

Intel® Enterprise Edition for Lustre* Software

Version 2.4.1.0
Release Notes

Copyright 2015, Intel Corporation.
World Wide Web: <http://www.intel.com>
December 22, 2015

Contents:

- [Product Overview](#)
 - [Intel® Enterprise Edition for Lustre* Software](#)
 - [Intel® Manager for Lustre* Software](#)
 - [Features](#)
 - [New Features and Improvements in this Release](#)
- [System Requirements and Installation Instructions](#)
- [Installing Updates to Intel® EE for Lustre* Software](#)
- [Known Issues](#)
- [Partner Installation Support](#)
- [Learn More](#)
- [Disclaimer and Legal Information](#)

Product Overview

Intel® Enterprise Edition for Lustre* Software

Intel® Enterprise Edition for Lustre* Software, when integrated with Linux, aggregates a range of storage hardware into a single Lustre* file system that is well-proven for delivering fast IO to applications across high-speed network fabrics, such as InfiniBand* and Ethernet.

Intel® EE for Lustre* Software is a global, single-namespace file system architecture that allows parallel access by many clients to all the data in the file system across many servers and storage devices. Designed to take advantage of the reliability features of enterprise-class storage hardware, Intel® EE for Lustre* Software supports availability features, including

redundant servers with storage failover. Metadata and data are stored on separate servers to allow each system to be optimized for the different workloads.

Intel® Manager for Lustre* Software

Intel® Manager for Lustre* software is purpose-built to simplify the deployment and management of Lustre* file systems. Intel® Manager for Lustre* Software reduces management complexity and costs, enabling storage administrators to exploit the performance and scalability of Lustre storage, and accelerate critical applications and work flows.

Intel® Manager for Lustre* Software greatly simplifies the creation and management of Lustre file systems, using either the graphical user interface (GUI) or a command line interface (CLI). The software dashboard lets you monitor one or more distributed Lustre file systems. Near real-time storage-monitoring lets you track Lustre file system usage, performance metrics, events, and errors at the Lustre level. Plug-ins provided by storage solution providers enable monitoring of hardware-level performance data, disk errors and faults, and other hardware-related information.

Features

GUI-based creation and management of Lustre file systems

The Intel® Manager for Lustre* Software provides a powerful, yet easy-to-use GUI that enables rapid creation of Lustre file systems. The GUI supports easy configuration for high availability and expansion, and enables performance monitoring and management of multiple Lustre file systems.

Graphical charts display near real-time performance metrics

Fully-configurable color charts display a variety of real-time performance metrics for single or multiple file systems, down to individual servers and targets, and reveal metrics including OST bandwidth, OST balance, read/write operations per OST (heat map), file system capacity, metadata operations, bandwidth, and various resource usage parameters.

Heat Map Charting with Job Statistics

The Heat Map, displayed on the Dashboard page, shows the level of read/write activity for each OST in all file systems. OSTs are displayed as rows, and consecutive time intervals are displayed as columns. One can monitor the level of read or write activity for a given OST over time. The job statistics feature displays the top five read and write job statistics for a given OST and time interval. This feature employs the jobstats feature in Lustre, and supports the creation of plug-ins to display user account, command line, job size, and job start/finish times.

Auto-configured high-availability clustering for server pairs

Pacemaker and Corosync are configured automatically when the system design follows configuration guidance. This removes the need for manually installing HA configuration files on storage servers and simplifies high-availability configuration.

PDU configuration and server outlet assignments support automatic failover

The PDU tab lets you configure and manage power distribution units. You can add a PDU and assign specific PDU outlets to specific servers. When you associate PDU failover outlets with servers using this tool, STONITH is automatically configured.

IPMI and BMC Configuration

As an alternative to PDU configuration, support for Intelligent Platform Management Interface and baseboard management controllers support server monitoring, high-availability configuration, and failover.

LNET Configuration

This feature assists in configuring LNET for a given server's network interface by setting the LNET network ID for that port. This feature requires a single LNET. You can configure multiple LNETs (i.e., with the use of routers), however in this release, additional LNETs cannot be configured from the GUI.

Support for Intel® Xeon Phi™ Coprocessor Clients

Intel® EE for Lustre* client software can be installed and configured to run on Intel® Xeon Phi™ Coprocessor clients. This means that the Intel® Xeon Phi™ Coprocessor clients can directly mount Lustre.

Simplified ISO-less installation and automated deployment streamlines installation

The installation strategy removes the need to manually install Intel® EE for Lustre* software on each server. Rather, Intel® EE for Lustre* software is quickly installed on the manager server, and from there required packages are automatically deployed to all storage servers. Storage servers and the manager server can run the same standard operating system as the rest of your estate.

Support for OpenZFS in Monitor Mode

Intel® EE for Lustre* software supports ZFS as a back-end file system replacement for ldiskfs. While Intel® Manager for Lustre* software is able to configure and manage high-availability Lustre storage solutions, Intel® EE for Lustre* software can discover and fully monitor ZFS file systems. Full, high-availability management of ZFS file systems is not supported.

Hierarchical Storage Management

This release of Intel® EE for Lustre* software incorporates support for hierarchical storage management (HSM). HSM provides tiered storage by allowing data to be moved between the file system and secondary, archival storage. The HSM framework included with this release of

Intel® EE for Lustre* software includes the POSIX copytool implementation, the worker agent, and the MDS Coordinator. HSM agents and copytool instances are configured via the Intel® Manager for Lustre* GUI.

Robinhood Policy Engine

The Robinhood policy engine, version 2.5.5, is included with Intel® EE for Lustre*, and can be used to manage policies for an HSM solution. Intel® Manager for Lustre* software performs the provisioning of the Robinhood agent server, which is performed via the manager GUI. Robinhood can be used with the HSM capabilities described above to automate HSM archiving and report generation. Robinhood also can be used by administrators for gathering information about your Lustre file system, such as the ranked list of the oldest files or the ranked list of the oldest files.

Automated Provisioning of Custom Lustre Service Nodes

This feature allows users to create custom profiles for new Lustre client types and, based on a given profile, deploy and install custom code to provide new services. The HSM copytool (above) is deployed in this way.

Apache Hadoop* adapter software

Intel® EE for Lustre* software incorporates the Apache Hadoop* adapter software. This Hadoop adapter for Lustre is compatible with the Apache Hadoop software, versions 2.3 and 2.5. Hadoop software allows users who run MapReduce jobs to bypass storing data in HDFS, and store the MapReduce output directly to Lustre instead. This allows the analytical processes direct access to scientific output instead of transferring data from the compute cluster storage system to another file system. Optimizations have also been made to the shuffle step in MapReduce to take advantage of Lustre's high speed network access to data. Many workloads will see an overall reduction in end-to-end processing time by using the Hadoop adapter with the Intel® EE for Lustre* file system. For more information, see *Installing Hadoop, the Hadoop Adapter for Intel® EE for Lustre*, and the Job Scheduler Integration*.

The following table shows version compatibility with Apache distributions.

Note: The Hadoop adapter plugin for this release of Intel® EE of Lustre* software is:

hadoop-lustre-plugin-3.0.0.

Hadoop Adapter Version	Target Hadoop Distributions	Java Version (Recommended)
hadoop-lustre-plugin-3.0.0	hadoop-2.3.0-apache hadoop-2.5.0-apache	Java Runtime Environment 1.7.0+

HPC Job Scheduler integration with MapReduce

Intel® EE for Lustre* Software incorporates the HPC *job scheduler integration* with MapReduce, supporting the Apache 2.5.0 distribution of Hadoop. This adapter for job schedulers allows you to integrate common resource schedulers into your cluster. You have the choice of installing the SLURM (Simple Linux Utility for Resource Management) job scheduler integration or the PBS (portable batch system) job scheduler integration. An integration guide is available: *Installing Hadoop, the Hadoop Adapter for Intel® EE for Lustre*, and the Job Scheduler Integration*.

Hadoop commonly uses Yarn to manage MapReduce jobs. However, virtually all HPC systems use a *job scheduler*, for example, SLURM. However, having two job schedulers, e.g., SLURM and Yarn, in a single system can cause problems. The HPC Job Scheduler integration with MapReduce replaces YARN with an interface to the main resource manager for the system. This allows MapReduce applications to be run as normal HPC jobs. The following table shows version compatibility with Apache distributions and SLURM and PBS.

Note: The Hadoop HPC job scheduler for this release of Intel® EE of Lustre* software is:
`hadoop-hpc-scheduler-3.0.0`

HPC Scheduler Version	Target Hadoop Distributions	SLURM Version	Java Version (Recommended)
hadoop-hpc-scheduler-3.0.0	hadoop-2.5.0-apache	SLURM 2.5.6	Java Runtime Environment 1.7.0+

Apache Hive compatibility

Hive is a data warehouse infrastructure built on top of Hadoop for providing data summarization, query, and analysis. Intel has tested the Hadoop adapter for Lustre provided with Intel® EE for Lustre* software for compatibility with Apache Hive version 2.5.

Apache Hbase compatibility

HBase is a non-relational, distributed database modeled after Google's BigTable and written in Java*. Hbase runs on top of HDFS (Hadoop Distributed File System). Intel® has tested the Hadoop adapter for Lustre provided with Intel® EE for Lustre* software for compatibility with Apache Hbase version 2.5.

Distributed Namespace

Distributed Namespace (DNE) allows the Lustre metadata to be distributed across multiple metadata servers. Intel® EE for Lustre* software supports DNE1 (as of release 2.3.0.0), which supports the use of multiple MDTs. This enables the size of the Lustre namespace and

metadata throughput to be scaled with the number of OSSs. This feature is supported at the Intel® Manager for Lustre* GUI.

Differentiated Storage Services

Differentiated Storage Services (DSS) allows I/O data to be classified, sometimes referred to as “hinting”. These hints pass seamlessly through Intel® EE Edition for Lustre* software, at which point data can be tiered and intelligently cached by the storage system. This enables a more efficient use of cache space, and decreases the likelihood of critical data being evicted when the cache fills. Intel® is working directly with storage and cache vendors to enable DSS hinting in Lustre appliances, and to provide optimized performance to Intel® EE Edition for Lustre* deployments with a mix of SSD and traditional storage.

New Features and Improvements in this Release

This release of Intel® EE for Lustre* software incorporates the following new features and improvements.

Note: All features listed next were also supported with Intel® EE for Lustre software release 2.4.0.0, except for the change to Linux OS version support. This change is noted for release 2.4.1.0, below.

Linux Operating System Support

Intel® EE for Lustre* Software release 2.4.1.0 may be installed on servers and clients running RHEL 6.7 or CentOS 6.7, and on clients running RHEL 7.2 or CentOS 7.2.

SLES Server and Client Support

This software version supports the installation of Lustre on clients running SUSE Linux Enterprise version 11 with SP4 and version 12.

Note: For SLES installations, Intel® Manager for Lustre* software is not supported or installed.

Robinhood Policy Engine Updated

The Robinhood policy engine has been updated to version 2.5.5. See [Robinhood Policy Engine](#).

OpenZFS updated

OpenZFS updated to version 0.6.4.2.

Note-worthy Lustre Tickets Resolved

- Support added for RHEL 7.2 Lustre clients (LU-7451). Renaming files is now handled correctly for files upgraded from Lustre 1.8.x releases or file-level backups (LU-5626).
- Robustness of read/write has been improved for large files with multiple stripes (resolved in Intel® EE for Lustre* Software release 2.4.0.0).

- Performance of parallel reads is now improved (LDEV-6) (resolved in Intel® EE for Lustre* Software release 2.4.0.0).
- DirectIO performance while increasing stripes is now improved (LDEV-141) (resolved in Intel® EE for Lustre* Software release 2.4.0.0).

Software Requirements and Installation Instructions

For information on installing Intel® Enterprise Edition for Lustre* software and its components, including system requirements for the manager server and storage servers, please see the *Intel® Enterprise Edition for Lustre* Software Installation Guide*.

Linux Operating System Support

- Intel® EE for Lustre* Software release 2.4.1.0 may be installed on servers and clients running RHEL 6.7 or CentOS 6.7, and on clients running RHEL 7.2 or CentOS 7.2.

Note: Before using the Red Hat or RHEL software referenced herein, please refer to Red Hat's website for more information, including without limitation, information regarding the mitigation of potential security vulnerabilities in the Red Hat software.

SLES Operating System Support

Intel® EE for Lustre software may be installed on file system servers and clients running SUSE Linux Enterprise version 11 with SP4 and version 12. Note that for SLES, the Intel® Manager for Lustre* software and dashboard are not supported or installed. Automatic configuration or monitoring of high availability with the Intel® Manager for Lustre* dashboard is not supported.

Intel® Manager for Lustre* Software is supported on:

- Google Chrome browser. Use the most current stable version.
- Mozilla Foundation Firefox browser. Use the most current stable version.

The Intel® Manager for Lustre* Software GUI is not supported on Internet Explorer*.

Known Issues

This section lists those tickets that the partner and customer should be made aware of.

Updating Intel® EE for Lustre* Software

Note: If you plan to update the Linux operating system to 6.7 or 7.2 (client) update the operating system first. Then update the Intel® EE for Lustre* Software.

For instructions on installing this update of Intel® EE for Lustre* software on the manager server (and all file system managed servers), see the Intel® EE for Lustre* Partner Installation Guide.

Note: A dependency in the software upgrade process may require consecutive version updates. Please see the Intel® EE for Luster* Partner Installation Guide for complete update instructions.

Note: A known issue (listed here as HYD-3017) can result in a server being made unavailable. This can happen if the server has been added to a Lustre file system, (using Intel® Manager for Lustre* software) and then the user decides to *Force Remove* the server from the file system. The Force Remove command should only be performed if the Remove command has been unsuccessful. Force Remove will remove the server from the Intel® Manager for Lustre* configuration, but will not remove Intel® Manager for Lustre* software from the server. All targets that depend on the server will also be removed without any attempt to unconfigure them. To completely remove the Intel® Manager for Lustre* software from the server (allowing it to be added to another Lustre file system), first contact technical support.

Ticket Number	Description
HYD-1423	UI breaks when there are storage resources in the DB whose plugins are unavailable.
HYD-1500	File system client count freezes when the file system has stopped - so may show connected clients to a stopped file system.
HYD-1643	Changes in size of block devices are not detected.
HYD-1740	Removing an HA OSS removes HA OSTs.
HYD-1891	<p>Removed storage devices still appear in Intel® Manager for Lustre* GUI until the agent or manager is restarted.</p> <p>Recommendation: If you remove a storage device from a server, you'll need to restart the agent or Intel® Manager for Lustre* Software.</p>
HYD-2195	Logs filtering ends with pop-up message "Processing..." when date/time fields are malformed.
HYD-2443	<p>When the agent deployment fails, there should not be a green check.</p> <p>Recommendation: Servers that are partially configured may show a green check, which may imply that the server is fully configured. After adding a server, check that the command completed successfully. On the GUI, check the command section in the panel at left to be sure the server was properly</p>

	added.
HYD-2559	MGS tries to mount over and over again.
HYD-2696	<p>The Advanced Options dialog on the file system creation page may not be cleared if the file system creation is unsuccessful.</p> <p>Recommendation: Be sure to check the values in this dialog if creating a file system after an unsuccessful attempt.</p>
HYD-2877	Additional output from local_settings.py can break the invocation of real-time.py
HYD-3017	<p>Corosync reports from multiple hosts can cause hosts to be marked offline / online repeatedly.</p> <p>Recommendation: This issue can result in a server being made unavailable. This can happen if the server has been added to a Lustre file system, (using Intel® Manager for Lustre* software) and then the user <i>Force Removes</i> the server from the file system. The Force Remove command should only be performed if the Remove command has been unsuccessful. Force Remove will remove the server from the software manager configuration, but will not remove Intel® Manager for Lustre* software from the server. Targets that depend on that server will also be removed without any attempt to unconfigure them. To completely remove the Intel® Manager for Lustre* software from the server (allowing it to be added to another Lustre file system), first contact technical support.</p>
HYD-3069	<p>MD RAID Array disk discovery is broken.</p> <p>Recommendation: Software RAID is not supported.</p>
HYD-3117	<p>Copytool is orphaned when coordinator is shut down and restarted.</p> <p>Recommendation: At Configuration menu, click HSM. At the HSM configuration page, stop and then restart the affected copytool instance. To avoid this issue, stop each copytool prior to shutting down the coordinator.</p>
HYD-3232	<p>The software makes a best attempt to configure iptables on the manager. However, if your system has a drastically different format than the default el6 iptables configuration, this configuration may fail. This will surface as an error during install.</p> <p>Recommendation: Use a fresh el6 iptables configuration when installing Intel® Manager for Lustre* software, or repair the configuration by manually.</p>

HYD-3250	<p>IML will not force you to set up Power Control before proceeding on to creating a file system, but having a managed file system without Power Control is not a supported configuration.</p> <p>Recommendation: Follow the software's online Help when creating HA file systems. Assign PDU outlets to servers before creating the file system.</p>
HYD-3278	<p>The RabbitMQ messaging broker is crashing under heavy load on the API.</p> <p>Recommendation: Shutdown and restart Intel® Manager for Lustre* software.</p>
HYD-3280	<p>Corosync blocks 10.10.10.x LNET addresses. After "add server" or "create file system" ping stopped working the LNET interface. It seems corosync consumes the entire 10.x.x.x addr space.</p> <p>Recommendation: File system running Intel® EE for Lustre* software cannot be installed on Class A subnets (the /8 subnets, in CIDR notation). It is recommended that /16 or /24 subnets be used (Class B and C, respectively).</p>
HYD-3357	<p>User no longer alerted that they need to upgrade servers. Following an upgrade of Intel® EE for Lustre* software, the alert to upgrade file system servers is not displayed.</p> <p>Recommendation: After upgrading the manager server, open Intel® Manager for Lustre* software, access the Server page and click Upgrade Servers.</p>
HYD-3689	<p>Not all OSTs are added when requested, when filtering the list.</p> <p>Recommendation: Refresh the page and add your OSTs without filtering, or use filtering, understanding that only those shown in the current filter will be added.</p>
HYD-3811	<p>Failover that occurs while the manager software is down is not detected when the manager software comes back up. Restarting the agent resolves this issue.</p>
HYD-3825	<p>Configuring corosync can use a bound port.</p> <p>The original ticket confirmed that corosync could inadvertently assign a server's port number to a port that is already bound (the selection range was to a port number less than 1024). Resolution of this ticket consisted of selecting a port number within the range 32767-65535. This has been implemented in this release. Conflict with a bound port may still occur, but this is extremely unlikely.</p>

HYD-3937	<p>Scanning ZFS FS with suspended pool results in error.</p> <p>Recommendation: In cases such that one or more of the ZFS pools in your ZFS based file system are in an suspended, degraded, or offline state, the software may not detect the file system. Any attempt may result in the software displaying an error message. Make sure your ZFS pools and file system are functioning properly before attempting to scan.</p>
HYD-3944	<p>After upgrading the software, the UI [version] may not reflect upgrade.</p> <p>Recommendation: Refresh browser page to update version number.</p>
HYD-4101	<p>Running some ZFS commands on a Storage Server while Intel® Manager for Lustre software is active may cause ZFS and Intel Manager for Lustre to deadlock due to the following defect:</p> <p>https://github.com/zfsonlinux/zfs/issues/2034.</p> <p>Recommendation: Avoid running ZFS commands on Intel® Manager for Lustre* storage servers, while the storage server is running the Intel® Manager for Lustre* Agent.</p>
HYD-4217	<p>For ZFS file systems, the manager software detects primary and secondary servers incorrectly.</p> <p>Recommendation: When discovering a ZFS file system, the servers that the target(s) are currently mounted on will be identified as the primary servers by the manager software. Accordingly, be sure that the targets are mounted on the primary servers before running Detect File System from the manager GUI. Note that one can forget and then Detect the file system again, if necessary.</p>
HYD-4248	<p>On rare occasions the Manager GUI will sometimes display a 502 proxy error.</p> <p>Recommendation: Reload the page using the browser reload button. No data or actions are lost because of this.</p>
HYD-4250	<p>In some situations the GUI will display an Updates Required alert even immediately after an install or update.</p> <p>Recommendation: This alert should be cleared by repeating the update process.</p>
HYD-4311	<p>When removing the Intel Manager for Lustre agent from a server the following</p>

	<p>entry is left in the rsyslog configuration.</p> <pre># added by chroma-agent \$PreserveFQDN on *. * @@127.0.0.1:515;RSYSLOG_ForwardFormat # added by chroma-agent</pre> <p>Recommendation: This change should be manually removed from the file.</p>
HYD-4335	<p>It is not possible to configure the power control from the command line interface. The GUI interface must be used for power control configuration.</p>
HYD-4366	<p>The REST interface will return incorrect metric results if the latest=true flag is used.</p> <p>Recommendation: Always use begin and end to request metrics from the REST interface.</p>
HYD-4497	<p>For very large file systems, when stopping the file system the manager software may prematurely timeout the action and indicate that the operation failed.</p> <p>Recommendation: Stopping the file system again will bring the manager software into synchronization with the actual file system state.</p>
HYD-4509	<p>The manager software is unable to stop all the copytool processes on the client.</p> <p>Recommendation: Manual intervention will be required on the copytool host to stop the copytool processes.</p>
HYD-4529	<p>After changing NID values on a running file system, the system status may contain alerts that cannot be dismissed.</p> <p>Recommendation: The manager software will need to be restarted using the chroma_config restart command.</p>
HYD-4628	<p>In multiple MDT configurations, the connected client count represents the sum of client connections to all of the MDTs.</p> <p>Recommendation: To get the number of clients physically connected to the file system, divide the apparent client count by the number of MDTs.</p>
HYD-4639	<p>Executing failover and/or failback commands for multiple targets in an HA pair, in parallel, will leave location constraints configured in Pacemaker that</p>

	<p>Intel® Manager for Lustre* software expects to be removed. This will cause the removal of targets (and by extension the removal of a file system) to fail, and may cause future failover and/or failback actions to fail.</p> <p>Recommendation: At the GUI and the CLI, and when using the REST API, avoid executing failover and/or failback commands in parallel for multiple targets in an HA pair. Rather, execute those commands in serial, preferably with a pause of two seconds after issuing the preceding command.</p>
HYD-4866	<p>The Intel® Manager for Lustre* Software must be able to run in environments where the locale is not en_US.UTF8. The software is currently unable to run on computers where the local environment is set to something other than English. This means that operators attempting to install or use the manager software on systems using non-English must explicitly set the environment of the command shell to en_US.UTF8 or the software will abort.</p> <p>Recommendation: Set the LANG environment variable as follows before installing Intel® Manager for Lustre* software:</p> <pre>export LANG=en_US.UTF-8</pre>
HYD-4975	<p>Upgrading Intel® EE for Luster and restarting the manager server will overwrite any changes previously made to the chroma-manager.conf template in /etc/httpd/conf.d/. Before upgrading your installation or restarting your manager node, make sure you backup any modifications to this file first.</p>

Partner Installation Support

For resellers experiencing problems during the initial install and validation of their configuration, please go to <https://jira.hpdd.intel.com/browse/IEEL> and create a support ticket. When reporting a problem, please be as specific as possible about its nature (installation, configuration, etc.), and include information about the hardware and any underlying software used in your configuration. Note that tickets filed at the above site are confidential and not publicly viewable.

Reasonable commercial efforts will be made to respond promptly to a support ticket, however response times will vary. If you encounter any issues regarding your request, please contact your account manager.

Learn More

For related documentation and to learn more about Lustre, Intel® Enterprise Edition for Lustre* Software, and enterprise-class support and services, visit the Intel® Lustre* Partner Portal:

<http://www.intel.com/content/www/us/en/software/intel-solutions-for-lustre-software-reseller-portal.html>

Disclaimer and Legal Information

Copyright 2015 Intel Corporation. All Rights Reserved.

Portions of the software contained or described herein and all documents related to said software (referred to herein as "Material") are owned by Intel Corporation or its supplier or licensors. "

Title to the Material remains with Intel Corporation or its suppliers and licensors. The Material contains trade secrets and proprietary and confidential information of Intel or its suppliers and licensors. The Material is protected by worldwide copyright and trade secret laws and treaty provisions. No part of the Material may be used, copied, reproduced, modified, published, uploaded, posted, transmitted, distributed, or disclosed in any way without Intel's prior express written permission. (Continued)

No license under any patent, copyright, trade secret or other intellectual property right is granted to or conferred upon you by disclosure or delivery of the Materials, either expressly, by implication, inducement, estoppel or otherwise. Any license under such intellectual property rights must be express and approved by Intel in writing.

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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Before using any third party software referenced herein, please refer to the third party software provider's website for more information, including without limitation, information regarding the mitigation of potential security vulnerabilities in the third party software.

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 go to: <http://www.intel.com/design/literature.htm>.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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