


 Lambda



Server

Quickstart Guide

Server

Last updated November 8, 2024

// 01

This Guide Will Walk You Through

This guide helps you verify that your server is functioning properly and is remotely accessible before you rack it in your data center or server closet. This guide assumes you're physically at your server with a keyboard and monitor, that you've completed the initial Ubuntu configuration, and that your server is connected to your network. You'll need the IP address of your LAN interface to access your server with SSH, and the IP address of your IPMI interface to access your server with IPMI.

1. Verifying that your server's GPUs, memory, and storage drives are detected.
2. Configuring your server for remote access via SSH and IPMI.

Note: If you need any help with these steps, or if you see something unexpected, contact Lambda Support before racking your server.

// 02

Verifying GPUs, Memory, and Storage

Use the `nvidia-smi` tool to monitor and manage your server's GPUs. To verify that all of the GPUs are detected, run `nvidia-smi -L`. This command outputs a list of all of the GPUs installed in your server, for example:

```
GPU 0: NVIDIA A10  
(UUID: GPU-1ecde72c-3e82-e60a-a61a-266f75992b88)
```

This example output shows a single NVIDIA A10 GPU.

To find basic information on server memory storage, use the `free` command.

To verify the amount of memory installed on your server, use the `free -hg` command. In the command output, look for the `Mem:` row under `total`.

To verify that your storage drives are detected, use the following command:

```
lsblk -de 7 -o name,model
```

This command provides output similar to the following:

```
Name nvmeOn1  
Model WDC PC SN730 SDBQNTY-1T00-1001
```

The output shows the device name and model (part number) of each storage drive detected in your server.

// 03

Configure Your Server for Remote Access via SSH and IPMI

In most cases, you'll use SSH to remotely administer your server. However, for low-level administration tasks, such as changing BIOS settings or reinstalling Ubuntu, you'll use your server's IPMI capabilities.

Install OpenSSH Server by running `sudo apt update && sudo apt -y install openssh-server`. From another computer on the same network as your server, try to SSH into your server using your server's LAN IP address, and the username and password you created as part of the initial Ubuntu configuration.

To use your server's IPMI capabilities, obtain the default IPMI password from your server's service tag. Install `ipmitool` by running `sudo apt update && sudo apt -y install ipmitool`. Run `ipmitool user list 1` to view the user list, and confirm that ID 2 is `admin` or `ADMIN`. To reset your server's IPMI password, run `ipmitool user setpassword 2`.

From another computer on the same network as your server, open a browser and go to the IP address for your server's IPMI. At the login page, enter `admin` or `ADMIN` for the username, depending on what appears in the user list (the command is case sensitive). Enter the password you set earlier. After entering these credentials, you will see your server's IPMI homepage.



QUESTIONS?
lambdalabs.com/support



NEED QUICK ACCESS TO MORE GPUS?
lambdalabs.com/service/gpu-cloud