



CONTENTS & NAVIGATION

2

Introduction
Getting an EDID

3

Choosing the correct ports

4

Modifying xorg.conf

5

Common errors



HOW TO SET UP A HEADLESS LINUX[®] SYSTEM TO USE WITH HP ZCENTRAL REMOTE BOOST

CONTENTS & NAVIGATION

2

Introduction
Getting an EDID

3

Choosing the correct ports

4

Modifying xorg.conf

5

Common errors

INTRODUCTION

In most circumstances, Remote Boost is used to connect to a remote system that has peripherals like a mouse, keyboard, and display. Sometimes, a system may be stored on a rack without any of these things. A system without a display connected to it is called a “headless” system. Remote Boost can function on these as well, but the machine needs a little configuration before it can function properly. In this paper, I will walk through the necessary steps to configure a headless Linux® sender system.

There are some packages required on the sender before we begin:

- A desktop environment
 - Minimally, the machine needs an X server and a login screen. This is provided by packages like gdm.
- The proprietary Linux® NVIDIA® driver. This provides several tools and features required to successfully set up a headless system.

This document will require you to use a Linux® command line, so you may need to know some basic commands. Here are all the commands you will use, along with a brief description:

- xrandr
 - This command will print out information about the displays connected to the system.
- nvidia-xconfig
 - This command will create a configuration file for the X server.
- mv <src> <dest>
 - This is the command to move or rename a file. It takes two arguments, the first is the file you’re moving, the second is the location you’re moving it to.
- cp <src> <dest>
 - This is the command to copy a file. Like mv, it takes two arguments, the file to copy and the location to copy it to.
- cd <dest>
 - This is the command to change the working directory. It takes one argument, the directory where you would like to go.
- sudo <cmd>
 - This is the command to run another command with administrator privileges. It can precede any other command.


GETTING AN EDID

EDID files describe the resolutions and capabilities of a monitor. In order to set up a headless system on Linux®, you need one of these files so that the system can pretend like it’s actually connected. There are many tools that can generate one of these EDID files.

All of these methods must be performed on a machine with physical monitors attached. You will be copying the attributes of those physical monitors to use on the headless machines later.

Using NVIDIA® X Server Settings

1. From the application menu, open “NVIDIA® X Server Settings”.
2. On the left-hand side, click on one of the connected monitors. It should appear under a menu titled GPU-0.
3. On the bottom right, click on the “Acquire EDID” button.
4. Choose a location to save your EDID file. Choose a name that corresponds with the monitor it came from.



CONTENTS & NAVIGATION

2

Introduction
Getting an EDID

3

Choosing the correct ports

4

Modifying xorg.conf

5

Common errors

4. You should always choose disconnected ports with the highest maximum clock speed. In the above example, that would be the ports DFP-0, DFP-2, DFP-4, and DFP-6.
 - A headless system would not show any ports as connected.
5. Write down the port names for use in the next step.

NOTE: Some ports on modern systems are “paired”, meaning only one of the two paired ports can be used to plug in our EDID file. While there’s no strict rule for determining paired ports, it’s generally safe to use all odd or all even ports.

NOTE: If you’ve previously configured a headless system with the option “UseDisplayDevice” “None”, these ports may not appear in the Xorg.0.log file. Remove this option and reboot your machine to find them.

MODIFYING XORG.CONF

Now that we have an EDID file and we know where we want to put it, we have to tell the X server to do that when it starts up. We do this by modifying the xorg.conf file. This will require root privileges.

All this work can be done with command **nvidia-xconfig** **g** from the command line. Your command will look something like this:

```
nvidia-xconfig -c /dev/null \  
-o test.conf \  
--custom-edid="/etc/X11/EDIDS/mymonitor.txt" \  
--connected-monitor="DFP-0" \  
--busid="PCI:5:0:0"
```

This is a pretty long command, so let’s break it down:

- The “**-c /dev/null**” option forces the command to generate an xorg.conf file from scratch. If you have other modifications in your xorg.conf file, you should not include this option.
- The “**-o test.conf**” option prevents the command from overwriting your current xorg.conf file. This way you can review your changes before you apply it to your system.
- The “**--custom-edid**” option plugs our EDID files into all the ports that begin with DFP. You should replace “**mymonitor.txt**” with the name of the EDID file you created in part 1.
 - If you want to apply an EDID file to a specific port rather than a set of ports, replace DFP with DFP-0, DFP-1, etc. like this:

```
--custom-edid="DFP-2:/etc/X11/EDIDS/mymonitor.txt"
```

– If you want to use different EDIDS with different ports, separate them with a semicolon, like this:

```
--custom-edid="DFP-0:/etc/X11/EDIDS/monitor1.txt; DFP-2:/etc/X11/EDIDS/monitor2.txt"
```

- The “**--connected-monitor**” option specifies that these ports are active and can be configured. You must specify individual ports here, not sets of ports like “DFP”.
 - If you want to activate multiple ports, separate them by commas, like this:

```
--connected-monitor="DFP-0, DFP-2"
```

- The “**--busid**” option isn’t required on some systems, but it’s best to do it anyway. You’ll have to find the PCI bus ID for your card. This can be done by running:

```
nvidia-xconfig --query-gpu-info
```

This command may print out an error message saying “VALIDATION ERROR: Data incomplete in /dev/null”. This error is expected, and will not cause a problem. After you’re satisfied with your configuration, apply it to your system by running the command:

```
mv test.conf /etc/X11/xorg.conf
```

And restart your system.

CONTENTS & NAVIGATION

2

Introduction
Getting an EDID

3

Choosing the correct ports

4

Modifying xorg.conf

5

Common errors

COMMON ERRORS

When modifying your xorg.conf file, it's easy to get your system into a state where the system fails to initialize properly. Trying to connect to this system with Remote Boost may result in the sender crashing or just displaying a black screen. It's also possible you can connect, and the resolutions associated with your EDID won't load properly. If any of these happen, check for these common mistakes:

- Check the spelling of all options to nvidia-xconfig. Note that the option "--connected-monitor" is NOT "--connected-monitors" if you're listing multiple.
- In the "--custom-edid" option, make sure there's a colon between the port and the path to the EDID file. There shouldn't be a space after the colon.
- If you're listing multiple EDID files in the "--custom-edid" option, make sure they're separated by a semicolon and a space.
- In the "--connected-monitor" option, multiple monitors must be separated by a comma and a space.
- Always use the ports listed in the Xorg.0.log file rather than those reported by tools like xrandr. The log file will show ports like "DFP-0", whereas xrandr may list the ports as "DP-0".
- If you've double-checked the command you ran and it still doesn't work, try changing the connected port. All cards are different, some ports may not work.

LET US HELP YOU CREATE AMAZING BUSINESS
SOLUTIONS TODAY

LEARN MORE

© Copyright 2021 HP Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel, Core, Thunderbolt, and Xeon are trademarks of Intel Corporation in the U.S. and other countries. NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

4AA7-9532ENW, February 2021.

