Frequently Asked Questions (FAQ) for Intel® Cluster Toolkit 3.2 31-Oct-2008

Intel Corporation

How do I use this FAQ? 31-Oct-2008

This FAQ for Intel® Cluster Toolkit 3.2 is divided into the following hyperlinked topics:

How do I use this FAQ? 31-Oct-2008
Disclaimer and Legal Information 22-Apr-2008
Product Capabilities 17-Oct-2008
Intel® Cluster Tool Acronyms 15-Oct-2007
Registration 21-Oct-2004
Installation 16-Oct-2008
Intel® MPI Library 26-Jul-2008
Intel® Debugger (IDB) 10-Oct-2008
Intel® Math Kernel Library 12-Jun-2008
Intel® Trace Analyzer and Collector 07-Oct-2008
Intel® MPI Benchmarks 07-Oct-2008
Uninstalling the Intel® Cluster Toolkit 12-Jun-2008
Version Compatibility 12-Jun-2008

Each topic has a calendar date associated with it indicating as to when it was last modified. Within each topic, there are fact items and/or questions which are also dated according to when those fact items or questions were last revised.

Back to Top

Disclaimer and Legal Information 22-Apr-2008

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

MPEG is an international standard for video compression/decompression promoted by ISO. Implementations of MPEG CODECs, or MPEG enabled platforms may require licenses from various entities, including Intel Corporation.

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

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

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

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

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Atom, Centrino Atom Inside, Centrino Inside, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skoool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2004-2008, Intel Corporation. All rights reserved.

Back to Top

Product Capabilities 17-Oct-2008

What is the Intel® Cluster Toolkit? 17-Oct-2008

The Intel® Cluster Toolkit 3.2 provides the high performance computing (HPC) user with four cluster-computing software tools bundled into a single integrated package with one license, one installation process, and one support service. The four component names in the toolkit are:

- Intel® MPI Benchmarks 3.2
- Intel® MPI Library 3.2
- Intel® Math Kernel Library 10.1
- Intel® Trace Analyzer and Collector 7.2

The operating systems that the Intel® Cluster Toolkit has been developed for are Linux*, Microsoft* Windows* Compute Cluster Server (also known as Microsoft Windows CCS*), and Microsoft* Windows* HPC Server 2008. HPC is an acronym for high-performance computing. In terms of the Intel Cluster Toolkit software for Windows*, consider references within this FAQ to Microsoft Windows CCS* OS and Microsoft* Windows* HPC Server 2008 OS **as interchangeable**.

Back to Top

What are the Intel® processor architectures that the Intel® Cluster Toolkit will run on? 06-Oct-2006

The Intel® processor architectures are:

Intel® Pentium® 4 processor, or
Intel® Xeon® processor, or
Intel® Itanium® 2 processor, or
Intel® Core™2 Duo processor (example of Intel® 64 (formerly Intel® EM64T) architecture)

Note that it is assumed that the processors listed above are configured into homogeneous clusters. Also, for the Microsoft Windows CCS* operating system, only processors based on the Intel® 64 architecture are supported.

Back to Top

What version of the Linux* operating system can I install the Intel® Cluster Toolkit on? 24-Sep-2008

Presently Intel® Cluster Toolkit is supported on:

OS Distributions	IA-32 Architecture	Intel® 64 Architecture		IA-64 Architecture
		32-Bit Applications	64-Bit Applications	
SGI* Propack* 5 for Linux*		S	S	S
Red Hat Enterprise Linux* 4.0	S	S	S	S
Red Hat Enterprise Linux* 5.0	S	S	S	S
SUSE Linux Enterprise Server* 9	S	S	S	S
SUSE Linux Enterprise Server* 10	S	S	S	S

S = Supported

Back to Top

What Microsoft* Windows* operating system versions can I install the Intel® Cluster Toolkit on? 17-Oct-2008

Presently Intel® Cluster Toolkit is supported on:

Microsoft* Windows* Distributions	IA-32 Architecture	Intel 64 Architecture		IA-64 Architecture
		32-Bit Applications	64-Bit Applications	
Microsoft* Windows* Compute Cluster Server (Microsoft Windows CCS*)	N/A	S	S	N/A
Microsoft* Windows* HPC Server 2008	N/A	S	S	N/A

S = Supported

Back to Top

Intel® Cluster Tool Acronyms 15-Oct-2007

Acronym	Definition
ABI	Application Binary Interface - describes the low-level interface an application program and the operating system, between an application and its libraries, or between component parts of an application.
BLACS	Basic Linear Algebra Communication Subprograms - provides a linear algebra oriented message passing interface for distributed memory computing platforms.
BLAS	Basic Linear Algebra Subroutines
DAPL	Direct Access Program Library - an Application Program Interface (API) for Remote Data Memory Access (RDMA).

DFT	Discrete Fourier Transform
Ethernet	Ethernet is the predominant local area networking technology. It is transports data over a variety of electrical or optical media. It transports any of several upper layer protocols via data packet transmissions.
GB	Gigabyte
ICT	Intel® Cluster Toolkit
ICTCE	Intel® Cluster Toolkit Compiler Edition
IMB	Intel® MPI Benchmarks
IP	Internet protocol
ITA or ita	Intel® Trace Analyzer
ITAC or itac	Intel® Trace Analyzer and Collector
ITC or itc	Intel® Trace Collector
MPD	Multi-purpose daemon protocol - a daemon that runs on each node of a cluster. These MPDs configure the nodes of the cluster into a “virtual machine” that is capable of running MPI programs.
MPI	Message Passing Interface - an industry standard, message-passing protocol that typically uses a two-sided send-receive model to transfer messages between processes.
NFS	The Network File System (acronym NFS) is a client/server application that lets a computer user view and optionally store and update file on a remote computer as though they were on the user's own computer. The user's system needs to have an NFS client and the other computer needs the NFS server. Both of them require that you also have TCP/IP installed since the NFS server and client use TCP/IP as the program that sends the files and updates back and forth.
PVM*	Parallel Virtual Machine
RAM	Random Access Memory

RDMA	Remote Direct Memory Access - this capability allows processes executing on one node of a cluster to be able to "directly" access (execute reads or writes against) the memory of processes within the same user job executing on a different node of the cluster.
RDSSM	TCP + shared memory + DAPL* (for SMP clusters connected via RDMA-capable fabrics)
RPM*	Red Hat Package Manager* - a system that eases installation, verification, upgrading, and uninstalling Linux packages.
ScaLAPACK	SCAlable LAPACK - an acronym for Scalable Linear Algebra Package or Scalable LAPACK.
shm	Shared memory only (no sockets)
SMP	Symmetric Multi-processor
ssm	TCP + shared memory (for SMP clusters connected via Ethernet)
STF	Structured Trace Format - a trace file format used by the Intel Trace Collector for efficiently recording data, and this trace format is used by the Intel Trace Analyzer for performance analysis.
TCP	Transmission Control Protocol - a session-oriented streaming transport protocol which provides sequencing, error detection and correction, flow control, congestion control and multiplexing.
VML	Vector Math Library
VSL	Vector Statistical Library

Back to Top

Registration 21-Oct-2004

When using the Intel web-based registration process, I encounter diagnostic error messages such as “Invalid password format” or “Invalid challenge answer”. What am I doing wrong? 21-Oct-2004

For the password and challenge answer the minimum number of characters must be 8. For example, if the challenge question would be: “What is your mother’s maiden name?”, and your mother’s maiden name is 5 characters, you might want to increase the text length by doubling the maiden name length. E.g., suppose the maiden name is “Smith”, then enter “SmithSmith”.

Back to Top

Installation 16-Oct-2008

I tried installing the Intel® Cluster Toolkit on the Ubuntu operating system and I could not. What am I doing wrong? 30-Jun-2007

The Intel® Cluster Toolkit should install just fine on the Ubuntu operating system, if you do not select the RPM install. You can do this by running the `install.sh` script and selecting the non-rpm choice, or you can use the `--nonrpm` option as part of the command-line for `install.sh`.

Back to Top

How come install packages for Intel® Cluster Toolkit 3.0.1 or greater are so large with respect to previous releases? 26-Jun-2007

The installation packages for Intel® Cluster Toolkit 3.0.1 or greater are much larger due to the integration of expanded Intel® Math Kernel Library functionality features. This will significantly increase the time required in doing a download of the Intel® Cluster Toolkit installation package. Installing the cluster toolkit on every node of a cluster may also take substantially longer. Note that during the installation process the installer may need approximately 4 gigabytes of temporary disk storage to manage the intermediate installation files.

Back to Top

For the Linux* operating system, does the Intel® Cluster Toolkit installer script which is called `install.sh` have any command-line options? 12-Jun-2008

The answer is yes. Simply type a shell command referencing `install.sh` such as:

```
install.sh --help | less
```

and you will see a list of options that look something like:

```
NAME
    install.sh - Install Intel(R) Cluster Toolkit for Linux*
3.2.

SYNOPSIS
    install.sh [options]

OPTIONS
    --help  Print this help and exit.

    --install-path=PATH
    --installpath=PATH
            Install the product to specified path.

    --temp-path=PATH
    --temppath=PATH
            Use specified folder for temporary files. By
default, /tmp/install.XXXXXX folder is used.

    --ict-update-path=PATH
    --ict_update_path=PATH
            Update an existing installation located at the
specified path.

    --license-path=PATH
    --licensepath=PATH
            Search specified path (a directory or a file)
for a license.

    --log-file=FILE
            Write log to the specified file.

    --nonroot
            Do not ask for root password.

    --nonrpm
            Perform non-RPM installation.

    --use-config=FILE
    --silent=FILE
            Perform silent installation by using responses
from FILE.

    --arch=x86
```

```
            Install x86-specific binaries on Intel(R) EM64T
platform.

COPYRIGHT
    Copyright 1999-2008, Intel Corporation. All Rights
Reserved.
```

Back to Top

For the Linux* operating system, how much free space storage do I need in `/tmp` to complete an install of the Intel® Cluster Toolkit? 18-Oct-2007

We recommend that you have at least two gigabytes of free space in `/tmp` when doing an install of the Intel® Cluster Toolkit. Also, the installer script `install.sh` has the command-line options:

```
--temp-path=PATH
```

or

```
--temppath=PATH
```

which can direct the use of temporary intermediate storage to another disk partition. Again if an alternative disk partition is used during installation, approximately 4 gigabytes of temporary disk storage should be available to manage the intermediate installation files.

Back to Top

Are there any prerequisite tools that need to be present prior to installing the Intel® Cluster Toolkit? 30-Nov-2006

For full functionality of the Intel® Cluster Toolkit software package on Linux*, you will need to acquire additional software resources, if they are not already installed on your system:

- For all of the Intel® processor architectures running Red Hat Enterprise Linux 3.0, Red Hat Enterprise Linux 4.0, Red Hat Enterprise Linux 5.0, SUSE Linux Enterprise Server 9, SUSE Linux 9.0, or SUSE Linux 9.1, SUSE Linux Enterprise Server 10, the version number on the Intel® Compilers should be 8.1 or greater.
- Perl is Open Source software where the Perl programming language has been optimized for scanning arbitrary text files,

extracting information from those text files, and generating printed reports based on that information. It is also a good language for many system administrative tasks. Perl is used in the Intel® Cluster Toolkit by the Intel® Trace Collector. You can download it for free as source code, or as a pre-compiled binary distribution. The Perl source code is available for distribution by setting a Web browser at the URL: http://www.perl.com/pub/a/language/info/software.html. At this URL, click on the hyperlink titled, "Stable Production Release".

- Python is Open Source software where the high-level built in data structures in combination with dynamic typing and dynamic binding, make it very attractive for rapid application development, as well as for use as a scripting or glue language to connect existing software components together. Python is an interpreted, interactive, object-oriented programming language. The current release of Python can be retrieved from the URL: http://www.python.org/download. If Python needs to be installed on your system, then be sure to retrieve the current production version. The Python programming language is used by Intel® MPI Library.
- The 4.2.2 version of RPM on Red Hat Enterprise Linux 3.0 for Itanium® 2 has a broken relocation feature. This will be a serious problem for users trying to do installs on clusters where there are shared devices. A recommended solution is for the user to upgrade to the latest release of RPM. A possible URL for retrieving a recent release of RPM that resolves this problem on the Itanium® 2 architecture is: http://www.redhat.com.
- Expect is a tool for automating interactive applications. You can download the expect software package from the following URL: http://expect.nist.gov/.
- PAPI* is an acronym for Performance API and it serves to gather information regarding performance counter hardware. Details can be found at the URL: http://icl.cs.utk.edu/papi/. The PAPI interface works in conjunction with Intel® Trace Analyzer and Collector.

Back to Top

On Itanium® 2-based systems running Red Hat Enterprise Linux* 3.0, the following symptom is occurring for RPM version 4.2.2 when doing an install with the `--nonroot` command-line option:

```
install --nonroot

the following symptom appears:

Installing Intel(R) MPI Library for Linux* version
p_3.2.011...
Extracting
files...    ##############################################
###   [100%]

Would you like to:

    Install:

    i.  Intel(R) MPI Library, Development Kit for Linux*
version 3.1-038

    x.  Exit

Your choice? (i/x) [i]: i

Cannot install Intel(R) MPI Library, Development Kit for
Linux* version 3.2-038 to
"/home/toolsteam/intel/ict/3.2.0.035/impi":
RPM version 4.2.2 prevent installation into a non-default
directory.
Installation failed.
WARNING: Intel(R) MPI Library for Linux* installer did
not create "/tmp/install.XX2svw1O/MPI.ini" file; the
product cannot be installed to all cluster nodes.
Press Enter to continue...
```

What am I doing wrong? 12-Jun-2008

The solution is to use an updated version of RPM.

Back to Top

In what instances do I need to use the `expect` software package? 30-Nov-2006

Within the “tar” package of the Intel Cluster Toolkit, there is an expect shell script file called `sshconnectivity.exp`. This `expect` shell script will create or update a `~/.ssh` directory on each node of the cluster beginning with the master node. If you do not have ssh-connectivity established on your Linux* cluster, the `sshconnectivity.exp` might be useful. As noted previously, you can download the expect software package from the following URL: http://expect.nist.gov/.

Back to Top

Where can I find TCL for Linux*? 21-Oct-2006

Any Linux implementation will probably have a TCL RPM where you should be able to download the TCL RPM off of that respective Linux vendor's website. For example the following web site might be useful for downloading TCL RPMs: http://rpm.pbone.net/index.php3/stat/3/limit/2/srodzaj/1/dl/40/search/tcl/dist%5b%5d/47/dist%5b%5d/44/dist%5b%5d/41/dist%5b%5d/37/dist%5b%5d/31/dist%5b%5d/32/dist%5b%5d/1/dist%5b%5d/28/dist%5b%5d/2/dist%5b%5d/3/dist%5b%5d/4/dist%5b%5d/5/dist%5b%5d/7/dist%5b%5d/42/dist%5b%5d/36/dist%5b%5d/35/dist%5b%5d/11/dist%5b%5d/12/dist%5b%5d/49/dist%5b%5d/45/dist%5b%5d/34/dist%5b%5d/13/dist%5b%5d/14/dist%5b%5d/15/dist%5b%5d/17/dist%5b%5d/18/dist%5b%5d/46/dist%5b%5d/48/dist%5b%5d/33/dist%5b%5d/19/dist%5b%5d/20/dist%5b%5d/23/dist%5b%5d/24/dist%5b%5d/25/dist%5b%5d/26/dist%5b%5d/27/dist%5b%5d/29/dist%5b%5d/30/dist%5b%5d/38/dist%5b%5d/39/dist%5b%5d/40/dist%5b%5d/43/dist%5b%5d/50/field%5b%5d/1/field%5b.

Back to Top

I am trying to build an `expect` executable based on a download from http://expect.nist.gov/ and I am encountering the following error:

. . .

```
checking for Tcl private headers... checking for tclInt.h... no
configure: error: Can't find Tcl private headers
```

. . .

when running the `configure` command. What am I doing wrong? 30-Mar-2006

It is assumed that `expect` and `Tcl` are both stored in a path such as `/usr/local/src`. For example, the paths might be:

```
/usr/local/src/tcl8.4.12          (actual version may be different)
/usr/local/src/expect-5.43        (actual version may be different)
```

The header file `tclInt.h` is located in the generic folder under a `Tcl` installation path such as `/usr/local/src/tcl8.4.12`.The `--with-tclinclude` option should be used with the configure command to

build `expect`. The configure command might look something like the following:

```
configure --with-tclinclude=/usr/local/src/tcl8.4.12/generic
```

Back to Top

For the Linux* operating system, I have established secure shell (SSH) connectivity on my cluster where each node has a public key. However, when I use any secure shell utility on the cluster, I am still prompted for a password. What am I doing wrong? 30-Nov-2006

Assuming that you have set up a NULL "`passphrase`" for your login account when using the `ssh-keygen` utility, for all systems in your cluster, and that you have distributed the public key for each node to all the other nodes (i.e., the public keys reside in the "`~/.ssh/authorized_keys`" file), then check the permissions of your home directory on your cluster. The "group" and "other" permission categories should only have read and execute accessibility for the home directory folder. If you activate a "write" permission setting for either the "group" or "other" categories, SSH will prompt you for the password. Also, the "`~/.ssh`" subdirectory should only have read, write, and execute permission for the login owner.

Note that the `expect` shell script `sshconnectivity.exp` will remove the write access capability on the group and other "permission categories" for the user's home directory folder.

Back to Top

Regarding a Linux* installation, how do I install the Intel® Trace Analyzer on Microsoft Windows XP*? 12-June-2008

Suppose that the installation staging area for the Intel Cluster Toolkit is:

```
/tmp/ict_staging_area
```

Within this directory after untarring the tar package you should see a folder that looks something like the following:

```
l_ict_<major>.<minor>.<update>.<package_num>
```

Recall that *<major>.<minor>.<update>.<package_num>* is a string such as:

`b_3.2.0.035`, where `b` is an acronym for beta

or

`p_3.2.0.035`, where `p` is an acronym for production

The *<package_num>* meta-symbol is a string such as `035`. This string indicates the package number. Change to that directory with the shell command:

```
cd l_ict_<major>.<minor>.<update>.<package_num>
```

Within this folder, there is Microsoft Windows* installer package that should like something like `w_itac_p_7.2.011.exe`. Transfer this package along with the license file that was issued to you for the Intel® Cluster Toolkit over to a Microsoft Windows XP* system and proceed to do an install of the Intel® Trace Analyzer on Windows. All of the features that are described for Intel® Trace Analyzer for Linux* are available for this Intel® Trace Analyzer for Microsoft Windows* application.

Back to Top

For the Linux* operating system, when I run the script `install.sh`, I am seeing diagnostics which look something like the following:

```
                    . . .
Cannot install Intel(R) Trace Analyzer And Collector
version 7.2p-029 into directory
"/usr/home/user01/intel/ict/3.2.0.035/itac/7.1":
The prefix directory path
"/usr/home/user01/intel/ict/3.2.0.035/itac/7.1" is owned
by cookie intel-ta_ipf version 7.2p-029.
Installation failed.
                    . . .
```

What am I doing wrong? 28-Apr-2008

This symptom occurs when a prior installation of the Intel® Cluster Toolkit into a directory such as `/usr/home/user01/intel/ict/3.2.0.035` was removed without

having issued the `uninstall.sh` command that for *this* example, resides in the directory `/usr/home/user01/intel/ict/3.2.0.035`. Try the following corrective action:

1) Run an `uninstall.sh` command that should look something like the following:

```
/usr/home/user01/intel/ict/3.2.0.035/uninstall.sh
```

2) Reissue the `install.sh` command as done previously
3) If you see a prompt of the following type:

```
Would you like to:

    Install:

    1.  Intel(R) Trace Collector version 7.2p-029
    2.  Intel(R) Trace Analyzer version 7.2p-029
    a.  All of the above.

    Uninstall:

    3.  intel-ta_ipf version 7.2p-029
        (non-rpm in
"/usr/home/user01/intel/ict/3.2.0.035/itac")
    4.  intel-tac_ipf version 7.1p-029
        (non-rpm in
"/usr/home/user01/intel/ict/3.2.0.035/itac")

    x.  Exit

Please type a selection (1/2/a/3/4/x) [a]:
```

Enter the digit 3 under the "`Uninstall`" menu selection. This will probably result in a message that looks something like:

```
"/usr/home/user01/intel/ict/3.2.0.035/itac" does not exist.
It looks like installation was damaged. Force
uninstallation.
Press Enter to continue...
```

4) For the above, repeat this process again for any item that appears in the "`Uninstall`" menu category prompt. You will eventually encounter messages that look something like:

```
Installing...
Installation successful.
To uninstall this package, run
"/usr/home/user01/intel/ict/3.2.0.035/itac/uninstall.sh".
```

5) Continue this process of doing the uninstall selection first, any time you encounter such subsequent prompts for the remaining software components(e.g., Intel® MPI Library):

```
Would you like to:

    Install:

    i.  Intel(R) MPI Library, Development Kit for Linux*
version 3.2p-011

    Uninstall:

    1.  intel-mpi-ipf version 3.2p-011
        (non-rpm in
"/usr/home/user01/intel/ict/3.2.0.035/impi")

    x.  Exit

Please type a selection (i/1/x) [i]:
                              . . .
```

6) Eventually you should see the following messages from the Intel® Cluster Toolkit installer:

```
                              . . .

   Installation successful.
   To uninstall this package, run
   "/usr/home/user01/intel/ict/3.2.0.035/impi/uninstall.sh".

   Completed cluster installation successfully.
```

Back to Top

When installing Intel® Cluster Toolkit 3.1 or greater, I am getting a diagnostic message regarding Intel® Math Kernel Library 10.0 that looks something like the following:

```
Installing Intel(R) Math Kernel Library 10.0 for Linux*
version p_10.0.3.020...

Detected operating system has not been tested:
  Red Hat Enterprise Linux AS release 4 (Nahant Update 3)
Install and program functionality may be incorrect.
Would you like to continue? ( Yes/No ) [ No ]:
```

What am I doing wrong? 12-Jun-2008

You are doing nothing wrong. Just simply type in “Yes” to the query and continue with the installation. Red Hat Enterprise Linux AS release 4 (Nahant Update 3) works fine with Intel® Math Kernel Library 10.0. This message has also been generated when installing Intel® Math Kernel Library 10.0 on SUSE 9.0 and SUSE 9.1, where the message can also be ignored by typing “Yes”.

Back to Top

During a Linux installation, I am using the the Linux installer option `--ict-update-path` in the following way:

```
./install.sh --ict-update-path=/opt/intel/ict/3.2.0.035
```

I am getting the diagnostic message:

```
ERROR: There is no valid Intel(R) Cluster Toolkit installation in the
specified directory.
```

What am I doing wrong? 05-May-2008

When you specify a Linux install option that looks something like the following:

```
--ict-update-path=/opt/intel/ict/3.2.0.035
```

on the command-line, the installer will try to install all the software components, but the base install directory will be different, and no links will be established for `ictvars.sh -> ictvars.01`. This defect will be resolved in an upcoming release.

Back to Top

During a Linux installation of the Intel® Cluster Toolkit, I am seeing a warning message that looks something like the following:

```
./Would you like to install the following?
    Intel(R) Math Kernel Library 10.1 for Linux*
<Enter> to continue, or x to exit:

Installing package...
```

```
WARNING: Unable to store license file due to insufficient
permissions.
Installation has been finished successfully.
```

What am I doing wrong? 16-Oct-2008

This is a warning message that is coming from the MKL sub-installer. It should be ignored. The MKL sub-installer shows such a message in case it was not able to save the license file into the directory `/opt/intel/licenses`. The master installer was able to complete the license file installation.

Back to Top

Intel® MPI Library 26-Jul-2008

When I attempt to run an MPI application on Microsoft Windows CCS*, via the `mpiexec` command, I am seeing an error message that looks something like:

```
Unable to connect to 'clusternode1:8678',
sock error: generic socket failure, error stack:
MPIDU_Sock_post_connect(1200): unable to connect to
clusternode1 on port 8678, exhausted all endpoints (errno -
1)
MPIDU_Sock_post_connect(1247): unable to connect to
clusternode1 on port 8678, No connection could be made
because the target machine actively refused it. (errno
10061)
```

What am I doing wrong? 18-Oct-2007

Open a DOS Window on the head node of the Microsoft Windows CCS* cluster and type the command:

```
clusrun smpd -status
```

This may provide a report that might look something like the following:

```
-------------- clusternode1 returns 805982518 --------------
-no smpd running on clusternode1

-------------- clusternode2 returns 671764790 --------------
no smpd running on clusternode2

-------------- clusternode3 returns 604655926 --------------
no smpd running on clusternode3
```

```
-------------- clusternode4 returns 873091382 --------------
no smpd running on clusternod4
```

where the number of nodes providing a status will be a function of the size of the cluster. If you observe the above, then from the same DOS Window issue one of the following commands:

```
clusrun smpd -install
```

or

```
clusrun smpd -regserver
```

You should see messages that look something like the following:

```
------------------- clusternode1 returns 0 -----------------
Intel(R) MPI Library Process Manager, Intel installed.

------------------- clusternode2 returns 0 -----------------
Intel(R) MPI Library Process Manager, Intel installed.

------------------- clusternode3 returns 0 -----------------
Intel(R) MPI Library Process Manager, Intel installed.

------------------- clusternode4 returns 0 -----------------
Intel(R) MPI Library Process Manager, Intel installed.
```

where again the number of nodes providing an installation message will be a function of the size of the cluster. You should now be able to launch an `mpiexec` command as before.

Back to Top

On Microsoft Windows CCS*, when I first login and open up my first DOS panel using the "Build Environment for Intel Cluster Toolkit" shortcut, and launch an `mpiexec` command which looks something like the following:

```
mpiexec -n 2 -machinefile
z:\cluster_file_share\machines.Windows my_mpi_prog.exe >
my_mpi_prog.report 2>&1
```

nothing happens. What am I doing wrong? 22-Oct-2007

You may want to open a second DOS panel with the "Build Environment for Intel Cluster Toolkit" shortcut, go to the working

directory where the `mpiexec` command was issued, and enter the command:

```
type <my-output-file>
```

where *<my-output-file>* might be something like `my_mpi_prog.report` as shown in the `mpiexec` command above. If you see a prompt in the text file which looks something like:

```
User credentials needed to launch processes:
account (domain\user) [clusternode1\user001]:
```

Either press the return key or enter a valid user account. This should be followed by a prompt for a password which looks something like:

```
account (domain\user) [clusternode1\user001]: password:
```

Enter an appropriate password. From here the application should proceed with execution of the `mpiexec` launch command.

Back to Top

When I run "`mpiifort test.f90`" I get the following error:

```
ERROR: could not determine icc directory for driver:
mpiifort
```

We are not using the Intel® C++ Compiler on our Linux* cluster system, nor does this error affect any of the other MPI compilation drivers. The Intel® Fortran Compiler works fine on its own. How can I fix this issue with Intel MPI Fortran compilation? 30-Nov-2006

The `mpiifort` compilation script calls a shell script named `compchk.sh`. A short term fix is to compile the application "`test.f90`" or "`test.f`" with "`-nocompchk`" as follows:

```
mpiifort -nocompchk test.f90
```

and

```
mpiifort -nocompchk test.f
```

Another solution is to install an Intel® C++ Compiler on the cluster system so as to resolve the apparent errors that are being detected by the shell script `compchk.sh`.

Back to Top

For the Linux* operating system, when compiling certain C++ applications using the Intel® MPI Library compilation driver **mpiicpc**, I encounter an error message that looks something like the following:

. . .

```
/opt/intel/ict/3.1.1/impi/3.1/include/mpicxx.h(45):
catastrophic error: #error directive: "SEEK_SET is #defined
but must not be for the C++ binding of MPI"
  #error "SEEK_SET is #defined but must not be for the C++
binding of MPI"
   ^

compilation aborted for my_mpi_program.cpp (code 4)
make: *** [my_mpi_program.o] Error 4
```

What am I doing wrong? 28-Apr-2008

For the application `my_mpi_program.cpp` try organizing the include files such that `mpi.h` comes before the `stdio.h`.

Back to Top

We have a cluster running the Linux* operating system, which has for example 3 nodes, where the **mpd.hosts** file contains:

```
$ cat mpd.hosts
node2
node3
```

The "master node", which is **"node1"** is not listed in the **mpd.hosts** file. The shell commands:

```
$ mpdboot -r ssh -n 2 -f ~/mpd.hosts
$ mpiexec -n 4 /bin/hostname
```

which are issued from **"node1"** reveal the following:

```
node1
node1
node2
node2
```

What am I doing wrong? 30-Nov-2006

By design, the "master node" (i.e., the node where the `mpiexec` command is launched from) needs to have a daemon running. If you want to run your application on `node2` and `node3`, you can use the `-host` options with the `mpiexec` command as follows:

```
$ mpdboot -r ssh -n 3 -f ~/mpd.hosts
$ mpiexec -n 2 -host node2 /bin/hostname : -n 2 -host node3
/bin/hostname
```

This will run the hostname command on nodes `node2` and `node3` only, even though there are three mpd daemons total.

Back to Top

On Microsoft Windows CCS*, I am running an MPI application with an **mpiexec** command such as the following:

```
$ mpiexec -n 4 ./my_prog.exe
```

and I immediately obtain a command-line prompt. No error messages are issued and I do not obtain any results. What am I doing wrong? 15-Oct-2007

During the link phase of building the application executable on Microsoft Windows CCS*, try adjusting the stack size setting for the executable by using the `/stack` option as is illustrated with the following compile and link command:

```
$ icl /D_CRT_SECURE_NO_DEPRECATE /Femy_prog
/I"%I_MPI_ROOT%"\em64t\include my_prog.c /link /stack:8000000
/LIBPATH:"%I_MPI_ROOT%\em64t\lib" impi.lib /NODEFAULTLIB:LIBCMTD.lib
```

For the example above, the value 8000000 sets the stack allocation in virtual memory to 8 million bytes. The default stack size is 1 MB.

Back to Top

On Microsoft Windows CCS*, I am running an MPI application with an **mpiexec** command such as the following:

```
$ mpiexec -n 4 ./my_prog.exe
```

and I am encountering a runtime error that looks something like the following:

```
forrtl: severe (29): file not found, unit 11, file
c:\Windows\system32\input.dat
```

I do not have an explicit reference to the folder path `c:\Windows\system32`. I am using a mapped drive which serves as a file share to all nodes of the cluster. What am I doing wrong? 26-Jul-2008

Reconfigure your `mpiexec` command by using the `-mapall` option:

```
$ mpiexec -n 4 -mapall ./my_prog.exe
```

Back to Top

Intel® Debugger (IDB) 10-Oct-2008

For Intel® MPI Library applications, what operating system or systems does the Intel® Debugger support? 10-Oct-2008

Intel® Debugger supports the debugging of Intel® MPI Library applications on Linux* OS.

Back to Top

To use the Intel® Debugger on Linux*, do I need to set any environment variables or set up any special source commands in my `.bashrc` or `.cshrc` file? 28-Apr-2008

Set the environment variable `IDB_HOME` to the folder path where the Intel Debugger executable, `idb`, resides. Also, you will want to source either `idbvars.sh` or `idbvars.csh` depending on which command-line shell you are using. For augmenting your `.bashrc` file, the Bourne* Shell or the Korn* Shell environment variable syntax for `IDB_HOME` and the sourcing of `idbvars.sh` might look something like:

```
export IDB_HOME=/opt/intel/idbe/10.1.015/bin
. ${IDB_HOME}/idbvars.sh
```

For augmenting your `.cshrc` file, the C Shell syntax should be something like:

```
setenv IDB_HOME /opt/intel/idbe/10.1.015/bin
source ${IDB_HOME}/idbvars.csh
```

Sourcing `idbvars.[c]sh` will update the `PATH` and `MANPATH` environment variables.

Back to Top

For the Linux* operating system, when I launch the Intel® Debugger through the Intel® MPI Library `mpiexec` command as follows:

```
mpiexec -idb -n 4 ./testcpp
```

the `mpiexec` command seems to hang or go into an infinite loop. What am I doing wrong? 30-Nov-2006

If you do not have a `~/.rhosts` file in your home directory (with permission bits set to 600), please create one. For each node in the computing cluster there should be a line in the `~/.rhosts` with the following information:

<hostname as echoed by the shell command hostname> <your username>

As an example, suppose there are 8 nodes in the computing cluster with the following host names:

```
clusternode1
clusternode2
clusternode3
clusternode4
clusternode5
clusternode6
clusternode7
clusternode8
```

Assuming that the names listed above make up your cluster, they could be added to your ~/`.rhosts` file with the following general syntax:

<hostname as echoed by the shell command hostname> <your username>

For the list of nodes above and assuming that your login name is user01, the contents of your ~/`.rhosts` file might be:

```
clusternode1 user01
```

```
clusternode2 user01
clusternode3 user01
clusternode4 user01
clusternode5 user01
clusternode6 user01
clusternode7 user01
clusternode8 user01
```

The permission bit settings of ~/.rhosts should be set to 600 using the chmod command. The shell command for doing this might be:

```
chmod 600 ~/.rhosts
```

Back to Top

For the Linux* operating system, the run command for the Intel® Debugger seems to cause my MPI application to hang. What am I doing wrong? 30-Nov-2006

The "run" command is disabled in MPI debugging. To continue the execution of the MPI application use "cont".

Back to Top

For the Linux* operating system, is there a rerun command when using the Intel® Debugger for an MPI application? 30-Nov-2006

Unfortunately, the rerun command is not yet supported within Intel Debugger. To restart, you will have to quit Intel Debugger and then re-enter the mpiexec command.

Back to Top

Intel® Math Kernel Library 12-Jun-2008

For the Linux* operating system, I am trying to build ScaLAPACK executables from the directory path /opt/intel/ict/3.2.0.035/mkl/10.0.3.020/tests/scalapack, where I am *not* explicitly setting the makefile variable LIBdir on the command-line. I encounter the following type of diagnostic error message:

. . .

```
mpiifort -o ../xslu_ipf_intelmpi_intel80_noopt
psludriver.o psluinfo.o psgetrrv.o psmatgen.o
pmatgeninc.o pslaschk.o pslafchk.o -
L/opt/intel/ict/3.2.0.035/mkl/10.0.3.020/tests/scalapack/
source/TESTING/../../lib/64 -
lmkl_scalapacktesting_intel80 -lmkl_scalapack -
lmkl_blacsF77init_intelmpi -lmkl_blacs_intelmpi -
lmkl_blacsF77init_intelmpi -lmkl_lapack -lmkl_ipf -lguide
-lpthread
ld: cannot find -lmkl_scalapacktesting_intel80
make[2]: *** [../xslu_ipf_intelmpi_intel80_noopt] Error 1
make[2]: Leaving directory
'/opt/intel/ict/3.2.0.035/mkl/10.0.3.020/tests/scalapack/
source/TESTING/LIN'
make[1]: *** [exe] Error 2
make[1]: Leaving directory
'/opt/intel/ict/3.2.0.035/mkl/10.0.3.020/tests/scalapack/
source/TESTING'
make: *** [lib64] Error 2
```

What am I doing wrong? 12-Jun-2008

For the makefile command-line, explicitly set the `LIBdir` variable to the path where the Intel® MKL resides. An example command-line *might be* something like:

```
gmake lib64 F=intel80
LIBdir=/opt/intel/ict/3.2.0.035/mkl/10.0.3.020/lib/64
```

Note that the make command above is applicable to Itanium® 2-based systems.

Back to Top

For Microsoft Windows CCS*, I am encountering the following link error message:

```
LINK : fatal error LNK1181: cannot open input file
'bufferoverflowu.lib'
```

when creating executables for ScaLAPACK and/or Cluster DFT. What am I doing wrong? 12-May-2008

Please source the `.bat` file:

```
vcvarsx86_amd64.bat
```

in your DOS command-line window where you are doing the Intel® Math Kernel Library experiments. This .bat file should be located in a `bin` subfolder within the Microsoft* Visual Studio* folder path and the DOS command for sourcing this file might look something like the following:

```
"C:\Program Files (x86)\Microsoft Visual Studio
8\VC\bin\x86_amd64\vcvarsx86_amd64.bat"
```

where the line above is contiguous.

Back to Top

Intel® Trace Analyzer and Collector 07-Oct-2008

I am trying to use the Intel® Trace Collector that comes with the Intel® Cluster Toolkit with an MPI library other than Intel® MPI Library. I am seeing a link error, where there are unresolved symbols. What am I doing wrong? 07-Oct-2008

At the present time the Intel® Trace Collector that comes with the Intel® Cluster Toolkit supports Intel® MPI Library and MPICH where the default library setting is Intel® MPI Library. To set the Intel® Trace Collector environment variables to support tracing of MPI applications that have been built with MPICH, use the following source command:

```
. .../itac/bin/itacvars.sh mpich
```

for Bourne Shell and for C Shell use:

```
source .../itac/bin/itacvars.csh mpich
```

where `...` is the full installation path to where the `itac` folder resides on your system.

To reference Intel® MPI Library, you can replace the word `mpich` with the value `mpi2` or `mpi3` when sourcing `itacvars.sh` or `itacvars.csh` for the respective shell command-line examples above.

Back to Top

For the Linux* operating system, when I instrument the application `test.cpp` with the Intel® MPI Library compilation driver `mpiicpc` and with the Intel® Trace Collector using the following command line:

```
mpiicpc test.cpp -g -L${VT_LIB_DIR} -lVT ${VT_ADD_LIBS} -o
testcpp_inst
```

I do not produce a trace file. What am I doing wrong? 20-Feb-2007

When compiling the C++ application `test.cpp` with the Intel® C++ Compiler (`mpiicpc`) and with Intel® Trace Collector requires an additional library called `-lmpiic`. The rule for this is that if one uses the MPI-2 C++ application programming interfaces (APIs) with the Intel® C++ Compiler, then Intel® Trace Collector cannot intercept the MPI calls as they are written. They have to be mapped to C function calls first, with the help of an MPI implementation-specific wrapper library, which with respect to the command-line, has to be placed in front of the Intel® Trace Collector library. The name of that wrapper library for the `mpiicpc` compilation driver is `-lmpiic`. Thus, for IA-32, Intel® 64, and Itanium® 2-base architectures the compilation protocol should be:

```
mpiicpc test.cpp -g -lmpiic -L${VT_LIB_DIR} -lVT
${VT_ADD_LIBS} -o testcpp_inst
```

Back to Top

For a Linux* cluster, I am trying to use my keyboard to generate chart displays for Intel® Trace Analyzer 6.0 or greater and nothing appears. For example, I pressed the keys Ctrl+Alt+E on the left side of the keyboard which is suppose to generate the Event Timeline chart. What am I doing wrong? 30-Nov-2006

The keys may have been remapped by the various software interfaces on your system. If the Ctrl, Alt, and E keys on the left side of the keyboard do not generate an Event Timeline chart, please try using the Ctrl, Alt, keys on the right side of the keyboard along with the E key.

Back to Top

For the Linux* operating system, when I run Intel® Trace Collector on my Linux* cluster I am seeing something like the following diagnostic:

. . .

```
[0] Intel Trace Collector INFO: Writing tracefile test.stf
in /tmp/test
[0] Intel Trace Collector WARNING: minimum clock increment
0.010000s is very high, please fix system setup to obtain
better traces (see ITC FAQ)
```

. . .

What am I doing wrong? 30-Mar-2006

One can work around this problem by enabling the TSC timer. This is done by specifying the "tsc" option at boot time. One can view this problem description from the Red Hat Knowledgebase* at the URL: http://kbase.redhat.com/faq/FAQ_79_3728.shtm

Back to Top

For the Linux* operating system, when compiling and linking an MPI application on a cluster based on an Intel® 64 (formerly Intel® EM64T) architecture a message similar to the following is appearing when attempting to link with the Intel® Trace Collector Library (i.e. `-lVT`):

```
/usr/bin/ld: skipping incompatible /opt/intel/ict/3.2.0.035/itac/lib/libVT.a
```

What am I doing wrong? 12-Jun-2008

The problem may be due to explicitly sourcing either the Bourne shell script `mpivars.sh` or the C Shell script `mpivars.csh` through a path such as:

```
/opt/intel/ict/3.2.0.035/impi/bin
```

When the MPI application is compiled with one of the Intel MPI Library compilation drivers, a 32-bit version of object files are created. However, during the link step, the linkage editor detects an incompatibility between the 32-bit object files and the 64-bit Intel® Trace Collector Library (i.e., `/opt/intel/ict/3.2.0.035/itac/lib/libVT.a`) and hence the error

message. To correctly build 64-bit object files on Intel® 64 (formerly Intel® EM64T) architectures using the Intel® MPI Library compilation drivers, one needs to explicitly source either `mpivars.sh` or `mpivars.csh` in a directory path such as:

```
/opt/intel/ict/3.2.0.035/impi/bin64
```

Another solution is to source one of the Intel® Cluster Toolkit shell script files `ictvars.sh` or `ictvars.csh` that exist in a directory path such as:

```
/opt/intel/ict/3.2.0.035/
```

The Bourne Shell source command might be something like:

```
. /opt/intel/ict/3.2.0.035/ictvars.sh
```

The C Shell source command might be something like:

```
source /opt/intel/ict/3.2.0.035/ictvars.csh
```

Back to Top

Intel® MPI Benchmarks 07-Oct-2008

For the Linux* operating system, when I use the following `gmake` command for the Intel® MPI Benchmarks:

```
gmake -f make_ict
```

I get the following error message:

```
make_ict:5: *** mpiicc is not defined through the PATH
environment variable setting. Please try sourcing an
Intel(r) Cluster Tools script file such as "mpivars.[c]sh"
or "ictvars.[c]sh".  Stop.
```

What am I doing wrong? 30-Nov-2006

Use of the makefile `make_ict` requires that `mpivars.[c]sh` or `ictvars.[c]sh` must be sourced. If you invoke the command above without having sourced one of these shell script files, a diagnostic message which looks like the above will appear.

Back to Top

How do I compile the Intel® MPI Benchmarks on Microsoft Windows CCS*? 07-Oct-2008

Within the Intel® MPI Benchmarks 3.2 folder configuration, there is a `WINDOWS` folder that contains Microsoft* Visual Studio* 2005 and 2008 project folders which allow you to use a pre-existing "`.vcproj`" project file in conjunction with Microsoft* Visual Studio* to build and run the associated Intel® MPI Benchmark application.

Within Microsoft Windows Explorer and starting at the Windows folder, you can go to one of the subfolders `IMB-EXT_VS_2005`, `IMB-EXT_VS_2008`, `IMB-IO_VS_2005`, `IMB-IO_VS_2008`, `IMB-MPI1_VS_2005`, or `IMB-MPI1_VS_2008` and click on the corresponding "`.vcproj`" file and open up Microsoft Visual Studio*

Back to Top

Uninstalling the Intel® Cluster Toolkit 12-Jun-2008

For the Linux* operating system, does the Intel® Cluster Toolkit uninstaller script which is called `uninstall.sh` have any options? 12-Jun-2008

The answer is yes. Simply type a shell command referencing `uninstall.sh` such as:

```
uninstall.sh --help | less
```

and you will see a list of options that look something like:

```
NAME
    uninstall.sh - Uninstall Intel(R) Cluster Toolkit for
Linux* 3.2.

SYNOPSIS
    uninstall.sh [options]

OPTIONS
    --help  Print this help and exit.

    --log-file=FILE
            Write log to the specified file.

    --single-node
```

```
    --singlenode
            Uninstall the product only from this node.

    --delete-update=UPDATE_NUMBER
            Delete update with the specified number.

    --list-update
            List all updates.

COPYRIGHT
    Copyright 1999-2008, Intel Corporation. All Rights
Reserved.
```

Back to Top

How do I uninstall the Intel® Cluster Toolkit on Microsoft Windows CCS*? 12-May-2008

Follow the start menu on the head node of the Microsoft Windows CCS* cluster system and click on the Control Panel selection. Proceed to click on the Add or Remove icons panel, and locate the entry for Intel® Cluster Toolkit. Remove the package from the cluster. You can use the F5 key to do a refresh on the Add or Remove Programs panel.

Back to Top

Version Compatibility 12-Jun-2008

For the Linux* operating system, is Intel® Cluster Toolkit 2.0 compatible with Intel® Cluster Toolkit 1.0? 30 Nov-2006

In general, the Intel® Cluster Toolkit is built around the current release of Intel® MPI Library. Intel® MPI Library 1.0 is binary incompatible with Intel® MPI Library 2.0. Therefore it is best to not attempt to mix the software tools that were packaged with the two releases.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 2.0.1 compatible with Intel® Cluster Toolkit 1.0 or Intel® Cluster Toolkit 2.0? 30-Nov-2006

Intel® Cluster Toolkit 2.0.1 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit 2.0.1 is compatible with Intel® Cluster Toolkit 2.0.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 3.0 compatible with Intel® Cluster Toolkit 1.0, or Intel® Cluster Toolkit 2.X? 30-Nov-2006

Intel® Cluster Toolkit 3.0 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit 3.0 is compatible with Intel® Cluster Toolkit 2.X.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 3.0.1 compatible with Intel® Cluster Toolkit 1.0, Intel® Cluster Toolkit 2.X, or Intel® Cluster Toolkit 3.X? 20-Feb-2007

Intel® Cluster Toolkit 3.0.1 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit 3.0.1 is compatible with Intel® Cluster Toolkit 2.X and Intel® Cluster Toolkit 3.X.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 3.1 compatible with Intel® Cluster Toolkit 1.0, Intel® Cluster Toolkit 2.X, or Intel® Cluster Toolkit 3.X? 26-Jun-2007

Intel® Cluster Toolkit 3.1 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit 3.1 is compatible with Intel® Cluster Toolkit 2.X and Intel® Cluster Toolkit 3.X.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 3.1.1 compatible with Intel® Cluster Toolkit 1.0, Intel® Cluster Toolkit 2.X, or Intel® Cluster Toolkit 3.X? 29-Feb-2008

Intel® Cluster Toolkit 3.1.1 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit 3.1.1 is compatible with Intel® Cluster Toolkit 2.X and Intel® Cluster Toolkit 3.X.

Back to Top

For the Linux* operating system, is Intel® Cluster Toolkit 3.2 compatible with Intel® Cluster Toolkit 1.0, Intel® Cluster Toolkit 2.X, or Intel® Cluster Toolkit 3.X? 12-Jun-2008

Intel® Cluster Toolkit 3.2 is not compatible with Intel® Cluster Toolkit 1.0. However, Intel® Cluster Toolkit Compiler Edition 3.2 is compatible with Intel® Cluster Toolkit 2.X and Intel® Cluster Toolkit 3.X.

Back to Top