

FreeFalcon AI Platform

System Requirements Guide

Overview

This document outlines the minimum and recommended hardware and software requirements for deploying the FreeFalcon AI Platform. The platform supports deployment on dedicated physical servers and virtual machines. While CPU-based processing is supported for lighter workloads, GPU acceleration is highly recommended for enterprise deployments.

Operating System

Operating System	Download
Ubuntu 22.04 LTS (Server or Desktop)	https://releases.ubuntu.com/jammy/ubuntu-22.04.5-desktop-amd64.iso
Ubuntu 24.04 LTS (Server or Desktop)	https://releases.ubuntu.com/noble/ubuntu-24.04.4-desktop-amd64.iso

Both Server and Desktop editions of Ubuntu 22.04 LTS and 24.04 LTS are fully supported. The download links above point to the Desktop images; Server images are available from the same release pages.

Installation Requirements

- Format all available disks using the **ext4** filesystem.
- Use an **SSD** as the system disk (minimum 128 GB).
- Use a dedicated **HDD** for event recording storage whenever possible.
- Separate system and recording storage for maximum stability and performance.

Minimum Hardware Requirements

Component	Minimum Requirement
CPU	Modern 4-core CPU, 2 GHz+, AVX2 or AVX-512
RAM	8 GB
System Disk	128 GB SSD
Event Storage	Dedicated HDD (recommended)
GPU	NVIDIA Pascal or newer (recommended)
File System	ext4

GPU Requirements

FreeFalcon AI strongly recommends utilizing an NVIDIA GPU for AI processing. GPU acceleration provides significantly higher throughput, lower latency, reduced CPU utilization and improved scalability.

Supported NVIDIA Architectures: Pascal, Volta, Turing, Ampere, Ada Lovelace, Hopper and Blackwell.

Network Ports

External Ports — the installer offers two port profiles. Choose one during installation.

Profile	HTTP	HTTPS
Standard (default)	80 TCP	443 TCP
Alternative	8000 TCP	8443 TCP

The platform is always served over HTTPS; HTTP requests are automatically redirected to HTTPS.

Internal Docker Service Ports (container-internal)

Port(s)	Service
80, 443	Internal Web Services
9001, 9100	Platform Services
6379	Redis Cache
3306, 33060	MySQL
4369	RabbitMQ Erlang Mapper
5671, 5672	RabbitMQ Messaging
15671, 15672	RabbitMQ Management
15691, 15692	RabbitMQ Metrics
25672	RabbitMQ Cluster Communication

These ports are used only inside the Docker network between platform services. They are not published to the host. In particular, MySQL is reachable only within the container network (on port 3306) and is not exposed on the host.

Client Workstation Requirements

The FreeFalcon AI client is entirely **web browser-based** and requires no software installation.

Supported operating systems: Windows, Linux, macOS.

Supported browsers: Google Chrome, Mozilla Firefox, Microsoft Edge, Apple Safari.

Physical Server vs Virtual Machine

FreeFalcon AI supports deployment on dedicated physical servers and virtual machines. Dedicated physical servers are recommended for production deployments requiring maximum AI performance, stability and scalability.

If virtualization is used, configure the VM with AVX2 support, direct CPU allocation whenever possible, GPU passthrough (where supported) and high-performance SSD storage.

Hardware Sizing

Hardware requirements depend on the number of cameras, camera resolution, frame rate, enabled AI analytics, AI processing frequency, event retention period and number of simultaneous users.

For accurate sizing, please download and use the **FreeFalcon AI Hardware Sizing Calculator**.

Best Practices

- Ubuntu 24.04 LTS for new deployments.
- SSD for operating system.
- Dedicated HDD for recordings.
- NVIDIA GPU acceleration.
- Gigabit or faster networking.
- Keep Ubuntu, Docker and NVIDIA drivers up to date.
- Restrict access to internal Docker service ports using a firewall.

Recommendation

FreeFalcon AI can operate on CPU-only hardware for smaller deployments. However, GPU acceleration is strongly recommended for production and enterprise-scale installations to maximize AI performance, support more simultaneous analytics and provide the best overall user experience.