

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) Scalable System Framework (Intel(R) SSF).

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) SSF is an architectural design foundation that enables development and deployment of a wide variety of high-performance, balanced solutions to support compute- and data-intensive workloads. Its standards-based programmability allows engineers to run diverse workloads on a broadly available, common infrastructure, making the power of high-performance computing accessible to businesses of all sizes.

<http://www.intel.com/ssf>

- 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 2019 Gold

- New 'clck' command simplifies execution with a single command.
- Added improved output messaging:
 - New compact summary output provided on screen.
 - Details of analysis provided in the output logfile.
 - Simplified scheme to assess issues as 'CRITICAL', 'WARNING', or 'INFORMATIONAL'
 - Added -R option for specifying where results are written.
 - Changed -o option to specify where log output is written.
- Added performance threshold checking for Intel(R) Xeon Phi(TM) Processor x205 Product Family.
- Added -X command line option, allowing a user to obtain a list of available framework definitions and their respective descriptions on data collected and analysis tests.
- Added the ability to mark two snapshots of a cluster state to identify changes. Currently supported with the following framework definitions: rpm_snapshot, hardware_snapshot, files_snapshot.
- Added a user option to collect any missing or old data before analysis.
- Added ability to collect data on a cluster that does not have pdsh if the Intel(R) MPI Library is installed.
- Added the ability to collect data without specifying a node file if nodes are allocated through SLURM.
- Added in support for validation of Intel(R) Select Solutions for Simulation and Modeling.
- New Intel(R) Cluster Checker API.

2.2 OLDER VERSIONS

2.2.1 Version 2018 Update 3

- Resolved an issue with running lshw without privileged access.

2.2.2 Version 2018 Update 2

- Added support for Intel(R) Select Solutions Framework Definitions.

2.2.3 Version 2018 Update 1

- Enabled provider configuration as part of Framework Definitions.
- Enabled latest Intel(R) Xeon(R) processors.
- Renamed predefined Framework Definition names to be more informative.
- Added support for the SGEMM benchmark.
- Improved memory checking.
- Enhanced user visible message output.
- Added checks for Intel(R) Omni-Path Fabric subnets.
- Bug fixes.

2.2.4 Version 2018

- Added support for Intel(R) Xeon(R) Scalable processors.
- Added Framework Definition feature to allow for customization of analysis.
- Added support for Intel(R) Turbo Boost Technology validation.
- Added support for analysis from multiple database sources.
- Updated samples and SDK.
- Converted documentation to online format.
- Enhanced Intel(R) Omni-Path Architecture validation.
- Added OpenFabrics Interfaces support.
- Enhanced user viewable message output.
- Bug fixes

2.2.5 Version 2017 Update 2

- Improved support for Intel(R) Xeon Phi(TM) Product Family x200 processors.
- Improved support for Intel(R) Omni-Path Architecture.
- Bug fixes.

2.2.6 Version 2017 Update 1

- Added Intel(R) Scalable System Framework Support.
- Deprecated support for Intel(R) Cluster Ready.
- Added additional support for Intel(R) Xeon Phi(TM) Product Family x200 processors.
- Removed heartbeat functionality.
- Bug fixes.

2.2.7 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>

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:

- CentOS 6 and 7
- Red Hat* Enterprise Linux* 6 and 7
- SUSE* Linux* Enterprise Server 11 and 12
- Ubuntu* 14.04, 16.04, and 17.04 (See Section 6 for known issues)

Runtimes:

- Intel(R) MPI Library

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

4. INSTALLATION NOTES

Intel(R) Cluster Checker is distributed as a standalone package.

To install package, run the following commands:

```
% tar -xzf l_clck_p_2019.0.<package#>.tgz -C /tmp
% cd /tmp/l_clck_p_2019.0.<package#>
% ./install.sh
```

Notes:

- The default Intel(R) Cluster Checker install path is /opt/intel/clck/2019.0.<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 about how to use, configure, and extend Intel(R) Cluster Checker. The User's Guide describes the basic usage models, contains information about specific configuration options, explains how to embed Intel(R) Cluster Checker functionality into other applications, shows how to add new checks to the tool, and demonstrates how to modify existing checks.

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 at <https://software.intel.com/en-us/intel-cluster-checker-support/documentation>.

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.
 - o When collecting data on Ubuntu*, if the installed "which" command does

not support `--skip-functions` and `--skip-alias`, a few providers will need additional configuration and a few providers will not run successfully. The following providers must be configured for the specification of absolute binary location:

- `cpuid`
- `cpupower`
- `dmesg`
- `ibstat`
- `lscpu`
- `numactl`
- `opahfirev`
- `opasmaquery`

Refer to Intel(R) Cluster Checker User Manual, Chapter 6 for details about specifying absolute binary paths for the above mentioned providers.

In addition, there are limitations to validating Intel(R) SSF compliance when running on Ubuntu.

It is not recommended to use Intel(R) Cluster Checker for Intel(R) SSF compliance when running on Ubuntu.

- 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 When using the Linux* boot parameter `isolcpus` with an Intel(R) Xeon Phi(TM) processor using default MPI settings, MPI based applications may fail. If possible, change or remove the `isolcpus` Linux* boot parameter. If this is not possible and you are using the Intel(R) MPI Library, you can try setting `I_MPI_PIN` to `off`. Refer to the Intel(R) Cluster Checker reference manual for details on specifying environment variables for tests.

o When run with `dgemm/dgemm_cpu_performance` or `stream/stream_memory_bandwidth_performance` framework, `"stream-outlier"` or `"dgemm-data-is-substandard"` may be observed as the corresponding provider scripts may not yield the expected performance with SNC-2/SNC-4 cluster mode and Flat memory mode configurations for Intel(R) Xeon Phi(TM) processor. There may be an issue with the kernel itself (BZ#1479763), documented at <https://access.redhat.com/errata/RHBA-2017:2581> If there are no corresponding diagnoses, the signs may be suppressed.

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

© 2019 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