# Intel® Cluster Checker Reference Manual

Version 3.0.1
August 3, 2015

# Disclaimer and Legal Information

## Optimization Notice

Intel 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

# Contents

# Chapter 1

# Analyzer

The behavior of Intel® Cluster Checker can be configured through an XML file that is specified as a command line configuration option. The default location of the file is `/opt/intel/clck/3.0.1/etc/clck.xml`

The `config` XML tag can be used to configure the behavior of the knowledge base.

## clck-checks

Specify the checks to be performed.

### XML syntax

```
<config>
  <clck-checks>
    <entry>CHECK</entry>
    <entry>CHECK</entry>
  </clck-checks>
</config>
```

### Description

The clck-checks tag specifies the checks that should be performed. Only checks that are specified will be performed.
Valid values for CHECK are given in Appendix A.
The entry tag can be repeated multiple times.

## clck-mode

Specify the mode under which Intel® Cluster Checker should run.

### XML syntax

```
<config>
  <clck-mode>MODE</clck-mode>
</config>
```

## Description

This specifies the mode that Intel® Cluster Checker should run under. Please see the User's Guide for more information.
Valid values for `MODE` are `certification`, `compliance` and `health`.

## data-age-threshold

Specify the maximum age of data points, in seconds, before a data point is considered "too old" for relevant analysis.

## XML syntax

```
<config>
  <data-age-threshold>NUMBER</data-age-threshold>
</config>
```

## Description

The `data-age-threshold` tag specifies the maximum age of data points, in seconds, before a data point is considered "too old" for relevant analysis.
The value should be an integer value greater than 0. The default value is 604800 seconds (1 week).

## dgemm-number-of-mad

Specify the number of median absolute deviations allowed before a dgemm value is considered an outlier.

## XML syntax

```
<config>
  <dgemm-number-of-mad>NUMBER</dgemm-number-of-mad>
</config>
```

## Description

The `dgemm-number-of-mad` tag specifies the number of allowed median absolute deviations before a dgemm value is considered an outlier.
The value should be an integer value greater than 0.

## dgemm-peak-fraction

Specify the minimum value of the ratio between the measured dgemm performance and theoretical peak value.

## XML syntax

```
<config>
  <dgemm-peak-fraction>NUMBER</dgemm-peak-fraction>
</config>
```

## Description

The `dgemm-peak-fraction` tag specifies the minimum value for the ratio between the measured dgemm performance and the theoretical peak performance. Any value below this will generate a sign.
The value should be an floating point value between $0$ and $1$.

## environment-blacklist

Specify the environment variable patterns that will be ignored for uniformity comparison.

## XML syntax

```
<config>
  <environment-blacklist>
    <entry>PATTERN</entry>
    <entry>PATTERN</entry>
  </environment-blacklist>
</config>
```

## Description

The environment-blacklist tag excludes environment variables from being compared for uniformity across the cluster.
The value within each `entry` tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.
The `entry` tag can be repeated multiple times.
Note that to exactly match meta characters, (^[.*(${\()+|?<>) , they should be escaped.

## hpl-number-of-mad

Specify the number of median absolute deviations allowed before an HPL value is considered an outlier.

## XML syntax

```
<config>
  <hpl-number-of-mad>NUMBER</hpl-number-of-mad>
</config>
```

## Description

The `hpl-number-of-mad` tag specifies the number of allowed median absolute deviations before an HPL value is considered an outlier.
The value should be an integer value greater than $0$.

## imb_pingpong-number-of-mad

Specify the number of median absolute deviations allowed before an HPL value is considered an outlier.

## XML syntax

```
<config>
   <imb_pingpong-number-of-mad>NUMBER</imb_pingpong-number-of-
       mad>
</config>
```

## Description

The imb_pingpong-number-of-mad tag specifies the number of allowed median absolute deviations before a PingPong latency or bandwidth value is considered an outlier.
The value should be an integer value greater than 0.

## kernel-blacklist

Specify the kernel parameter patterns that will be ignored for uniformity comparisons.

## XML syntax

```
<config>
   <kernel-blacklist>
      <entry>PATTERN</entry>
      <entry>PATTERN</entry>
   </kernel-blacklist>
</config>
```

## Description

The kernel-blacklist tag excludes kernel parameters from being compared for uniformity across the cluster.
The value within each entry tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.
The entry tag can be repeated multiple times.
Note that to exactly match meta characters, (^[.*(${\()+|?<>) , they should be escaped.

## lshw-blacklist

Specify the lshw output patterns that will be ignored for uniformity comparison.

## XML syntax

```
<config>
   <lshw-blacklist>
      <entry>PATTERN</entry>
      <entry>PATTERN</entry>
   </lshw-blacklist>
</config>
```

## Description

The lshw-blacklist tag excludes lshw items from being compared for uniformity across the cluster.

The value within each `entry` tag is interpreted as a POSIX* matching regular expression. If this value is not a valid POSIX* regular expression, then no filtering will be done.

The `entry` tag can be repeated multiple times.

Note that to exactly match meta characters, `(^[.*(${\()+|?<>)`, they should be escaped.

## lshw-uniformity-threshold

Specify the threshold ratio for checking the uniformity of `lshw` entries.

## XML syntax

```
<config>
  <lshw-uniformity-threshold>NUMBER</lshw-uniformity-threshold>
</config>
```

## Description

The `lshw` uniformity check determines whether each `lshw` entry is uniform across the cluster. If the percentage of nodes that share the same `lshw` entry value is above the value specified for the `lshw-uniformity-threshold` tag, then that value is considered uniform in that cluster. If the percentage of nodes that share the same `lshw` entry value is below the uniformity threshold, then a sign is generated.

The value should be an floating point value between $0$ and $1$.

## memory-uniformity-threshold

Specify the maximum allowable deviation, in bytes, from the median memory size before a memory size is considered non-uniform.

## XML syntax

```
<config>
  <memory-uniformity-threshold>NUMBER</memory-uniformity-
      threshold>
</config>
```

## Description

The `memory-uniformity-threshold` tag specifies the maximum allowable deviation, in bytes, from the median memory size before a memory size is considered non-uniform.

Any value greater than $0$ can be used for this tag. The default value is 268435456 bytes (256 MB).

## ntp-offset-threshold

Specify the maximum offset value that an NTP peer is allowed to have.

**XML syntax**

```
<config>
  <ntp-offset-threshold>NUMBER</ntp-offset-threshold>
</config>
```

**Description**

The `ntp-offset-threshold` tag specifies the maximum offset value that an
NTP peer is allowed to have before a sign is generated.
Any floating point value can be used for this tag.

# Chapter 2

# Data Providers

The data provider configuration mechanism is environment variables. The name space `CLCK_PROVIDER_PROVIDER-NAME` is reserved for each provider, where the `PROVIDER-NAME` is the name of the corresponding provider. Each provider can define its own variables within its name space, and each provider is responsible for reading its own environment variables.

While these variables can be set in the users shell initialization files, e.g., `.bashrc`, the recommended way is to define these variables in the configuration file that is read and propagated by `clck-collect` and `clckd`.

The XML configuration file hides the raw environment variable names. Instead, the environment variable gets constructed by combining the XML tag names.

For example, the following construct -

```xml
<configuration>
  <provider>
    <hpl_cluster>
      <fabrics>value</fabrics>
    </hpl_cluster>
  </provider>
</configuration>
```

will generate the environment variable `CLCK_PROVIDER_HPL_CLUSTER_FAB-RICS`, which would then be made available to the `HPL_CLUSTER` provider.

Note, the value is enclosed within double quotes when passed to the provider. Consequently, if the value provided in the XML contains double quotes, it will be ignored and the environment variable will not be constructed.

Additionally, XML configuration files can also be used to enable running the same provider multiple times with disjoint sets of configuration options. For example, the following XML configuration would run the `HPL_CLUSTER` twice, once with `CLCK_PROVIDER_HPL_CLUSTER_FABRICS=shm:tcp` and once with `CLCK_-PROVIDER_HPL_CLUSTER_FABRICS=shm:dapl`

```xml
<configuration>
  <provider>
    <hpl_cluster>
      <optionset id="1">
        <fabrics>shm:tcp</fabrics>
      </optionset>

      <optionset id="2">
        <fabrics>shm:dapl</fabrics>
      </optionset>
    </hpl_cluster>
  </provider>
</configuration>
```

The "id" attribute is required. For each run, the value of the corresponding "id" attribute will be stored in the database under the "OptionID" column. The value of this attribute is purely arbitrary and is not meaningful other than to distinguish one set of options from another in the output columns of the database. However, the value must not contain tabs, newlines, parantheses, ampersands, vertical bars, the symbol '<', the symbol '~', spaces, semicolons or double quotes.

If no "optionset" tags are specified, the "OptionID" column will contain the string "default".

Additionally, if any tags are specified outside of "optionset" tags, the corresponding environment variables will be made available for every single "optionset" specified.

## 2.1   hpl_cluster

### CLCK_PROVIDER_HPL_CLUSTER_FABRICS

Select the network fabric to be used.

### Environment variable syntax

```
CLCK_PROVIDER_HPL_CLUSTER_FABRICS=value
```

### XML syntax

```
<configuration>
  <provider>
    <hpl_cluster>
      <fabrics>value</fabrics>
    </hpl_cluster>
  </provider>
</configuration>
```

### Arguments

value

The network fabric to use.

### Description

This value directly maps to the I_MPI_FABRICS Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

### CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER

Select the DAPL provider to use.

### Environment variable syntax

```
CLCK_PROVIDER_HPL_CLUSTER_DAPL_PROVIDER=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_cluster>
      <dapl_provider>value</dapl_provider>
    </hpl_cluster>
  </provider>
</configuration>
```

**Arguments**

value

> The DAPL provider to use.

**Description**

This value directly maps to the `I_MPI_DAPL_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

**CLCK_PROVIDER_HPL_CLUSTER_OFA_ADAPTER_NAME**

Select the OFA adapter name to use.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_CLUSTER_OFA_ADAPTER_NAME=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_cluster>
      <ofa_adapter_name>value</ofa_adapter_name>
    </hpl_cluster>
  </provider>
</configuration>
```

**Arguments**

value

> The OFA adapter name to use.

**Description**

This value directly maps to the `I_MPI_OFA_ADAPTER_NAME` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

**CLCK_PROVIDER_HPL_CLUSTER_TCP_NETMASK**

Select the TCP netmask value to use.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_CLUSTER_TCP_NETMASK=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_cluster>
      <tcp_netmask>value</tcp_netmask>
    </hpl_cluster>
  </provider>
</configuration>
```

**Arguments**

value

The TCP netmask value to use.

**Description**

This value directly maps to the `I_MPI_TCP_NETMASK` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

**CLCK_PROVIDER_HPL_CLUSTER_TMI_PROVIDER**

Select the TMI provider to use.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_CLUSTER_TMI_PROVIDER=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_cluster>
      <tmi_provider>value</tmi_provider>
    </hpl_cluster>
  </provider>
</configuration>
```

**Arguments**

value

The TMI provider to use.

**Description**

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

**CLCK_PROVIDER_HPL_CLUSTER_OPTIONS**

Any additional desired options to be provided to the mpirun command.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_CLUSTER_OPTIONS=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <hpl_cluster>
      <options>value</options>
    </hpl_cluster>
  </provider>
</configuration>
```

**Arguments**

value

   Any options that are passed verbatim to the mpirun command.

**Description**

The value of this tag is passed as is to the mpirun command.

## 2.2   hpl_pairwise

**CLCK_PROVIDER_HPL_PAIRWISE_FABRICS**

Select the network fabric to be used.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_PAIRWISE_FABRICS=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <hpl_pairwise>
      <fabrics>value</fabrics>
    </hpl_pairwise>
  </provider>
</configuration>
```

**Arguments**

value

   The network fabric to use.

### Description

This value directly maps to the `I_MPI_FABRICS` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

### CLCK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER

Select the DAPL provider to use.

### Environment variable syntax

```
CLCK_PROVIDER_HPL_PAIRWISE_DAPL_PROVIDER=value
```

### XML syntax

```xml
<configuration>
  <provider>
    <hpl_pairwise>
      <dapl_provider>value</dapl_provider>
    </hpl_pairwise>
  </provider>
</configuration>
```

### Arguments

value

The DAPL provider to use.

### Description

This value directly maps to the `I_MPI_DAPL_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

### CLCK_PROVIDER_HPL_PAIRWISE_OFA_ADAPTER_NAME

Select the OFA adapter name to use.

### Environment variable syntax

```
CLCK_PROVIDER_HPL_PAIRWISE_OFA_ADAPTER_NAME=value
```

### XML syntax

```xml
<configuration>
  <provider>
    <hpl_pairwise>
      <ofa_adapter_name>value</ofa_adapter_name>
    </hpl_pairwise>
  </provider>
</configuration>
```

**Arguments**

value

> The OFA adapter name to use.

**Description**

This value directly maps to the `I_MPI_OFA_ADAPTER_NAME` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

**CLCK_PROVIDER_HPL_PAIRWISE_TCP_NETMASK**

Select the TCP netmask value to use.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_PAIRWISE_TCP_NETMASK=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_pairwise>
      <tcp_netmask>value</tcp_netmask>
    </hpl_pairwise>
  </provider>
</configuration>
```

**Arguments**

value

> The TCP netmask value to use.

**Description**

This value directly maps to the `I_MPI_TCP_NETMASK` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.

When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

**CLCK_PROVIDER_HPL_PAIRWISE_TMI_PROVIDER**

Select the TMI provider to use.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_PAIRWISE_TMI_PROVIDER=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_pairwise>
      <tmi_provider>value</tmi_provider>
    </hpl_pairwise>
  </provider>
</configuration>
```

**Arguments**

value

   The TMI provider to use.

**Description**

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

**CLCK_PROVIDER_HPL_PAIRWISE_OPTIONS**

Any additional desired options to be provided to the mpirun command.

**Environment variable syntax**

```
CLCK_PROVIDER_HPL_PAIRWISE_OPTIONS=value
```

**XML syntax**

```
<configuration>
  <provider>
    <hpl_pairwise>
      <options>value</options>
    </hpl_pairwise>
  </provider>
</configuration>
```

**Arguments**

value

   Any options that are passed verbatim to the mpirun command.

**Description**

The value of this tag is passed as is to the mpirun command.

## 2.3   imb_pingpong

**CLCK_PROVIDER_IMB_PINGPONG_FABRICS**

Select the network fabric to be used.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_FABRICS=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <fabrics>value</fabrics>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

value

 The network fabric to use.

**Description**

This value directly maps to the `I_MPI_FABRICS` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate fabric.

**CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER**

Select the DAPL provider to use.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_DAPL_PROVIDER=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <dapl_provider>value</dapl_provider>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

value

 The DAPL provider to use.

**Description**

This value directly maps to the `I_MPI_DAPL_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate DAPL provider.

### CLCK_PROVIDER_IMB_PINGPONG_OFA_ADAPTER_NAME

Select the OFA adapter name to use.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_OFA_ADAPTER_NAME=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <ofa_adapter_name>value</ofa_adapter_name>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

value

> The OFA adapter name to use.

**Description**

This value directly maps to the `I_MPI_OFA_ADAPTER_NAME` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

### CLCK_PROVIDER_IMB_PINGPONG_TCP_NETMASK

Select the TCP netmask value to use.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_TCP_NETMASK=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <tcp_netmask>value</tcp_netmask>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

value

> The TCP netmask value to use.

**Description**

This value directly maps to the `I_MPI_TCP_NETMASK` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate OFA adapter name.

**CLCK_PROVIDER_IMB_PINGPONG_TMI_PROVIDER**

Select the TMI provider to use.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_TMI_PROVIDER=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <tmi_provider>value</tmi_provider>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

  value

        The TMI provider to use.

**Description**

This value directly maps to the `I_MPI_TMI_PROVIDER` Intel® MPI Library environment variable. Refer to the Intel® MPI Library Reference Manual for more information and recognized values.
When this variable is not set, Intel® MPI Library automatically chooses the most appropriate TMI provider.

**CLCK_PROVIDER_IMB_PINGPONG_OPTIONS**

Any additional desired options to be provided to the mpirun command.

**Environment variable syntax**

```
CLCK_PROVIDER_IMB_PINGPONG_OPTIONS=value
```

**XML syntax**

```xml
<configuration>
  <provider>
    <imb_pingpong>
      <options>value</options>
    </imb_pingpong>
  </provider>
</configuration>
```

**Arguments**

value

>    Any options that are passed verbatim to the mpirun command.

**Description**

The value of this tag is passed as is to the mpirun command.

# Chapter 3

# Data Collection

The following options alter the behavior of data collection.

## CLCK_COLLECT_DATABASE_BUSY_TIMEOUT

Specify the amount of time to wait for a database lock to become available.

### Environment variable syntax

```
CLCK_COLLECT_DATABASE_BUSY_TIMEOUT=value
```

### Arguments

value

> The number of milliseconds to wait for a database lock to become available before giving up. The value must be greater than 0. The default value is 60,000 milliseconds.

### Description

When inserting a new row into the database, the database is locked and any concurrent write attempts are prevented. This value specifies the amount of time that the concurrent write(s) should wait for the database to be unlocked before giving up. If the timeout expires and the database is still locked, the concurrent write(s) will not be successful and the data will be lost.

## CLCK_COLLECT_DATABASE_CLOSE_DELAY

Specify the amount of time to wait after data collection has finished for data to arrive.

### Environment variable syntax

```
CLCK_COLLECT_DATABASE_CLOSE_DELAY=value
```

### Arguments

value

> The number of seconds to wait after data collection has finished for any remaining data to be accumulated. The value must be greater than 0. The default value is 1 second.

**Description**

All data that is in the accumulate queue will always be written to the database, but some data may still be "on the wire" when data collection has finished. This option provides a method to wait an additional amount of time for data to be received by the accumulate server before exiting. Clusters with very slow networks or a very large number of nodes may need to increase this value from the default.
This option is only recognized when using on-demand data collection.

**CLCK_COLLECT_DATABASE_VFS_MODULE**

Specify the SQLite* VFS module.

**Environment variable syntax**

CLCK_COLLECT_DATABASE_VFS_MODULE=value

**Arguments**

value = unix

Uses POSIX advisory locks when locking the database. Note that the implementation of POSIX advisory locks on some filesystems, e.g., NFS, is incomplete and/or buggy. This value should usually only be selected when the database is located on a local filesystem.

value = unix-dotfile

Uses dot-file locking when locking the database. This value usually works around filesystem implementation issues related to POSIX advisory locks. This is the default value.

value = unix-excl

Obtains and holds an exclusive lock on the database file. All concurrent database operations will be prevented while the lock is held. This value may help in the event of database errors during data collection of if collected data is missing from the database.

value = unix-none

No locking is used. This option should only be used if there is a guarantee that only a single writer will modify the database at any given time. Otherwise, this value can easily result in database corruption if two or more processes are writing to the database concurrently.

**Description**

The SQLite* OS interface layer, or "VFS", can be selected at runtime. The VFSes differ primarily in the way they handle file locking. See http://www.sqlite.org/vfs.html for more information.
This option is only recognozed when using on-demand data collection.

# Appendix A

# List of Checks

| Check | Description | Default State |
|-------|-------------|---------------|
| all_to_all | IP address consistency | Enabled |
| cpu | Intel® Cluster Ready CPU compliance | Enabled |
| dgemm | Floating point performance | Enabled |
| environment | Environment variables | Enabled |
| ethernet | Ethernet driver uniformity and wellness | Enabled |
| heartbeat | Verify data is recent | Enabled |
| hpl | High Performance Linpack* | Enabled |
| icr_cluster | Intel® Cluster Ready minimum node count compliance | Enabled |
| icr_version | Intel® Cluster Ready version compliance | Enabled |
| imb_pinpgong | MPI performance | Enabled |
| infiniband | InfiniBand* uniformity and wellness | Enabled |
| intel_cluster_runtimes | Intel® Cluster Ready runtime library compliance | Enabled |
| iozone | Disk I/O performance | Enabled |
| java | Java* uniformity and functionality | Enabled |
| kernel | Linux* kernel | Enabled |
| kernel_param | Kernel parameter uniformity | Enabled |
| libraries | Intel® Cluster Ready runtime library compliance | Enabled |
| lshw | Hardware uniformity | Enabled |
| memory | Memory compliance | Enabled |
| miccheck | Intel® Xeon Phi™ coprocessor uniformity and wellness | Enabled |
| micinfo | Intel® Xeon Phi™ coprocessor uniformity and wellness | Enabled |
| mount | Mount point compliance and uniformity | Enabled |
| mpi_local | Single-node MPI functionality | Enabled |
| mpi_internode | Multi-node MPI functionality | Enabled |
| ntp | Clock synchronization | Enabled |
| offload_phi | Intel® Xeon Phi™ coprocessor functionality | Enabled |
| perl | Perl* compliance, uniformity, and functionality | Enabled |
| process | Process table | Enabled |
| python | Python* compliance, uniformity, and functionality | Enabled |
| rpm | RPM uniformity | Enabled |
| shells | Shell compliance | Enabled |

| Check | Description | Default State |
|---|---|---|
| storage | Disk capacity | Enabled |
| stream | Memory bandwidth performance | Enabled |
| tcl | Tcl compliance, uniformity, and functionality | Enabled |
| x11_tools | X11 tool compliance | Enabled |