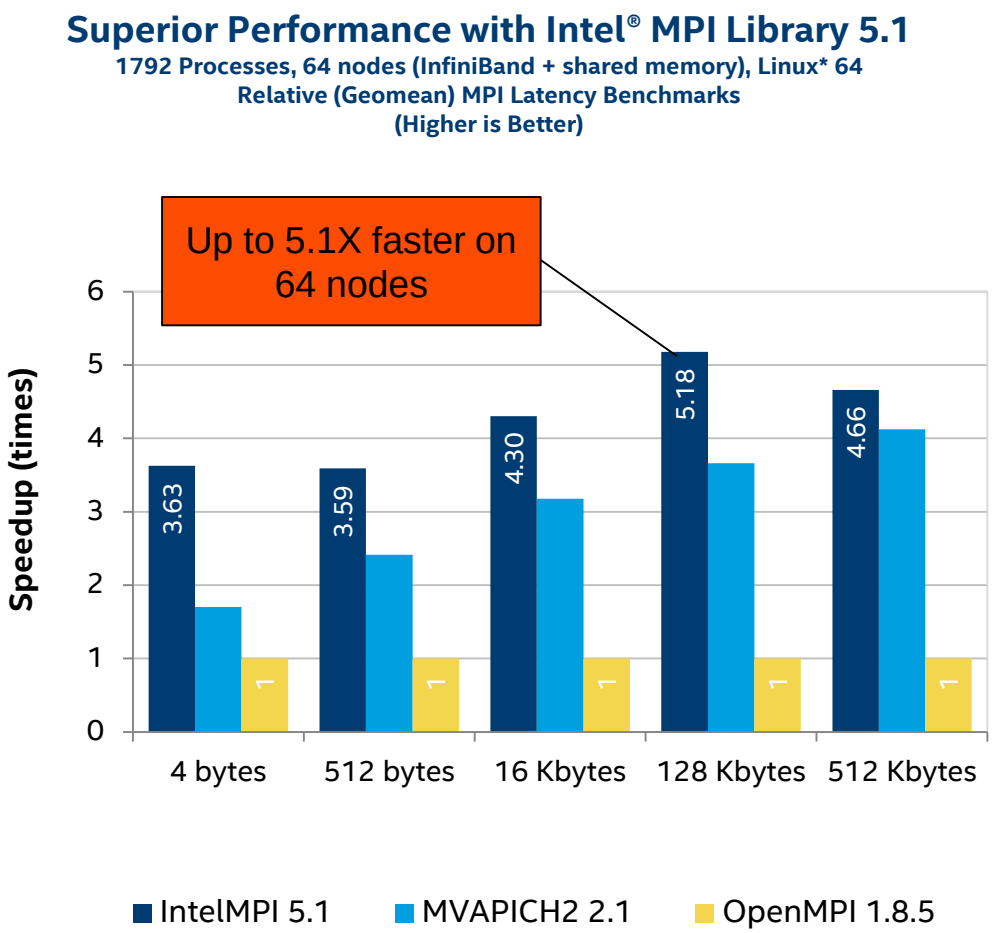

Configuration Info:

Hardware: CPU: Dual Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz; 64 GB RAM. Interconnect: Mellanox Technologies MT27500 Family [ConnectX-3].

Software: RedHat* RHEL 6.5; OFED 3.5-2; Intel® MPI Library 5.1 Intel® MPI Benchmarks 4.1 (Built with Intel® C++ Compiler XE 15.0.3 for Linux*);

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

Superior MPI performance on Windows* 64

96 Processes, 8 nodes (IP over InfiniBand + shared memory)

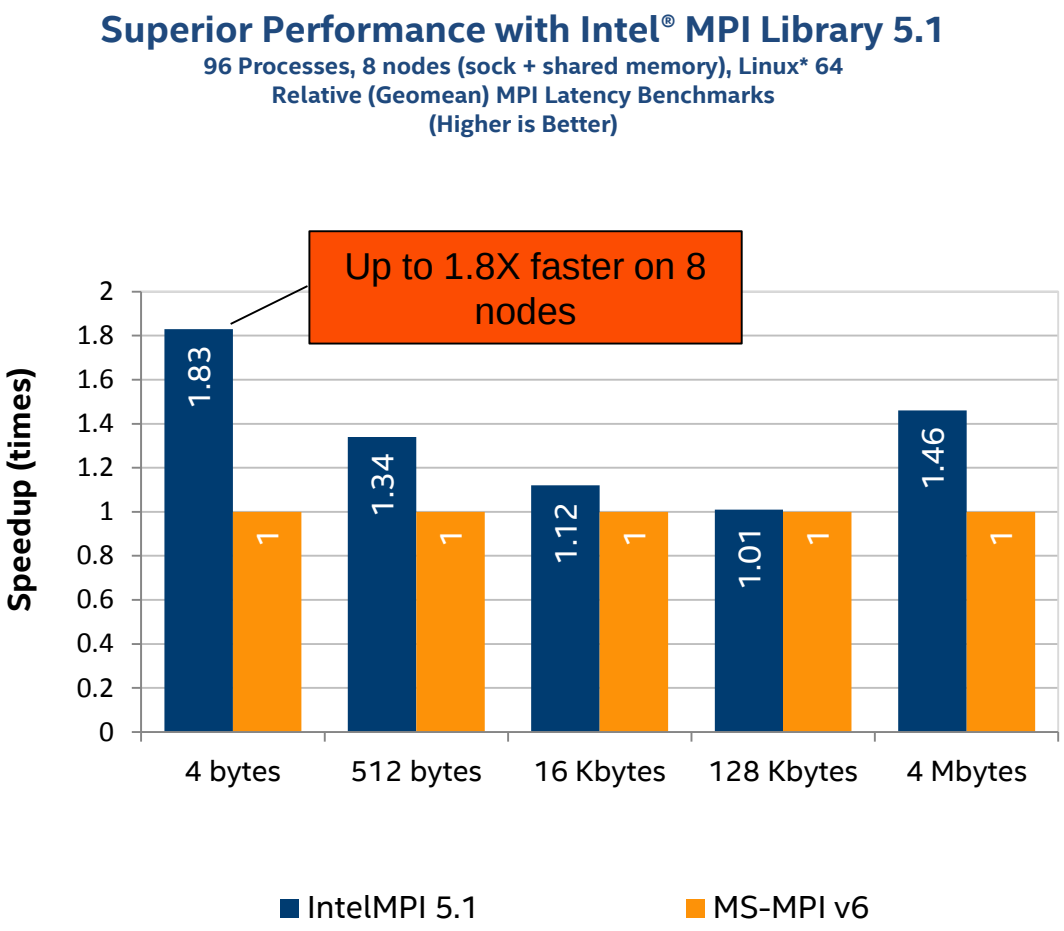

Configuration Info:

HOST Hardware : Intel® Xeon® CPU X5680 @ 3.33GHz, RAM 24GB; Interconnect: InfiniBand, Mellanox ConnectX VPI (MT26428) QDR;

Software: Microsoft Windows Server 2008 HPC Edition, OFED 3.2, Intel® C/C++ Compiler XE 15.0.3, Intel® MPI Benchmarks 4.1.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

Superior MPI performance on Xeon Phi™

64 Processes, 8 nodes (InfiniBand + shared memory)

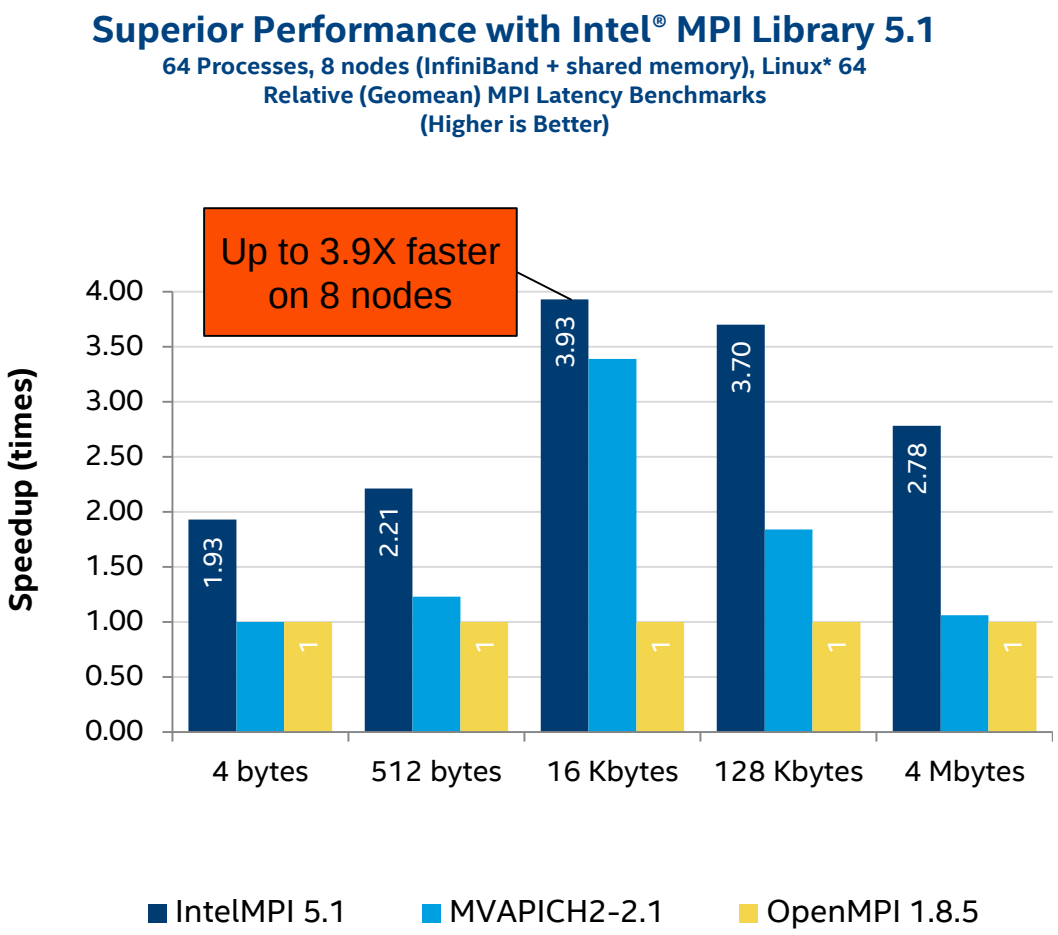

Configuration Info:

HOST Hardware : Intel® Xeon® CPU E5-2680 @ 2.70GHz, RAM 64GB; Interconnect: InfiniBand, Mellanox Technologies MT27500 Family [ConnectX-3];

MIC Hardware: Intel® Xeon Phi™ coprocessor SE10/7120 series (rev 20) 1238095 kHz; 61 cores. RAM: 15872 MB per card.

Software: RHEL 6.5, OFED 3.12-1, MPSS Version: 3.5, Intel® C/C++ Compiler XE 15.0.3, Intel® MPI Benchmarks 4.1.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. * Other brands and names are the property of their respective owners. Benchmark Source: Intel Corporation.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

Intel® MPI Library 5.1
What's New

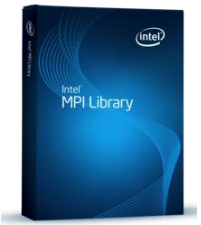

Ease of use

- Brand new Troubleshooting section for quicker issue resolution

Performance & Scaling

- Support for OpenFabrics* Interface (OFI*) v1.0
- Additional features for faster automatic tuning

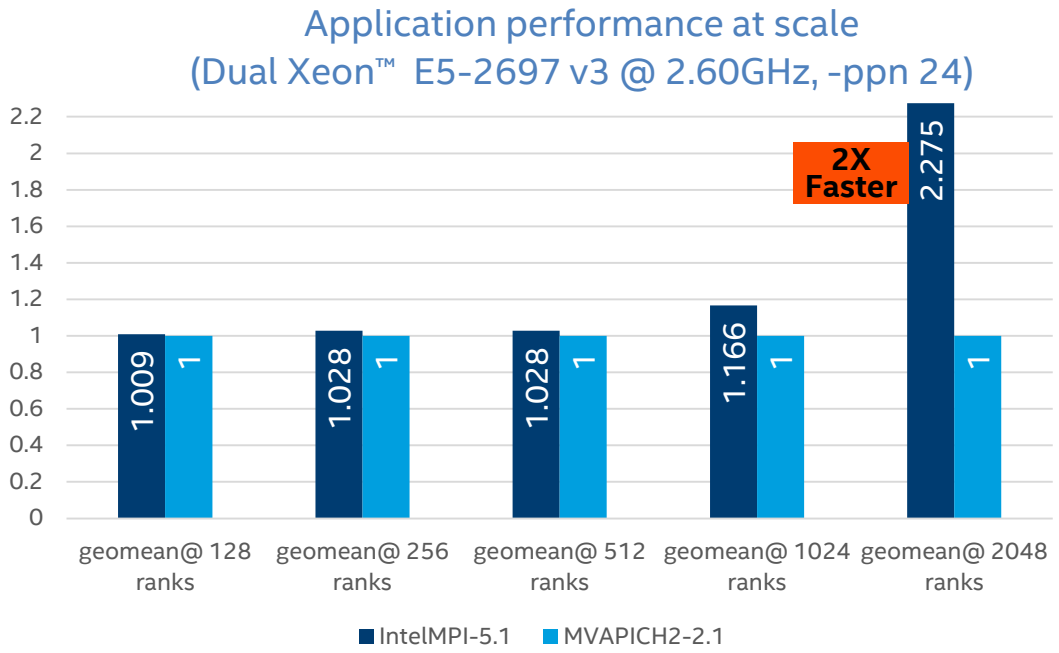

Based on the following set of tests: SpecMPI-2.0/104.milc, SpecMPI-2.0/127.wrf2 (Built with Intel® C++ Compiler XE 15.0.3.187 for Linux*);
Hardware: CPU: Dual Intel® Xeon E5-2697v3@2.60Ghz; 64 GB RAM. Interconnect: Mellanox Technologies* MT27500 Family [ConnectX*-3] FDR.
Software: RedHat* RHEL 6.5; OFED 3.5-2-MIC-rc3; Intel® MPI Library 5.1, MVAPICH2-2.1.
Environment: MV2_IBA_HCA=mlx4_0 (mvapich2 only).

Performance Tuning Tools for Distributed Applications

Intel® Trace Analyzer and Collector

Tune cross-node MPI

- Visualize MPI behavior
- Evaluate MPI load balancing
- Find communication hotspots

Intel® VTune™ Amplifier XE

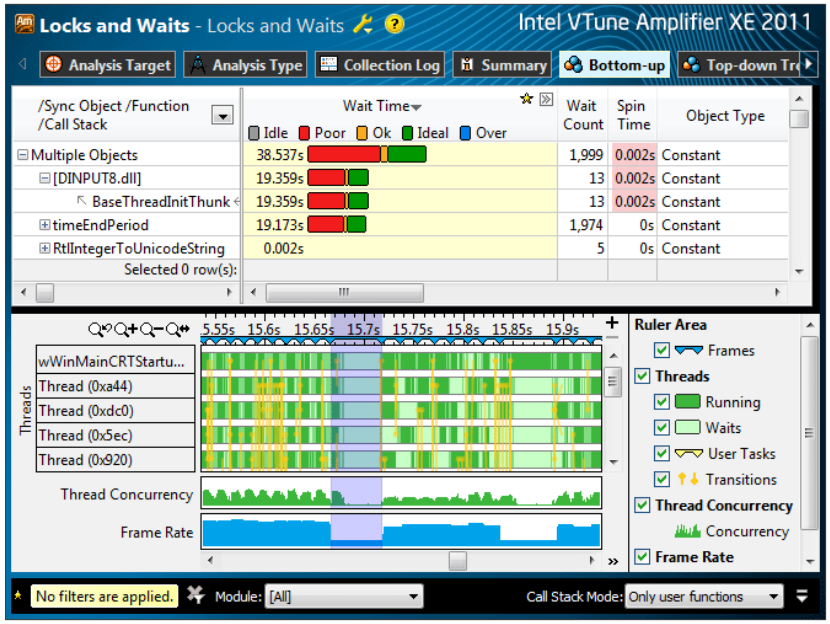

Tune single node threading

- Visualize thread behavior
- Evaluate thread load balancing
- Find thread sync bottlenecks

Intel® Trace Analyzer and Collector

Value Proposition

What	• **Intel's High Performance MPI Communications Profiler & Analyzer for Scalable HPC Development**
Why	• **Scale Performance –** Perform on More Nodes • **Scale Forward –** Multicore and Manycore Ready • **Scale Efficiently –** Tune & Debug on More Nodes
How	• **Visualize** - Understand parallel application behavior • **Evaluate -** Profiling statistics and load balancing • **Analyze** – Automated analysis of common MPI issues • **Identify –** Communication hotspots

Intel® Trace Analyzer and Collector Overview

Intel® Trace Analyzer and Collector helps the developer:

- Visualize and understand parallel application beh
- Evalua nd load balancing
- Identify communication hotspots

Features

- Event-based approach
- Low overhead
- Excellent scalability
- Powerful aggregation and filte
- Performance Assistance and Imbalance Tuning
- NEW in 9.1: MPI Performance Snapshot

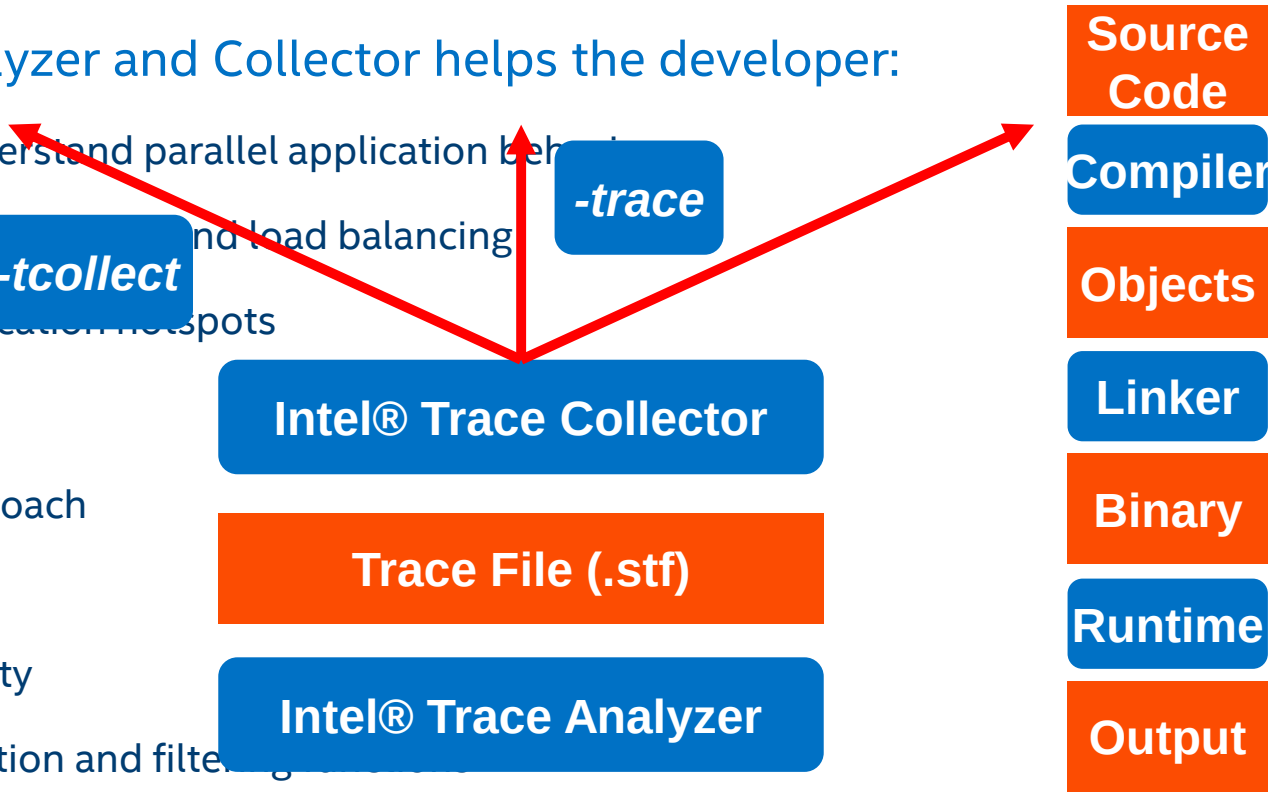

Using the Intel® Trace Analyzer and Collector is … Easy!

Step 1

Run your binary and create a tracefile

```
$ mpirun -trace -n 2 ./test
```

Step 2

View the Results:

```
$ traceanalyzer &
```

Intel® Trace Analyzer and Collector

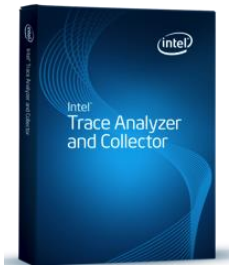

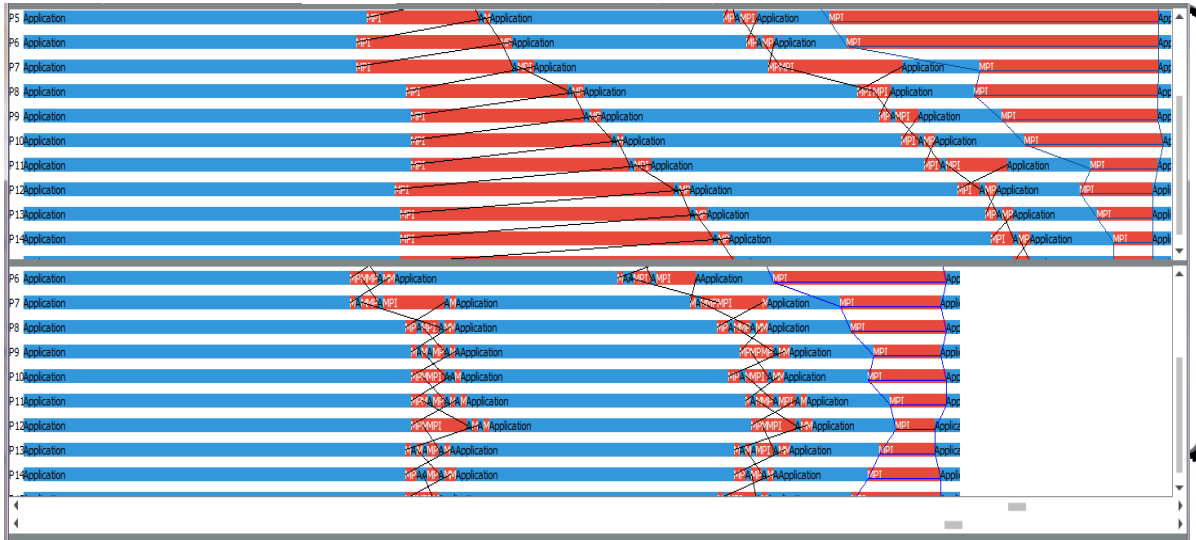

Compare the event timelines of two communication profiles

Blue = computation
Red = communication

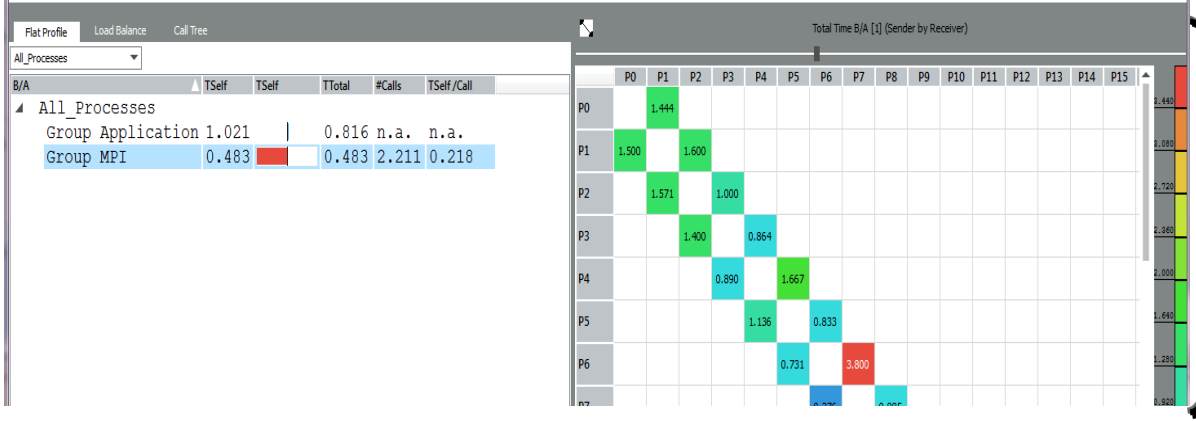

Chart showing how the MPI processes interact

Improving Load Balance: Real World Case

Too high load on Host
= too low load on coprocessor

Collapsed data per node and coprocessor card

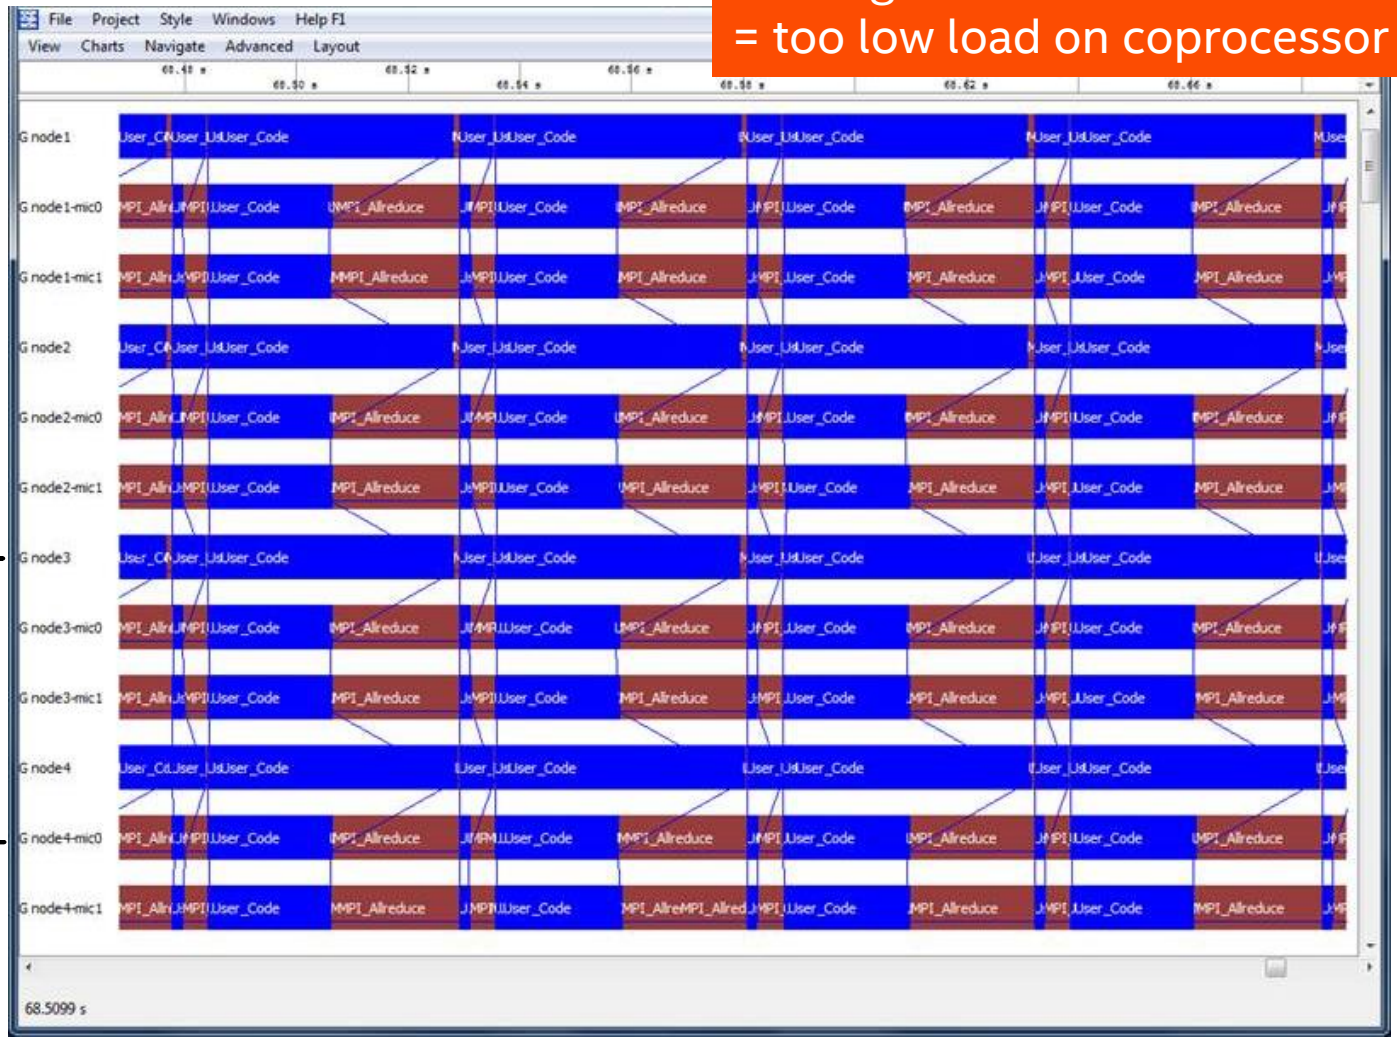

Host
16 MPI procs x
1 OpenMP thread

Coprocessor
8 MPI procs x
28 OpenMP threads

Improving Load Balance: Real World Case

Too low load on Host
= too high load on coprocessor

Collapsed data per node and coprocessor card

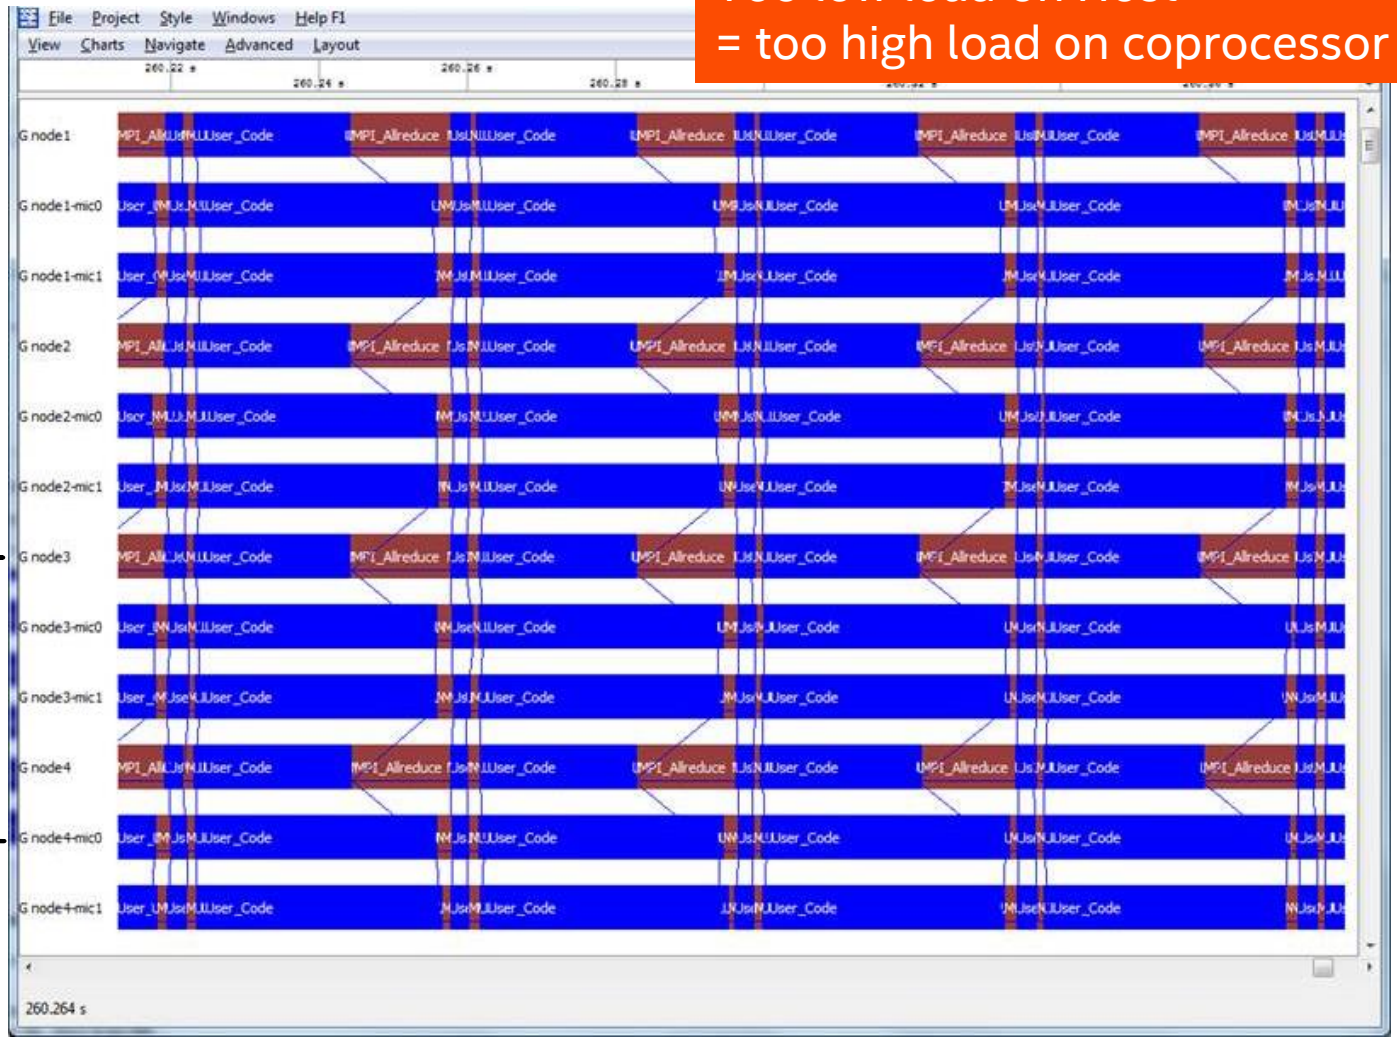

Host
16 MPI procs x
1 OpenMP thread

Coprocessor
24 MPI procs x
8 OpenMP threads

Improving Load Balance: Real World Case

Perfect balance
Host load = Coprocessor load

Collapsed data per node and coprocessor card

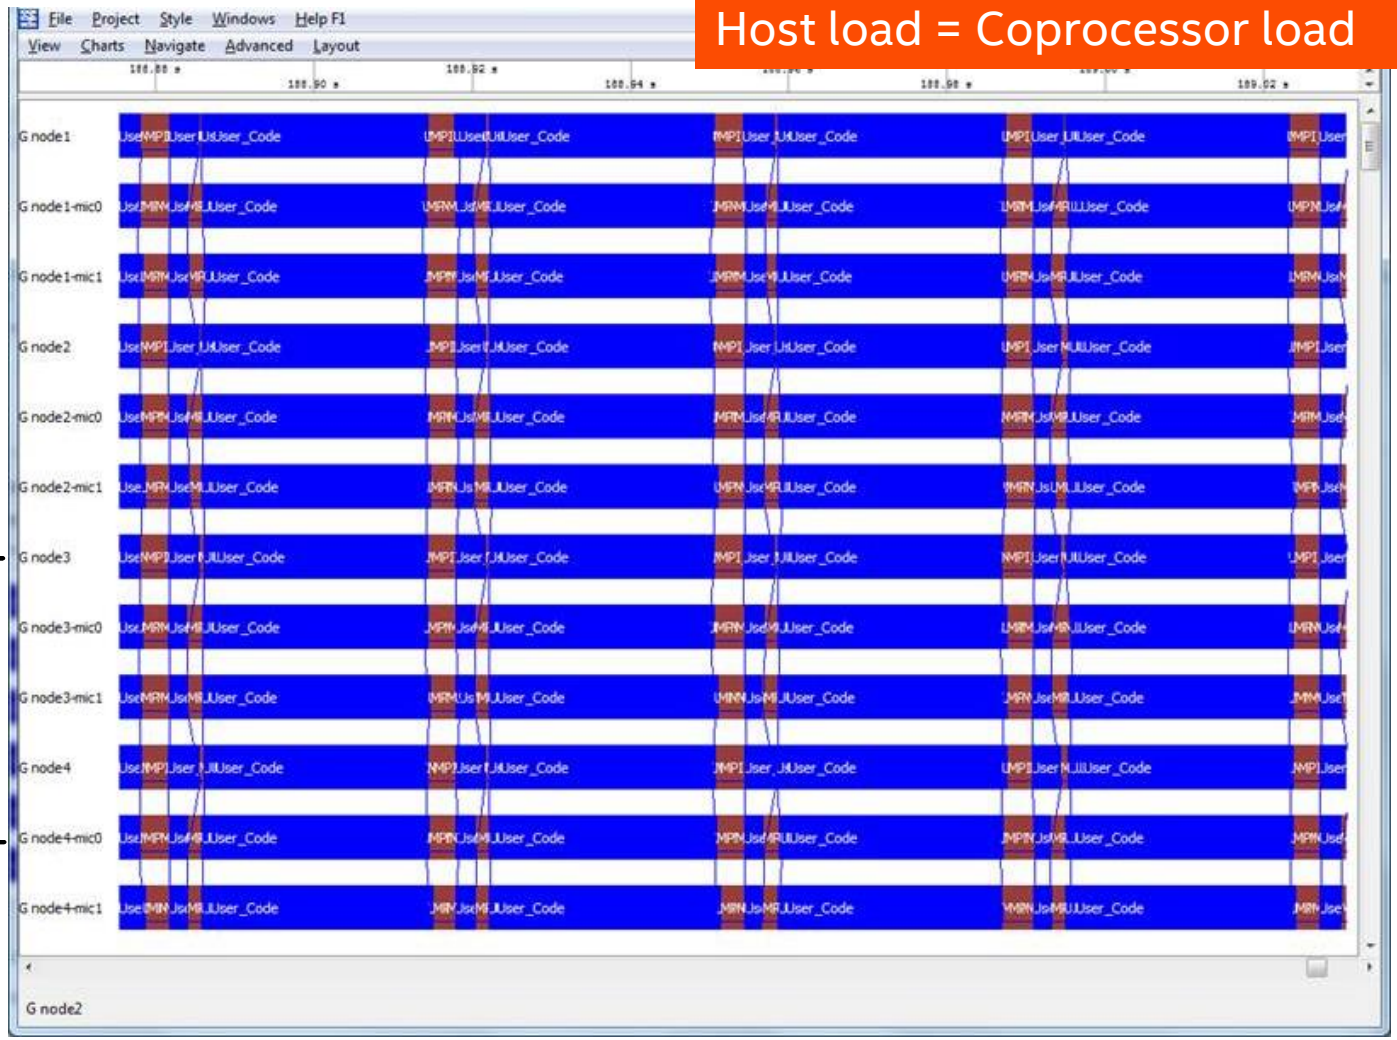

Host
16 MPI procs x
1 OpenMP thread

Coprocessor
16 MPI procs x
12 OpenMP thrds

Ideal Interconnect Simulator (Idealizer)

Helps to figure out application's imbalance simulating its behavior in the "ideal communication environment"

Actual trace

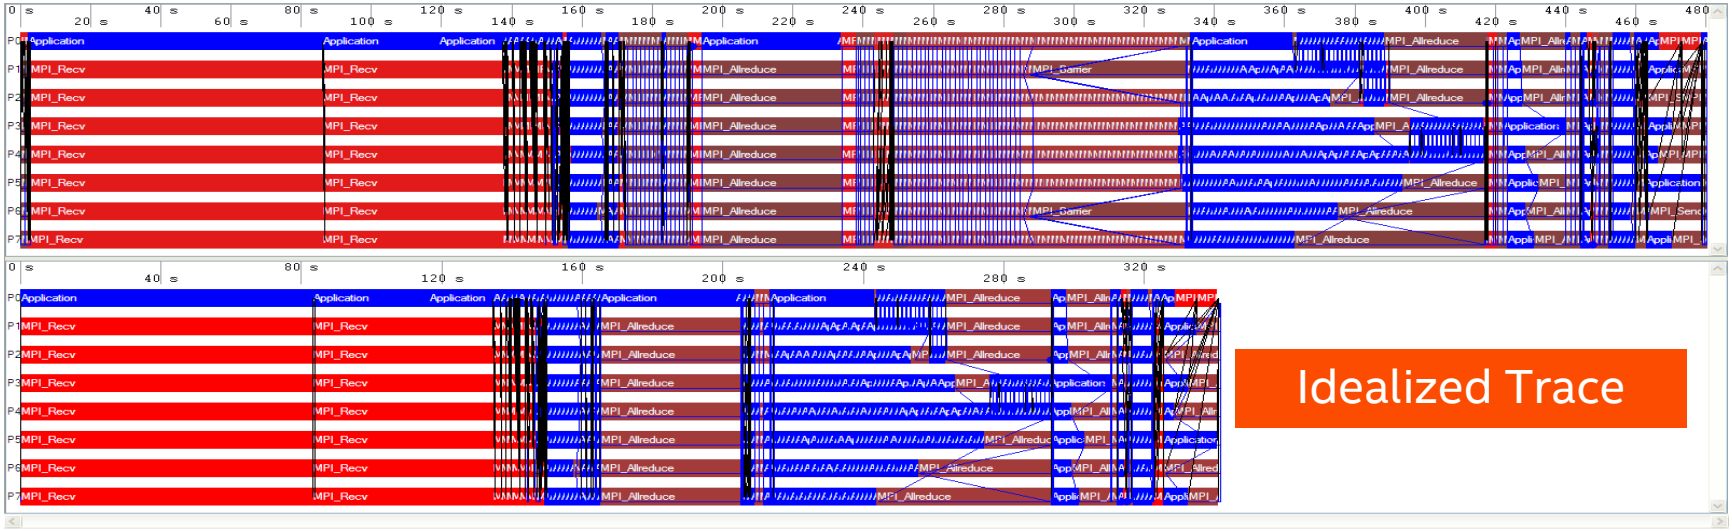

Idealized Trace

Easy way to identify application bottlenecks

MPI Performance Assistance

Automatic Performance Assistant

- Detect common MPI performance issues
- Automated tips on potential solutions

Automatically detect performance issues and their impact on runtime

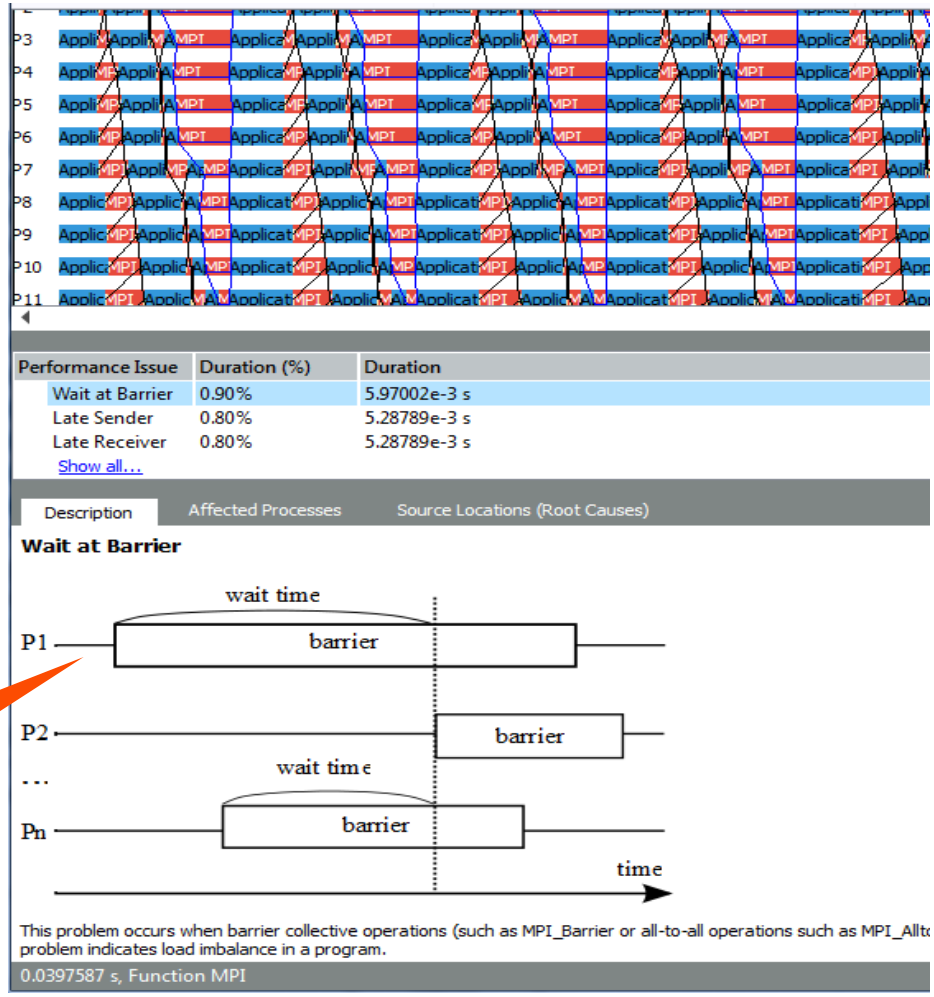

MPI Performance Snapshot
High capacity MPI profiler

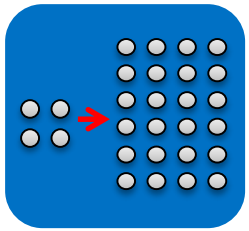

Lightweight
Low overhead profiling for 100K+ Ranks

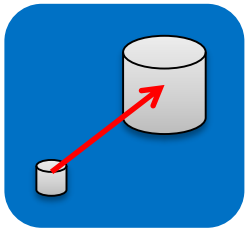

Scalability
Performance variation at scale can be detected sooner

Identifying Key Metrics
Shows PAPI counters and MPI/OpenMP imbalances

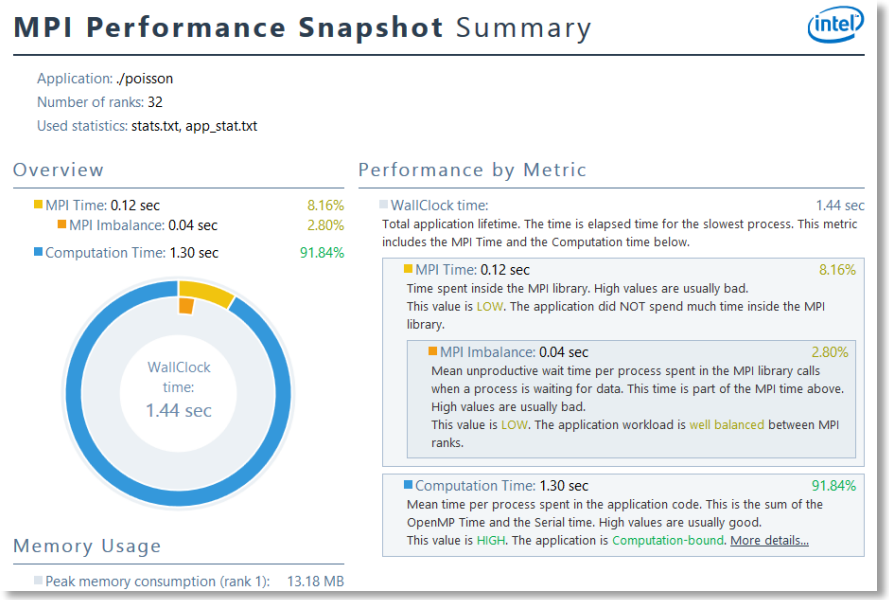

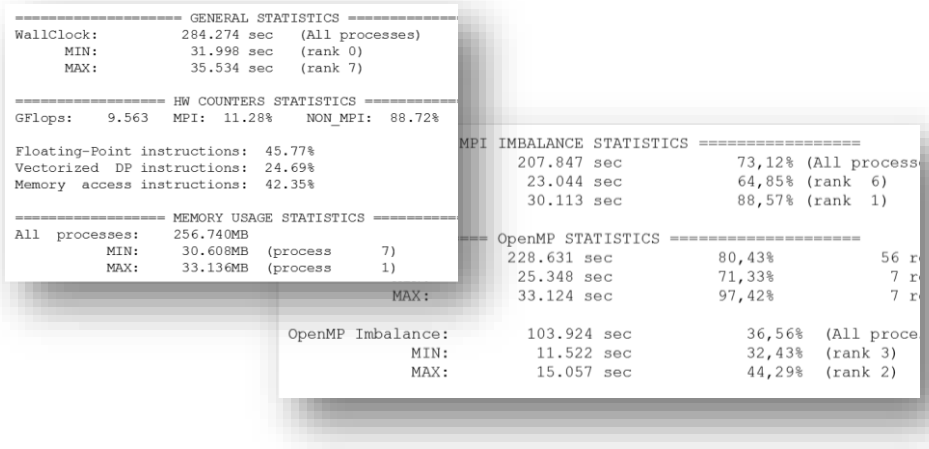

MPI Correctness Checking

Highlights: Checks and pin-point hard to find run-time errors

- Unique feature to identify run-time errors
- Displays the correctness (parameter passing) of MPI communication for more robust and reliable MPI based HPC applications

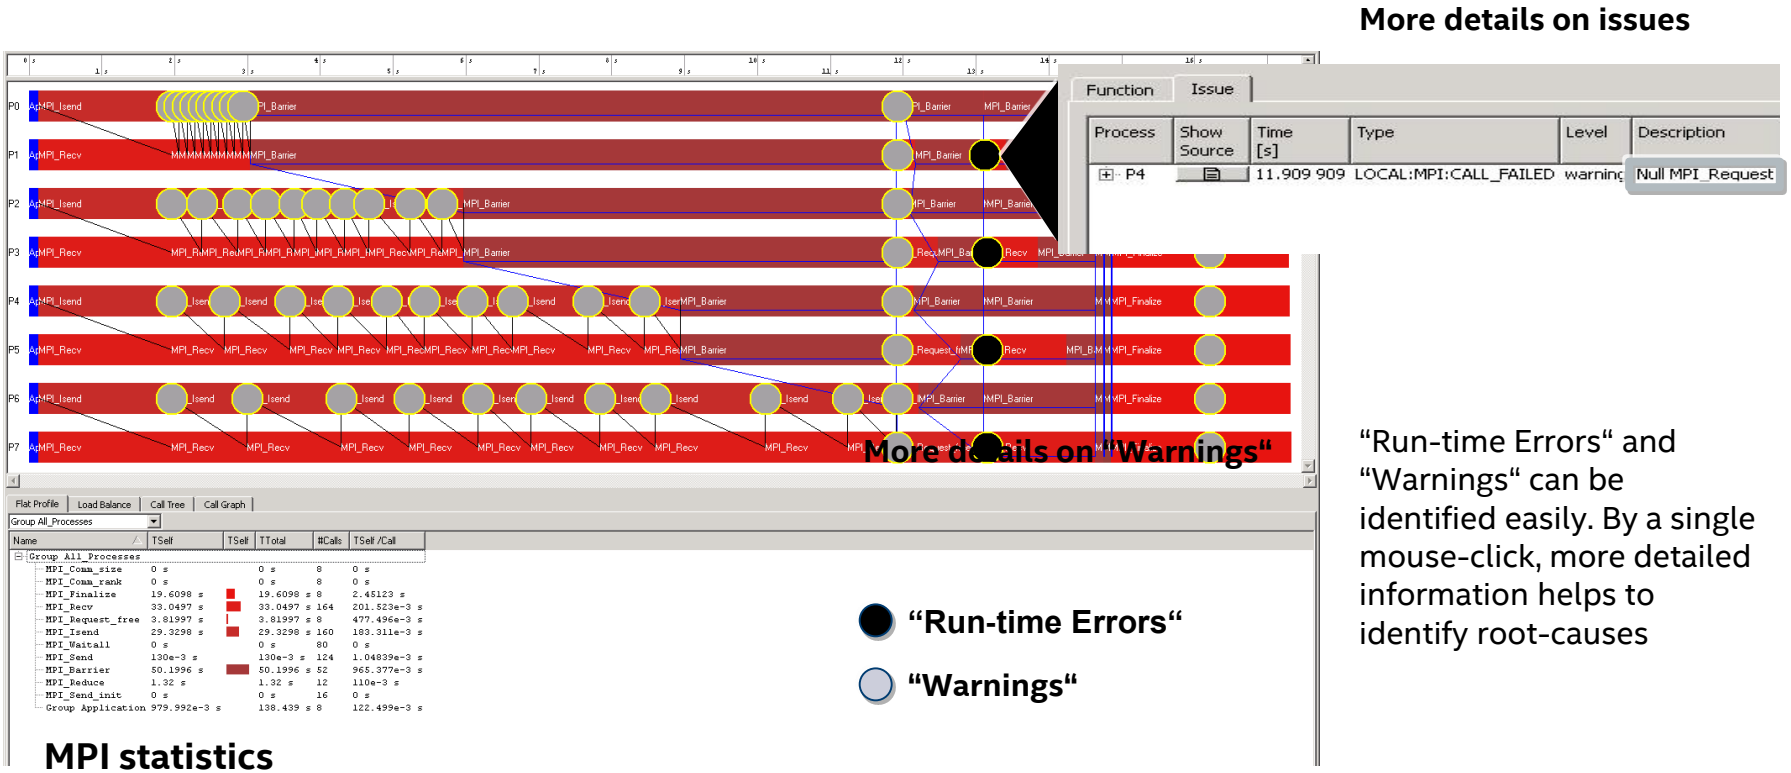

"Run-time Errors" and "Warnings" can be identified easily. By a single mouse-click, more detailed information helps to identify root-causes

Intel® MPI Benchmarks 4.1

Standard benchmarks with OSI-compatible CPL license

- Enables testing of interconnects, systems, and MPI implementations
- Comprehensive set of MPI kernels that provide performance measurements for:
  - Point-to-point message-passing
  - Global data movement and computation routines
  - One-sided communications
  - File I/O
  - Supports MPI-1.x, MPI-2.x, and MPI-3.x standards

What's New:

Introduction of new benchmarks

- Measure cumulative bandwidth and message rate values

The Intel® MPI Benchmarks provide a simple and easy way to measure MPI performance on your cluster

Online Resources

Intel® MPI Library product page

- www.intel.com/go/mpi

Intel® Trace Analyzer and Collector product page

- www.intel.com/go/traceanalyzer

Intel® Clusters and HPC Technology forums

- http://software.intel.com/en-us/forums/intel-clusters-and-hpc-technology

Intel® Xeon Phi™ Coprocessor Developer Community

- http://software.intel.com/en-us/mic-developer

Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Optimization Notice

Intel’s compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804


intel®