

FreeFalcon AI Platform

Installation Guide - Part 1

1. Overview

This guide describes the installation procedure for the FreeFalcon AI Platform on Ubuntu Linux.

The installation is performed using the supplied installation package and automated installation script.

The installer automatically installs and configures:

- Docker Engine
- Docker Compose
- FreeFalcon AI Platform
- Required platform services
- Databases
- Message broker
- AI services

No manual Docker installation or configuration is required.

2. Prerequisites Checklist

Requirement	Status
Ubuntu 22.04 LTS or Ubuntu 24.04 LTS installed	☐
NVIDIA Driver 535 or newer installed	☐
NVIDIA GPU detected by nvidia-smi	☐
NVIDIA runtime libraries installed	☐
NVIDIA Container Toolkit 1.14 or newer installed	☐
Required network ports available	☐
FreeFalcon AI installer copied to the server	☐

3. Supported Operating Systems

- Ubuntu Server 22.04 LTS
- Ubuntu Desktop 22.04 LTS
- Ubuntu Server 24.04 LTS
- Ubuntu Desktop 24.04 LTS

Both Desktop and Server editions are fully supported.

4. Verify Required Network Ports

External Ports:

Standard profile: 80 TCP - HTTP, 443 TCP - HTTPS

Alternative profile: 8000 TCP - HTTP, 8443 TCP - HTTPS

Internal Ports:

80, 443, 9001, 9100, 6379, 3306, 33060, 4369, 5671, 5672, 15671, 15672, 15691, 15692, 25672

Note: these are container-internal ports used only between platform services on the Docker network. MySQL is reachable only within the container network (port 3306) and is not exposed on the host.

Verify:

```
sudo ss -tulpn
```

or

```
sudo netstat -tulpn
```

Example:

```
Netid State Recv-Q Send-Q Local Address:Port
```

```
tcp LISTEN 0 4096 0.0.0.0:22
```

```
tcp LISTEN 0 4096 127.0.0.1:631
```

```
udp UNCONN 0 0 0.0.0.0:68
```

5. NVIDIA Driver

Install NVIDIA Driver version 535 or newer.

Official guide:

<https://docs.nvidia.com/datacenter/tesla/driver-installation-guide/>

Verify:

```
nvidia-smi
```

Example:

```
NVIDIA-SMI 570.172.08
```

```
Driver Version: 570.172.08
```

```
CUDA Version: 12.8
```

```
GPU: NVIDIA GeForce RTX 3060
```

Memory: 0 MiB / 12288 MiB

This confirms the GPU and CUDA are available.

6. NVIDIA Runtime Libraries

The NVIDIA runtime libraries used by FreeFalcon AI must correspond to the installed NVIDIA driver version.

Required packages include:

- libnvidia-compute
- libnvidia-decode

These packages are typically installed automatically together with the NVIDIA driver.

For additional information, refer to:

<https://packages.ubuntu.com/>

Verify Installed Libraries

Run:

```
dpkg -l | grep libnvidia
```

Example output:

```
ii libnvidia-compute-570  570.172.08-0ubuntu0.24.04.1
ii libnvidia-decode-570  570.172.08-0ubuntu0.24.04.1
ii libnvidia-encode-570  570.172.08-0ubuntu0.24.04.1
ii libnvidia-extra-570   570.172.08-0ubuntu0.24.04.1
```

Verify that the installed package versions match the installed NVIDIA driver version.

7. NVIDIA Container Toolkit

FreeFalcon AI utilizes Docker containers for AI inference.

To allow Docker containers to access the NVIDIA GPU, install:

NVIDIA Container Toolkit 1.14 or newer

Official installation guide:

<https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/latest/install-guide.html>

Verify Installation

Run:

```
nvidia-container-cli --version
```

Expected output:

```
version: 1.17.8
build date: 2025-03-12T14:31:22Z
build revision: xxxxxxxxxx
build compiler: gcc 11.4
build platform: linux/amd64
```

Any recent version compatible with Container Toolkit 1.14 or newer is supported.

8. Verify Prerequisites

Before proceeding with the installation, verify that:

- NVIDIA Driver is installed
- GPU is detected by `nvidia-smi`
- NVIDIA Container Toolkit is installed
- NVIDIA runtime libraries are installed

9. Copy the Installation Package

Copy the FreeFalcon AI installation package to the Ubuntu server.

Example:

```
FreeFalcon_AI_Installer.zip
```

Extract the archive:

```
unzip FreeFalcon_AI_Installer.zip
```

Navigate into the extracted directory:

```
cd FreeFalcon_AI_Installer
```

10. Switch to the Root User

Before running the installer, switch to the root account.

```
sudo su
```

11. Make the Installation Script Executable

Grant execute permission if necessary.

```
chmod +x install.sh
```

12. Start the Installation

Run the installer:

```
./install.sh
```

The installer automatically performs the following actions:

- Installs Docker Engine
- Installs Docker Compose
- Downloads all required Docker images
- Deploys the FreeFalcon AI Platform
- Configures networking
- Creates Docker containers
- Configures databases
- Configures Redis
- Configures RabbitMQ
- Deploys AI services

No manual Docker configuration is required.

13. Select Port Profile

At the beginning of the installation, you will be prompted to select the network port profile.

FreeFalcon AI Port Selection

- 1) Standard - HTTP 80 / HTTPS 443
 - 2) Alternative - HTTP 8000 / HTTPS 8443
- Select ports [1]:

Option 1 – Standard (Default)

Web interface on HTTP port 80 and HTTPS port 443. Choose this when ports 80/443 are free (this is the default).

Option 2 – Alternative

Web interface on HTTP port 8000 and HTTPS port 8443. Choose this when ports 80/443 are already in use.

Note: regardless of the selected profile, the platform is always served over HTTPS. Plain HTTP requests are automatically redirected to HTTPS. After selecting the desired option, the installation continues automatically.

14. Automatic Installation

After selecting the port profile, the installation proceeds without further user interaction.

Depending on the server hardware and Internet connection, the installation may take several minutes.

15. Verify GPU Utilization

After AI Engine Initialization has completed, verify that the platform is utilizing the NVIDIA GPU.

Run:

```
nvidia-smi
```

Expected output:

```
+-----+
| NVIDIA-SMI 570.172.08      Driver Version: 570.172.08   CUDA Version: 12.8   |
|=====+=====+=====+=====+=====+=====+=====+=====+=====+=====+
| GPU Name                   Memory Usage | GPU-Util    |
|=====+=====+=====+=====+=====+=====+=====+=====+=====+=====+
| 0 NVIDIA GeForce RTX 3060  4265 MiB / 12288 MiB | 74%         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Processes:

PID	Type	Process Name	GPU Memory
25483	C	python3	3920 MiB
26177	C	freefalcon-ai	345 MiB

The presence of running AI processes and allocated GPU memory confirms that FreeFalcon AI is successfully utilizing the NVIDIA GPU.

16. Access the FreeFalcon AI Platform

Once the installation and AI Engine Initialization are complete, open a supported web browser.

If the Standard profile was selected:

`https://<server-ip>` (HTTP port 80 is automatically redirected to HTTPS 443)

If the Alternative profile was selected:

`https://<server-ip>:8443` (HTTP port 8000 is automatically redirected to HTTPS 8443)

Replace <server-ip> with the IP address or hostname of your FreeFalcon AI server.

Supported client operating systems:

- Windows
- Linux
- macOS

Supported web browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge
- Apple Safari

No client software installation is required.

17. Installation Complete

Congratulations!

Your FreeFalcon AI Platform has been successfully installed and is ready for configuration and AI analytics deployment.

For hardware sizing recommendations, deployment best practices, and system requirements, refer to the FreeFalcon AI System Requirements Guide.