

inspur

inspur



M6

Inspur® M6 Server
Product Brochure

Inspur Electronic Information Industry Co.,Ltd.
en.inspur.com Email: iei-hwtjb@inspur.com

Inspur Systems, Inc.
www.inspursystems.com Email: sales@inspursystems.com

Inspur Germany GmbH
en.inspur.com Email: eusales@inspur.com



Compute Power Defines Productivity

Overview

Currently, digital transformation in the world is featured with accelerated innovation and the digital economy drives future growth through new business formats and models. The essence of the digital economy is the wide application of the latest ICT technologies that are penetrating every link of the economic activities, such as production, logistics, distribution, exchange, and consumption. On the one hand, technology empowers business efficiency and upgrading; on the other hands, it induces the generation and flow of massive data from all links. In the connected era, data is the core asset to bring the most value to the digital economy.

Computing infrastructure (servers, storage, networks, etc.) is already one of the core instruments to process the numerous data and facilitates the implementation of cloud computing, big data and AI. Relying on continuous support, enterprises are blessed with even convenient and efficient transaction and circulation of commodities. This also injects power into the innovation and upgrading of all kinds of new business models, such as video streaming, social networking, e-commerce, sharing economy. Compute power has become the core production factor in the digital economy of a nation.

To help customers compute, process, analyse, store and migrate massive amounts of data, and to provide persistent power for the day-to-day operation of enterprises, Inspur M6, a new generation server, provides industry-leading solutions highlighted by versatility, robustness, and agility. It serves the customer in various application scenarios, such as multi-cloud, edge computing, and AI.

Four Features

Break-thru Innovations

- In a limited space, continuously increase the computing density, provide top-notch I/O scalability, and eliminate the bottleneck of data transmission.
- 1U space supports 32 x NVMe SSDs, meaning outstanding storage and transmission efficiency to satisfy high-density workloads with stringent requirements for latency and bandwidth.

Precision Engineering

- After 28-month hard work, Inspur M6 comes into being with a full improvement in designs, components and systems. New features like environment-awareness, re-defined power standards, automated operations & maintenance tools make the server smarter and easier to use.

Security

- **Hardware:** Double-layer PSU protection and BIOS/BMC redundancy designs ensure absolute system security in terms of power supply and remote access.
- **Firmware:** FPGA enhances security with the Hardware Root of Trust and enables lifecycle protection to BIOS, BMC and other firmware.
- **Systems:** Space isolation, industrial components selection, redundant development process, system derating and other measures are leveraged to guarantee at least 8-year trouble-free and stable system operations.

Openness

- Inspur, a leader in open computing and a key member of the three major open computing organizations, supports various open standards and has been committed to building an open computing ecosystem. From products, standards to leading projects, Inspur constantly enhances its influence in the open computing ecosystem.

Inspur® M6 Product Family

Contents

A Computing Server

- 06 NF5280M6
- 07 NF5180M6
- 08 NF5270M6
- 09 NF5260M6/FM6



NF5280M6



NF5180M6



NF5270M6



NF5260M6/FM6

B Storage Server

- 10 NF5466M6
- 11 NF5266M6



NF5466M6



NF5266M6

C Multi-Node Server

- 12 i24M6
- 13 i48M6
- 14 i24LM6



i24M6



i48M6



i24LM6

D Mission Critical Server

- 15 NF8260M6
- 16 NF8480M6



NF8260M6



NF8480M6

E AI Server

- 17 NF5468M6



NF5468M6

F Rack-scale Systems

- 18 ORS6000S



ORS6000S

G Infrastructure Management Software Suite

- 19 ISPIM
- 20 ISREST, ISA
- 21 ISQP, ISIB



ISPIM



ISREST



ISA



ISQP



ISIB

NF5280M6

Rackmount server for all scenarios



NF5280M6 is a 2U dual-socket high-end flagship rackmount server, featured with robust computing performance and ultimate compatibility and scalability. Meeting the configuration requirements of various industries, it is suitable for data analysis and processing, distributed storage in deep learning training, etc.

Features

Versatile and full-scenario coverage

■ 2U space and flexible design allow both general-purpose and heterogeneous computing in storage and I/O applications.

Data security and operation reliability

■ The latest PFR/SGX technology and double flash system brings system reliability to a higher level.

Leading designs

■ Release more I/O resources through 13×PCIe expansions, RAID-enabled M.2/E1.S, and storage/OCP daughter card upgrade.

Specifications

Model	NF5280M6	
Processor	2×Intel® Xeon® Ice Lake scalable processors	
Memory	32×memory slots, up to DDR4-3200, up to 16×Intel® Optane™ PMem	
Configuration	General	AI
Storage	Front	24×2.5-inch SATA/SAS/NVMe (hot-plug) 12×2.5-inch SATA/SAS (hot-plug)
	Rear	4×2.5-inch SATA/SAS/NVMe SSDs (hot-plug) 4×3.5-inch SATA/SAS/NVMe SSDs (hot-plug) + 2×E1.S or M.2 SSDs
	Built-in	4×3.5-inch SATA/SAS (hot-plug)
I/O Expansion	Up to 13×PCIe slot 1×RAID Mezz card 1×OCP3.0×16 card	4×FHFL or 8×FHHL GPUs/graphics cards 1×RAID Mezz card 1×OCP3.0×16 card
	PSU	1+1 redundant power supplies; 550W/800W/1300W/1600W/2000W standard CRPS power supply
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>	

NF5180M6

High-density rackmount server



NF5180M6 is a 1U dual-socket flagship server that is designed to maximize performance, scalability, and density. It is suitable for HPC and virtualization and various computing-intensive scenarios and meets the deployment requirement of high-density data centers.

Features

All-flash and density optimization

■ Fully adapt to the latest storage media in the industry, support up to 32×hot-pluggable E1.S SSDs. ■ 1U space enables an extreme cache pool of up to 256T and 10-million-level IOPS for business application.

Flexible I/O expansion

■ Onboard dual-port 10Gb or 1Gb NIC available to release I/O resources. ■ On-demand PCIe customization meets the requirement of different customers.

High-density ultimate cooling

■ Introduce EVAC to reduce the TCO of the data center and avoid potential risks of performance balance caused by uneven heat dissipation.

Specifications

Model	NF5180M6		
Processor	2×Intel® Xeon® Ice Lake scalable processors		
Memory	32×memory slots, up to DDR4-3200, up to 16×Intel® Optane™ PMem		
Storage	General	12×3.5-inch	32×E1.S
	Font	Front	Front
	Rear	Rear	Built-in
I/O Expansion	Up to 3×standard PCIe slots, 1×OCP3.0 card		
PSU	1+1 redundancy power supplies, 550W/800W/1300W		
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>		

NF5270M6

Rackmount server optimized for enterprise applications



In the spirit of streamlined design, NF5270M6 is a mid-range 2U server tailored for small-scale virtualization, database, office OA and other scenarios.

Features

- | | |
|--|--|
| <p>Flexible configuration and simple structure</p> <ul style="list-style-type: none"> Support up to 25×2.5-inch SAS / SATA SSDs in the front and 4×2.5-inch SAS / SATA SSDs at the rear. Support up to 12×NVMe SSDs, optional OCP3.0 and multiple network interfaces (1G / 10G / 25G / 100G), up to 5×PCIe4.0 expansion (4 ×PCIe×16 + 1×PCIe×8). | <p>Optimized design and noise reduction</p> <ul style="list-style-type: none"> Thanks to the effective design of the air hood and combined with fan cooling, the machine achieves the same performance as the previous version and has the noise reduced by 20%. |
| <p>Easy to use and maintain</p> <ul style="list-style-type: none"> Modular design ensures easy assembly. Toolless disassembly and maintenance save efforts. | |

Specifications

Model	NF5270M6
Processor	2×Intel® Xeon® Ice Lake scalable processors
Memory	16×memory slots, up to DDR4-3200, up to 12×Intel® Optane™ PMem
Storage	<p>Front 24×2.5-inch SATA/SAS/NVMe SSDs (up to 12×NVMe per system), hot-plug 25×2.5-inch SATA/SAS SSDs (hot-plug) 12×2.5-inch SATA/SAS SSDs (hot-plug) 12×3.5-inch SATA/SAS/NVMe SSDs (hot-plug)</p> <p>Rear 4×2.5-inch SATA/SAS/NVMe SSDs (hot-plug) 4×2.5-inch SATA/SAS/NVMe SSDs (hot-plug) + 2×E1.S or M.2 SSDs</p>
I/O Expansion	Up to 5×standard PCIe, 1×RAID Mezz card, 1×OCP3.0 (optional), and dual-port 1Gb NIC, 4×GPUs
PSU	1+1 redundant power supplies, 550W/800W/1300W Platinum
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

NF5260M6/FM6

Rackmount server optimized for data centers



NF5260M6/FM6 is Inspur's first dual-socket server that supports on-demand front and rear I/O. With a flexible and modular design that can respond to various Internet applications, it is an ideal choice for data centers.

Features

- | | |
|---|---|
| <p>Flexible I/O design</p> <ul style="list-style-type: none"> Non-onboard PCIe slot design satisfies more non-PCIe I/O customizations; supports up to 2×hot-pluggable OCPs, free up more I/O resources and ensure uninterrupted network services. | <p>Ultimate architecture design</p> <ul style="list-style-type: none"> Unique front I/O design proposes innovative measures to isolate the cold and hot aisles in the data center, which increases the lifecycle of heat-sensitive components, facilitates maintenance, and effectively solves the wiring problem in the data center. |
| <p>Flexible deployment mode</p> <ul style="list-style-type: none"> Centralized power supply modules and CRPS power supplies can be switched freely. Support the stand-alone node and rack-scale delivery and ensure hassle-free deployment. | <p>Innovative cooling scheme</p> <ul style="list-style-type: none"> The reserved liquid cooling solution reduces PUE in the data center to the limit; air cooling scheme is available as needed. |

Specifications

Model	NF5260M6	NF5260FM6
Processor	2×Intel® Xeon® Ice Lake scalable processors	
Memory	32×memory slots, up to DDR4-3200, up to 16×Intel® Optane™ PMem	
Storage	<p>Front: 12×3.5-inch SAS/SATA/NVMe SDDs (hot-plug) 24×2.5-inch SAS/SATA/NVMe SSDs (16×NVMe SSDs) (hot-plug); No HDD.</p> <p>Rear: 4×2.5-inch SAS/SATA/NVMe SSDs (2×NVMe SSDs) (hot-plug)</p> <p>Built-in: 2×M.2 SSDs</p>	<p>Front: 9×3.5" SAS/SATA/NVMe SSD (2×NVMe SSD) (hot-plug); 16×2.5-inch SATA/SAS/NVMe SSDs (hot-plug); 14×2.5-inch SATA/SAS/SSD + 2×NVMe SSDs; 8×2.5-inch SAS/SATA/NVMe SSDs (4×NVMe SSDs) (hot-plug)</p> <p>Built-in: 2×M.2 SSDs</p>
I/O Expansion	<p>Internal RAID controller card</p> <p>Rear I/O: 2×OCP3.0 + 4×PCIe×16 1×OCP3.0+2×PCIe×16 + 1×PCIe×8 4×PCIe×16 + 2×PCIe×8</p>	<p>Front I/O: 1×OCP3.0+2×PCIe×8 + 1×PCIe×16 for 9×3.5-inch 1×OCP3.0+2×PCIe×8 + 1×PCIe×16 for 16×2.5-inch 1×OCP3.0+4×PCIe×16 + 2×PCIe×8 for 8×2.5-inch</p> <p>Front I/O: 2×built-in riser cards: ×8×2 + ×16×1</p>
PSU	1+1 redundancy power supplies, 550W/800W/1300W, centralized power supply	
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>	

NF5466M6

General-purpose rackmount storage server



NF5466M6 is a 4U dual-socket storage server featured by high storage capacity, robust computing performance and excellent I/O expansion. It is designed for scenarios such as warm/cold data storage, video surveillance storage, big data storage, cloud storage pool, etc.

Features

High storage capacity

■ Support up to 46×3.5-inch drives and an extraordinary storage capacity of up to 828TB (18TB per drive) to meet the needs of local data storage.

Excellent I/O scalability

■ Excellent I/O scalability: Support up to 13 x PCIe expansions to further enhance I/O scalability. ■ Leverage a variety of storage, network, and GPU modules to serve different application scenarios. ■ Flexible configuration tailored as needed eliminate I/O bottlenecks and improve overall performance.

Robust computing performance

■ Adopt the new generation of processors with a maximum performance increase of 40%. Support 2×FHFL GPUs or 8×FHHL GPUs to satisfy the needs in IVA scenario and enable efficient local data analysis and real-time AI implementation.

Specifications

Model	NF5466M6
Processor	2×Intel® Xeon® Ice Lake scalable processors
Memory	32×memory slots, up to DDR4-3200, RDIMM/Intel® Optane™ PMem/NVDIMM
Storage	Front 24×3.5-inch SATA/SAS SSDs (hot-plug) Rear 16×3.5-inch SATA/SAS SSDs (hot-plug) 4×2.5-inch SATA/SAS/NVMe SSDs (hot-plug) 2×SATA M.2 or 2×E1.S SSDs Built-in 6×3.5-inch SATA/SAS SSDs in the middle Up to 2×TF cards Up to 16×U.2 NVMe SSDs
I/O Expansion	Up to 13×PCIe slots, including 1×PCIe RAID controller card and 1×OCP3.0 slot
PSU	1+1 redundant platinum power supplies High-voltage DC and titanium power supply
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

NF5266M6

High-density rackmount storage server



NF5266M6 is a 2U dual-socket storage server that adopts an innovative three-tier storage architecture. It has improvement and optimization in storage and computing performance, flexible configuration, and intelligent management. It is suitable for big data, CDN, hyper-converged, distributed storage, etc.

Features

High-density storage

■ 2U space supports up to 26×3.5-inch drives. Compared with the traditional 2U12 drive model, the storage density is increased by 116%. ■ Excellent cache capacity (SSD:HDD=1:3) meets the requirement of massive and high-density storage and large cache size.

High bandwidth

■ I/O expansion capability is increased by 75%. ■ The transfer rate of PCIe4.0 is doubled and reaches 16GT/s. ■ Support the 200Gb network and ensure a low-latency experience.

Robust performance

■ 2×scalable processors enable even outstanding CPU computing power. ■ Compared with the previous generation, this machine has the performance boosted by up to 44%, ideal for demanding scenario of computing power in big data analysis.

Intelligent environment sensing

■ Software and hardware design enables real-time monitoring of the temperature, airflow, and air pressure of the operating environment, and intelligently cooling strategy; improves the service life of the drive and ensures the storage server is in the best operating state.

Specifications

Model	NF5266M6
Processor	2×Intel® Xeon® Ice Lake scalable processors
Memory	32×memory slots, up to DDR4-3200, RDIMM/Intel® Optane™ PMem/NVDIMM
Storage	Front up to 24×3.5-inch SAS/SATA SSDs (hot-plug) Rear up to 24×2.5-inch SAS/SATA/NVMe/E1.S SSDs or 2×3.5-inch drives (hot-plug)
I/O Expansion	Up to 7×standard PCIe slots, 1×OCP3.0 slot
PSU	1+1 redundant platinum power supplies High-voltage DC and titanium power supply
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

C Multi-Node Server

i24M6

High-density computing multi-node server



i24M6, a multi-node high-density server, perfectly expresses high-density, efficiency, reliability, and intelligence. The compact 2U space reduces the cost and input in three dimensions: data center resources, energy consumption, and deployment costs, which is the ideal solution to reduce the TCO of the data center.

Features

High density and efficiency through optimized architecture

■ 2U space with 4× dual-socket nodes means a computing density three times higher than the 2U rackmount server. ■ Modular node design enables fast deployment with the overall time expenses reduced by over 50%.

Stable, reliable and intelligent management

■ All nodes share redundant power supplies and fans, which ensures stable system operation. ■ Power Capping helps to achieve machine-level efficiency and energy saving. ■ Support CMC/BMC and cascading management.

Ultimate performance and flexible adaptation

■ In all-flash mode, the entire server supports the configuration of 24 × NVMe SSDs, accelerating operation by 10 times. ■ 2 × built-in M.2 SSDs support both software and hardware RAID. ■ It also supports liquid cooling for processors and various I/O expansion.

Specifications

Model	i24M6 / NS5160M6
Processor	2 × Intel® Xeon® Ice Lake scalable processors
Memory	16 × memory slots, up to DDR4-3200, Intel® Optane™ PMem
Storage	24 × 2.5-inch SATA / SAS / NVMe SSDs (hot-plug) 12 × 3.5-inch SATA / SAS / NVMe SSDs (hot-plug) Up to 2 × built-in M.2 SSDs
I/O Expansion	Single node supports 1 × OCP3.0 slot and 2 × PCIe slots
PSU	2 × 2000W power supplies (Platinum)
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

C Multi-Node Server

i48M6

High-density multi-node server with load balancing



i48M6 is a multi-node modular server optimized for high-density data centers. Based on the latest Intel Whitley platform, it can deploy 2 to 8 × dual-socket computing/storage nodes in a standard 4U rack, which stands for flexible scalability and meets the requirement of various applications. It is suitable for all kinds of scenarios, such as HPC, cloud computing, big data analysis, and mass storage.

Features

Ultimate computing and storage density

■ 8 × NS5480M6 (computing nodes) can be deployed in the 4U space, which is 4 times the deployment density of a standard 2U rackmount server. ■ 72 × 3.5-inch drives can be deployed in the 4U space (storage node - NS5486M6 and drive tray - NS5486JD).

I/O balancing and enhanced application performance

■ I/O evenly distributed on two CPUs to maximize the application efficiency of the NUMA architecture and reduce response delay; the overall performance improved by over 20%; 4U space supports up to 48 × I/O expansions to satisfy the data transmission in large-scale applications.

Modular design for flexible expansion

■ Adopt full modular design and flexible expansion of different modules, such as chassis, node, front and rear I/O, management module, and integrated network. ■ Unified architecture and flexible design ensure the smooth expansion of data centers from small to large scale, saving the initial infrastructure investment to the greatest extent.

Precision engineering for secure and reliable operation

■ Redundant power supplies and cooling fans are shared by the system. ■ Inspur's unique control technologies and advanced air-cooling solution enable the optimal working condition to save energy and allow stable operation at the system level. ■ Unique drive drive tray and advanced slide rail design greatly improve maintenance convenience and data security.

Specifications

NS5480M6: high-density computing node	
Processor	2 × Intel® Xeon® scalable processors
Memory	16 × memory slots, up to DDR4-3200, Intel® Optane™ PMem
Storage	Up to 4 × 2.5-inch SATA/SAS/SSD/NVMe SSDs (hot-plug), 2 × M.2 SSDs
PCIe Expansion	Up to 4 × standard PCIe slots, 1 × OCP slot, 1 × RAID card slot
NS5484M6: storage/computing-balanced node	
Processor	2 × Intel® Xeon® scalable processors
Memory	16 × memory slots, up to DDR4-3200, Intel® Optane™ PMem
Storage	Up to 2 × 2.5-inch SATA/SAS/SSD/NVMe SSDs (hot-plug), 12 × 3.5/2.5-inch drives (hot-plug), 2 × M.2 SSDs
PCIe Expansion	Up to 3 × standard PCIe slots, 1 × OCP slot, 1 × RAID card slot
NS5486M6: mass storage node	
Processor	2 × Intel® Xeon® scalable processors
Memory	16 × memory slots, up to DDR4-3200, Intel® Optane™ PMem
Storage	Up to 2 × 2.5-inch SATA/SAS/SSD/NVMe SSDs (hot-plug), 36 × 3.5/2.5-inch drives (hot-plug), 2 × M.2 SSDs
PCIe Expansion	Up to 3 × standard PCIe slots, 1 × OCP slot, 1 × RAID card slot
NS5486JD: storage expansion module	
Storage	Built-in: 1 × PCIe SAS/RAID daughter card Drive tray: 36 × 3.5-inch SATA/SAS/SSDs (hot-plug and 2.5-inch available)
Chassis	
PSU	2+2 / 3+1 redundant power supplies 1300W/1600W/2000W Platinum 100V-240V AC, 240V HVDC
Fan	5 × fan modules, N+1 redundancy, backflow preventer design
PCIe Expansion	8 × PCIe slots (hot-plug) at the rear
Management/Network	Rear: 2 × redundant CMC modules for centralized management of the entire chassis Rear: aggregation switch module

C Multi-Node Server

i24LM6

High density liquid cooled server



i24LM6 is Inspur's latest 2U-4Node liquid cooled rack server optimized for high density datacenter and HPC, with flexible configuration of air and liquid cooling. It is suitable for a wide variety of compute-intensive workloads including HPC, high-performance data analytics and more. With industry's first support of 3rd Gen Intel® Xeon® Scalable processors, the liquid cooling design of CPU, DIMM and VR are fully optimized to deliver unmatched compute and cooling performance. Based on i24LM6, Inspur launches the unique Air/Liquid Cooling CDU solution to help quickly deploy liquid cooling without upgrading original facilities, greatly increase deployment efficiency and reduce operational costs.

Features

Unrivaled performance

- Support 4x 2-socket compute nodes in 2U, industry's highest density
- Up to 8x 3rd Gen Intel® Xeon® Scalable processors, enabling overclocking under liquid cooling with a compute performance increase of 10%
- Each node support 16x 3200MT/s DIMM slots, featuring ultra-high compute density and storage capacity

Highly-efficient liquid cooling

- Warm water cooling, heat emission efficiency up to 80%
- liquid cooling covering CPU, DIMM and VR, PUE≤1.1
- Reduce specialized cooling facilities in datacenter and related cost

Product Specification

Chassis Model	i24LM6	Node Model	NS5160LM6
Form Factor	2U 4-Node	Form Factor	1U, 2-Socket Compute Node
Storage	Conf.1: 24×2.5" NVMe+ 8×2.5" SAS/SATA Conf.2: 16×2.5" SAS/SATA/NVMe+8×2.5" SAS/SATA Conf.3: 12×3.5" SAS/SATA	Processor	2×3rd Gen Intel® Xeon® Scalable processor
PSU	support 1+1 redundancy power supplies, 2×2000W hot-plug 80Plus Platinum, support DPC	Storage	Conf.1: 4×2.5" SAS/SATA/NVMe + 2×2.5" SAS/SATA Conf.2: 6×2.5" NVMe Conf.3: 3×3.5" SAS/SATA 2×SATA/PCIe M.2
Cooling	Chassis: 4×80mm system fan, N+1 Redundant Node: Liquid cooling	Memory	16× memory slots, up to DDR4-3200, RDIMM/LRDIMM
Working Temperature	5°C-35°C. Refer to the <i>Technical White Paper</i> .	I/O Expansion	1×PCIe 4.0 x16 slot 1×OCP3.0 Card, support NCSI 1×RAID Mezz card

D Mission Critical Server

NF8260M6

High-density quad-socket rackmount server



NF8260M6 is a 2U quad-socket high-density rackmount server that takes into account computing performance, physical resources and TCO optimization. It supports various scenarios such as virtualization and memory database. With the perfect combination of density, efficiency and reliability in a limited space, it is suitable for large and medium-sized enterprises and the Internet industry.

Features

Extreme computing density

- Based on the latest generation of the Intel Cedar Island platform, the 2U space can accommodate a maximum of 4 processors to provide sufficient computing density with tight computing resources.

Green and energy-saving

- Multiple fan wall systems and radiators are available to match different business loads, avoid unnecessary energy consumption, and minimize TCO.

Cost-effective scalability

- 6 x onboard PCIe3.0 at the rear cover 90% of the demand and support on-demand I/O expansion.

Specifications

Model	NF8260M6
Processor	2/4×Intel®Xeon® Cooper Lake scalable processors
Memory	48×memory slots, up to DDR4-3200 and 24×Intel®Optane™ PMem
Storage	Front up to 24 or 25×2.5-inch drives (optional) 24×U.2 NVMe SSDs Built-in up to 2×M.2 SSD, 2×Micro SD cards
I/O Expansion	Up to 12×PCIe 3.0 slots 1×OCP3.0×16 NIC
PSU	1+1 redundancy power supplies (up to Titanium), 800W/1300W/1600W/2000W standard CRPS Platinum
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

D Mission Critical Server

NF8480M6

Expansion-optimized quad-socket rackmount server



NF8480M6, a high-end flagship model of Inspur, is a 4U quad-socket rackmount server. It is suitable for applications such as large transaction databases, memory databases, and ERP. With improved scalability and performance to target the insufficient optimization capacity of small and medium-sized customers, it serves better I/O-intensive application scenarios and safeguards the critical business of customers.

Features

Intensive application scenarios

- I/O-balance design even distribute I/O resources, suitable for high demanding of network and I/O-intensive application scenarios.
- Support Bfloat16 to accelerates data analysis and data mining.

Dynamic expansion

- 19 x PCIe slots, up to 50 x 2.5-inch drives, and 24 x NVMe SSDs, flexibly adaptable for various application scenarios.

Secure and reliable

- Multiple redundancy design. ■ BIOS/BMC chip-level redundancy design for fault-tolerant solutions from component to solution level.
- System reliability improved significantly through the optimization of fault-prone components.

Specifications

Model	NF8480M6
Processor	2/4× Intel® Xeon® Cooper Lake scalable processors
Memory	48× memory slots, up to DDR4-3200 and 24× Intel® Optane™ PMem
Storage	Front up to 49 or 50× 2.5-inch drives(optional) 24× U.2 NVMe SSDs Built-in up to 2× M.2 SSD, 2× Micro SD cards
I/O Expansion	Up to 18× PCIe 3.0 slots 1× OCP3.0× 16 NIC
PSU	N+N redundancy power supplies (up to Titanium), 550W/800W/1300W/1600W/2000W standard CRPS Platinum
Working Temperature	5°C-45°C. Refer to the <i>Technical White Paper</i>

E AI Server

NF5468M6

Ultra-flexible AI server



NF5468M6 is the latest addition in the best-selling Inspur NF5468-series of versatile AI servers. It supports the most AI accelerators and home-grown AI chips, making it ideal for a wide range of AI applications including AI cloud, IVA, video processing and so much more. NF5468M6 offers ultra-high storage capacity and the unique function of switching topologies between Balance, Common and Cascade in one click, which helps to flexibly adapt to various needs for AI application performance optimization.

Features

Leading architecture

- Support up to 20x PCIe GPUs/accelerators in 4U. ■ Integrate the latest NVIDIA A40 and MIG-empowered A100 GPU. ■ Support AMD Instinct MI100 powered by the latest Matrix Core technology.
- Support 200G HDR InfiniBand Switch and 100Gb OCP 3.0 Card.

Supreme performance

- Support 2x 3rd Gen Intel® Xeon® Scalable processors, TDP 270W, 3UPI. ■ 4GPU non-blocking communication, CPU-to-GPU communication without touching PCIe Switch. ■ 8x hot swap NVMe SSD, providing outstanding I/O performance. ■ Up to 8x 350W GPUs in the PCIe 4.0-empowered system with 2x bandwidth.

Flexible configuration

- Perfectly compatible with the most AI accelerators and home-grown AI chips. ■ Unique function of switching topologies in one click for various AI applications. ■ Three different GPU configurations designed for a variety of AI applications.

Outstanding ecosystem

- Extensive array of AI accelerator partners across the world including NVIDIA, AMD, Intel, Cambricon and more. ■ Leading deep learning frameworks covering TensorFlow, PyTorch, PaddlePaddle, MXNet and so much more. ■ Powerful AI development toolkits, for example, Inspur AIStation, algorithm and application optimization services.
- Wide variety of AI solutions across industries contributed by Inspur's AI ecosystem partners.

Specifications

Model	NF5468M6-P	NF5468M6-T	NF5468M6-V
Processor	2× 3rd Gen Intel® Xeon® Scalable processor		
GPU	8× A100, A40, MI100 FHFL Double Width PCIe GPU	4× A100, A40, MI100 FHFL Double Width PCIe GPU	16× A10 FHFL Single Width PCIe GPU
Memory	32× memory slots, up to DDR4-3200, RDIMM/LRDIMM		
Storage	24× 2.5" or 12× 3.5" SAS/SATA (up to 8× NVMe SSD), 2× M.2 SATA SSD	16× 2.5" or 12× 3.5" SAS/SATA (up to 2× NVMe SSD), 2× M.2 SATA SSD	24× 2.5" or 12× 3.5" SAS/SATA (up to 8× NVMe SSD), 2× M.2 SATA SSD
I/O Expansion	4× PCIe 4.0× 16 slot 2× RAID Mezz card	2× PCIe 3.0× 8 slot, 2× PCIe 4.0× 8slot 2× RAID Mezz card	4× PCIe 4.0 x16 slot 2× RAID Mezz card
PSU	2+2 redundant, 4× 1600W / 2000W / 2200W / 3000W 80Plus Platinum hot-swap PSU		
Working Temperature	5°C-35°C. Refer to the <i>Technical White Paper</i> .		

F Rack-scale Systems

ORS6000S

ODCC-standard
open computing
rack-scale server



ORS6000S is an ODCC3.0-based appliance server that is developed to continuously improve the resource utilization of data centers and reduce the overall cost.

Features

Open architecture and high-density deployment

- Based on the ODCC3.0, ORS6000S supports a flexible combination of different nodes such as computing and balancing nodes. A single cabinet can accommodate a maximum of 32 x dual-socket nodes.
- Compared with ordinary rackmount servers, its computing density is increased by 100%. Also, it is compatible with standard rackmount switches and integrates network, computing, storage and other functions in a single machine, which further allows flexible configuration for various needs.

Centralized power supply and energy-saving

- Leverage Busbar and power supply modules to support centralized power supply at the machine level.
- Instead of the distributed power supply of the conventional rackmount server, ORS6000S enables centralized power supply with N+N redundancy, which satisfies the power supply demands of multiple nodes while balancing the rack-level load output. The design achieves the maximum efficiency at 50% of loading and has an operating peak efficiency of approximately 94% at full load.

Integrated delivery, efficient deployment, operations and maintenance

- Servers and switches are integrated into the rack at the factory.
- Support unified purchase. Convenient deployment and rack-level delivery greatly reduce the time for installation and deployment and increases the delivery speed by 10 times.
- Front node maintenance, toolless hot-pluggable module, and optimized wiring design reduce the difficulty of system operations and maintenance.


Specifications

Model	ORS6000S
Rack	Standard 42U Rack (1200mm×600mm×2100mm)
Balancing Node	SN5264FM6 <ul style="list-style-type: none"> 2× Intel® Xeon® Ice Lake scalable processors 32× DDR4 RDIMM/(16× RDIMM + 16× Intel® Optane™ PMem) Up to 12× 3.5-inch HDDs + 2× NVMe SSDs, 2× onboard SATA M.2 SSDs 1× OCP 3.0
Computing Node	SN5160FM6 <ul style="list-style-type: none"> 2× Intel® Xeon® Ice Lake scalable processors 32× DDR4 RDIMM/(16× RDIMM + 16× Intel® Optane™ PMem) 8× 2.5-inch NVMe SSDs, 2× onboard SATA M.2 SSDs 1× OCP 3.0
Power Supply Module	In the middle of the enclosure, N+N redundancy, centralized power supply
Network Module	3× 1U rackmount switches (1Gb or 10Gb)
Management Module	Integrated RMC module in the middle of the enclosure for power supply and node status monitoring

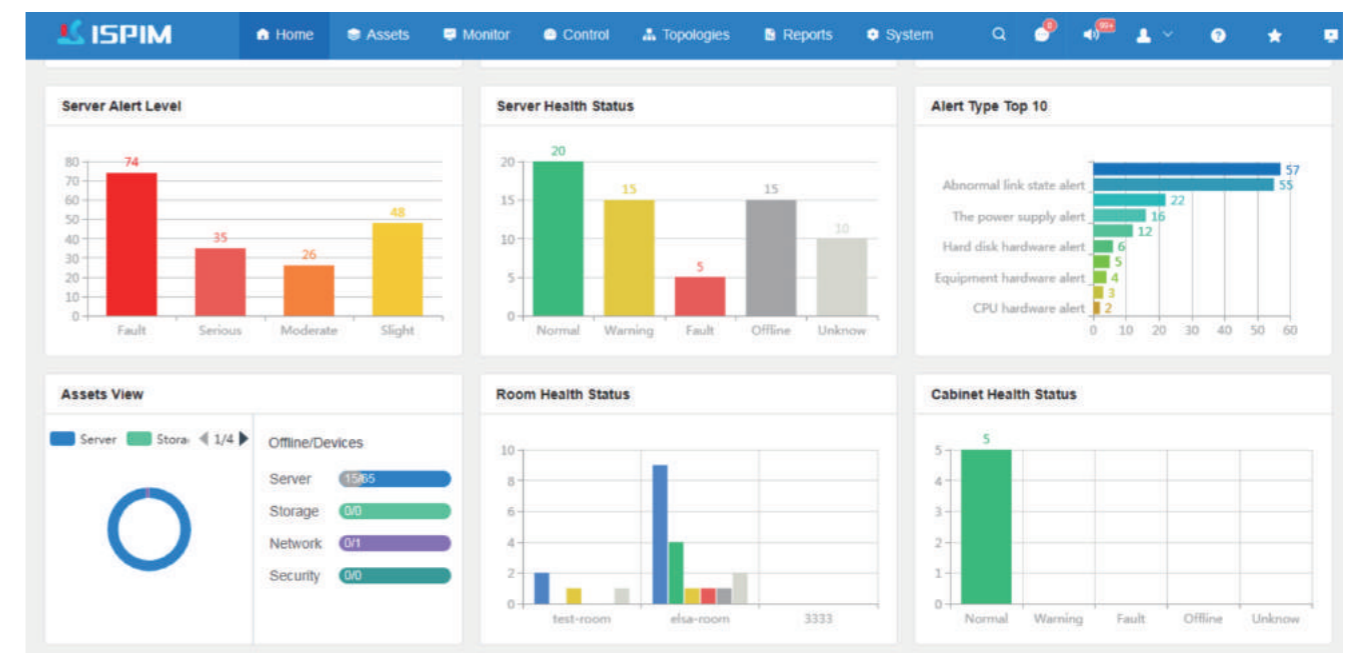
G Infrastructure Management Software Suite

ISPIM

Inspur Server Physical
Infrastructure Manager



ISPIM (Inspur Server Physical Infrastructure Manager) is a new generation of infrastructure operation and maintenance management platform for industry data centers, providing users with leading and efficient management solutions at the data center level. ISPIM enables unified resource management, in-depth fault diagnosis, second-level performance monitoring, intelligent energy management, automated 3D topology, and stateless auto-deployment, and realizes unified operation and maintenance of the server, storage, network, security device, and edge device. Consequently, this effectively helps enterprises improve operation and maintenance efficiency, reduce related costs, and ensure the safe, reliable, and stable operation of data centers.



Features

Unified asset management

- Automated information organization and statistical analysis of the server, storage, switch and other devices realize life-cycle management of assets.

Visual topology

- Leverage automated topology for device network links and 3D topology for the enclosure and rack to dynamically display the temperature, power consumption, and status of the facility.

Intelligent energy management

- Support AI-enabled aggregation and analysis of device indicators (such as temperature, airflow, CUPS) and provide intelligent power consumption strategies and suggestions for optimization.

Extensive monitoring

- Adopt both active and passive network monitoring and provide 24/7 real-time monitoring and alarming.
- Support to locate root causes of failures and in-band performance indicating in seconds.

Efficient management in the stateless environment


- Enable alarming and auto-repair for firmware deviations from baseline and support full firmware upgrade of the device.

Third-party integration

- Provide standard RESTful, SNMP and other northbound interfaces to connect to third-party management platforms.

ISREST

Inspur Server RESTful Tool



ISREST (Inspur Server RESTful Tool) is used for daily maintenance of the server and converts the odata into the command line format that is easy to read and use. Manage the server through ISREST query and setting commands, in line with the new network governance model.

Features

Simplified management protocols

- Integrate Redfish and IPMI protocols to provide a unified and comprehensive mechanism for server hardware query and setting.

Lightweight and convenient operation and maintenance

- Link control enables timely resource release after the BMC operation is completed, ensuring minimal resource occupation.

Cross-platform and multi-model compatibility

- Support standard APIs and compatible with Inspur M6/M5 families, which is convenient for third-party software integration.

ISA

Inspur Server Automation



ISA (Inspur Server Automation) is used for the daily maintenance of servers and can effectively reduce the burden of human labour. By connecting to the out-of-band management network, ISA enables BIOS/ BMC batch configuration and upgrade, batch server inspection, batch log collection, batch power operation, batch RAID configuration, etc.

Features

Flexible deployment, lightweight operation

- ISA is a lightweight operation and maintenance tool that can be used on laptops without installation. ISA is compatible with mainstream Windows and Linux OS and a variety of commonly used operating environments, ensuring limited resources occupancy.

Minimal operation, rich templates

- ISA GUI enables intuitive display and operations, such as automated device query, task parameter setting, and batch task execution.
- Perform by following the prompt messages. ISA allows to export and import templates, streamlining configuration and saving effort.

Absolute security, reliable out-of-band connection

- ISA clients perform the batch operation and maintenance by connecting to the out-of-band management network. Through the connection with the ISBMC management network port, ISA is separated from the network without collecting customer's business information, greatly improving the operation and maintenance security.

Various interfaces and convenient expansion

- With the help of the manual and ISA CLI, operation and maintenance personnel can easily integrate ISA into their proprietary management software or connect to third-party software, avoiding any waste of software investment.

ISQP

Inspur Server Quick Provisioning



ISQP (Inspur Server Quick Provisioning) is used for stand-alone server maintenance. It supports intelligent OS installation, RAID configuration, hardware fault diagnosis, BIOS/BMC firmware upgrade, FRU information refresh, HDD secure erase, etc. Thus, ensure simple operations of device deployment, reuse, scrap, etc.

Features

Automated OS installation

- Support stand-alone installation for Windows and Linux OS, and automated installation of the chipset driver, network adapter driver and other drivers.

Firmware upgrade

- BIOS, BMC and other firmware upgrades; FRU information refreshment.

HDD secure erase

- Completely erase all information on HDDs during device reuse and scrap.

RAID configuration

- Enable disk configuration and management, array card management, VD management, etc.

Hardware diagnosis

- Display hardware information and generate the diagnostic report, explaining BIOS status, BMC status, CPU, HDD, etc.

ISIB

Inspur Server Intelligent Boot



ISIB (Inspur Server Intelligent Boot) allows PXE-based OS batch deployment and supports Windows Server, RedHat, CentOS, SUSE, Ubuntu, etc. It serves the OS batch installation after device installation. Also, ISIB has constructed a unified management platform for bare metal servers, covering resource management, firmware upgrade, BIOS/BMC/RAID/hardware configuration, out-of-band management, system deployment and migration, logging, etc. This effectively helps users improve operation and maintenance efficiency, reduce related costs, and ensure reliable and stable operation of data centers.

Features

One-stop services

- From the perspective task management, it allows one-stop deployment from firmware upgrade, hardware configuration, to OS installation or migration.

Tiered architecture

- Modular design and API call facilitate the connection with the enterprise's in-house platform.

Custom settings

- Meet various needs for customization.

Automated OS installation

- Two installation modes are available, that is, batch installation and mirroring/cloning.

Accurate log system

- Support to trace the execution results of a single command.