

CONTENTS

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

1. OVERVIEW

Intel® Cluster Checker verifies the configuration and performance of Linux*-based clusters and checks the cluster's compliance with the Intel® Select Solutions for Simulation and Modeling.

1.1. RELATED PRODUCTS AND SERVICES

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

These are some of the products related to Intel® Cluster Checker:

- o The Intel® 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® MPI Library for Linux*, the Intel® Trace Analyzer and Collector for Linux*, and the Intel® 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® processor-based systems.

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

2. NEW FEATURES

2.1 WHAT'S NEW IN VERSION 2019 Update 10.0

- Added feature to propagate all user environment variables to all nodes when using pdsh for remote execution, with the exception of the following system variables: \$HOSTNAME, \$HOST, \$HOSTTYPE, \$MACHTYPE, \$OSTYPE, \$PMI_RANK, \$PMI_SIZE, \$PMI_FD, \$MPI_LOCALRANKID, \$MPI_LOCALNRANKS, \$DISPLAY, \$SHLVL, \$_, \$BASH_FUNC. This is the default setting.
- To suppress propagation of any environment variables to any nodes when using pdsh add command line option '-e' or 'export CLCK_TURN_OFF_ENV_PROPAGATION=true'
- the framework definition 'environment_variables_uniformity' can be used to inspect a list of all environment variables found on all nodes. (clckdb -D <path>/clck.db -provider printenv)
- Correct paths for the Intel® MPI Benchmarks providers to locate IMB binaries and removed imb_wrapper in imb_pingpong as it was not being used.
- Modifications to the messages provided when *-data-missing is reported by the analyzer. This is to improve the recommended remedies.
- Fix to correctly check for Intel Python 3, dropping check for Intel Python 2 with Intel® Parallel Studio XE 2020 runtimes or later.
- Fix for false positive signs from clck-analyzer about missing libraries for offload. These libraries are no longer required.
- Fix to execution of the data collection for the framework definition "mpi_local_functionality".
- Fix to standalone execution of framework definition "infiniband_admin".
- Extended node groups feature to work on relative framework definition paths.
- Improvements to several messages provided on 'data-is-too-old' or '*-data-missing-remedy' as reported by the analyzer.
- Corrected online documentation for hpcg_cluster framework definition.
- Improvement to the data collection of the "df" data provider.
- Added environment support for zsh and ksh in vars.sh script.
- Update SQLite to version 3.32.3 to fix:
 - CVE-2020-15358
- Update libxml2 to fix:
 - CVE-2020-24977
- Update zeromq to version 4.3.3 to fix:
 - CVE-2020-15166

2.2 OLDER VERSIONS

2.2.1 Version 2019 Update 9.0

- Further enhancements and improvements to the console output and reports were made including the following:
 - Additional details are provided when using the -n / --include-node analyzer filter parameter.
 - The output formatting reverts to previous versions when the new beta feature for node group analysis is not activated, -g / --groupfile.

- When using the new beta feature for node group analysis filtering, the ordering of groups will be alphabetical. Numerics will be treated as a string value and generated node groups will be printed last.
- Further improvements to the node group feature output were made including the following:
 - Better debugging and error messages for configuration set up.
 - Removed superfluous spacing.
 - Draw more attention to the line listing any runtime issues.
 - Removed duplicated information in node group specific output and added descriptors on the node group sections.
- Replace example node group configuration file with a better example.
- The '-h' help flag output has been revised
 - The '-n' flag has a changed description in the help output. Provided examples in the man pages.
- Better handling of Intel® MPI Benchmarks and OSU Micro-Benchmarks providers.
- Added a configurable timeout to collect extensions.
 - Timeout can be set as an environment variable or in the main configuration file.
 - Also documented the default timeout value in the configuration file.
- Added linking via RPATH:
 - Distributed common libraries moved to a separate 'lib_common' folder.
- Update SQLite to version to fix:
 - CVE-2020-11655 and CVE-2020-11656
 - CVE-2020-13434 and CVE-2020-13435
 - CVE-2020-13630, CVE-2020-13631 and CVE-2020-13632

2.2.2 Version 2019 Update 8.0

- Extensive changes made to the formatting of the output to enhance readability and parsing of the analysis.
 - Grouping of issues by type; functionality, performance, uniformity
 - Summary report on the number of each type of issue
 - Separation of Cluster Checker execution issues from system environmental issues
 - Detailed analysis and recommendations are detailed in an analysis log file
- An option to redirect the generated reports from Cluster Checker into a JSON formatted file to enable separate reporting and analysis by user tools
 - To enable, include `<entry>clck_json</entry>` as part of the `<postproc_extensions>` block of the cluster checker configuration XML file
 - default file `<INSTALL_PATH>/etc/clck.xml`
- Extended the collection and analysis capabilities of Intel® Cluster Checker to include the OSU Micro-Benchmark Suite (must be downloaded and installed separately: <http://mvapich.cse.ohio-state.edu/benchmarks/>). The OSU Collectives Benchmark Suite includes point-to-point, blocking and non-blocking collectives benchmark tests that verifies the functionality of the MPI functions.
- Introduction of new Intel® MPI blocking and non-blocking collectives benchmark tests for the verification of functionality of Intel® MPI. The results are collected and analyzed.

- Introduction of a new beta feature to perform analysis on specific node groups defined by the user or system administrator in a nodegroup file. This beta feature only applies to the analysis of the nodes, not the collection of data on the nodes. An example node group file is available in <INSTALL_PATH>/etc/example_group_file.xml. The documentation provides further details on this new beta feature.
- CVE fixes for libxml2 - CVE-2020-7595 and CVE-2019-20388

2.2.3 Version 2019 Update 7.0

- Improved error messaging covering:
 - when no high speed fabrics are detected for MPI
 - incorrect detection during analysis for select solutions priv plus
 - time out of collecting information using mpi.so collector
- Fix to correctly identify Intel Python 3.7.4
- Common Vulnerabilities and Exposures (CVE) for SQLite3: Updated SQLite3 package to sqlite3.31.0 to address a group of CVEs identified in late 2019.

2.2.4 Version 2019 Update 6.0

- Added environment modules support for Intel® Cluster Checker.
 - The Environment Module file is found <INSTALL_PATH>/clck/2019.7/env/modulefile/
 - The command: module use <INSTALL_PATH>/clck/2019.7/env/modulefile/ will add the clck to your module environment
 - module av will show what modules are available to be loaded via module load
- Added patch to address SQLite CVE-2019-9937.
- Improved clarity and response output for certain checks.
- Improved handling and analysis of data from multiple databases. Analysis is performed on the most recent data collected when there exists multiple collections of the same data types.

2.2.5 Version 2019 Update 5.0

- New default tests with faster execution.
- New predefined in-depth test sets made for user or admin specific analysis.
- Enhanced summary output lists brief facts on nodes and issues.
- Troubleshooting tests on prerequisites for Intel® MPI Library.
- Verifies the uniformity of the BIOS and management firmware settings for the Intel® Server Boards through Intel's System Configuration Utility (syscfg).
- Support for the latest Intel processors (Intel® Xeon® Platinum 9200 Processor Family).
- Support for validation of Intel® Select Solutions for Simulation & Visualisation.
- New -t option to set data age threshold.
- Bug fixes and improvements, CVE updates.

2.2.6 Version 2019 Update 4.0

- Enhanced functionality for testing memory uniformity.
- Added flexibility on checks for memlock limits to InfiniBand and Intel® Omni-Path Architecture (Intel® OPA) checks.
- Improved support for diskless clusters.
- Improved messages and bug fixes.

2.2.7 Version 2019 Update 3.5

- Added support for checking of second-generation Intel® Xeon® Scalable Processors by privileged or non-privileged users.
- Updated support for validation of Intel® Select Solutions for Simulation and Modeling to include the second-generation Intel® Xeon® Scalable Processor solution.
- Added support for checking Intel® Optane(TM) DC Persistent Memory configurations and uniformity.
- Included support for second-generation Intel® Xeon® Scalable Processor with the Intel HPC Platform Specification.
- Added checking for the Intel® Parallel Studio XE 2019.0 runtimes.

2.2.8 Version 2019 Update 2.1

- Updated support for validation of Intel® Select Solutions for Simulation and Modeling.
- Added in support for Intel HPC Platform Specification 2018.0.
- Intel® Cluster Checker 2019 Update 2.1 includes functional and security updates. Users should update to the latest version.

2.2.9 Version 2019 Update 2.0

- Intel® Cluster Checker 2019 Update 2 includes functional and security updates. Users should update to the latest version.

2.2.9 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® 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® 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® Select Solutions for Simulation and Modeling.
- New Intel® Cluster Checker API.

2.2.2 OLDER RELEASE NOTES

- The release notes for older, versions of Cluster Checker can be found at: <https://software.intel.com/en-us/articles/intel-cluster-checker-release-notes-and-new-features>

3. SYSTEM REQUIREMENTS

The following sections describe hardware and software requirements.

3.1. HARDWARE

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

3.2. SOFTWARE

Operating Systems:

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

Runtimes:

- Intel® MPI Library

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

4. WHERE TO FIND THE RELEASE

Intel® Cluster Checker can be installed with Intel® Parallel Studio XE, standalone via Intel® Registration Center: <https://registrationcenter.intel.com> or as a standalone package via the Intel® YUM repository.

See the Installation section of the User Guide for more information.

5. INSTALLATION NOTES

The default Intel® Cluster Checker install path is: `/opt/intel/clck/2019.10`

Intel® Cluster Checker is distributed as a standalone package.

To install package, run the following commands:

```
% tar -xzf <package name> -C /tmp
% cd /tmp/<package directory>
% ./install.sh
```

Notes:

- Intel® 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.
- To install a local copy on each node, repeat the package installation for each node.

6. DOCUMENTATION

This release of Intel® Cluster Checker includes the following documentation:

The Getting Started Guide walks through using Intel® Cluster Checker for the first time.

The Intel® Cluster Checker User's Guide contains information about how to use, configure, and extend Intel® Cluster Checker. The User's Guide describes the basic usage models, contains information about specific configuration options, explains how to embed Intel® Cluster Checker functionality into other applications, shows how to add new checks to the tool, and demonstrates how to modify existing checks.

The Intel® Cluster Checker API reference describes the API that may be used to embed Intel® Cluster Checker functionality into other

software programs.

The documentation can be found at:

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

7. KNOWN LIMITATIONS AND TROUBLESHOOTING

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

- Data collection behavior and functionality
 - o Intel® MPI Benchmarks (imb_*) providers will output false positives when run with MPICH.
 - o When executing data collection as root, the following framework definitions can just hang with no message, but maybe terminated by a single ctrl-c or terminate by a provider timeout:
 - mpi_multinode_functionality
 - hpl_cluster_performance
 - imb_* (Intel® MPI Benchmarks)
 - osu_* (OSU Micro-Benchmarks)
 - o When using mpi.so collector option (not default) with Intel MPI Library 2019 update 3, there is a known issue which will cause data to not be collected. Please use an earlier or later version of Intel MPI Library.
 - o Currently Cluster Checker will not collect data for the HPCG benchmarks correctly if Intel® MPI and MPICH (www.mpich.org) are both installed on the environment being tested. The test environment will look for and execute the Intel® MPI optimized binary for HPCG and thus reset the environment variables for MPICH. Discovery and handling of this limitation will be corrected in future versions of Intel® Cluster Checker.
 - o imb_pingpong_fabric_performance framework definition when launched with an odd number of nodes through Slurm, with MPI as the collector mechanism (mpi.so), will report no-data for the last server assigned to the slurm job. Workaround involves using a nodefile to specifically test the last server where 'no-data' was reported with another server in the infrastructure.
 - o The compute node hostname identified in the nodefile must match the hostname reported by the either the uname or hostname utility on the compute node itself. Deviations in the hostnames, or use of fully qualified domain names in either the nodefile or the compute node, may impact or produce inaccurate uniformity percentages and counts and be reported as a failure or warning by Cluster Checker.
 - o Please note that for execution of HPCG benchmarks (such as in the checks hpcg_single and hpcg_cluster) on non-standard install path for the Intel®

MPI Library and Intel® Math Kernel Library (Intel® MKL) runtime, libraries must be installed and be exported in the LD_LIBRARY_PATH on the system.

- o Use of the latest runtime libraries for Intel® MPI Library and Intel® Math Kernel Library is required to ensure compatibility with Intel® Cluster Checker.
- o If the temporary directory used during collection is located on a shared file system, the directory will not be deleted.
- 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 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 Intel® Cluster Checker uses the command "ldconfig -p" as well as the environment variable LD_LIBRARY_PATH to detect the presence of required libraries. In order for Intel® Cluster Checker to detect required libraries, they must be present in the LD_LIBRARY_PATH or the result of "ldconfig -p". (Applies to the Framework Definitions second-gen-xeon-sp_user, second-gen-xeon-sp_priv, intel_hpc_platform_compat-hpc-2018.0, intel_hpc_platform_sdvis-core-2018.0, and intel_hpc_platform_second-gen-xeon-sp-2019.0)
- o In order for Intel® Cluster Checker to detect the Intel® Distribution for Python*, it must be in the user's PATH. (Applies to the Framework Definitions second-gen-xeon-sp_user, second-gen-xeon-sp_priv, intel_hpc_platform_compat-hpc-2018.0, and intel_hpc_platform_second-gen-xeon-sp-2019.0)
- o If Intel® Parallel Studio is sourced before the Intel® Distribution for Python* in the user's environment, Intel® Cluster Checker is unable to detect all the required libraries for Intel® MPI Library. (Applies to the Framework Definitions second-gen-xeon-sp_user, second-gen-xeon-sp_priv, intel_hpc_platform_compat-hpc-2018.0, and intel_hpc_platform_second-gen-xeon-sp-2019.0)
- o The detected version of Intel® MPI Library is used to determine whether

Intel® Cluster Checker checks for Intel® Parallel Studio 2018 or 2019. If the Intel® MPI Library version does not match the version of the rest of Intel® Parallel Studio, the wrong set of libraries will be checked. (Applies to the Framework Definition `intel_hpc_platform_compat-hpc-2018.0`)

- o Intel® Cluster Checker can only detect the version of the Intel® Fortran Compiler version with Intel® Parallel Studio 2017 or later. (Applies to the Framework Definitions `second-gen-xeon-sp_user`, `second-gen-xeon-sp_priv`, `intel_hpc_platform_compat-hpc-2018.0`, and `intel_hpc_platform_second-gen-xeon-sp-2019.0`)
- o In addition, there are limitations to validating Intel® Select Solutions compliance when running on Ubuntu. It is not recommended to use Intel® Cluster Checker for Intel® Select Solutions 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® 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® MPI Library, you can try setting `I_MPI_PIN` to off. Refer to the Intel® 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® 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.
- o The sign `paraview-missing` fires despite ParaView* being present on the system. (Applies to the Framework Definition `intel_hpc_platform_sdvis-cluster-2018.0`)

If you did not register Intel® Cluster Checker during installation, please do so at the Intel® 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.

9. 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

© 2020 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