

## Inspur NF5488A5

#1 AI Training Server



Inspur NF5488A5 is an exciting new addition to the comprehensive Inspur AI portfolio. Ranking first in single server performance in MLPerf o.7, it supports 8 fully interconnected NVLink 3.0 NVIDIA A100 GPUs, and 2 AMD Rome processors enabling PCIe4.0 in a 4U space, NF5488A5 delivers outstanding AI performance of 5 petaFLOPS and ultra-high data throughput, making it ideal choice for a series of AI scenarios such as intelligent image, video and voice processing, financial analysis and virtual assistant.

### Extreme AI Performance

Equipped with 2 AMD Rome CPU, it enables PCIe4.0 and offers superb AI performance of 5 petaFLOPS

### Extreme Hardware Design

It can be deployed in a wide range of data centers with its 4U design, greatly helps to lower cost and increase operation efficiency.

### Extreme CPU/GPU-GPU Communication

Comparing the previous generation, the transformation speed is doubled. It has a high bandwidth of GPU-GPU Communication of 600GB/s.

### Accelerate AI Application

It is suitable for various AI applications such as intelligent image, video and voice processing, financial analysis and virtual assistant.

## Product Specification

<b>Model</b>	<b>NF5488A5</b>
GPU	1* HGX A100 8GPU
Processor	2*AMD® EPYC® Rome/Milan Processor (Up to 240W TDP)
Memory	32*DDR4 2933MT/s RDIMMs//LRDIMM, up to 4 TB
Storage	8*2.5" SAS/SATA (or 4*NVMe+4* SATA/SAS) , 4* NVMe M.2, 2* SATA M.2
Built-in PCIe	4*low profile PCIe x16
Front I/O	2*USB 3.0 port, 1*VGA port, 1*RJ45 management port
Ethernet card	10G Ethernet Optical Interface
Management	Built-in Aspeed2500 BMC module, support IPMI, SOL, KVM Over IP, VM, etc.
RAID	Support 3508 Raid Card
Operating system	Red Hat Enterprise 7.8 64bit,CentOS 7.8; Ubuntu 18.04
Cooling	N+1 Redundant hot swap fans
PSU	3+1 Redundant PSUs
Chassis	4U, W*H*D 448mm*175.5mm*850mm