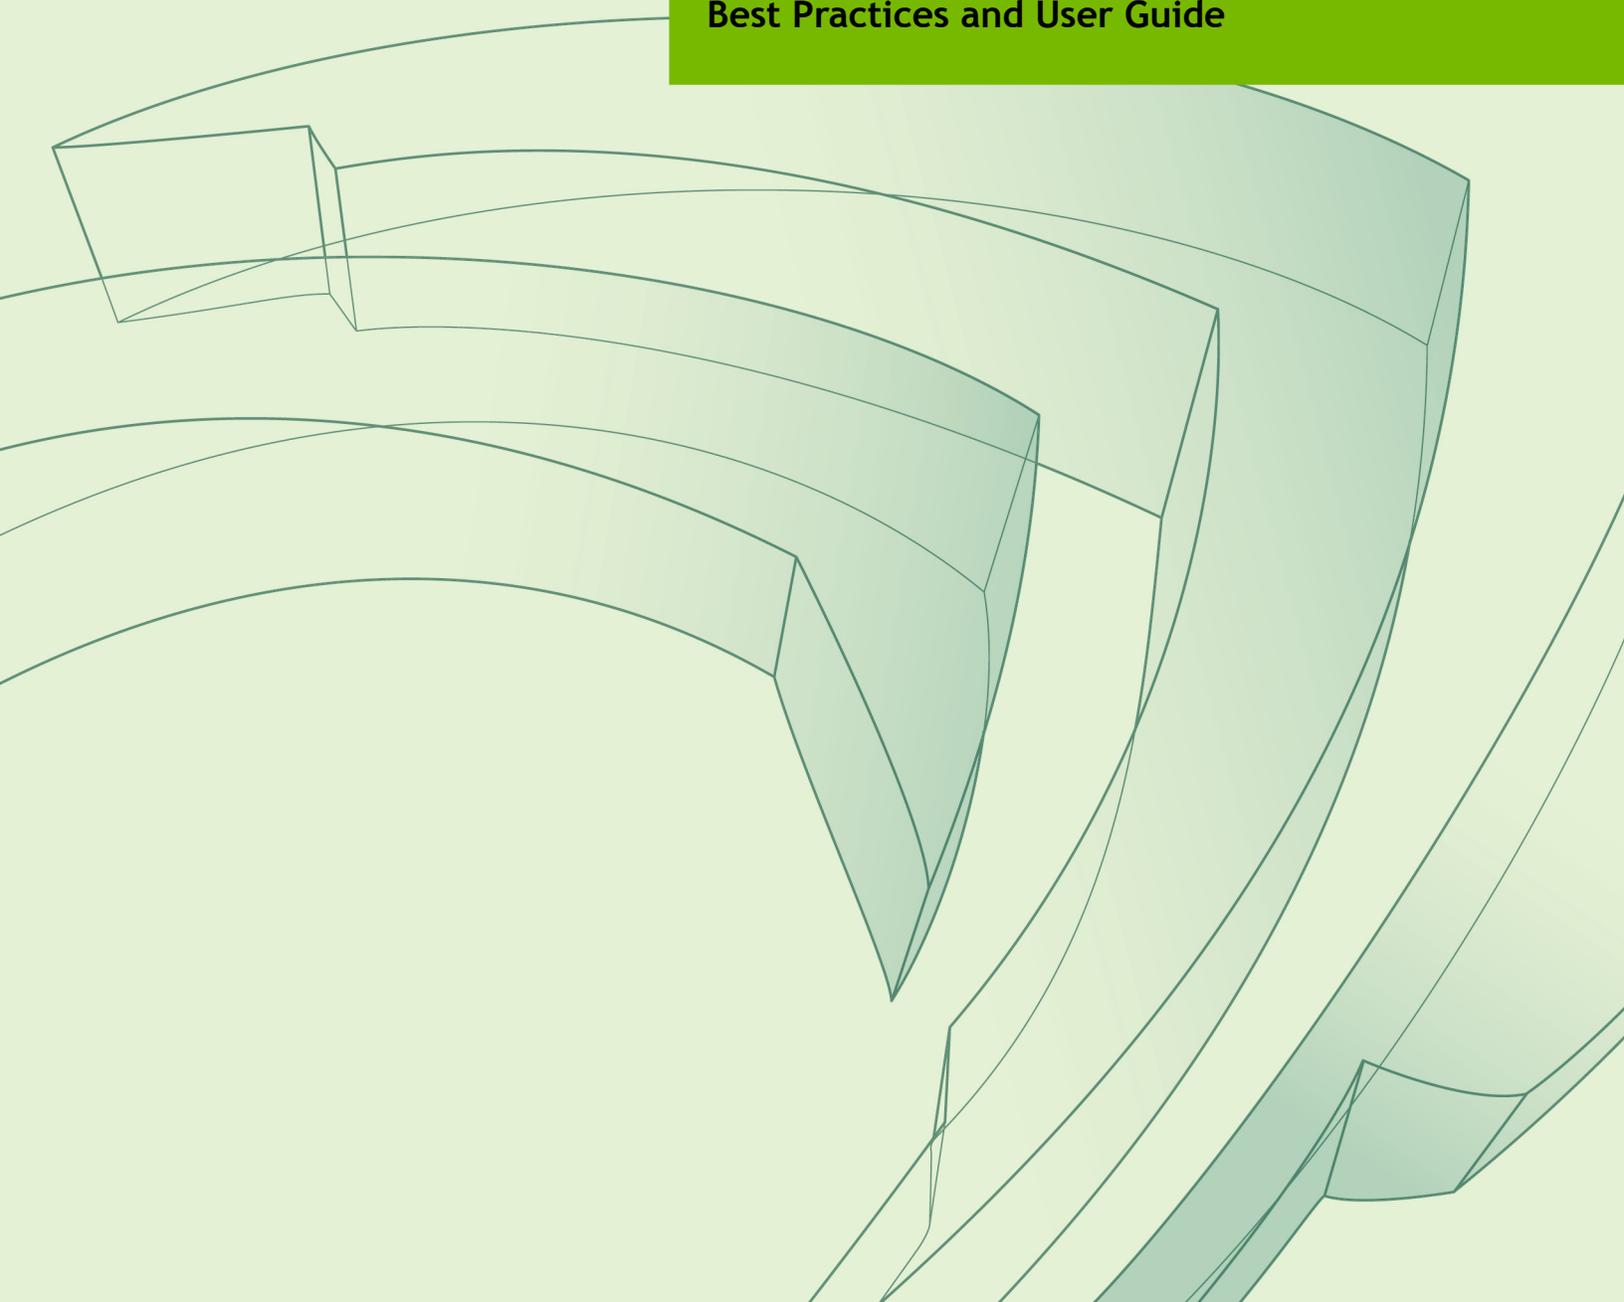




# NVIDIA VALIDATION SUITE

vR352 | May 2015

**Best Practices and User Guide**



# TABLE OF CONTENTS

|                                                           |           |
|-----------------------------------------------------------|-----------|
| <b>Chapter 1. Overview.....</b>                           | <b>1</b>  |
| 1.1. NVVS Goals.....                                      | 1         |
| 1.2. Beyond the Scope of the NVIDIA Validation Suite..... | 2         |
| 1.3. Dependencies.....                                    | 2         |
| 1.4. Supported Products.....                              | 2         |
| <b>Chapter 2. Using NVVS.....</b>                         | <b>3</b>  |
| 2.1. Command line options.....                            | 3         |
| 2.2. Usage Examples.....                                  | 4         |
| 2.3. Configuration file.....                              | 5         |
| 2.4. Global parameters.....                               | 6         |
| 2.5. GPU parameters.....                                  | 6         |
| 2.6. Test parameters.....                                 | 7         |
| <b>Chapter 3. Overview of Plugins.....</b>                | <b>9</b>  |
| 3.1. Deployment Plugin.....                               | 9         |
| 3.2. GPU Bandwidth Plugin.....                            | 10        |
| <b>Chapter 4. Test Output.....</b>                        | <b>17</b> |
| 4.1. JSON Output.....                                     | 17        |

# Chapter 1.

## OVERVIEW

The NVIDIA Validation Suite (NVVS) is the system administrator and cluster manager's tool for detecting and troubleshooting common problems affecting NVIDIA® Tesla™ GPUs in a high performance computing environments. NVVS focuses on software and system configuration issues, diagnostics, topological concerns, and relative performance.

### 1.1. NVVS Goals

The NVIDIA Validation Suite is designed to:

1. Provide a system-level tool, in production environments, to assess cluster readiness levels before a workload is deployed.
2. Facilitate multiple run modes:
  - ▶ Interactive via an administrator or user in plain text.
  - ▶ Scripted via another tool with easily parseable output.
3. Provide multiple test timeframes to facilitate different preparedness or failure conditions:
  - ▶ Quick tests to use as a readiness metric
  - ▶ Medium tests to use as an epilogue on failure
  - ▶ Long tests to be run by an administrator as post-mortem
4. Integrate the following concepts into a single tool to discover deployment, system software and hardware configuration issues, basic diagnostics, integration issues, and relative system performance.
  - ▶ Deployment and Software Issues
    - ▶ NVML library access and versioning
    - ▶ CUDA library access and versioning
    - ▶ Software conflicts
  - ▶ Hardware Issues and Diagnostics
    - ▶ PCIe interface checks
    - ▶ Framebuffer and memory checks
    - ▶ Compute engine checks

- ▶ Integration Issues
  - ▶ PCIe replay counter checks
  - ▶ Topological limitations
  - ▶ Permissions, driver, and cgroups checks
  - ▶ Basic power and thermal constraint checks
- 5. Provide troubleshooting help
- 6. Easily integrate into *Cluster Scheduler* and *Cluster Management* applications
- 7. Reduce downtime and failed GPU jobs

## 1.2. Beyond the Scope of the NVIDIA Validation Suite

NVVS is not designed to:

1. Provide comprehensive hardware diagnostics
2. Actively fix problems
3. Replace the field diagnosis tools. Please refer to <http://docs.nvidia.com/deploy/hw-field-diag/index.html> for that process.
4. Facilitate any RMA process. Please refer to <http://docs.nvidia.com/deploy/rma-process/index.html> for those procedures.

## 1.3. Dependencies

- ▶ This version of the NVIDIA Validation Suite is designed to support the target environments supported by the GPU Deployment Kit (GDK) from whence it came. Please refer to the GDK README for version information. The GDK and supporting documentation can be located at <https://developer.nvidia.com/gpu-deployment-kit>.
- ▶ NVVS requires the standard C++ runtime library with GLIBCXX of at least version 3.4.5 or greater.

## 1.4. Supported Products

The NVIDIA Validation Suite supports Tesla GPUs running on 64-bit Linux (bare metal) operating systems. NVIDIA® Tesla™ Line:

- ▶ All Kepler architecture GPUs
- ▶ All Maxwell architecture GPUs

# Chapter 2.

## USING NVVS

The various command line options of NVVS are designed to control general execution parameters, whereas detailed changes to execution behavior are contained within the configuration files detailed in the next chapter.

### 2.1. Command line options

The various options for NVVS are as follows:

| Short option | Long option           | Description                                                                                                                                                |
|--------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>-a</i>    | <i>--appendLog</i>    | When generating a debug logfile, do not overwrite the contents of a current log. Used in conjunction with the <i>-d</i> and <i>-l</i> options.             |
| <i>-c</i>    | <i>--config</i>       | Specify the configuration file to be used. The default is <code>/etc/nvidia-validation-suite/nvvs.conf</code>                                              |
|              | <i>--configless</i>   | Run NVVS in a configless mode. Executes a "long" test on all supported GPUs.                                                                               |
| <i>-d</i>    | <i>--debugLevel</i>   | Specify the debug level for the output log. The range is 0 to 5 with 5 being the most verbose. Used in conjunction with the <i>-l</i> flag.                |
| <i>-g</i>    | <i>--listGpus</i>     | List the GPUs available and exit. This will only list GPUs that are supported by NVVS.                                                                     |
| <i>-l</i>    | <i>--debugLogFile</i> | Specify the logfile for debug information. This will produce an encrypted log file intended to be returned to NVIDIA for post-run analysis after an error. |
|              | <i>--quiet</i>        | No console output given. See logs and return code for errors.                                                                                              |

| Short option    | Long option                  | Description                                                                                                                                                  |
|-----------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-p</code> | <code>--pluginpath</code>    | Specify a custom path for the NVVS plugins.                                                                                                                  |
| <code>-s</code> | <code>--scriptable</code>    | Produce output in a colon-separated, more script-friendly and parseable format.                                                                              |
|                 | <code>--specifiedtest</code> | Run a specific test in a configless mode. Multiple word tests should be in quotes.                                                                           |
|                 | <code>--statsonfail</code>   | Output statistic logs only if a test failure is encountered.                                                                                                 |
| <code>-t</code> | <code>--listTests</code>     | List the tests available to be executed through NVVS and exit. This will list only the readily loadable tests given the current path and library conditions. |
| <code>-v</code> | <code>--verbose</code>       | Enable verbose reporting.                                                                                                                                    |
|                 | <code>--version</code>       | Displays the version information and exits.                                                                                                                  |
| <code>-h</code> | <code>--help</code>          | Display usage information and exit.                                                                                                                          |

## 2.2. Usage Examples

To display the list of GPUs available on the system.

```
user@hostname
$ nvvs -g

NVIDIA Validation Suite (version 352.00)

Supported GPUs available:
[0000:01:00.0] -- Tesla K40c
[0000:05:00.0] -- Tesla K20c
[0000:06:00.0] -- Tesla K20c
```

An example "quick" test (explained later) using a custom configuration file.

```
user@hostname
$ nvvs -c Tesla_K40c_quick.conf

NVIDIA Validation Suite (version 352.00)

Software
Blacklist ..... PASS
NVML Library ..... PASS
CUDA Main Library ..... PASS
CUDA Toolkit Libraries ..... PASS
Permissions and OS-related Blocks ..... PASS
Persistence Mode ..... PASS
Environmental Variables ..... PASS
```

To output an encrypted debug file at the highest debug level to send to NVIDIA for analysis after a problem.

```

user@hostname
$ nvvs -c Tesla_K40c_medium.conf -d 5 -l debug.log

NVIDIA Validation Suite (version 352.00)

Software
  Blacklist ..... PASS
  NVML Library ..... PASS
  CUDA Main Library ..... PASS
  CUDA Toolkit Libraries ..... PASS
  Permissions and OS-related Blocks ..... PASS
  Persistence Mode ..... PASS
  Environmental Variables ..... PASS
Hardware
  Memory GPU0 ..... PASS
Integration
  PCIe ..... FAIL
    *** GPU 0 is running at PCI link width 8X, which is below the minimum
    allowed link width of 16X (parameter:
min_pci_width)"

```

The output file, debug.log would then be returned to NVIDIA.

## 2.3. Configuration file

The NVVS configuration file is a [YAML](#)-formatted (e.g. human-readable JSON) text file with three main stanzas controlling the various tests and their execution.

The general format of a configuration file consists of:

```

%YAML 1.2
---

globals:
  key1: value
  key2: value

test_suite_name:
- test_class_name1:
  test_name1:
    key1: value
    key2: value
    subtests:
      subtest_name1:
        key1: value
        key2: value
  test_name2:
    key1: value
    key2: value
- test_class_name2:
  test_name3:
    key1: value
    key2: value

gpus:
- gpuset: name
  properties:
    key1: value

```

```

key2: value
tests:
  name: test_suite_name

```

There are three distinct sections: *globals*, *test\_suite\_name*, and *gpus* each with its own subsection of parameters and as is with any YAML document, **indentation is important** thus if errors are generated from your own configuration files please refer to this example for indentation reference.

## 2.4. Global parameters

| Keyword                  | Value Type | Description                                                                                                                                                                                                                  |
|--------------------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| logfile                  | String     | The prefix for all detailed test data able to be used for post-processing.                                                                                                                                                   |
| logfile_type             | String     | Can be <i>json</i> , <i>text</i> , or <i>binary</i> . Used in conjunction with the logfile global parameter. Default is JSON.                                                                                                |
| scriptable               | Boolean    | Accepts <i>true</i> , or <i>false</i> . Produces a script-friendly, colon-separated output and is identical to the <i>-s</i> command line parameter.                                                                         |
| serial_override          | Boolean    | Accepts <i>true</i> , or <i>false</i> . Some tests are designed to run in parallel if multiple GPUs are given. This parameter overrides that behavior serializing execution across all tests.                                |
| require_persistence_mode | Boolean    | Accepts <i>true</i> , or <i>false</i> . Persistence mode is a prerequisite for some tests, this global overrides that requirement and should only be used if it is not possible to activate persistence mode on your system. |

## 2.5. GPU parameters

The *gpus* stanza may consist of one or more *gpuset*s which will each match zero or more GPUs on the system based on their *properties*(a match of zero will produce an error).

GPUs are matched based on the following criteria with their configuration file keywords in parenthesis:

- ▶ Name of the GPU, i.e. Tesla K40c (*name*)
- ▶ Brand of the GPU, i.e. Tesla (*brand*)
- ▶ A comma separated list of indexes (*index*)
- ▶ The GPU UUID (*uuid*)
- ▶ or the PCIe Bus ID (*busid*)

The matching rules are based off of exclusion. First, the list of supported GPUs is taken and if no *properties* tag is given then all GPUs will be used in the test. Because a UUID or PCIe Bus ID can only match a single GPU, if those properties are given then only that GPU will be used if found. The remaining properties, *index*, *brand*, and *name* work in an "AND" fashion such that, if specified, the result must match at least one GPU on the system for a test to be performed.

For example, if *name* is set to "Tesla K40c" and *index* is set to "0" NVVS will error if index 0 is not a Tesla K40c. By specifying both *brand* and *index* a user may limit a test to specific "Tesla" cards for example. **In this version of NVVS, all matching GPUs must be homogeneous.**

The second identifier for a *gpuset* is *tests*. This parameter specifies either the suite of tests that a user wishes to run or the test itself.

At present the following suites are available:

- ▶ *Quick* -- meant as a pre-run sanity check to ensure that the GPUs are ready for a job. Currently runs the Deployment tests described in the next chapter.
- ▶ *Medium* -- meant as a quick, post-error check to make sure that nothing very obvious such as ECC enablement or double-bit errors have occurred on a GPU. Currently runs the Deployment, Memory, and PCIe/Bandwidth tests.

An individual test can also be specified. Currently the keywords are: *Memory*, *PCIe*, and *PCIe*. Please see the "custom" section in the next subchapter to configure and tweak the parameters when this method is used.

## 2.6. Test parameters

The format of the NVVS configuration file is designed for extensibility. Each test suite above can be customized in a number of ways described in detail in the following chapter for each test. Individual tests belong to a specific class of functionality which, when wanting to customize specific parameters, must also be specified.

The classes and the respective tests they perform are as follows:

| Class name  | Tests      | Brief description                                                                                                      |
|-------------|------------|------------------------------------------------------------------------------------------------------------------------|
| Software    | Deployment | Checks for various runtime libraries, persistence mode, permissions, environmental variables, and blacklisted drivers. |
| Hardware    | Memory     | Allocate as much memory as possible and perform read/write operations.                                                 |
| Integration | PCIe       | Test host to GPU, GPU to host, and P2P (if possible) bandwidth.                                                        |

Some tests also have subtests that can be enabled by using the *subtests* keyword and then hierarchically adding the subtest parameters desired beneath. An example would be the PCIe Bandwidth test which may have a section that looks similar to this:

```
long:  
- integration:  
  pcie:  
    test_unpinned: false  
    subtests:  
      h2d_d2h_single_pinned:  
        min_bandwidth: 20  
        min_pci_width: 16
```

When only a specific test is given in the GPU set portion of the configuration file, both the suite and class of the test are *custom*. For example:

```
%YAML 1.2  
---  
  
globals:  
  logfile: nvvs.log  
  
custom:  
- custom:  
  targeted performance:  
    test_duration: 60  
  
gpus:  
- gpuset: all_K40c  
  properties:  
    name: Tesla K40c  
  tests:  
    - name: targeted performance
```

# Chapter 3.

## OVERVIEW OF PLUGINS

The NVIDIA Validation Suite consists of a series of plugins that are each designed to accomplish a different goal.

### 3.1. Deployment Plugin

The deployment plugin's purpose is to verify the compute environment is ready to run Cuda applications and is able to load the NVML library.

#### Preconditions

- ▶ `LD_LIBRARY_PATH` must include the path to the cuda libraries, which for version X.Y of Cuda is normally `/usr/local/cuda-X.Y/lib64`, which can be set by running `export LD_LIBRARY_PATH=/usr/local/cuda-X.Y/lib64`
- ▶ The linux nouveau driver must not be running, and should be blacklisted since it will conflict with the nvidia driver

#### Configuration Parameters

*None at this time.*

#### Stat Outputs

*None at this time.*

#### Failure

The plugin will fail if:

- ▶ The corresponding device nodes for the target GPU(s) are being blocked by the operating system (e.g. cgroups) or exist without r/w permissions for the current user.
- ▶ The NVML library `libnvidia-ml.so` cannot be loaded
- ▶ The Cuda runtime libraries cannot be loaded

- ▶ The **nouveau** driver is found to be loaded
- ▶ Any pages are pending retirement on the target GPU(s)

## 3.2. GPU Bandwidth Plugin

The GPU bandwidth plugin's purpose is to measure the bandwidth and latency to and from the GPUs and the host.

### Preconditions

- ▶ This plugin must be run as root or be granted the ability to set application clocks on supporting hardware (K10 or higher). To allow non-root to change application clocks on device *deviceIndex*, as root, first enable persistence mode with `nvidia-smi -i deviceIndex -pm 1` and then allow non-root to change application clocks with `nvidia-smi -i deviceIndex -acp 0`
- ▶ If the GPUs the plugin is being run on are K80 or higher, the plugin must be run as root or be granted the ability to disable auto boosted clocks. Non-root is allowed to set auto boost by default on K80 or higher. To enable non-root setting auto boost on device *deviceIndex*, as root, first enable persistence mode with `nvidia-smi -i deviceIndex -pm 1` and then allow non-root to change auto boosted clocks with `nvidia-smi -i deviceIndex --auto-boost-permission=0`

### Sub Tests

The plugin consists of several self-tests that each measure a different aspect of bandwidth or latency. Each subtest has either a pinned/unpinned pair or a p2p enabled/p2p disabled pair of identical tests. Pinned/unpinned tests use either pinned or unpinned memory when copying data between the host and the GPUs.

Each sub test is represented with a tag that is used both for specifying configuration parameters for the sub test and for outputting stats for the sub test. P2p enabled/p2p disabled tests enable or disable GPUs on the same card talking to each other directly rather than through the PCIe bus.

| Sub Test Tag              | Pinned/Unpinned          | Description                                      |
|---------------------------|--------------------------|--------------------------------------------------|
|                           | P2P Enabled/P2P Disabled |                                                  |
| h2d_d2h_single_pinned     | Pinned                   | Device <-> Host Bandwidth, one GPU at a time     |
| h2d_d2h_single_unpinned   | Unpinned                 | Device <-> Host Bandwidth, one GPU at a time     |
| h2d_d2h_concurrent_pinned | Pinned                   | Device <-> Host Bandwidth, all GPUs concurrently |

| Sub Test Tag                   | Pinned/Unpinned          | Description                                                                                                                                                                                                     |
|--------------------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                | P2P Enabled/P2P Disabled |                                                                                                                                                                                                                 |
| h2d_d2h_concurrent_unpinned    | Unpinned                 | Device <-> Host Bandwidth, all GPUs concurrently                                                                                                                                                                |
| h2d_d2h_latency_pinned         | Pinned                   | Device <-> Host Latency, one GPU at a time                                                                                                                                                                      |
| h2d_d2h_latency_unpinned       | Unpinned                 | Device <-> Host Latency, one GPU at a time                                                                                                                                                                      |
| p2p_bw_p2p_enabled             | P2P Enabled              | Device <-> Device bandwidth one GPU pair at a time                                                                                                                                                              |
| p2p_bw_p2p_disabled            | P2P Disabled             | Device <-> Device bandwidth one GPU pair at a time                                                                                                                                                              |
| p2p_bw_concurrent_p2p_enabled  | P2P Enabled              | Device <-> Device bandwidth, concurrently, focusing on bandwidth between GPUs between GPUs likely to be directly connected to each other -> for each (index / 2) and (index / 2)+1                              |
| p2p_bw_concurrent_p2p_disabled | P2P Disabled             | Device <-> Device bandwidth, concurrently, focusing on bandwidth between GPUs between GPUs likely to be directly connected to each other -> for each (index / 2) and (index / 2)+1                              |
| 1d_exch_bw_p2p_enabled         | P2P Enabled              | Device <-> Device bandwidth, concurrently, focusing on bandwidth between gpus, every GPU either sending to the gpu with the index higher than itself (l2r) or to the gpu with the index lower than itself (r2l) |
| 1d_exch_bw_p2p_disabled        | P2P Disabled             | Device <-> Device bandwidth, concurrently, focusing on bandwidth between gpus, every GPU either sending to the gpu with the index higher than itself                                                            |

| Sub Test Tag             | Pinned/Unpinned          | Description                                                |
|--------------------------|--------------------------|------------------------------------------------------------|
|                          | P2P Enabled/P2P Disabled |                                                            |
|                          |                          | (l2r) or to the gpu with the index lower than itself (r2l) |
| p2p_latency_p2p_enabled  | P2P Enabled              | Device <-> Device Latency, one GPU pair at a time          |
| p2p_latency_p2p_disabled | P2P Disabled             | Device <-> Device Latency, one GPU pair at a time          |

### Configuration Parameters- Global

| Parameter Name   | Type  | Default | Value Range    | Description                                                                                                                                               |
|------------------|-------|---------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| test_pinned      | Bool  | True    | True/False     | Include subtests that test using pinned memory.                                                                                                           |
| test_unpinned    | Bool  | True    | True/False     | Include subtests that test using unpinned memory.                                                                                                         |
| test_p2p_on      | Bool  | True    | True/False     | Include subtests that require peer to peer (P2P) memory transfers between cards to occur.                                                                 |
| test_p2p_off     | Bool  | True    | True/False     | Include subtests that do not require peer to peer (P2P) memory transfers between cards to occur.                                                          |
| max_pcie_replays | Float | 80.0    | 1.0 - 100000.0 | Maximum number of PCIe replays to allow per GPU for the duration of this plugin. This is based on an expected replay rate <8 per minute for PCIe Gen 3.0, |

| Parameter Name | Type | Default | Value Range | Description                                                                                           |
|----------------|------|---------|-------------|-------------------------------------------------------------------------------------------------------|
|                |      |         |             | assuming this plugin will run for less than a minute and allowing 10x as many replays before failure. |

### Configuration Parameters- Sub Test

| Parameter Name     | Default (Range)                  | Affected Sub Tests                                                                                              | Description                                                                                                     |
|--------------------|----------------------------------|-----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| min_bandwidth      | Null<br>(0.0 - 100.0)            | h2d_d2h_single_pinned,<br>h2d_d2h_single_unpinned,<br>h2d_d2h_concurrent_pinned,<br>h2d_d2h_concurrent_unpinned | Minimum bandwidth in GB/s that must be reached for this sub-test to pass.                                       |
| max_latency        | 100,000.0<br>(0.0 - 1,000,000.0) | h2d_d2h_latency_pinned,<br>h2d_d2h_latency_unpinned                                                             | Latency in microseconds that cannot be exceeded for this sub-test to pass.                                      |
| min_pci_generation | 1.0<br>(1.0 - 3.0)               | h2d_d2h_single_pinned,<br>h2d_d2h_single_unpinned                                                               | Minimum allowed PCI generation that the GPU must be at or exceed for this sub-test to pass.                     |
| min_pci_width      | 1.0<br>(1.0 - 16.0)              | h2d_d2h_single_pinned,<br>h2d_d2h_single_unpinned                                                               | Minimum allowed PCI width that the GPU must be at or exceed for this sub-test to pass. For example, 16x = 16.0. |

### Stat Outputs - Global

| Stat Name         | Stat Scope | Type  | Description                                                                                            |
|-------------------|------------|-------|--------------------------------------------------------------------------------------------------------|
| pcie_replay_count | GPU        | Float | The per second reading of PCIe replays that have occurred since the start of the GPU Bandwidth plugin. |

### Stat Outputs -Sub Test

Stats for the GPU Bandwidth test are also output on a test by test basis, using the sub test name as the group name key. The following stats sections are organized by sub test.

#### h2d\_d2h\_single\_pinned/h2d\_d2h\_single\_unpinned

| Stat Name      | Type  | Description                                                                                   |
|----------------|-------|-----------------------------------------------------------------------------------------------|
| <i>N_h2d</i>   | Float | Average bandwidth from host to device for device <i>N</i>                                     |
| <i>N_d2h</i>   | Float | Average bandwidth from device to host for device <i>N</i>                                     |
| <i>N_bidir</i> | Float | Average bandwidth from device to host and host to device at the same time for device <i>N</i> |

#### h2d d2h concurrent pinned/h2d d2h concurrent unpinned

| Stat Name      | Type  | Description                                                                                   |
|----------------|-------|-----------------------------------------------------------------------------------------------|
| <i>N_h2d</i>   | Float | Average bandwidth from host to device for device <i>N</i>                                     |
| <i>N_d2h</i>   | Float | Average bandwidth from device to host for device <i>N</i>                                     |
| <i>N_bidir</i> | Float | Average bandwidth from device to host and host to device at the same time for device <i>N</i> |
| sum_bidir      | Float | Sum of the average bandwidth from device to host and host to device for all devices.          |
| sum_h2d        | Float | Sum of the average bandwidth from host to device for all devices.                             |
| sum_d2h        | Float | Sum of the average bandwidth from device to host for all devices.                             |

#### h2d d2h latency pinned/h2d d2h latency unpinned

| Stat Name    | Type  | Description                                             |
|--------------|-------|---------------------------------------------------------|
| <i>N_h2d</i> | Float | Average latency from host to device for device <i>N</i> |
| <i>N_d2h</i> | Float | Average latency from device to host for device <i>N</i> |

| Stat Name       | Type  | Description                                                                                 |
|-----------------|-------|---------------------------------------------------------------------------------------------|
| <i>N</i> _bidir | Float | Average latency from device to host and host to device at the same time for device <i>N</i> |

p2p\_bw\_p2p\_enabled/p2p\_bw\_p2p\_disabled

| Stat Name          | Type  | Description                                                                                           |
|--------------------|-------|-------------------------------------------------------------------------------------------------------|
| <i>N</i> _M_onedir | Float | Average bandwidth from device <i>N</i> to device <i>M</i> , copying one direction at a time.          |
| <i>N</i> _M_bidir  | Float | Average bandwidth from device <i>N</i> to device <i>M</i> , copying both directions at the same time. |

p2p\_bw\_concurrent\_p2p\_enabled/p2p\_bw\_concurrent\_p2p\_disabled

| Stat Name          | Type  | Description                                                                                                     |
|--------------------|-------|-----------------------------------------------------------------------------------------------------------------|
| l2r_ <i>N</i> _M   | Float | Average bandwidth from device <i>N</i> to device <i>M</i>                                                       |
| r2l_ <i>N</i> _M   | Float | Average bandwidth from device <i>M</i> to device <i>N</i>                                                       |
| bidir_ <i>N</i> _M | Float | Average bandwidth from device <i>M</i> to device <i>N</i> , copying concurrently                                |
| r2l_sum            | Float | Sum of average bandwidth for all right ( <i>M</i> ) to left ( <i>N</i> ) copies                                 |
| r2l_sum            | Float | Sum of average bidirectional bandwidth for all right ( <i>M</i> ) to left ( <i>N</i> ) and left to right copies |

1d\_exch\_bw\_p2p\_enabled/1d\_exch\_bw\_p2p\_disabled

| Stat Name     | Type  | Description                                                 |
|---------------|-------|-------------------------------------------------------------|
| l2r_ <i>N</i> | Float | Average bandwidth from device <i>N</i> to device <i>N+1</i> |
| r2l_ <i>N</i> | Float | Average bandwidth from device <i>N</i> to device <i>N-1</i> |

| Stat Name | Type  | Description                            |
|-----------|-------|----------------------------------------|
| l2r_sum   | Float | Sum of all l2r average bandwidth stats |
| r2l_sum   | Float | Sum of all l2r average bandwidth stats |

### p2p\_latency\_p2p\_enabled/p2p\_latency\_p2p\_disabled

| Stat Name  | Type  | Description                                             |
|------------|-------|---------------------------------------------------------|
| <i>N_M</i> | Float | Average latency from device <i>N</i> to device <i>M</i> |

### Failure

The plugin will fail if:

- ▶ The latency exceeds the configured threshold for relevant tests.
- ▶ The bandwidth cannot exceed the configured threshold for relevant tests.
- ▶ If the number of PCIe retransmits exceeds a user-provided threshold.

# Chapter 4.

## TEST OUTPUT

The output of tests can be collected by setting the "logfile" global parameter which represents the prefix for the detailed outputs produced by each test. The default type of output is JSON but text and binary outputs are available as well. The latter two are meant more for parsing and direct reading by custom consumers respectively so this portion of the document will focus on the JSON output.

### 4.1. JSON Output

The JSON output format is keyed based off of the "stats" keys given in each test overview from Chapter 3. These standard JSON files can be processed in any number of ways but two example Python scripts have been provided to aid in visualization in the default installation directory.. The first is a JSON to comma-separated value script (json2csv.py) which can be used to import key values in to a graphing spreadsheet. Proper usage would be:

```
user@hostname
$ python json2csv.py -i stats_targeted_performance.json -o stats.csv -k
  gpu_temperature,power_usage
```

Also provided is an example Python script that uses the [pygal](#) library to generate readily viewable scalar vector graphics charts (json2svg.py), able to be opened in any browser. Proper usage would be:

```
user@hostname
$ python json2svg.py -i stats_targeted_performance.json -o stats.svg -k
  gpu_temperature,power_usage
```

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© 2014-2015 NVIDIA Corporation. All rights reserved.