

Intel® MPI Library for Linux* OS

Developer Guide

Contents

Legal Information	4
1. Introduction	5
1.1. Introducing Intel® MPI Library.....	5
1.2. Conventions and Symbols.....	5
1.3. Related Information	5
2. Installation and Prerequisites	7
2.1. Installation.....	7
2.2. Prerequisite Steps.....	7
3. Compiling and Linking	9
3.1. Compiling an MPI Program	9
3.1.1. Compiling an MPI/OpenMP* Program	9
3.1.2. Adding Debug Information.....	9
3.1.3. Test MPI Programs.....	10
3.2. Compilers Support.....	10
3.3. ILP64 Support.....	10
4. Running Applications	12
4.1. Running an MPI Program	12
4.2. Running an MPI/OpenMP* Program	12
4.3. MPMD Launch Mode	13
4.4. Fabrics Control	14
4.4.1. Selecting Fabrics.....	14
4.4.2. libfabric* Support.....	14
4.4.3. OFI* Providers Support.....	15
4.5. Selecting Library Configuration	17
4.6. Job Schedulers Support.....	18
4.6.1. Altair* PBS Pro*, TORQUE*, and OpenPBS*	18
4.6.2. IBM* Platform LSF*	18
4.6.3. Parallelnavi NQS*	19
4.6.4. SLURM*	19
4.6.5. Univa* Grid Engine*	19
4.6.6. SIGINT, SIGTERM Signals Intercepting.....	19
4.6.7. Controlling Per-Host Process Placement	19
4.7. Controlling Process Placement.....	19
4.7.1. Specifying Hosts	20
4.7.2. Using Machine File	20
4.7.3. Using Argument Sets	21
4.8. Java* MPI Applications Support	21
4.8.1. Running Java* MPI applications.....	21
4.8.2. Development Recommendations	22
5. Debugging.....	23
5.1. Debugging.....	23
5.1.1. GDB*: The GNU* Project Debugger.....	23
5.1.2. DDT* Debugger	23

- 5.2. Using -gtool for Debugging..... 23
- 6. Analysis and Tuning24**
 - 6.1. Displaying MPI Debug Information..... 24
 - 6.2. Tracing Applications 25
 - 6.2.1. High-Level Performance Analysis 25
 - 6.2.2. Tracing Applications 25
 - 6.3. Interoperability with Other Tools through -gtool..... 26
 - 6.4. MPI Tuning..... 27
- 7. Additional Supported Features28**
 - 7.1. Asynchronous Progress Control..... 28
 - 7.2. Multiple Endpoints Support 28
 - 7.2.1. MPI_THREAD_SPLIT Programming Model..... 28
 - 7.2.2. Threading Runtimes Support..... 30
 - 7.2.3. Code Change Guide 30

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 known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel technologies features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

Intel, the Intel logo, VTune, Xeon, and Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

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

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

Copyright 2003-2018 Intel Corporation.

This software and the related documents are Intel copyrighted materials, and your use of them is governed by the express license under which they were provided to you (**License**). Unless the License provides otherwise, you may not use, modify, copy, publish, distribute, disclose or transmit this software or the related documents without Intel's prior written permission.

This software and the related documents are provided as is, with no express or implied warranties, other than those that are expressly stated in the License.

1. Introduction

The *Intel® MPI Library Developer Guide* explains how to use the Intel® MPI Library in some common usage scenarios. It provides information regarding compiling, running, debugging, tuning, and analyzing MPI applications, as well as troubleshooting information.

This *Developer Guide* helps a user familiar with the message passing interface start using the Intel® MPI Library. For full information, see the *Intel® MPI Library Developer Reference*.

1.1. Introducing Intel® MPI Library

The Intel® MPI Library is a multi-fabric message-passing library that implements the Message Passing Interface, version 3.1 (MPI-3.1) 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-3.1 functions as your needs dictate

1.2. Conventions and Symbols

The following conventions are 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).

1.3. 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 Developer Reference](#) 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® Software Documentation Library](#)
- [Intel® Software Products Support](#)

2. Installation and Prerequisites

2.1. Installation

If you have a previous version of the Intel® MPI Library for Linux* OS installed, you do not need to uninstall it before installing a newer 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_<version>.<update>.<package#>/linux/mpi`.

There are two 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>/intel64` directory.

2.2. Prerequisite Steps

Before you start using any of the Intel® MPI Library functionality, make sure to establish the proper environment for Intel MPI Library. Follow these steps:

1. Set up the Intel MPI Library environment. Source the `mpivars.[c]sh` script:

```
$ . <installdir>/intel64/bin/mpivars.sh
```

By default, `<installdir>` is

```
/opt/intel/compilers_and_libraries_<version>.<update>.<package>/linux/mpi.
```

2. To run an MPI application on a cluster, Intel MPI Library needs to know names of all its nodes. Create a text file listing the cluster node names. The format of the file is one name per line, and the lines starting with `#` are ignored. To get the name of a node, use the `hostname` utility.

A sample host file may look as follows:

```
$ cat ./hosts
# this line is ignored
clusternode1
clusternode2
clusternode3
clusternode4
```

3. For communication between cluster nodes, in most cases Intel MPI Library uses the SSH protocol. You need to establish a passwordless SSH connection to ensure proper communication of MPI processes. Intel MPI Library provides the `sshconnectivity.exp` script that helps you do the job. It automatically generates and distributes SSH authentication keys over the nodes.

The script is located at

`/opt/intel/parallel_studio_xe_<version>.<update>.<package>/bin` by default. Run the script and pass the previously created host file as an argument.

If the script does not work for your system, try generating and distributing authentication keys manually.

After completing these steps, you are ready to use Intel MPI Library.

3. Compiling and Linking

3.1. Compiling an MPI Program

This topic describes the basic steps required to compile and link an MPI program, using the Intel® MPI Library SDK.

To simplify linking with MPI library files, Intel MPI Library provides a set of compiler wrapper scripts with the `mpi` prefix for all supported compilers. To compile and link an MPI program, do the following:

1. Make sure you have a compiler in your `PATH` environment variable. For example, to check if you have the Intel® C Compiler, enter the command:

```
$ which icc
```

If the command is not found, add the full path to your compiler into the `PATH`. For Intel® compilers, you can source the `compilervars.[c]sh` script to set the required environment variables.

2. Source the `mpivars.[c]sh` script to set up the proper environment for Intel MPI Library. For example, using the Bash* shell:

```
$ . <installdir>/intel64/bin/mpivars.sh
```

3. Compile your MPI program using the appropriate compiler wrapper script. For example, to compile a C program with the Intel® C Compiler, use the `mpiicc` script as follows:

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

You will get an executable file `myprog` in the current directory, which you can start immediately. For instructions of how to launch MPI applications, see [Running an MPI Program](#).

NOTE

By default, the resulting executable file is linked with the multi-threaded optimized library. If you need to use another library configuration, see [Selecting Library Configuration](#).

For details on the available compiler wrapper scripts, see the *Developer Reference*.

See Also

Intel® MPI Library Developer Reference, section *Command Reference > Compiler Commands*

3.1.1. Compiling an MPI/OpenMP* Program

To compile a hybrid MPI/OpenMP* program using the Intel® compiler, use the `-qopenmp` option. For example:

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

This enables the underlying compiler to generate multi-threaded code based on the OpenMP* pragmas in the source. For details on running such programs, refer to [Running an MPI/OpenMP* Program](#).

3.1.2. Adding Debug Information

If you need to debug your application, add the `-g` option to the compilation command line. For example:

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

This adds debug information to the resulting binary, enabling you to debug your application. Debug information is also used by analysis tools like Intel® Trace Analyzer and Collector to map the resulting trace file to the source code.

3.1.3. Test MPI Programs

Intel® MPI Library comes with a set of source files for simple MPI programs that enable you to test your installation. Test program sources are available for all supported programming languages and are located at `<installdir>/test`, where `<installdir>` is `/opt/intel/compilers_and_libraries_<version>.x.xxx/linux/mpi` by default.

3.2. Compilers Support

Intel® MPI Library supports the GCC* and Intel® compilers out of the box. It uses binding libraries to provide support for different `glibc` versions and different compilers. These libraries provide C++, Fortran 77, Fortran 90, and Fortran 2008 interfaces.

The following binding libraries are used for GCC* and Intel® compilers:

- `libmpicxx.{a|so}` – for g++ version 3.4 or higher
- `libmpifort.{a|so}` – for g77/gfortran interface for GCC and Intel® compilers

Your application gets linked against the correct GCC* and Intel® compilers binding libraries, if you use one of the following compiler wrappers: `mpicc`, `mpicxx`, `mpifc`, `mpif77`, `mpif90`, `mpigcc`, `mpigxx`, `mpiicc`, `mpiiipc`, or `mpiifort`.

For other compilers, PGI* and Absoft* in particular, there is a binding kit that allows you to add support for a certain compiler to the Intel® MPI Library. This binding kit provides all the necessary source files, convenience scripts, and instructions you need, and is located in the `<install_dir>/binding` directory.

To add support for the PGI* C, PGI* Fortran 77, Absoft* Fortran 77 compilers, you need to manually create the appropriate wrapper script (see instructions in the binding kit *Readme*). When using these compilers, keep in mind the following limitations:

- Your PGI* compiled source files must not transfer `long double` entities
- Your Absoft* based build procedure must use the `-g77`, `-B108` compiler options

To add support for the PGI* C++, PGI* Fortran 95, Absoft* Fortran 95, and GNU* Fortran 95 (4.0 and newer) compilers, you need to build extra binding libraries. Refer to the binding kit *Readme* for detailed instructions.

3.3. ILP64 Support

The term ILP64 denotes that integer, long, and pointer data entities all occupy 8 bytes. This differs from the more conventional LP64 model, in which only long and pointer data entities occupy 8 bytes while integer entities occupy 4 bytes. More information on the historical background and the programming model philosophy can be found, for example, in http://www.unix.org/version2/whatsnew/lp64_wp.html

Intel® MPI Library provides support for the ILP64 model for Fortran applications. To enable the ILP64 mode, do the following:

Use the Fortran compiler wrapper option `-i8` for separate compilation and the `-ilp64` option for separate linking. For example:

```
$ mpiifort -i8 -c test.f
$ mpiifort -ilp64 -o test test.o
```

For simple programs, use the Fortran compiler wrapper option `-i8` for compilation and linkage. Specifying `-i8` will automatically assume the ILP64 library. For example:

```
$ mpiifort -i8 test.f
```

When running the application, use the `-ilp64` option to preload the ILP64 interface. For example:

```
$ mpirun -ilp64 -n 2 ./myprog
```

The following limitations are present in the Intel MPI Library in regard to this functionality:

- Data type counts and other arguments with values larger than $2^{31} - 1$ are not supported.
- Special MPI types `MPI_FLOAT_INT`, `MPI_DOUBLE_INT`, `MPI_LONG_INT`, `MPI_SHORT_INT`, `MPI_2INT`, `MPI_LONG_DOUBLE_INT`, `MPI_2INTEGER` are not changed and still use a 4-byte integer field.
- Predefined communicator attributes `MPI_APPNUM`, `MPI_HOST`, `MPI_IO`, `MPI_LASTUSED`, `MPI_TAG_UB`, `MPI_UNIVERSE_SIZE`, and `MPI_WTIME_IS_GLOBAL` are returned by the functions `MPI_GET_ATTR` and `MPI_COMM_GET_ATTR` as 4-byte integers. The same holds for the predefined attributes that may be attached to the window and file objects.
- Do not use the `-i8` option to compile MPI callback functions, such as error handling functions, or user-defined reduction operations.
- Do not use the `-i8` option with the deprecated functions that store or retrieve the 4-byte integer attribute (for example, `MPI_ATTR_GET`, `MPI_ATTR_PUT`, etc.). Use their recommended alternatives instead (`MPI_COMM_GET_ATTR`, `MPI_COMM_SET_ATTR`, etc).
- If you want to use the Intel® Trace Collector with the Intel MPI Library ILP64 executable files, you must use a special Intel Trace Collector library. If necessary, the `mpiifort` compiler wrapper will select the correct Intel Trace Collector library automatically.
- There is currently no support for C and C++ applications.

4. Running Applications

4.1. Running an MPI Program

Before running an MPI program, place it to a shared location and make sure it is accessible from all cluster nodes. Alternatively, you can have a local copy of your program on all the nodes. In this case, make sure the paths to the program match.

Run the MPI program using the `mpirun` command. The command line syntax is as follows:

```
$ mpirun -n <# of processes> -ppn <# of processes per node> -f <hostfile> ./myprog
```

For example:

```
$ mpirun -n 4 -ppn 2 -f hosts ./myprog
```

In the command line above:

- `-n` sets the number of MPI processes to launch; if the option is not specified, the process manager pulls the host list from a job scheduler, or uses the number of cores on the machine.
- `-ppn` sets the number of processes to launch on each node; if the option is not specified, processes are assigned to the physical cores on the first node; if the number of cores is exceeded, the next node is used.
- `-f` specifies the path to the host file listing the cluster nodes; alternatively, you can use the `-hosts` option to specify a comma-separated list of nodes; if hosts are not specified, the local node is used.
- `myprog` is the name of your MPI program.

The `mpirun` command is a wrapper around the `mpiexec.hydra` command, which invokes the Hydra process manager. Consequently, you can use all `mpiexec.hydra` options with the `mpirun` command.

For the list of all available options, run `mpirun` with the `-help` option, or see the *Intel® MPI Library Developer Reference*, section *Command Reference > Hydra Process Manager Command*.

NOTE

The commands `mpirun` and `mpiexec.hydra` are interchangeable. However, you are recommended to use the `mpirun` command for the following reasons:

- You can specify all `mpiexec.hydra` options with the `mpirun` command.
- 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 `mpirun` when an MPI program is running under a batch scheduler or job manager.

See Also

[Controlling Process Placement](#)
[Job Schedulers Support](#)

4.2. Running an MPI/OpenMP* Program

To run a hybrid MPI/OpenMP* program, follow these steps:

1. Make sure the thread-safe (debug or release, as desired) Intel® MPI Library configuration is enabled (this is the default behavior). To switch to such a configuration, source `mpivars.[c]sh` with the appropriate argument. For example:

```
$ source mpivars.sh release_mt
```

2. Set the `I_MPI_PIN_DOMAIN` environment variable to specify the desired process pinning scheme. The recommended value is `omp`:

```
$ export I_MPI_PIN_DOMAIN=omp
```

This sets the process pinning domain size to be equal to `OMP_NUM_THREADS`. Therefore, if for example `OMP_NUM_THREADS` is equal to 4, each MPI process can create up to four threads within the corresponding domain (set of logical processors). If `OMP_NUM_THREADS` is not set, each node is treated as a separate domain, which allows as many threads per MPI process as there are cores.

NOTE

For pinning OpenMP* threads within the domain, use the Intel® compiler `KMP_AFFINITY` environment variable. See the Intel compiler documentation for more details.

3. Run your hybrid program as a regular MPI program. You can set the `OMP_NUM_THREADS` and `I_MPI_PIN_DOMAIN` variables directly in the launch command. For example:

```
$ mpirun -n 4 -genv OMP_NUM_THREADS=4 -genv I_MPI_PIN_DOMAIN=omp ./myprog
```

See Also

Intel® MPI Library Developer Reference, section *Tuning Reference > Process Pinning > Interoperability with OpenMP**.

4.3. MPMD Launch Mode

Intel® MPI Library supports the multiple programs, multiple data (MPMD) launch mode. There are two ways to do this.

The easiest way is to create a configuration file and pass it to the `-configfile` option. A configuration file should contain a set of arguments for `mpirun`, one group per line. For example:

```
$ cat ./mpmd_config
-n 1 -host node1 ./io <io_args>
-n 4 -host node2 ./compute <compute_args_1>
-n 4 -host node3 ./compute <compute_args_2>
$ mpirun -configfile mpmd_config
```

Alternatively, you can pass a set of options to the command line by separating each group with a colon:

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

The examples above are equivalent. The `io` program is launched as one process on `node1`, and the `compute` program is launched on `node2` and `node3` as four processes on each.

When an MPI job is launched, the working directory is set to the working directory of the machine where the job is launched. To change this, use the `-wdir <path>`.

Use `-env <var> <value>` to set an environment variable for only one argument set. Using `-genv` instead applies the environment variable to all argument sets. By default, all environment variables are propagated from the environment during the launch.

4.4. Fabrics Control

The Intel® MPI Library switched from the Open Fabrics Alliance* (OFA) framework to the Open Fabrics Interfaces* (OFI) framework and currently supports libfabric*.

NOTE

The supported fabric environment has changed since Intel® MPI Library 2017 Update 1. The `dapl`, `tcp`, `tmi`, and `ofa` fabrics are now deprecated.

OFI is a framework focused on exporting communication services to applications. OFI is specifically designed to meet the performance and scalability requirements of high-performance computing (HPC) applications running in a tightly coupled network environment. The key components of OFI are application interfaces, provider libraries, kernel services, daemons, and test applications.

Libfabric is a library that defines and exports the user-space API of OFI, and is typically the only software that applications deal with directly. The libfabric's API does not depend on the underlying networking protocols, as well as on the implementation of the particular networking devices, over which it may be implemented. OFI is based on the notion of application centric I/O, meaning that the libfabric library is designed to align fabric services with application needs, providing a tight semantic fit between applications and the underlying fabric hardware. This reduces overall software overhead and improves application efficiency when transmitting or receiving data over a fabric.

4.4.1. Selecting Fabrics

Intel® MPI Library enables you to select a communication fabric at runtime without having to recompile your application. By default, it automatically selects the most appropriate fabric based on your software and hardware configuration. This means that in most cases you do not have to bother about manually selecting a fabric.

However, in certain situations specifying a particular communication fabric can boost performance of your application. The following fabrics are available:

Fabric	Network hardware and software used
shm	Shared memory (for intra-node communication only).
ofi	OpenFabrics Interfaces* (OFI)-capable network fabrics, such as Intel® True Scale Fabric, Intel® Omni-Path Architecture, InfiniBand* and Ethernet (through OFI API).

Use the `I_MPI_FABRICS` environment variable to specify a fabric. The description is available in the *Developer Reference*, section *Tuning Reference > Fabrics Control*.

4.4.2. libfabric* Support

By default, `mpivars.[c]sh` sets the environment to libfabric shipped with the Intel MPI Library.

To disable this, use the `I_MPI_OFI_LIBRARY_INTERNAL` environment variable or the `-ofi_internal` option passed to the `mpivars.[c]sh` script:

```
$ source <installdir>/intel64/bin/mpivars.sh -ofi_internal=0 # do not set the
environment to libfabric from the Intel MPI Library

$ source <installdir>/intel64/bin/mpivars.sh -ofi_internal=1 # set the environment
to libfabric from the Intel MPI Library

$ source <installdir>/intel64/bin/mpivars.sh # a short form of -ofi-internal=1
```

NOTE

Set the `I_MPI_DEBUG` environment variable to 1 before running an MPI application to see the libfabric version and provider.

Example

```
$ export I_MPI_DEBUG=1
$ mpiexec -n 1 IMB-MPI1 -help
[0] MPI startup(): libfabric version: 1.5.0
[0] MPI startup(): libfabric provider: psm2
...
```

See Also

For more information, see [Working with libfabric* on Intel® MPI Library Cluster Systems](#).

Intel® MPI Library Developer Reference, section *Environment Variables Reference > Environment Variables for Fabrics Control > OFI-capable Network Fabrics Control*

4.4.3. OFI* Providers Support

Intel® MPI Library supports psm2, sockets, verbs, and RxM OFI* providers. Each OFI provider is built as a separate dynamic library to ensure that a single libfabric* library can be run on top of different network adapters. To specify the path to provider libraries, set the `FI_PROVIDER_PATH` environment variable.

To get a full list of environment variables available for configuring OFI, run the following command:

```
$ fi_info -e
```

psm2

The PSM2 provider runs over the PSM 2.x interface supported by the Intel® Omni-Path Fabric. PSM 2.x has all the PSM 1.x features, plus a set of new functions with enhanced capabilities. Since PSM 1.x and PSM 2.x are not application binary interface (ABI) compatible, the PSM2 provider works with PSM 2.x only and does not support Intel True Scale Fabric.

The following runtime parameters can be used:

Name	Description
<code>FI_PSM2_INJECT_SIZE</code>	Define the maximum message size allowed for <code>fi_inject</code> and <code>fi_tinject</code> calls. The default value is 64.
<code>FI_PSM2_LAZY_CONN</code>	Control the connection mode established between PSM2 endpoints that OFI endpoints are built on top of. When set to 0 (eager connection mode), connections are established when addresses are inserted into the address vector. When set to 1 (lazy connection mode), connections are established when addresses are used the first time in communication.
	<p>NOTE</p> <p>Lazy connection mode may reduce the start-up time on large systems at the expense of higher data path overhead.</p>

sockets

The sockets provider is a general purpose provider that can be used on any system that supports TCP sockets. The provider is not intended to provide performance improvements over regular TCP sockets, but rather to allow developers to write, test, and debug application code even on platforms that do not have high-performance fabric hardware. The sockets provider supports all libfabric provider requirements and interfaces.

The following runtime parameters can be used:

Name	Description
FI_SOCKETS_IFACE	Define the prefix or the name of the network interface. By default, it uses any.

verbs

The verbs provider enables applications using OFI to be run over any verbs hardware (InfiniBand*, iWarp*, and so on). It uses the Linux Verbs API for network transport and provides a translation of OFI calls to appropriate verbs API calls. It uses librdmacm for communication management and libibverbs for other control and data transfer operations.

The verbs provider uses RxM utility provider to emulate `FI_EP_RDM` endpoint over verbs `FI_EP_MSG` endpoint by default. The verbs provider with `FI_EP_RDM` endpoint can be used instead of RxM by setting the `FI_PROVIDER=^ofi_rxm` runtime parameter.

The following runtime parameters can be used:

Name	Description
FI_VERBS_INLINE_SIZE	Define the maximum message size allowed for <code>fi_inject</code> and <code>fi_tinject</code> calls. The default value is 64.
FI_VERBS_IFACE	Define the prefix or the full name of the network interface associated with the verbs device. By default, it is <code>ib</code> .
FI_VERBS_MR_CACHE_ENABLE	Enable Memory Registration caching. The default value is 0. NOTE Set the value to 1 to increase bandwidth for medium and large messages.

Dependencies

The verbs provider requires `libibverbs` (v1.1.8 or newer) and `librdmacm` (v1.0.16 or newer). If you are compiling libfabric from source and want to enable verbs support, it is essential to have the matching header files for the above two libraries. If the libraries and header files are not in default paths, specify them in the `CFLAGS`, `LDFLAGS`, and `LD_LIBRARY_PATH` environment variables.

RxM

The RxM (RDM over MSG) provider (`ofi_rxm`) is a utility provider that supports `FI_EP_RDM` endpoint emulated over `FI_EP_MSG` endpoint of the core provider.

RxM provider requires the core provider to support the following features:

- MSG endpoints (FI_EP_MSG)
- FI_MSG transport (to support data transfers)
- FI_RMA transport (to support rendezvous protocol for large messages and RMA transfers)
- FI_OPT_CM_DATA_SIZE of at least 24 bytes

The following runtime parameters can be used:

Name	Description
FI_OFI_RXM_BUFFER_SIZE	Define the transmit buffer size/inject size. Messages of smaller size are transmitted via an eager protocol and those above would be transmitted via a rendezvous protocol. Transmitted data is copied up to the specified size. By default, the size is 16k.
FI_OFI_RXM_SAR_LIMIT	Control the RxM SAR (Segmentation and Reassembly) protocol. Messages of greater size are transmitted via rendezvous protocol.
FI_OFI_RXM_USE_SRX	Control the RxM receive path. If the variable is set to 1, the RxM uses Shared Receive Context of the core provider. The default value is 0.
	<p>NOTE</p> <p>Setting this variable to 1 improves memory consumption, but may increase small message latency as a side-effect.</p>

4.5. Selecting Library Configuration

You can specify a particular configuration of the Intel® MPI Library to be used, depending on your purposes. This can be a library optimized for multi-threading debug or release version with the global or per-object lock.

To specify the configuration, source the `mpivars.[c]sh` script with `release`, `debug`, `release_mt`, or `debug_mt` argument. For example:

```
$ . <installdir>/intel64/bin/mpivars.sh release
```

You can use the following arguments:

Argument	Definition
<code>release</code>	Set this argument to use multi-threaded optimized library (with the global lock). This is the default value
<code>debug</code>	Set this argument to use multi-threaded debug library (with the global lock)
<code>release_mt</code>	Set this argument to use multi-threaded optimized library (with per-object lock for the thread-split model)
<code>debug_mt</code>	Set this argument to use multi-threaded debug library (with per-object lock for the thread-split model)

NOTE

You do not need to recompile the application to change the configuration. Source the `mpivars.[c]sh` script with appropriate arguments before an application launch.

Alternatively, if your shell does not support sourcing with arguments, you can use the `I_MPI_LIBRARY_KIND` environment variable to set an argument for `mpivars.[c]sh`. See the *Intel® MPI Library Developer Reference* for details.

4.6. Job Schedulers Support

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*

The support is implemented in the `mpirun` wrapper script. `mpirun` detects the job scheduler under which it is started by checking specific environment variables and then chooses the appropriate method to start an application.

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

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

An example of a batch job script may look as follows:

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

4.6.2. IBM* Platform LSF*

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

For example, to submit a job, run the command:

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

4.6.3. 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, `/usr/bin/plesh` is used as remote shell by the process manager during startup.

4.6.4. 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.

For example, to submit a job, run the command:

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

4.6.5. 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.

4.6.6. 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:

1. Use the following command to see available queues:

```
$ qconf -sql
```

2. Execute the following command to modify the queue settings:

```
$ qconf -mq <queue_name>
```

3. Find `terminate_method` and change signal to `SIGTERM`.
4. Save queue configuration.

4.6.7. Controlling Per-Host Process Placement

When using a job scheduler, by default Intel MPI Library uses per-host process placement provided by the scheduler. This means that the `-ppn` option has no effect. To change this behavior and control process placement through `-ppn` (and related options and variables), use the `I_MPI_JOB_RESPECT_PROCESS_PLACEMENT` environment variable:

```
$ export I_MPI_JOB_RESPECT_PROCESS_PLACEMENT=off
```

4.7. Controlling Process Placement

Placement of MPI processes over the cluster nodes plays a significant role in application performance. Intel® MPI Library provides several options to control process placement.

By default, when you run an MPI program, the process manager launches all MPI processes specified with `-n` on the current node. If you use a job scheduler, processes are assigned according to the information received from the scheduler.

4.7.1. Specifying Hosts

You can explicitly specify the nodes on which you want to run the application using the `-hosts` option. This option takes a comma-separated list of node names as an argument. Use the `-ppn` option to specify the number of processes per node. For example:

```
$ mpirun -n 4 -ppn 2 -hosts node1,node2 ./testc
Hello world: rank 0 of 4 running on node1
Hello world: rank 1 of 4 running on node1
Hello world: rank 2 of 4 running on node2
Hello world: rank 3 of 4 running on node2
```

To get the name of a node, use the `hostname` utility.

An alternative to using the `-hosts` option is creation of a host file that lists the cluster nodes. The format of the file is one name per line, and the lines starting with `#` are ignored. Use the `-f` option to pass the file to `mpirun`. For example:

```
$ cat ./hosts
#nodes
node1
node2
$ mpirun -n 4 -ppn 2 -f hosts ./testc
```

This program launch produces the same output as the previous example.

If the `-ppn` option is not specified, the process manager assigns as many processes to the first node as there are physical cores on it. Then the next node is used. That is, assuming there are four cores on `node1` and you launch six processes overall, four processes are launched on `node1`, and the remaining two processes are launched on `node2`. For example:

```
$ mpirun -n 6 -hosts node1,node2 ./testc
Hello world: rank 0 of 6 running on node1
Hello world: rank 1 of 6 running on node1
Hello world: rank 2 of 6 running on node1
Hello world: rank 3 of 6 running on node1
Hello world: rank 4 of 6 running on node2
Hello world: rank 5 of 6 running on node2
```

NOTE

If you use a job scheduler, specifying hosts is unnecessary. The processes manager uses the host list provided by the scheduler.

4.7.2. Using Machine File

A machine file is similar to a host file with the only difference that you can assign a specific number of processes to particular nodes directly in the file. Contents of a sample machine file may look as follows:

```
$ cat ./machines
node1:2
node2:2
```

Specify the file with the `-machine` option. Running a simple test program produces the following output:

```
$ mpirun -machine machines ./testc
Hello world: rank 0 of 4 running on node1
Hello world: rank 1 of 4 running on node1
```

```
Hello world: rank 2 of 4 running on node2
Hello world: rank 3 of 4 running on node2
```

4.7.3. Using Argument Sets

Argument sets are unique groups of arguments specific to a particular node. Combined together, the argument sets make up a single MPI job. You can provide argument sets on the command line, or in a configuration file. To specify a node, use the `-host` option.

On the command line, argument sets should be separated by a colon ':'. Global options (applied to all argument sets) should appear first, and local options (applied only to the current argument set) should be specified within an argument set. For example:

```
$ mpirun -genv I_MPI_DEBUG=2 -host node1 -n 2 ./testc : -host node2 -n 2 ./testc
```

In the configuration file, each argument set should appear on a new line. Global options should appear on the first line of the file. For example:

```
$ cat ./config
-genv I_MPI_DEBUG=2
-host node1 -n 2 ./testc
-host node2 -n 2 ./testc
```

Specify the configuration file with the `-configfile` option:

```
$ mpirun -configfile config
Hello world: rank 0 of 4 running on node1
Hello world: rank 1 of 4 running on node1
Hello world: rank 2 of 4 running on node2
Hello world: rank 3 of 4 running on node2
```

See Also

[Controlling Process Placement with the Intel® MPI Library \(online article\)](#)
[Job Schedulers Support](#)

4.8. Java* MPI Applications Support

Intel® MPI Library provides an experimental feature to enable support for Java MPI applications. Java bindings are available for a subset of MPI-2 routines. For a full list of supported routines, refer to the Developer Reference, section *Miscellaneous > Java* Bindings for MPI-2 Routines*.

4.8.1. Running Java* MPI applications

Follow these steps to set up the environment and run your Java* MPI application:

1. Source `mpivars.sh` from the Intel® MPI Library package to set up all required environment variables, including `LIBRARY_PATH` and `CLASSPATH`.
2. Build your Java MPI application as usual.
3. Update `CLASSPATH` with the path to the `jar` application or pass it explicitly with the `-cp` option of the `java` command.
4. Run your Java MPI application using the following command:

```
$ mpirun <options> java <app>
```

where:

- o `<options>` is a list of `mpirun` options
- o `<app>` is the main class of your Java application

For example:

```
$ mpirun -n 8 -ppn 1 -f ./hostfile java mpi.samples.Allreduce
```

Sample Java MPI applications are available in the `<install_dir>/test` folder.

4.8.2. Development Recommendations

You can use the following tips when developing Java* MPI applications:

- To reduce memory footprint, you can use Java direct buffers as buffer parameters of collective operations in addition to using Java arrays. This approach allows you to allocate the memory out of the JVM heap and avoid additional memory copying when passing the pointer to the buffer from JVM to the native layer.
- When you create Java MPI entities such as `Group`, `Comm`, `Datatype`, and similar, memory is allocated on the native layer and is not tracked by the garbage collector. Therefore, this memory must be released explicitly. Pointers to the allocated memory are stored in a special pool and can be deallocated using one of the following methods:
 - `entity.free()`: frees the memory backing the `entity` Java object, which can be an instance of `Comm`, `Group`, etc.
 - `AllocablePool.remove(entity)`: frees the memory backing the `entity` Java object, which can be an instance of `Comm`, `Group`, etc.
 - `AllocablePool.cleanUp()`: explicitly deallocates the memory backing all Java MPI objects created by that moment.
 - `MPI.Finalize()`: implicitly deallocates the memory backing all Java MPI objects and that has not been explicitly deallocated by that moment.

5. Debugging

This section explains how to debug MPI applications using the debugger tools

5.1. Debugging

Intel® MPI Library supports the GDB* and Allinea* DDT debuggers for debugging MPI applications. Before using a debugger, make sure you have the application debug symbols available. To generate debug symbols, compile your application with the `-g` option.

5.1.1. 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 on 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 process.

5.1.2. DDT* Debugger

You can debug MPI applications using the Allinea* DDT* debugger. Intel does not provide support for this debugger, you should obtain the support from Allinea*. According to the DDT documentation, DDT supports the Express Launch feature for the Intel® MPI Library. You can debug your application as follows:

```
$ ddt mpirun -n <# of processes> [<other mpirun arguments>] <executable>
```

If you have issues with the DDT debugger, refer to the DDT documentation for help.

5.2. Using `-gtool` for Debugging

The `-gtool` runtime option can help you with debugging, when attaching to several processes at once. Instead of attaching to each process individually, you can specify all the processes in a single command line. For example:

```
$ mpirun -n 16 -gtool "gdb:3,5,7-9=attach" ./myprog
```

The command line above attaches the GNU* Debugger (GDB*) to processes 3, 5, 7, 8 and 9.

See Also

Intel® MPI Library Developer Reference, section *Command Reference > Hydra Process Manager Command > Global Options > gtool Options*

6. Analysis and Tuning

Intel® MPI Library provides a variety of options for analyzing MPI applications. Some of these options are available within the Intel MPI Library, while some require additional analysis tools. For such tools, Intel MPI Library provides compilation and runtime options and environment variables for easier interoperability.

6.1. Displaying MPI Debug Information

The `I_MPI_DEBUG` environment variable provides a convenient way to get detailed information about an MPI application at runtime. You can set the variable value from 0 (the default value) to 1000. The higher the value, the more debug information you get. For example:

```
$ mpirun -genv I_MPI_DEBUG=2 -n 2 ./testc
[1] MPI startup(): Internal info: pinning initialization was done
[0] MPI startup(): Internal info: pinning initialization was done
...
```

NOTE

High values of `I_MPI_DEBUG` can output a lot of information and significantly reduce performance of your application. A value of `I_MPI_DEBUG=5` is generally a good starting point, which provides sufficient information to find common errors.

By default, each printed line contains the MPI rank number and the message. You can also print additional information in front of each message, like process ID, time, host name and other information, or exclude some information printed by default. You can do this in two ways:

- Add the '+' sign in front of the debug level number. In this case, each line is prefixed by the string `<rank>#<pid>@<hostname>`. For example:

```
$ mpirun -genv I_MPI_DEBUG=+2 -n 2 ./testc
[0#3520@clusternode1] MPI startup(): Multi-threaded optimized library
...
```

To exclude any information printed in front of the message, add the '-' sign in a similar manner.

- Add the appropriate flag after the debug level number to include or exclude some information. For example, to include time but exclude the rank number:

```
$ mpirun -genv I_MPI_DEBUG=2,time,norank -n 2 ./testc
11:59:59 MPI startup(): Multi-threaded optimized library
...
```

For the list of all available flags, see description of `I_MPI_DEBUG` in the *Developer Reference*.

To redirect the debug information output from `stdout` to `stderr` or a text file, use the `I_MPI_DEBUG_OUTPUT` environment variable:

```
$ mpirun -genv I_MPI_DEBUG=2 -genv I_MPI_DEBUG_OUTPUT=/tmp/debug_output.txt -n 2
./testc
```

Note that the output file name should not be longer than 256 symbols.

See Also

Intel® MPI Library Developer Reference, section *Miscellaneous > Other Environment Variables > I_MPI_DEBUG*

6.2. Tracing Applications

Intel® MPI Library provides a variety of options for analyzing MPI applications. Some of these options are available within the Intel MPI Library, while some require additional analysis tools. For such tools, Intel MPI Library provides compilation and runtime options and environment variables for easier interoperability.

Intel® MPI Library provides tight integration with the Intel® Trace Analyzer and Collector, which enables you to analyze MPI applications and find errors in them. Intel® MPI Library has several compile- and runtime options to simplify the application analysis. Apart from the Intel Trace Analyzer and Collector, there is also a tool called Application Performance Snapshot intended for a higher level MPI analysis.

Both of the tools are available as part of the Intel® Parallel Studio XE Cluster Edition. Before proceeding to the next steps, make sure you have these product installed.

6.2.1. High-Level Performance Analysis

For a high-level application analysis, Intel provides a lightweight analysis tool Application Performance Snapshot (APS), which can analyze MPI and non-MPI applications. The tool provides general information about the application, such as MPI and OpenMP* utilization time and load balance, MPI operations usage, memory and disk usage, and other information. This information enables you to get a general idea about the application performance and identify spots for a more thorough analysis.

Follow these steps to analyze an application with the APS:

1. Set up the environment for the compiler, Intel MPI Library and APS:

```
$ source
<psxe_installdir>/compilers_and_libraries_<version>.<update>.<package>/linux/
bin/compilervars.sh intel64
$ source
<psxe_installdir>/compilers_and_libraries_<version>.<update>.<package>/linux/
mpi/intel64/bin/mpivars.sh
$ source <psxe_installdir>/performance_snapshots/apsvars.sh
```

2. Run your application with the `-aps` option of `mpirun`:

```
$ mpirun -n 4 -aps ./myprog
```

APS will generate a directory with the statistics files `aps_result_<date>-<time>`.

3. Launch the `aps-report` tool and pass the generated statistics to the tool:

```
$ aps-report ./aps_result_<date>-<time>
```

You will see the analysis results printed in the console window. Also, APS will generate an HTML report `aps_report_<date>_<time>.html` containing the same information.

For more details, refer to the *Application Performance Snapshot User's Guide*.

6.2.2. Tracing Applications

To analyze an application with the Intel Trace Analyzer and Collector, first you need generate a trace file of your application, and then open this file in Intel® Trace Analyzer to analyze communication patterns, time utilization, etc. Tracing is performed by preloading the Intel® Trace Collector profiling library at runtime, which intercepts all MPI calls and generates a trace file. Intel MPI Library provides the `-trace (-t)` option to simplify this process.

Complete the following steps:

1. Set up the environment for the Intel MPI Library and Intel Trace Analyzer and Collector:

```
$ source <mpi_installdir>/intel64/bin/mpivars.sh
$ source <itac_installdir>/intel64/bin/itacvars.sh
```

- Trace your application with the Intel Trace Collector:

```
$ mpirun -trace -n 4 ./myprog
```

As a result, a trace file `.stf` is generated. For the example above, it is `myprog.stf`.

- Analyze the application with the Intel Trace Analyzer:

```
$ traceanalyzer ./myprog.stf &
```

The workflow above is the most common scenario of tracing with the Intel Trace Collector. For other tracing scenarios, see the Intel Trace Collector documentation.

See Also

[Application Performance Snapshot User's Guide](#)
[Intel® Trace Collector User and Reference Guide](#)

6.3. Interoperability with Other Tools through `-gtool`

To simplify interoperability with other analysis tools, Intel® MPI Library provides the `-gtool` option (also available as the `I_MPI_GTOOL` environment variable). By using the `-gtool` option you can analyze specific MPI processes with such tools as Intel® VTune™ Amplifier XE, Intel® Advisor, Valgrind* and other tools through the `mpiexec.hydra` or `mpirun` commands.

Without the `-gtool` option, to analyze an MPI process with the VTune Amplifier, for example, you have to specify the relevant command in the corresponding argument set:

```
$ mpirun -n 3 ./myprog : -n 1 amplxe-cl -c advanced-hotspots -r ah -- ./myprog
```

The `-gtool` option allows you to specify a single analysis command for all argument sets (separated by colons ':') at once. Even though it is allowed to use `-gtool` within a single argument set, it is not recommended to use it in several sets at once and combine the two analysis methods (with `-gtool` and argument sets).

For example, to analyze processes 3, 5, 6, and 7 with the VTune Amplifier, you can use the following command line:

```
$ mpirun -n 8 -gtool "amplxe-cl -collect hotspots -r result:3,5-7" ./myprog
```

The `-gtool` option also provides several methods for finer process selection. For example, you can easily analyze only one process on each host, using the `exclusive` launch mode:

```
$ mpirun -n 8 -ppn 4 -hosts node1,node2 -gtool "amplxe-cl -collect hotspots -r result:all=exclusive" ./myprog
```

You can also use the `-gtoolfile` option to specify `-gtool` parameters in a configuration file. All the same rules apply. Additionally, you can separate different command lines with section breaks.

For example, if `gtool_config_file` contains the following settings:

```
env VARIABLE1=value1 VARIABLE2=value2:3,5,7-9; env VARIABLE3=value3:0,11
env VARIABLE4=value4:1,12
```

The following command sets `VARIABLE1` and `VARIABLE2` for processes 3, 5, 7, 8, and 9 and sets `VARIABLE3` for processes 0 and 11, while `VARIABLE4` is set for processes 1 and 12:

```
$ mpirun -n 16 -gtoolfile gtool_config_file a.out
```

Using `-gtool` for Debugging

The `-gtool` runtime option can help you with debugging, when attaching to several processes at once. Instead of attaching to each process individually, you can specify all the processes in a single command line. For example:

```
$ mpirun -n 16 -gtool "gdb:3,5,7-9=attach" ./myprog
```

The command line above attaches the GNU* Debugger (GDB*) to processes 3, 5, 7, 8 and 9.

NOTE

Do not use the `-gdb` and `-gtool` options together. Use one option at a time.

See Also

Intel® MPI Library Developer Reference, section *Command Reference > mpiexec.hydra > Global Options > gtool Options*

6.4. MPI Tuning

Intel® MPI Library provides a tuning utility `mpitune`, which allows you to automatically adjust Intel® MPI Library parameters, such as collective operation algorithms, to your cluster configuration or application. The tuner iteratively launches a benchmarking application with different configurations to measure performance and stores the results of each launch. Based on these results, the tuner generates optimal values for the parameters that are being tuned.

For example, you can use the tuner to identify the best `MPI_Bcast` algorithms (specified with `I_MPI_ADJUST_BCAST`) for specific message sizes. The result may look as follows:

```
I_MPI_ADJUST_BCAST=2:0-0;1:1-64;2:65-509;1:510-8832;3:8833-0
```

Use the resulting variable value with the application launch to achieve performance gain.

All tuner parameters should be specified in a configuration file, including the library parameter to be tuned, the application launch command and other parameters. You can modify the sample configuration file `etc/tune_cfg/example` according to your environment (for example, to specify other collective operations to be tuned) and use it for tuning as follows:

```
$ mpitune --config <installdir>/etc/tune_cfg/config
```

NOTE

The `mpitune` usage model has changed since the 2018 release. Tuning parameters should now be specified in a configuration file rather than as command-line options.

For details on the `mpitune` configuration options, refer to the Developer Reference, section *Command Reference > mpitune*.

7. Additional Supported Features

7.1. Asynchronous Progress Control

Intel® MPI Library supports asynchronous progress threads that allow you to manage communication in parallel with application computation and, as a result, achieve better communication/computation overlapping. This feature is supported for the `release_mt` and `debug_mt` versions only.

NOTE

Asynchronous progress has a full support for MPI point-to-point operations, blocking collectives, and a partial support for non-blocking collectives (`MPI_Ibcast`, `MPI_Ireduce`, and `MPI_Iallreduce`).

To enable asynchronous progress, pass 1 to the `I_MPI_ASYNC_PROGRESS` environment variable. You can define the number of asynchronous progress threads by setting the `I_MPI_ASYNC_PROGRESS_THREADS` environment variable. The `I_MPI_ASYNC_PROGRESS_ID_KEY` variable sets the MPI info object key that is used to define the progress `thread_id` for a communicator.

Setting the `I_MPI_ASYNC_PROGRESS_PIN` environment variable allows you to control the pinning of the asynchronous progress threads. In case of N progress threads per process, the first N logical processors from the list will be assigned to the threads of the first local process, while the next N logical processors - to the second local process and so on.

Example

For example, if the thread affinity is 0, 1, 2, 3 with 2 progress threads per process and 2 processes per node, then the progress threads of the first local process are pinned to logical processors 0 and 1, while the progress threads of the second local process are pinned to processors 2 and 3.

The code example is available below or in the `async_progress_sample.c` file in the `doc/examples` subdirectory of the package.

For more information on environment variables, refer to the Intel® MPI Library Developer Reference, section *Environment Variable Reference > Environment Variable Reference for Asynchronous Progress Control*.

7.2. Multiple Endpoints Support

The traditional MPI/OpenMP* threading model has certain performance issues. Thread safe access to some MPI objects, such as requests or communicators, requires an internal synchronization between threads; the performance of the typical hybrid application, which uses MPI calls from several threads per rank, is often lower than expected.

The PSM2 Multiple Endpoints (Multi-EP) support in the Intel® MPI Library makes it possible to eliminate most of the cross-thread synchronization points in the MPI workflow, at the cost of some limitations on what is allowed by the standard `MPI_THREAD_MULTIPLE` thread support level. The Multi-EP support, implemented with `MPI_THREAD_SPLIT` (thread-split) programming model, implies several requirements to the application program code to meet, and introduces a few runtime switches. These requirements, limitations, and usage rules are discussed in the sections below.

7.2.1. MPI_THREAD_SPLIT Programming Model

This feature is supported for the `release_mt` and `debug_mt` versions only.

The communication patterns that comply with the thread-split model must not allow cross-thread access to MPI objects to avoid thread synchronization and must disambiguate message matching, so that threads could be separately addressed and not more than one thread could match the message at the same time. Provided that, the user must notify the Intel MPI Library that the program complies with thread-split model, that is, it is safe to apply the high performance optimization.

Each `MPI_THREAD_SPLIT`-compliant program can be executed correctly with a thread-compliant MPI implementation under `MPI_THREAD_MULTIPLE`, but not every `MPI_THREAD_MULTIPLE`-compliant program follows the `MPI_THREAD_SPLIT` model.

This model allows MPI to apply optimizations that would not be possible otherwise, such as binding specific hardware resources to concurrently communicating threads and providing lockless access to MPI objects.

Since `MPI_THREAD_SPLIT` is a non-standard programming model, it is disabled by default and can be enabled by setting the environment variable `I_MPI_THREAD_SPLIT`. If enabled, the threading runtime control must also be enabled to enable the programming model optimizations (see [Threading Runtimes Support](#)).

Setting the `I_MPI_THREAD_SPLIT` variable does not affect behavior at other threading levels such as `SINGLE` and `FUNNELED`. To make this extension effective, request the `MPI_THREAD_MULTIPLE` level of support at `MPI_Init_thread()`.

NOTE

Thread-split model has support for MPI point-to-point operations and blocking collectives.

MPI_THREAD_SPLIT Model Description

As mentioned above, an `MPI_THREAD_SPLIT`-compliant program must be at least a thread-compliant MPI program (supporting the `MPI_THREAD_MULTIPLE` threading level). In addition to that, the following rules apply:

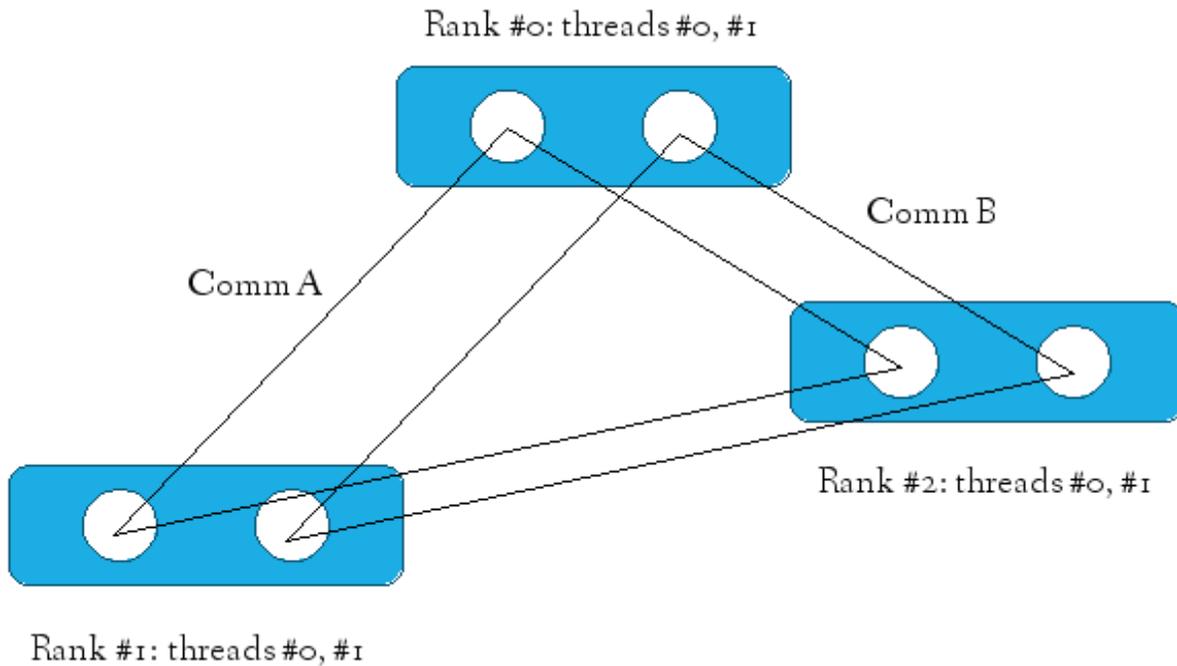
1. Different threads of a process must not use the same communicator concurrently.
2. Any request created in a thread must not be accessed by other threads, that is, any non-blocking operation must be completed, checked for completion, or probed in the same thread.
3. Communication completion calls that imply operation progress such as `MPI_Wait()`, `MPI_Test()` being called from a thread don't guarantee progress in other threads.

The model implies that each process thread has a distinct logical thread number `thread_id`. `thread_id` must be set to a number in the range 0 to NT-1, where NT is the number of threads that can be run concurrently. `thread_id` can be set implicitly, or your application can assign it to a thread. Depending on the assignment method, there are two usage submodels:

1. Implicit model: both you and the MPI implementation know the logical thread number in advance via a deterministic thread number query routine of the threading runtime. The implicit model is only supported for OpenMP* runtimes via `omp_get_thread_num()`.
2. Explicit model: you pass `thread_id` as an integer value converted to a string to MPI by setting an MPI Info object (referred to as info key in this document) to a communicator. The key `thread_id` must be used. This model fits task-based parallelism, where a task can be scheduled on any process thread.

The `I_MPI_THREAD_ID_KEY` variable sets the MPI info object key that is used to explicitly define the `thread_id` for a thread (`thread_id` by default).

Within the model, only threads with the same `thread_id` can communicate. To illustrate it, the following communication pattern complies to the `MPI_THREAD_SPLIT` model: Suppose Comm A and Comm B are two distinct communicators, aggregating the same ranks. The system of these two communicators will fit the `MPI_THREAD_SPLIT` model only if all threads with `thread_id #0` use Comm A, while all threads with `thread_id #1` use Comm B.



7.2.2. Threading Runtimes Support

The MPI thread-split programming model has special support of the OpenMP* runtime, but you may use any other threading runtime. The support differs in the way you communicate with the MPI runtime to set up `thread_id` for a thread and the way you set up the number of threads to be run concurrently. If you choose the OpenMP runtime support, make sure you have the OpenMP runtime library, which comes with the recent Intel® compilers and GNU* gcc compilers, and link the application against it.

The support is controlled with the `I_MPI_THREAD_RUNTIME` environment variable. Since the threading runtime support is a non-standard functionality, you must enable it explicitly using the `generic` or `openmp` argument for the OpenMP* runtime support.

You can set the maximum number of threads to be used in each process concurrently with the `I_MPI_THREAD_MAX` environment variable. This helps the MPI implementation allocate the hardware resources efficiently. By default, the maximum number of threads per rank is 1.

OpenMP* Threading Runtime

The OpenMP runtime supports the both implicit and explicit submodels. By default, the Intel MPI Library assumes that `thread_id` is set with the `omp_get_thread_num()` function call defined in the OpenMP standard. This scenario corresponds to the implicit submodel. You can use the explicit submodel by setting the `thread_id` info key for a communicator, which is particularly useful for OpenMP tasks.

By default, the maximum number of threads is set with the `omp_get_max_threads()` function. To override this function, set either `I_MPI_THREAD_MAX` or `OMP_NUM_THREADS` environment variable.

7.2.3. Code Change Guide

The example in this section shows you one of the ways to change a legacy program to effectively use the advantages of the `MPI_THREAD_SPLIT` threading model. The original code is available in the `thread_split.cpp` file in the `doc/examples` subdirectory of the package.

In the original code, the functions `work_portion_1()`, `work_portion_2()`, and `work_portion_3()` represent a CPU load that modifies the content of the memory pointed to by the `in` and `out` pointers. In this particular example, these functions perform correctness checking of the `MPI_Allreduce()` function.

Changes Required to Use the OpenMP* Threading Model

1. To run MPI functions in a multithreaded environment, `MPI_Init_thread()` with the argument equal to `MPI_THREAD_MULTIPLE` must be called instead of `MPI_Init()`.
2. According to the `MPI_THREAD_SPLIT` model, in each thread you must execute MPI operations over the communicator specific to this thread only. So, in this example, the `MPI_COMM_WORLD` communicator must be duplicated several times so that each thread has its own copy of `MPI_COMM_WORLD`.

NOTE

The limitation is that communicators must be used in such a way that the thread with `thread_id n` on one node communicates only with the thread with `thread_id m` on the other. Communications between different threads (`thread_id n` on one node, `thread_id m` on the other) are not supported.

3. The data to transfer must be split so that each thread handles its own portion of the input and output data.
4. The barrier becomes a two-stage one: the barriers on the MPI level and the OpenMP level must be combined.
5. Check that the runtime sets up a reasonable affinity for OpenMP threads. Typically, the OpenMP runtime does this out of the box, but sometimes, setting up the `OMP_PLACES=cores` environment variable might be necessary for optimal multi-threaded MPI performance.

Changes Required to Use the POSIX Threading Model

1. To run MPI functions in a multithreaded environment, `MPI_Init_thread()` with the argument equal to `MPI_THREAD_MULTIPLE` must be called instead of `MPI_Init()`.
2. You must execute MPI collective operation over a specific communicator in each thread. So the duplication of `MPI_COMM_WORLD` should be made, creating a specific communicator for each thread.
3. The info key `thread_id` must be properly set for each of the duplicated communicators.

NOTE

The limitation is that communicators must be used in such a way that the thread with `thread_id n` on one node communicates only with the thread with `thread_id m` on the other. Communications between different threads (`thread_id n` on one node, `thread_id m` on the other) are not supported.

4. The data to transfer must be split so that each thread handles its own portion of the input and output data.
5. The barrier becomes a two-stage one: the barriers on the MPI level and the POSIX level must be combined.
6. The affinity of POSIX threads can be set up explicitly to reach optimal multithreaded MPI performance.