# Graphic processing unit troubleshooting

Common GPU adapter issues

# Graphic processing unit overview

Failed GPUs can present in multiple ways. GPUs or vGPU-equipped VMs might:

- Hang or become unresponsive
- Crash and restart on other hosts
- Fail to start

GPU troubleshooting actions can be broken down into a few general steps:

- Verify the system matches the HX best recipe.
- Ensure that the NVIDIA GPU drivers are installed. The driver for the NVIDIA GPU card is not packaged with Nutanix Foundation and must be downloaded from NVIDIA and installed manually.
- If a GPU needs to be replaced, the documentation for both ESXi and AHV can be found in the following article: Nutanix KB 6465.

A full list of KBs referencing NVIDIA can be found on the Nutanix Support Portal.

# Common troubleshooting commands for GPUs

| Purpose | Command | Component |
|---|---|---|
| Confirming GPU installation | `lspci \| grep -i display` | Hypervisor |
| Printing out which VMs are using which GPUs | `gpuvm` | Hypervisor |
| Confirming GPU configuration | `esxcli hardware pci list -c 0x0300 -m 0xff` | Hypervisor |
| Checking if Xorg is running | `/etc/init.d/xorg status` | Hypervisor |
| Manually starting Xorg | `/etc/init.d/xorg start` | Hypervisor |
| Checking Xorg logging | `cat /var/log/Xorg.log \| grep -E "GPU\|nv"` | Hypervisor |
| Verifying the VIB installation | `esxcli software vib list \| grep NVIDIA` | Virtual GPU Manager/Resource Manager |
| Confirming the VIB is loading | `esxcfg-module -l \| grep nvidia` | Virtual GPU Manager/Resource Manager |
| Manually loading the VIB | `esxcli system module load -m nvidia` | Virtual GPU Manager/Resource Manager |
| Verifying the module is loading | `cat /var/log/vmkernel.log \| grep NVRM` | Virtual GPU Manager/Resource Manager |
| Checking the vGPU management | `nvidia-smi` | Virtual GPU Manager/Resource Manager |

# Querying GPU information

- To query the VBIOS version of each device: `#nvidia-smi --query-gpu=gpu_name,gpu_bus_id,vbios_version --format=csv`
- To monitor the hypervisor-side GPU metrics (this query will work for both ESXi and XenServer): `#nvidia-smi --query-gpu=timestamp,name,pci.bus_id,driver_version,pstate,pcie.link.gen.max,pcie.link.gen.current,temperature.gpu,utilization.gpu,utilization.memory,memory.total,memory.free,memory.used --format=csv -l 5`
- Use nvidia-smi for logging GPU usage
  - Short-term logging
    - Add the `-f <filename>` option to redirect the output to a file.
    - Prepend the `-t <seconds>` timeout to run the query for <seconds> and then stop logging.
  - Long-term logging
    - Create a shell script to automate the creation of the log file with timestamp data added to the filename and query parameters.
    - Add a custom cron job to /var/spool/cron/crontabs to call the script at the required intervals.
- Issue `#nvidia-smi --help-query-gpu` to get a complete list of the query arguments.