

CONTENTS

1. OVERVIEW
 2. NEW FEATURES
 3. SYSTEM REQUIREMENTS
 4. INSTALLATION NOTES
 5. DOCUMENTATION
 6. KNOWN LIMITATIONS AND TROUBLESHOOTING
 7. TECHNICAL SUPPORT
 8. DISCLAIMER AND LEGAL INFORMATION
-

1. OVERVIEW

Intel(R) Cluster Checker verifies the configuration and performance of Linux*-based clusters and checks the cluster's compliance with the Intel(R) Cluster Ready Specification.

1.1. RELATED PRODUCTS AND SERVICES

Information about Intel(R) software development products is available at <http://www.intel.com/software/products>.

These are some of the products related to Intel(R) Cluster Checker:

- o Intel(R) Cluster Ready is a program designed to provide a common platform solution to build clustered applications. It helps organizations build, seamlessly deploy, and manage compatible and consistent applications using a clustered foundation. This program can significantly simplify the work of platform integrators, applications developers, and cluster component vendors.

<http://www.intel.com/go/cluster>

- o The Intel(R) C++ and Fortran Compilers include advanced optimization and multithreading capabilities, highly optimized performance libraries, and analysis tools for creating fast reliable multithreaded applications.

<http://www.intel.com/software/products/compilers>

- o The Intel(R) MPI Library for Linux*, the Intel(R) Trace Analyzer and Collector for Linux, and the Intel(R) Math Kernel Library Cluster Edition for Linux* are the most awarded development tools. They create, analyze, and optimize high-performance applications on clusters of Intel(R) processor-based systems.

<http://www.intel.com/software/products>

2. NEW FEATURES

2.1 WHAT'S NEW IN VERSION 2017

- Added support for Intel(R) Xeon Phi(TM) Product Family x200 processors.
- Added support for the Lustre* file system.
- Significantly improved analysis performance for large databases.
- Extended the API to provide the ability to collect data.
- The separate analyzer (clck.xml) and data collection (clckd.xml) configuration files were merged into a single file (clck.xml). Previous configuration files are not compatible.
- Allows more granular suppressions. See the section on suppressions in the User's Guide for more information.
- The asynchronous data collection daemons, clckd and clck-serverd, have been replaced by a plugin to the Open Resilient Cluster Manager* (ORCM). See orcm/README for more information. The ORCM plugin is a technical preview feature; please see the Known Limitation section for its current limitations.
- Databases from previous versions of the product are incompatible with version 2017 due to database schema changes.
- The samples have been moved online and are no longer distributed as part of the product. They are now available at the URL below:
<https://software.intel.com/en-us/product-code-samples?topic=20903>

2.2 OLDER VERSIONS

2.2.1 Version 3.1.2

- Added support for Intel(R) Omni-Path Architecture.
- Added "intel" prefix to RPM packages.
- Bug fixes

2.2.2 Version 3.1.1

- Support for certification of Intel(R) Cluster Ready specification 1.4
- The analyzer now supports validation against a schema for its XML configuration file.
- Checks that were not generating signs when the required data in the database was more than 7 days old now correctly do so.
- When using on-demand data collection, a pre-check is run which verifies that all nodes provided via the nodefile are accessible. If this check fails, on-demand collection will not proceed.
- Improved outlier detection in the analysis of performance data.
- Bug fixes

2.2.3 Version 3.1.0

- Analysis sets may be used to combine together multiple checks.
- On-demand data collection and analysis can now be stopped with an interrupt signal (such as CTRL-C) cleanly.
- Fixed an issue where comments in the <provider> tag of etc/clckd.xml were generating errors.
- When using on-demand data collection, the nodefile no longer needs to be located in a shared, writeable directory.
- When using on-demand data collection, the provider configuration files can now be owned by any user, not necessarily the user running the data collection.
- The 'compute' and 'enhanced' node roles may no longer be used for the same node.
- Signs are now generated if the latest data in the database is more than 7 days old.
- The path to the OFED tools is now configurable for the ibstat, ibv_devinfo and ofedinfo providers through environment variables set via the <provider> tag in etc/clckd.xml.
- The MPI-based providers now accept the <dapl_provider_list> tag in etc/clckd.xml, which maps to the I_MPI_DAPL_PROVIDER_LIST environment variable.
- Improved output and warning messages.

2.2.4 Version 3.0.1

- Added clck-analyze command line options to enable / disable checks, set the checking mode, and control the output.
- Updated the samples and Developer's Guide documentation.
- Bug fixes

2.2.5 Version 3.0

Intel(R) Cluster Checker 3.0 has been redesigned and re-implemented as an expert system, improving scalability, extensibility, ease of use, and diagnostic capability.

The key changes relative to the previous versions of Intel(R) Cluster Checker are:

- Data collection and analysis have been split into separate steps.

Note: the 'clck' script is provided to emulate the combined data collection and analysis behavior of previous versions.
- Data is stored in a database rather than being discarded after each run.

- The root causes of some issues can be diagnosed and reported, with recommendations to remedy the problem.
- An API is provided to embed Intel(R) Cluster Checker capabilities into user code.
- The ability to extend and enhance the checking capability is provided.

3. SYSTEM REQUIREMENTS

The following sections describe hardware and software requirements.

3.1. HARDWARE

- Intel(R) Xeon(R) processor (Intel(R) 64 architecture)
- 1 GB of RAM recommended
- 160 MB of free hard disk space required for installation

3.2. SOFTWARE

Operating Systems:

- Red Hat* Enterprise Linux* 6 and 7
- SUSE* Linux* Enterprise Server 11 and 12
- CentOS 6 and 7

Runtimes:

- Intel(R) MPI Library

Note: while the full SDK versions of these components fulfill the requirement, only the runtime is required.

4. INSTALLATION NOTES

Intel(R) Cluster Checker is distributed as part of Intel(R) Parallel Studio XE 2017 Cluster Edition and also as a standalone package.

To install Intel(R) Cluster Checker as part of Intel(R) Parallel Studio XE 2017 Cluster Edition, please refer to the Intel(R) Parallel Studio XE 2017 Cluster Edition documentation.

To install the standalone package, run the following commands:

```
% tar -xzf l_clck_p_2017.0.<package#>.tgz -C /tmp
% cd /tmp/l_clck_p_2017.0.<package#>
% ./install.sh
```

Notes:

- The default Intel(R) Cluster Checker install path is /opt/intel/clck/2017.0.<package#>.
- Intel(R) Cluster Checker needs to be installed on all nodes. This can either be accomplished either by installing into a shared directory or by installing a local copy on each node. Both options are supported by the installer.

5. DOCUMENTATION

This release of Intel(R) Cluster Checker includes the following documentation:

The Getting Started Guide walks through using Intel(R) Cluster Checker for the first time.

The Intel(R) Cluster Checker User's Guide contains information on how to use Intel(R) Cluster Checker. It describes the basic usage models and the basic steps using the tool.

The Intel(R) Cluster Checker Developer's Guide contains information on how to extend Intel(R) Cluster Checker. It describes how to embed Intel(R) Cluster Checker functionality into other applications and how to add new checks to the tool or modify existing checks.

The Intel(R) Cluster Checker Reference Manual contains information on specific configuration options. It describes environment variables and configuration options that influence checking behavior.

The Intel(R) Cluster Checker API reference describes the API that may be used to embed Intel(R) Cluster Checker functionality into other software programs.

The documentation can be found in the doc subdirectory.

6. KNOWN LIMITATIONS AND TROUBLESHOOTING

The following is a list of known issues in this release.

- Data collection behavior and functionality

- o If the temporary directory used during collection is located on a shared file system, the directory will not be deleted.
- o The ORCM plugin is a technical preview feature.
- o Databases located on NFS file systems mounted with the "nolock" option are not supported. Not all data from concurrent data collection instances per database will be written to the database and the database may become corrupted. A single data collector instance per database can usually be used successfully in this case.
- o The error "Error: disk I/O error" may be generated when accessing a database located on a Lustre file system. The Lustre file system must be mounted with the "-o flock" option.
- o The 'iozone' data provider does not execute correctly on diskless clusters.
- o If collecting data as root, the value of the CLCK_SHARED_TEMP_DIR environment variable must be set to the fully-qualified path of a directory accessible on all nodes.

- Analysis behavior and functionality

- o Clusters containing dual port InfiniBand* adapters where the second port is unused should suppress the 'infiniband-port-physical-state-not-linkup' and 'infiniband-port-state-not-active' signs. See Chapter 4 of the User's Guide for more information on how to suppress signs.
- o The Intel(R) Cluster Ready Specification requires that the runtime libraries of the Intel tools are installed in specific, non-default locations. Either download and install the separate Intel(R) Cluster Runtimes, or, to fulfill these requirements with Intel(R) Parallel Studio XE 2017 Cluster Edition, manually create symbolic links using the following commands.

```
% mkdir -p /opt/intel/cce/17.0.0
% ln -s /opt/intel/bin /opt/intel/cce/17.0.0/bin
% ln -s /opt/intel/lib/intel64 /opt/intel/cce/17.0.0/lib

% mkdir -p /opt/intel/fce/17.0.0
% ln -s /opt/intel/bin /opt/intel/fce/17.0.0/bin
% ln -s /opt/intel/lib/intel64 /opt/intel/fce/17.0.0/lib

% mkdir -p /opt/intel/cmkl
% ln -s /opt/intel/mkl /opt/intel/cmkl/2017.0.0
```

7. TECHNICAL SUPPORT

If you did not register Intel(R) Cluster Checker during installation, please do so at the Intel(R) Software Development Products Registration Center at <http://registrationcenter.intel.com>. Registration entitles you to free technical support, product updates and upgrades for the duration of the support term.

For information about how to find Technical Support, Product Updates, User Forums, FAQs, tips and tricks, and other support information, please visit: <http://www.intel.com/software/products/support/>

Note: If your distributor provides technical support for this product, please contact them for support rather than Intel.

8. DISCLAIMER AND LEGAL INFORMATION

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice.

Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, the Intel logo, Xeon, and Xeon Phi 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

© 2016 Intel Corporation.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804