



NVIDIA Mosaic Technology

User's Guide



Table of Contents

Mosaic Technology.....	3
Terminology.....	3
System Requirements	6
Mosaic Technology.....	7
Connect Your Displays	7
NVIDIA Control Panel.....	7
Configuring Mosaic	9
Configure Display and 3D Settings Prior to Enabling Mosaic.....	9
Setting up Mosaic with the NVIDIA control panel.....	10
Create Mosaic Configuration	11
Select Topology.....	11
Select Displays	13
Arrange Displays	14
Adjust Overlap and Bezel Correction	15
Setting up Mosaic with the Command Line Utility	16
Troubleshooting and Notes	19
Hot Plug Persistence and Mosaic	19
Multi-Display Configuration Upon Disabling Mosaic	19
Display Resolutions Upon Disabling Mosaic	19
Use of HDMI or Display Port Display Adapters.....	19
Appendix A – Restrictions When Using Mosaic	20
Matching Displays	20
Maximum Number of Displays in Single System	20
Maximum Resolution	20
Minimum Resolution	20
Multiple Mosaic GRID.....	20
Non-Mosaic Displays.....	20



Mosaic Technology

NVIDIA® Mosaic™ Technology enables an easy way to scale up multiple high-resolution monitors in perfect synchronization. It is supported on NVIDIA RTX™ and NVIDIA Quadro solutions and presents multiple displays as a single logical GPU display to the operating system and applications.

Terminology

Control Display

Often referred to as the +1 display, Control Display is an individual display that is not defined as part of the Mosaic and must be connected to a GPU that is not defined as part of the Mosaic. Depending on your requirements, it may or may not need to be synchronized with the Mosaic GRID.

External Sync

External Sync refers to a sync signal that is generated external to the system and is used as a reference to synchronize the frame update of displays attached to the system.

Framelock

Framelock refers to when multiple displays are synchronized so that the start and end of the frame aligns for all displays.

LDA mode

LDA mode is an acronym referring to the Linked Display Adapter mode which is a feature of the Windows operating system where multiple GPUs are placed into a specific configuration. When in LDA mode GPUs are able to use more efficient interconnect to other GPUs in the LDA group. With modern explicit APIs such as Vulkan or DirectX 12 applications can individually address all GPUs in the LDA group as well as share data between any number of GPUs within the group.

Mosaic

Mosaic creates one large desktop of monitors and appears as a single Desktop to the operating system. Taskbars will span across the monitors and applications will maximize across the Mosaic configuration.

Mosaic GRID

A GRID refers to a set of displays in a Mosaic Group and usually defined as a set of rows and columns arranging the displays.

For example, a 2x4 GRID will have two rows and four columns of and a 3x3 GRID will have three rows and three columns.

Mosaic Group



A Mosaic Group refers to a set of GPUs configured in Mosaic mode. The GPUs in a Mosaic Group may have displays attached or not depending on configuration. When a Mosaic Group contains GPUs that do not have attached displays, it is usually referred to as the *Maximum GPU Topology*. When a Mosaic Group contains only GPUs that have displays attached, it is referred to as the “*Minimum GPU Topology*”.

NVLINK

NVLink is a NVIDIA technology that provides a high-bandwidth, low latency interconnect between GPUs.

Present Barrier

Conceptually, a Present Barrier refers to a combination of *swapgroup* and a *swapbarrier*, meaning that a present barrier refers to a mechanism where multiple displays on a single system or multiple systems update the displayed frame at exactly the same time. The term Present Barrier is also used to define an API interface within a graphics API, such as Vulkan or DirectX 12.

Rasterlock

Rasterlock refers to when multiple displays are synchronized so that each pixel on each display is aligned to the same pixel and all other displays.

SDI

SDI is an acronym for *Serial Digital Interface* which refers to a *Society of Motion Picture and Television Engineers (SMPTE)* standard that describes how to transport video in a digitally encoded format.

SLI

SLI is an acronym referring to an NVIDIA technology called *Scalable Link Interface* that allows multiple GPUs to be linked together to improve rendering performance. SLI implements multiple algorithms to improve performance, such as *Alternate Frame Rendering (AFR)* or *Split Frame Rendering (SFR)*.

- Explicit SLI

Graphics APIs have evolved to give applications significantly more control over GPU resources and so Vulkan and DirectX 12 are often referred to as *explicit* APIs, meaning that an application explicitly controls the GPU instead of the driver. This explicit application control also applies to SLI where with Vulkan and DirectX 12, applications can implement whatever scaling algorithm is best for their needs. This is referred to as *explicit SLI*.

- Implicit SLI

The SLI SFR and AFR algorithms were a GPU driver feature and transparent to applications, meaning that applications would benefit from improved rendering performance without modification. This is called *implicit SLI*.

Smart Clone

View large mosaic on a larger monitor as a control display with pan and scan or scaling. It cannot be on any of the GPUs that are participating in the Mosaic.

Swap Barrier

A *swapbarrier* refers to a mechanism where multiple displays or *swapgroups* on multiple systems are synchronized so that all displays across all systems update the displayed frame at exactly the same time. The term *swapbarrier* is also used to define an API interface within a graphics API, such as OpenGL.

Swap Group

A *swapgroup* refers to a group of displays in a single system where all displays in the group will update the displayed frame at exactly the same time. The term *swapgroup* is also used to define an API interface within a graphics API, such as OpenGL.

SwapReady

Swapready is a hardware signal created by the GPU used to synchronize a group of GPUs so that all the GPUs in the group update the display image at exactly the same time, regardless of how long each GPU took to render.

Synchronized Mosaic

Creates one large desktop of monitors and appears as a single desktop to the operating system with all of the monitors synchronized to the same timing. This will ensure no “tearing” across monitors on the setup.

- For two GPU configurations with NVLINK, GPUs can share the sync timing signal needed to synchronize Mosaic across the two GPUs. This helps eliminate “tearing” when using two GPUs in the Mosaic.
- For two plus GPU configurations, a Quadro Sync II card is needed to synchronize Mosaic across all the GPUs.

Bi-Level & Tri-level Frame Sync

Bi-level and Tri-level sync refers to industry-standard sync types for Standard Definition (SD) and High Definition (HD) digital video, respectively, and represents the start of a digital video frame. Both bi-level and tri-level syncs are often referred to as Composite sync meaning that the sync and video data are transmitted on the same video signal.

System Requirements

Table 1.1 Supported Desktop Workstation Products

Desktop Workstation GPU	Mosaic with sync-capability
RTX 6000 Ada Generation / RTX 4000 SFF Ada Generation / RTX A4000 / Quadro RTX 4000	• One GPU • Up to eight GPUs with up to two Quadro Sync II cards
RTX A6000 / RTX A5000 / Quadro RTX 8000 / Quadro RTX 6000 / Quadro RTX 5000	• One GPU • Two GPUs with NVLink Bridge • Up to eight GPUs with up to two Quadro Sync II cards

Installing a single Quadro Sync II board enables up to four GPUs to be synchronized or frame locked. To scale to up to eight GPUs, two Quadro Sync II boards can be installed in the same system. For more information on Quadro Sync, go to:

<https://www.nvidia.com/en-us/design-visualization/solutions/quadro-sync/>

NOTE: RTX 4000 class and above NVIDIA RTX Laptop GPUs support Mosaic depending on the OEM configurations. Please consult OEMs if Mosaic is supported.

Table 1.2 Supported Operating Systems

Operating System
Microsoft Windows 10 and 11
Linux

Mosaic Technology

There are two ways to enable Mosaic. The first is using the NVIDIA Control Panel and the second is using the Mosaic Command Line utility. Both are covered in this User Guide and this section explains how to access each.

Connect Your Displays

Connect all the displays that you intend to use to the NVIDIA RTX GPUs

- If a Mosaic GRID crosses two or more GPUs, then the connectors on these GPUs that are not assigned to the Mosaic configuration are not available for use.
- Within a Mosaic GRID all display input connections must be of the same type (for example: if Display Port to HDMI adapters are used, all connections for displays in the Mosaic GRID should use the same Display Port to HDMI adapter, and the same display ports must be used on each card).

NOTE: NVIDIA recommended display adapters are listed at:

https://nvidia.custhelp.com/app/answers/detail/a_id/4449/~nvidia-recommended-display-adapters

- The visual characteristics of displays vary by manufacturer and model from the same manufacturer. NVIDIA recommends using the same make and model for all displays in a Mosaic GRID.

NOTE: If your display devices do not send an EDID to the GPU, Mosaic will not be set. In this case, you must apply a file-based EDID for all of the displays that are part of the Mosaic configuration.

NVIDIA Control Panel

The NVIDIA Control Panel can be accessed in two ways:

1. Right clicking on the desktop and selecting, as shown in Figure 1.
2. Clicking the Windows Start button and selecting the NVIDIA Control Panel, as shown in Figure 2.




Figure 1. Launching the NVIDIA control panel from right click on the Windows desktop




Figure 2. Launching the NVIDIA control panel from clicking the Windows Start button and selecting the NVIDIA Control Panel

After launching the NVIDIA Control Panel, the Mosaic settings are accessed by the *Set up Mosaic* item in the index list on the left-hand side of the Control Panel, as shown in Figure 3.




Figure 3. NVIDIA Control Panel showing the *Set up Mosaic* function

Configuring Mosaic

There are several settings that should be configured before enabling Mosaic mode.

Configure Display and 3D Settings Prior to Enabling Mosaic

The display settings that you set on that display are applied to all the displays after enabling Mosaic.

Open the NVIDIA Control Panel and configure any display or 3D settings before enabling Mosaic.

Display settings that should be set prior to enabling Mosaic can include but are not limited to:

- Desktop color settings
- Flat panel scaling
- Video and TV settings

3D settings that should be set prior to enabling Mosaic may include but are not limited to:

- Stereo
- Stereo Display Mode

Setting up Mosaic with the NVIDIA control panel

1. From the *Select a Task* navigation pane, under **Workstation**, click **Set Up Mosaic** to open the associated page.




Figure 6. NVIDIA Control Panel Set Up Mosaic Page

2. If not already expanded, click the *Sync Capability Information* icon (+) to expand the information box and see which GPUs in your system are able to synchronize multiple displays with each other.

NOTE: This box appears only if there is at least one sync-capable GPU in your system.

To view this information, move the mouse over the corresponding icon in the Sync Capability column:

- **Sync not supported:** Low end and entry level GPUs are not sync-capable.
- **Single GPU Sync:** Displays on this single GPU can be synchronized with each other, but not across other GPUs.
- **Multi GPU Sync/SLI Bridge:** Displays can be synchronized with other displays across multiple GPUs using an SLI bridge/NVLink.
- **Multi GPU Sync/Quadro Sync:** Displays can be synchronized with other displays across multiple GPUs using a Quadro Sync/Sync II card.

- **Multi GPU Sync/SLI Bridge/Quadro Sync:** Displays can be synchronized with other displays across multiple GPUs using a Quadro Sync/Sync II card as well as SLI bridge /NVLink.
-

Create Mosaic Configuration

Click **Create a new configuration** and then follow the steps in the Mosaic Configuration wizard as described in the following sections.

You can create more than one Mosaic configuration depending on the displays that are connected. You must click **Create a new configuration** and then follow the wizard for each Mosaic configuration that you create.

Select Topology




Figure 7. Select Topology Example

1. Select the number of displays, the row x column topology, and the orientation of the displays.

By default, the number of physically connected displays appears under **Number of displays**. The “Orientation of displays” options are:

- Landscape [default]
 - Portrait [90 deg rotation]

 - Landscape (flipped) [180 deg rotation]
 - Portrait (flipped) [270 deg rotation]
2. Optionally enter a configuration name for this Mosaic configuration.
 3. Check **I am using recommended connections for the selected topology**, if your connections match the appropriate connectivity specified in the recommended connections document, otherwise leave unchecked.
 - If you are using the recommended connections, then click **Enable Mosaic** and you will be directed to stage 2 of the *Arrange displays* tab. If there are problems with your connections, then you will be directed to the *Select displays* tab where you can make corrections.
 - If you are not using recommended connections, then click **Next** to proceed to the *Select displays* tab.
 - Minimum GPU Topology
 - > Only GPUs with displays attached are part of the Mosaic Group.
 - > This is the default option for Mosaic setup and will ignore any GPUs that do not have a display attached.
 - Maximum GPU Topology option:
 - > All GPUs are part of the Mosaic Group.

NOTE: Developers can improve rendering in some applications if they are written to take advantage of this option. GPUs in the Mosaic Group will also be in the same LDA group, so GPU to GPU transfers will be enhanced in this configuration. Most applications today will not see a performance improvement. Minimum GPU Topology is the default configuration.

Select Displays




Figure 8. Select Displays Example

- Select the displays to use, the refresh rate, and the resolution per display.
- Icons for each selected display will appear in the *Selected display sources* box.

Arrange Displays




Figure 9. Arrange Displays Example

- Arrange the display sources to match your physical displays by either dragging the icons from the Selected display source box to the desired position in the Topology box, or by entering the number for each display into the desired position in the Topology box.
- When completed, click **Apply**. Mosaic is activated.
- With Mosaic activated, click Next to proceed to the Adjust overlap and bezel correction tab, or click Finish if no overlap and bezel correction is needed.
- Upon clicking Finish, the wizard closes, and the *Set up Mosaic* page shows your settings.

Adjust Overlap and Bezel Correction




Figure 10. Example Adjust Overlap and Bezel Correction Tab




Figure 11. Bezel and Overlap Correction

You can adjust the horizontal and/or vertical edge gap as needed. For example, you can introduce gaps between edges to accommodate bezeled displays, or you can overlap edges.

NOTE: Edge overlap is available only with sync-capable GPUs.

1. Select whether to use the same setting for all edges, or whether to use separate settings for each edge.
2. To adjust an edge, click the edge from the Select edges box, then make the adjustments on the Overlap/bezel correction values for selected edges section, then click **Apply**.

While bezel correction is often referred to as "negative overlap," the bezel correction pixel values are presented as positive.

3. Click **Finish** when adjustments are completed.


Upon clicking **Finish**, the wizard closes and the *Set up Mosaic* page shows your settings for the new Mosaic configuration.

Setting up Mosaic with the Command Line Utility

The following steps provide instructions on using the ConfigureMosaic.exe utility:

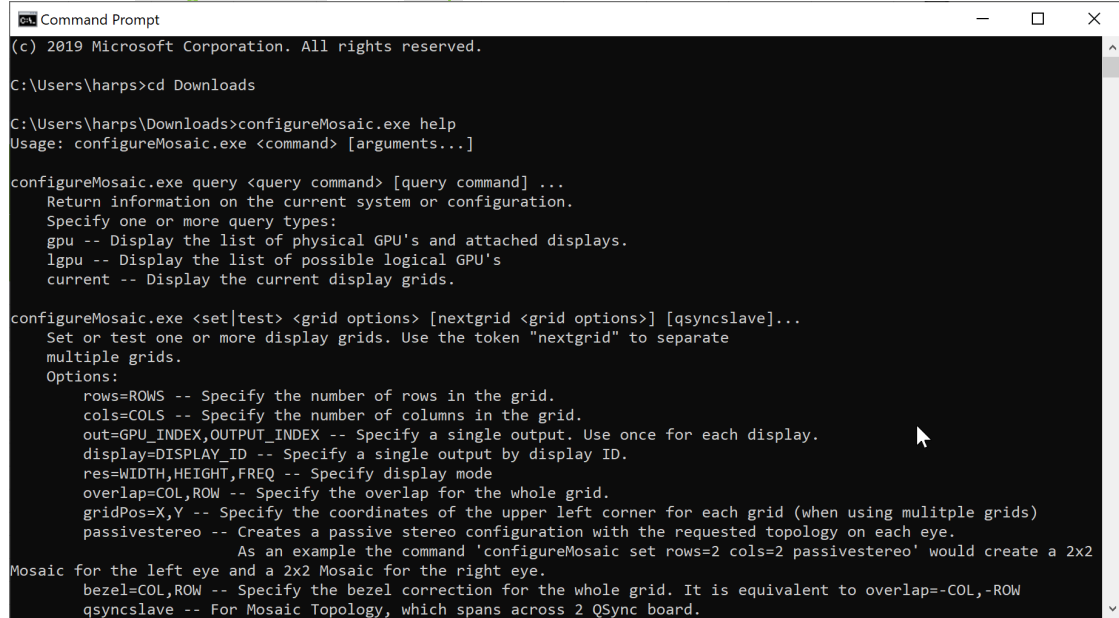
1. Download the Mosaic Utility from the NVIDIA driver homepage on www.nvidia.com.

NOTE: This will be located under the **Quadro Advanced Options** bullet. Under the **Utilities** subheading you will find the **Mosaic Utility** for download



2. Download “configureMosaic.exe,” using the “Download Now” button.
3. The configureMosaic.exe utility must be run in a command window with administrator user privileges.
4. To find a list of commands for the ConfigureMosaic.exe utility run the following command:

C:\PATH\configureMosaic.exe help



```

(c) 2019 Microsoft Corporation. All rights reserved.

C:\Users\harps>cd Downloads

C:\Users\harps\Downloads>configureMosaic.exe help
Usage: configureMosaic.exe <command> [arguments...]


configureMosaic.exe query <query command> [query command] ...
    Return information on the current system or configuration.
    Specify one or more query types:
    gpu -- Display the list of physical GPU's and attached displays.
    lgpu -- Display the list of possible logical GPU's
    current -- Display the current display grids.

configureMosaic.exe <set|test> <grid options> [nextgrid <grid options>] [qsyncslave]...
    Set or test one or more display grids. Use the token "nextgrid" to separate
    multiple grids.
    Options:
        rows=ROWS -- Specify the number of rows in the grid.
        cols=COLS -- Specify the number of columns in the grid.
        out=GPU_INDEX,OUTPUT_INDEX -- Specify a single output. Use once for each display.
        display=DISPLAY_ID -- Specify a single output by display ID.
        res=WIDTH,HEIGHT,FREQ -- Specify display mode
        overlap=COL,ROW -- Specify the overlap for the whole grid.
        gridPos=X,Y -- Specify the coordinates of the upper left corner for each grid (when using multiple grids)
        passivestereo -- Creates a passive stereo configuration with the requested topology on each eye.
        As an example the command 'configureMosaic set rows=2 cols=2 passivestereo' would create a 2x2
        Mosaic for the left eye and a 2x2 Mosaic for the right eye.
        bezel=COL,ROW -- Specify the bezel correction for the whole grid. It is equivalent to overlap=-COL,-ROW
        qsyncslave -- For Mosaic Topology, which spans across 2 QSync board.
  
```

5. A basic setup for this would be a command based on the number of rows, columns, and desired resolution. If no resolution flag is set, the command will use the default native resolution of the monitors in the setup. For example:

C:\Path\ConfigureMosaic.exe set rows=2, cols=4 res=3840,2160,60

will result in a Mosaic GRID of 8 monitors in total with 2 horizontal rows and 4 vertical columns.



To set an overlap or bezel correction, use the "overlap" or "bezel" options which will be applied to all rows and columns. Bezel correction is equivalent to negative overlap. For example:

```
C:\Path\ConfigureMosaic.exe set rows=2, cols=4 res=3840,2160,60 overlap=100,200
```

Will set an overlap of 100 pixels between rows and 200 pixels between columns.

The "overlapcol" and "overlaprow" options specify the overlap between each column and each row. You can specify a different overlap value for each row and column. Any overlaps that you do not specify will be zero. For example:

```
C:\Path\ConfigureMosaic.exe set rows=2, cols=4 res=3840,2160,60 overlaprow=100  
overlapcol=50,100,50
```

Will set an overlap of 100 pixels between rows and an overlap of 50, 100, 50 pixels between the four columns.

Similarly, the "bezelcol," and "bezelrow" options work in the same way as "overlapcol," and "overlaprow" options detailed above.

For full details and a listing of all options, refer to the configureMosaic.exe help as detailed above.

Troubleshooting and Notes

Hot Plug Persistence and Mosaic

- When hot-unplugging a Mosaic display, Mosaic is disabled only when all the displays in the Mosaic topology are unplugged.
- When hot-plugging a display, Mosaic is restored only when all the displays in the previous Mosaic topology are connected.

Multi-Display Configuration Upon Disabling Mosaic

With multiple displays connected in a Mosaic configuration, after disabling Mosaic the multi-display configuration is not restored to the previous non-Mosaic configuration.

It is the current driver policy to enable extended mode on any connected Mosaic displays when Mosaic is disabled. This is consistent with the default Windows policy of activating all connected displays in extended mode.

The resulting configuration may not match any configuration that was set up prior to enabling Mosaic. After disabling Mosaic, you can set up a different display configuration using the *NVIDIA Control Panel->Set Up Multiple Displays* page.

Display Resolutions Upon Disabling Mosaic

When disabling Mosaic, the following occurs with connected displays:

- Displays that were part of that Mosaic configuration retain the same resolution that was set when in Mosaic mode.
- Connected displays that were not part of a Mosaic configuration are automatically activated at their respective native resolutions.

Use of HDMI or Display Port Display Adapters

Occasionally, on some Mosaic setups, use of some display adapters and extenders may cause some blinking, flickering or result in the Mosaic setup not working properly.

- If you are experiencing issues with setting up the Mosaic configuration, blinking or flickering, refer to our recommended display adapter list [here](#).

Appendix A – Restrictions When Using Mosaic

Matching Displays

All displays must have at least one display resolution and refresh rate (including all advanced timing parameters) in common.

Maximum Number of Displays in Single System

Mosaic supports up to 32 displays in a single system with eight NVIDIA RTX GPUs.

Maximum Resolution

These supported GPUs have a maximum desktop limitation of 16,384 x 16,384 pixels. For proper operation, do not configure display settings that will total greater than 16,384 pixels in either the horizontal or vertical direction. For example, a 1 x 4 array consisting of panels of 6400 x 1600 pixel output each would have a total pixel output dimension of 25,600 x 1600. This would exceed the 16,384 limitation in the horizontal direction.

The total pixel output for a configuration is the total horizontal pixels by total vertical pixels. Total size on the Set-up Mosaic page calculates the total pixel dimension for you. The total pixels displayed is a combination of the Display configuration, resolution setting, and edge overlap settings.

Minimum Resolution

Mosaic also requires individual displays to use a resolution greater than 800 x 600 pixel output. If a display in your configuration is set to one of the smaller pixel output resolutions (e.g., 640 x 480 or 800 x 600), the system will revert to single display mode.

Multiple Mosaic GRID

On Windows, it is possible to create multiple Mosaic mode GRIDs using both the NVIDIA control panel and ConfigureMosaic.exe utility, where the first GRID can span multiple GPUs and subsequent GRIDs are contained on one GPU.

Non-Mosaic Displays

- Users can assign some displays to a Mosaic configuration and then use the NVIDIA Control Panel->Set Up Multiple Display page to enable another display separate from the Mosaic setup.

- Within the same system, it is possible to mix some physical displays in Mosaic with displays that aren't in Mosaic mode. This will be expanded on in the Using Mosaic mode section.
- It is not recommended to mix synchronized displays in Mosaic mode on the same GPU with non-synchronized displays, not in Mosaic mode. The limitation will be expanded upon in the using Mosaic mode section of this document.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS AND 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. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/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.

DOCUMENT

Copyright

© 2021-2023 NVIDIA CORPORATION & AFFILIATES. All rights reserved.