

Enabling NVIDIA GPU Support for the LPR Module in Cluebase VMS

Prerequisites

NVIDIA GPU support requires the NVIDIA drivers and NVIDIA Container Toolkit to be installed on the host system.

Install and configure NVIDIA Container Toolkit before proceeding:

[NVIDIA Container Toolkit Documentation](#)

You can verify that Docker has access to the NVIDIA runtime using:

```
docker info | grep -i nvidia
```

You can also verify GPU availability on the host system with:

```
nvidia-smi
```

1. Load the Docker image

Copy the `roadar_lpr_websocket.tar.gz` archive containing the Docker image to the server.

Load the Docker image using the following command:

```
sudo docker load -i roadar_lpr_websocket.tar.gz
```

2. Stop the existing `vms-lpr` container

Navigate to the directory where Cluebase VMS is installed and stop the `vms-lpr` container if it is currently running:

```
sudo docker compose down vms-lpr
```

3. Update the `docker-compose.yml` configuration

Open the `docker-compose.yml` file in a text editor.

Replace the following configuration block:

```
vms-lpr:
  image: vcloudaiorg/vcloudai-vms-lpr:latest
  restart: always
  container_name: vms-lpr
  network_mode: host
  extra_hosts:
    - host.docker.internal:host-gateway
  depends_on:
    - vms-server
  volumes:
    - ./static/lpr:/data
  environment:
    WS_HOST: host.docker.internal:${WS_SERVER_PORT}
    API_HOST: 127.0.0.1
    API_PORT: ${ROADAR_PORT}
    STATE_FILE: /data/state.json
    LICENSE_SERVICE_HOST: 127.0.0.1
    LICENSE_SERVICE_PORT: 32433
```

with the following GPU-enabled configuration:

```
vms-lpr:
  image: roadar_lpr_websocket
  restart: always
  container_name: vms-lpr
  network_mode: host
  extra_hosts:
    - host.docker.internal:host-gateway
  depends_on:
    - vms-server
  runtime: nvidia
  deploy:
    resources:
      reservations:
        devices:
          - driver: nvidia
            count: all
            capabilities: [gpu, compute, video]
  volumes:
    - ./static/lpr:/data
  environment:
    - WS_HOST=host.docker.internal:${WS_SERVER_PORT}
    - API_HOST=127.0.0.1
    - API_PORT=${ROADAR_PORT}
    - STATE_FILE=/data/state.json
    - LICENSE_SERVICE_HOST=127.0.0.1
    - LICENSE_SERVICE_PORT=32433
    - TRIAL_TIME=0
    - NVIDIA_VISIBLE_DEVICES=all
    - NVIDIA_DRIVER_CAPABILITIES=compute,video,utility
```

Save the file and exit the text editor.

4. Start the `vms-lpr` container

Run the following command to start the container with NVIDIA GPU support enabled:

```
sudo docker compose up -d vms-lpr
```

5. Verify NVIDIA GPU support inside the container

To verify that the container has access to the NVIDIA GPU, run:

```
sudo docker exec -it vms-lpr nvidia-smi
```

If GPU support is configured correctly, the command will display information about the installed NVIDIA GPU(s).

Revision #2

Created 2026-05-26 18:26:35 UTC by Zalim

Updated 2026-05-26 18:43:34 UTC by Zalim