

Gateway Setup Guide

📌 Overview

Lumeo [Gateway](#) can be setup on a wide variety of hardware running on your premises, in public clouds or in the Lumeo cloud. This guide describes how to set it up on your hardware.

Helpful References :

[Gateway](#) documentation

[Recommended Gateway Hardware](#)

Lumeo Gateway Installer

Interactive Install

Set up a new Gateway by running the following command on it and following the installer prompts:

```
bash <(wget -qO- https://link.lumeo.com/setup)
```

This command can also be found in Console under Gateways → Add Gateway. The installer offers you a bunch of options to optimize your installation in the interactive model. Installing the Gateway using this installer will also set it up to auto update whenever a new Gateway version is released.

Unattended Install

Gateway installer supports the following environment variables to allow for bulk, unattended installations:

Env Variable	Description	Notes
--------------	-------------	-------

Env Variable	Description	Notes
LUMEO_APP_ID	Specify the Application ID to connect this gateway to. See API section to find your Application ID.	Currently only supported for host (non-containerized) installations or Advanced Installation (see section below)
LUMEO_API_KEY	Provide the API Key for the specific Application you want to connect this gateway to. See API section to find your API Key.	Currently only supported for host (non-containerized) installations or Advanced Installation (see section below)
NO_PROMPT	Specify this env variable to provision a Gateway using the default settings.	

Examples:

Install and provision gateway to a specific Application Id.

```
LUMEO_APP_ID="app_id" LUMEO_API_KEY="api_key" bash <(wget -qO- https://link.lumeo.com/s
```

Manual Container Install

You can also install Lumeo containers directly using the following commands, but then you will be responsible for updating the container yourself when new Gateway versions are released.

The interactive or unattended install methods described above also let you use containers, but in addition, will set your containers up for automated upgrades when new versions are released. Use the Manual container install method when you would like control over upgrades.

Nvidia Jetson

```
docker run -d --restart unless-stopped \
  -v lumeo-container-name:/var/lib/lumeo \
  --env LUMEO_APP_ID="<app_id>" --env LUMEO_API_KEY="<api_key>" \
  --network host --gpus all lumeo/gateway-nvidia-jetson:latest
```

x86 + DGPU

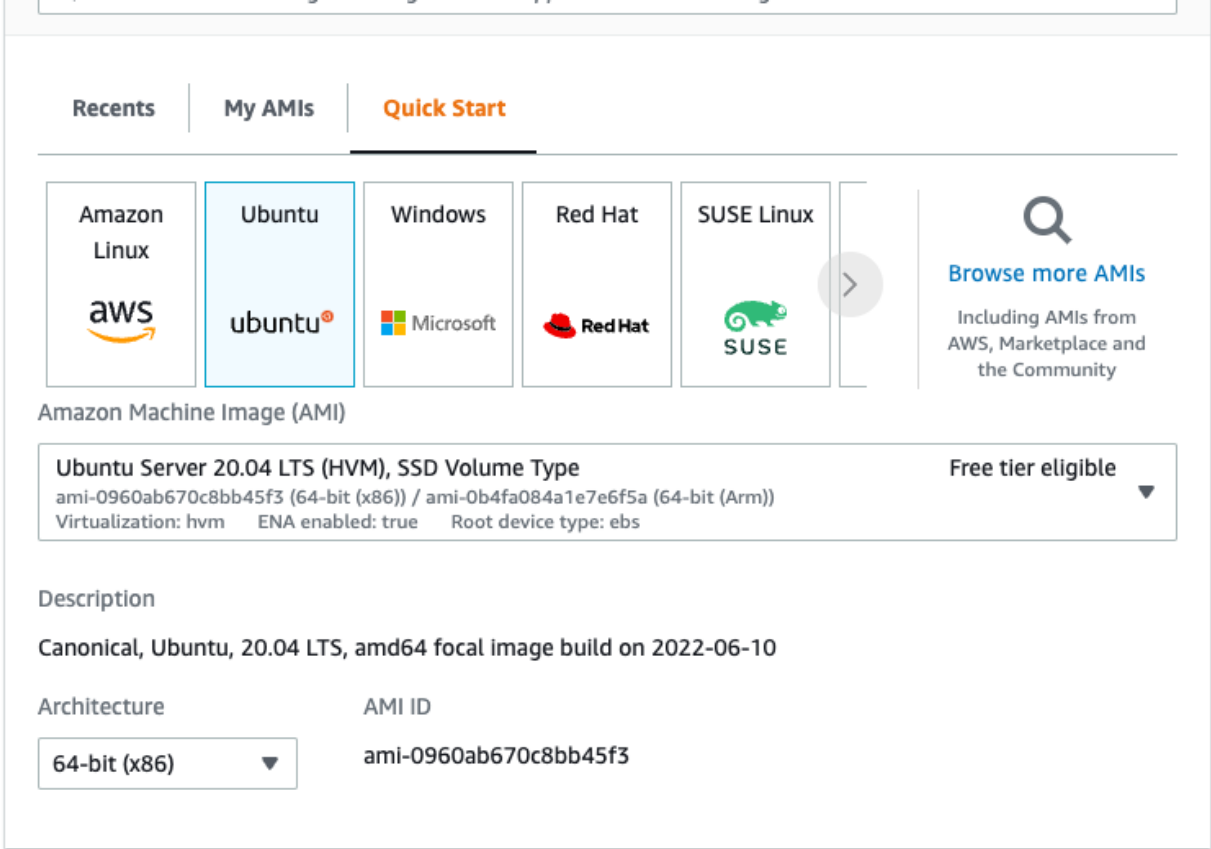
```
docker run -d --restart unless-stopped \
  -v lumeo-container-name:/var/lib/lumeo \
  --env LUMEO_APP_ID="<app_id>" --env LUMEO_API_KEY="<api_key>" \
  --network host --gpus all lumeo/gateway-nvidia-dgpu:latest
```

AWS - Lumeo Gateway Installer

Before using the Lumeo Gateway Installer on your Nvidia GPU enabled AWS EC2 machines, follow these steps :

In the EC2 console, configure and launch machine instance with the following configuration:

- Ubuntu 20.04 or 22.04 OS (x86)
- At least 50 GB of Disk space
- Instance type must have an Nvidia GPU (XXdn.XXXXX are Nvidia GPUs)
 - <https://docs.aws.amazon.com/dlami/latest/devguide/gpu.html>
 - G4dn.xlarge is a single T4 which is suitable and recommended for most applications



The screenshot displays the AWS Management Console interface for selecting an Amazon Machine Image (AMI) and an instance type. The 'Quick Start' tab is active, showing a row of AMI options: Amazon Linux, Ubuntu (selected), Windows, Red Hat, and SUSE Linux. To the right, there is a 'Browse more AMIs' link. Below the AMI selection row, the details for the selected 'Ubuntu Server 20.04 LTS (HVM), SSD Volume Type' AMI are shown, including its AMI ID 'ami-0960ab670c8bb45f3' and a note that it is 'Free tier eligible'. The architecture is set to '64-bit (x86)'. Below the AMI details, the 'Instance type' section is visible, showing 'g4dn.xlarge' as the selected instance type. The console also provides information about the instance type, such as 'Family: g4dn', '4 vCPU', and '16 GiB Memory'.

Recents | My AMIs | **Quick Start**

Amazon Linux | **Ubuntu** | Windows | Red Hat | SUSE Linux | [Browse more AMIs](#)

Including AMIs from AWS, Marketplace and the Community

Amazon Machine Image (AMI)

Ubuntu Server 20.04 LTS (HVM), SSD Volume Type Free tier eligible

ami-0960ab670c8bb45f3 (64-bit (x86)) / ami-0b4fa084a1e7e6f5a (64-bit (Arm))

Virtualization: hvm ENA enabled: true Root device type: ebs

Description

Canonical, Ubuntu, 20.04 LTS, amd64 focal image build on 2022-06-10

Architecture AMI ID

64-bit (x86) ami-0960ab670c8bb45f3

▼ **Instance type** [Info](#)

Instance type

g4dn.xlarge

Family: g4dn 4 vCPU 16 GiB Memory

On-Demand Linux pricing: 0.526 USD per Hour

On-Demand Windows pricing: 0.71 USD per Hour

[Compare instance types](#)

After launching the virtual, log in, execute the following commands to configure the machine instance and install Lumeo Gateway software:

```
sudo apt update
```

Install nvidia drivers

```
sudo apt install nvidia-compute-utils-495 nvidia-dkms-495 nvidia-driver-495 nvidia-kern
```

Verify the nvidia drivers are properly installed, execute the command:

```
nvidia-smi
```

To see output like this

Shell

```
ubuntu@aws-lumeo-gateway:~$ nvidia-smi
```

```
Sat Jun  4 01:10:01 2022
```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
NVIDIA-SMI		510.73.05		Driver Version: 510.73.05			CUDA Version: 11.6		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
GPU	Name		Persistence-M		Bus-Id		Disp.A		Volatile
Fan	Temp	Perf	Pwr:Usage/Cap		Memory-Usage		GPU-Util		Uncorr. ECC
									Compute M.
									MIG M.
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====									
0	Tesla T4		Off		00000000:00:1E.0		Off		0
N/A	37C	P8	15W / 70W		0MiB / 15360MiB		0%		Default
									N/A
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Processes:									
GPU		GI	CI	PID	Type	Process name		GPU Memory	
		ID	ID					Usage	
=====									
No running processes found									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									

If you receive an error, reboot the instance as it maybe required for different instance types and the drivers to load properly.

Finally execute the installer interactively and selecting the default option which is container.

```
bash <(wget -q0- https://link.lumeo.com/setup)
```

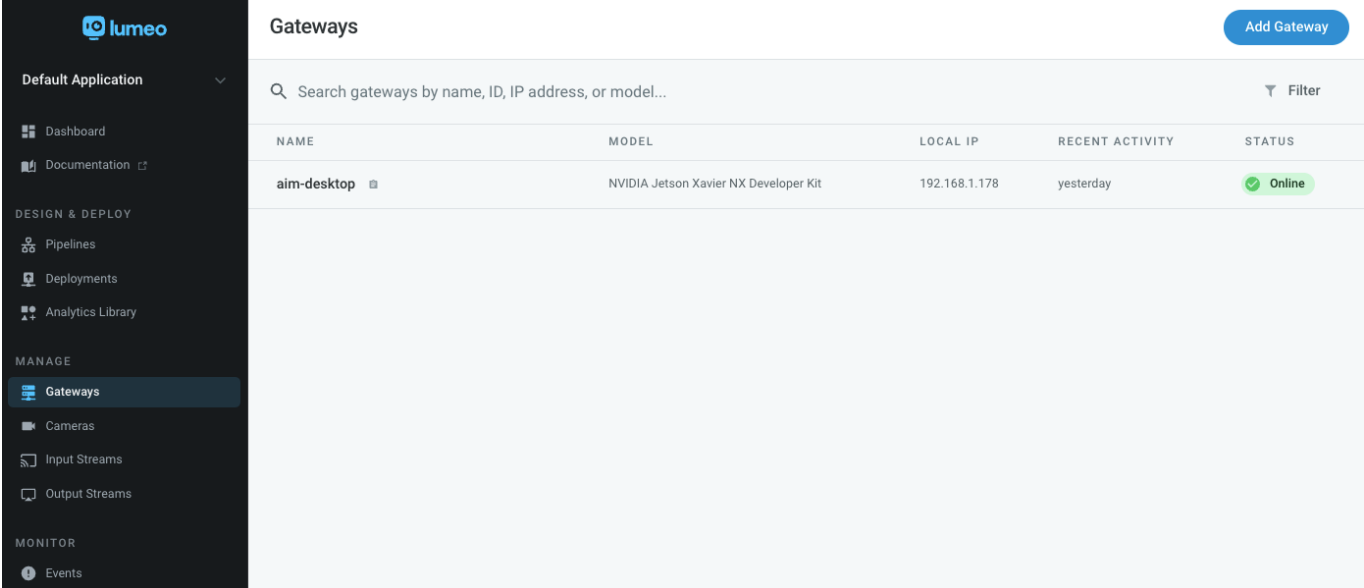
AWS - AMI Images

Using a Lumeo published AMI image makes it easy to setup a Lumeo Gateway hosted in your AWS account and provision it to your Lumeo account.

Instructions : Coming soon.

Post Install

Once setup, return to the Console to see it show up in the Gateways list.



The screenshot shows the Lumeo console interface. On the left is a dark sidebar with the Lumeo logo and a navigation menu. The main area is titled 'Gateways' and features a search bar, a filter icon, and a table of gateway information.

NAME	MODEL	LOCAL IP	RECENT ACTIVITY	STATUS
aim-desktop	NVIDIA Jetson Xavier NX Developer Kit	192.168.1.178	yesterday	Online

You need to ensure the Gateway is able to talk to Lumeo's cloud services. More info in [Security & Network](#) section.

Updated 6 days ago

← Recommended Gateway Hardware

Jetson Tips: Monitoring and Jetpack Upgrade →

Did this page help you? ☐ Yes ☐ No