Datasheet

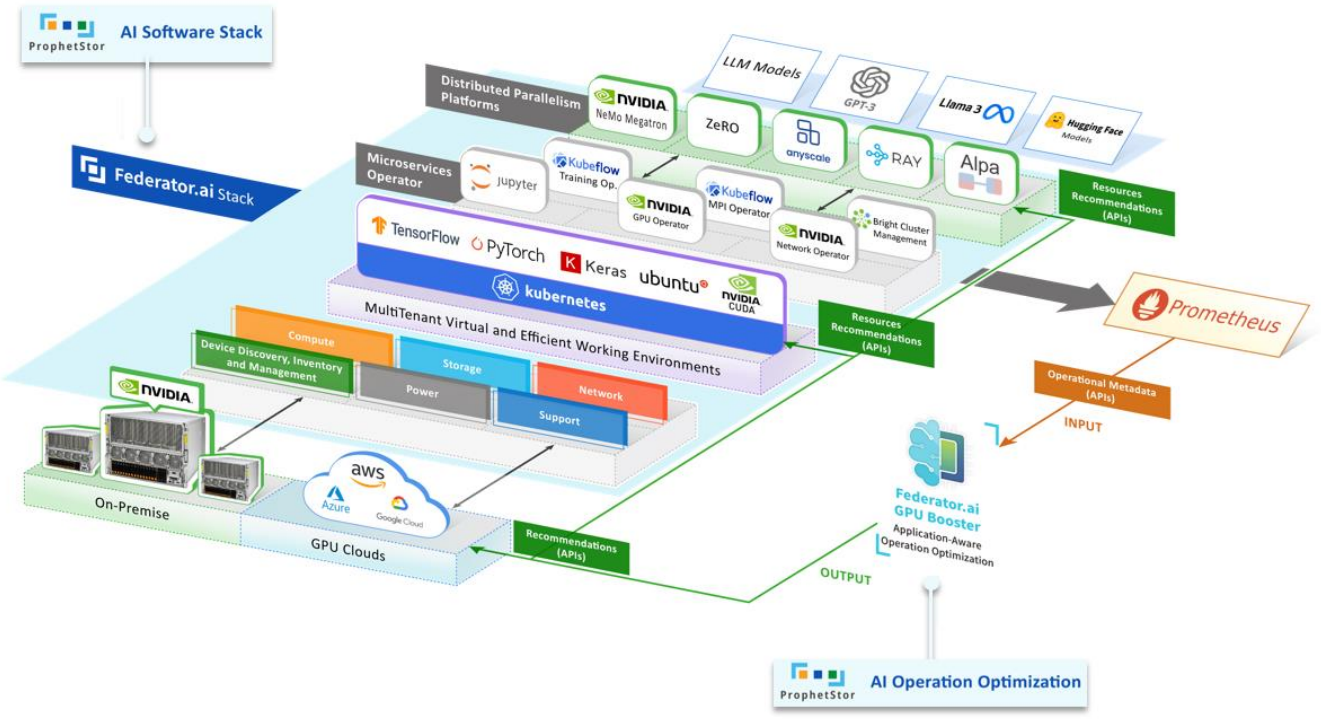
Federator.ai Stack™

Accelerating AI Innovation with a Streamlined GPU Ecosystem Installation

- **A comprehensive suite of GPU optimization tools that prevents reduced performance, compatibility issues, and installation errors.**
- **Rapid time-to-online setup, deploying the GPU-powered AI training ecosystem in tens of minutes.**
- **An included step-by-step video guide walks through the software installation process for efficient AI/ML training.**

Transforming AI Workflows

Federator.ai Stack is engineered to empower AI and ML innovations, offering a comprehensive suite of tools for optimizing GPU computing resources with a streamlined installation process that completes in only tens of minutes. Tailored for researchers, data scientists, and IT professionals, it streamlines the deployment and management of AI applications, ensuring you can focus on breakthroughs while we handle the complexities of infrastructure optimization.



How Federator.ai GPU Booster Works to Optimize GPU Resource Usage for AI/ML Training

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

Benefits

- **All-in-One AI Tool Suite:** Federator.ai Stack provides everything necessary for productive AI/ML training— from containerized platforms and open-source monitoring to powerful transformers for large model training—all in a single, up-to-date installation.
- **Seamless Integration with the NVIDIA Ecosystem:** Integrated with leading high-end NVIDIA GPUs, including A100, H100, H200, and GB200, Federator.ai Stack enables efficient parallel AI/ML training across a heterogeneous GPU environment that supports multiple GPU generations.
- **Simple, Guided Installation Process:** Federator.ai Stack offers a streamlined, guided installation that ensures a complete and up-to-date setup, preventing performance issues, compatibility challenges, and security vulnerabilities due to missing or outdated components.
- **Quick Time to Deployment:** Federator.ai Stack brings scalable AI clusters online in tens of minutes, enabling organizations to focus on innovation with the out-of-the-box Federator.ai GPU Booster rather than on setup.

Prerequisites

Before initiating the installation, ensure your system meets the following hardware requirements:

- **CPU:** Intel Xeon E3 or above
- **Memory:** Minimum of 64GB
- **GPU:** Nvidia H100/H200/GB200 recommended
- **Network:** At least 1 NIC
- **Local Storage:** 1TB SSD recommended
- **Persistent Storage:** (Optional) 500GB NFS storage

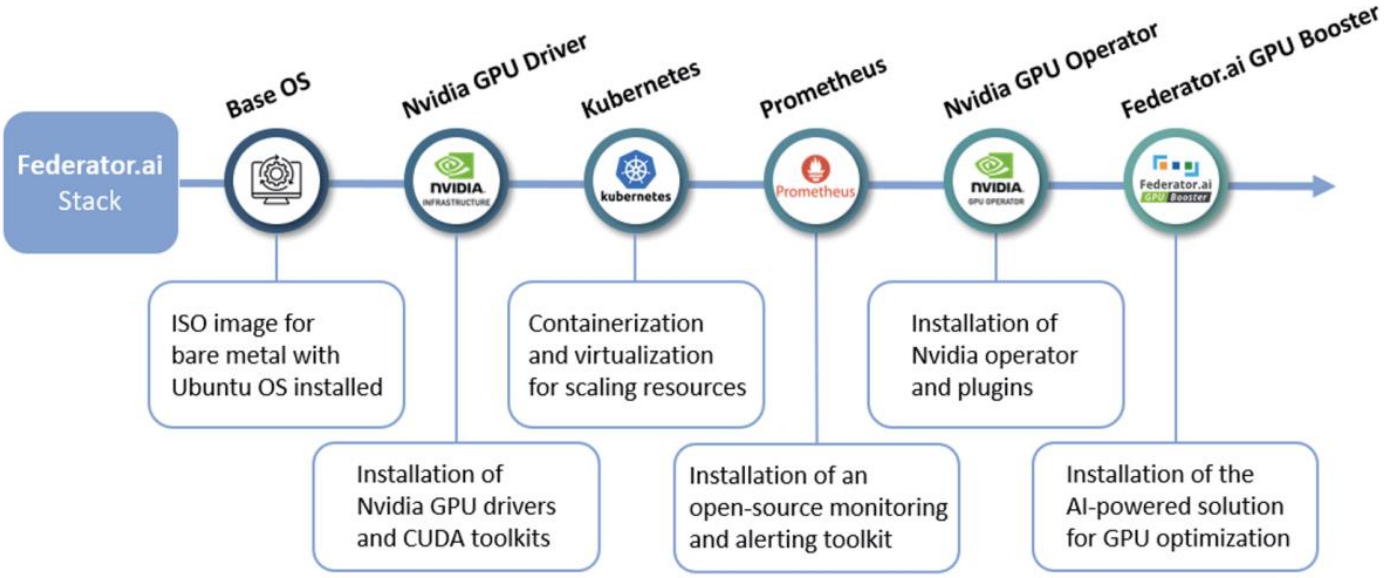
Installation Process

An ISO image is provided to install the base OS for a bare metal GPU server. For a GPU server that has already been installed with a base OS, the Federator.ai Stack launcher script is provided to ease the installation process.

- **Base OS:** For a bare metal GPU server, boot the GPU server from the Federator.ai Stack ISO image, and an Ubuntu base OS will be installed.

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

- **Nvidia GPU Driver:** Nvidia GPU drivers and CUDA Toolkits will be installed on Ubuntu OS.
- **Kubernetes:** A Kubernetes cluster will be created if there is no existing Kubernetes cluster on the GPU server.
- **Prometheus:** An open-source monitoring and altering toolkit for monitoring microservices and containers. This will be installed in the Kubernetes cluster.
- **Nvidia GPU Operator:** The Nvidia GPU Operator for Kubernetes, which includes the Nvidia device plugin, will be installed in the Kubernetes cluster.
- **Federator.ai GPU Booster:** Federator.ai GPU Booster will be installed in the Kubernetes cluster.



The Installation process of Federator.ai Stack

Software Packages

The Federator.ai Stack includes the following:

- **Nvidia GPU Drivers and CUDA Toolkits:** Essential tools and drivers for GPU acceleration.
- **Automatic Kubernetes Cluster Creation:** For seamless deployment and scaling of containerized applications.
- **Nvidia GPU Operator:** Automates GPU resource management within Kubernetes environments.
- **Prometheus:** Collects Kubernetes system and container metrics as well as Nvidia GPU metrics.

Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com

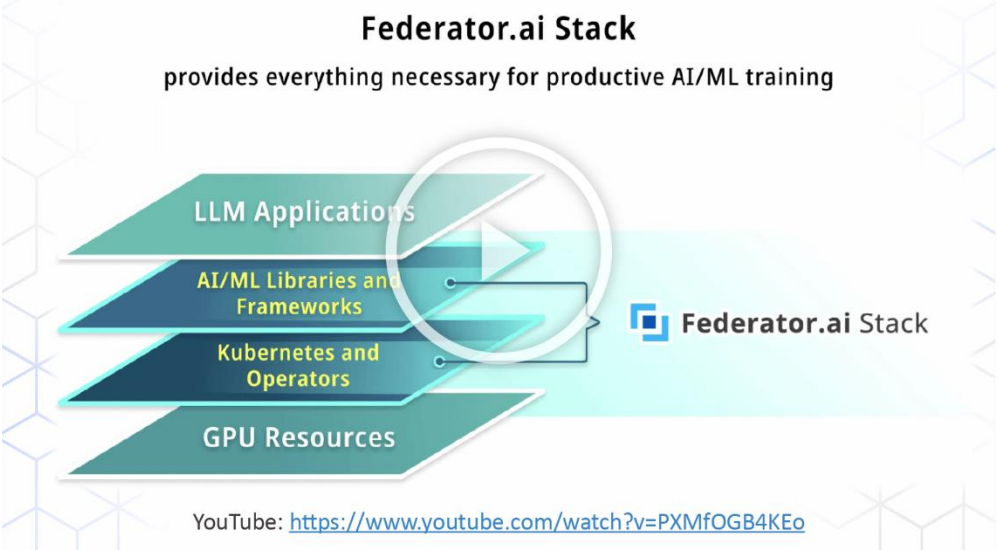
- **Federator.ai GPU Booster:** Enhances GPU efficiency for AI and ML workloads.
- **Optional AI/ML Frameworks and Toolkits:** A selection of frameworks and tools for AI development, ready for integration.
  - Kubeflow
  - KubeRay
  - Nvidia Nemo Framework

Included Software Versions

With Federator.ai Stack, the following versions are deployed:

- **Base OS:** Ubuntu Server 22.04 LTS
- **Kubernetes:** v1.30
- **Nvidia Driver:** 560.35.03
- **CUDA Toolkit:** 12.6
- **Nvidia GPU Operator:** v24.6.2
- **Prometheus:** 33.2.1
- **Federator.ai GPU Booster:** v5.2.1

Video | Federator.ai Stack optimizes the Time-to-Online of GPU servers



Copyright © 2012-2026 ProphetStor Data Services, Inc. All rights reserved. | https://prophetstor.com