



SYSTEMS WITH 1 TB OR MORE OF SYSTEM MEMORY

DA-07387-001_v01 | August 2014

Application Note



DOCUMENT CHANGE HISTORY

DA-07387-001_v01

Version	Date	Authors	Description of Change
01	August 20, 2014	MB, SM	Initial Release

SYSTEMS WITH 1 TB OR MORE OF SYSTEM MEMORY

On systems with 1 TB or more of system memory the behavior is as follows:



Note: Refer to the *Memory Addressable by NVIDIA Graphics Cards Application Note* (DA-06286-001) for information regarding the amount of system memory that is addressable by various NVIDIA graphics cards.

On Linux OSs, the NVIDIA Linux driver attempts to identify the scenario where the host system has more memory than a given graphics processing unit (GPU) can address (which is 1 TB on current generation GPUs). If this scenario is detected, the NVIDIA driver will drop back to allocations from the 4 GB DMA zone to avoid address truncation. This means that the driver will use the `__GFP_DMA32` flag and limit itself to memory addresses below the 4 GB boundary. This is done on a per-GPU basis, so limiting one GPU will not limit other GPUs in the system. For example, if an NVIDIA® Quadro® K6000 (which can address 1 TB) and a Quadro 5000 (which can address 512 GB) are installed in a system with 1 TB of memory, then the Quadro K6000 operates normally with no limitations, while the Quadro 5000 falls back to the 4 GB limit. This is the behavior on R331 drivers starting with 331.93 and R340 drivers starting with 340.28.

On VMware ESXi 5.1 and ESXi 5.5, the NVIDIA VMware VIB drivers mimic the behavior of the NVIDIA Linux drivers, for example, they will fall back to the 4 GB limit when there is more than 1 TB of system memory.

On Citrix Xen 6.2, there is a 1 TB limit on the hypervisor and a 128 GB limit on VMs.

On Windows OSs (Server 2008 R2, Server 2012, and Server 2012 R2) with an NVIDIA® Tesla® card in the default TCC (Tesla Compute Cluster) mode on systems with 1 TB or more system memory, after the NVIDIA Windows driver installer is run and completed,

the Microsoft Device Manager will show a yellow bang for the GPU device with a “This device cannot start. (Code 10)” error and hence NVIDIA Tesla cards in TCC mode will not operate on Windows systems with 1 TB or more of system memory. See Figure 1.

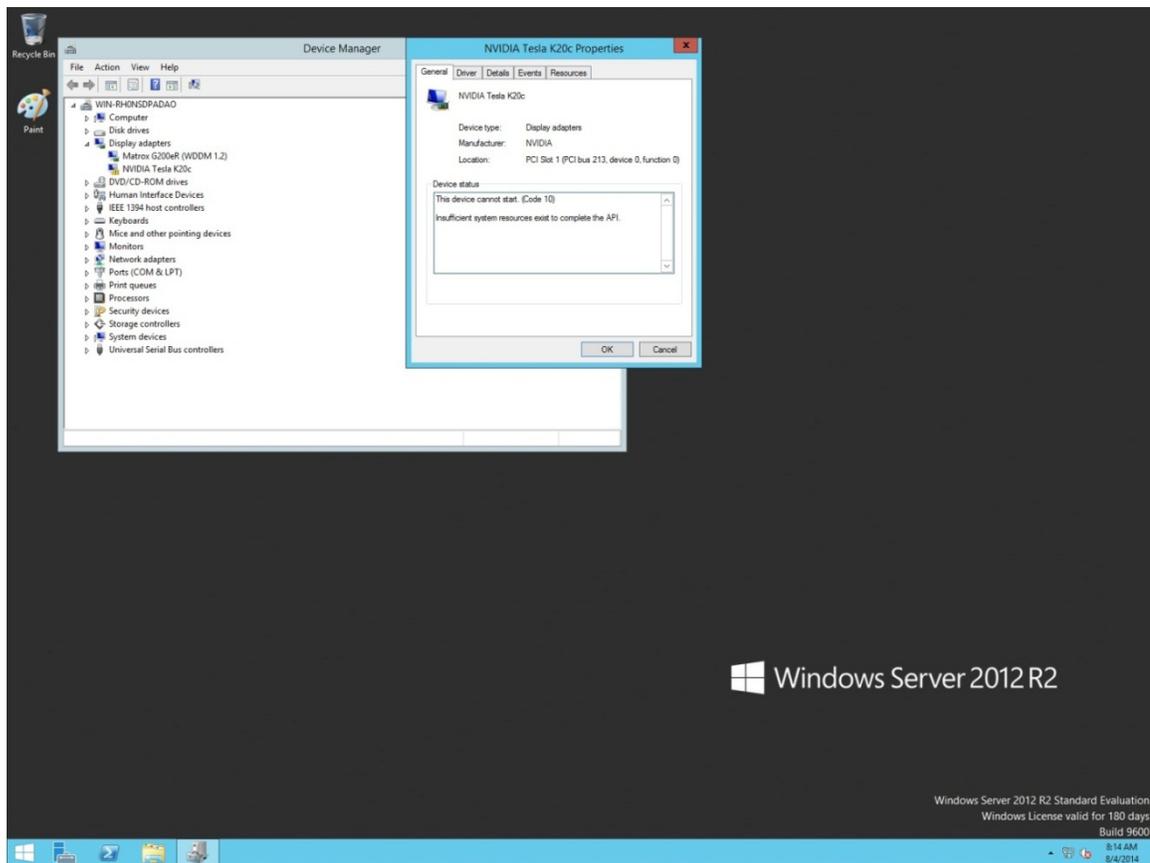


Figure 1. Code 10 Error Message

On Windows OSs (Server 2008 R2, Server 2012, and Server 2012 R2) with an NVIDIA Tesla, Quadro, or an NVIDIA GRID™ card in WDDM (Windows Display Driver Model) mode on systems with 1 TB or more system memory, after the NVIDIA Windows driver installer is run and completed, the Microsoft Device Manager will indicate “This device is working properly” even though the system will not be functional. See Figure 2.

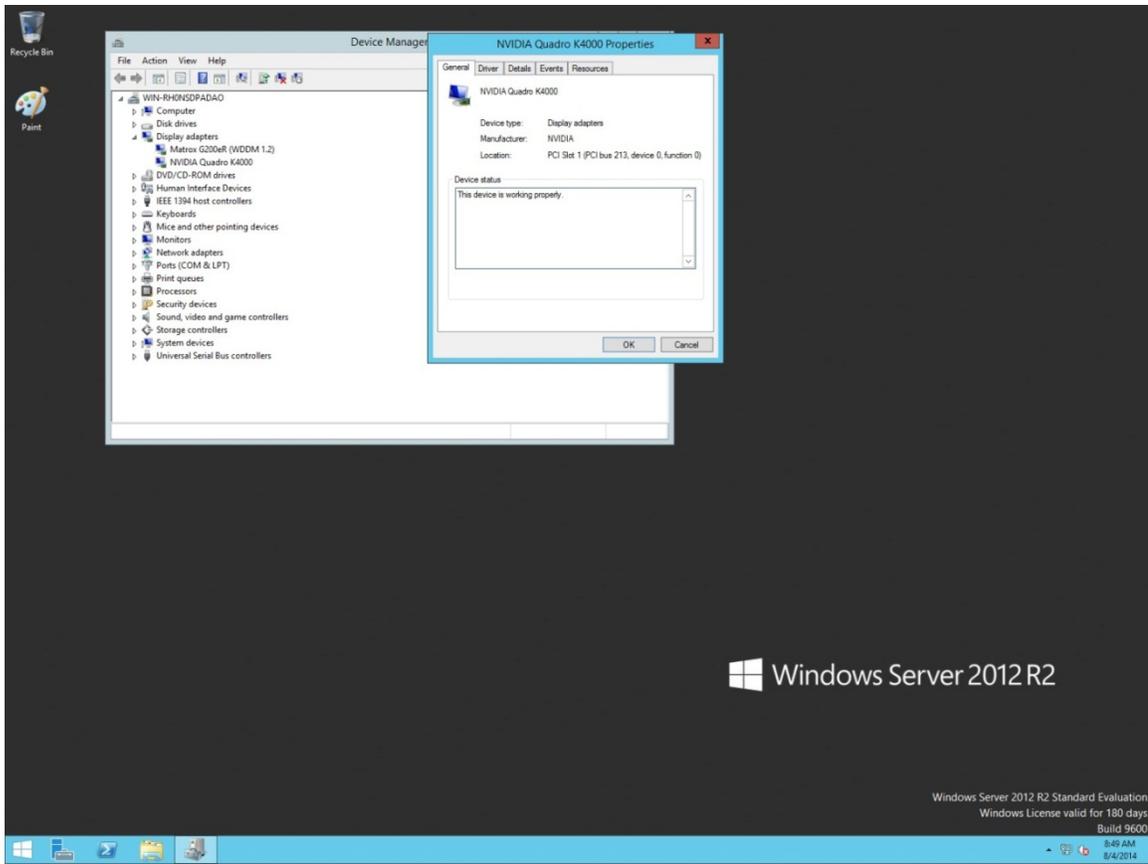


Figure 2. Windows Display Driver Model

However, after the system is rebooted, the Microsoft Device Manager will indicate “Windows has stopped this device because it has reported problems (Code 43).”

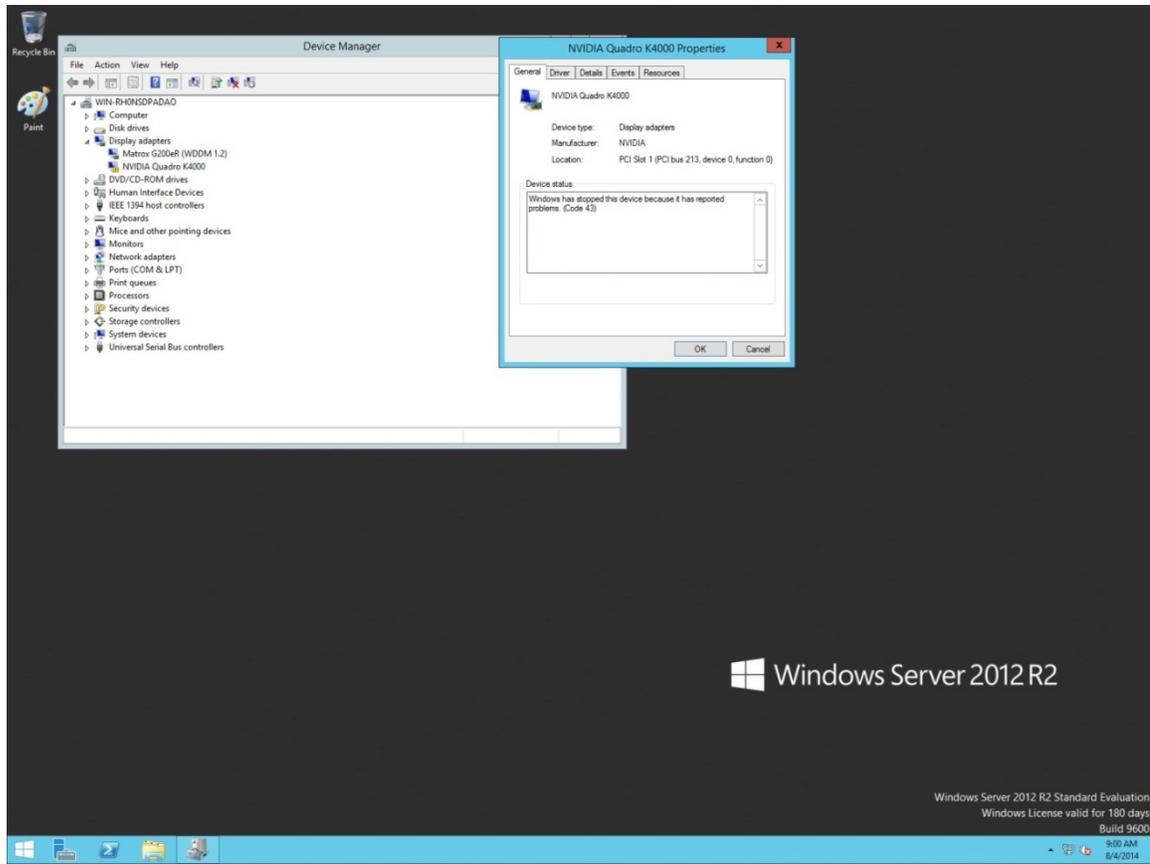


Figure 3. Code 43 Error Message

On systems with multiple NVIDIA cards, multiple reboots may be required before the Microsoft Device Manager will indicate “Windows has stopped this device because it has reported problems (Code 43).”

This is the behavior on Windows R331 drivers starting with 333.44 and R340 drivers starting with 340.66.

Notice

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. 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 publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification 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. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regard to the purchase of the NVIDIA product referenced in this specification.

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 these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 specification. NVIDIA does not accept any 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 specification, 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 specification. 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 specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

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. 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 NVIDIA terms and conditions of sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA GRID, Quadro, and Tesla 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.

Copyright

© 2014 NVIDIA Corporation. All rights reserved.