Message Passing Toolkit: Release Notes

Document Number 007–3689–003

St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

The MPI implementation for the CRAY T3E system is derived from the implementation of MPI for the CRAY T3D system developed at Edinburgh Parallel Computing Centre. The software is supplied to Cray Research under license from The University of Edinburgh.

PVM (Parallel Virtual Machine) is based on software that was developed by the Oak Ridge National Laboratory, the University of Tennessee, and Emory University. This work was supported in part by the Applied Mathematical Sciences subprogram of the Office of Energy research, U.S. Department of Energy, in part by the National Science Foundation, and in part by the State of Tennessee.

Contents

Page

Page

Introduction [1]

This document provides an overview of the Cray Message Passing Toolkit and Message Passing Toolkit for IRIX (MPT) 1.3 release. MPT is a software package that supports interprocess data exchange for applications that use concurrent, cooperating processes on a single host or on multiple hosts. Data exchange is done through *message passing*, which is the use of library calls to request data delivery from one process to another or between groups of processes.

The MPT 1.3 release is supported on the following platforms:

- Cray PVP systems running UNICOS release 10.0 or later. The MPT 1.3 release requires a bugfix package to be installed on UNICOS systems running release 10.0.0.0 or 10.0.0.1. The bugfix package, `MPT12_OS_FIXES`, is available through the `getfix` utility. It is also available from the anonymous FTP site `ftp.cray.com` in directory `/pub/mpt/fixes/MPT12_OS_FIXES`.
- CRAY T3E systems running UNICOS/mk release 1.5 or later.
- Silicon Graphics MIPS based systems running IRIX release 6.2 or later.

  IRIX 6.2 systems running Parallel Virtual Machine (PVM) require the POSIX patch set and any patches recommended by the patch set.

  IRIX 6.2 systems running Message Passing Interface (MPI) require the kernel rollup patch 1650 or later.

  IRIX 6.3 systems running MPI require the kernel rollup patch 2328 or later.

  IRIX systems running MPI applications must also be running Array Services software version 3.0 or later.

  IRIX systems using the Miser feature must be running IRIX 6.5 or later and Array Services version 3.1 or later.

  MPI jobs that use the checkpoint and restart feature require IRIX 6.4 plus kernel patch 2839 and CPR patch 2689, or IRIX 6.5. In addition, CPR patch 3017 fixes IRIX 6.4 and IRIX 6.5 problems that exist with the `ash` (array session handler) option of the `cpr`(1) command. Array Services 3.1 or later is also required for this feature.

  Activation of the MPI usability improvements feature is dependent on Array Services 3.2 plus patch 3532.

These release notes include information about the following:

- New features
- Compatibilities and differences
- Documentation
- Customer services
- Release package contents

The remainder of this chapter discusses the following topics:

- MPT components
- MPT software
- Product description
- Distribution of these release notes
- Reader comments

1.1 MPT Components

The MPT 1.3 package contains the following components and accompanying documentation:

- Parallel Virtual Machine (PVM)
- Message Passing Interface (MPI)
- Logically shared, distributed memory (SHMEM) message passing routines

The PVM and MPI components of MPT are based on industry standards for the message passing programming model. These components also contain enhancements specific to Silicon Graphics and Cray Research systems. For a description of MPI, see the *Message Passing Toolkit: MPI Programmer's Manual*. For a description of PVM, see the *Message Passing Toolkit: PVM Programmer's Manual*.

The SHMEM library is packaged with MPT on UNICOS and IRIX systems only. On UNICOS/mk systems, the SHMEM library is delivered with the CrayLibs package. SHMEM man pages are printed in the *Application Programmer's Library Reference Manual*.

1.2 MPT Software

For UNICOS and UNICOS/mk systems, software included in MPT was designed to be used with the Cray Programming Environment 3.0 release or later. The `Modules` software package is used to support the installation of both the Programming Environment and MPT. To use the MPT software, load the `mpt` module in addition to loading the Programming Environment module. For information on using modules, see *Installing Programming Environment Products*, or, if the Programming Environment has already been installed on your system, see the online ASCII file `/opt/ctl/doc/README`. After you have initialized the modules, enter the following command to access the MPT software:

```
module load mpt
```

To unload the `mpt` module, enter the following command:

```
module unload mpt
```

For IRIX systems, MPT software can be installed in an alternate location for use with the modules software package. If MPT software has been installed on your system for use with modules, you can access the software with the `module` command shown in the previous example. If MPT has not been installed for use with modules, the software resides in default locations on your system (`/usr/include`, `/usr/lib`, `/usr/array/PVM`, and so on), as in previous releases. For further information, see *Installing MPT for Use with Modules*, in the `relnotes` for modules.

1.3 Product Description

The MPT package contains the following subcomponents:

- An optimized version of Message Passing Interface (MPI) for IRIX systems, UNICOS systems, and UNICOS/mk systems. This version is based on the MPI 1.2 specification from the Message Passing Interface Forum. It supports high-speed communications among the processors on each system. It also supports high-speed communications between UNICOS systems of like architecture, and between IRIX systems running IRIX 6.2 or later.
- An optimized version of Parallel Virtual Machine (PVM) for IRIX systems, UNICOS systems, and UNICOS/mk systems. This version is based on Oak Ridge National Laboratories (ORNL) version 3.3.10. It supports high-speed communications among the processors on each system. It also supports high-speed communications between heterogeneous systems.

- Logically shared, distributed memory (SHMEM) message passing capability using the `shmem_get`(3) and `shmem_put`(3) functions on IRIX systems and UNICOS systems. (SHMEM on UNICOS/mk systems is delivered with CrayLibs.) This high-performance library supports communications among the processors on each system. For more information, see the `intro_shmem`(3) man page.

MPI is a standard specification for message passing libraries, allowing portable message passing programs in the Fortran and C languages. MPI was created by the Message Passing Interface Forum (MPIF). MPIF is not sanctioned or supported by any official standards organization. Its goal was to develop a widely used standard for writing message passing programs. The implementation of this standard includes a library (`libmpi.a` for UNICOS and UNICOS/mk systems and `libmpi.so` for IRIX systems), a run-time command (`mpirun`(1)), and a library that allows profiling of message passing applications (`libpmpi.a` for UNICOS and UNICOS/mk systems and `libpmpi.so` for IRIX systems). On UNICOS and IRIX systems, additional commands and libraries are available to help you with debugging.

PVM is a software project that was developed jointly by Oak Ridge National Laboratories (ORNL), the University of Tennessee, and Emory University. PVM consists of a main library, a user-level daemon (`pvmd3`(1)), a console (`pvm`(1)), and some additional commands and libraries. The main PVM library name differs depending on the platform.

On IRIX systems, the PVM library name is `libpvm3.so` (formerly, `libpvm.so`). An additional library, named `libpvm3-i8.so` (formerly, `libpvm-i8.so`), has Fortran `INTEGER*8` interfaces so that 64-bit integers can be used. Symbolic links from `libpvm.so` to `libpvm3.so` and from `libpvm-i8.so` to `libpvm3-i8.so` are provided with this release for compatibility. The links will be removed in a future version of MPT.

On UNICOS and UNICOS/mk systems, the PVM library is named `libpvm3.a`.

The PVM and MPI libraries provide communications and synchronization functions that are necessary for writing distributed applications. For example, you can add calls that cause one task to send a message to another, or to receive a message, or to wait until another task is finished. The PVM software supports heterogeneous systems by automatically converting data. The MPI software is supported between UNICOS systems of the same architecture and within a partition on UNICOS/mk systems. This software also supports communication across all machines running IRIX 6.2 or later.

SHMEM message passing is another form of distributed programming. It differs from PVM and MPI message passing in that it uses one-sided

communication (that is, one processing element (PE) can send or receive data from another PE without the knowledge of that PE).

To use PVM, MPI, or SHMEM message passing directly, you must change your source code to add the appropriate calls.

The CRAY T3D versions of PVM and SHMEM message passing are part of the Cray Fortran and Cray Standard C Programming Environment releases. MPT does not include these implementations, but does include Cray PVP support that you need when using PVM to distribute applications that run partly on CRAY T3D systems. MPI support for CRAY T3D systems is provided through a third-party product available from Edinburgh Parallel Computing Centre (EPCC). On CRAY T3E systems, PVM and MPI are part of MPT, and SHMEM is part of the Programming Environment.

1.4 Distribution of These Release Notes

You can access these release notes electronically, as follows:

- On IRIX systems, after you have installed the online release notes (the `relnotes` subsystem), you can view the release notes on your screen.
  - If you have a graphics system, select `Release Notes` from the `Help` submenu of the Toolchest. This invokes the `grelnotes` graphical browser for the online release notes. For options to this command, see the `grelnotes`(1) man page.
  - If you do not have a graphics system, you can use the `relnotes`(1) command to access the online release notes.
- ASCII and PostScript files are also available on the following systems:
  - Cray CRInform system, which is an online information and problem-reporting system for Cray customers. For more information on CRInform, see Section 5.5, page 17.
  - Cray `craypark` system (formerly the `hydra` system) in the `/home/craypark/release_docs` directory. The `craypark` system is available only to Cray service personnel.

1.5 Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and part number of the document with your comments.

You can contact us in any of the following ways:

- Send electronic mail to the following address:

  `techpubs@sgi.com`

- Send a facsimile to the attention of "Technical Publications" at fax number +1 650 932 0801.

- Use the Suggestion Box form on the Technical Publications Library World Wide Web page:

  `http://techpubs.sgi.com/library/`

- Call the Technical Publications Group, through the Technical Assistance Center, using one of the following numbers:

  For Silicon Graphics IRIX based operating systems: 1 800 800 4SGI

  For UNICOS or UNICOS/mk based operating systems or CRAY Origin2000 systems: 1 800 950 2729 (toll free from the United States and Canada) or +1 651 683 5600

- Send mail to the following address:

  Technical Publications
  Silicon Graphics, Inc.
  2011 North Shoreline Boulevard, M/S 535
  Mountain View, California 94043–1389

We value your comments and will respond to them promptly.

New Features [2]

This chapter describes the new MPT 1.3 features.

2.1 MPI Scaling to Support 48x128 Clusters

MPI on IRIX systems can now support up to 48 hosts and 128 processes per host. The internal MPI buffering scheme on IRIX has also changed. In MPT 1.3, separate pools of buffers are used to handle interhost and intrahost MPI communications.

2.2 Multiboard Adapter Selection

On IRIX systems, MPI now automatically detects multiple HIPPI network adapters and uses as many of them as possible when sending messages among hosts. The multiboard feature uses a "round robin" selection scheme in choosing the next available adapter over which to send the current message. The message is sent entirely over one adapter.

2.3 Thread-Safe MPI

On IRIX systems, MPI calls are thread-safe. MPI processes can be multi-threaded. Each thread associated with a process can issue MPI calls. Thread-safe MPI supports both POSIX threads and share group processes (known as sprocs). This feature is dependent on MPI being run on an IRIX 6.5 system or later.

2.4 MPI Collectives Optimization

The `MPI_Barrier`(3) routine is optimized to utilize fetch operations (known as fetch-op) hardware when available. This results in dramatically reduced barrier synchronization times.

2.5 Compile-time Interface Checking for Fortran MPI Subroutine Calls

An MPI interface definition module called `mpi_interface` has been added to permit Fortran 90 programmers to verify correctness of MPI subroutine calls at

compile time. To activate interface checking, specify the `-auto_use mpi_interface` option on the `f90` command line, as in the following example:

```
f90 -64 -auto_use mpi_interface program.f -lmpi
```

This feature depends on new versions of the Fortran 90 compiler. You must use Fortran 90 compiler release 7.2.1 or higher. Previous versions cannot process the `mpi_interface` definition module, nor can they process the `-auto_use` command line option.

For installations in which MPT 1.3 is installed in an alternate location, this feature also depends on new versions of the `mpt` and `MIPSpro` module files from the Modules Software package. You must get the module files that were packaged with Modules Software version 2.2.1.1 or later. To determine if your Modules Software is of the proper level, type the following command:

```
versions modules
```

For more details of the use of the `-auto_use` option to activate compile-time checking, see the `MPI`(1) command.

2.6 MPI Statistics

New MPI statistics routines allow you to view MPI internal buffer management usage. With the `-stats` option specified on the `mpirun`(1) command, each process prints statistics about the amount of data sent with MPI calls during the `MPI_Finalize` process. Routines for resetting and retrieving statistics are also available. For more information, see the `MPI_SGI_stat`(3) man page. These routines are available only on IRIX systems.

2.7 MPI I/O

MPT 1.3 contains the ROMIO implementation of MPI I/O, in which a rich API for performing I/O in a message passing application is defined. This feature is available on IRIX and UNICOS/mk systems.

2.8 XMPI and XPVM

With the MPT 1.3 release, support for the XMPI and XPVM utilities has been removed.

2.9 Cray TotalView Support

With the MPT 1.3 release, support for startup of MPI jobs on UNICOS systems is available with Cray TotalView version 3.0.0.16.

2.10 Message Queue Display

For IRIX systems, support for MPI message queue display with the Dolphin ToolWorks TotalView Multiprocess Debugger has been added. For more details, see the ToolWorks page at the following URL:

`http://www.dolphinics.com/toolworks`

2.11 MPI Usability Improvements

On IRIX systems, MPI start-up error diagnostics replace the error message, `MPI: could not run executable`. The new diagnostics are more descriptive than the former message, returned by the `mpirun`(1) command. This feature is dependent on Array Services 3.2 and patch 3532.

Compatibilities and Differences [3]

This chapter describes the compatibility issues that users should consider when purchasing MPT version 1.3.

3.1 XMPI and XPVM Retired

The XMPI and XPVM performance analysis tools have been retired. MPI program performance analysis can be accomplished through the use of SpeedShop and MPI statistics. For more information on MPI statistics, see the `MPI_SGI_stat`(3) man page.

Documentation [4]

This chapter describes the documentation that supports the MPT 1.3 release. It contains information about the following:

- Online information access
- Documentation included with the release

4.1 Online Information Access

The following types of online information products are available to MPT 1.3 customers:

- Silicon Graphics publications information at the following URLs:

  `http://techpubs.sgi.com/library`

  `http://techpubs.sgi.com/infosearch`

  These websites contain information that allows you to browse documents online, order documents, and send feedback to Silicon Graphics.

  Cray also has documents available online at the following URL:

  `http://www.cray.com/swpubs`

- The Cray DynaWeb server, which allows you to view manuals online using a World Wide Web browser such as Netscape or Mosaic. Please see your Cray system administrator for the local URL of the Cray DynaWeb server.
- The *User Publications Catalog*, describes the availability and content of all Cray hardware and software documents for customers. Silicon Graphics customers who subscribe to the Cray Inform (CRInform) program can access this information on the CRInform system.
- Man pages, which describe a particular element of the operating system or a compatible product. To see a detailed description of a particular command or routine, use the `man`(1) command.
- A message system, available on UNICOS and UNICOS/mk systems, which provides explanations of error messages. To see an explanation of a message, use the `explain`(1) command.

- Cray online glossary, available on UNICOS and UNICOS/mk systems, which explains the terms used in a manual. To get a definition, use the `define`(1) command.
- IRIS InSight, available on IRIX systems, which allows you to view manuals online. For a description of IRIS InSight, see the `insight`(1) man page

4.2 Documentation Included with the Release

Most documentation included with the MPT 1.3 release is available through the Cray DynaWeb server, through IRIS InSight, and in printed form. The *SHMEM Library Ready Reference* is available in printed form only. For information on accessing these documents, see Section 4.1, page 13.

The following documentation is available to MPT 1.3 customers:

- *Message Passing Toolkit: MPI Programmer's Manual*
- *Message Passing Toolkit: PVM Programmer's Manual*
- *Message Passing Toolkit: Release Notes*
- *SHMEM Library Ready Reference*

The man pages for MPI and PVM are available in online form only. SHMEM man pages are printed in the *SHMEM Library Ready Reference* and in the *Application Programmer's Library Reference Manual*. The following publications are available for purchase through the Minnesota Distribution Center. For ordering information, see Section 6.4.1, page 28.

- *Application Programmer's Library Reference Manual*
- *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, by Gropp, Lusk, and Skjellum, publication TPD–0011

Customer Services [5]

This chapter describes the MPT support policy and the following customer services that are available to support the MPT 1.3 release:

- MPT World Wide Web page
- Training support
- Software problem reporting and resolution process
- Cray Research CRInform program
- *Pipeline* and the *Pipeline Supercomputing Supplement*
- Additional resources

5.1 Support Policy

Silicon Graphics offers standard support for problems in the following categories:

- Build problems
- Run-time problems specific to Silicon Graphics and Cray computer systems

Fixes to problems will be made available through one of the following mechanisms:

- Upgrade utility
- Major release
- Minor release

5.1.1 MPI Support Policy

Silicon Graphics will address MPI problems in a timely fashion.

5.1.2 PVM Support Policy

Silicon Graphics will help customers resolve PVM problems by forwarding general PVM problems to developers at Oak Ridge National Laboratories (ORNL) and the University of Tennessee, and assisting in their resolution.

Problems with Silicon Graphics enhancements to PVM will be addressed in a timely fashion.

5.1.3 SHMEM Library Support Policy

Silicon Graphics will address SHMEM message passing problems in a timely fashion.

5.2 MPT World Wide Web Page

Current product information regarding MPT is available online through World Wide Web sites. Access to the information is available at the following URLs:

```
http://www.cray.com/products/software/pe/mpt.html
```

```
http://www.sgi.com/Products/software/mpt.html
```

5.3 Training Support

For information about Silicon Graphics courses, training office locations, current class schedules, and training services listed by country, contact us in one of the following ways:

- On the World Wide Web at the following URL:

  ```
  http://www.sgi.com/support/custeducation.html
  ```

- Call one of the following numbers:

  1 800 800 4744 (toll free in the United States or Canada)

  +1 651 683 3825 (from all other locations)

5.4 Software Problem Reporting and Resolution Process

If you experience problems with MPT 1.3, contact your service representative. Your service representative will work with you to resolve the problem. If you choose to have full-time or part-time on-site support, your on-site support personnel are your primary contacts for service. If you have elected not to have on-site support, please call the Customer Support Center and report your problem to them, or submit a request for technical assistance (RTA) through the CRInform program (see Section 5.5).

When you report problems, use the product name most appropriate for the problem and refer to MPT for all products. Use any of the following product names:

- PVM
- MPI
- LIBSMA

For UNICOS and UNICOS/mk customers who need current information on the status of reported problems, see the Software Problem Report (SPR) section of the CRInform program.

5.5 CRInform Program

The Cray Inform (CRInform) program is a World Wide Web-based information and problem-reporting service for UNICOS and UNICOS/mk customers. Using the CRInform program, you can do the following:

- Report software problems
- Request technical assistance
- Communicate directly with other UNICOS and UNICOS/mk customers
- Read about software problems similar to yours reported at other sites
- Learn about solutions to various problems
- Find information about classes
- Read about new products, and more

The CRInform program automatically logs as news items those events that are pertinent to your site, so you do not have to search through the system for new information. The logged events include changes in SPR or RTA activity, new orderable software, new issues of the *Pipeline Supercomputing Supplement*, new field notices (FNs), new software release documents, new software problem fix information, new marketing information, and new CRInform program information. You can also get automatic e-mail notification of any or all of the news items.

Version 4.0 of the CRInform program is available through the World Wide Web. You need access to the CRInform Web server and a browser (such as Mosaic, Netscape, or Lynx), which allows you to view information or make service

requests. You can use your own site's browser, or use either the Mosaic or the Lynx browser available on the `crinform` system.

5.6 *Pipeline* and the *Pipeline Supercomputing Supplement*

Customers who have a support contract receive *Pipeline*, the Silicon Graphics customer newsletter. Customers who have a support contract for a Cray system also receive the *Pipeline Supercomputing Supplement*. *Pipeline* provides product and support information about Silicon Graphics workstations and servers, and the *Pipeline Supercomputing Supplement* provides product and support information about Cray supercomputers. *Pipeline* and the *Pipeline Supercomputing Supplement* are both published six times a year (January/February, March/April, and so on). *Pipeline* is available on the World Wide Web in Supportfolio Online (`http://support.sgi.com/`), and the *Pipeline Supercomputing Supplement* is available in CRInform (`http://crinform.cray.com/`) (see Section 5.5, page 17).

5.7 Additional Resources

Along with the standard support described in the preceding sections, users of PVM and MPI have a variety of other resources available to help them use these products. These resources should supplement, not replace, standard Silicon Graphics support. In particular, users are encouraged to use Silicon Graphics support when they encounter what may be a PVM or MPI bug, even if the problem does not apply to Silicon Graphics systems. However, these other resources can be helpful in gaining a better understanding of PVM and MPI and how they can be used. The following sections describe some of these resources.

5.7.1 MPI Standard

You can access the MPI standard in any of the following ways:

- As online PostScript or hypertext on the World Wide Web at the following URL:

  `http://www.mcs.anl.gov/mpi`

- As a journal article in the fall 1994 issue of the *Journal of Supercomputing Applications*

5.7.2 News Groups

The `comp.parallel.pvm` and `comp.parallel.mpi` Internet news groups focus on PVM and MPI, respectively, and related products. The PVM and MPI developers use these news groups to announce new releases of the products and to discuss related topics, such as the PVM User Group meeting (see Section 5.7.5). Users are also free to use the news groups to interact with one another by responding to postings made by others.

To access `comp.parallel.pvm` or `comp.parallel.mpi`, you must have access to the Internet and a news reader program. If necessary, contact your site administrator for additional information.

5.7.3 World Wide Web Servers

Several PVM-related web servers are available. There is one for `netlib` (see Section 5.7.7, page 20) that includes the PVM source code, other material on PVM, and additional software packages and subjects. The URL for this server is as follows:

```
http://www.netlib.org/pvm/index.html
```

Another Web page specific to PVM can be found at the following URL:

```
http://www.epm.ornl.gov/pvm/pvm_home.html
```

5.7.4 MPI User Group Meetings

You can obtain information about upcoming MPI developers conferences and user group meetings through `comp.parallel.mpi` or through the MPI web page listed in Section 5.7.1, page 18.

5.7.5 PVM User Group Meetings

There have been several PVM user group meetings held in the past, and there are plans to continue these in the future. These meetings have been quite successful, attracting a wide range of presentations from the PVM development team, vendors working on PVM, and users working on PVM-related tools or using PVM in applications.

Announcements of PVM user group meetings are posted in the news group `comp.parallel.pvm`.

5.7.6 PVM E-Mail Support

The PVM developers have established the following e-mail alias, which they use to answer questions and provide support, as time permits:

```
pvm@msr.epm.ornl.gov
```

It is suggested that users first try Silicon Graphics support or contacts within the PVM news group when asking questions. This e-mail address, however, can be used when needed.

The MPI Forum e-mail discussions, and both current and earlier versions of the standard, are available from the `netlib` server. For more information about `netlib`, see the following section.

5.7.7 `netlib` Source

Public domain versions of PVM and MPI are available from a source called `netlib`. If your site is using the versions of these products contained in MPT, you will not need to obtain the source code from `netlib`. However, the following additional resources available from `netlib` might be useful:

- Test and sample programs contained in the public domain releases that are not included in the Silicon Graphics release
- Papers and presentations by the PVM and MPI developers that describe PVM, MPI, and related products
- Presentations from past PVM user group meetings

You can obtain information from `netlib` in the following ways:

- Use the Web, which allows you to obtain the files directly. This is perhaps the easiest and most convenient method. The PVM page at `netlib` is located at the following URL:

  ```
  http://www.netlib.org/pvm/index.html
  ```

  The MPI page at `netlib` is located at the following URL:

  ```
  http://www.netlib.org/mpi/index.html
  ```

- Send electronic mail to `netlib@ornl.gov` (or `netlib@research.att.com`), and include the following in the `Subject:` line or message body:

  ```
  send index from pvm3
  ```

or

```
send index from mpi
```

Once you receive the index, similar `send` messages will return parts of PVM or MPI. This assumes the ORNL or AT&T host can determine how to return your e-mail. If not, use anonymous `ftp`.

- Use anonymous `ftp` from `netlib2.cs.utk.edu`.
- Use anonymous `rcp` from `netlib2.cs.utk.edu`. For example:

  `rcp anon@netlib2.cs.utk.edu:pvm3/`*your-local-file*

  or

  `rcp anon@netlib2.cs.utk.edu:mpi/`*your-local-file*

  You can use the following commands to obtain file lists:

  ```
  rsh netlib2.cs.utk.edu -l anon ls pvm3
  ```

  or

  ```
  rsh netlib2.cs.utk.edu -l anon ls mpi
  ```

- Use the `xnetlib` browsing tool (which you can obtain from `netlib` by using one of the preceding methods).
- For access in Europe, use the duplicate collection in Oslo, as follows:

  Internet: `netlib@nac.no`

  EARN/BITNET: `netlib%nac.no@norunix.bitnet`

  X.400: `s=netlib; o=nac; c=no`

  EUNET/uucp: `nac!netlib`

- For access in the Pacific, use the following address, located at the University of Wollongong, NSW, Australia:

  `netlib@draci.cs.uow.edu.au`

Release Package [6]

This chapter contains the following information about the MPT 1.3 release package:

- Summary of hardware and software requirements
- Licensing information
- List and description of the contents of the release package
- Ordering instructions

6.1 Hardware and Software Requirements

To build and use MPT software on Cray PVP systems, CRAY T3E systems, and MIPS based systems, you must have the following software:

- UNICOS release 10.0 or later for Cray PVP systems; UNICOS/mk release 1.5 or later for CRAY T3E systems; and IRIX release 6.2 or later for MIPS based systems.

  The MPT 1.3 release requires a bugfix package to be installed on UNICOS systems running release 10.0.0.0 or 10.0.0.1. The bugfix package, `MPT12_OS_FIXES`, is available through the `getfix` utility. It is also available from the anonymous FTP site `ftp.cray.com` in directory `/pub/mpt/fixes/MPT12_OS_FIXES`.

  IRIX 6.2 systems running PVM require the POSIX patch set and any patches recommended by the patch set.

  IRIX 6.2 systems running MPI require the kernel rollup patch 1650 or later.

  IRIX 6.3 systems running MPI require the kernel rollup patch 2328 or later.

- For UNICOS and UNICOS/mk systems, Cray Programming Environment releases 3.0 or later.

- For UNICOS and UNICOS/mk systems, the `Modules` software package. Because the Cray Programming Environment 3.0 releases use the `Modules` software package, in addition to having the Cray Programming Environment installed on your system, you must have the `PrgEnv` module loaded before you can use MPT. Modules are not required for IRIX systems.

For information on using modules, see *Installing Programming Environment Products*.

- For IRIX systems running MPI applications, Array Services release 3.0 or later. For UNICOS systems, Array Services is included in the required UNICOS versions.
- For MPI jobs using the checkpoint and restart feature, IRIX 6.4 plus kernel patch 2839 and CPR patch 2689, or IRIX 6.5. For IRIX 6.4 and IRIX 6.5 problems with the `ash` (array session handler) option of the `cpr` command, CPR patch 3017.
- For IRIX 6.5 systems using checkpoint and restart or Miser features, Array Services release 3.1 or later.
- For IRIX systems using the MPI usability improvements, Array Services release 3.2 or later.

6.2 Licensing Information

To order MPT 1.3, customers must pay for the product and agree to the stated terms and conditions that relate to the product. The following sections provide names and addresses of licensing contacts.

6.2.1 Licensing contacts for customers in the U.S. and Canada

For more information on the licensing and pricing of MPT 1.3, customers in the United States and Canada should see their account representative or field contract negotiator.

6.2.2 Licensing Contacts for Customers outside of the U.S. and Canada

Customers outside of the United States and Canada may obtain further licensing and export information by contacting their account representative or the individuals listed in the following sections.

6.2.2.1 European Regional Contract Administrators/Specialists

Customers in European sales regions can direct licensing inquiries to the following regional Contract Administrators/Specialists:

United Kingdom, Poland, and Russia:

Nick Jacobs
Silicon Graphics Ltd.
1530 Arlington Business Park
Theale
Reading
Berkshire, RG7 4SB
England
Telephone: +44 118 925 7041
Fax: +44 118 925 7716
E-mail: `niknik@reading.sgi.com`

Germany:

Thomas Wegener
Silicon Graphics GmbH
Am Hochacker 3
85630 Grasbrunn-Neukeferloh
Germany
Telephone: +49 89 461080
Fax: +49 89 46108 222
E-mail: `thomasw@munich.sgi.com`

France:

Cecile Goachet or Patricia Guillerm-Brillet
Silicon Graphics France
21 rue Albert Calmette
78351 Jouy en Josas
France
Telephone: +33 01 34 88 80 00
Fax: +33 01 34 65 96 19
E-mail: `cgoachet@paris.sgi.com` or `pguiller@paris.sgi.com`

All other parts of Europe:

Simon Locke
Silicon Graphics Ltd.
1530 Arlington Business Park
Theale
Reading
Berkshire, RG7 4SB
England
Telephone: +44 118 925 7049
Fax: +44 118 925 7946
E-mail: `simonl@reading.sgi.com`

6.2.2.2 Japan Contract Negotiators

Customers in Japan can direct licensing inquiries to Joel Lee at the following address:

Joel Lee
Japan Legal Counsel
Nihon Silicon Graphics Cray K. K.
P. O. Box 5011 Yebisu Garden Place Tower
4-20-3, Ebisu, Shibuya-ku, Tokyo 150
Japan
Telephone: +81 3 5488 1819
Fax: +81 3 5420 7020
E-mail: `jlee@nsg.sgi.com`

6.3 Release Package Contents

The MPT 1.3 release package includes the following:

- Cray Message Passing Toolkit release CD
- *Message Passing Toolkit: Release Notes* (this manual)
- *Message Passing Toolkit: PVM Programmer's Manual*
- *Message Passing Toolkit: MPI Programmer's Manual*
- *Installing Programming Environment Products*
- *SHMEM Library Ready Reference*

MPT 1.3 consists of the following software:

- PVM
- MPI
- SHMEM `libsma` library (on UNICOS and IRIX systems)
- Header files
- PVM, MPI, and SHMEM man pages
- `Modules` software

The MPT 1.3 release CD for IRIX systems also includes the Array Services release 3.2 package, product code SC4-PCAS-3.0.

6.4 Ordering the MPT 1.3 Release Package

MPT 1.3 is distributed by order to sites that have agreed to the stated terms and conditions that relate to the product (for licensing details, see Section 6.2, page 24). The most current revision of the release package is supplied.

You can use one of the following methods to order MPT 1.3:

- Customers outside of the United States and Canada should contact their account representative or contract negotiator for ordering information.
- Customers who subscribe to the CRInform program can order software release packages electronically by using the `Order Cray Software` option.

- Customers can contact the order desk at the Minnesota Distribution Center by telephone (+1 651 683 5907 or 1 800 284 2729 extension 35907) or through electronic mail (`orderdsk@cray.com`).

Software will be shipped by ground service or 5-day international service unless otherwise requested.

Sites outside the United States may be required to provide a customer-signed Letter of Assurance before this software can be shipped. Address questions about which customers must sign Letters of Assurance or which software requires such a letter to Alan Benfell at the following address:

Alan Benfell
International Administration
Silicon Graphics, Inc.
655F Lone Oak Drive
Eagan, MN 55121 USA
Telephone: +1 651 683 7476 or 1 800 284 2729
Fax: +1 651 683 7297
E-mail: `benfa@sgi.com`

6.4.1 Obtaining Publications

The MPT 1.3 release package includes the publications indicated in Section 6.3, page 27.

Silicon Graphics maintains publications information at the following URLs:

`http://techpubs.sgi.com/library`

`http://techpubs.sgi.com/infosearch`

These websites contain information that allows you to browse documents online, order documents, and send feedback to Silicon Graphics.

Cray also has documents available online at the following URL:

`http://www.cray.com/swpubs`

The *User Publications Catalog* describes the availability and content of all Cray hardware and software documents for customers. Customers who subscribe to the Cray Inform (CRInform) program can access this information on the CRInform system.

To order a document, call the Minnesota Distribution Center at +1 651 683 5907. Silicon Graphics employees may send electronic mail to `orderdsk@sgi.com` (UNIX system users).

Customers who subscribe to the CRInform program can order software release packages electronically by using the `Order Cray Software` option.

Customers outside of the United States and Canada should contact their local service organization for ordering and documentation information.

Updated Release Information [A]

This appendix describes information about MPT update releases, including information about significant bug fixes and features. The update release information that appears in this appendix is available only through use of the `relnotes`(1) command and is not available via `insight`(1).

In subsequent major, minor, or revision releases, the update release information is deleted from this appendix and included with the list of new features or in other appropriate parts of the manual.

Index

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-3689-003.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
  - On the Internet: techpubs@sgi.com
  - For UUCP mail (through any backbone site): *[your_site]*!sgi!techpubs
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 650-932-0801
- To send your comments by **traditional mail**, use this address:

  Technical Publications
  Silicon Graphics, Inc.
  2011 North Shoreline Boulevard, M/S 535
  Mountain View, California 94043-1389