



# NVIDIA TESLA TROUBLESHOOTING GUIDE

TB-07344-001\_v01 | July 9, 2014

**Technical Note**



## DOCUMENT CHANGE HISTORY

TB-07344-001\_v01

Version	Date	Authors	Description of Change
01	07/09/2014	OC/DB	Initial release

# NVIDIA TESLA TROUBLESHOOTING GUIDE

## INTRODUCTION

This technical note is provided to OEMs and customers using NVIDIA Tesla GPUs, and is intended to assist in troubleshooting system failures to determine whether they are the result of an issue with the GPU.

*If a failure is determined to follow the GPU, please return it to NVIDIA for an RMA replacement.*

The following sections describe procedures to follow to determine whether a failure follows the GPU, depending on whether diagnostic tools can be run on the GPU or not.

- ▶ [If Diagnostics Can be Run on the GPU](#)
- ▶ [If Diagnostics Cannot be Run on the GPU](#)

## IF DIAGNOSTICS CAN BE RUN ON THE GPU

If the failure is of a nature that allows tools/diagnostics to be run on the GPU, then perform the following immediately upon observing the failure:

1. Run **nvidia-bug-report.sh**, which will collect configuration data about the system.
2. **IMPORTANT!:** Save the resulting log file (**nvidia-bug-report.log.tgz**) to submit with the RMA request.
3. Power cycle (reboot) the node.
4. Run **nvidia-bug-report.sh** a second time.
5. **IMPORTANT!:** Save the second log file to submit with the RMA request, using a different file name to differentiate it from the file saved in Step 2.
6. Review the collected **nvidia-bug-report** log files for any NVRM: Xid error messages, as documented here: <http://docs.nvidia.com/deploy/xid-errors/index.html>.

- If **Xid errors 13 or 31** are reported, there is likely an issue with the application coding.

Take the actions recommended here: [http://docs.nvidia.com/deploy/xid-errors/index.html#topic\\_5](http://docs.nvidia.com/deploy/xid-errors/index.html#topic_5).

Run the application in **cuda-memcheck** to detect potential application bugs. Once any application-level bugs identified have been resolved, verify whether the failure still occurs with the corrected application.

- If **other Xid errors** are reported, or if **cuda-memcheck** failed to resolve the issue, then verify whether the failure occurs with the latest driver from <http://www.nvidia.com>.

If the failure occurs with the latest NVIDIA driver, run **nvidia-bug-report.sh** to capture configuration data with the new driver.

**IMPORTANT!:** Save the resulting log file to submit with the RMA request.

7. Run **nvidia-healthmon**.

**nvidia-healthmon** is a tool for detecting and troubleshooting common problems affecting NVIDIA Tesla GPUs in a high performance computing environment and contains limited hardware, as well as software and system configuration, diagnostic capabilities. The latest version can be downloaded from <http://developer.nvidia.com/cuda-downloads>

Detailed instructions for running **nvidia-healthmon** on a failing GPU, as well as interpreting the results, can be found here:

[http://docs.nvidia.com/deploy/pdf/Healthmon User Guide.pdf](http://docs.nvidia.com/deploy/pdf/Healthmon%20User%20Guide.pdf)

*If nvidia-healthmon reports configuration problems* - follow the recommendations made by nvidia-healthmon to resolve the problems.

*If nvidia-healthmon returns error code 0 (no problems detected) OR it reports a problem with the Tesla hardware* - run the latest version of the NVIDIA Field Diagnostic (available from the system OEM).

- a Follow the instructions provided by the field diagnostic (for example, if ECC memory protection had previously been disabled, the user will be requested to re-enable ECC and re-run the diagnostic).
- b **IMPORTANT!:** Save the resulting log file.  
  
If the diagnostic completes successfully, the generated log file name is in the following format: **fieldiag\_<PASS/FAIL/CONFIG>\_<serial number>.log**  
  
If the diagnostic crashes, the generated log file name is **fieldiag.log**.
- c If the field diagnostic fails, arrange for RMA replacement of the board, and include all log files with the RMA request.

## IF DIAGNOSTICS CANNOT BE RUN ON THE GPU

Follow these instructions if the failure is such that tools or diagnostics cannot be run on the GPU - for example, if the GPU is not detected, or the failure is a system hang.

1. Verify that external power is properly connected to the GPU board.
2. Check whether the GPU board is fully seated in the PCIe slot.
3. Determine whether the failure follows the GPU board by:
  - Verifying whether the failure occurs with a replacement GPU board, and
  - Verifying whether the suspect board fails in a different system/node.
4. **IMPORTANT!:** If the failure follows the board in the above steps, then with the RMA submission provide a clear written description of the actions that were taken to demonstrate that the failure followed the particular Tesla board to be returned as an RMA.

### **Notice**

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.

### **Trademarks**

NVIDIA, the NVIDIA logo, 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.