

IBM Spectrum MPI  
Version 10 Release 1

## *Installation*





IBM Spectrum MPI  
Version 10 Release 1

## *Installation*



**Note**

Before using this information and the product it supports, read the information in "Notices" on page 25.

This edition applies to version 10, release 1, modification 0 of IBM Spectrum MPI (product number 5725-G83) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2016.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Tables</b> . . . . .	<b>v</b>
<b>About this information</b> . . . . .	<b>vii</b>
Who should use this information . . . . .	vii
Conventions and terminology used in this information . . . . .	vii
Prerequisite and related information . . . . .	viii
How to send your comments . . . . .	ix
<b>Chapter 1. Introducing IBM Spectrum MPI 10.1.</b> . . . . .	<b>1</b>
IBM Spectrum MPI components . . . . .	1
<b>Chapter 2. Planning to install the IBM Spectrum MPI software</b> . . . . .	<b>3</b>
IBM Spectrum MPI installation requirements . . . . .	3
IBM Spectrum MPI hardware requirements . . . . .	3
IBM Spectrum MPI software requirements . . . . .	3
Debug programs for IBM Spectrum MPI installation . . . . .	5
IBM Spectrum MPI Limitations . . . . .	5
Information for the system administrator . . . . .	5
Software compatibility within workstation clusters . . . . .	5
Node resources . . . . .	5
Setting the maximum number of processes per user . . . . .	6
Setting the memory limit per user . . . . .	6
File systems . . . . .	6
User IDs on remote nodes . . . . .	6
IBM Spectrum MPI user authorization . . . . .	6
Tuning your Linux system for more efficient parallel job performance . . . . .	6
Understanding the effects of changing the SMT or HT mode . . . . .	8
<b>Chapter 3. Installing IBM Spectrum MPI</b> . . . . .	<b>11</b>
Install the IBM Spectrum MPI software step-by-step . . . . .	11
View the release notes before installation . . . . .	11
Step 1: Obtain software packages . . . . .	11
Step 2: Set root authority . . . . .	11
Step 3: Install RPMs and accept the license packages . . . . .	11
Viewing the release notes after installation . . . . .	12
Migration installation . . . . .	13
Resolving installation errors . . . . .	13
<b>Chapter 4. Performing installation-related tasks</b> . . . . .	<b>15</b>
Removing a software component . . . . .	15
Finding installed components . . . . .	15
<b>Chapter 5. Understanding how installing IBM Spectrum MPI alters your system</b> . . . . .	<b>17</b>
Understanding the IBM Spectrum MPI directory structure . . . . .	17
<b>Chapter 6. Additional information for the system administrator</b> . . . . .	<b>19</b>
Configuring key-based login with OpenSSH . . . . .	19
<b>Chapter 7. Using example applications</b> . . . . .	<b>21</b>
<b>Accessibility features for IBM Spectrum MPI.</b> . . . . .	<b>23</b>
<b>Notices</b> . . . . .	<b>25</b>
Programming interface information . . . . .	27
MPI support statement . . . . .	27
Open MPI license . . . . .	27
Trademarks . . . . .	28
Terms and conditions for product documentation. . . . .	29
IBM Online Privacy Statement . . . . .	30
<b>Index</b> . . . . .	<b>31</b>



---

## Tables

1. Conventions . . . . .	vii	3. Network tuning . . . . .	7
2. Additional software requirements . . . . .	4	4. Network tuning: ARP entries . . . . .	8





---

## About this information

<b>Disclaimer:</b>
--------------------

The functions or features found herein may not be available on all operating systems or platforms and do not indicate the availability of these functions or features within the IBM® product or future versions of the IBM product. The development, release, and timing of any future features or functionality is at IBM's sole discretion. IBM's plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. The information mentioned is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The information may not be incorporated into any contract and it should not be relied on in making a purchasing decision.
---

This information describes how to install IBM Spectrum™ MPI, IBM's implementation of Open MPI 2.0.

IBM Spectrum MPI is a complete MPI implementation, based on the Open MPI open source project, and is designed to comply with all the requirements of the Message Passing Interface standard, MPI: A Message-Passing Interface Standard, Version 3.1, University of Tennessee, Knoxville, Tennessee, June 4, 2015.

For information about Open MPI, and to obtain official Open MPI documentation, refer to the Open MPI web site ([www.open-mpi.org](http://www.open-mpi.org)).

For more information about using IBM Spectrum MPI, refer to the *IBM Spectrum MPI: User's Guide*.

---

## Who should use this information

This information is intended for system programmers and administrators, who plan and install IBM Spectrum MPI.

---

## Conventions and terminology used in this information

Table 1 shows the conventions used in this information:

*Table 1. Conventions*

Convention	Usage
<b>bold</b>	Environment variables.
monospace	Examples and information that the system displays, command line options, file names, pathnames.
<b>bold monospace</b>	Command names and parameter names.
<i>italic</i>	<i>Italic</i> words or characters represent variable values that you must supply.  Italics are also used for unit titles, the first use of a term, and general emphasis in text.
<key>	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, <Enter> refers to the key on your terminal or workstation that is labeled with the word <i>Enter</i> .

Table 1. Conventions (continued)

Convention	Usage
\	In command examples, a backslash indicates that the command or coding example continues on the next line. For example: <pre>mkcondition -r IBM.FileSystem -e "PercentTotUsed &gt; 90" \ -E "PercentTotUsed &lt; 85" -m d "FileSystem space used"</pre>
{item}	Braces enclose a list from which you must choose an item in format and syntax descriptions.
[item]	Brackets enclose optional items in format and syntax descriptions.
<Ctrl-x>	The notation <Ctrl-x> indicates a control character sequence. For example, <Ctrl-c> means that you hold down the control key while pressing <c>.
item...	Ellipses indicate that you can repeat the preceding item one or more times.
	<ul style="list-style-type: none"> <li>• In <i>synopsis</i> statements, vertical lines separate a list of choices. In other words, a vertical line means <i>Or</i>.</li> <li>• In the margin of the document, vertical lines indicate technical changes to the information.</li> </ul>

## Prerequisite and related information

IBM Spectrum MPI is a member of the IBM Spectrum Computing family ([www.ibm.com/systems/spectrum-computing/](http://www.ibm.com/systems/spectrum-computing/)).

The IBM Spectrum MPI library consists of:

- *IBM Spectrum MPI: Installation*, GC27-8264-00
- *IBM Spectrum MPI: User's Guide*, GC27-8265-00

To access the most recent IBM Spectrum MPI documentation in PDF and HTML format, refer to IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)), on the web.

The IBM Spectrum MPI books are also available in PDF format from the IBM Publication Center ([www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss](http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss)), on the web.

It is easiest to locate a book in the IBM Publications Center by supplying the book's publication number. The publication number for each of the IBM Spectrum MPI books is listed after the book title in the preceding list.

IBM Platform LSF<sup>®</sup> (Load Sharing Facility) also works in conjunction with IBM Spectrum MPI. The LSF publications can be found in IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)) and the IBM Publication Center ([www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss](http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss)).

## Terminology

For terms and definitions related to IBM Spectrum MPI, see IBM Terminology ([www.ibm.com/software/globalization/terminology/](http://www.ibm.com/software/globalization/terminology/)).

---

## How to send your comments

Your feedback is important in helping to provide the most accurate, high-quality information. If you have comments about this information or other IBM Spectrum MPI documentation, go to the IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)) and use the **Provide feedback** links.



---

## Chapter 1. Introducing IBM Spectrum MPI 10.1

The IBM Spectrum MPI licensed program product is a set of software components that help you develop, debug, analyze, and run C, C++, or FORTRAN programs on a cluster of networked servers.

Before you install IBM Spectrum MPI, be familiar with its software components. For a brief overview, see *IBM Spectrum MPI components* following.

For the latest information, always review the IBM Spectrum MPI product release notes, available from the following websites:

Passport Advantage® website at <http://www.ibm.com/software/howtobuy/passportadvantage/>

IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter))

---

### IBM Spectrum MPI components

IBM Spectrum MPI is made up of various components for submitting and running jobs in a high-performance cluster environment.

The IBM Spectrum MPI components are:

#### **Message passing and collective communication API subroutine libraries**

Libraries that contain subroutines that help application developers parallelize their code. These libraries are described in *IBM Parallel Environment Runtime Edition: MPI Programming Guide*. For additional information about MPI, see <http://www.open-mpi.org>.

#### **Optimized Collectives Library (libcollectives)**

An optimized collectives library with both one-sided and active messages collectives.

#### **Parallel Active Messaging Interface (PAMI)**

A messaging API that supports point-to-point communications.

#### **IBM Spectrum MPI examples**

A set of example programs, included in the product package.

#### **IBM Spectrum MPI documentation**

The IBM Spectrum MPI library consists of:

- *IBM Spectrum MPI: Installation*, GC27-8264-00
- *IBM Spectrum MPI: User's Guide*, GC27-8265-00

You can view, search, and print the most recent IBM Spectrum MPI documentation in PDF and HTML format. See IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)) for more information.

The IBM Spectrum MPI books are also available in PDF format from the IBM Publication Center ([www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss](http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss)).



---

## Chapter 2. Planning to install the IBM Spectrum MPI software

When you are planning to install the IBM Spectrum MPI software, you need to ensure that you meet all of the necessary system requirements. You also need to think about what your programming environment is and the strategy for using that environment.

---

### IBM Spectrum MPI installation requirements

Various system requirements for installing and running the IBM Spectrum MPI software, include requirements for hardware, software, disk space, and file sets.

#### IBM Spectrum MPI hardware requirements

For information on the hardware that is supported by IBM Spectrum MPI, Version 10.1, refer to the current *announcement letter* at IBM Offering Information ([http://www.ibm.com/common/ssi/index.wss?request\\_locale=en](http://www.ibm.com/common/ssi/index.wss?request_locale=en)).

#### IBM Spectrum MPI software requirements

The software that is required for IBM Spectrum MPI includes various components plus in some cases, extra software. You need to decide which components to install on your system based on the features you plan to use. You might also need to install some additional products or components, based on how you plan to use IBM Spectrum MPI.

##### Operating system requirements

For information on the operating system requirements that IBM Spectrum MPI Version 10.1 supports, refer to the current *announcement letter* at IBM Offering Information ([http://www.ibm.com/common/ssi/index.wss?request\\_locale=en](http://www.ibm.com/common/ssi/index.wss?request_locale=en)).

The following file sets are required:

- `ibm_smpi-10.1.0.0-rh6.x86_64.rpm`
- `ibm_smpi-10.1.0.0-rh7.ppc64le.rpm`
- `ibm_smpi-10.1.0.0-rh7.x86_64.rpm`
- `ibm_smpi_kt-10.1.0.0-rh7-ppc64le.tbz2`
- `ibm_smpi_lic_s-10.1-rh6.x86_64.rpm`
- `ibm_smpi_lic_s-10.1-rh7.ppc64le.rpm`
- `ibm_smpi_lic_s-10.1-rh7.x86_64.rpm`

## Additional software requirements for IBM Spectrum MPI installation

Table 2. Additional software requirements

If you plan to:	This software is required:	Things to consider:
Compile parallel executables	<p>A working C or FORTRAN compiler. IBM Parallel Environment Runtime Edition supports parallel program development that uses the following compilers.</p> <p>The IBM compilers for Linux on Power Systems™ servers include:</p> <ul style="list-style-type: none"> <li>• IBM C/C++ compiler, V13.1.4, or later</li> <li>• IBM FORTRAN compiler, V15.1.4, or later</li> </ul> <p>The GNU compilers for Linux include:</p> <ul style="list-style-type: none"> <li>• C compiler</li> <li>• C++ compiler, V4.4.7</li> <li>• FORTRAN compiler</li> </ul> <p>The Intel compilers for Linux include:</p> <ul style="list-style-type: none"> <li>• C compiler, V12.1.5</li> <li>• C++ compiler</li> <li>• FORTRAN compiler</li> </ul>	A working C or FORTRAN compiler is needed to build parallel applications.
Provide resource management and scheduling functions by way of IBM Platform LSF to submit and run batch applications and manage network resources.	<p>LSF Version 9.1.1 (or later). If you want the latest level of task affinity support, you must use LSF 9.1.1.1 (or later).</p> <p>On x86-based servers, the LSF RPMs are found in these tar files:</p> <ul style="list-style-type: none"> <li>• lsf9.1.1_linux2.6-glibc2.3-x86_64.tar.Z</li> <li>• lsf9.1.1_lsfinstall_linux_x86_64.tar.Z</li> </ul> <p>Install the following On Power Systems servers:</p> <ul style="list-style-type: none"> <li>• lsf9.1.1_linux2.6-glibc2.3-ppc64.tar.Z</li> <li>• lsf9.1.1_lsfinstall_linux_ppc64.tar.Z</li> </ul>	See <i>Installing IBM Platform LSF on UNIX and Linux</i> for more information about installing LSF.
Analyze or tune program performance.	The IBM Parallel Performance Toolkit for POWER® (5765-PD2). The required file sets are ppedev.hpct and ppedev.runtime.	The IBM Parallel Performance Toolkit for POWER provides hpctView and the HPCT plug-in for Eclipse PTP for performance and trace data visualization. The required file sets are hpctView and ppedev-update.



Table 2. Additional software requirements (continued)

If you plan to:	This software is required:	Things to consider:
Choose to have IBM Spectrum MPI generate lightweight core files instead of standard core files.	The <b>GDB</b> command, which is included with the GDB RPM.	GDB is the GNU Project Debugger. For more information, see the GNU Project Debugger Web site ( <a href="http://www.gnu.org/software/gdb/">http://www.gnu.org/software/gdb/</a> ).

## Debug programs for IBM Spectrum MPI installation

IBM Spectrum MPI supports the Allinea DDT and TotalView debuggers. For more information, see Chapter 8, *Debugging and profiling applications*, in *IBM Spectrum MPI: User's Guide*.

## IBM Spectrum MPI Limitations

Some IBM Spectrum MPI product options and related software are subject to certain limitations, as explained in this section.

**Note:** Always check the release notes for current limitations from the following website: IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter)).

## Information for the system administrator

For system administrators, it is important to understand software compatibility for IBM Spectrum MPI and how to plan out your node resources. You also need to determine which nodes in your cluster require which file sets.

### Software compatibility within workstation clusters

For all workstations *within a workstation cluster*, the same release level (including maintenance levels) of IBM Spectrum MPI software is required. (Running the same release level ensures that an individual application can run on any workstation in the cluster.)

### Node resources

How you plan your node resources vary according to whether you are installing IBM Spectrum MPI with or without a resource manager.

#### On a cluster that uses a resource manager

The system administrator uses a resource manager to partition nodes into *pools* or *features* or both, to which they assign names or numbers and other information. The workstation from which parallel jobs are started is called the *home node* and it can be any workstation on the LAN.

#### On a cluster without a resource manager

On an IBM Power Systems cluster without a resource manager, you assign nodes or servers to the following categories:

- *Home node* (workstation from which parallel jobs are started) for running **mpirun**
- Nodes or servers for developing and compiling applications
- Nodes or servers that run applications in the parallel environment

You must identify the nodes or servers that run as execution nodes by name in a host list file.

## Setting the maximum number of processes per user

An operating system limits the number of processes that can be created by a single user. IBM Spectrum MPI requires this limit to be at least four times the maximum number of tasks on a node to support both running and debugging parallel programs.

On Linux, the limit is controlled by the `nproc` attribute in `/etc/security/limits.conf`, and can be changed only by root.

## Setting the memory limit per user

An operating system limits the amount of memory available per user. IBM Spectrum MPI requires this limit be increased. On Linux, the limit is controlled only by the system administrator (root) and by the `memlock` attribute in `/etc/security/limits.conf`. The value should be changed to at least 2 GB for all users, for both the soft and hard limits, as follows:

```
* soft memlock 2097152
* hard memlock 2097152
```

**Note:** Users cannot use the Linux `ulimit` command to set or change limits.

---

## File systems

IBM Spectrum MPI file sets are installed in the `/opt_ibm_spectrum_mpi` directory by default. The installation directory is relocatable.

If you do not use a shared file system, you need to copy the user's executable files to the other nodes. If you are managing your cluster with xCAT, you can use the xCAT commands `xdsh` and `xdcp`. For more information about `xdsh` and `xdcp` refer to the xCAT documentation repository (<http://xcat-docs.readthedocs.io>).

---

## User IDs on remote nodes

On each remote node, the system administrator must set up a user ID, other than a root ID, for each user on each remote node who is executing serial or parallel applications.

Each user must have an account on all nodes where a job runs. Both the user name and user ID must be the same on all nodes. Also, the user must be a member of the same named group on the home node and the remote nodes.

---

## IBM Spectrum MPI user authorization

IBM Spectrum MPI supports running jobs under the secure shell (ssh) or the remote shell (rsh). If you are using ssh to connect to a remote host, in order for `mpirun` to operate properly, it is recommended that you set up a passphrase for passwordless login. For more information, see the Open MPI web site ([www.open-mpi.org](http://www.open-mpi.org)).

---

## Tuning your Linux system for more efficient parallel job performance

The Linux default network and network device settings might not produce optimum throughput (bandwidth) and latency numbers for large parallel jobs. The information that is provided describes how to tune the Linux network and certain network devices for better parallel job performance.

This information is aimed at private networks with high-performance network devices such as the Gigabit Ethernet network, and might not produce similar results for 10/100 public Ethernet networks.

The following table provides examples for tuning your Linux system for better job performance. By following these examples, it is possible to improve the performance of a parallel job that runs over an IP network.

Table 3. Network tuning

Network Tuning Factors	Tuning for the current boot session	Modifying the system permanently
<b>arp_ignore</b> - With <b>arp_ignore</b> set to 1, a device answers only to an ARP request if the address matches its own.	echo '1' > /proc/sys/net/ipv4/conf/all/arp_ignore	Add this line to the /etc/sysctl.conf file: net.ipv4.conf.all.arp_ignore = 1
<b>arp_filter</b> - With <b>arp_filter</b> set to 1, the kernel answers only to an ARP request if it matches its own IP address.	echo '1' > /proc/sys/net/ipv4/conf/all/arp_filter	Add this line to the /etc/sysctl.conf file: net.ipv4.conf.all.arp_filter = 1
<b>rmem_default</b> - Defines the default receive window size.	echo '1048576' > /proc/sys/net/core/rmem_default	Add this line to the /etc/sysctl.conf file: net.core.rmem_default = 1048576
<b>rmem_max</b> - Defines the maximum receive window size.	echo '2097152' > /proc/sys/net/core/rmem_max	Add this line to the /etc/sysctl.conf file: net.core.rmem_max = 2097152
<b>wmem_default</b> - Defines the default send window size.	echo '1048576' > /proc/sys/net/core/wmem_default	Add this line to the /etc/sysctl.conf file: net.core.wmem_default = 1048576
<b>wmem_max</b> - Defines the maximum send window size.	echo '2097152' > /proc/sys/net/core/wmem_max	Add this line to the /etc/sysctl.conf file: net.core.wmem_max = 2097152
<b>Set device txqueuelen</b> - Sets each network device, for example, eth0, eth1, and on.	/sbin/ifconfig <i>device_interface_name</i> txqueuelen 4096	Not applicable
<b>Turn off device interrupt coalescing</b> - To improve latency.	See sample script. This script must be run after each reboot.	Not applicable

This sample script unloads the e1000 Gigabit Ethernet device driver and reloads it with interrupt coalescing disabled:

```
#!/bin/ksh
Interface=eth0
Device=e1000
Kernel_Version=`uname -r`
ifdown ${Interface}
rmmod ${Device}
insmod /lib/modules/${Kernel_Version}/kernel/drivers/net/${Device}/${Device}.ko \
InterruptThrottleRate=0,0,0
ifconfig ${Interface}
exit $?
```

MPI jobs use shared memory to handle intranode communication. You might need to modify the system default for *allowable maximum shared memory size* to allow a large MPI job to successfully enable shared memory usage. It is recommended that you set the system *allowable maximum shared memory size* to 256 MB or larger for supporting large MPI jobs.

To modify this limit for the current boot session, run the following command as root:

```
echo "268435456" > /proc/sys/kernel/shmmax
```

To modify this limit permanently, add the following line to the `/etc/sysctl.conf` file and reboot the system:

```
kernel.shmmax = 268435456
```

*Table 4. Network tuning: ARP entries*

Network tuning factors	Tuning for the current boot session and updating it into the boot image
<b>gc_thresh3</b> - The value in <code>/proc/sys/net/ipv4/neigh/default/gc_thresh3</code> should be the maximum number of compute operating system nodes, plus 300.	<code>echo "5300" &gt;/proc/sys/net/ipv4/neigh/default/gc_thresh3</code>
<b>gc_thresh2</b> - The value in <code>/proc/sys/net/ipv4/neigh/default/gc_thresh2</code> should be 100 less than <b>gc_thresh3</b> .	<code>echo "5200" &gt;/proc/sys/net/ipv4/neigh/default/gc_thresh2</code>
<b>gc_thresh1</b> - The value in <code>/proc/sys/net/ipv4/neigh/default/gc_thresh1</code> should be 100 less than <b>gc_thresh2</b> .	<code>echo "5100" &gt;/proc/sys/net/ipv4/neigh/default/gc_thresh1</code>
<b>gc_interval</b> - The ARP garbage collection interval on the compute nodes should be high so that it does not process ARP cleanup.	<code>echo "1000000000" &gt;/proc/sys/net/ipv4/neigh/default/gc_interval</code>
<b>gc_stale_time</b> - The ARP stale time should be set high so that it does not get discarded.	<code>echo "2147483647" &gt;/proc/sys/net/ipv4/neigh/default/gc_stale_time</code>
<b>base_reachable_time_ms</b> - The ARP valid entry time (in milliseconds) should be set high so that it does not get discarded.	<code>echo "2147483647" &gt;/proc/sys/net/ipv4/neigh/default/base_reachable_time_ms</code>

DNS caching should be enabled to minimize runtime host name resolution, especially if LDAP is also enabled in the cluster.

---

## Understanding the effects of changing the SMT or HT mode

Simultaneous multi-threading (SMT) is a function of Power Systems servers that allows multiple logical CPUs to share physical core. This same function for Intel is called hyper-threading (HT). The SMT and HT settings can be changed by the system administrator at any time, and it is important to understand how changing these settings affects cpusets and running jobs.

For cpusets, note the following:

- After lowering the SMT mode or disabling HT for a core:
  - The newly disabled CPUs are removed from the existing cpuset
  - `sched_setaffinity()` automatically skips the newly disabled CPUs
  - The newly disabled CPUs are removed from the binding list of running processes.
- When raising the SMT mode or enabling HT for a core, the newly enabled CPUs are not added to the existing cpuset.
- If the SMT or HT mode is changed, the CPU numbering remains the same.

For jobs that are currently running, note the following:

- Lowering the SMT mode or disabling the HT for a core means that the job that is running is updated by the kernel to use the most recently available CPUs.
- When raising the SMT mode or enabling HT for a core, the existing *per-job* cpuset and the job that is running are not affected. However, the system administrator needs to regenerate the CPUs contained in the *parallel\_jobs* cpuset so that the newly enabled CPUs will be available for future jobs.



---

## Chapter 3. Installing IBM Spectrum MPI

---

### Install the IBM Spectrum MPI software step-by-step

The step-by-step procedure for installing the IBM Spectrum MPI software on an IBM Power® System or x86, Little Endian mode server follows.

#### View the release notes before installation

Before you install the IBM Spectrum MPI package, you might want to look at the release notes. The release notes might contain some special or additional information about installing the package and are available from the following websites:

Passport Advantage website at <http://www.ibm.com/software/howtobuy/passportadvantage/>

IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter))

#### Step 1: Obtain software packages

This step consists of copying the installation images off the distribution medium and exporting the installation directory that provides the installation images available for mounting. You must complete this step if any of the machines in your cluster do not have the proper installation device to read the distribution medium.

Install the IBM Spectrum MPI installation image. Use the **rpm** command.

To install IBM Spectrum MPI using the **rpm** command, first refer to the list of required RPMs. Then, install the following components in the order shown:

1. IBM Spectrum MPI license RPM
2. IBM Spectrum MPI RPM

#### Copy software packages to the system

To copy the IBM Spectrum MPI software off the distribution medium, follow these instructions:

**INSERT**

the distribution medium in the installation device.

**COPY** the packages with the `.rpm` extension to the local file system.

#### Step 2: Set root authority

You must have root authority to install the package.

#### Step 3: Install RPMs and accept the license packages

You must install and accept the IBM Spectrum MPI license packages on each node of a cluster to successfully install IBM Spectrum MPI. The license must be installed before or concurrently with the IBM Spectrum MPI component base RPM. Note also that IBM Spectrum MPI checks the current license during run time.

By default, the license and product packages are installed in the `/opt/ibm/spectrum_mpi` directory. You can choose an alternate path to install packages by passing the command line flag `--prefix` to the `rpm` command, specifying the alternate path.

**Note:** If you choose an alternate path, the same path must be used for both the license and product packages.

The `rpm` command attempts to create the alternate installation path if it does not exist.

You can choose how to process the IBM Spectrum MPI license. For unattended installation, the license agreement can be accepted automatically when the license package is installed. Alternately, the license agreement can be viewed and accepted manually after the license package is installed.

The license agreement can be accepted automatically, without viewing it, as follows:

1. Select the license package appropriate for your platform (this example uses the Power Systems platform).
2. Login as root.
3. Set the environment variable `IBM_SPECTRUM_MPI_LICENSE_ACCEPT=yes`.
4. Enter: `rpm -i ibm_smpi_lic_s-10.1-rh7.ppc64le.rpm`. The license is not displayed, but it is accepted automatically.

The license agreement can be viewed and accepted manually as follows:

1. Select the license package appropriate for your platform (this example uses the Power Systems platform).
2. Login as root.
3. Set the environment variable `IBM_SPECTRUM_MPI_LICENSE_ACCEPT=no`.
4. Enter: `rpm -i ibm_smpi_lic_s-10.1-rh7.ppc64le.rpm`.
5. When installation is complete, the license acceptance script must be invoked to review and accept the license. The script can be found under the installation path directory at `lap_se/bin/accept_spectrum_mpi_license.sh`.

To install the IBM Spectrum MPI RPMs, you can use the Linux `rpm` command. In general, IBM supports the `-i`, `-U`, and `-e` RPM options only.

1. Select the product package appropriate for your platform (this example uses the RHEL6 x86\_64 platform).
2. Login as root.
3. If you have not done so previously, install the IBM Spectrum MPI license package and accept the license agreement.
4. Enter: `rpm -i ibm_smpi-10.1.0.0-rh6.x86_64.rpm` to install the package.

## Viewing the release notes after installation

After the package is installed always review the IBM Spectrum MPI product release notes, available from the following websites:

Passport Advantage website at <http://www.ibm.com/software/howtobuy/passportadvantage/>

IBM Knowledge Center ([www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter))



---

## Migration installation

IBM Spectrum MPI is a new product. Migration from previous releases of IBM MPI or IBM Parallel Environment is not supported.

**CAUTION:**

**Any MPI application needs to be recompiled and rebuilt with new MPI libraries.**

---

## Resolving installation errors

If an installation error occurs, system administrators might need to resolve the situation by removing the IBM Spectrum MPI software by using **rpm** command options (such as the **-noscripts** option), correcting the issue, and then reinstalling the software.

Installation errors can be caused by:

- Lack of sufficient disk space to install the software. In such cases, refer to the information about the required disk space for installation, and expand the file systems or remove unneeded files, to allow for sufficient disk space.



---

## Chapter 4. Performing installation-related tasks

After installation of IBM Spectrum MPI, a number of tasks might need to be performed from time to time.

---

### Removing a software component

You can remove any of the IBM Spectrum MPI RPMs, 1 RPM at a time, by using the **rpm -e** command with the name of the RPM you want to remove. However, because the IBM Spectrum MPI components depend on each other, you cannot randomly delete any of these RPMs; they must be removed in the reverse order in which they were installed. The RPMs must be deleted in the following order:

1. IBM Spectrum MPI RPM
2. IBM Spectrum MPI license RPM

---

### Finding installed components

To determine which IBM Spectrum MPI product RPMs, or which IBM Spectrum MPI license RPM is installed, you can use the **rpm -qa** command with the **grep** command.

To determine which IBM Spectrum MPI product RPMs are installed, enter the following command:

```
rpm -qa | grep smpi-
```

Similar output appears:

```
ibm_smpi-10.1.0.0-rh7.ppc64le.rpm  
ibm_smpi-10.1.0.0-rh7.x86_64.rpm  
ibm_smpi-10.1.0.0-rh6.x86_64.rpm
```

**Note:** Depending on the platform, only one of these RPMs is installed.



---

## Chapter 5. Understanding how installing IBM Spectrum MPI alters your system

Your system is altered when you install IBM Spectrum MPI.

---

### Understanding the IBM Spectrum MPI directory structure

Installing IBM Spectrum MPI creates directories in the default installation path:

```
/opt/ibm/spectrum_mpi
```

These directories include:

```
[c712f6n01][/opt/ibm/spectrum_mpi]>ls  
bin etc examples include lap_se lib profilesupport properties share  
[c712f6n01][/opt/ibm/spectrum_mpi]>
```



---

## Chapter 6. Additional information for the system administrator

---

### Configuring key-based login with OpenSSH

To use key-based login with OpenSSH, users must generate SSH key files and add the contents of their public key file to the `authorized_keys` file in the `$HOME/.ssh` directory. If key-based login is not already configured for end users, the system administrator performs the following tasks:

1. If the `$HOME/.ssh` directory does not exist, create it, as follows:

```
mkdir $HOME/.ssh
```

```
chmod 700 $HOME/.ssh
```

2. If the end user does not have SSH keys, generate them with the **ssh-keygen** command (enter **CR** at all prompts).
3. If the end user's public key is not added to the `$HOME/.ssh/authorized_keys` file, add it as follows:

```
cat $HOME/.ssh/public_key_file >> $HOME/.ssh/authorized_keys
```

Where *public\_key\_file* is the `.pub` file that is generated by the **ssh-keygen** command.





---

## Chapter 7. Using example applications

IBM Spectrum MPI provides example programs that demonstrate basic communication between ranks, I/O, and thread-safe programming concepts. Both C and FORTRAN compilers must be installed to compile and run the example programs.

For more information, see:

`/opt/ibm/spectrum_mpi/examples.`



---

## Accessibility features for IBM Spectrum MPI

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

### Accessibility features

IBM Spectrum MPI includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

IBM Spectrum MPI uses the latest W3C Standard, WAI-ARIA 1.0 ([www.w3.org/TR/wai-aria/](http://www.w3.org/TR/wai-aria/)), to ensure compliance with US Section 508 ([www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards)) and Web Content Accessibility Guidelines (WCAG) 2.0 ([www.w3.org/TR/WCAG20/](http://www.w3.org/TR/WCAG20/)). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by IBM Spectrum MPI.

The IBM Spectrum MPI online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at [http://www.ibm.com/support/knowledgecenter/doc/kc\\_help.html#accessibility](http://www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

### Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

### IBM and accessibility

For more information about the commitment that IBM has to accessibility, see IBM Accessibility ([www.ibm.com/able](http://www.ibm.com/able)).



---

## Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not

been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_.

---

## Programming interface information

### MPI support statement

The IBM Spectrum MPI product is a complete MPI implementation, based on the Open MPI open source project, designed to comply with all the requirements of the Message Passing Interface standard, MPI: A Message-Passing Interface Standard, Version 3.1, University of Tennessee, Knoxville, Tennessee, June 4, 2015. If you believe that IBM Spectrum MPI does not comply with the MPI-3.1 standard, please contact IBM Service.

### Open MPI license

Most files in this release are marked with the copyrights of the organizations who have edited them. The copyrights below are in no particular order and generally reflect members of the Open MPI core team who have contributed code to this release. The copyrights for code used under license from other parties are included in the corresponding files.

Copyright (c) 2004-2010 The Trustees of Indiana University and Indiana University Research and Technology Corporation. All rights reserved.

Copyright (c) 2004-2010 The University of Tennessee and The University of Tennessee Research Foundation. All rights reserved.

Copyright (c) 2004-2010 High Performance Computing Center Stuttgart, University of Stuttgart. All rights reserved.

Copyright (c) 2004-2008 The Regents of the University of California. All rights reserved.

Copyright (c) 2006-2010 Los Alamos National Security, LLC. All rights reserved.

Copyright (c) 2006-2010 Cisco Systems, Inc. All rights reserved.

Copyright (c) 2006-2010 Voltaire, Inc. All rights reserved.

Copyright (c) 2006-2011 Sandia National Laboratories. All rights reserved.

Copyright (c) 2006-2010 Sun Microsystems, Inc. All rights reserved. Use is subject to license terms.

Copyright (c) 2006-2010 The University of Houston. All rights reserved.

Copyright (c) 2006-2009 Myricom, Inc. All rights reserved.

Copyright (c) 2007-2008 UT-Battelle, LLC. All rights reserved.

Copyright (c) 2007-2010 IBM Corporation. All rights reserved.

Copyright (c) 1998-2005 Forschungszentrum Juelich, Juelich Supercomputing Centre, Federal Republic of Germany

Copyright (c) 2005-2008 ZIH, TU Dresden, Federal Republic of Germany

Copyright (c) 2007 Evergrid, Inc. All rights reserved.

Copyright (c) 2008 Chelsio, Inc. All rights reserved.

Copyright (c) 2008-2009 Institut National de Recherche en Informatique. All rights reserved.

Copyright (c) 2007 Lawrence Livermore National Security, LLC. All rights reserved.

Copyright (c) 2007-2009 Mellanox Technologies. All rights reserved.

Copyright (c) 2006-2010 QLogic Corporation. All rights reserved.

Copyright (c) 2008-2010 Oak Ridge National Labs. All rights reserved.

Copyright (c) 2006-2010 Oracle and/or its affiliates. All rights reserved.

Copyright (c) 2009 Bull SAS. All rights reserved.  
Copyright (c) 2010 ARM ltd. All rights reserved.  
Copyright (c) 2010-2011 Alex Brick . All rights reserved.  
Copyright (c) 2012 The University of Wisconsin-La Crosse. All rights reserved.  
Copyright (c) 2013-2016 Intel, Inc. All rights reserved.  
Copyright (c) 2011-2014 NVIDIA Corporation. All rights reserved.

\$COPYRIGHT\$

Additional copyrights may follow

\$HEADER\$

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer listed in this license in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

The copyright holders provide no reassurances that the source code provided does not infringe any patent, copyright, or any other intellectual property rights of third parties. The copyright holders disclaim any liability to any recipient for claims brought against recipient by any third party for infringement of that parties intellectual property rights.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.



Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

---

## Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

### Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

### Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

---

## IBM Online Privacy Statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled “Cookies, Web Beacons and Other Technologies”, and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

---

# Index

## A

- accessibility features for this product 23
- administrator, additional information for 19
  - configuring key-based login 19
- API subroutine libraries, described 1

## C

- components 1
- conventions and terminology vii

## D

- debug programs
  - installation 5

## E

- errors, resolving for installation 13
- Example applications, additional 21

## F

- file systems 6

## H

- hardware requirements 3

## I

- IBM Spectrum MPI documentation, described 1
- IBM Spectrum MPI feature
  - how installation alters system 17
- install
  - license packages 11
  - RPMs 11
- installation errors, resolving 13
- installation requirements 3
- installation-related tasks, performing 15
  - finding installed components 15
  - removing a software component 15
- installing
  - copy software packages 11
  - viewing the release notes after installation 12
- installing IBM Spectrum MPI 11
  - IBM Spectrum MPI for Linux 13
  - resolving installation errors 13
  - understanding directory structure 19
  - understanding how installing IBM Spectrum MPI alters your system 17

## L

- limitations 5

## M

- memory, setting the limit 6
- migration installation 13

## N

- node resources 5

## O

- Operating system requirements 3

## P

- performance tuning, Linux system 7
- planning to install IBM Spectrum MPI 3
- processes, setting the maximum number 6

## R

- release notes 11
- removing a software component 15
- requirements
  - installation 3
- resources, node 5
- root authority
  - set 11

## S

- software packages
  - obtain 11
- software requirements 3, 4
- software, planning to install IBM Spectrum MPI 3
- system administrator, information for 5
- system partitioning 5

## T

- tuning, Linux system 7

## U

- user authorization 6
- user IDs on remote nodes 6







Product Number: 5725-G83

Printed in USA

GC27-8264-00

