R925 2U All-round AI Server

The R925 is a 2U dual-socket high-density rack server based on 4th/5th Gen Intel® Xeon® Scalable processors, boasting flexible resource configuration and wide industry adaptability. Built specifically to configure resources flexibly based on service requirements to meet application needs of various industries, it stands out in ensuring optimal performance for application scenarios such as machine learning, AI inference, cloud computing, high-performance computing, cloud storage, big data, virtualization, office systems, and more.



Extremely Cost-Effective AI Computing Server

Specifications

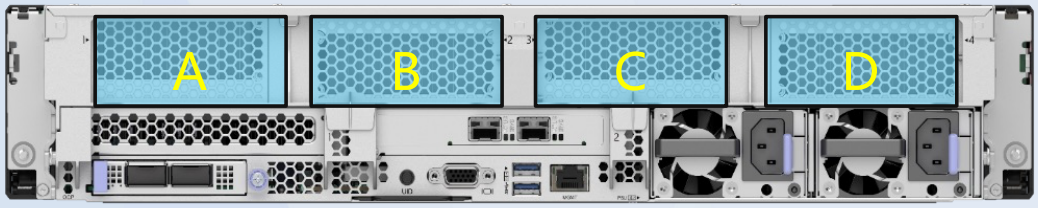
Model	R925
Processor	Supports 2*4th/5th Gen Intel Xeon® Scalable Processors (TDP 385W)
Memory	Up to 32*DDR5 slots, up to 5600MHz
RAID	Supports BCM 95 series RAID cards.
Storage	Up to 12*3.5" hot-swappable drives (12*NVMe compatible) Internal: 2*SATA/PCIe M.2
Network Interfaces	Network card: 1*OCP 3.0, with NCSI Optional OCP 3.0 interfaces: 2*GE, 2*10GE, 2*25GE, 2*100GE, 2*200GE
GPU	Supports up to 4*double-slot GPU or 4*single-slot GPU

Additional Interface	Rear: 1*RJ-45 management port, 2*USB 3.0, 1*VGA Front:2*USB 3.0 + 1*VGA
Power Supply	Supports 2000W/2200W/2700W CRPS high-efficiency Platinum PSU(hot-swappable, 1+1 redundant) Voltage: 100-240V AC (refer to nameplate)
Fan	6*hot-swappable fans, supports N+2 redundant
Management	Integrated BMC with 1Gbps RJ-45 port. Supports Redfish, SNMP, IPMI 2.0, HTML5/iKVM remote control. BMC dual-flash backup
OS Support	Windows Server, RHEL, SLES, CentOS, Ubuntu, VMware ESXi, etc.
Operating Temperature	5℃ ~ 35℃
Chassis Dimensions	770mm(L)*482mm(W) *87mm(H)

Product View



3GPU (3*double-slot or 4*single-slot GPU) riser model



4GPU (4*double-slot GPU) riser model

Product Features

Strong computing power, more flexible intelligent computing matching

Elevate your IT efficiency with 2U dual-socket high-density performance. Equipped with 4th/5th Gen Intel® Xeon® Scalable processors, it combines flexible resource configuration with strong industry adaptability—all within a compact 2U dual-socket form factor. Tailored specifically to configure resources flexibly based on service requirements to meet application needs of various industries, it stands out in excelling at scenarios including machine learning, AI inference, cloud computing, high-performance computing, and more.

High Reliability and Security

Cutting system downtime by 69%, the proprietary AI-powered memory fault self-healing technology guarantees uninterrupted, stable system operation.

Efficient Energy Saving

Proprietary smart algorithms minimize power draw for fans and CPUs to further reduce energy use.

With Intelligent Service Awareness & Dynamic Load Adjustment, the CPU operating frequency is dynamically tuned in real time according to actual service workloads, optimizing performance and energy efficiency intelligently.

Intelligent O&M

Delivering 30× higher upgrade efficiency, the tool-free design enables fast component assembly and disassembly to shorten O&M time. Out-of-band visual management supports precise remote localization of physical device status and information, enabling intelligent full-lifecycle services for the product.