Using a GPU as a Console Display Adapter on a Lenovo ThinkSystem Server running Linux

Explains how to use a GPU to drive local video displays

Provides step-by-step instructions on how to configure Linux

Covers both runlevel 3 text mode and X server mode

Provides examples using an NVIDIA GPU

Xiaochun Li

LENOVO PRESS

Abstract

All Lenovo® ThinkSystem™ servers have an onboard video adapter to drive the video display attached to the server's VGA port. Many ThinkSystem servers also support a Graphics Processing Unit (GPU), and some of these GPUs offer additional video ports (typically DisplayPort or HDMI).

GPUs on servers are typically used to offload tasks off the server CPU, including AI, VDI, and rendering tasks. Customers may want to also use the GPU to drive one or more monitors locally attached to the server, however, most Linux Server OS releases do not easily enable such local display support.

This paper provides guidance on enabling GPUs for use with local displays. The paper is for Linux administrators wishing to use a GPU installed in a ThinkSystem server to drive one or more local video monitors.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction

The local console of Linux server operating systems can run in text mode, where the server is run on runlevel 3, or in graphical mode, where X server is used to provide graphical services to the local display.

Our testing environment was based on the following servers:

- ThinkSystem ST550 with the NVIDIA P2000 GPU
- ThinkSystem SR630 with the NVIDIA P2000 GPU
- ThinkSystem SR650 with the NVIDIA P620 GPU

We eveluated various versions of Red Hat Enterprise Linux installed on these servers, RHEL7.3, RHEL7.4, RHEL7.5, RHEL7.6, and RHEL7.7

Enabling a GPU in text mode

This section describes how to enable a NVIDIA GPU to operate in text mode in the Linux console. For graphics cards provided by other manufacturers, the steps may be slightly different.

Some open-source drivers included in the released server OS have defects. For example, the nouveau drivers have issues with RandR 1.3 libs. Therefore, it is advised that you use the proprietary driver offered by the graphics card manufacturer.

To install the proprietary driver based on RHEL7.x, perform the following steps:

1. Download the appropriate NVIDIA driver for your graphics controller from the GPU vendor web site:

   http://www.nvidia.com

2. Run the following commands to install the proprietary driver of NVIDIA. The driver cannot be installed if the X server is running on the system, so ensure that the system is started in text mode (runlevel 3):

   ```
   # init 3
   # sh nvidia_filename.run
   ```

   Ensure that this driver is saved in the local disk of the target system. Installing from an external device, such as a flash drive, will cause known issues such as an installation failure.

3. Edit Grub 2 to blacklist the nouveau driver. Edit /etc/default/grub and add the following to the GRUB_CMDLINE_LINUX line:

   ```
   rd.driver.blacklist=nouveau nouveau.modeset=0 fbcon=map:10
   ```

   The kernel parameter configuration `rd.driver.blacklist=nouveau nouveau.modeset=0` blacklists the nouveau driver module to disable it from getting loaded at boot from initramfs.

   For the kernel parameter configuration `fbcon=map:10`, the value 10 is the correspondence between fb and tty, indicating which graphics card driver gets mapped to which console. Therefore, you can use hot keys such as Ctrl+Alt+F1 to switch the display console to the graphics card that we want to display under runlevel 3.

4. Rebuild the grub.cfg file by running one of the following commands:

   Legacy BIOS:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

UEFI BIOS:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

5. Edit the /etc/modprobe.d/blacklist.conf file and add the following line to it, so that this is added into initramfs at rebuild:

```
blacklist nouveau
```

6. Back up the current initramfs and build a new one as follows:

```
# mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r)-nouveau.img
# dracut /boot/initramfs-$(uname -r).img $(uname -r)
```

7. Run the following command to set the default runlevel of the text mode in systemd:

```
# systemctl set-default multi-user.target
```

8. Restart the system.

The system should not load the nouveau module now at boot.

Enabling a GPU when logging in to Linux via the GUI

Most Linux OS releases will start the X server when using a GUI as the local display. The xorg.conf file provided by the X server is the configuration file for Xorg, and it needs to be modified for X server to work on a ThinkSystem GPU. This file supports various configuration parameters to configure the graphics cards, monitors, input devices, and the layout of the Xorg server on the server.

For further details about the xorg.conf file, run the following man command on the terminal:

```
man xorg.conf
```

All the following operations are performed based on RHEL7.x.

Checking the graphics card information in the server

The first step is determine the PCIe configuration of the GPU, as follows:

1. Obtain the bus IDs of available graphics cards using the **lspci** command, as follows:

```
[root@localhost ~]# lspci |grep -i vga
02:00.0 VGA compatible controller: Matrox Electronics Systems Ltd. MGA G200e [Pilot]
ServerEngines (SEP1) (rev 42)
5b:00.0 VGA compatible controller: NVIDIA Corporation GP106GL [Quadro P2000] (rev a1)
[root@localhost ~]#
```

Figure 1 Finding the bus IDs

In the output of the command, the numbers at the beginning of the entries indicate the bus IDs of the cards that are currently available on the machine. In our example, this is 30:00.0. The output is in hexidecimal.

2. Determine what kinds of driver a graphics card has been loaded with, using the following command.

```
[root@localhost ~]# lspci -vnn |grep -i VGA -A 20
02:00.0 VGA compatible controller [0300]: Matrox Electronics Systems Ltd. MGA G200e
[Pilot] ServerEngines (SEP1) [102b:0522] (rev 42) (prog-if 00 [VGA controller])
        Subsystem: Emulex Corporation Device [19a2:0101]
        Flags: bus master, fast devsel, latency 0, IRQ 16, NUMA node 0
        Memory at d8000000 (32-bit, non-prefetchable) [size=16M]
        Memory at d9a10000 (32-bit, non-prefetchable) [size=16K]
        Memory at d9000000 (32-bit, non-prefetchable) [size=8M]
        Expansion ROM at d9a00000 [disabled] [size=64K]
        Capabilities: [dc] Power Management version 2
        Capabilities: [54] MSI: Enable- Count=1/1 Maskable- 64bit-
        Kernel driver in use: mgag200
        Kernel modules: mgag200

--
5b:00.0 VGA compatible controller [0300]: NVIDIA Corporation GP106GL [Quadro P2000]
[10de:1c30] (rev a1) (prog-if 00 [VGA controller])
        Subsystem: NVIDIA Corporation Device [10de:11b3]
        Physical Slot: 1
        Flags: bus master, fast devsel, latency 0, IRQ 143, NUMA node 0
        Memory at ec000000 (32-bit, non-prefetchable) [size=16M]
        Memory at 22fe0000000 (64-bit, prefetchable) [size=256M]
        Memory at 22ff0000000 (64-bit, prefetchable) [size=32M]
        I/O ports at b000 [size=128]
        Expansion ROM at ed000000 [disabled] [size=512K]
        Capabilities: [60] Power Management version 3
        Capabilities: [68] MSI: Enable- Count=1/1 Maskable- 64bit+
        Capabilities: [78] Express Legacy Endpoint, MSI 00
        Capabilities: [100] Virtual Channel
        Capabilities: [250] Latency Tolerance Reporting
        Capabilities: [128] Power Budgeting <?>
        Capabilities: [420] Advanced Error Reporting
        Capabilities: [600] Vendor Specific Information: ID=0001 Rev=1 Len=024 <?>
        Capabilities: [900] #19
        Kernel driver in use: nvidia
        Kernel modules: nouveau, nvidia_drm, nvidia
```

Figure 2 Command to determine the type of driver

Notes:

- For the add-in graphics card, if you have installed the NVIDIA proprietary driver instead of the open-source nouveau driver, the nouveau driver must be disabled.
- To disable the nouveau driver, add the kernel command-line parameter configuration: `modprobe.blacklist=nouveau`. Refer to the steps in “Enabling a GPU in text mode” on page 3.

2.2 Editing the Xorg configuration file

The configuration file is located at /etc/X11/xorg.conf on the server. To edit it, you must be a root user. If the file does not exist, create one.

Figure 3 shows a complete xorg.conf file.

```
Section "Device"
     Identifier  "card0"
     Driver      "mgag200"
     BusID       "PCI:02:00:00"
EndSection
Section "Device"
     Identifier  "card1"
     Driver      "nvidia"
     BusID       "PCI:91:00:00"
EndSection

Section "Monitor"
     Identifier  "monitor0"
EndSection
Section "Monitor"
     Identifier  "monitor1"
EndSection

Section "Screen"
     Identifier  "screen0"
     Device      "card0"
     Monitor     "monitor0"
EndSection
Section "Screen"
     Identifier  "screen1"
     Device      "card1"
     Monitor     "monitor1"
EndSection

Section "ServerLayout"
     Identifier  "default"
     Screen 0    "screen1"
EndSection
Section "ServerLayout"
     Identifier  "default1"
     Screen 1    "screen0"
EndSection
```

Figure 3 xorg.conf file

The terms used in the xorg.conf file are as follows:

- A Device is a graphics controller or GPU
- A Monitor is the phyical video display that is connected to the server
- A Screen is a logical representation of the Device and Monitor together
- A ServerLayout is a logical representation of the Screen, Monitor, and Device

Add or modify some of the entries in the xorg.conf file as follows:

1. Modify the file so that it contains two Device sections. We have two devices because there are two controllers as shown in Figure 2 on page 5.

a. The Identifier entry specifies a unique name for the graphics device, which can be anything that will help you identify the device.

b. The Driver entry specifies the name of the video driver to be used.

c. The BusID entry specifies where the graphics card is located. A bus ID is in decimal format, instead of hexadecimal format.

> **Note:** If you use the Bus ID from the `lspci` command as shown in Figure 1 on page 4, you will have to convert the Bus ID from hexidecimal to decimal.

Note: The Identifier and Driver entries in the Device section are mandatory. Other entries are either generic or device specific, and are optional. The Device sections are referenced in the Screen sections.

2. Create two Monitor sections. We are creating two monitors because we have two displays connected to the GPU. There can be more than two monitors, depending on what you have connected.

   — The Identifier entry specifies a unique name for the monitor, which can be anything that will help you identify the monitor.

   Note: Only the Identifier entry is mandatory. Create such an entry for every monitor that you have connected.

3. Create two Screen sections. The number of screens should match the number of devices and the number of monitors.

   The configuration file may contain multiple Screen sections, and there must be at least one entry for each screen that is used. A Screen is a logical representation of the combination of a graphics card and a display device. The Screen sections are referenced in the ServerLayout sections.

   a. The Identifier entry specifies a unique name for the screen.

   b. The Device entry specifies a graphics card name, which is represented by the Identifier entry in the corresponding Device section.

   c. The Monitor entry specifies a monitor name, which is represented by the Identifier entry in the corresponding Monitor section.

   Note: Among these entries, only the Identifier and Device entries are mandatory. Other optional entries are not elaborated in this guide. For these entries not elaborated, their default values are used instead.

4. Create the last section that gathers the preceding sections, the ServerLayout section. A ServerLayout section is a logical representation of the Screen, Monitor, and Device sections, which form complete Xorg configuration. There can be multiple ServerLayout sections in the xorg.conf file, and there must be at least one Identifier entry and one Screen entry in each ServerLayout section.

   a. The Identifier entry specifies a unique name for the ServerLayout.

   b. The Screen 0 or Screen 1 entry specifies a screen name, which is represented by the Identifier entry in the corresponding Screen section.

   Note: The display cannot be output to two graphics cards at the same time, so you can modify the Screen 0 entry of the ServerLayout section to determine the graphics card you want to choose. Some Add-in graphics cards could not display when you choose the Onboard card as the screen.

Once you have finished configuring the /etc/X11/xorg.conf file, save and exit, and then reboot the OS and press F1 to enter System Setup.

2.3 Switching to the graphics card

The steps to activate the GPU are as follows:

1. In System Setup (F1 at boot), enter the UEFI configuration for activating the Add-in graphics card. The default active video is Onboard Device.
2. From the BIOS setup menu path, select **System Settings** → **Devices and I/O ports** → **Active Video** to activate the Add-in graphics card as shown in Figure 4 on page 8.

Figure 4 Devices and I/O Ports in System Setup

3. After activating Add-in Device, save and exit the BIOS setup menu, and then enter the Linux OS. The Add-in graphics card should now display the X server GUI.

References

For additional information, see these web pages:

- ArchWiki page for Xorg

  https://wiki.archlinux.org/index.php/Xorg

- Super User page, “How can I get multiple video cards to work on Linux?”

  https://superuser.com/questions/117239/how-can-i-get-multiple-video-cards-to-work-on-linux

Author

Xiaochun Li is a Linux engineer at the Lenovo Data Center Group in Beijing, China. He specializes in development related to Linux kernel storage and memory management, as well as kernel DRM. Before joining in Lenovo, he had been working for INSPUR as an OS engineer for several years. With seven years of industry experience, he now focuses on Linux kernel RAS, storage, and memory subsystems.

Thanks to the following people for their contributions to this project:

- Song Shang, Lenovo test engineer for Linux enablement

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on September 27, 2019.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp1228

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo® Lenovo(logo)® ThinkSystem™

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.