

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

<https://www.intel.com/content/www/us/en/developer/tools/oneapi/cluster-checker.html>

1.1. RELATED PRODUCTS AND SERVICES

Some of the products related to Intel® Cluster Checker include:

- o The Intel® oneAPI Base Toolkit contains the Intel® oneAPI DPC++/C++ Compiler and the Intel® oneAPI Math Kernel Library for Linux*, which include advanced optimization and multithreading capabilities, highly optimized performance libraries, and analysis tools for creating fast, reliable multithreaded applications.

<https://software.intel.com/content/www/us/en/develop/tools/oneapi/base-toolkit.html>

- o The Intel® oneAPI HPC Toolkit contains the Intel® MPI Library for Linux*, the Intel® Trace Analyzer and Collector for Linux*, and the Intel® Fortran Compiler Classic, all of which are among the most awarded development tools. They create, analyze, and optimize high-performance applications on clusters of Intel® processor-based systems.

<https://software.intel.com/content/www/us/en/develop/tools/oneapi/hpc-toolkit.html>

More information about Intel® oneAPI Toolkits software development products is available at:

<https://software.intel.com/content/www/us/en/develop/tools/oneapi.html>

2. NEW FEATURES

2.1 WHAT'S NEW IN VERSION 2021.7 (2021.7.0)

- Support added for the LSF scheduler
- Improved handling of invalid postprocessor extension input
- More concise output of numerical node names
- Improvements to the Select Solutions analysis and reporting
- Minor fixes and tweaks to provider functionality and parameters
- Update SQLite3* to latest version (version 3.39.2)
- Update libxml2 to latest version (version 2.9.14)

2.2 WHAT'S NEW IN VERSION 2021.6 (2021.6.0)

- Added support for directories as valid arguments for the '-o' command line option
- Support for Intel® Select Solutions for Professional Visualization 2022.0 including new framework definitions:
 - select_solutions_provis_benchmarks_base_2022.0
 - select_solutions_provis_benchmarks_plus_2022.0
 - select_solutions_provis_user_base_2022.0
 - select_solutions_provis_user_plus_2022.0
- Added support for checking allocated logical cores using the Intel(R) MPI Library cpuinfo tool
- Minor updates to providers for more robust functionality or to provide more information on runtime issues
 - Updates to provider parsing for tools with updated output.
- Update SQLite3*(v3.38.0), libxml2(v2.9.13), IOzone(v493) to latest version.

2.3 WHAT'S NEW IN VERSION 2021.5 (2021.5.1)

- Added checks to handle library name/version changes for the Intel® oneAPI Runtime libraries
 - Updated related output messages to improve clarity regarding any major components that were not found
- Updated linking to Intel® Math Kernel Library used by HPCG providers to work with the version included in Intel® oneAPI Toolkit 2022.1 release
 - Added better dependency error reporting in HPCG provider scripts
- Update SQLite3* to latest version (versions 3.37.2)

2.4 WHAT'S NEW IN VERSION 2021.5 (2021.5.0)

- Added support to discover the type(s) of fabric in use on a system and update the Cluster Checker environment and/or configuration variables accordingly
- Relocatable execution warning file; will now go into the same directory as the log output file specified with the '-o' option.
- Command line options added for collection methods/extensions
 - '-p' flag added for setting collection method

- '-i' flag added for checking which collection methods are available in the installation
- Minor bug fixes and output improvements

2.5 WHAT'S NEW IN VERSION 2021.4 (2021.4.0)

- Added a dependency tree for framework definitions to display the hierarchy of included framework definitions
 - Added tree at the top of '-X <FWD name>' output
 - Added tree as a section in the results output file
- Environment variable path expansion handling added to Framework Definitions
- Bug fixes and improvements
- Improved providers and parsing to account for newest output changes
- Updated third-party SQLite3*, IOzone*, CZMQ*, and ZeroMQ* libraries to latest stable versions
- A CVE with SQLite3 was patched and fixed, CVE-2021-36690.

2.6 WHAT'S NEW IN VERSION 2021.3 (2021.3.1)

Intel® Cluster Checker 2021.3 includes the following enhancements and bug fix:

- Improvements to the environment modulefile and vars.sh script
- Fixed a bug related to MPI workload execution. Benchmarks using Intel® MPI Library or MPICH will maintain the environment (variables) when running with the MPI collect extension
- Cleanup of unneeded environment variables from providers:
 - CLCK_PROVIDER_*_DAPL_PROVIDER_LIST (for setting I_MPI_DAPL_PROVIDER_LIST)
 - CLCK_PROVIDER_*_DAPL_PROVIDER (for setting I_MPI_DAPL_PROVIDER)
 - CLCK_PROVIDER_*_OFA_ADAPTER_NAME (for setting I_MPI_OFA_ADAPTER_NAME)
 - CLCK_PROVIDER_*_TMI_PROVIDER (for setting I_MPI_TMI_PROVIDER)
- Software Development Kit: Added missing header file to SDK include directory
- Latest libxml2 2.9.12 included addressing CVE-2021-3541 & previously patched CVE items.
- Analysis of IMB_pingpong extension uses latest collected data for analysis
- Environment propagation no longer propagates the local-only PROFILEREAD (readonly variable) and LC_* (locale settings) environment variables

2.7 WHAT'S NEW IN VERSION 2021.2 (2021.2.2)

This latest Intel® Cluster Checker release adds:

- Support for Third Generation Intel® Xeon® Scalable processors
- Support for Intel® HPC Platform Specification 2.0
- Support for Intel® Select Solutions for Simulation and Modeling 2021.0 including new framework definitions:
 - select_solutions_network_performance
 - select_solutions_sim_mod_benchmarks_plus_2021.0
 - select_solutions_sim_mod_priv_plus_2021.0
 - select_solutions_sim_mod_user_plus_2021.0
- A CVE with SQLite3 was patched and fixed, CVE-2021-20227.

2.8 WHAT'S NEW IN VERSION 2021.2 (2021.2.1)

- Improved execution of HPL and HPCG with mpirun launch command on workload managed systems (e.g. Slurm)
- Intel® Ethernet 800 series support (intel_ethernet800_* frameworks)
 - driver checks and version uniformity
 - required libraries
 - memlock unlimited check
 - PCIe checks
- setvars scripts compatible with mksh environments

2.9 WHAT'S NEW IN VERSION 2021.1 (2021.1.1)

Intel(R) Cluster Checker Version 2021.1 represents several significant improvements from the last major revision of Intel(R) Cluster Checker Version 2019. Many of the modifications and improvements listed below have been introduced in subsequent updates to Intel(R) Cluster Checker Version 2019 over the past year and included in beta releases for Version 2021. The following is a compilation of all the significant changes organized by category (not by update or beta release).

What's New Section Contents

- 2.9.1 Improvements to Output Formatting, Error Reports and JSON Support
- 2.9.2 New Support for OSU, MPICH, and MPI Workloads and Analysis
- 2.9.3 New and Modified Framework Definitions
- 2.9.4 New Feature for Node Group Analysis and Reporting
- 2.9.5 Support for Linux Environment Modules
- 2.9.6 Improved Support for ksh and zsh shells
- 2.9.7 Propagation of Environment Variables
- 2.9.8 Changes to Default Temporary File Permissions
- 2.9.9 Update to Checks on \$TMPDIR
- 2.9.10 Updates of Third Party Components to Correct CVEs

2.9.1 Improvements to Output Formatting, Error Reports and JSON Support

Intel(R) Cluster Checker has greatly improved its output formatting and error reporting. These changes have been introduced over the past year in the Update packages, and during the beta release for oneAPI. Intel(R) Cluster Checker has adopted a new summarized report upon completion, with details from the analysis available in the `clck_report.log` (or designated `-o <filename>`), and details about the execution environment for Intel(R) Cluster Checker in the `clck_execution_warnings.log` file.

2.9.2 New Support for OSU and Intel(R) MPI workloads and analysis

Intel(R) Cluster Checker now provides a wide array of workloads and test suites available for the evaluation and continued performance analysis of the compute environment. Support is now extended to the execution of Ohio State University (OSU) workloads. While the specific OSU workloads are not provided as part of

the standard Intel(R) Cluster Checker installation, the OSU workloads can be installed separately and used by Intel(R) Cluster Checker for the evaluation and analysis of the environment. Many of the OSU workloads support the MPICH messaging protocols.

2.9.3 New and modified framework definitions

Several new framework definitions have been added to Intel(R) Cluster Checker to aid in the validation and analysis of clustered environments, while others have been updated to provide a cleaner interface and better understanding of the frameworks.

These new and modified framework definitions include the following:

```
health_base
health_user
health_exteded_user
health_admin
health_extended_admin
osu_*
```

For a full listing of the available framework definitions execute the following on the command line:

```
clck -X list
```

For details about a specific framework definition either refer to the reference section of the User's Guide or perform the following on the command line:

```
clck -X <framework definition name>
example: clck -X health_base
```

2.9.4 New Feature for Node Group Analysis and Reporting

Intel(R) Cluster Checker has introduced the ability to analyze discrete groups of compute nodes. The new feature utilizes a command line option '-g' and an xml file designating the grouping of the compute nodes. The grouping could be based upon physical characteristics of the nodes, location, or other features. A compute node can belong in multiple node groups at the same time. The new feature only applies to the analysis of the nodes. Collection of the data for the specific framework definitions is across all designated nodes. Please see the User's Guide for more details on implementing this feature, the xml file setup and example, and interpretation of the analysis results.

2.9.5 Support for Linux Environment Modules

Intel(R) Cluster Checker now supports the use of Linux environment modules for setting up the proper runtime environment. The environment modules support provides an alternative to the current setup requirements found in the clckvars.[sh|csh] script.

Please see the User's Guide for additional details and setup requirements.

2.9.6 Improved Support for ksh and zsh shells

Intel(R) Cluster Checker has expanded the number of shell environments that can be utilized for the collection and analysis of clustered systems by supporting ksh (korn shell) and zsh (z-shell). Specific details, setup, and exemptions are included in the User's Guide for Intel(R) Cluster Checker.

2.9.7 Propagation of Environment Variables

Intel(R) Cluster Checker has now added environment variable propagation with its default pdsh-based remote execution on compute nodes. This enables proper operation of data collection on compute nodes within consistent environment settings. A new command line option '-e' was added to turn off the environment variable propagation for backward compatibility.

2.9.8 Changes to Default Temporary File Permissions

Intel(R) Cluster Checker has modified its default temporary file permissions to remove the "world-writable" flag on all files and the temporary directory used by Intel(R) Cluster Checker to store temporary files before analysis. This change is to provide better security to users and ensure only those authorized to access Intel(R) Cluster Checker results are allowed to do so. Intel(R) Cluster Checker still allows users and groups already authorized to execute the tools to read and write files, but limits access by others.

2.9.9 Update to Checks on \$TMPDIR

Intel(R) Cluster Checker has modified the check for file permissions to not require any 'world' permissions for \$TMPDIR.

2.9.10 Updates of Third Party Components to Correct CVEs

Intel(R) Cluster Checker leverages several third-party components to assist in the collection, analysis, and generation of reports. To maintain security commitments to Intel(R) and our customers, Intel(R) Cluster Checker is regularly scanned using industry-standard tools for exposure of Common Vulnerable Exploits in the third-party components. It is important for customers using Intel(R) Cluster Checker to maintain their implementations at the latest level to not only receive the latest improvements to the collection and analysis, but also to receive the most recent updates to these third-party components.

2.10 OLDER VERSIONS

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

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
- Red Hat* Enterprise Linux* 8
- Rocky Linux* 8
- SUSE* Linux* Enterprise Server 15 SP2
- SUSE* Linux* Enterprise Server 15 SP3

3.3. ADDITIONAL REQUIREMENTS

Provider Requirements:

- HPCG benchmark providers (hpcg_single and hpcg_cluster):
 - Intel® MPI Library
 - Intel® Math Kernel Library (Intel® MKL)

Extension Requirements:

- For MPI-based collect (mpi.so):
 - Implementation of MPI (Intel® MPI Library is recommended)

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® oneAPI HPC Toolkit.

See the Installation section of the User Guide or the Installation documentation for the Intel® oneAPI HPC Toolkit for more information.

5. INSTALLATION NOTES

The default Intel® Cluster Checker install path is:
/opt/intel/oneapi/clck/2021.7.0

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 on 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://www.intel.com/content/www/us/en/developer/tools/oneapi/cluster-checker-documentation.html>

7. KNOWN LIMITATIONS AND TROUBLESHOOTING

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

- Data collection behavior and functionality
 - o An issue when launching Cluster Checker over MPI with SLURM 20.11.X has been identified. If mpi.so is used to launch clck through a slurm job submission, any MPI application used by Cluster Checker will not

start. This is caused by a change in how SLURM grants access to assigned resources. Please refer to the SLURM RELEASE NOTES for details <https://slurm.schedmd.com/archive/slurm-20.11.6/news.html>

To work around this issue, the SLURM_OVERLAP environment variable can be set to "1":

```
export SLURM_OVERLAP=1
```

- o When `imb_pingpong_fabric_performance` framework definition is launched with an odd number of nodes through Slurm with MPI as the collector mechanism (`mpi.so`), it 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 will normally be successful 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 In order for Intel® Cluster Checker to detect required libraries, they must be present in the user's `LD_LIBRARY_PATH` or output by the "`ldconfig -p`" command. (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 Intel® TrueScale InfiniBand(TM) does not have driver support for Operating Systems supported by oneAPI. As a result collection of TrueScale data is not supported with Intel® Cluster Checker.
- o There are known limitations to validating Intel® Select Solutions compliance when running on Ubuntu. It is therefore not recommended to use Intel® Cluster Checker for Intel® Select Solutions compliance in this configuration.

- Analysis behavior and functionality

- o `select_solutions_network_performance` framework definition requires a minimum of 4 nodes. If run with 3 or fewer nodes, a warning is issued that does not explain the failure. This will be fixed in future releases.
- 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 the Intel® MPI Library is being used, one option is to set `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)

- Installation Limitation

- o The 2021.2 release conflicts with 2021.1 release if it was installed by YUM/DNF/Zypper package manager. Remove the 2021.1 release before installing 2021.2 release. To install both releases side-by-side, use web installer tool.

8. TECHNICAL SUPPORT

Intel® Cluster Checker is part of the Intel® oneAPI HPC toolkit. Information on priority support for Intel® oneAPI Base and HPC Toolkits is available here: <https://software.intel.com/en-us/oneapi/support>

Community forum based support is available through the forums: <https://community.intel.com/t5/Intel-oneAPI-HPC-Toolkit/bd-p/oneapi-hpc-toolkit>

Information on how to locate Technical Support, Product Updates, User Forums, FAQs, tips and tricks, and other support information: <https://software.intel.com/en-us/oneapi/support>

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

© 2022 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