

# Intel® MPI Library for Linux\*

## Getting Started Guide

---

The Intel® MPI Library is a multi-fabric message passing library that implements the Message Passing Interface, v2 (MPI-2) specification. Use it to switch between interconnection fabrics without re-linking.

This *Getting Started Guide* explains how to use the Intel® MPI Library to compile and run a simple MPI program. This guide also includes basic usage examples and troubleshooting tips.

To quickly start using Intel® MPI Library, print this short guide and walk through the example provided.

## Contents

1	About this Document .....	4
1.1	Intended Audience .....	4
1.2	Using Doc Type Field.....	4
1.3	Conventions and Symbols.....	4
1.4	Related Information.....	5
2	Using the Intel® MPI Library.....	6
2.1	Usage Model.....	6
2.2	Before You Begin.....	6
2.3	Quick Start.....	6
2.4	Compiling and Linking.....	7
2.5	Setting up MPD Daemons .....	7
2.6	Selecting a Network Fabric .....	8
2.7	Running an MPI Program .....	9
3	Troubleshooting .....	10
3.1	Testing Installation.....	10
3.2	Troubleshooting MPD Setup .....	10
3.3	Compiling and Running a Test Program .....	11
4	Next Steps .....	13

## Disclaimer and Legal Notices

THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

The software described in this document may contain software defects which may cause the product to deviate from published specifications. Current characterized software defects are available on request.

This document as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Developers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Improper use of reserved or undefined features or instructions may cause unpredictable behavior or failure in developer's software code when running on an Intel processor. Intel reserves these features or instructions for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from their unauthorized use.

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Chips, Core Inside, Dialogic, EtherExpress, ETOX, FlashFile, i386, i486, i960, iCOMP, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, MMX logo, Optimizer logo, OverDrive, Paragon, PDCharm, Pentium, Pentium II Xeon, Pentium III Xeon, Performance at Your Command, Pentium Inside, skool, Sound Mark, The Computer Inside., The Journey Inside, VTune, Xeon, Xeon Inside and Xircom are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2003-2007, Intel Corporation. All rights reserved.

# 1 About this Document

---

The goal of this *Getting Started Guide* is to provide you with a complete command and tuning reference for the Intel® MPI Library.

## 1.1 Intended Audience

This *Getting Started Guide* tells a first time user how to install and use the Intel MPI Library.

## 1.2 Using Doc Type Field

This *Reference Manual* contains the following sections:

### Document Organization

Section	Description
Section 1 About this Document	Section 1 introduces this document
Section 2 Using the Intel® MPI Library	Section 2 describes how to use the Intel® MPI Library
Section 3 Troubleshooting	Section 3 outlines first-aid troubleshooting actions

## 1.3 Conventions and Symbols

The following conventions are used in this document.

### Conventions and Symbols used in this Document

<i>This type style</i>	Document or product names
<a href="#">This type style</a>	Hyperlinks
<code>This type style</code>	Commands, arguments, options, file names
<code>THIS_TYPE_STYLE</code>	Environment variables
<code>&lt;this type style&gt;</code>	Placeholders for actual values
<code>[ items ]</code>	Optional items
<code>{ item   item }</code>	Selectable items separated by vertical bar(s)
<b>(SDK only)</b>	Only for the Software Development Kit (SDK) users

## 1.4 Related Information

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

[Product Web Site](#)

[Intel® MPI Library support](#)

[Intel® Cluster Tools Products](#)

[Intel® Software Development Products](#)

## 2 Using the Intel® MPI Library

---

### 2.1 Usage Model

Using the Intel® MPI Library involves the following steps:

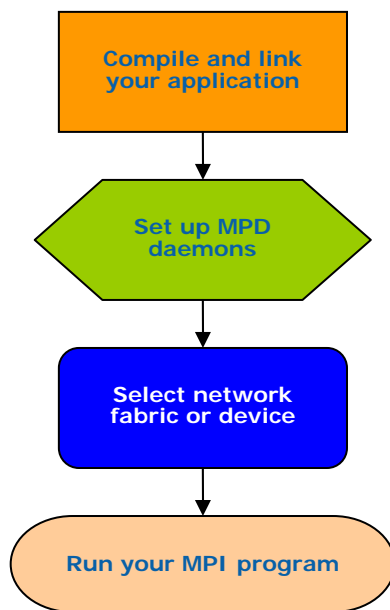


Figure 1: Flowchart representing the usage model for working with Intel MPI Library.

### 2.2 Before You Begin

Before using the Intel® MPI Library, ensure that the library, scripts, and utility applications are installed. See the product *Intel® MPI Library for Linux\* Installation Guide* for installation instructions.

### 2.3 Quick Start

1. Use the `source` shell command for getting proper environment settings from the `mpivars.[c]sh` scripts included with the Intel MPI Library. It is located in the `<installdir>/bin` directory, or, for Intel® 64 in 64-bit mode, the `<installdir>/bin64` directory.
2. Create an `mpd.hosts` text file that lists the nodes in the cluster using one host name per line.
3. **(SDK only)** Make sure that you have compiler in your `PATH`.
4. **(SDK only)** Compile test program using the appropriate compiler driver. For example:  

```
$ mpicc -o test <installdir>/test/test.c
```
5. Execute the test using the `mpirun` command.  

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

6. See the rest of this document and *Intel® MPI Library Reference Manual* for more details.

## 2.4 Compiling and Linking

### (SDK only)

To compile and link an MPI program with the Intel MPI Library:

1. Ensure that the underlying compiler and related software appear in your `PATH`.  
If you use Intel® compilers, ensure that the compiler library directories appear in `LD_LIBRARY_PATH` environment variable.  
For example, for Intel® C++ Compiler version 10.1 and Intel® Fortran Compiler version 10.1, execute the appropriate setup scripts:  

```
/opt/intel_cc_10.1/bin/iccvars.[c]sh, and  
/opt/intel_fc_10.1/bin/ifortvars.[c]sh
```
2. Compile your MPI program via the appropriate `mpi` command.  
For example, use the `mpicc` command to compile C code using the GNU\* C compiler as follows:  

```
$ mpicc <installdir>/test/test.c
```

where `<installdir>` is a full path to installed package.

All supported compilers have equivalent commands that use the prefix `mpi` on the standard compiler command. For example, the Intel MPI Library command for Intel Fortran® Compiler from version 8.0 up (`ifort`) is `mpiifort`.

## 2.5 Setting up MPD Daemons

The Intel MPI Library uses a Multi-Purpose Daemon (MPD) job startup mechanism. In order to run programs compiled with the `mpicc` (or related) commands, set up MPD daemons.

Always start and maintain your own set of MPD daemons, as opposed to having the system administrator start up the MPD daemons once for use by all users on the system. This setup enhances system security and gives you flexibility in controlling your execution environment.

To set up MPD daemons:

1. Set up environment variables with appropriate values and directories, for example, in the `.cshrc` or `.bashrc` files.
  - Ensure that the `PATH` variable includes the `<installdir>/bin` directory, or, for Intel® 64 in 64-bit mode, the `<installdir>/bin64` directory. Use the `mpivars.[c]sh` scripts included with the Intel MPI Library to set up this variable.
  - Ensure that the `PATH` variable includes the directory for Python\* version 2.2 or higher.
  - **(SDK only)** If you are using Intel® compilers, ensure that the `LD_LIBRARY_PATH` variable contains the directories for the compiler library. Set this variable by using the `*vars.[c]sh` scripts included with the compiler.
  - Set any additional environment variables your application uses.

**NOTE:** Steps two through four are optional.

2. Create a `$HOME/.mpd.conf` file. To set up your MPD password, enter the following into this file:  

```
secretword=<mpd secret word>
```

Do not use any Linux\* login password. An arbitrary `<mpd secret word>` string only controls access to the MPD daemons by various cluster users.

3. Set protection on the `$HOME/.mpd.conf` file using the `chmod` command so that only you have read and write privileges:

```
$ chmod 600 $HOME/.mpd.conf
```

- Verify that you can observe the `PATH` settings and `.mpd.conf` contents through `rsh` on all nodes of the cluster. For example, use the following commands with each `<node>` in the cluster:

```
$ rsh <node> env
```

```
$ rsh <node> cat $HOME/.mpd.conf
```

Make sure that every node, rather than only one of them, can connect to any other node.

If your cluster uses `ssh` instead of `rsh`, look into the Notes section below.

- Create an `mpd.hosts` text file that lists the nodes in the cluster using one host name per line.
- Shut down the eventual MPD daemons using the `mpdallexit` command:

```
$ mpdallexit
```

- Use the `mpdboot` command to start up the MPD daemons:

```
$ mpdboot -n <#nodes>
```

The file `$PWD/mpd.hosts` will be used by default if it is present. If there is no host file, the `mpdboot` command will start one MPD daemon on the local machine.

- Use the `mpdtrace` command to determine the status of the MPD daemons:

```
$ mpdtrace
```

The output should be a list of nodes that are currently running MPD daemons. This list should match the contents of the `mpd.hosts` file.

**NOTE:** If your cluster uses `ssh` instead of `rsh`, make sure that every node can connect to any other node via `ssh` without a password. Look into your system manuals for details.

**NOTE:** If your cluster uses `ssh` instead of `rsh`, add the `-r ssh` or `-- rsh=ssh` option to the `mpdboot` invocation string.

## 2.6 Selecting a Network Fabric

The Intel® MPI Library dynamically selects different fabrics for communication between MPI processes. To select specific fabric combination, set the `I_MPI_DEVICE` environment variable to one of the following values:

<code>I_MPI_DEVICE</code> values	Supported fabric
<code>sock</code>	TCP/Ethernet*/sockets
<code>shm</code>	Shared memory only (no sockets)
<code>ssm</code>	TCP + shared memory (for SMP clusters connected via Ethernet*)
<code>rdma[:&lt;provider&gt;]</code>	InfiniBand*, Myrinet* (via specified DAPL* provider)
<code>rdssm[:&lt;provider&gt;]</code>	TCP + shared memory + DAPL* (for SMP clusters connected via RDMA-capable fabrics)

Ensure that the selected fabric is available. For example, use `shm` only when all the processes can communicate with each other via shared memory. Use `rdma` only when all processes can communicate with each other via a single DAPL provider.



## 2.7 Running an MPI Program

To launch programs linked with the Intel MPI Library, use the `mpiexec` command:

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

Use the `-n` option to set the number of processes. This is the only obligatory option for the `mpiexec` command.

If you are using a network fabric as opposed to the default fabric, use the `-genv` option to specify a value to be assigned to the `I_MPI_DEVICE` variable.

For example, to run an MPI program using the `shm` fabric, type in the following command:

```
$ mpiexec -genv I_MPI_DEVICE shm -n <# of processes> ./a.out
```

For the `rdma` capable fabric, use the following command:

```
$ mpiexec -genv I_MPI_DEVICE rdma -n <# of processes> ./a.out
```

You can select any supported device. For more information, see [Selecting a Network Fabric](#) above.

If you successfully ran your application using the Intel MPI Library, it should run as with other MPI libraries. You can now move your application from one cluster to another and use different fabrics between the nodes without re-linking. If you encounter problems, see [Troubleshooting](#) for possible solutions.

## 3 Troubleshooting

Use the following sections to troubleshoot problems with installation, setup, and running applications using the Intel® MPI Library.

### 3.1 Testing Installation

To ensure that the Intel® MPI Library is installed and functioning, complete the general testing, compile and run a test program.

To test the installation:

1. Verify that you have Python\* version 2.2 or higher in your `PATH`:

```
$ rsh <nodename> python -V
```

If this command returns an error message or a value lower than 2.2, install Python\* version 2.2 or higher, and make sure that you have it in your `PATH`.

2. Check for the presence of a Python\* XML module such as `python-xml*` or `libxml2-python*`:

```
$ rpm -qa | grep python-xml
$ rpm -qa | grep libxml2-python
```

Install the missing module if the output does not include the name “python-xml” or “libxml2-python” and a version number.

3. Check for the presence of an XML parser such as `expat*` or `pyxml*`:

```
$ rpm -qa | grep expat
$ rpm -qa | grep pyxml
```

Install the missing module if the output does not include the name “expat” or “pyxml” and a version number.

4. Verify that `<installdir>/bin` (`<installdir>/bin64` for Intel® 64 in 64-bit mode) is in your `PATH`:

```
$ rsh <nodename> which mpiexec
```

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

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

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

You should start a `mpd` ring before executing the `mpiexec` command. You should see the correct directories for these path variables for each node you test. If not, call the appropriate `*vars.[c]sh` scripts. For example, for Intel® C++ Compiler version 10.1 use the following source command:

```
$ . /opt/intel_cc_10.1/bin/iccvars.sh
```

5. In some unusual circumstances, you may need to include the `<installdir>/lib` directory (`<installdir>/lib64` for Intel® 64 in 64-bit mode) in your `LD_LIBRARY_PATH`. To verify your `LD_LIBRARY_PATH` settings, use the command:

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

### 3.2 Troubleshooting MPD Setup

Check if it is possible to run the `mpd` command on the local machine. Do the following:

```
# mpd &
# mpdtrace
```

```
# mpdallexit
```

The output of `mpdtrace` should show the hostname of the machine you are running on. If this is not the case, or if you cannot start up the MPD, check that the installation was correct and the environment was set up properly.

The next troubleshooting steps assume that MPD daemons are set up and running. To start your diagnosis, verify that MPD daemons are running on all expected nodes using:

```
mpdtrace
```

The output lists all MPD daemons running or indicates an error. If some desired nodes are missing from the output list of `mpdtrace`, do the following:

1. Try to restart the MPD daemons using the following commands:
  - Kill all running MPD daemons:
 

```
$ mpdallexit
```
  - For each node, ensure all daemons were killed:
 

```
$ rsh <nodename> ps -aef | grep python
$ rsh <nodename> kill -9 <remaining python processes>
```
  - Reboot the MPD daemons. Be sure to use the appropriate configuration options and host file:
 

```
$ mpdboot [<options>]
```
  - Confirm that all expected MPD daemons are now running:
 

```
$ mpdtrace
```
2. If the output of the `mpdtrace` command is not still indicating that all expected MPD daemons are running, follow the next steps:
  - Kill and restart the MPD daemons as described in step 1, adding the debug and verbose options to the `mpdboot` command:
 

```
$ mpdboot -d -v [<options>]
```

 Note the `rsh` commands in the output from step a. For example:
 

```
launch cmd= rsh -n <nodename> '<installdir>/bin/mpd \
-h <nodename> -p <port-number> --ncpus=<ncpus> -e -d'
```
  - Copy and paste the line of the output from the `rsh` command up to the end of line. For example:
 

```
$ rsh -n <nodename> '<installdir>/bin/mpd \
-h <nodename> -p <port-number> --ncpus=<ncpus> -e -d'
```
  - Execute the edited `rsh` command. Use the resulting output to diagnose and correct the underlying problem. For example, the most common problems include:
    - Failure of the `rsh` command to contact `<nodename>`.
    - Other failure of the `rsh` command, for example, a system setup problem.
    - The `<installdir>/bin/mpd` command could not be found or could not be executed.

## 3.3 Compiling and Running a Test Program

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

1. **(SDK only)** Compile a test program included with the product release as follows:
 

```
$ cd <installdir>/test
$ mpicc test.c
```
2. If you are using InfiniBand\*, Myrinet\*, or other RDMA-capable network hardware and software, verify that everything is functioning.
3. Run the test program with all available configurations on your cluster.
  - Test the `sock` device using:
 

```
$ mpiexec -n 2 -env I_MPI_DEBUG 2 -env I_MPI_DEVICE sock ./a.out
```

You should see one line of output for each rank, as well as debug output indicating the `sock` device used.

- Test the `ssm` devices using:

```
$ mpiexec -n 2 -env I_MPI_DEBUG 2 -env I_MPI_DEVICE ssm ./a.out
```

You should see one line of output for each rank, as well as debug output indicating the `ssm` device used.

- Test any other fabric devices using:

```
$ mpiexec -n 2 -env I_MPI_DEBUG 2 -env I_MPI_DEVICE <device> ./a.out
```

where `<device>` can be `shm`, `rdma`, or `rdssm`

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

**NOTE:** The `<installdir>/test` directory in the Intel MPI Library Development Kit contains other test programs in addition to `test.c`

## 4 Next Steps

---

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

**Release Notes** include key product details. See the Release Notes for updated information on requirements, technical support, and known limitations.

For more information see Websites:

[Product Web Site](#)

[Intel® MPI Library support](#)

[Intel® Cluster Tools Products](#)

[Intel® Software Development Products](#)