



Intel® Cluster Toolkit 3.2.2

Release Notes

(Revision 20091105)

Contents

- Overview
- Product Contents
- New Features
- Resolved Issues
- System Requirements
- Installation Notes
- Installation on Linux* OS
- Installation on Windows* OS
- Documentation
- Known Limitations
- Technical Support
- Disclaimer and Legal Information

Overview

Intel® Cluster Toolkit 3.2.2 for Linux* OS and Microsoft* Windows* Compute Cluster Server OS* (Microsoft Windows CCS*) accelerate parallel software development on homogenous cluster systems based on IA-32, or IA-64 (Itanium®), or Intel® 64 architectures. The Microsoft Windows CCS operating system only supports Intel® 64 architecture. In terms of the Intel® Cluster Toolkit software for Windows*, consider references within this document to Microsoft Windows CCS OS, and Microsoft* Windows* HPC Server 2008 OS **as interchangeable**. The Microsoft Windows CCS and Microsoft Windows HPC Server 2008 operating systems only support Intel® 64 architecture.

Intel Cluster Toolkit 3.2.2 supports application development using Intel® MPI Library with optimized parallel libraries, performance analysis, and benchmarks. Intel Cluster Toolkit 3.2.2 saves software developers time and improves execution performance on distributed computing systems.

Intel Cluster Toolkit 3.2.2 for Linux and Microsoft Windows CCS supports crucial parts of the message-passing interface (MPI) application development process including:

- Intel MPI Library 3.2 Update 2 which is approaching conformity with the Message Passing Interface 2 Standard (MPI-2), enables multiple interconnect solutions with a single implementation. Intel seeks to be a software leader in MPI and open standards.
- The Intel® Trace Analyzer and Collector 7.2 Update 2 (ITAC)
 - Intel® Trace Collector 7.2 Update 2 (ITC) provides event-based tracing in cluster applications through an instrumentation library that causes low-overhead in execution. The trace information provides performance data, statistics, multi-threaded events, and automatic instrumentation of user binaries on IA-32 architectures.
 - The Intel® Trace Analyzer 7.2 Update 2 (ITA) provides visual analysis of application activities gathered by the Intel Trace Collector. This software component has been completely rewritten.
 - A message checking component of the Intel Trace Collector provides a novel MPI correctness technology which detects errors with data types, buffers, communicators, point-to-point messages and collective operations, deadlocks, and data corruption.
- Application tuning with optimized mathematical library functions from Intel® Math Kernel Library 10.2 Update 2 (Intel® MKL 10.2 Update 2) that includes ScaLAPACK solvers and Cluster DFTs (Discrete Fourier Transforms).

- Intel® MPI Benchmarks 3.2 make it easy to gather performance information about a cluster system.

Note that when the system administrator or user completes the Intel Cluster Toolkit 3.2.2 installation process, there will be a file called `Doc_Index.htm` in the `doc` folder on the master node of the cluster. This file can be used as a documentation map to navigate to various information resources pertaining to the Intel Cluster Toolkit. Additional information about the exact location of `Doc_Index.htm`, and its content are further described in the [Installation Notes](#) section of this release notes document.

[Back to Top](#)

Product Contents

Below is a table of product components, including documentation file names for those components.

Tool	Version	Manual	Revision
Intel® MPI Library	3.2 Update 2	Doc_Index.html for Linux* Doc_Index.htm for Microsoft* Windows* CCS*	3.2 Update 2
Intel® Trace Analyzer and Collector	7.2 Update 2	ITA_Reference_Guide.pdf ITC_Reference_Guide.pdf	7.2 Update 2
Intel® Math Kernel Library	10.2 Update 2	mkl_documentation.htm	10.2 Update 2
Intel® MPI Benchmarks	3.2	IMB_Users_Guide.pdf	3.2

[Back to Top](#)

New Features

Intel® Cluster Toolkit 3.2.2

- With one installation session, the Intel Cluster Tools installer on Linux* OS or Windows* OS will install:
 - Intel® MPI Benchmarks 3.2
 - Intel® MKL 10.2 Update 2
 - Intel® MPI Library 3.2 Update 2
 - Intel® Trace Analyzer and Collector 7.2 Update 2
- In Section 4.2 of the Intel® Cluster Toolkit for Linux* OS Getting Started Guide, which is titled, "How to Set Up MPD Daemons on Linux*", there is a description on how to use the `mpdboot` option: `--parallel-startup`
- In Section 4.10 of the Intel® Cluster Toolkit for Linux* OS Getting Started Guide, which is titled, "Extended File I/O System Support on Linux* OS", there is a description on how to use the environment variables `I_MPI_EXTRA_FILESYSTEM` and `I_MPI_EXTRA_FILESYSTEM_LIST` in conjunction with the `mpiexec` and `mpirun` commands associated with Intel® MPI Library for Linux* OS
- In Section 7.3 of the Intel® Cluster Toolkit for Linux* OS Getting Started Guide, and Section 7.3 of the Intel® Cluster Toolkit for Microsoft* Windows* Compute Cluster Server OS Getting Started Guide, which are respectively titled, "Experimenting with the High Performance Linpack Benchmark", recipes are provided on how to build High Performance Linpack executables
 - SUSE Linux Enterprise Server* 11 is now supported
 - Please see the section of this document titled, "Installation on Linux* OS" for some possible installation issues and workarounds involving Linux OS on Itanium® architecture

Intel® MPI Library 3.2 Update 2 for Linux* OS

This release includes the following updates from the Intel® MPI Library 3.2 for Linux* OS (see product documentation for more details):

- Intel® Compiler 11.1 Update 3 Support

- Increased number of supported communicators (up to 16384 simultaneously created communicators. Prior to the Intel® MPI Library 3.2 Update 2 release, the limit was 1024 communicators)
- Up to 1.5 times out-of-the-box performance improvement for mpirun and mpdboot
- Up to 3 times faster startup through the mpdboot --parallel-startup option
- Abstract-Device Interface for I/O (ADIO) virtualization for file system integration (Linux). Up to ten times performance improvement in file I/O for:
 - Panasas* ActiveScale File System (PanFS)
 - Parallel Virtual File System*, Version 2 (Pvfs2) through the I_MPI_EXTRA_FILESYSTEM and I_MPI_EXTRA_FILESYSTEM_LIST environment variables

Intel® MPI Library 3.2 Update 2 for Windows* OS

The Intel® MPI Library 3.2 Update 2 for Windows* OS includes the following updates compared to the Intel® MPI Library 3.2 Update 1 for Windows* (see product documentation for more details):

- Intel® Compiler 11.1 Update 3 Support
- Increased number of supported communicators (up to 16384 simultaneously created communicators. Prior to the Intel® MPI Library 3.2 Update 2 release, the limit was 1024 communicators)

Intel® Math Kernel Library (MKL) 10.2 Update 2

New in Intel® MKL 10.2 Update 2 for Linux* OS and for Windows* OS

- Performance
 - BLAS level 1 and level 2 performance improvements:
 - 32-bit operating systems
 - Intel® Xeon® processor 5300 series
 - Level 1: CAXPY, (C,Z)DOT(U,C) improved by 25%, *SCAL improved by 5-10%, and ZCOPY improved by up to 3 times
 - Level 2: *GER improved by 30-40%, (C,Z)HER/(C,Z)HPR improved by 10-20%, and (C,Z)TBMV/(C,Z)TBSV improved by 10%
 - Intel® Xeon® processor 5400 series
 - Level 1: (C,Z)AXPY, (C,Z)DOT(U,C) improved by 20%, (C,Z)SCAL improved by 10%, and ZCOPY improved by up to 2 times
 - Level 2: *GER improved by 30-40%, (C,Z)HER/(C,Z)HPR improved by 10-20%, and (C,Z)TBMV/(C,Z)TBSV improved by 10%
 - Intel® Xeon® processor 5500 series
 - Level 1: (C,Z)AXPY, (C,Z)DOT(U,C) improved by 20%, *SCAL improved by 30%, ZCOPY improved by up to 3 times, and *ASUM improved by 15%
 - Level 2: *GER improved by 30-40%, (C,Z)HER/(C,Z)HPR improved by 10-20%, and (C,Z)TBMV/(C,Z)TBSV improved by 10%
 - 64-bit operating systems
 - Intel® Xeon® processor 5500 series
 - Level 1: (C,Z)AXPY, (C,Z)DOT(U,C) improved by 10%, and (S,SC)ASUM improved by 5%
 - Level 2: *GER, (C,Z)HER/(C,Z)HPR, (C,Z)TBMV/(C,Z)TBSV improved by 10%
 - Improved performance of out-of-place double-precision matrix arithmetic routines MKL_Domatcopy and MKL_Domatadd transpose function on Intel® Core™ i7 processors and Intel® Xeon® processor 5400 and 5500 series
 - LAPACK
 - Improved performance of parallel ?POTRF by up to 1.1 times for the upper case and 1.14 times for lower case on medium sizes (N approximately 1000)
 - Parallel versions of ?GEBRD, ?SYTRD and ?HETRD for bi- and tridiagonalizations are improved by up to 1.4 times on cache boundary and out-of-cache sizes.
 - Improved the parallel ?STEDC divide and conquer eigensolvers by up to 1.6 times
 - PARDISO OOC
 - Performance for symmetric positive definite matrices was greatly improved — up to 5 times on internal benchmarks
 - VSL
 - Improved the performance of double precision Sobol basic random number generator (BRNG) by up to 2.9 times for dimensions ≥ 16 on 45nm High-k Intel® Core™2 processors and Intel® Xeon® processor 5500 series
 - VML
 - Improved the v(s,d)Pow function by up to 2.6 times (35% improvement on average) on Intel® Core™2 processors, 45nm High-k Intel® Core™2 processors, and the Intel® Xeon® processor 5500 series
 - 20% and 40% performance improvement of v(s,d)Ceil/Trunc/Floor functions on 45nm High-k Intel® Core™2 processors and Intel® Xeon® processor 5500 series respectively.
 - 15% performance improvement of vsSin/Cos/SinCos functions on Intel® 64

- 35% performance improvement of vdSin/Cos/SinCos functions on Intel® Core™2 processor, 45nm High-k Intel® Core™2 processors, and Intel® Xeon® processor 5500 series
- o FFTs
 - Improved scalability of 1D, single precision, complex FFTs through better cache use and improved vectorization on Intel® Xeon® processors
 - 50% performance improvement of small 3D, M*K*L less than 32000, complex FFT on Xeon® processor 5500 series
 - Interface
 - o Moved Trust-Region (TR) Solver, Iterative Sparse Solver (ISS), GNU Multiple Precision (GMP) components from solver libraries to standard MKL libraries so now solver libraries are empty (for backward compatibility).
 - o Added dynamic support for the Intel MKL progress routine feature
 - o Added support for 64-bit integer parameters (ILP64) in FFTW MPI wrappers.
 - o Fortran interface was added to out-of-place double-precision matrix arithmetic routines MKL_Domatcopy and MKL_Domatadd transpose function
 - Fixed Issues
- o BLAS
 - Fixed DGEMM threading on Intel® Xeon® processor 5500 series for $N=4*K+(2 \text{ or } 3)$ threads, where $K=2,3,4...$
 - Fixed dispatching for the Intel® Pentium® 4 processor (default code was loaded)
 - Fixed Custom DLL builder for the 32-bit Windows* stdcall interface
- o LAPACK
 - Fixed user ability to redefine xerbla in LAPACK for dynamic libraries
 - Resolved intermittent issue with workspace size returned by LAPACK routines due to losing digits caused by integer (32bit) to float (24bit mantissa) conversion on very large problems. Issue is fixed by using ceiling rounding for such conversions.
- o PARDISO OOC
 - Extended the set of matrices which can be solved on IA32. The number of nonzero in the lower triangular matrix should be less than $2*10^9$.
- o VML
 - Improved accuracy of vdHypot function on Intel® Pentium® III and non-Intel processors (32-bit libraries)
 - Fixed bug in v(s,d)Modf/Rint/NearbyInt functions for +/-0 results.
 - Improved accuracy for the vdSqrt HA function to within .52 ulps
 - Fixed bug in low accuracy (LA) v(s,d)Inv and v(s,d)Div functions for denormalized results
 - Fixed bug in v(z,c)Mul/Div/MulByConj functions for in-place data calls (input and output arrays are the same)
 - Improved accuracy for v(s,d)CdfNorm and fixed bugs in special values behavior
 - Improved accuracy of LA version of the vsPow2o3 function
 - o Improved accuracy and fixed bugs in special value behavior for the v(s,d)Powx functions

Intel® Trace Analyzer and Collector 7.2 Update 2

New in Intel® Trace Analyzer and Collector 7.2 Update 2

- Intel® Compiler 11.1 Update 3 support
- Intel® Trace Analyzer (ITA) adaption to Qt4 behavior

[Back to Top](#)

Resolved Issues

Bug Number	Title of Bug Report
802	[DOC] broken links in documentation
891	[INS] License Duplication in folder /opt/intel/licenses on Linux
905	[INS] IMB install failure after rejecting cluster components install
916	[SWC][FNC] No Intel Copyright notice on mpicl, mpifc
931	Add note to install section of ICT Release Notes
935	Create article on how to install cluster tools on Debian
948	[DOC] Invalid hyperlinks in Doc Index
950	[DOC] Invalid columns formatting in "Getting started guide"
951	[DOC] Invalid links to MPI documentation in "Getting Started Guide"

[Back to Top](#)

System Requirements

Processor System Requirements

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

Note that it is assumed that the processors listed above are configured into homogeneous clusters.

[Back to Top](#)

Disk-Space Requirements

20 GBs of disk space (minimum)

Note that during the installation process the installer may need up to 4 gigabytes of temporary disk storage to manage the intermediate installation files.

[Back to Top](#)

Operating System Requirements

OS Distributions	IA-32 Architecture	Intel® 64 Architecture		IA-64 Architecture
		32-Bit Applications	64-Bit Applications	
Red Hat Enterprise Linux* 4.0	S	S	S	S
Red Hat Enterprise Linux* 5.0	S	S	S	S
SUSE Linux Enterprise Server* 10	S	S	S	S
SUSE Linux Enterprise Server* 11	S	S	S	S
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](#)

Memory Requirements

2 GB of RAM (minimum)

[Back to Top](#)

Intel® Professional Edition C++ Compiler and Intel® Fortran Compilers

For all of the Intel® architectures, the version number on the Intel compilers should be 9.1 or greater.

[Back to Top](#)

Installation Notes

Setting Up Secure Shell on Linux* OS

Installation of the Intel Cluster Toolkit 3.2.2 assumes that the homogenous cluster has `ssh` connectivity.

1. Make sure `ssh` is installed on your computing cluster. To do this type the shell command:

```
which ssh
```

If it does not exist, you can acquire `ssh` from the URL:

<http://www.openssh.org>

2. Create an authentication key as follows:

```
ssh-keygen -t rsa
```

This will generate a private and public key pair. The private key should be saved in:

```
~/.ssh/identity
```

and the public key should be saved in:

```
~/.ssh/identity.pub
```

3. Authorize access by placing the contents of the public key into the `~/.ssh/authorized_keys` file. All keys listed in that file are allowed access. One way to do this is to issue the shell command:

```
cat ~/.ssh/identity.pub >> ~/.ssh/authorized_keys
```

If the computing system that you are connecting to does not share a common file system, then `~/.ssh/identity.pub` should be concatenated to the `~/.ssh/authorized_keys` file of the computing system you will be connecting to. Secure Shell will insist that the file `authorized_keys` has its permissions set so that it is not group writable, so do the following:

```
chmod go-rwx ~/.ssh/authorized_keys
```

This step avoids the need to enter your password each time you want to run a secure shell command.

4. In order to avoid typing in your pass phrase each time `ssh` is invoked, an `ssh-agent` needs to be created and your pass phrase added. This is done as follows:

```
ssh-agent $SHELL
```

```
ssh-add
```

[Back to Top](#)

Checklist in Case There Are Problems on Linux* OS

1. Make sure that the host names listed in the `machines.LINUX` file are also listed in the `/etc/ssh_known_hosts` file on your network or your `~/.ssh/known_hosts` file in your home directory.
2. It is important that `/tmp` has permissions set to 377, with root as owner and group 0.
3. `openssh` has a `-v` flag option which is very useful for tracking down handshaking problems.
4. If you encounter the following message:

```
"scp: FATAL: Executing ssh1 in compatibility mode failed (Check that scp1 is in your PATH). Lost connection."
```

Then there is an incompatibility with the secure shell protocol between the master node and at least one of the other nodes listed in the `machines.LINUX` file. The solution is to install compatible versions of secure shell on all nodes of the cluster.

[Back to Top](#)

Alternative Solution to Establishing Secure Shell Connectivity on Linux* OS

Within the "tar" package of the Intel Cluster Toolkit, there is an `expect` shell script file called `sshconnectivity.exp`. "expect" is a tool for automating interactive applications. To run "sshconnectivity.exp", the `expect` runtime software needs to be installed on your Linux system. To make sure that the `expect` runtime software is properly installed, type:

```
which expect
```

If you encounter a "Command not found." error message, you can download the `expect` software package from the following URL:

<http://expect.nist.gov/>

The `expect` shell script file called "sshconnectivity.exp" can be used to help you establish secure shell connectivity on a cluster system. The syntax for the command is:

```
./sshconnectivity.exp machines.LINUX
```

This `expect` shell script will create or update a `~/ .ssh` directory on each node of the cluster beginning with the master node which must be the first name listed in the `machines.LINUX` file. This script will prompt you for your password twice.

```
Enter your user password:
Re-enter your user password:
```

To provide security each time you enter your user password, asterisks will appear in lieu of the password text. Upon successful completion of the script, the following message fragment will appear:

```
Node count = 4
Secure shell connectivity was established on all nodes.
```

A log of the transactions for this script will be recorded in:

```
/tmp/sshconnectivity.<login-name>.log
```

where `<login-name>` is a meta-symbol for your actual login.

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. If this is not done, a password prompt will continue to be issued for any secure shell activity.

[1]

This process is demonstrated by the following complete graph (Figure 1) illustration where a vertex in the graph represents a cluster computing node, and an edge between two vertices connotes that the two cluster computing nodes have exchanged public keys for secure shell connectivity. Secure shell connectivity is intended to provide secure, encrypted communication channels between two or more cluster nodes over an insecure network.

The script `sshconnectivity.exp` will call the appropriate secure shell utilities to generate a private key and a public key for each node of the cluster.

For the complete graph example in Figure 1, suppose there are nodes (vertices) 1 to n in the cluster. For a given node i , nodes 1 to $i - 1$ and nodes $i + 1$ to n are provided with the public key from node i . The user's public keys for a given node will be stored in the `~/ .ssh` folder associated with the user's home directory for that computing node. Since there are $n - 1$ edges to a given node, that node will have $n - 1$ public keys in the `~/ .ssh` folder that were provided by the other $n - 1$ nodes in the cluster. The example in Figure 1 represents a computing cluster that has at total of 5 nodes. The edges connecting a node indicate that that node has received 4 public keys from the remaining computing nodes. Also looking out from a given node indicates that the given node has provided its own public key to the remaining nodes that are reachable via the 4 edge paths.

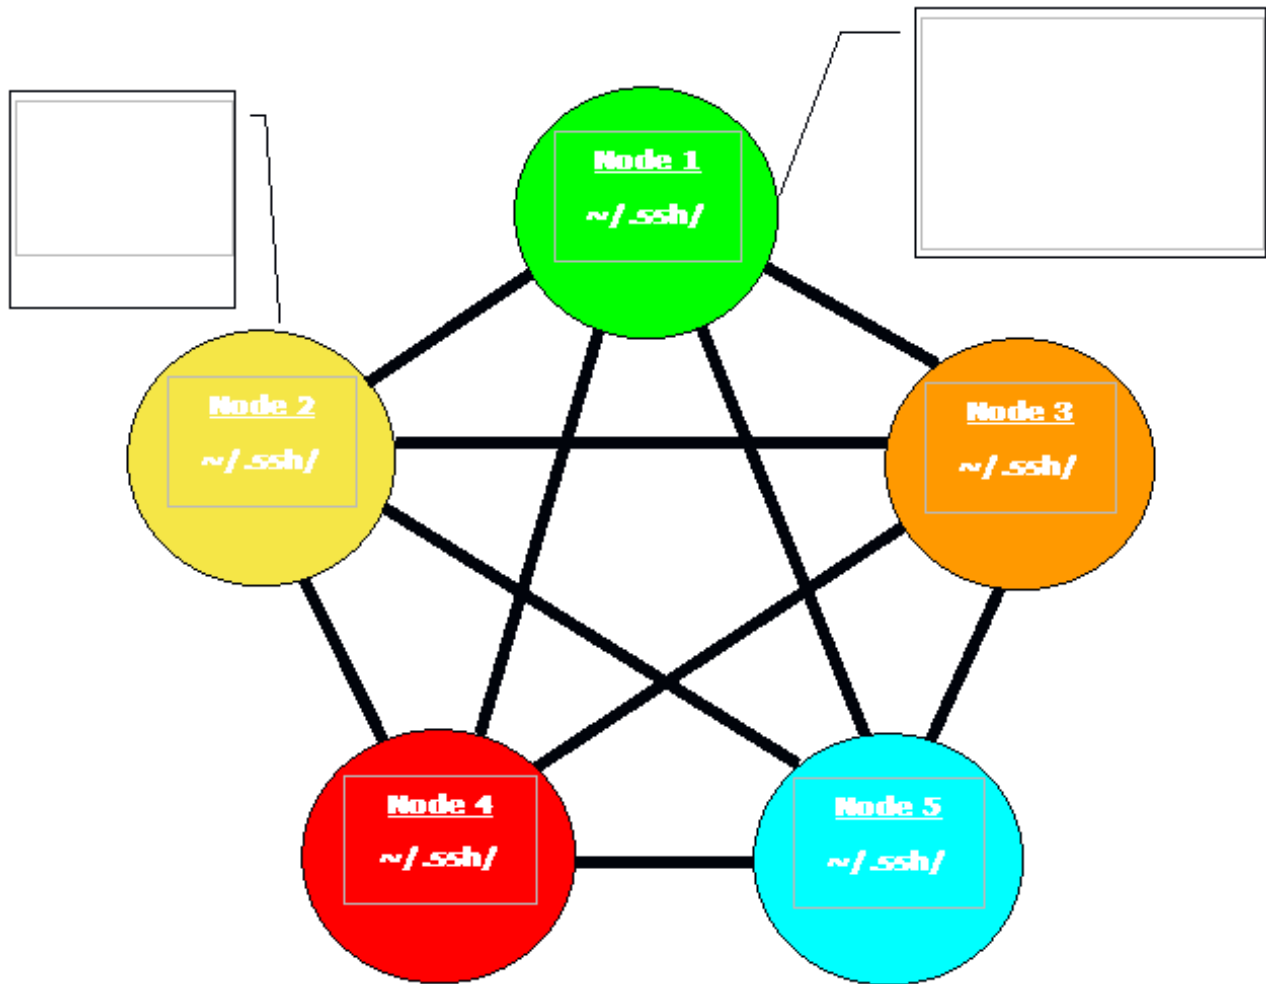


Figure 1 – Illustration of Secure Shell Connectivity for a Computing Cluster

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. If this is not done, a password prompt will continue to be issued for any secure shell activity.

[Back to Top](#)

The Intel Cluster Toolkit installation process on Linux is comprised of eight basic steps. The installation process on Microsoft Windows CCS alternatively uses the Microsoft installer model. The Intel Cluster Toolkit 3.2.2 package for Linux and Microsoft Windows CCS consist of the following components:

Software Component	Default Installation Directory on Linux	Default Installation Directory on Windows
Intel Math Kernel Library 10.2	/opt/intel/ict/3.2.2.0x/mkl where "0xx" is a sequence of three digits denoting the build number.	c:\program files (x86)\intel\ict\3.2.2.0xx\mkl where "0xx" is a sequence of three digits denoting the build number.

Intel MPI Library 3.2 Update 2	/opt/intel/ict/3.2.2.0xx/impi where "0xx" is a sequence of three digits denoting the build number.	c:\program files (x86)\intel\ict\3.2.2.0xx\mpi where "0xx" is a sequence of three digits denoting the build number.
Intel MPI Benchmarks 3.2	/opt/intel/ict/3.2.2.0xx/imb where "0xx" is a sequence of three digits denoting the build number.	c:\program files (x86)\intel\ict\3.2.2.0xx\imb where "0xx" is a sequence of three digits denoting the build number.
Intel Trace Analyzer and Collector 7.2 Update 2	/opt/intel/ict/3.2.2.0xx/itac where "0xx" is a sequence of three digits denoting the build number.	c:\program files (x86)\intel\ict\3.2.2.0xx\itac where "0xx" is a sequence of three digits denoting the build number.

Note that the Intel Cluster Toolkit installer will automatically make the appropriate selection of binaries, scripts, and text files from its installation archive based on the Intel processor architecture of the host system where the installation process is initiated. You do not have to worry about selecting the correct software component names for the given Intel® architecture. Also note that the values such as 035 in 3.2.2.035 and 037 in 3.2.2.037 are *example build number values* and may vary with your actual installation.

Recall that you as a user of the Intel Cluster Toolkit on Linux, Microsoft Windows CCS, or both may need assistance from your system administrator in installing the associated software packages on your cluster system, if the installation directory requires system administrative write privileges (e.g. /opt/intel on Linux). This assumes that your login account does not have administrative capabilities.

[Back to Top](#)

Installation on Linux* OS

Here are some important installation notes that should be read before beginning an installation on Linux OS: The 4.2.2 version of RPM on Red Hat Enterprise Linux* 3.0 OS for Itanium 2 architecture has a broken relocation feature. This will be a serious problem for users trying to install 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>

Secondly, on Itanium® architecture for some OSes (for example, Red Hat Enterprise Linux 4.0 Update 7 OS, and Red Hat Enterprise Linux 5.0 Update 2 OS), if the Intel® Cluster Toolkit installer is installing the Intel® Math Kernel Library using the RPM registration process, the cluster toolkit installer will automatically revert to the non-RPM mode (i.e., use the installer command-line option `--nonrpm`) to complete the install of the Intel® Math Kernel Library. In these situations an informational message like the following will appear in the installer log file and on the user's screen during installation:

The RPM package management Utility on this system does not allow installation of non-native software packages at non-default locations. Installation of Intel Math Kernel Library (MKL) will continue without using the RPM registration database.

The above symptom does not seem to appear on SLES 10 OS and SLES 11 OS for Itanium architecture.

To begin installation on Linux*:

1. A `machines.LINUX` file will either need to be created, or an existing `machines.LINUX` file can be used by the Intel Cluster Toolkit installer to deploy the appropriate software packages from the toolkit amongst the nodes of the cluster. This `machines.LINUX` file contains a list of the computing nodes (i.e. the hostnames) for the cluster. The format is one hostname per line:

`hostname`

The hostname should be the same as the result from the Linux* command "hostname". An example of the content for the file `machines.LINUX`, where a contrived cluster consists of eight nodes might be:

`clusternode1`

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

A line of text above is considered a comment line if column 1 contains the "#" symbol. It is always assumed that the first node in the list is the master node. The remaining nodes are the compute nodes. The text `clusternode1` and `clusternode2`, for example, represent the names of two of the nodes in a contrived computing cluster. The contents of the `machines.LINUX` file can also be used by users to construct an `mpd.hosts` file for the multi-purpose daemon (MPD) protocol. The MPD protocol is used for running MPI applications that utilize Intel MPI Library.

2. In preparation for doing the installation, the user may want to create a staging area. On the system where the Intel Cluster Toolkit software components are to be installed, it is recommended that a staging area be constructed in a directory such as `/tmp`. An example folder path staging area might be:

```
/tmp/ict_staging_area
```

where `ict_staging_area` is an acronym for Intel Cluster Toolkit staging area.

3. Upon registering for Intel Cluster Toolkit 3.2.2, you will receive a serial number (e.g., C111-12345678) for this product. Your serial number can be found within the email receipt of your product purchase. Go to the [Intel Registration Center](#) site and provide the product serial number information. Once admission has been granted into the registration center, a user will be able to access the Intel® Premier web pages for software support.

4. The license for the Intel Cluster Toolkit license file that is provided to the user should be placed in a location pointed to by the `INTEL_LICENSE_FILE` environment variable. Do not change the file name as the ".lic" extension is critical. Common locations for the attached license file are:

```
<installation_path>/licenses
```

For example, on the cluster system where the Intel Cluster Toolkit software is to be installed, all licenses for Intel-based software products might be placed in:

```
/opt/intel/licenses
```

It is also imperative that the user and/or the system administrator set the environment variable `INTEL_LICENSE_FILE` to the directory path where the Intel software licenses will reside *prior* to doing an installation of the Intel Cluster Toolkit. For Bourne* Shell or Korn* Shell the syntax for setting the `INTEL_LICENSE_FILE` environment variable might be:

```
export INTEL_LICENSE_FILE=/opt/intel/licenses
```

For C Shell, the syntax might be:

```
setenv INTEL_LICENSE_FILE /opt/intel/licenses
```

5. Patrons can place the Intel Cluster Toolkit software package into the staging area folder.

6. The installer package for the Intel Cluster Toolkit has the following general nomenclature:

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

where `<major>.<minor>.<update>.<package_num>` is a string such as:

`b_3.2.2.xxx`, where `b` is an acronym for beta

or

p_3.2.2.xxx, where p is an acronym for production

The *package_num* meta-symbol is a string such as 035. This string indicates the package number.

The command:

```
tar -xvzf l_ict_<major>.<minor>.<update>.<package_num>.tar.gz
```

will create a subdirectory called l_ict_<major>.<minor>.<update>.<package_num>. Change to that directory with the shell command:

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

For example, suppose the installation package is called l_ict_p_3.2.2.035.tar.gz. In the staging area that has been created, type the command:

```
tar -xvzf l_ict_p_3.2.2.035.tar.gz
```

This will create a subdirectory called l_ict_p_3.2.2.035. Change to that directory with the shell command:

```
cd l_ict_p_3.2.2.035
```

7. As mentioned in the release notes section titled, "[Alternative Solution to Establishing Secure Shell Connectivity](#)", the expect shell script file called "sshconnectivity.exp" can be used to help the user establish secure shell connectivity on a cluster system. "sshconnectivity.exp" is also located within the l_ict_<version>.<release> directory staging area. The syntax for the "sshconnectivity.exp" command is:

```
./sshconnectivity.exp machines.LINUX
```

This expect shell script will create or update a ~/.ssh directory on each node of the cluster beginning with the master node which must be the first name listed in the machines.LINUX file. This script will prompt the user for a cluster password twice.

```
Enter your cluster password:
Re-enter your cluster password:
```

Each time the user enters the cluster password, asterisks will appear in lieu of the password text so as to provide security. Upon successful completion of the script, the following message fragment will appear:

```
...
Node count = 4
Secure shell connectivity was established on all nodes.
...
```

A log of the transactions for this script will be recorded in:

```
/tmp/sshconnectivity.<login-name>.log
```

where <login-name> is a meta-symbol for the user's actual login.

8. Once secure shell connectivity is established, type a variation of the install command as demonstrated by the table below and follow the prompts issued by this install script.

Installation command	Is root password required initially?	Installer prompts to be aware of	Default installation area
./install.sh	Yes		/opt/intel/ict/...

./install.sh --nonroot	No	<p>We recommend that you install the software using RPM (option 1). This will require root password. If you do not have root password, you can do a local installation in your home folder by choosing option 2 below.</p> <p>Which of the following would you like to do ?</p> <ol style="list-style-type: none"> 1. Install the software using RPM (root password required) - Recommended. 2. Install the software without using RPM database (root password not required). <p>x. Exit</p> <p>Please make a selection :</p>	./intel/ict/... in your home directory
./install.sh --nonrpm	Yes	<p>If you do not have the root password, you can do a local installation in your home folder by choosing option 2 below.</p> <p>Which of the following would you like to do?</p> <ol style="list-style-type: none"> 1. Install the software using RPM (root password required) - Recommended. 2. Install the software without using RPM database (root password not required). <p>x. Exit.</p> <p>Please make a selection: (1/2/x) [2]</p>	./intel/ict/... in your home directory, if option 2 is selected
./install.sh --nonroot --nonrpm	No		./intel/ict/... in your home directory

Note that the Intel MPI Benchmarks are only installed on the master node.

By default, the global root directory for the installation of the Intel Cluster Toolkit is:

`/opt/intel/ict/<major>.<minor>.<update>.<package_num>`

where `<major>`, `<minor>`, `<update>`, and `<package_num>` are integers. An example would be 3.2.2.035.

Within the folder path `/opt/intel/ict/<major>.<minor>.<update>.<package_num>` you will find the text files:

```
ictvars.csh
```

```
ictvars.sh
```

and

```
ictsupport.txt
```

If you are using Bourne Shell or Korn Shell for the login session, you should type:

```
./ictvars.sh
```

and for a login session that uses C Shell, you should type:

```
source ./ictvars.csh
```

The file called:

```
ictsupport.txt
```

contains the Package ID and Package Contents information. Please use the information in `ictsupport.txt` when submitting customer support requests.

For the default installation path, an index file, an FAQ file, and the user's guide are located in the directory path:

```
/opt/intel/ict/<major>.<minor>.<update>.<package_num>/doc
```

where as mentioned above, `<major>`, `<minor>`, `<update>`, and `<package_num>` are integers. A complete default folder path to the documentation directory might be:

```
/opt/intel/ict/3.2.2.035/doc
```

The name of the index file is:

```
Doc_Index.htm
```

The index file can be used to navigate to the FAQ, the release notes, the user's guide, and an internet accessible Intel Cluster Toolkit Tutorial. The index file will also provide links to Intel Trace Analyzer and Collector documentation, Intel MPI Library documentation, Intel MKL documentation, and Intel MPI Benchmarks documentation.

The name of the FAQ file is:

```
HelpMe_FAQ.htm
```

The name of the user's guide file is:

```
Getting_Started.htm
```

By default, the local version of the release notes is located in the directory path:

```
/opt/intel/ict/<major>, <minor>, <update>, and <package_num>/release_notes
```

The name of the release notes file is:

```
Release_Notes.htm
```

With the default installation settings, the Intel® Cluster Toolkit support file (`ictsupport.txt`) is installed at:

[Back to Top](#)

Installation on Windows* OS

The installer package for the Intel Cluster Toolkit has the following general nomenclature:

w_ict_<major>, <minor>, <update>, and <package_num>.exe

where <major>, <minor>, <update>, and <package_num> is a string such as:

b_3.2.2.xxx, where b is an acronym for beta

or

p_3.2.2.xxx, where p is an acronym for production

The <package_num> meta-symbol is a string such as 037. This string indicates the package number.

To install Intel® Cluster Toolkit on Windows simply go to the staging area where "w_ict_p_3.2.2.xxx.exe" is located, where "xxx" is a sequence of three digits (i.e., the <package_num>), and do the following:

1. Double click on the ".exe" file
2. Follow the instructions that are provided by the installation wizard. The installation wizard will install the components of the Intel® Cluster Toolkit on all nodes of the cluster.

[Back to Top](#)

Documentation

Below is a list of documents related to the Intel® Cluster Toolkit.

Document Name	File Name	Type of Information in the Document	Comment
Intel® Cluster Toolkit Release Notes	Release_Notes.htm	Information about this release and how to use the library with various compilers. Includes: <ul style="list-style-type: none"> • Overview • System Requirements • Setting up Secure Shell on Linux • Installation • FAQ – Frequently Asked Questions • Technical Support and Feedback • Copyright and Legal Information 	For details about the Intel® Cluster Toolkit, please see the <i>Intel® Cluster Toolkit Getting Started Guide 3.2.2.</i>
Intel® Cluster Toolkit Getting Started Guide	Getting_Started.htm	Contains detailed information about: <ul style="list-style-type: none"> • Intel® MKL 10.2 Update 2 • Intel® MPI Benchmarks 3.2 • Intel® MPI Library 3.2 Update 2 • Intel® Trace Analyzer and Collector 7.2 Update 2 	For implementation-specific information about the toolkit, please see the <i>Intel® Cluster Toolkit Release Notes.</i>

Intel® Cluster Toolkit FAQ	HelpMe_FAQ.htm	Contains a list of frequently asked questions (FAQ) regarding the Intel® Cluster Toolkit 3.2.2.	This file may contain helpful hints when using the Intel® Cluster Toolkit 3.2.2.
Intel® Cluster Toolkit Tutorial	Note that this link is not relevant to an Intel® Cluster Tools Beta program that is in preparation for a major release.	Contains detailed information about: <ul style="list-style-type: none"> • Intel® MKL 10.2 Update 2 • Intel® MPI Benchmarks 3.2 • Intel® MPI Library 3.2 Update 2 • Intel® Trace Analyzer and Collector 7.2 Update 2 	This is a web-based version of the Intel® Cluster Toolkit Getting Started Guide. It may contain information that is more recent than that of the local version of the Intel® Cluster Toolkit Getting Started Guide that has been installed on the user's computing cluster. The tutorial and the user's guide will have a revision number located in the title page such as "(Revision 20091105)". The larger the number, the more current the document revision.

[Back to Top](#)

Known Limitations

1. On Itanium® 2 processor-based systems, the following symptom may occur for RPM version 4.2.2 when doing an install with the `--nonroot` option:

```
install.sh -nonroot
```

the following symptom appears:

```
Installing Intel(R) MPI Library for Linux* version p_3.2.2.038...
Extracting files... ##### [100%]
```

Would you like to:

Install:

i. Intel(R) MPI Library, Development Kit for Linux* version 3.2.2-038

x. Exit

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

```
Cannot install Intel(R) MPI Library, Development Kit for Linux* version 3.2-038 to "/home/tooluser/intel/impi/3.2.2":
```

```
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.
```

XX2svw10/MPI.ini" file; the product cannot be installed to all cluster nodes.
Press Enter to continue...

The solution is to use an updated version of RPM.

[Back to Top](#)

2. There have been situations where during the installation process, /tmp has filled up. We recommend that you have at least one gigabyte 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 intermediate storage to another disk partition referenced by <PATH>.

[Back to Top](#)

3. If during installation you see a diagnostic which looks something like:

...

```
Cannot install Intel(R) Trace Analyzer And Collector version 7.2.2p-029 into directory
"/usr/home/user01/intel/itac/7.2.2":
The prefix directory path "/usr/home/user01/intel/itac/7.2.2" is owned by cookie intel-
ta_ipf version 7.2.2p-029.
Installation failed.
```

...

try the following corrective action when the installer completes:

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

```
/usr/home/user01/intel/ict/3.2.2.035/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 7.2.2p-029
2. Intel(R) Trace Analyzer version 7.2.2p-029
- a. All of the above.

```
Uninstall:
```

3. intel-ta_ipf version 7.2.2p-029
(non-rpm in "/usr/home/user01/intel/itac/7.2.2")
4. intel-tac_ipf version 7.2.2p-029
(non-rpm in "/usr/home/user01/intel/itac/7.2.2")
- 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/itac/7.2.2" 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/itac/7.2.2/uninstall.
sh".
```

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

Would you like to:

Install:

i. Intel(R) MPI Library, Development Kit for Linux* version 3.2.2p-038

Uninstall:

1. intel-mpi-ipf version 3.2.2p-038
(non-rpm in "/usr/home/user01/intel/ict/3.2.2.035/impi/")

x. Exit

Please type a selection (i/l/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/3.2.2.035/impi/
uninstall".
```

Completed cluster installation successfully.

This "Installation failed" symptom occurs when a prior installation of the Intel® Cluster Toolkit into a directory such as /usr/home/user01/intel/ict/3.2.2.035 was removed without having issued the uninstall.sh command that for this example, resides in the directory /usr/home/user01/intel/ict/3.2.2.035.

[Back to Top](#)

4. If when installing Intel® Cluster Toolkit 3.2.2, you see a diagnostic message regarding Intel® Math Kernel Library 10.2 Update 2 that looks something like the following:

```
Installing Intel(R) Math Kernel Library 10.2 for Linux* version p_10.2.2.020...
```

```
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 ]:
```

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 10.2. This message has also been generated when installing Intel® Math Kernel Library 10.2 on SUSE* 9.0 and SUSE* 9.1, where the message can also be ignored by typing "Yes".

[Back to Top](#)

5. If during installation, you encounter the error message, 'The prefix directory path "/usr/home/user01/intel/itac/7.2.2" is owned by cookie intel-ta_ipf version 7.2.2p-029.', you can resolve this problem by editing or removing the files /opt/intel_sdp_products.db and /root/intel/intel_sdp_products.db.

[Back to Top](#)

Technical Support

Your feedback is very important to us. To receive technical support for the tools provided in this product and technical information including FAQ's and product updates, you need to register for an Intel Premier Support account at the [Registration Center](#).

NOTE: Registering for support varies for release product or pre-release products (alpha, beta, etc) - only released software products have support web pages at <http://software.intel.com/sites/support/>.

To register for an account, please visit the Intel® Registration Center web site at <http://www.intel.com/software/products/registrationcenter/index.htm>. If you have forgotten your password, please email a request to: quadsupport@mailbox.intel.com. Please do not email your technical issue to this email address.

The product support web site, <http://software.intel.com/en-us/articles/intel-cluster-toolkit-support-resources/>, provides top technical issues, [FAQs & Known Issues](#), [Documentation and Training](#), and product errata.

[Back to Top](#)

Submitting Issues

To submit an issue via the Intel Premier Support website, please perform the following steps:

1. Ensure that Java* and JavaScript* are enabled in your browser.
2. Go to <https://premier.intel.com/>.
3. Type in your Login and Password. Both are case-sensitive.
4. Click the "Submit Issues" button in the left margin.
5. Read the Confidentiality Statement and click the "I Accept" button.
6. Click on the "Go" button next to the "Product" drop-down list.
7. Click on the "Submit Issue" link in the left navigation bar.
8. Choose "Development Environment (tools, SDV, EAP)" from the "Product Type" drop-down list.
9. If this is a software or license-related issue choose "Intel(R) Cluster Toolkit" from the "Product Name" drop-down list.
10. Enter your question and complete the fields in the web-page windows that follow to successfully submit the issue.

Please follow these guidelines when forming your problem report or product suggestion:

1. Describe your difficulty or suggestion. For problem reports, please be as specific as possible (e.g., including compiler and link command line options), so that we may reproduce the problem. Please include a small test case if possible.
2. Describe your system configuration information. Be sure to include specific information that may be applicable to your setup: operating system, name and version number of installed applications, and anything else that may be relevant to helping us address your concern.

[Back to Top](#)

Disclaimer and Legal Information

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, skool, 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 © 2004-2009, Intel Corporation. All rights reserved.

[Back to Top](#)

[1]

A mathematical definition of a complete graph in graph theory is a simple graph where an edge connects every pair of vertices. The complete graph on n vertices has n vertices and $n(n - 1)/2$ edges, and is denoted by K_n . Each vertex in the graph has degree $n - 1$. All complete graphs are their own cliques (a maximal complete graph). A graph of this type is maximally connected because the only vertex cut which disconnects the graph is the complete set of vertices.