

Intel® MPI Library for Linux* OS

User's Guide

Contents

Legal Information	4
1. Introduction.....	6
1.1. Introducing Intel® MPI Library	6
1.2. Intended Audience	7
1.3. Conventions and Symbols.....	7
1.4. Related Information	7
2. Usage Model.....	9
3. Installation and Licensing	10
3.1. Installing Intel® MPI Library.....	10
3.2. Licensing for Intel® MPI Library Runtime Environment and Software Development Kit	10
4. Compiling and Linking.....	11
4.1. Compiling an MPI Program	11
4.2. Adding Debug Symbols.....	11
4.3. Other Compilers Support	11
5. Running Applications	13
5.1. Running an MPI Program	13
5.2. Intel® MPI Library Configurations.....	13
5.3. Multi-threaded Applications	14
5.4. Selecting Fabrics	14
5.4.1. TCP Socket Connection	14
5.4.2. Shared Memory	14
5.4.3. Shared Memory and DAPL* Connection	14
5.4.4. Shared memory and TMI*	15
5.4.5. Shared memory and OFA*	15
5.4.6. Shared memory and OFI*	15
5.4.7. Multi-rail capability.....	15
5.4.8. I_MPI_FABRICS.....	16
6. Debugging and Testing	18
6.1. GDB*: The GNU* Project Debugger	18
6.2. TotalView* Debugger	18
6.3. DDT* Debugger	19
6.4. Logging	19
6.4.1. Getting Debug Information	19
6.4.2. Tracing an Application	19
6.4.3. Checking Correctness	19
6.4.4. Gathering Statistics.....	20
6.5. Testing the Installation.....	20
6.5.1. Compiling and Running a Test Program	20
7. Process Management	22
7.1. Selecting a Process Manager	22
7.2. Scalable Process Management System (Hydra)	22
7.3. Multipurpose Daemon* (MPD*).....	22
7.4. Controlling MPI Process Placement.....	22
8. Tuning with mpitune Utility	24
8.1. Cluster-Specific Tuning	24

8.2. Application-Specific Tuning.....	24
8.3. Setting Time Limit	25
8.4. Setting a Fabrics List	25
8.5. Setting a Range for the Number of Processes	25
8.6. Setting a Limitation for Hosts Usage	25
8.7. Restoring mpitune from the Last Saved Session	25
8.8. Tuning Applications Manually.....	25
9. Job Schedulers Support	26
9.1. Altair® PBS Pro®, TORQUE®, and OpenPBS®	26
9.2. IBM® Platform LSF®	26
9.3. Parallelnavi NQS®	26
9.4. SLURM®	27
9.5. Univa® Grid Engine®	27
9.6. SIGINT, SIGTERM Signals Intercepting.....	27
10. General Cluster Considerations	28
10.1. Defining which Nodes to Use.....	28
10.2. Password-less ssh Connection	28
10.3. Heterogeneous Systems and Jobs	29
11. Troubleshooting.....	30
11.1. General Troubleshooting Procedures	30
11.2. Examples of MPI Failures	30
11.2.1. Communication Problems	30
11.2.2. Environment Problems	33
11.2.3. Other Problems.....	35
12. Using the Intel® MPI Library with the Intel® Many Integrated Core (Intel® MIC) Architecture	37
12.1. Libraries	37
12.2. Multiple Cards	37
12.3. Using Intel® MPI Library on Intel® Xeon Phi™ Coprocessor	37
12.3.1. Building an MPI Application.....	37
12.3.2. Running an MPI Application	38

Legal Information

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

MPEG-1, MPEG-2, MPEG-4, H.261, H.263, H.264, MP3, DV, VC-1, MJPEG, AC3, AAC, G.711, G.722, G.722.1, G.722.2, AMRWB, Extended AMRWB (AMRWB+), G.167, G.168, G.169, G.723.1, G.726, G.728, G.729, G.729.1, GSM AMR, GSM FR are international standards promoted by ISO, IEC, ITU, ETSI, 3GPP and other organizations. Implementations of these standards, or the standard enabled platforms may require licenses from various entities, including Intel Corporation.

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.

Intel, the Intel logo, BlueMoon, BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Cilk, Core Inside, E-GOLD, Flexpipe, i960, Intel, the Intel logo, Intel AppUp, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Insider, the Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel Sponsors of Tomorrow., the Intel Sponsors of Tomorrow. logo, Intel StrataFlash, Intel vPro, Intel XScale, Intel True Scale Fabric, InTru, the InTru logo, the InTru Inside logo, InTru soundmark, Itanium, Itanium Inside, MCS, MMX, MPSS, Moblin, Pentium, Pentium Inside, Puma, skool, the skool logo, SMARTi, Sound Mark, Stay With It, The Creators Project, The Journey Inside, Thunderbolt, Ultrabook, vPro Inside, VTune, Xeon, Xeon Phi, Xeon Inside, X-GOLD, XMM, X-PMU and XPOSYS are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Java is a registered trademark of Oracle and/or its affiliates.

Bluetooth is a trademark owned by its proprietor and used by Intel Corporation under license.

Intel Corporation uses the Palm OS* Ready mark under license from Palm, Inc.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

© 2015 Intel Corporation. Portions (PBS Library) are copyrighted by Altair Engineering, Inc. and used with permission. All rights reserved.

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

1. Introduction

The *Intel® MPI Library for Linux* OS User Guide* explains how to use the Intel® MPI Library in some common usage scenarios. It provides information regarding compiling, linking, running, and debugging MPI applications, as well as information on integration within a cluster environment.

This *User's Guide* contains the following sections

Document Organization

Section	Description
Section 1. Introduction	Introduces this document
Section 2. Usage Model	Presents the usage model for working with the Intel® MPI Library
Section 3. Installation and Licensing	Describes the installation process and provides information about licensing
Section 4. Compiling and Linking	Gives instructions about how to compile and link MPI applications
Section 5. Running Applications	Describes the steps for running an application
Section 6. Debugging and Testing	Explains how to start an application under a debugger
Section 7. Process Management	Gives information about process managers and how to set up password-less ssh connections
Section 8. Tuning with mpitune Utility	Describes how to use the mpitune utility to find optimal settings for the Intel® MPI Library.
Section 9. Job Schedulers Support	Describes integration with job schedulers
Section 10. General Cluster Considerations	Discusses general considerations for clusters related to MPI usage
Section 11. Troubleshooting	Provides general troubleshooting steps and examples
Section 12. Using the Intel® MPI Library with the Intel® Many Integrated Core (Intel® MIC) Architecture	Describes some special considerations when using the Intel® MPI Library with the Intel® MIC architecture

1.1. Introducing Intel® MPI Library

The Intel® MPI Library is a multi-fabric message passing library that implements the Message Passing Interface, version 3.0 (MPI-3.0) specification. It provides a standard library across Intel® platforms that:

- Delivers best in class performance for enterprise, divisional, departmental and workgroup high performance computing. The Intel® MPI Library focuses on improving application performance on Intel® architecture based clusters.
- Enables you to adopt MPI-3.0 functions as your needs dictate

The product comprises the following main components:

- *Runtime Environment (RTO)* includes the tools you need to run programs, including scalable process management system (Hydra*), Multipurpose Daemon* (MPD), supporting utilities, shared (.so) libraries, and documentation.
- *Software Development Kit (SDK)* includes all of the Runtime Environment components plus compilation tools, including compiler drivers such as `mpicc`, include files and modules, static (.a) libraries, debug libraries, and test codes.

1.2. Intended Audience

This *User's Guide* helps an experienced user to start using the Intel® MPI Library and contains brief descriptions of the main functionality as a set of how-to instructions and examples. For full information, see *Intel® MPI Library Reference Manual for Linux* OS*.

1.3. Conventions and Symbols

The following conventions are used in this document.

Table 1.3-1 Conventions and Symbols used in this Document

<i>This type style</i>	Document or product names
<code>This type style</code>	Commands, arguments, options, file names
<code>THIS_TYPE_STYLE</code>	Environment variables
<code><this type style></code>	Placeholders for actual values
<code>[items]</code>	Optional items
<code>{ item item }</code>	Selectable items separated by vertical bar(s)
(SDK only)	For Software Development Kit (SDK) users only

1.4. Related Information

To get more information about the Intel® MPI Library, explore the following resources:

- *Intel® MPI Library Release Notes* for updated information on requirements, technical support, and known limitations.
- *Intel® MPI Library Reference Manual* for in-depth knowledge of the product features, commands, options, and environment variables.

- [*Intel® MPI Library for Linux* OS Knowledge Base*](#) for additional troubleshooting tips and tricks, compatibility notes, known issues, and technical notes.

For additional resources, see:

[*Intel® MPI Library Product Web Site*](#)

[*Intel Product Support*](#)

[*Intel® Cluster Tools Products Website*](#)

[*Intel® Software Development Products Website*](#)

2. Usage Model

Using the Intel® MPI Library involves the following steps:

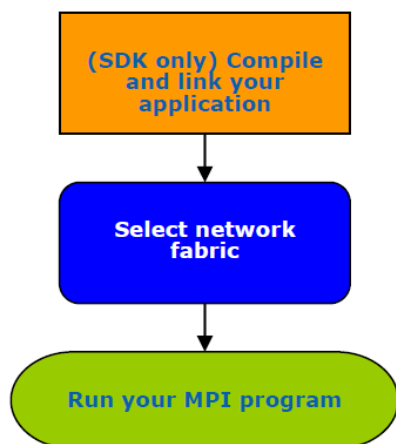


Figure 1: Flowchart representing the usage model for working with the Intel® MPI Library.

3. Installation and Licensing

This section describes the installation process and provides information about licensing for Intel® MPI Library Runtime Environment (RTO) and Software Development Kit (SDK).

3.1. Installing Intel® MPI Library

If you have a previous version of the Intel® MPI Library for Linux* OS installed, you do not need to uninstall it before installing the latest version.

Extract the `l_mpi[-rt]_p_<version>.<package_num>.tar.gz` package by using following command:

```
tar -xvzf l_mpi[-rt]_p_<version>.<package_num>.tar.gz
```

This command creates the subdirectory `l_mpi[-rt]_p_<version>.<package_num>`.

To start installation, run `install.sh`. The default installation path for the Intel® MPI Library is `/opt/intel/compilers_and_libraries_2016.<update>.<package#>/linux/mpi`.

There are 2 different installations:

- RPM-based installation- this installation requires root password. The product can be installed either on a shared file system or on each node of your cluster.
- Non-RPM installation- this installation does not require root access and it installs all scripts, libraries, and files in the desired directory (usually `$HOME` for the user).

Scripts, include files, and libraries for different architectures are located in different directories. By default, you can find binary files and all needed scripts under `<installdir>/<arch>` directory. For example, for Intel®64 architecture, `<arch>` is `bin64`.

NOTE

To use the full functionality of the library on Intel®64 architecture, set the environment from the `<installdir>/bin64` directory. You can use the provided script files to simplify setting the environment variables. Source the file `<installdir>/bin[64]/mpivars.[c]sh` to set appropriate values for your desired development environment.

For more information on installation, see *Intel® MPI Library for Linux* OS Installation Guide*. You can also find information about how to install the product in silent mode and some useful installer options.

3.2. Licensing for Intel® MPI Library Runtime Environment and Software Development Kit

There are two different licensing options:

- Intel® MPI Library Runtime Environment (RTO) license. The license covers everything you need to run Intel® MPI Library-based applications and is free and permanent.
- Intel® MPI Library Software Development Kit (SDK) license. This license covers all of Runtime Environment components as well as the compilation tools: compiler drivers (`mpiicc`, `mpicc`, and so on), files and modules, static (`.a`) libraries, debug libraries, trace libraries, and test sources. This license is fee-based, with several options described in the product end-user license agreement (EULA).

4. Compiling and Linking

This section gives instructions about how to compile and link different kinds of your Intel® MPI Library applications, and details on different debugging and compiler support options. The following topics are included in this section:

- [Compiling an MPI program](#)
- [Adding debug symbols](#)
- [Other compilers support](#)

4.1. Compiling an MPI Program

This topic describes the basic steps required to compile and link an MPI program, when using only the Intel® MPI Library Development Kit. To compile and link an MPI program with the Intel® MPI Library:

1. Source the appropriate `mpivars.[c]sh` script to get the proper environment settings (using Bash* shell):

```
$ . <installdir>/bin64/mpivars.sh
```

2. Compile your MPI program using the appropriate `mpixxx` compiler command. For example, to compile a program written in C, use the `mpiicc` command as follows:

```
$ mpiicc <installdir>/test/test.c -o testc
```

You'll get an executable file `testc` in the current directory which can be started immediately. For description of how to launch your application, see [Running Applications](#) in this document.

NOTE

By default, the executable file `testc` is linked with the multi-threaded optimized Intel MPI Library. If you need to use other configuration, see [Intel® MPI Library Configurations](#).

Other supported compilers have an equivalent command that uses the prefix `mpi` on the standard compiler command. For the full list of supported compilers, see the Compiler Commands topic in *Intel® MPI Library Reference Manual for Linux* OS*.

4.2. Adding Debug Symbols

If you need to debug your application, add the `-g` option. In this case debug information is added to the binary. You can use any debugger to debug the application.

```
$ mpiicc -g test.c -o testc
```

4.3. Other Compilers Support

Intel® MPI Library provides binding libraries to support different operating systems (different `glibc*` versions) and different compilers. These libraries provide C++, F77, F90 interfaces.

GNU* and Intel® Compilers binding libraries:

- `libmpicxx.{a|so}` – for `g++` ≥ 3.4 ;
- `libmpifort.{a|so}` – for `g77/gfortran` interface for GNU* and Intel® Compilers;

Your application will be linked against correct GNU* and Intel* Compilers binding libraries if you are using `mpicc, mpicxx, mpifc, mpif77, mpif90, mpigcc, mpigxx, mpiicc, mpiicpc` or `mpiifort` compiler commands.

For third-party compilers, there is a binding kit which allows you to add support for a certain compiler to the Intel* MPI Library for Linux* OS.

NOTE

The Intel* MPI Library supports Intel* compilers as well as GNU* compilers out of the box. See the *Intel* MPI Library Release Notes* for more details.

The Intel* MPI Library supports PGI* C, PGI* Fortran 77, Absoft* Fortran 77 compilers out of the box, with the following caveats:

- Your PGI* compiled source files must not transfer long double entities
- Your Absoft* based build procedure must use the `-g77, -B108` compiler options
- Install and select the right compilers
- Ensure that the respective compiler runtime is installed on all nodes

You have to build extra Intel* MPI Library binding libraries if you need the support for PGI* C++, PGI* Fortran 95, Absoft* Fortran 95 and GNU* Fortran 95 higher than version 4.0 bindings.

This binding kit provides all the necessary source files, convenience scripts, and instructions you need.

The binding kit and detailed description are located in `binding` directory. To get access to the binding, submit a request to the Intel* MPI Library for Linux* OS product at the Intel* Premier Support site.

5. Running Applications

After you have compiled and linked your application, you are ready to run your MPI applications. This topic describes the steps for running applications.

5.1. Running an MPI Program

The easiest way to run an MPI program is using the `mpirun` command:

```
$ mpirun -n <# of processes> ./myprog
```

This command invokes the `mpiexec.hydra` command which uses the Hydra Process Manager by default. Use the `mpiexec.hydra` options on the `mpirun` command line.

NOTE

The commands `mpirun` and `mpiexec.hydra` are interchangeable. You are recommended to use the `mpirun` command to run Intel MPI programs for the following reasons:

1. You can specify all `mpiexec.hydra` options with the `mpirun` command.
2. The `mpirun` command detects if the MPI job is submitted from within a session allocated using a job scheduler like PBS Pro* or LSF*. Thus, you are recommended to use the `mpirun` command when an Intel MPI program is running under batch schedulers or job managers.

Use the `-n` option to set the number of MPI processes. If the `-n` option is not specified, the process manager pulls the host list from a job scheduler, or uses the number of cores on the machine.

By default, the ssh protocol is used for communication between nodes. If you are using rsh instead, use `-r rsh` option:

```
$ mpirun -r rsh -n <# of processes> ./myprog
```

For a successful run, configure password-less ssh connections for all nodes. For more details, see the Job Startup Commands topic in *Intel® MPI Library Reference Manual for Linux® OS*.

After you successfully run your application using the Intel® MPI Library, you can move your application from one cluster to another and use different fabrics between the nodes without re-linking. If you encounter problems, see [Debugging and Testing](#) for possible solutions.

5.2. Intel® MPI Library Configurations

To configure your Intel® MPI Library, source the script `mpivars.[c]sh` with appropriate arguments. For example:

```
$ . <installdir>/bin64/mpivars.sh release
```

You can use the following arguments in this command. The multi-threaded optimized Intel® MPI Library is used by default.

Argument	Definition
release	Set this argument to use single-threaded optimized Intel® MPI Library.
debug	Set this argument to use single-threaded debug Intel® MPI Library.
release_mt	Set this argument to use multi-threaded optimized Intel® MPI Library.

debug_mt	Set this argument to use multi-threaded debug Intel® MPI Library.
----------	---

NOTE

If you want to use different configuration of Intel® MPI Library, run the `mpivars.[c]sh` script with appropriate arguments before an application launch. You do not need to recompile applications.

5.3. Multi-threaded Applications

To run OpenMP* application and pin threads inside the domain, make sure the `KMP_AFFINITY` environment variable is set to use the corresponding OpenMP* feature.

Run the application:

```
$ mpirun -genv OMP_NUM_THREADS 4 -n <# of processes> ./myprog
```

For more details see the Interoperability with OpenMP* topic in the *Intel® MPI Library Reference Manual for Linux* OS*.

5.4. Selecting Fabrics

By default, the Intel® MPI Library selects a network fabric based on the list of fabrics specified in `I_MPI_FABRICS_LIST`. To select a specific network fabric combination, use the `-genv` option to assign a value to the `I_MPI_FABRICS` variable. You can also assign a value using the export command.

If the specified fabric is not available, Intel® MPI Library will go down the list specified in `I_MPI_FABRICS_LIST` and select the next available fabric.

You can disable this fallback behavior by using the `I_MPI_FALLBACK` variable:

```
$ export I_MPI_FALLBACK=0
```

By default, the fallback is enabled. If `I_MPI_FABRICS` is set, the fallback will be disabled.

5.4.1. TCP Socket Connection

Use the following command to run an MPI program over TCP sockets using the available Ethernet connection on the cluster. The program does not use shared memory within a node:

```
$ mpirun -genv I_MPI_FABRICS tcp -n <# of processes> ./myprog
```

5.4.2. Shared Memory

Use the following command to run an MPI program over the shared-memory fabric (shm) only:

```
$ mpirun -genv I_MPI_FABRICS shm -n <# of processes> ./myprog
```

5.4.3. Shared Memory and DAPL* Connection

To use shared memory for intra-node communication and the Direct Access Programming Library* (DAPL*) layer for inter-node communication, use the following command:

```
$ mpirun -genv I_MPI_FABRICS shm:dapl -n <# of processes> ./myprog
```

Use the `I_MPI_DAPL_UD` environment variable to enable connectionless DAPL User Datagrams* (DAPL UD*):

```
$ export I_MPI_DAPL_UD=enable
```

```
$ mpirun -genv I_MPI_FABRICS shm:dapl -n <# of processes> ./myprog
```

This is the default method if no fabric options are selected.

5.4.4. Shared memory and TMI*

To use shared memory for intra-node communication and the Tag Matching Interface* (TMI*) for inter-node communication, use the following command (make sure that you have libtmi.so library in the search path of the ldd command):

```
$ mpirun -genv I_MPI_FABRICS shm:tmi -n <# of processes> ./myprog
```

This is the recommended method if using Intel® Omni Scale Fabric (formerly Intel® True Scale) or the Myricom* MX interface.

5.4.5. Shared memory and OFA*

To select shared memory for intra-node communication and OpenFabrics* Enterprise Distribution (OFED) verbs for inter-node communication, use the following command:

```
$ mpirun -genv I_MPI_FABRICS shm:ofa -n <# of processes> ./myprog
```

This is the recommended method if using the Open Fabrics* Enterprise Distribution (OFED*) software stack.

5.4.6. Shared memory and OFI*

To select shared memory for intra-node communication and OpenFabrics Interface* (OFI*) for inter-node communication, use the following command:

```
$ mpirun -genv I_MPI_FABRICS shm:ofi -n <# of processes> ./myprog
```

This is the recommended method if using the OpenFabrics Interface* (OFI*) software stack.

5.4.7. Multi-rail capability

If your cluster is equipped with several connection cards or multi-port cards, you can improve bandwidth of communications using the following settings:

```
$ export I_MPI_FABRICS=shm:ofa
$ export I_MPI_OFA_NUM_ADAPTERS=<num>
```

Where <num> is the number of connection adapters (1 by default).

If connection cards have several ports, you can specify the number of ports using the following setting:

```
$ export I_MPI_OFA_NUM_PORTS=<num>
```

For more details see the Fabrics Control topic in *Intel® MPI Library Reference Manual for Linux® OS*.

After you successfully run your application using the Intel MPI Library over any of the fabrics described, you can move your application from one cluster to another and use different fabrics between the nodes without re-linking. If you encounter problems, see [Debugging and Testing](#) for possible solutions.

Additionally, using `mpirun` is the recommended practice when using a resource manager, such as PBS Pro* or LSF*.

For example, to run the application in the PBS environment, follow these steps:

1. Create a PBS launch script that specifies number of nodes requested and sets your Intel MPI Library environment. For example, create a `pbs_run.sh` file with the following content:

```
#PBS -l nodes=2:ppn=1
#PBS -l walltime=1:30:00
#PBS -q workq
#PBS -V
# Set Intel MPI environment
mpi_dir=<installdir>/<arch>/bin
```

```
cd $PBS_O_WORKDIR
source $mpi_dir/mpivars.sh
# Launch application
mpirun -n <# of processes> ./myprog
```

2. Submit the job using the PBS `qsub` command:

```
$ qsub pbs_run.sh
```

When using `mpirun` under a job scheduler, you do not need to determine the number of available nodes. Intel MPI Library automatically detects the available nodes through the Hydra process manager.

5.4.8. I_MPI_FABRICS

This topic is an excerpt from the *Intel® MPI Library Reference Manual for Linux® OS* which provides further details on the `I_MPI_FABRICS` environment variable.

Select a particular network fabric to be used for communication.

Syntax

```
I_MPI_FABRICS=<fabric>|<intra-node fabric>:<inter-nodes fabric>
```

Where

- `<fabric> := {shm, dapl, tcp, tmi, ofa, ofi}`
- `<intra-node fabric> := {shm, dapl, tcp, tmi, ofa, ofi}`
- `<inter-nodes fabric> := {dapl, tcp, tmi, ofa, ofi}`

Arguments

Argument	Definition
<code><fabric></code>	Define a network fabric
<code>shm</code>	Shared-memory
<code>dapl</code>	DAPL-capable network fabrics, such as InfiniBand*, iWarp*, Dolphin*, and XPMEM* (through DAPL*)
<code>tcp</code>	TCP/IP-capable network fabrics, such as Ethernet and InfiniBand* (through IPoIB*)
<code>tmi</code>	Network fabrics with tag matching capabilities through the Tag Matching Interface (TMI), such as Intel® True Scale Fabric and Myrinet*
<code>ofa</code>	Network fabric, such as InfiniBand* (through OpenFabrics* Enterprise Distribution (OFED*) verbs) provided by the Open Fabrics Alliance* (OFA*)
<code>ofi</code>	OFI (OpenFabrics Interfaces*)-capable network fabric including Intel® True Scale Fabric, and TCP (through OFI* API)

For example, to select the winOFED* InfiniBand* device, use the following command:

```
$ mpirun -n <# of processes> \ -env I_MPI_FABRICS shm:dapl <executable>
```


For these devices, if `<provider>` is not specified, the first DAPL* provider in the `/etc/dat.conf` file is used. The `shm` fabric is available for both Intel® and non-Intel microprocessors, but it may perform additional optimizations for Intel microprocessors than for non-Intel microprocessors.

NOTE

Ensure the selected fabric is available. For example, use `shm` only if all the processes can communicate with each other through the availability of the `/dev/shm` device. Use `dapl` only when all processes can communicate with each other through a single DAPL provider.

6. Debugging and Testing

This topic explains how to debug MPI applications with different debugger tools.

6.1. GDB*: The GNU* Project Debugger

Use the following command to launch the GDB* debugger with Intel® MPI Library:

```
$ mpirun -gdb -n 4 ./testc
```

You can work with the GDB debugger as you usually do with a single-process application. For details about how to work with parallel programs, see the GDB documentation at <http://www.gnu.org/software/gdb/>.

You can also attach to a running job with

```
$ mpirun -n 4 -gdba <pid>
```

Where <pid> is the process ID for the running MPI rank.

6.2. TotalView* Debugger

Intel® MPI Library supports the use of the TotalView* debugger from Rogue Wave* Software, Inc. To debug an MPI program, add `-tv` to the global `mpirun` arguments, as in

```
$ mpirun -tv -n 4 ./testc
```

NOTE

In case of ssh communication, you need to set the `TVDSVRLAUNCHCMD` environment variable to `ssh`.

You will get a popup window from TotalView asking whether you want to start the job in a stopped state. If so, when the TotalView window appears, you may see assembly code in the source window. Click on `main()` in the stack window (upper left) to see the source of the main function. TotalView shows that the program (all processes) are stopped in the call to `MPI_Init()`. When debugging with TotalView using the above startup sequence, you need to exit TotalView before restarting an Intel MPI Library job.

To debug with TotalView an enable restarting the session, use the following command line:

```
$ totalview python -a 'which mpirun' -tvsu <mpirun_args> <prog> <prog_args>
```

If you have TotalView 8.1.0 or later, you can use a feature called indirect launch.

1. Invoke TotalView as:

```
$ totalview <prog> -a <prog_args>
```

2. Select the **Process/Startup Parameters** command.
3. Choose the **Parallel** tab in the resulting dialog box and choose **MPICH2** as the parallel system.
4. Set the number of tasks using the **Tasks** field.
5. Enter other needed `mpirun` arguments into the **Additional Starter Arguments** field.

If you want to be able to attach to a running MPI job using TotalView, you must use the `-tvsu` option in the `mpirun` command when starting the job. Using this option adds a barrier inside `MPI_Init()` and hence may affect startup performance slightly. After all tasks have returned from `MPI_Init()`, there is no performance degradation incurred from this option.

6.3. DDT* Debugger

You can debug MPI applications using DDT* debugger. Intel does not provide support for this debugger. You should obtain the support from Allinea*. According to the DDT User Guide at <http://www.allinea.com/products/ddt-support/>, you can use the `-tv` option to run DDT with certain TotalView* variables set beforehand.

```
$ export TOTALVIEW=DDT_INSTALLATION_PATH/bin/ddt-debugger-mps
$ mpirun -np 4 -tv ./your_app
```

If you have problems with the DDT debugger, see DDT documentation for help.

6.4. Logging

Sometimes debugging an application is not effective and you prefer to use logging instead. There are several ways to get logging information from running applications.

6.4.1. Getting Debug Information

Environment variable `I_MPI_DEBUG` provides a very convenient way to get information from an MPI application at runtime. You can set value of this variable from 0 (the default value) to 1000. The higher the value, the more debug information you get.

```
$ mpirun -genv I_MPI_DEBUG 5 -n 8 ./my_application
```

NOTE

High values of `I_MPI_DEBUG` can output a lot of information and significantly reduce performance of an application. A value of `I_MPI_DEBUG=5` is generally a good starting point, which provides sufficient information to find common errors. See the `I_MPI_DEBUG` description in *Intel® MPI Library Reference Manual for Linux* OS* for more details.

6.4.2. Tracing an Application

Use the `-t` or `-trace` option to link the resulting executable files against the Intel® Trace Collector library. This has the same effect as when `-profile=vt` is used as an argument to `mpiicc` or another compiler script.

```
$ mpiicc -trace test.c -o testc
```

To use this option, you need to:

- Install the Intel® Trace Analyzer and Collector first. The tool is distributed as part of the Intel® Parallel Studio XE Cluster Edition bundle only.
- Include the installation path of the Intel® Trace Collector in the `VT_ROOT` environment variable. Set `I_MPI_TRACE_PROFILE` to the `<profile_name>` environment variable to specify another profiling library. For example, set `I_MPI_TRACE_PROFILE` to `vtfs` to link against the fail-safe version of the Intel® Trace Collector.

6.4.3. Checking Correctness

Use `-check_mpi` option to link the resulting executable file against the Intel® Trace Collector correctness checking library. This has the same effect as when `-profile=vtmc` is used as an argument to `mpiicc` or another compiler script.

```
$ mpiicc -profile=vtmc test.c -o testc
```

Or

```
$ mpiicc -check_mpi test.c -o testc
```

To use this option, you need to:

- Install the Intel® Trace Analyzer and Collector first. The tool is distributed as part of the Intel® Parallel Studio XE Cluster Edition bundle only.
- Include the installation path of the Intel® Trace Collector in the `VT_ROOT` environment variable. Set `I_MPI_CHECK_PROFILE` to the `<profile_name>` environment variable to specify another checking library.

For more information on the Intel® Trace Analyzer and Collector, see the documentation provided with this product.

6.4.4. Gathering Statistics

If you want to collect statistics on MPI functions used in your application, you can set the `I_MPI_STATS` environment variable to a number, between 1 to 10. This environment variable controls the amount of statistics information collected and the output to the log file. By default, no statistics are gathered.

For details, see Statistics Gathering Mode in *Intel® MPI Library Reference Manual for Linux* OS*.

6.5. Testing the Installation

To ensure that the Intel® MPI Library is installed and functioning correctly, complete the general testing below, in addition to compiling and running a test program.

To test the installation (on each node of your cluster):

1. Verify that `<installdir>/<arch>/bin` is in your `PATH`:

```
$ ssh <nodename> which mpirun
```

You should see the correct path for each node you test.

(SDK only) If you use the Intel® Composer XE packages, verify that the appropriate directories are included in the `PATH` and `LD_LIBRARY_PATH` environment variables

```
$ mpirun -n <# of processes> env | grep PATH
```

You should see the correct directories for these path variables for each node you test. If not, call the appropriate `compilervars.[c]sh` script. For example, for the Intel® Composer XE 2015 use the following source command:

```
$ . /opt/intel/composer_xe_2015/bin/compilervars.sh intel64
```

2. In some unusual circumstances, you need to include the `<installdir>/<arch>/lib` directory in your `LD_LIBRARY_PATH`. To verify your `LD_LIBRARY_PATH` settings, use the command:

```
$ mpirun -n <# of processes> env | grep LD_LIBRARY_PATH
```

6.5.1. Compiling and Running a Test Program

To compile and run a test program, do the following:

1. **(SDK only)** Compile one of the test programs included with the product release as follows:

```
$ cd /opt/intel/samples_2016/en/mpi
```

```
$ mpiicc -o myprog test.c
```

2. If you are using InfiniBand*, Myrinet*, or other RDMA-capable network hardware and software, verify that everything is functioning correctly using the testing facilities of the respective network.
3. Run the test program with all available configurations on your cluster.

- Test the TCP/IP-capable network fabric using:

```
$ mpirun -n 2 -genv I_MPI_DEBUG 2 -genv I_MPI_FABRICS tcp ./myprog
```

You should see one line of output for each rank, as well as debug output indicating the TCP/IP-capable network fabric is used.

- Test the shared-memory and DAPL-capable network fabrics using:

```
$ mpirun -n 2 -genv I_MPI_DEBUG 2 -genv I_MPI_FABRICS shm:dapl ./myprog
```

You should see one line of output for each rank, as well as debug output indicating the shared-memory and DAPL-capable network fabrics are being used.

- Test any other fabric using:

```
$ mpirun -n 2 -genv I_MPI_DEBUG 2 -genv I_MPI_FABRICS <fabric> ./myprog
```

where *<fabric>* is a supported fabric. For more information, see *Selecting Fabrics*.

For each of the `mpirun` commands used, you should see one line of output for each rank, as well as debug output indicating which fabric was used. The fabric(s) should agree with the `I_MPI_FABRICS` setting.

The `/opt/intel/samples_2016/en/mpi` directory in the Intel® MPI Library Development Kit contains other test programs in addition to `test.c`

7. Process Management

This topic describes the process managers included with the Intel® MPI Library:

- [Selecting a process manager](#)
- [Scalable process management system \(Hydra\)](#)
- [Multipurpose Daemon* \(MPD*\)](#)
- [Controlling MPI process placement](#)

7.1. Selecting a Process Manager

The `mpirun` script uses the process manager specified by the `I_MPI_PROCESS_MANAGER` variable. By default, `mpirun` selects the Hydra Process Manager. Setting `I_MPI_PROCESS_MANAGER` to `hydra` will explicitly select Hydra, and setting it to `mpd` will explicitly select MPD. The process manager can also be selected by directly calling the appropriate `mpiexec` file: `mpiexec.hydra` for Hydra or `mpiexec` for MPD.

7.2. Scalable Process Management System (Hydra)

Hydra is a simplified, scalable process manager. Hydra will check for known resource managers to determine where the processes may be run and to distribute the processes among the targets using proxies on each host. These proxies will be used for process launching, cleanup, I/O forwarding, signal forwarding, and other tasks.

You can start Hydra by using `mpiexec.hydra`. See [Scalable Process Management System \(Hydra\) Commands](#) topic for a detailed list of options in the *Intel® MPI Library Reference Manual*.

NOTE

Multipurpose daemon* (MPD) has been deprecated starting from Intel® MPI Library 5.0 release. To start parallel jobs, use the scalable process management system (Hydra).

7.3. Multipurpose Daemon* (MPD*)

MPD stands for Multipurpose Daemon. This is the Intel® MPI Library process management system for starting parallel jobs, which have to run on all nodes. MPDs gather information about the system and hardware, as well as communicate with each other to exchange required information. For example, an MPD ring is required for correct pinning under the MPD process manager.

NOTE

Multipurpose daemon* (MPD) has been deprecated starting with Intel® MPI Library 5.0 release. Convert to using the scalable process management system (Hydra) instead to start parallel jobs.

7.4. Controlling MPI Process Placement

The `mpirun` command controls how the ranks of the processes are allocated to the nodes of the cluster. By default, the `mpirun` command uses group round-robin assignment, putting consecutive MPI process on all processor ranks of a node. This placement algorithm may not be the best choice for your application, particularly for clusters with symmetric multi-processor (SMP) nodes.

Suppose that the geometry is `<#ranks> = 4` and `<#nodes> = 2`, where adjacent pairs of ranks are assigned to each node (for example, for two-way SMP nodes). To see the cluster nodes, enter the command:

```
cat ~/mpd.hosts
```

The results should look as follows:

```
clusternode1  
clusternode2
```

To equally distribute four processes of the application on two-way SMP clusters, enter the following command:

```
mpirun -perhost 2 -n 4 ./myprog.exe
```

The output for the `myprog.exe` executable file may look as follows:

```
Hello world: rank 0 of 4 running on clusternode1  
Hello world: rank 1 of 4 running on clusternode1  
Hello world: rank 2 of 4 running on clusternode2  
Hello world: rank 3 of 4 running on clusternode2
```

Alternatively, you can explicitly set the number of processes to be executed on each host through the use of argument sets. One common use case is when employing the master-worker model in your application. For example, the following command equally distributes the four processes on `clusternode1` and on `clusternode2`:

```
mpirun -n 2 -host clusternode1 ./myprog.exe : -n 2 -host clusternode2  
./myprog.exe
```

See Also

You can get more details in the *Local Options* topic of the *Intel® MPI Library Reference Manual* for Linux* OS.

You can get more information about controlling MPI process placement online at [Controlling Process Placement with the Intel® MPI Library](#).

8. Tuning with mpitune Utility

This section describes how to use the `mpitune` utility to find optimal settings for the Intel® MPI Library:

- [Cluster-specific tuning](#)
- [Application-specific tuning](#)
- [Setting time limit](#)
- [Setting a fabrics list](#)
- [Setting a range for the number of processes](#)
- [Setting a limitation for hosts usage](#)
- [Restoring mpitune from the last saved session](#)
- [Tuning applications manually](#)

8.1. Cluster-Specific Tuning

Intel® MPI Library has more than 100 parameters. The defaults are set for common usage and generally provide good performance for most clusters and most applications. However, if you want to get even higher performance, you can use the `mpitune` utility. This utility uses the Intel® MPI Benchmarks (IMB) as a benchmark program running tests several times with different parameters and searching for the best ones. Start the `mpitune` utility with the following command:

```
$ mpitune
```

Then, start your application with the `-tune` option to enable the tuned settings:

```
$ mpirun -tune -perhost 8 -n 64 ./your_app
```

For best results, run `mpitune` with write access permissions for `<installdir>/<arch>/etc` which is the location for tuned parameters. If you do not have write access, a new configuration file will be saved in your current directory.

By default, `mpitune` uses the Intel® MPI Benchmarks (IMB) as benchmark program. Alternatively, you can substitute with your benchmark of choice by using the following command:

```
$ mpitune -test \"your_benchmark -param1 -param2\"
```

You can then apply the new settings as described in this topic.

The Intel® MPI Benchmarks executable files, which are more optimized for Intel microprocessors than for non-Intel microprocessors, are used by default. This may result in different tuning settings on Intel microprocessors than on non-Intel microprocessors.

8.2. Application-Specific Tuning

Use the `mpitune` utility to find optimal settings for your specific application.

```
$ mpitune --application \"your_app\" --output-file yourapp.conf
```

Where `\"your_app\"` is the exact command line you use to start your application. For example:

```
\"mpitune --application \"./my_test\" --output-file $PWD/my_test.conf
```

Tuned settings are saved in `yourapp.conf` file. To apply them, call `mpirun` as:

```
$ mpirun -tune $PWD/yourapp.conf -perhost 8 -n 64 ./your_app
```

`\"your_app\"` is not only executable file but any script which can be started as a separate process.

NOTE

The script should not change the `I_MPI_*` variables.

8.3. Setting Time Limit

The process of tuning can take a lot of time. Due to the varying factors involved in each cluster and application setup, the tuning time can be unpredictable.

To restrict the tuning time, you can set the `--time-limit` option. For example, to limit the tuning to 8 hours (480 minutes), run the following command:

```
$ mpitune --time-limit 480
```

The time unit used is minutes.

8.4. Setting a Fabrics List

To define the fabrics to be tested, use the `fabrics-list` option.

```
$ mpitune --fabrics-list shm,ofa,dapl
```

The available fabrics are: `shm:dapl,shm:tcp,shm,dapl,shm:ofa,shm:tmi,ofa,tmi,tcp`.

8.5. Setting a Range for the Number of Processes

To limit the number of processes running on one node, you can use the `perhost-range min:max` option. For example: the following command defines the number of MPI ranks on each node, between 4 and 8:

```
$ mpitune --perhost-range 4:8
```

8.6. Setting a Limitation for Hosts Usage

To limit the number of nodes on which tuning will be performed, use the `--host-range min:max` option. For example, the following command will restrict running on 8 to 16 nodes only:

```
$ mpitune --host-range 8:16
```

8.7. Restoring mpitune from the Last Saved Session

Sometimes an unexpected event can occur during the `mpitune`. In this case, you can use the intermediate saved in a `mpituner_session_<session-id>.mts` file. To restart `mpitune` from the last saved session:

```
$ mpitune --session-file ./mpituner_session_<session-id>.mts
```

Where `<session-id>` is the Unix* timestamp of the moment tuner started.

8.8. Tuning Applications Manually

There is a family of `I_MPI_ADJUST_*` environment variables that allow you to manually tune the collective operations of the Intel® MPI Library. By setting a range of message sizes and choosing different algorithms, you can improve the performance of your application. For more information, see the `I_MPI_ADJUST` Family topic in *Intel® MPI Library Reference Manual for Linux* OS* for details.

9. Job Schedulers Support

The Intel® MPI Library supports the majority of commonly used job schedulers in the HPC field.

The following job schedulers are supported on Linux* OS:

- Altair* PBS Pro*
- Torque*
- OpenPBS*
- IBM* Platform LSF*
- Parallelnavi* NQS*
- SLURM*
- Univa* Grid Engine*

On Linux* OS, this support is implemented in the `mpirun` wrapper script. `mpirun` determines the job scheduler under which it was started by checking specific environment variables and then chooses the appropriate method to start an application.

9.1. Altair* PBS Pro*, TORQUE*, and OpenPBS*

If you use one of these job schedulers, with the `$PBS_ENVIRONMENT` exists and the value is `PBS_BATCH` or `PBS_INTERACTIVE`, `mpirun` uses `$PBS_NODEFILE` as a machine file for `mpirun`. You do not need to specify the `-machinefile` option directly.

Example of a batch job script:

```
#PBS -l nodes=4:ppn=4
#PBS -q queue_name
cd $PBS_O_WORKDIR
mpirun -n 16 ./myprog.exe
```

9.2. IBM* Platform LSF*

If you use the IBM* Platform LSF* job scheduler, and the `$LSB_MCPU_HOSTS` is set, it will be parsed to get the list of hosts for the parallel job. `$LSB_MCPU_HOSTS` does not store the main process name; thus the local hostname will be added to the top of the hosts list. Based on this hosts list, a machine file for `mpirun` is generated with a unique name: `/tmp/lsf_${username}.$$`. The machine file is removed when the job is completed.

Example:

```
$ bsub -n 16 mpirun -n 16 ./myprog.exe
```

9.3. Parallelnavi NQS*

If you use Parallelnavi NQS* job scheduler and the `$ENVIRONMENT`, `$QSUB_REQID`, `$QSUB_NODEINF` options are set, the `$QSUB_NODEINF` file is used as a machine file for `mpirun`. Also `-r plesh` is used as remote shell by the process manager during startup.

9.4. SLURM*

If the `$SLURM_JOBID` is set, the `$SLURM_TASKS_PER_NODE`, `$SLURM_NODELIST` environment variables will be used to generate a machine file for `mpirun`. The name of the machine file is `/tmp/slurm_${username}.$$`. The machine file will be removed when the job is completed.

Example:

```
$ srun -N2 --nodelist=host1,host2 -A
$ mpirun -n 2 myprog.exe
```

9.5. Univa* Grid Engine*

If you use the Univa* Grid Engine* job scheduler and the `$PE_HOSTFILE` is set, then two files will be generated: `/tmp/sge_hostfile_${username}.$$` and `/tmp/sge_machifile_${username}.$$`. The latter is used as the machine file for `mpirun`. These files are removed when the job is completed.

9.6. SIGINT, SIGTERM Signals Intercepting

If resources allocated to a job exceed the limit, most job schedulers terminate the job by sending a signal to all processes.

For example, Torque* sends `SIGTERM` three times to a job and if this job is still alive, `SIGKILL` will be sent to terminate it.

For Univa* Grid Engine*, the default signal to terminate a job is `SIGKILL`. Intel® MPI Library is unable to process or catch that signal causing `mpirun` to kill the entire job. You can change the value of the termination signal through the following queue configuration:

- Use the following command to see available queues:

```
$ qconf -sql
```
- Execute the following command to modify the queue settings:

```
$ qconf -mq <queue_name>
```
- Find "terminate_method" and change signal to `SIGTERM`
- Save queue configuration

10. General Cluster Considerations

This topic discusses general considerations for clusters related to MPI usage:

- Defining which nodes to use
- Password-less ssh connection
- Heterogeneous system and jobs

10.1. Defining which Nodes to Use

By default, Intel® MPI Library looks for a file called `mpd.hosts`. This file should contain a list of all available nodes on the cluster which can be used for your application. The format of the `mpd.hosts` file is a list of node names, one name per line. Blank lines and the portions of any lines that follow a `#` character are ignored.

You can specify the full path to this file by using the `-f` option.

When running under a supported job scheduler, using the `-f` option is unnecessary as the hosts will be determined by the scheduler.

10.2. Password-less ssh Connection

When a process is started remotely, `ssh` is used to launch the processes by default. Without a password-less SSH capability enabled, your password will be requested when launching jobs. A script is provided with the Intel® MPI Library installation package that will automatically generate and distribute SSH keys for a user. The script is named `sshconnectivity.exp` and is located in the main folder after extracting the tarball. If the script does not work for your system, the keys can be generated and distributed manually by following these steps:

1. Generate a public key

```
local> ssh-keygen -t dsa -f .ssh/id_dsa
```

When you are prompted for a password, leave it blank by pressing the <enter> key

Two new files `id_dsa` and `id_dsa.pub` are created in the `.ssh` directory. The latter one is the public key.

2. Distribute the public key to remote nodes

Go to the `.ssh` directory. Copy the public key to the remote machines.

```
local> cd .ssh
local> scp id_dsa.pub user@remote:~/.ssh/id_dsa.pub
```

Log into the remote machine and go to the `.ssh` directory on the remote side.

```
local> ssh user@remote
remote> cd .ssh
```

Add the client's public key to the known public keys on the remote server.

```
remote> cat id_dsa.pub >> authorized_keys
remote> chmod 640 authorized_keys
remote> rm id_dsa.pub
remote> exit
```

Next time you log into the remote server, you will not be prompted for a password.

NOTE

ssh setup depends on the ssh client distribution.

10.3. Heterogeneous Systems and Jobs

All clusters are not homogeneous. All jobs are not homogeneous. The Intel® MPI Library is able to run multiple sets of commands and arguments in one command line through two different methods.

The easiest option for running multiple commands in two methods is by creating a configuration file and defining the `-configfile` option. A configuration file contains a set of arguments to `mpirun`, one group per line.

```
-n 1 -host node1 ./io <io_args>
-n 4 -host node2 ./compute <compute_args_1>
-n 4 -host node3 ./compute <compute_args_2>
```

Alternatively, a set of options can be passed on the command line by separating each group with ":".

```
mpirun -n 1 -host node1 ./io <io_args> : -n 4 -host node2 ./compute
<compute_args_1> : -n 4 -host node3 ./compute <compute_args_2>
```

When a process is launched, the working directory will be set to the working directory of the machine where the job was launched. To change this, use the `-wdir <path>`.

Use `-env <var> <value>` to set an environment variable to a value for only one process group. Using `-genv` instead will apply the environment variable to all process groups. By default, all environment variables are propagated from the environment at launch.

11. Troubleshooting

This section provides the following troubleshooting information:

- General Intel® MPI Library troubleshooting procedures
- Typical MPI failures with corresponding output messages and behavior when a failure occurs
- Recommendations on potential root causes and solutions

11.1. General Troubleshooting Procedures

When faced with errors or failures during usage of the Intel® MPI Library, take the following general troubleshooting steps:

1. Check the *System Requirements* section and the *Known Issues* section in *Intel® MPI Library Release Notes*.
2. Check accessibility of the hosts. Run a simple non-MPI application (for example, the `hostname` utility) on the problem hosts using `mpirun`.

Example:

```
$ mpirun -ppn 1 -n 2 -hosts node01,node02 hostname
node01
node02
```

This may help reveal an environmental problem (such as, the MPI remote access tool is not configured properly), or a connectivity problem (such as, unreachable hosts).

3. Run the MPI application with debug information enabled by setting the environment variable `I_MPI_DEBUG=6` and/or `I_MPI_HYDRA_DEBUG=on`. Increase the integer value of debug level to get more detailed information. This action helps narrow down to the problematic component.
4. If you have the availability, download and install the latest version of Intel MPI Library from the [official product page](#) and check if your problem persists.
5. If the problem still persists, you can submit a ticket via [Intel® Premier Support](#).

11.2. Examples of MPI Failures

This section provides examples of typical MPI failures including descriptions, output messages, and related recommendations. The following problems which may cause MPI failures are discussed in this section:

- Communication problems
- Environmental problems
- Other problems

11.2.1. Communication Problems

Communication problems with the Intel® MPI Library are usually caused by a signal termination (`SIGTERM`, `SIGKILL`, or other signals). Such terminations may be due to a host reboot, receiving an unexpected signal, out-of-memory (OOM) manager errors and others.

To deal with such failures, you need to find out the reason for the MPI process termination (for example, by checking the system log files).

Example 1

Symptom/Error Message

```
[50:node02] unexpected disconnect completion event from [41:node01]
```

and/or

```
=====
= BAD TERMINATION OF ONE OF YOUR APPLICATION PROCESSES
= PID 20066 RUNNING AT node01
= EXIT CODE: 15
= CLEANING UP REMAINING PROCESSES
= YOU CAN IGNORE THE BELOW CLEANUP MESSAGES
=====
```

The exact node and the MPI process reported in the table may not reflect the one where the initial failure had occurred.

Cause

One of MPI processes is terminated by a signal (for example, SIGTERM or SIGKILL) on node01. The MPI application was run over the `dapl` fabric.

Solution

Try to find out the reason of the MPI process termination. This may be a host reboot, receiving an unexpected signal, OOM manager errors and others. Check the system log files.

Example 2

Symptom/Error Message

```
rank = 26, revents = 25, state = 8
Assertion failed in file ../../src/mpid/ch3/channels/nemesis/netmod/tcp/socksm.c
at line 2969: (it_plfd->revents & POLLERR) == 0
internal ABORT - process 25
Fatal error in PMPI_Alltoall: A process has failed, error stack:
PMPI_Alltoall(1062).....: MPI_Alltoall(sbuf=0x9dd7d0, scount=64, MPI_BYTE,
rbuf=0x9dc7b0, rcount=64, MPI_BYTE, comm=0x84000000) failed
MPIR_Alltoall_impl(860)...:
MPIR_Alltoall(819).....:
MPIR_Alltoall_intra(360)...:
dequeue_and_set_error(917): Communication error with rank 2rank = 45, revents =
25, state = 8
Assertion failed in file ../../src/mpid/ch3/channels/nemesis/netmod/tcp/socksm.c
at line 2969: (it_plfd->revents & POLLERR) == 0
internal ABORT - process 84
...
Fatal error in PMPI_Alltoall: A process has failed, error stack:
PMPI_Alltoall(1062).....: MPI_Alltoall(sbuf=MPI_IN_PLACE, scount=-1,
MPI_DATATYPE_NULL, rbuf=0x2ba2922b4010, rcount=8192, MPI_INT, MPI_COMM_WORLD)
failed
MPIR_Alltoall_impl(860)...:
MPIR_Alltoall(819).....:
MPIR_Alltoall_intra(265)...:
MPIC_Sendrecv_replace(658):
dequeue_and_set_error(917): Communication error with rank 84
...
```

and/or:

```
=====
= BAD TERMINATION OF ONE OF YOUR APPLICATION PROCESSES
```

```
= PID 21686 RUNNING AT node01
= EXIT CODE: 15
= CLEANING UP REMAINING PROCESSES
= YOU CAN IGNORE THE BELOW CLEANUP MESSAGES
=====
```

The exact node and the MPI process reported in the table may not reflect the one where the initial failure had occurred.

Cause

One of MPI processes is terminated by a signal (for example, `SIGTERM` or `SIGKILL`). The MPI application was run over the `tcp` fabric. In such cases, hang of the MPI application is possible.

Solution

Try to find out the reason of the MPI process termination. This may be a host reboot, receiving an unexpected signal, OOM manager errors and others. Check the system log files.

Example 3

Symptom/Error Message

```
[mpiexec@node00] control_cb (../../pm/pmiserv/pmiserv_cb.c:773): connection to
proxy 1 at host node01 failed
[mpiexec@node00] HYDT_dmxu_poll_wait_for_event
(../../tools/demux/demux_poll.c:76): callback returned error status
[mpiexec@node00] HYD_pmci_wait_for_completion
(../../pm/pmiserv/pmiserv_pmci.c:501): error waiting for event
[mpiexec@node00] main (../../ui/mpich/mpiexec.c:1063): process manager error
waiting for completion
```

Cause

The remote `pmi_proxy` process is terminated by the `SIGKILL` (9) signal on node01.

Solution

Try to find out the reason of the `pmi_proxy` process termination. This may be a host reboot, receiving an unexpected signal, OOM manager errors and others. Check the system log files.

Example 4

Symptom/Error Message

```
Failed to connect to host node01 port 22: No route to host
```

Cause

One of the MPI compute nodes (`node01`) is not available on the network. In such cases, hang of the MPI application is possible.

Solution

Check the network interfaces on the nodes and make sure the host is accessible.

Example 5

Symptom/Error Message

```
Failed to connect to host node01 port 22: Connection refused
```

Cause

The MPI remote node access mechanism is `SSH`. The `SSH` service is not running on node01.

Solution

Check the state of the `SSH` service on the nodes.

11.2.2. Environment Problems

Environmental errors may happen when there are problems with the system environment, such as mandatory system services are not running, shared resources are unavailable and so on.

When you encounter environmental errors, check the environment. For example, verify the current state of important services.

Example 1

Symptom/Error Message

```
librdmacm: Warning: couldn't read ABI version.  
librdmacm: Warning: assuming: 4  
librdmacm: Fatal: unable to get RDMA device list
```

or:

```
CMA: unable to get RDMA device list  
librdmacm: couldn't read ABI version.  
librdmacm: assuming: 4
```

Cause

The `OFED*` stack is not loaded. The application was run over the `dapl` fabric. In such cases, hang of the MPI application is possible.

Solution

See the `OFED*` documentation for details about `OFED*` stack usage.

Example 2

Symptom/Error Message

```
[0] MPI startup(): Multi-threaded optimized library  
[1] DAPL startup(): trying to open DAPL provider from I_MPI_DAPL_PROVIDER: ofa-  
v2-mlx4_0-1  
[0] DAPL startup(): trying to open DAPL provider from I_MPI_DAPL_PROVIDER: ofa-  
v2-mlx4_0-1  
[1] MPI startup(): DAPL provider ofa-v2-mlx4_0-1  
[1] MPI startup(): dapl data transfer mode  
[0] MPI startup(): DAPL provider ofa-v2-mlx4_0-1  
[0] MPI startup(): dapl data transfer mode
```

In such cases, hang of the MPI application is possible.

Cause

The Subnet Manager (`opensmd*`) service is not running. The application was run over the `dapl` fabric. The following output is provided when you set `I_MPI_DEBUG=2`.

Solution

Check the current status of the service. See the `OFED*` documentation for details on `opensmd*` usage.

Example 3

Symptom/Error Message

```
node01-mic0:MCM:2b66:e56a0b40: 2379 us(2379 us): scif_connect() to port 68,  
failed with error Connection refused  
node01-mic0:MCM:2b66:e56a0b40: 2494 us(115 us): open_hca: SCIF init ERR for  
mlx4_0
```

```
Assertion failed in file
../../../../src/mpid/ch3/channels/nemesis/netmod/dapl/dapls_module_init.c
at line 761: 0
internal ABORT - process 0
```

Cause

The `mpxyd` daemon (CCL-proxy) is not running. The application was run over the `dapl` fabric. In such cases, hang of the MPI application is possible.

Solution

Check the current status of the service. See the DAPL* documentation for details on `mpxyd` usage.

Example 4

Symptom/Error Message

```
node01-mic0:SCM:2b94:14227b40: 201 us(201 us): open_hca: ibv_get_device_list()
failed
node01-mic0:SCM:2b94:14227b40: 222 us(222 us): open_hca: ibv_get_device_list()
failed
node01-mic0:CMA:2b94:14227b40: 570 us(570 us): open_hca: getaddr_netdev ERROR:No
such device. Is ib0 configured?
...
Fatal error in MPI_Init: Other MPI error, error stack:
MPIR_Init_thread(784).....:
MPID_Init(1326).....: channel initialization failed
MPIDI_CH3_Init(141).....:
dapl_rc_setup_all_connections_20(1386): generic failure with errno = 872609295
getConnInfoKVS(849).....: PMI_KVS_Get failed
```

Cause

The `ofed-mic` service is not running. The application was run over the `dapl` fabric. In such cases, hang of the MPI application is possible.

Solution

Check the current status of the service. See the Intel® MPSS documentation for details on `ofed-mic` usage.

Example 5

Symptom/Error Message

```
pmi_proxy: line 0: exec: pmi_proxy: not found
```

Cause

The Intel® MPI Library runtime scripts are not available. A possible reason is that the shared space cannot be reached. In such cases, hang of the MPI application is possible.

Solution

Check if the shared path is available across all the nodes.

Example 6

Symptom/Error Message

```
[0] DAPL startup: RLIMIT_MEMLOCK too small
[0] MPI startup(): dapl fabric is not available and fallback fabric is not
enabled
```

or:

```
node01:SCM:1c66:3f226b40: 6815816 us: DAPL ERR reg_mr Cannot allocate memory
```

Cause

Wrong system limits: the max locked memory is too small. The application was run over the `dapl` fabric.

Solution

Check the system limits and update them if necessary. The following example shows the correct system limits configuration:

```
$ ulimit -a
core file size (blocks, -c) 0
data seg size (kbytes, -d) unlimited
scheduling priority (-e) 0
file size (blocks, -f) unlimited
pending signals (-i) 256273
max locked memory (kbytes, -l) unlimited
max memory size (kbytes, -m) unlimited
open files (-n) 1024
pipe size (512 bytes, -p) 8
POSIX message queues (bytes, -q) 819200
real-time priority (-r) 0
stack size (kbytes, -s) unlimited
cpu time (seconds, -t) unlimited
max user processes (-u) 1024
virtual memory (kbytes, -v) unlimited
file locks (-x) unlimited
```

Example 7

Symptom/Error Message

```
Are you sure you want to continue connecting (yes/no)? The authenticity of host
'node01 (<node01_ip_address>)' can't be established.
```

This message may repeat continuously until manual interruption.

Cause

The MPI remote node access mechanism is `SSH`. `SSH` is not configured properly: unexpected messages appear in the standard input (`stdin`).

Solution

Check the `SSH` connection to the problem node.

Example 8

Symptom/Error Message

```
Password:
```

Cause

The MPI remote node access mechanism is `SSH`. `SSH` is not password-less. In such cases, hang of the MPI application is possible.

Solution

Check the `SSH` settings: password-less authorization by public keys should be enabled and configured.

11.2.3. Other Problems

Example 1

Symptom/Error Message

```
cannot execute binary file
```

Cause

Wrong binary executable file format or architecture.

This error occurs when you run a binary executable file build for `x86_64` architecture on a `k10m` node (for example, on an Intel® Xeon Phi™ coprocessor). In such cases, hang of the MPI application is possible.

Solution

Verify the correctness of the binary file and the command line options.

Example 2

Symptom/Error Message

```
node01.9234ipath_userinit: assign_context command failed: Invalid argument
node01.9234Driver initialization failure on /dev/ipath (err=23)
```

Cause

Intel® True Scale IBA resource exhaustion. The MPI application was run over the `tmi` fabric.

Depending on Intel® True Scale Fabric hardware, PSM* may not support CPU over-subscription of the node. The maximum amount of processes which can be run on the node is limited and depends on a combination of the particular Intel® True Scale Fabric hardware and the amount of CPU cores.

Solution

Limit the number of MPI processes per node.

12. Using the Intel® MPI Library with the Intel® Many Integrated Core (Intel® MIC) Architecture

Using the Intel® MPI Library in combination with an Intel® MIC Architecture card is similar to using another node, but there are a few special considerations. This topic provides the information on these special considerations.

12.1. Libraries

The Intel® MIC Architecture uses different binaries and libraries, and these must be present on the card. In order to copy the appropriate files to the card, you can use the following commands:

```
(host)$ scp <installdir>/mic/bin/* host0-mic0:/bin/
(host)$ scp <installdir>/mic/lib/* host0-mic0:/lib64/
```

This assumes that the hostname of the card is `host0-mic0`. Any additional libraries needed by the application can be copied in a similar manner.

12.2. Multiple Cards

To use multiple cards for a single job, the Intel® Manycore Platform Software Stack (Intel® MPSS) needs to be configured for peer-to-peer support (see the Intel® MPSS documentation for details) and the host(s) needs to have IP forwarding enabled.

```
(host)$ sudo sysctl -w net.ipv4.ip_forward=1
```

Each host/card should be able to ping every other host/card and the launching host should be able to connect to every target, as with a classic cluster.

12.3. Using Intel® MPI Library on Intel® Xeon Phi™ Coprocessor

Intel® MPI Library for the Intel® Many Integrated Core Architecture (Intel® MIC Architecture) supports only the Intel® Xeon Phi™ coprocessor (previously codenamed: Knights Corner).

12.3.1. Building an MPI Application

To build an MPI application for the host node and the Intel® Xeon Phi™ coprocessor, follow these steps:

1. Establish the environment settings for the compiler and for the Intel® MPI Library:

```
$ . <install-dir>/compilers_and_libraries/linux/bin/compilervars.sh
intel64

$ . <install-dir>/compilers_and_libraries/linux/mpi/intel64/bin/mpivars.sh
```

2. Build your application for Intel® Xeon Phi™ coprocessor:

```
$ mpiicc -mmic myprog.c -o myprog.mic
```

3. Build your application for Intel® 64 architecture:

```
$ mpiicc myprog.c -o myprog
```

12.3.2. Running an MPI Application

To run an MPI application on the host node and the Intel® Xeon Phi™ coprocessor, do the following:

1. Ensure that NFS is properly set up between the hosts and the Intel® Xeon Phi™ coprocessor(s). For information on how to set up NFS on the Intel® Xeon Phi™ coprocessor(s), visit the Intel® Xeon Phi™ coprocessor developer community at <http://software.intel.com/en-us/mic-developer>.
2. Use the `I_MPI_MIC_POSTFIX` environment variable to append the `.mic` postfix extension when running on the Intel® Xeon Phi™ coprocessor.

```
$ export I_MPI_MIC_POSTFIX=.mic
```

3. Make sure your `~/mpi.hosts` file contains the machine names of your Intel® Xeon® host processors and the Intel® Xeon Phi™ coprocessor(s). For example:

```
$ cat ~/mpi.hosts  
clusternode1  
clusternode1-mic0
```

4. Launch the executable file from the host.

```
$ export I_MPI_MIC=on  
$ mpirun -n 4 -hostfile ~/mpi.hosts ./myprog
```

NOTE

You can also use the `-configfile` and `-machinefile` options.

To run the application on Intel® Xeon Phi™ coprocessor only, follow the steps described above and ensure that `mpi.hosts` contains only the Intel® Xeon Phi™ coprocessor name.

See Also

You can get more details in the Intel® Xeon Phi™ Coprocessor Support topic of the Intel® MPI Library Reference Manual for Linux* OS.

You can get more information about using Intel® MPI Library on Intel® Xeon Phi™ coprocessor online at [How to run Intel® Xeon Phi™ Coprocessor](#).