



## Platform MPI

### Performance, Simplicity and Reliability

#### Highlights:

- Workload manager, hardware, and OS-independent
- Fully integrated with Platform LSF – resulting in the only fully supported workload management + MPI solution on the market today
- Only commercial MPI implementation not based on an open source implementation
- Multi-core optimization
- Checkpoint / restart
- Suspend / resume

#### Benefits:

- Lowers the total cost of ownership through standardization and economies of scale.
- Reduces business risk through professionally developed and supported commercial-grade software.
- Accelerates the time to results through increased performance, scalability, stability and reliability.
- Consistently outperforms other MPI implementations and delivers higher application node-count scalability
- Provides a single consistent interface for interconnects, operating systems and architectures reducing the complexity of application development, testing, verification and support.
- Increases throughput for systems with redundant InfiniBand HCAs
- No need to recompile applications with other compatible MPI software.
- 40% performance improvement in CFD applications based on industry standard benchmarks.  
<http://www.spec.org/mpi2007/results/>

#### Ideal for:

- Enterprises that develop and deploy in-house applications
- Independent software vendors that want to improve the performance of their applications
- Application specialists who are looking for increased throughput

#### Enhance HPC application performance

Platform MPI is a fully integrated message passing interface (MPI) solution that enables users to take advantage of the leading interconnect technologies to build high performance applications, while simplifying the number of binary distributions required. Platform MPI communication libraries result in applications that leverage the latest high performance technologies while simplifying the development complexity, the testing challenges, and the support burden of multiple software distributions.

#### Simplicity

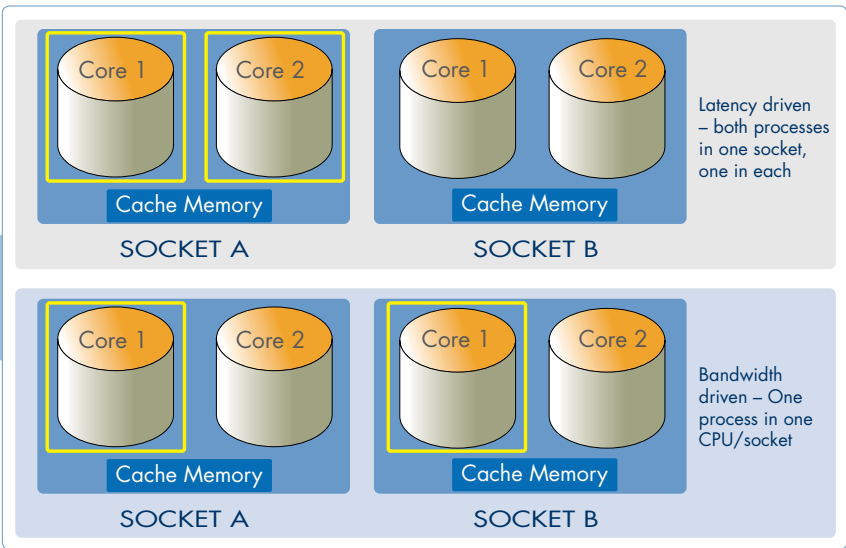
High performance parallel application development is complicated enough without the additional burden of managing support for a growing list of hardware architectures that leverage different interconnects. Platform MPI simplifies development, testing, and maintenance through its unique interconnect-independent architecture.

Through the use of dynamic runtime selectable libraries, developers are able to distribute a single binary executable that supports a full range of end user interconnects. This reduces the number of binaries that have to be produced, tested, maintained, and distributed to customers, while increasing the overall reliability of the application.

Comparing Platform MPI and Other Leading MPI Software using SPEC MPI2007 on Dell cluster

Application		Number of Ranks	Other MPI		Platform MPI		Throughput improvement when using Platform MPI over other-MPI
			Seconds	Score	Seconds	Score	
Fluid Dynamics	115.fds4	64	550	3.55	388	5.03	142%
Quantum Chromodynamics	104.milc	64	508	3.08	387	4.05	131%
Heat Transfer	128.GAPgeofem	64	446	4.63	345	5.99	129%
Weather Forecasting	127.wrf2	64	1362	5.72	1066	7.32	128%
Hydrodynamics	132.zeusmp2	64	657	4.72	537	5.78	122%
Hydrodynamics	129.tera_if	64	643	4.3	548	5.05	117%
Matrix Decomposition	137.lu	64	968	3.8	836	4.4	116%
Fluid Dynamics	107.leslie3d	64	1328	3.93	1182	4.42	112%
Molecular Dynamics	126.lammps	64	802	3.63	744	3.92	108%
Oceanography	121.pop2	64	1715	2.41	1596	2.59	107%
Quantum Chemistry	130.socorro	64	1061	3.6	1004	3.8	106%
Ray Tracing	122.tachyon	64	587	4.77	581	4.82	101%
Computational Electromagnetics	113.GemsFDTD	64	1077	5.86	1065	5.92	101%
<b>SPEC MPI™_base 2007</b>			<b>4.04</b>		<b>4.71</b>		<b>117%</b>

Policy-based process-to-core affinity



## Key Features

### Multi-core optimization

Platform MPI intelligently distributes processes across cores within a system based on policies delivering up to 35% performance improvement. Platform MPI has specific optimizations and tuning parameters to maximize performance for multi-core environments. This allows applications to leverage the right balance of processors, memory availability, and network capacity to achieve optimal performance. In conjunction with workload manager core balancing, a doubling of throughput has been demonstrated with a crash test benchmark (LS-Dyna, neon\_revised\_revisited).

### Suspend Resume

Platform MPI has the ability to allow the workload manager to suspend a running job so that a higher priority job can execute on those resources. At the completion of this higher priority job, the workload manager can then resume the suspended job on these same resources.

### Checkpoint restart

Checkpoint restart supports pre-emptive scheduling of jobs in the cluster. Regular checkpoints can be performed on long-running, but lower priority jobs. When a higher priority job enters the queue, the lower priority job can be stopped, freeing up resources to perform the higher priority job. Later, the suspended job can be restarted when resources become available again. In a future release, this feature will also protect long-running jobs from failure – in this case, the jobs files will also be checkpointed, and the job can be migrated to an available resource and restarted from the last checkpoint.

### Automatic selection of optimal network

Through the use of a priority network list built from system configuration files, user environment variables, user command line options and library hard-coded defaults, Platform MPI can dynamically select the optimal network connection between each node and each other node within a cluster at runtime.

### Multithread safe

Multithreaded applications can fully exploit Platform MPI and multiple threads can simultaneously request services and conduct communication.

### Linux® command line replication

Command line arguments to the application are automatically provided to all MPI processes, avoiding tedious parsing and broadcasting of parameters to other MPI processes.

### MIMD support

The Multiple Instruction – Multiple Data (MIMD) model is supported through provisions that launch different executables which constitute the whole MPI application.

### Tracing and monitoring

MPI-related monitoring presentation can be selected through environment variables to determine presentation of timing and trace information. Regular expressions, single or groups of functions and other options can be selected for monitoring.

### Support for popular debuggers

Platform MPI fully supports Etnus TotalView® analysis tools, Allinea's distributed debugging tool (DDT), and standard GNU gdb.

### Total Platform Solution

Platform offers a total solution for high performance computing, which includes Platform Manager for cluster provisioning, monitoring and upgrading, and Platform MPI for a high performance, easy to use MPI solution. Platform also offers the consulting, training and support to ensure your success. Our total solution provides the industry's fastest, most functional and scalable products for high performance clustering.

# Platform™

Platform Computing provides software that dynamically connects IT resources to workload demand according to business policies. Over 2,000 of the world's largest organizations rely on our solutions to improve IT productivity and reduce data center costs. Platform has strategic relationships with Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat®, and SAS®. Building on 16 years of market leadership, Platform continues to help data centers be more efficient, responsive and dynamic. Visit [www.platform.com](http://www.platform.com)

World Headquarters  
Platform Computing Inc.  
3760 14th Avenue  
Markham, Ontario  
Canada L3R 3T7  
Tel: +1 905 948 8448  
Fax: +1 905 948 9975  
Toll-free tel: 1 877 528 3676  
[info@platform.com](mailto:info@platform.com)

Sales - Headquarters  
Toll-free tel: 1 877 710 4477  
Tel: +1 905 948 8448  
  
North America  
New York: +1 646 290 5070  
San Jose: +1 408 392 4900  
Detroit: +1 248 359 7820

Europe  
Basingstoke: +44 (0) 1256 883756  
London: +44 (0) 20 7977 1480  
Paris: +33 (0) 1 41 10 09 20  
Düsseldorf: +49 2102 61039 0  
Munich: +49 89 517397 52  
Oslo: +44 1256 883756  
[info-europe@platform.com](mailto:info-europe@platform.com)

Asia-Pacific  
Beijing: +86 10 82276000  
Xi'an: +86 029 87607400  
[asia@platform.com](mailto:asia@platform.com)  
Tokyo: +81 (0)3-6302-2901  
[info-japan@platform.com](mailto:info-japan@platform.com)  
Singapore: +65 6307 6590  
[lliew@platform.com](mailto:lliew@platform.com)