NF8480M5 Technical White Paper

Version No: **3.2**

Date of issue: June 30, 2020

Inspur and "浪潮" are registered trademarks of Inspur Group.

Windows is a registered trademark of Microsoft.

Intel and Xeon are registered trademarks of Intel.

Other trademarks belong to their respective registered companies.

Table of Contents

1 Product Overview

Inspur NF8480M5 server is a high-end four-way rack server designed by Inspur for the business application needs of governments, communications, finance, energy and various large enterprises based on the new Intel® Xeon® scalable computing platform. Boasting powerful computing capability, scalability, excellent RAS features and energy efficiency ratio, this product is an ideal choice for in-memory databases, ERP, CRM, business intelligence analysis systems, large-scale virtualized applications and data-intensive applications. It is especially suitable for key business applications with high requirements for system performance, scalability and stability.

Figure 1-1 Picture of NF8480M5

2 Features

For different application scenarios, NF8480M5 maintains the consistent high-quality and reliable performance of Inspur servers, applying the extreme design concept in performance, scalability, usability, manageability, etc.

Performance:

- NF8480M5 supports Intel® Xeon® scalable processors (Skylake, Cascade Lake). A single CPU has up to 28 cores and 56 threads, supports a maximum TDP 205W CPU. It also supports the intra-processor full interconnection technology, which greatly reduces intra-processor communication latency. A full direct connection architecture is adopted between processors to realize direct communication.

- It is scalable with 48 2933 MT/s DDR4 RAMs of RDIMM and LRDIMM types, which can provide excellent speed, high availability and up to 6TB memory capacity.

- It supports 24 Intel® OPTANE™ PMEM, with a single PMem capacity up to 512G, which can satisfy key applications with high memory requirements.

- It can support 12 hot-swappable NVME SSDs with a maximum read performance of 3.5GB/s.

Scalability:

- It supports up to 24 3.5-inch or 2.5-inch local hard disks, with two built-in M.2 SSDs that can be used as startup disks.

- It supports up to 16 PCIE 3.0 standard expansion slots, and the Riser Module can be flexibly configured according to the different needs of customers for IO so as to be applied in different IO scenarios.

- Supporting Intel integrated I/O technology, it can integrate PCI Express 3.0 controllers into Intel® Xeon® scalable processors, thus significantly reducing I/O latency and improving overall system performance.

- It supports the OCP daughter card design, the free switch between OCP and PHY network adapters, and provides multiple network interface options including 1G, 10G, 25G, thus offering a more flexible network structure for applications.

Usability:

NF8480M5 provides a dozen of functions to enhance usability and improve system running time:

- Key vulnerable parts are redundant; the power supply supports multiple redundancy modes including 1+1/2+1/2+2/3+1 modes; and the system fan supports N+1 redundancy.

- Equipped with hot-swappable SAS/SATA hard disks that support RAID 0/1/10/5/50/6/60 levels and power-down data protection for supercapacitors, and provide RAID Cache.

- NVME hard disks can support RAID functions through Intel VROC technology, and can support RAID 0/1/5/10 levels by matching different raid keys.

- Equipped with the SCSA super current safety protection technology, it boasts ultra-low impedance power supply, low line loss, high reliability, reduced footprint, improved maintenance efficiency and heat dissipation efficiency, and improved system reliability and availability.

- Based on a humanized design concept, the entire system can be maintained without tools. The front 3.5-inch hard disk backplane adds a structural reinforcement frame, which is easy for disassembly and greatly reduces the time and cost of operation and maintenance.

- The reliability after applying SSDs is much higher than that of traditional mechanical hard disks. Therefore it can prolong system running time.

● After applying the next-generation Inspur server management system, technicians can offer guidance to equipment through the Web management interface and fault diagnosis LEDs, and mark faulty devices through the UID indicators on the front panel, and quickly pin down the components where faults have occurred or are occurring. In this way, the maintenance will be simplified, problems will be solved faster and the system usability will be improved.

Manageability

● An intelligent server management chip is embedded to support IPMI2.0 and Redfish management modes, and realize complete remote system monitoring, remote KVM, virtual media and other management functions.

● The BMC ROM can support dual-mirrored redundancy, and it can be started from the backup boot partition if the upgrade of the primary boot partition fails.

● Inspur's power management technology can help users accurately monitor and control system power consumption in real time. With the help of the exclusive Power Telemetry technology, the Power Thermal Aware Solution and the Node manager 4.0 technology, it enables full energy consumption control and further improves the energy efficiency performance of the entire IT system.

● It can obtain the low-level debugging communication such as CPU registers with local USB interfaces and BMC management chips, and can remotely diagnose server faults through the management network.

● The built-in oscilloscope technology detects key signal status of the server hardware in real time, saves state transitions, and assists engineers in problem solving.

Energy efficiency

● 80 PLUS platinum power modules of different power levels are provided, and the efficiency of the power modules under the 50% load is as high as 94%.

● It supports active PFC (Power Factor Correction) to improve the utilization rate of the mains by the power supply.

● The system cooling fan can regulate the speed by partitions or in PID (Proportional-Integral-Derivative) mode, while supporting CPU-based intelligent frequency regulation for energy saving and consumption reduction.

● The fully-optimized system cooling design and the efficient and energy-saving cooling fan can reduce the energy consumption in system heat dissipation.

Security

● It realizes firmware encryption to prevent illegal writing of unknown firmware, and support BIOS digital signature.

● The built-in TPM 2.0 hardware encryption chip can flexibly select algorithms according to user requirements.

● In terms of hardware design, the upper cover of the chassis can be locked and the background alarm can be set.

3 Logic Architecture

NF8480M5 supports 2 or 4 Intel® Xeon® scalable processors, with each CPU supporting 6 memory channels and each channel supporting two memory slots, so a maximum of 48 DDR4 DIMMs are supported. Processors are fully interconnected through the UPI bus, with the transmission rate as high as 10.4GT/s.

It provides up to 16 PCIE3.0 standard slots, and low-speed IO interfaces such as VGA, USB3.0, and

serial ports to meet the needs of users in various application scenarios.

The processor is connected to three PCI-E Riser cards through the PCI-E bus, and supports different PCI-E slots through different PCI-E Risers. The OCP A/B connector on the motherboard is connected to CPU1 through the PCI-E bus, while the OCP C connector is connected to the PCH. Figure 3-1 shows the logic block diagram of NF8480M5.

Figure 3-1 NF8480M5 motherboard logic block diagram-
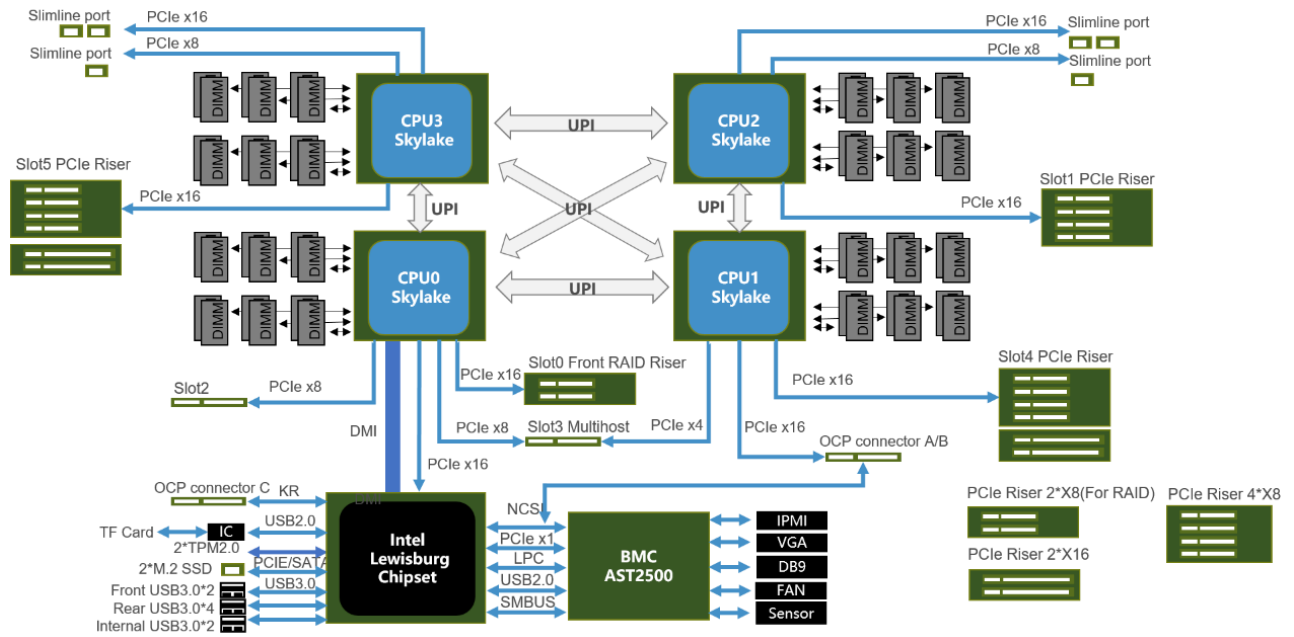
4 Product Profile

4.1 Front panel

4.1.1 Front view of 24 × 3.5"/2.5" front panel

Figure 4-1 Front view
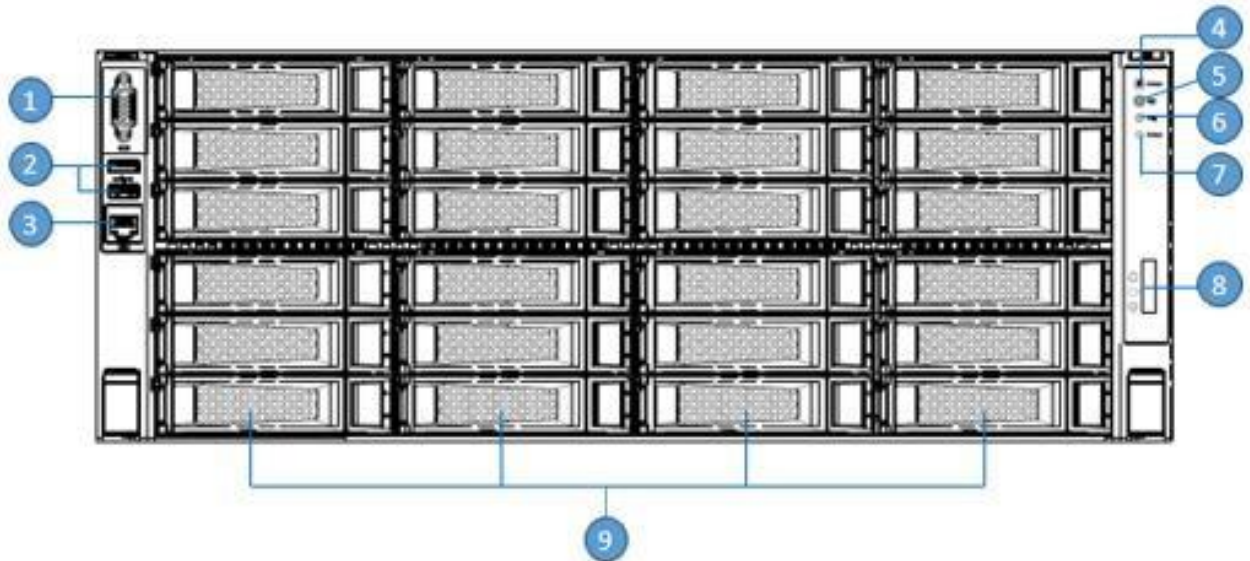
No.	Module Name	No.	Module Name
1	VGA interface	6	Reset button
2	USB 3.0 interface	7	System health indicator
3	System serial port	8	OLED LCD diagnostic screen
4	Power button	9	Front hard disk
5	UID lamp		

4.1.2 OLED LCD diagnostic screen

Figure 4-2 OLED LCD diagnostic screen
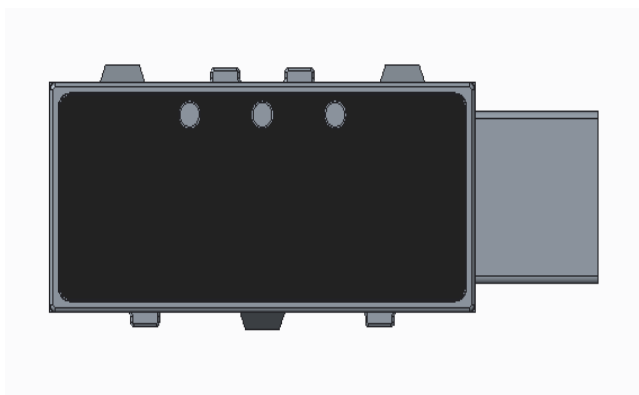
Features:

The OLED LCD diagnostic screen includes the display function and the setting function. The display function is mainly used to display ambient temperature, power consumption of the whole machine, device name, device PN/SN, IP and MAC of dedicated ports and shared ports, and system error information. The setting function is mainly used to modify the static IP of BMC.

4.1.3 Indicators on hard disk carrier

Figure 4-3 Indicators on hard disk carrier
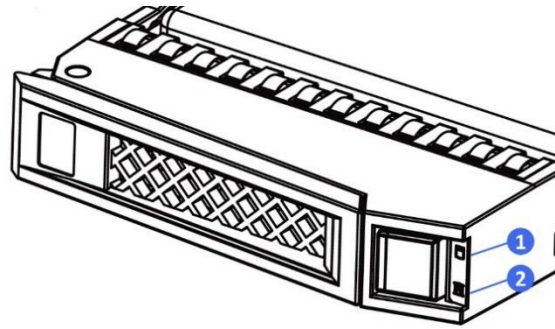
No.	Module Name	Description
1	Hard disk fault indicator	Steady red: The hard disk is faulty. Steady blue: The hard disk is being located. Flashing blue: RAID rebuilding
2	Hard disk status indicator	Steady green: The hard disk is normal. Blinking green: Data is being read from or written to the hard disk.

4.2 Interior top view

Figure 4-4 Internal schematic diagram of server
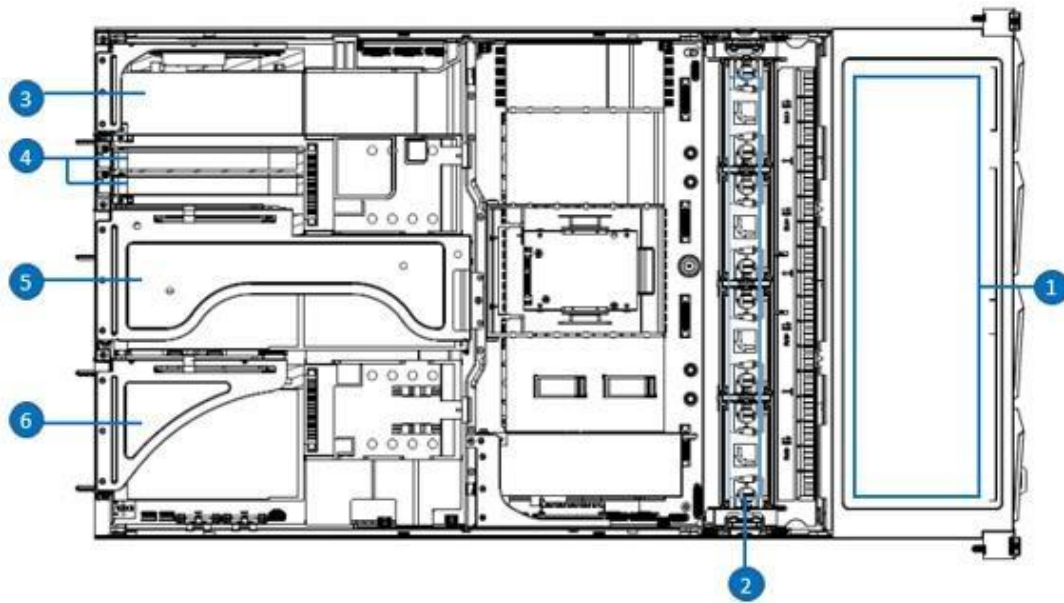
No.	Module name	No.	Module name
1	Front 3.5" * 24 hard disk bay	4	2 PCI-E X8 slots, supporting hot swap
2	8 * 8038 fans	5	4 PCI-E X8 or 2 PCI-E X16 riser cards
3	4 PCI-E X8 riser cards	6	4 PCI-E X8 or 2 PCI-E X16 riser cards

4.3 Rear panel

Figure 4-5 View of normal rear panel
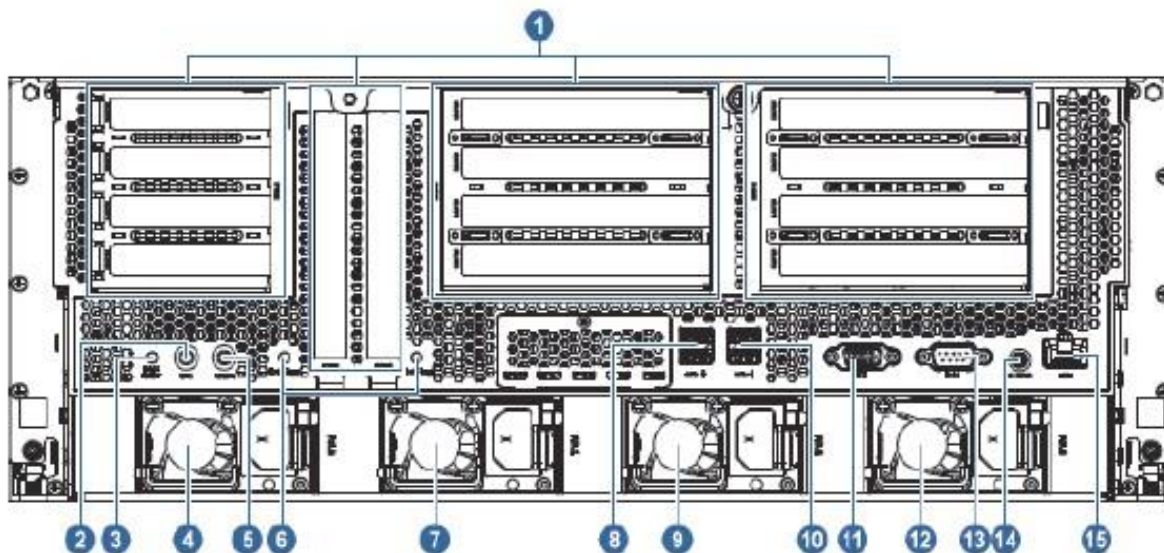
No.	Module name	No.	Module name
1	PCI-E expansion	9	PSU2
2	UID indicator	10	USB3.0 interface (2|3)
3	BMC reset button	11	VGA interface
4	PSU0	12	PSU3
5	Power button	13	Serial port
6	Start button for hot-swappable components	14	BMC serial port
7	PSU1	15	IPMI management port
8	USB3.0 interface (0|1)		

4.4 Motherboard layout

Figure 4-6 - Motherboard layout
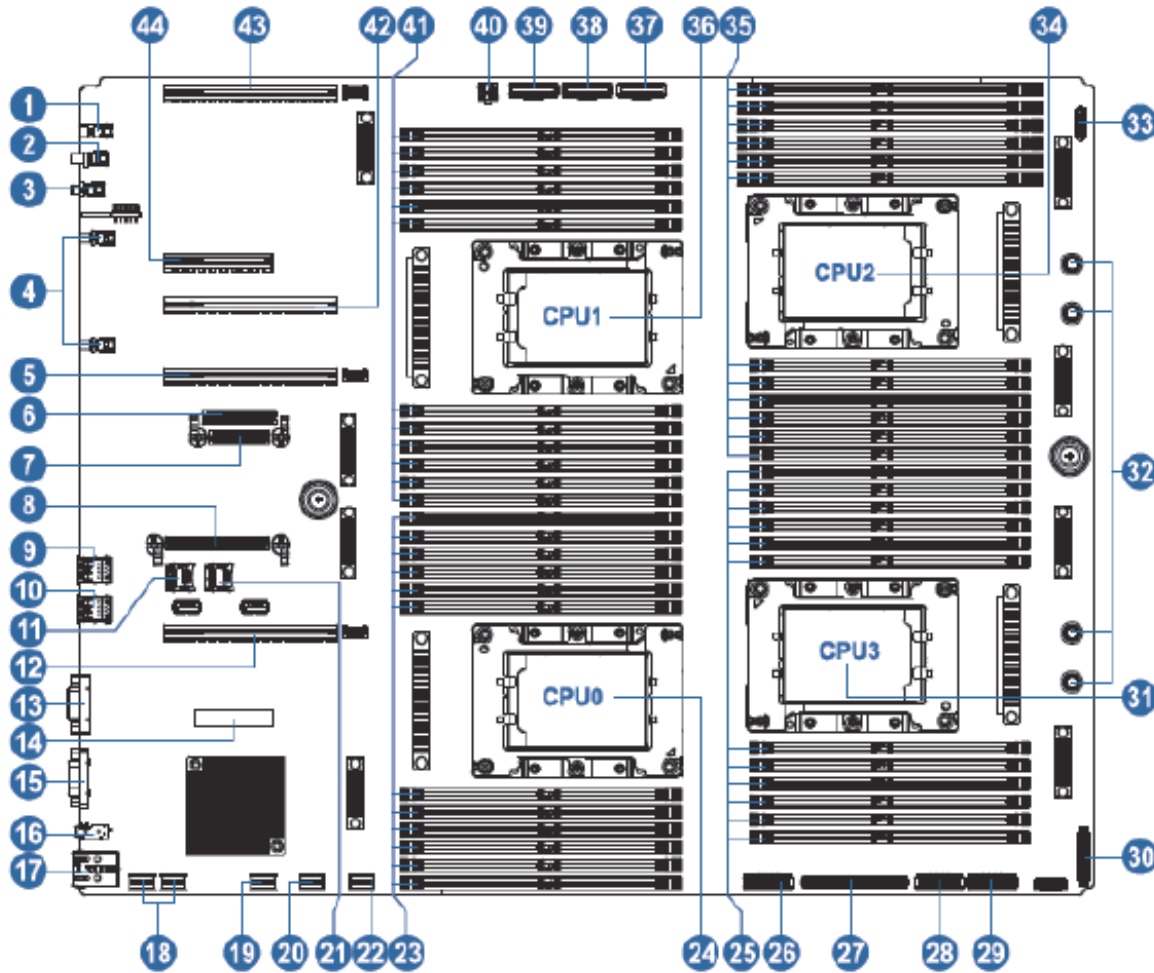
No.	Module name	No.	Module name
1	BMC reset button	23	Memory slot (CPU0)
2	UID indicator	24	CPU0
3	Power button	25	Memory slot (CPU3)
4	Hot-swappable components	26	SLIM LINE2 interface (CPU3)
5	PCIE4 slot (CPU1)	27	PCIE0 slot (CPU0)
6	OCPB card slot (CPU1)	28	SLIM LINE0 interface (CPU3)
7	OCPC card slot	29	SLIM LINE 1 interface (CPU3)
8	OCPA card slot (CPU1)	30	Power board control signal line interface
9	USB3.0 interface (0|1)	31	CPU3
10	USB3.0 interface (2|3)	32	Motherboard power supply port
11	BMC TF card slot	33	Front control board interface
12	PCIE5 slot (CPU3)	34	CPU2
13	VGA interface	35	Memory slot (CPU2)
14	TPM/TCM card slot	36	CPU1
15	Serial port	37	SLIM LINE0 interface (CPU2)
16	BMC serial port	38	SLIM LINE 1 interface (CPU2)
17	IPMI management network port	39	SLIM LINE 2 interface (CPU2)
18	SATA interface	40	Standby power supply interface
19	M.2 hard disk Interface 0	41	Memory slot (CPU1)

20	M.2 hard disk Interface 1	42	PCIE3 slot (CPU0)
21	System TF card slot	43	PCIE1 slot (CPU2)
22	Front control board interface	44	PCIE2 slot (CPU0)

4.5 Riser card

Figure 4-7 Schematic diagram of riser card
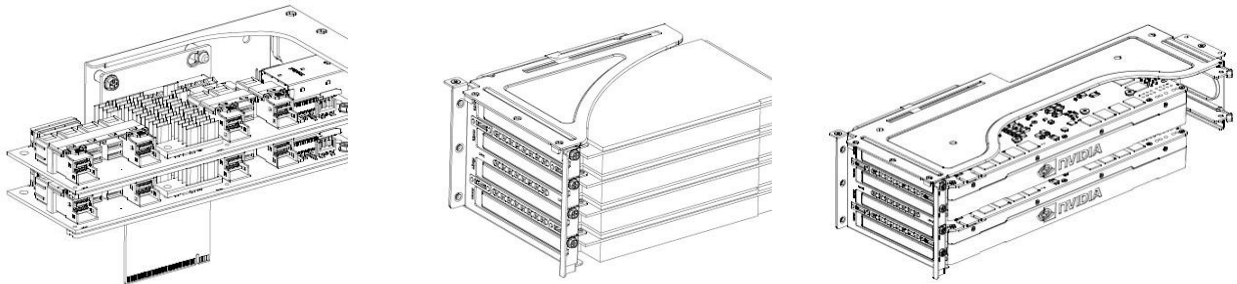
RAID card riser   Half-height half-length riser card   Full-height full-length riser card

Model	Position 1	Position 2	Position 3
RAID card Riser (X8*2)	PCIE 0	N/A	N/A
PCIE Riser 1 (full height X8*4)	N/A	PCIE 4	PCIE 5
PCIE Riser 1 (half height X8*4)	PCIE 1	N/A	N/A
PCIE Riser 2 (X16*2)	N/A	PCIE 4	PCIE 5

Note: Please refer to Figure 4-6: Motherboard Layout for the location of PCIE0. Please refer to the Figure 4-8 below for the location of PCIE1, PCIE4 and PCIE5.

Figure 4-8 PCIE slot distribution
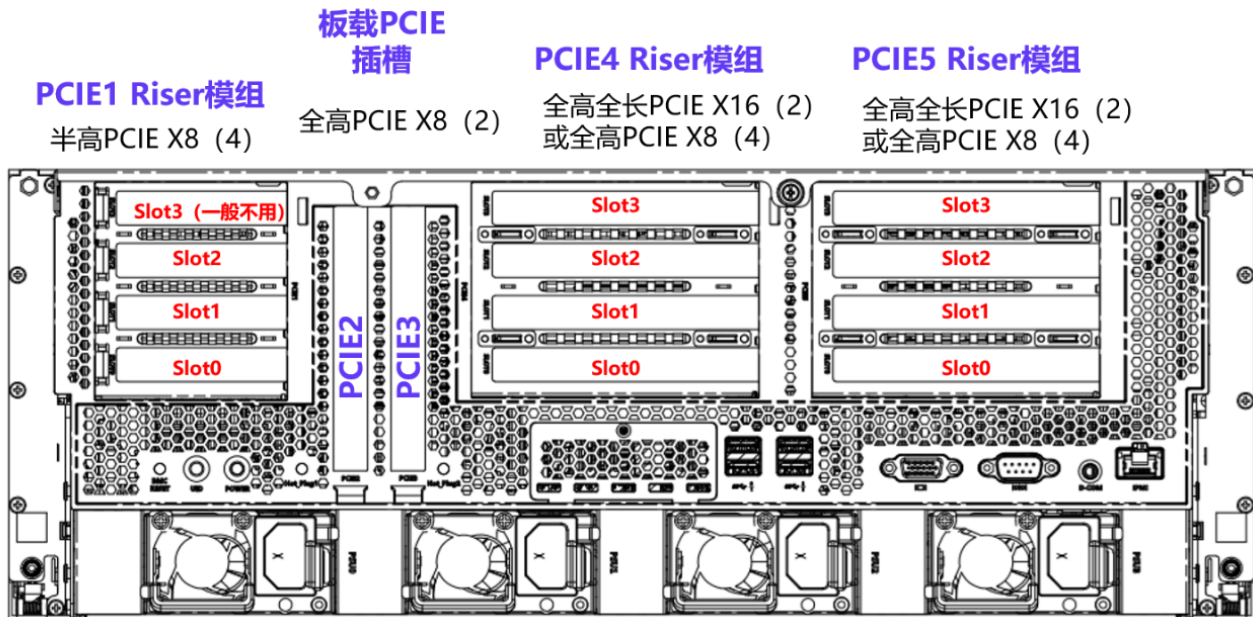
PCIE slots and corresponding CPU				
PCIE	Slot location	Signal source	Slot specifications	Remarks
PCIE0 Raid Card Rise Template	-	CPU0	2 个 PCIeX8	Dedicated to RAID card or Expander card
	-			
Onboard PCIE Slot	PCIE2	CPU0	2 个 PCIeX8	Directly connected to motherboard
	PCIE3	CPU0		
PCIe4 Riser Module	Slot0	CPU1	4 PCIE x8 full-height half-length cards or 2 PCIE x16 full-height full-length cards	The signals of slot2 and slot3 need to be led to the Riser through the slimline cables of the two x8 slimline interfaces on the motherboard
	Slot1	CPU1		
	Slot2	CPU2		
	Slot3	CPU2		
PCIe1 Riser Module	Slot0	CPU2	4 PCIE x8 half-height half-length cards	The signals of slot2 and slot3 need to be led to the Riser through the slimline cables of the two x8 slimline interfaces on the motherboard
	Slot1	CPU2		
	Slot2	CPU2		
	Slot3	CPU3		
PCIe5 Riser Module	Slot0	CPU3	4 PCIE x8 full-height half-length cards or 2 PCIE x16 full-height full-length cards	The signals of slot2 and slot3 need to be led to the Riser through the slimline cables of the two x8 slimline interfaces on the motherboard
	Slot1	CPU3		
	Slot2	CPU3		
	Slot3	CPU3		

4.6 Hard disk backplane

There are two hard disk backplanes for NF8480M5, which are ordinary hard disk backplane (3.5" × 12 SAS) and NVME hard disk backplane (3.5" × 12 NVME), as shown in the figure below.
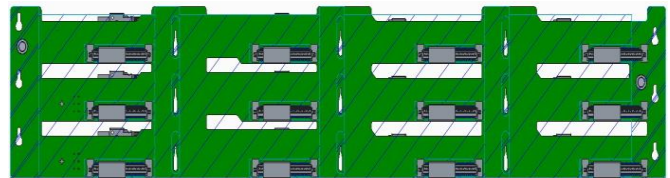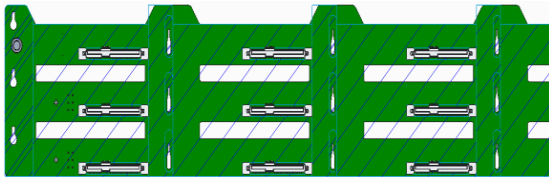
Schematic diagram of 3.5" × 12 SAS backplane   Schematic diagram of 3.5" × 12 NVME backplane

Table 4-1 Relationship between RAID/SAS/Expander card and hard disk configuration

Number of hard disks	Hard disk backplane	RAID card	Expander card
1≤HDD≤8	1 SAS	1	0
9≤HDD≤12	1 SAS	1	1
13≤HDD≤24	2 SAS	1	1
1≤NVME≤12	1NVME	0	0
1≤NVME 1≤NVME+ HDD≤12 and 1≤HDD≤8	1NVME	1	0
1≤NVME, and 9≤HDD	1NVME+1SAS	1	1

Installation principle of the hard disk backplane: If one hard disk backplane is configured, install it on the upper part of the front window of the chassis

Note: (1) A maximum of one NVME hard disk backplane is supported; if the configuration is one NVME backplane + one normal backplane, the NVME backplane is installed on the top

(2) As Expander cards are about to be phased out in subsequent stages, if there are more than 8 HDDs, 16i or 24i raid cards are preferred.

5 System Specifications

Table 5-1 Product parameters

Component	Description
Specifications	4U rack type
Processor	2 or 4 next-generation Intel® Xeon® scalable processors: Supports up to 28 cores (2.7GHz) A maximum frequency of 3.8GHz (4 cores) Interconnection of 3 UPI buses, with the transmission rate of a single link up to 10.4GT/s Supports a TDP up to 205 W.
Chipset	Intel C621/C624/C627
Memory	Support up to 48 DDR4 2400/2666/2933 MT/s RAMs Two CPUs support 24 DIMMs, and four CPUs support 48 DIMMs Support RDIMM / LRDIMM / OPTANE™ PMem
Storage	Support up to 24 3.5/2.5 inch SAS/SATA HDD/SSD hard disks Support up to 12 U.2 NVME hard disks Support up to two built-in M.2 SSD hard disks, and support soft RAID 0/1 Support up to two SD cards, BMC_TF and SYS_TF cards
Storage controller	Equipped with configurable high-performance RAID 0/1/5/6/10/1E/50/60 and supports power failure protection module
Network interface	Support 1 PHY card or OCP daughter card, with the flexible expansion of multiple network configurations. PHY card supports 1G/10G OCP card supports 1G/10G/25G, supports NCSI function, and realizes the access to BMC management system through Sharelink technology Standard PCI-E Ethernet card: support 1G/10G/25G/40G/100G
Display controller	Integrated in the AST2500 chip, with the maximum resolution of 1920X1200 and the cache capacity of 64MB
I/O expansion slot	Support up to 16 standard PCIE3.0 slots and 1 OCP/PHY card slot
Interface	USB interface: 2 front USB 3.0 interfaces, 4 rear USB3.0 interfaces Serial interface: 1 front RJ45 serial port, 1 rear DB9 system serial port VGA interface: 1 front VGA interface, 1 rear VGA interface
Fan	8 hot-swappable N+1 redundant 8038 fans (dual rotors)
Power supply	Support 4 platinum 550W/800W/1300W/1600W power modules Support 100-240V AC, 336V high voltage DC or -48V low voltage DC power supply Support multiple redundancy modes including 1+1/2+1/2+2/3+1. Please refer to "6.9 Power Supply" for details
System management	1 rear RJ45 independent IPMI management network port Support remote management control, such as IPMI2.0, KVM over IP (HTML5 /Java), SOL, SNMP, Redfish A standard OLED LCD diagnostic screen for system fault diagnosis
Operating system	Redhat 6/ Redhat 7/ centos/ OracleLinux/ Windows Server/ ESXi/ Suse/ InCloud Sphere/ Ubuntu/ Galaxy Kylin/ Debian Please refer to "6.10 Operating System" for details

Table 5-2 Physical specifications

Item	Specifications
Dimensions	With hangers: W (width) 447.6mm; H (height) 175.5mm; D (depth) 812.9 mm Without hangers: W (width) 447.6mm; H (height) 175.5mm; D (depth) 785.4mm
Weight	3.5*24 configuration Host (without packaging): 52.6kg Gross weight (including packaging): 67.4kg (including packaging + guide rail + accessory box)
Temperature	Working temperature: 5~45℃ [1,2,3]; Storage temperature (with packaging): -40~+70℃ Storage temperature (without packaging): -40~+55℃
Humidity	Working humidity: 10%~90% R.H. Storage humidity (with packaging): 10%~93% R.H. Storage humidity (without packaging): 10%~93%R.H.
Noise (Bels)	Idle

(Sound power) 4,5,6,7	LWAd: 4.98B general low-end configuration; 4.98 B general configuration LpAm: 52.2 dB general low-end configuration; 52.2 dB general configuration Operating LWAd: 6.35B general low-end configuration; 6.62 B general configuration LpAm: 49.8 dB general configuration; 56.4 dB general configuration
Height	Operating temperature from 0°C to 40°C at 0 to 914 meters (3000 feet); Operating temperature from 10°C to 32°C at 914 to 2133 meters (7000 feet)

Table 5-3 Operating temperature specifications

Model	Maximum working temperature 30℃	Maximum working temperature 35℃	Maximum working temperature 40℃	Maximum working temperature 45℃
Full range of models	√	X	X	X
Full range of models (Except V100)	√	√	X	X
Full range of models (12pcs below 3.5 HDD)	√	√	X	X
Non-GPU configuration models	√	√	√	X
Low-end configuration models (140W CPU, no GPU, no OCP, no PCI-SSD, no optical port network adapter)	√	√	√	√

Note:

1. Not all configurations support a working temperature of 5~45℃, and the GPU configuration supports a working temperature of 10~30℃

2. Standard Operating Temperature

The sea level is 10° to 35°C (50° to 95°F); the temperature drops by 1.0°C for every 305 meters above the sea level (1.8°F for every 1000 feet); the maximum height is 3050 meters (10000 feet), and it cannot be exposed to direct sunlight. The maximum rate of change is 20°C/HR (36°F/HR). The altitude and maximum temperature change rate will vary due to different system configurations;

If the fan fails or the ambient temperature is above 30°C (86°F), the system performance may decrease.

3. Extended ambient working temperature

For some of the approved configurations, the supported system inlet range is extended to 5° to 10℃ (41° to 50°F) and 35° to 45°C (95° to 104°F) at sea level. The temperature decreases by 1.0°C for every 175m (1.8°F per 574ft) increased between 900m (2953ft) and 3050m (10,000ft) above sea level.

For some of the approved configurations, the supported system inlet range is extended to 35° to 45°C (104° to 113°F) at sea level. The temperature decreases by 1.0°C for every 125m (1.8°F per 410ft) increased between 900m (2953ft) and 3050m (10,000ft) above sea level.

If it runs within the extended environmental working range or the fan fails, system performance may decrease.

4. This document lists the weighted acoustic power (LWAd) and weighted acoustic pressure (LpAm) when the product runs at 23°C, It complies with the ISO7779 (ECMA 74) noise measurement standard, and is declared according to ISO 9296 (ECMA 109). The listed acoustic levels are applicable to general shipment configurations; other options may cause volume increase. Please contact your sales representative for more details.

5. The listed acoustic levels are measured only by the specific test configurations. The acoustic levels vary according to the system configuration. The value is subject to change without further notice and is for reference only.

6. The test evaluation based on samples (types) conforms to the cited product specifications. This product or portfolio is eligible to carry the appropriate compliance mark and statement.

7. The listed acoustic levels are for standard shipping configurations, while other system configurations may cause the noise to increase.

Table 5-4 Safety & EMC

Safety	IEC 60950-1:2005 (Second Edition); Am1:2009 + Am2:2013 IEC 60950-1:2005 EN 60950-1:2006+A11:2009+A1:2010+A12:2011+A2:2013 GB4943.1-2011 UL 60950-1 and CAN/CSA C22.2 No. 60950-1-07 Standard for Information Technology Equipment-Safety-Part 1: General Requirements TC 004/2011 IS 13252 (PART 1): 2010/ IEC 60950-1: 2005
EMC	GB/T9254-2008 (idt CISPR 22: 2006) GB17625.1-2012 (idt IEC 61000-3-2: 2009) EN 55032:2015 EN 61000-3-2:2014 EN 61000-3-3:2013 EN 55024:2010+A1:2015 EN 55035:2017 AS/NZS CISPR 32:2015 CFR 47 FCC Part 15 subpart B, 2018 ICES-003 ISSUE 6:2016 TC 020/2011 KN32 KN35

Table 5-5 Industry Standard Compliance

ACPI 6.1 Compliant
PCI-E 3.0 Compliant
WOL Support
SMBIOS 3.1
UEFI 2.6
Redfish API
IPMI 2.0
Microsoft® Logo certifications
PXE Support
Advanced Encryption Standard (AES)
SNMP v3
TLS 1.2
ACtive Directory v1.0

6 Compatibility List

※ The Compatibility List was updated in July 2020. For the latest compatibility and product part models not reflected in this manual, please consult the technical personnel of Inspur for details.

6.1 Processor

NF8480M5 supports four Intel Xeon scalable processors

Table 6-1 CPU list

Intel Platinum 8100&8200 series

Model	Cores	Number of threads	Fundamental frequency (GHz)	Maximum turbo frequency (GHz)	Cache (MB L3)	Maximum memory size	UPI	TDP (W)
8280L	28	56	2.70	4.0	38.5	4.5 TB	3	205
8280M	28	56	2.70	4.0	38.5	2 TB	3	205
8280	28	56	2.70	4.0	38.5	1 TB	3	205
8270	26	52	2.70	4.0	35.75	1 TB	3	205
8268	24	48	2.90	3.90	35.75	1 TB	3	205
8276M	28	56	2.20	4.0	38.5	2 TB	3	165
8276	28	56	2.20	4.0	38.5	1 TB	3	165
8260M	24	48	2.40	3.90	35.75	2 TB	3	165
8260	24	48	2.40	3.90	35.75	1 TB	3	165
8256	4	8	3.80	3.90	16.5	1 TB	3	105
8253	16	32	2.20	3.00	22	1 TB	3	125
8180	28	56	2.50	3.80	38.5	768 GB	3	205
8176M	28	56	2.10	3.80	38.5	1.5 TB	3	165
8176	28	56	2.10	3.80	38.5	768 GB	3	165
8170	26	52	2.10	3.70	35.75	768 GB	3	165
8168	24	48	2.70	3.70	33	768 GB	3	205
8164	26	52	2.00	3.70	35.75	768 GB	3	150
8160	24	48	2.10	3.70	33	768 GB	3	150
8158	12	24	3.00	3.70	24.75	768 GB	3	150
8156	4	8	3.60	3.70	16.5	768 GB	3	105

Intel Gold 6100&6200 series

Model	Cores	Number of threads	Fundamental frequency (GHz)	Maximum turbo frequency (GHz)	Cache (MB L3)	Maximum memory size	UPI	TDP (W)
6254	18	36	3.1	4	24.75	1 TB	3	200
6252N	24	48	2.3	3.6	35.75	1 TB	3	150
6252	24	48	2.1	3.7	35.75	1 TB	3	150
6248	20	40	2.5	3.9	27.5	1 TB	3	150
6246	12	24	3.3	4.2	24.75	1 TB	3	165
6244	8	16	3.6	4.4	24.75	1 TB	3	150
6242	16	32	2.8	3.9	22	1 TB	3	150
6240	18	36	2.6	3.9	24.75	1 TB	3	150
6238	22	44	2.1	3.7	30.25	1 TB	3	140
6234	8	16	3.3	4	24.75	1 TB	3	130
6230N	20	40	2.3	3.5	27.5	1 TB	3	125

6230	20	40	2.1	3.9	27.5	1 TB	3	125
6226	12	24	2.7	3.7	19.25	1 TB	3	125
6154	18	36	3	3.7	24.75	768 GB	3	200
6152	22	44	2.1	3.7	30.25	768 GB	3	140
6150	18	36	2.7	3.7	24.75	768 GB	3	165
6148	20	40	2.4	3.7	27.5	768 GB	3	150
6146	12	24	3.2	4.2	24.75	768 GB	3	165
6144	8	16	3.5	4.2	24.75	768 GB	3	150
6142	16	32	2.6	3.7	22	768 GB	3	150
6140M	18	36	2.3	3.7	24.75	1.5 TB	3	140
6140	18	36	2.3	3.7	24.75	768 GB	3	140
6138	20	40	2	3.7	27.5	768 GB	3	125
6136	12	24	3	3.7	24.75	768 GB	3	150
6134	8	16	3.2	3.7	24.75	768 GB	3	130
6132	14	28	2.6	3.7	19.25	768 GB	3	140
6130	16	32	2.1	3.7	22	768 GB	3	125
6128	6	12	3.4	3.7	19.25	768 GB	3	115
6126	12	24	2.6	3.7	19.25	768 GB	3	125

Intel Gold 5100&5200 series

Model	Cores	Number of threads	Fundamental frequency (GHz)	Maximum turbo frequency (GHz)	Cache (MB L3)	Maximum memory size	UPI	TDP (W)
5220	18	36	2.2	3.9	24.75	1 TB	2	125
5218N	16	32	2.3	3.7	22	1 TB	2	110
5218	16	32	2.3	3.9	22	1 TB	2	125
5217	8	16	3	3.7	11	1 TB	2	115
5215	10	20	2.5	3.4	13.75	1 TB	2	85
5122	4	8	3.6	3.7	16.5	768 GB	2	105
5120	14	28	2.2	3.2	19.25	768 GB	2	105
5118	12	24	2.3	3.2	16.5	768 GB	2	105
5117	14	28	2	2.8	19.25	768 GB	2	105
5115	10	20	2.4	3.2	13.75	768 GB	2	85

6.2 Memory

NF8480M5 supports up to 48 DDR4 RAMs. Each processor supports 6 memory channels, and each channel supports 2 memory slots. It supports RDIMM / LRDIMM / Optane^{TM} PMem (persistent memory). The following memory protection technologies are supported:

- ECC (Error Correcting Code:)

- Memory Mirroring

- Memory Rank Sparing

- SDDC (Single Device Data Correction)

- ADDDC (Adaptive Double- Device Data Correction)

- PPR (Power up - Post Package Repair)

Table 6-2 Memory list

Category	Capacity	Frequency	Data width	Organization
RDIMM	16GB	2400	×72	1R×4

RDIMM	16GB	2666	×72	1R×4
RDIMM	16GB	2933	×72	1R×4
RDIMM	32GB	2400	×72	2R×4
RDIMM	32GB	2666	×72	2R×4
RDIMM	32GB	2933	×72	2R×4/2R×8
RDIMM	64GB	2400	×72	2R×4
RDIMM	64GB	2666	×72	2R×4
RDIMM	64GB	2933	×72	2R×4
LRDIMM	64GB	2400	×72	4R×4
LRDIMM	64GB	2666	×72	4R×4
LRDIMM	64GB	2933	×72	4R×4
RDIMM	128GB	2933	×72	4R×4

The memory slot layout and installation principles are shown in the following figure and Table 6-3:

Table 6-3 Insertion method of ordinary RAM 错误!链接无效。

Table 6-4 Insertion method of PMem RAM

		4CPU																																																
		CPU0												CPU1												CPU2												CPU3											
DRAM Number	PMem Number	CH0		CH1		CH2		CH3		CH4		CH5		CH0		CH1		CH2		CH3		CH4		CH5		CH0		CH1		CH2		CH3		CH4		CH5		CH0		CH1		CH2		CH3		CH4		CH5	
		D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1	D0	D1
24	24	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓	V	✓
24	8	V	✓	V		V		V		V	✓	V		V		V	✓	V		V		V	✓	V		V		V	✓	V		V		V	✓	V		V		V	✓	V		V		V	✓	V	
16	8	V		V		✓		V		V		V		V		V		✓		V		V		V		V		V		✓		V		V		V		V		V		✓		V		V		V		✓
24	16	V	✓	V		V		V		✓	V	V	✓	V		V		V	✓	V		V	✓	V	✓	V		V	✓	V		V		V	✓	V		V	✓	V		V		V	✓	V		V	✓
32	8	V	V	V	V	✓		V	V	V	V	✓		V	V	V	V			V	V	V	V	✓		V	V	V	V		✓	V	V	V	V	✓		V	V	V	V	✓		V	V	V	V		✓

Note:

1. RAMs of different types (RDIMM, LRDIMM) and different specifications (capacity, bit width, rank, height, etc.) are not allowed to be mixed in one server;

2. The maximum memory capacity can be achieved when four processors are installed. When two processors are used, the maximum memory capacity is half of the displayed capacity;

3. When NVDIMM is used, a technical review needs to be submitted;

4. Please note that the installation method for 12-RAM configuration with 4 CPUs and the installation method for 6-RAM configuration with 2 CPUs are different from others;

5. When Optane™ PMem memory is used:

It must be accompanied by RDIMM RAMs, and 16G, 32G RDIMM RAMs are selectable;

It can only match with Cascade 5/6/8 series CPUs, the number of RAMs is 4;

It cannot be used together with 1RANK8 memory.

6.3 Storage

NF8480M5 supports up to 24 3.5-inch or 2.5-inch hard disks, up to 2 12*3.5-inch hard disk backplanes, and a SAS/SATA interface. The whole unit supports up to 12 U.2 NVME SSDs.

6.3.1 SATA/SAS hard disk model

Table 6-5 Hard disk options

Model	RPM	Capacity	Maximum number supported
2.5 SAS	7.2K	1T/2T	24
	10K	300G/600G/900G/1.2T/1.8T/2.4T	24
	15K	300GB/600GB/900GB	24
2.5 SATA	7.2K	1T/2T	24
3.5 SAS	7.2K	1T/2T/4T/6T/8T/10T/12T/14T	24
3.5 SATA	7.2K	1T/2T/4T/6T/8T/10T/12T/14T/16T	24

Note: NF8480M5 supports a mixed combination of four different SAS/SATA HDDs

6.3.2 SSD model

Table 6-6 SSD options

Model	Capacity
SATA SSD	240G
SATA SSD	400G
SATA SSD	480G
SATA SSD	960G
SATA SSD	1.92T
SATA SSD	3.84T
SATA SSD	7.68T
SAS SSD	480G
SAS SSD	800G
SAS SSD	960G
SAS SSD	1.6T
SAS SSD	1.92T
SAS SSD	3.2T
SAS SSD	3.84T
SAS SSD	6.4T

Note: NF8480M5 supports a mixed combination of two different SSDs

6.3.3 U.2 NVME SSD

Table 6-7 U.2 NVME SSD

Model	Capacity	Maximum quantity
U.2 NVME SSD	1T	12
U.2 NVME SSD	1.6T	12
U.2 NVME SSD	2T	12
U.2 NVME SSD	3.2T	12
U.2 NVME SSD	4T	12

U.2 NVME SSD	6.4T	12
U.2 NVME SSD	7.68T	12
U.2 NVME SSD	8T	12

Note: The optional NVME SSD is not supported when two processors are installed.

6.3.4 PCI-E NVME SSD

Table 6-8 PCI-E NVME SSD options

Model	Capacity
PCI-E NVME SSD	1.5T
PCI-E NVME SSD	1.6T
PCI-E NVME SSD	2T
PCI-E NVME SSD	3.2T
PCI-E NVME SSD	4T
PCI-E NVME SSD	6.4T

Note: NF8480M5 supports up to 4 optional PCI-E SSDs.

6.3.5 M.2 hard disk

Table 6-9 M.2 hard disk

Model	Capacity
M.2 SATA SSD	240G
M.2 SATA SSD	480G
M.2 SATA SSD	960G
M.2 SATA SSD	1.92T

6.4 RAID/SAS card

Table 6-10 8i RAID/SAS card

SAS card	Inspur	SAS card_INSPUR_SAS3008+IT+PCIE3.0
		SAS card_INSPUR_PM8222_SmartHBA_8_SAS3_PCIE3
		SAS card_INSPUR_PM8222_HBA_8_SAS3_PCIE3
RAID card	Inspur	RAID card_INSPUR_SAS3108_2GB_SAS12G_PCIE3
		RAID card_INSPUR_SAS3108_4GB_SAS12G_PCIE3
		RAID card_Inspur_PM8204_RA_8_2GB_SAS3_PCIE3
		RAID card_Inspur_PM8204_RA_8_4GB_SAS3_PCIE3
		RAID card_INSPUR_ PM8060_2GB_SAS12G_PCIE3.0
	LSI	RAID card_L_8R0_9361-8i_1GB_HDM12G_PCIE3.0
		RAID card_L_8R0_9361-8i_2GB_HDM12G_PCIE3.0
		RAID card_L_8R0_9460-8i_2GB_HDM12G_PCIE3
		RAID card_L_8R0_9460-8i_4GB_HDM12G_PCIE3

Table 6-11 16i and 24i RAID cards

16i RAID card	LSI	RAID card_L_16R0_9361-16i_2GB_HDM12G_PCIE3
		RAID card_L_16R0_9460-16i_4GB_HDM12G_PCIE3
24i RAID card	Adaptec	RAID card_AD_24R0_3154-24i_4GB_ HDM12G_PCIE3
	LSI	RAID card_L_24R0_9361-24i_4GB_HDM12G_PCIE3

Note: The 8i raid card, the 16i raid card or the 24i raid card cannot be mixed. A maximum of 8i raid cards can be supported.

6.5 Network adapter

Table 6-12 PHY card

Type	Model & Description	Speed rate	Number of interfaces
PHY card	NIC_10Gbps_2Port_LC_Inspur_CS4227_PHY_SM	10G	2
	NIC_10Gbps_2Port_RJ45_Inspur_X557_PHY	10G	2
	NIC_10Gbps_4Port_LC_Inspur_CS4223_PHY_MM	10G	4
	NIC_10Gbps_4Port_LC_Inspur_CS4223_PHY_SM	10G	4
	NIC_10Gbps_2Port_LC_Inspur_CS4227_PHY_MM	10G	2
	NIC_1Gbps_4Port_RJ45_Inspur_1543_PHY	1G	4

Note: Gigabit PHY cards are not backward compatible with 100M switches. A PHY card does not support NCSI function. If you need this function, please choose OCP network adapters.

Table 6-13 OCP card

Type	Model & Description	Speed rate	Number of interfaces
OCP	NIC_1Gbps_2Port_RJ45_OEM-PurchasedParts_I350_OCP2.0	1G	2
	NIC_10Gbps_1Port_LC_Inspur_82599_OCP2.0_MM	10G	1
	NIC_25Gbps_1Port_LC_Mellanox_CX4_OCP2.0_MM	25G	1
	NIC_25Gbps_1Port_LC_Inspur_CX4LX_OCP2.0_MM	25G	1
	NIC_25Gbps_2Port_LC_Inspur_CX4LX_OCP2.0_MM	25G	2
	NIC_25Gbps_2Port_LC_Mellanox_CX4_OCP2.0_MM	25G	2
	NIC_25Gbps_2Port_LC_Mellanox_CX5_OCP2.0_SM_3	25G	2
	NIC_25Gbps_2Port_LC_Mellanox_CX5_OCP2.0_MM	25G	2
	NIC_25Gbps_2Port_LC_Mellanox_CX5_OCP2.0_3	25G	2

Table 6-14 Standard PCI-E network adapter

Type	Model & Description	Speed rate	Number of interfaces
PCI-E network adapter	NIC_1Gbps_2Port_RJ45_Intel_I350_PCIe	1G	2
	NIC_1Gbps_4Port_RJ45_Inspur_I350_PCIe	1G	2
	NIC_1Gbps_4Port_RJ45_Intel_I350_PCIe	1G	4
	NIC_10Gbps_1Port_LC_Inspur_82599_PCIe_MM	10G	1
	NIC_10Gbps_1Port_LC_Inspur_82599_PCIe_SM	10G	1
	NIC_10Gbps_1Port_LC_Intel_82599_PCIe_MM	10G	1
	NIC_10Gbps_2Port_LC_Broadcom_BCM57412_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Broadcom_BCM57412_PCIe_SM	10G	2
	NIC_10Gbps_2Port_LC_Cavium_BCM57840_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Inspur_82599_PCIe	10G	2
	NIC_10Gbps_2Port_LC_Inspur_82599_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Inspur_82599_PCIe_SM	10G	2
	NIC_10Gbps_2Port_LC_Inspur_X710_PCIe	10G	2
	NIC_10Gbps_2Port_LC_Inspur_X710_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Inspur_X710_PCIe_SM	10G	2
	NIC_10Gbps_2Port_LC_Intel_82599_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Intel_X710_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Mellanox_CX3PRO_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Mellanox_CX3PRO_PCIe_SM	10G	2
	NIC_10Gbps_2Port_LC_Mellanox_CX4LX_PCIe	10G	2
	NIC_10Gbps_2Port_LC_Mellanox_CX4LX_PCIe_MM	10G	2
	NIC_10Gbps_2Port_LC_Mellanox_CX4LX_PCIe_SM	10G	2
	NIC_10Gbps_2Port_LC_Solarflare_SFC9240_PCIe_MM_Base	10G	2

NIC_10Gbps_2Port_RJ45_Intel_X540_PCIe	10G	2
NIC_10Gbps_2Port_RJ45_Intel_X550_PCIe	10G	2
NIC_10Gbps_4Port_LC_Intel_X710_PCIe_MM	10G	4
NIC_10Gbps_4Port_LC_Intel_X710_PCIe_SM	10G	4
NIC_25Gbps_1Port_LC_Mellanox_CX4LX_PCIe	25G	1
NIC_25Gbps_1Port_LC_Mellanox_CX4LX_PCIe_MM	25G	1
NIC_25Gbps_2Port_LC_Broadcom_BCM57414_PCIe_MM	25G	2
NIC_25Gbps_2Port_LC_Mellanox_CX4LX_PCIe	25G	2
NIC_25Gbps_2Port_LC_Mellanox_CX5_PCIe	25G	2
NIC_25Gbps_2Port_LC_Mellanox_CX5_PCIe_MM	25G	2
NIC_25Gbps_2Port_LC_Mellanox_CX5_PCIe_SM	25G	2
NIC_40Gbps_1Port_LC_Mellanox_CX3PRO_PCIe_MM	40G	1
NIC_40Gbps_2Port_LC_Intel_XL710_PCIe_MM	40G	2
NIC_40Gbps_2Port_LC_Mellanox_CX4_PCIe	40G	2
NIC_40Gbps_2Port_LC_Mellanox_CX4_PCIe_MM	40G	2
NIC_100Gbps_1Port_LC_Mellanox_CX4_PCIe_MM	100G	1
NIC_100Gbps_2Port_LC_Mellanox_CX4_PCIe_MM	100G	2
NIC_100Gbps_2Port_LC_Mellanox_CX5_PCIe_MM	100G	2

Note: A maximum of four self-developed X8 network adapters are supported, and up to 3 types of self-developed network adapters are allowed to be mixed, but they can only be inserted into riser card slots instead of onboard PCIE slots;

A maximum of 6 self-developed network adapters of each type are supported and the mixing of up to three types is allowed.

A 100G network adapter must be installed in the PCIEX16 slot;

An external network adapter does not support NCSI function. If you need this function, please choose OCP network adapters.

6.6 FC HBA card

Table 6-15 FC HBA card

Type	Model & Description	Speed rate	Number of interfaces
FC HBA card	HBA_1Port_FC8Gb_Qlogic	8Gb/s	1
	HBA_1Port_FC8Gb_Emulex	8Gb/s	1
	HBA_2Port_FC8Gb_Qlogic	8Gb/s	2
	HBA_2Port_FC8Gb_Emulex	8Gb/s	2
	HBA_1Port_FC16Gb_Qlogic	16Gb/s	1
	HBA_1Port_FC16Gb_Emulex_LPE31000-AP	16Gb/s	1
	HBA_2Port_FC16Gb_Emulex	16Gb/s	2
	HBA_2Port_FC16Gb_Qlogic	16Gb/s	2
	HBA_1Port_FC32G_Qlogic	32Gb/s	1
	HBA_2Port_FC32G_Qlogic	32Gb/s	2
	HBA_2Port_FC32G_Emulex	32Gb/s	2

Note: A maximum of 4 HBA cards of each type are supported, and the mixing of up to two types is allowed.

6.7 HCA card

Table 6-16 HCA card

Type	Model & Description	Speed rate	Number of interfaces
HCA card	HCA_1Port_EDR100G_M	100Gbps	1

	HCA_1Port_EDR100G_I	100Gbps	1
	HCA_2Port_EDR100G_M	100Gbps	2
	HCA_1Port_HDR100G_M	100Gbps	1

Note: A maximum of 2 HCA cards of each type are supported, and the mixing of up to two types is allowed.

A 100G HCA card must be installed in the PCIEX16 slot;

6.8 GPU

Table 6-17 GPU

Type	Model & Description	Maximum quantity
GPU card	GPU_Nvidia_Tesla-M60_16G	2
	GPU_Nvidia_Tesla-P100-PCIE_16G	2
	GPU_Nvidia_Tesla-P40_24G	2
	GPU_Nvidia_Tesla-P4_8G	2
	GPU_Nvidia_Tesla-V100-PCIE_32G	2
	GPU_Nvidia_Tesla-T4_16G	2
	GPU_Nvidia_Tesla-V100-PCIE_16G	2
Graphics card	Graphics_P620_2G	2
	Graphics_P400_2G_PG178	2
	Graphics_P5000_16G	2
	Graphics_P4000_8G	2
	Graphics_P620_2G	2
	Graphics_P400_2G_Ver3.0	2
	Graphics_P4000_8G	2

Note: The GPU has an X16 bus and needs to be installed in the X16 slot of the Riser card.

6.9 Power supply

The power supply adopts Intel's standard CRPS, and general electrical and structural design; supports hot swap and N+N redundancy; and can load up to 4 power supplies. The power supply supports tool-free disassembly and assembly, and is automatically locked when inserted into a server. The CRPS power supply meets 80PLUS platinum efficiency and provides a variety of output power. Users can choose power supplies of different power rates according to specific configurations.

● Support the following rated AC 110V~230V power supply and DC 240V power supply, as well as multiple redundancy modes including 1+1/2+1/2+2/3+1 modes

550W platinum power supply: 550W (110VAC), 550W (230VAC), 550W (240VDC for China)

800 W platinum power supply: 800 W (110VAC), 800 W (230VAC), 800 W (240VDC for China)

1300 platinum power supply: 1000W (110VAC), 1300W (230VAC), 1300W (240VDC for China)

1600 W platinum power supply: 1000 W (110VAC), 1600 W (230VAC), 1600 W (240VDC for China)

Note: The rated power of 1300W/1600W will drop to 1000W at the rated 110VAC power supply.

Input voltage range:

110VAC~230VAC: 90V ~ 264V

240VDC: 190V ~ 300V

● Support the following rated DC 240V~336V power supply, as well as multiple redundancy modes including 1+1/2+1/2+2/3+1 modes

550W 336VDC power supply: 550W (240VDC), 550W (336VDC)

800W 336VDC power supply: 800W (240VDC), 800W (336VDC)

Input voltage range:

240VDC~336VDC: 190V ~ 400V

- Support the following rated DC -48V power supply, as well as multiple redundancy modes including 1+1/2+1/2+2/3+1 modes

800W -48VDC power supply: 800W (-48VDC)

Input voltage range:

-48VDC: -40V ~ -72V

6.10 Operating system

Table 6-18 Operating System

OS manufacturer	OS version
Windows	Windows_Server_2012R2
	Windows Server 2016
	Windows server 2019
Red Hat	Red Hat Enterprise 6.9
	Red Hat Enterprise 7.4
	Red Hat Enterprise 7.6
SUSE	SUSE12.2
	SUSE12.4
Centos	Centos_7.0
	Centos_7.4
	Centos_7.5
	Centos_7.6
Debian	Debian_9.9
Oracle Linux	Oracle Linux_6.10
	Oracle Linux_7.6
ESXi 6.X	Vmware Esxi_6.0
	Vmware Esxi_6.5
Ubuntu16	Ubutu_Server_14.04
	Ubutu_Server_18.04
InCloud Sphere	Incloud5.5
	InCloud sphere Enterprise Edition 5.6
Galaxy Kylin	Kylin-4.0.2

Note: The Cascade Lake CPU supports windows 2016 and 2019, redhat7.5 and higher version. The Skylake CPU supports all the operating systems mentioned above.

7 Configuration Considerations

- Support up to 2 GPUs;

- When the number of PCIe is greater than 9 and less than or equal to 13, or the number of self-developed network adapters is greater than 7 and less than or equal to 11, the PCIe device will conflict with NVME;

- Inspur is only responsible for the configuration of the initial shipment. For example, a 2GPU configuration will not be accompanied by additional GPUs. Any faulty GPU subsequently purchased by customers are not supported;

- Cascade Lake CPU can only run on Linux 7.6 or higher operating systems

- When the number HDD is greater than 10 or the number of SSD is greater than 12, a 9361-16i/24i RAID card is recommended instead of an Expander card for expansion

- The random read performance will decrease if it is calculated that there are more than 4 SSDs, and the read / write bandwidth will decrease if there are more than 8 SSDs, which is even worse than directly connecting to an 8i RAID card. Therefore, the combination of Expander card (82885T) with SSD is not recommended.

- The full performance will be exerted if the number of HDDs is less than or equal to 15, while the read / write bandwidth performance will decrease if the number is greater than 15.

- It is recommended to select hard disks of the same brand/model, and it is not recommended to mix up hard disks of different models.

8 System Management

NF8480M5 integrates the next-generation BMC intelligent management system. Inspur BMC intelligent management system is a remote server management system independently developed by Inspur. It is compatible with IPMI2.0, the management benchmark in the server industry, and features highly reliable and more intelligent hardware monitoring and management. The main features of Inspur BMC intelligent management system include:

● Support Intelligent Platform Management Interface (IPMI)

● Support redirection of keyboard, mouse, video and text console

● Support remote virtual media

● Support Redfish protocol

● Support simple network management protocol (SNMP)

● Support login into BMC via web browser

Table 8-1 Specifications of BMC Intelligent Management System

Specifications	Description
Management interface	Support a variety of management interfaces, meet various system integration methods; can be integrated into any standard management system, and support the following interfaces: IPMI CLI SNMP HTTPS Redfish
Fault detection	Provide rich fault detection functions to accurately locate hardware faults
Alarm management	Support alarm management and multiple formats of alarm reporting including SNMP Trap (v1/v2c/v3), Email Alert, and syslog service to ensure 7/24 highly reliable operation.
Virtual KVM	Provide convenient remote maintenance measures, and no on-site operation is required even when the system fails.
Virtual media	Support to virtualize local media devices or mirrors, USB devices, and folders into media devices of remote servers, simplify the complexity of operating system installation.
Web UI	Support visual image interface, and the setup and query tasks can be completed quickly through a simple click on the interface.
Screen shot	View screen snapshots via WebGui or Restful interface
Software dual-image backup	When the software completely running currently crashes, it can be started from the backup image.
Support intelligent power management	Power capping technology helps you easily increase deployment density, while dynamic energy saving technology helps you effectively reduce operating costs.
IPv6	Support IPv6 functions to facilitate the construction of a full IPv6 environment and provide you with abundant IP address resources.
NC-SI function	Support NC-SI (Network Controller Sideband Interface) function to allow you to easily access to BMC system through business network port.
Hardware watchdog timer	When the BMC has no response for a period longer than the safety setting time, the control fan enters the full speed protection mode
Power supply control	on/off/cycle/status
UID remote control	The UID lamp of a single machine is manually lit, which facilitates you to find equipment in the machine room
Firmware upgrade	BMC/BIOS can be upgraded
Serial port redirection	Serial port I/O in IPMI session redirection system based on IP
Storage information viewing	Display Raid logical array information and the information of corresponding physical disks under logical array

9 Certification

As of April 2020, NF8480M5 has obtained the following certifications:

Region	Certified item	Certification logo	Mandatory/Voluntary	Description
China	3C		Mandatory	
	Environment label		Voluntary	
International mutual recognition	CB		Voluntary	
EU	CE		Mandatory	
USA	FCC		Mandatory	
	UL		Voluntary	
	Energy star		Voluntary	
Russia	CU certification		Mandatory	
	Information Security	N/A	Mandatory	
Korea	E-Standby Energy Efficiency Certification		Mandatory	
	KC certification		Mandatory	
Australia	RCM		Mandatory	

10 Support and Service

Global service hotline:

- 1-844-860-0011 (toll free)

- 1-646-517-4966 (direct line)

- Service email: serversupport@inspur.com

Information required from customers:

- Name

- Tel

- E-mail

- Product model

- Product service SN number

- Problem description

11 Description of New Technologies

11.1 Intel Scalable Architecture

Intel's next-generation Xeon processor based on Skylake and Cascade Lake architectures will begin to adopt a new Mesh interconnect architecture design on the chip design architecture to replace the traditional Ring interconnect design, so as to improve CPU access latency and support higher memory bandwidth requirements. At the same time, featured by low power consumption, it allows the processor to operate at a low processor clock speed and at a relatively lower voltage environment, so as to provide better performance improvement and improve energy efficiency. Compared with products of previous generations, the overall performance of Intel Xeon scalable processors has increased by 1.65 times and OLTP warehouse load has increased by 5 times compared with current system.

11.2 Intel VROC Technology

Intel VROC technology represents Virtual RAID on CPU and is specially designed for enterprise RAID solutions based on NVME SSD. The biggest advantage is that it can directly manage connections to Intel's scalable PCI-E channels without using a dedicated RAID HBA.

11.3 QAT Technology

Intel® QuickAssist Technology (Intel® QAT) accelerates applications by speeding up computationally intensive operations. It provides a software-based foundation for security, authentication and compression, and can significantly improve the performance and efficiency of standard platform solutions. It is embodied in the following aspects:

In the cloud domain, it increases the throughput of applications, adds hardware acceleration for network security, routing, storage and big data, and maximizes CPU utilization.

Accelerate SSL/TLS on the network, allowing higher-performance encrypted communications in a secure network and higher platform application efficiency.

With regard to big data, the compressed file system data blocks support faster analysis, enable faster Hadoop* runtime for big data and reduce processor requirements, and can complete various operations with low latency, thus improving overall performance.

11.4 RAM Mirroring

It can prevent uncorrectable RAM errors and avoid system failures. In this mode, the system maintains two copies of all data. If an uncorrectable RAM error occurs, the system will automatically retrieve the sound data from the mirrored (redundant) copy. The system will continue to operate normally without user intervention. By providing additional mirrored redundancy in the memory system, it provides maximum protection for the entire system and prevents memory failures that occur during the correction of ECC, SDDC, DDDC and online spare memory.

11.5 Memory Rank Sparing

It provides protection against persistent DRAM failures. It tracks too many correctable errors and copies the contents from unhealthy ranks to available backup ranks before multi-bit or persistent single-bit errors occur, which may lead to uncorrectable errors in the future. It does not recognize or disable a single failed DRAM. Instead it disables DIMMs or ranks. Since DIMMs or ranks are required to perform sparing, this technology reduces the total amount of available memory by virtue of the memory used for sparing. Each DIMM of spare parts can only handle one failure. The DIMMs that may receive fatal/uncorrectable memory errors will be automatically removed from operation, thereby reducing system downtime.

11.6 Power Awareness Technology

It can monitor the real-time power consumption of CPU, RAM, HDD, fan and other high-power components in the system, and provide power consumption data for operation and maintenance.

12 Relevant Documents

For more information, please click the following link:

http://www.inspur.com

The website provides some resources to help customers solve problems and learn our products, including product manuals, drivers and firmware.

13 Trademark

Inspur and its logo belong to Inspur Group Co., Ltd. All other trademarks and trade names mentioned in this document are the property of their respective owners.