



LAND



SEA



AIR

AV2000

CONDUCTION LIQUID COOLED
MULTI CPU-GPU SERVER



- Thermal Dissipation Coefficient : 15KW
- Liquid Cooled Platform : Gun Drilled -10 x 600 x 10mm Pass (\varnothing x L x H)
- CPU: Intel 3rd Gen. Xeon Scalable(ICELAKE)[®] Platinum 8380 (40 x cores, 2.3 GHz, 270W)
- GPU: 2 x NVIDIA[®] A100 80GB (6912 Cuda Cores, HBM 2e) PCIe Gen.4 x 16
- 2TB RDIMM ECC DDR4-3200MHz
- 8x 10GBase-T Ethernet
- 1x 1GBase-T Ethernet for IPMI LAN
- Storage (1) : RAID 1 - 2TB by Gen 4.0 NVMe U.2
- Storage (2) : RAID 10 - S-SATA 32TB by SATA SSD
- Dimension: 480 x 600 x 132 mm (WxDxH)



INDEX

1. INTRODUCTION

2. Intel Ice Lake platform

-NVIDIA A100 GPU

-Quad 10GbE X710-T4

-High Point RAID SSD-7502

3. Conduction Liquid Cooling Plate (C.L.C.P.)

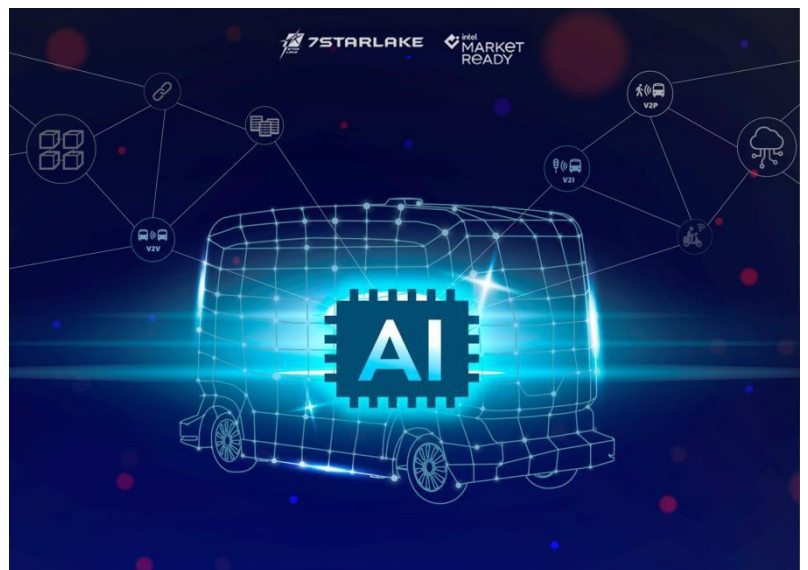
4. How AV2000 Support 32TB

S-SATA RAID 10, AND 2TB NVMe RAID 1

5. Specification

1. Introduction

The global transformation is rapidly scaling the demands for flexible computer, networking, and storage. Future workloads will necessitate infrastructures that can seamlessly scale to support immediate responsiveness and widely diverse performance requirements. The exponential growth of data generation and consumption, the rapid expansion of cloud-scale computing and 5G networks, and the convergence of high-performance computing (HPC) and artificial intelligence (AI) into new usages requires that today's data centers and networks evolve now— or be left behind in a highly competitive environment.



7Starlake's AV2000 Conduction Liquid Cooled AI inference GPU Server features 3rd Gen. Intel Xeon Scalable (ICE LAKE) ® Platinum 8380 processor (40 Cores, 2.3 GHz, 270W) with 2 x NVIDIA A100 GPU, 2TB RDIMM ECC DDR4-3200MHz and 2TB by Gen 4.0 NVMe U.2 (RAID1), 32TB by SATA SSD (RAID10), to provide the seamless performance foundation for the data centric era from the multi-cloud to intelligent edge, and back.

AV2000 enables a new level of consistent, pervasive, and breakthrough performance in new AI inference to implement machine learning and deep learning. Combining stunning inference performance, powerful CPU and expansion capability, it is the perfect ruggedized platform for versatile edge AI applications.

With innovative liquid cooled design, AV2000 is built in most advanced " Gun Drilled ", which with 10 pipes (each pipe 5mm x 5mm x π x 800mm) to dissipate max 15KW heat.



2. Intel's 10nm Ice Lake Xeon With 40 Cores , PCIe 4.0

Today's scientific discoveries are fueled by innovative algorithms, new sources and volumes of data, and advances in compute and storage. Machine learning, deep learning, and AI converge the capabilities of massive compute with the flood of data to drive next-generation applications, such as autonomous systems and self-driving vehicles. Recognizing this demand, 7STARLAKE employs 3rd Gen Intel® Xeon® Scalable processors that are built specifically for the flexibility to run complex AI workloads on the same hardware as existing workloads to be the CPU of AV2000.

Intel Ice Lake is based on 3rd Generation Intel Xeon Ice Lake Scalable Processor in LGA4189 Socket with C621A chipset, up to 40 cores, 270W TDP. Intel Ice Lake-SP processors will be based on the 10nm+ process node. The main highlight of Ice Lake-SP processors will be support for PCIe Gen 4 and 8-channel DDR4 memory.

CAN FD
 A100
 A100
 Quad 10GbE
 Quad 10GbE
 2TB RAID 1 NVMe
 S-SATA 32TB RAID 10

A1

PEAK
PCAN-PCI Express FD Dual Channel

A2

NVIDIA
A100 TENSOR CORE GPU

A3

NVIDIA
A100 TENSOR CORE GPU

A4

Intel
X710T4BLK Quad10Gbit Ethernet
Converged Network Adapter RJ45

A5

Intel
X710T4BLK Quad10Gbit Ethernet
Converged Network Adapter RJ45

A6

Highpoint
SSD7502 PCIe 4.0 x16 2 port
M.2 NVMe RAID AIC RAID

A7

Seagate S-SATA 32TB RAID 10
16TB 7200RPM 3.5" Enterprise (ST16000NM001G)

NVIDIA A100 GPU

NVIDIA A100 Tensor Core GPU delivers unprecedented acceleration at every scale to power the world's highest-performing elastic data centers for AI, data analytics, and HPC. Powered by the NVIDIA Ampere Architecture, A100 is the engine of the NVIDIA data center platform. A100 provides up to 20X higher performance over the prior generation and can be partitioned into seven GPU instances to dynamically adjust to shifting demands. Available in 40GB and 80GB memory versions, A100 80GB debuts the world's fastest memory bandwidth at over 2 terabytes per second (TB/s) to run the largest models and datasets.

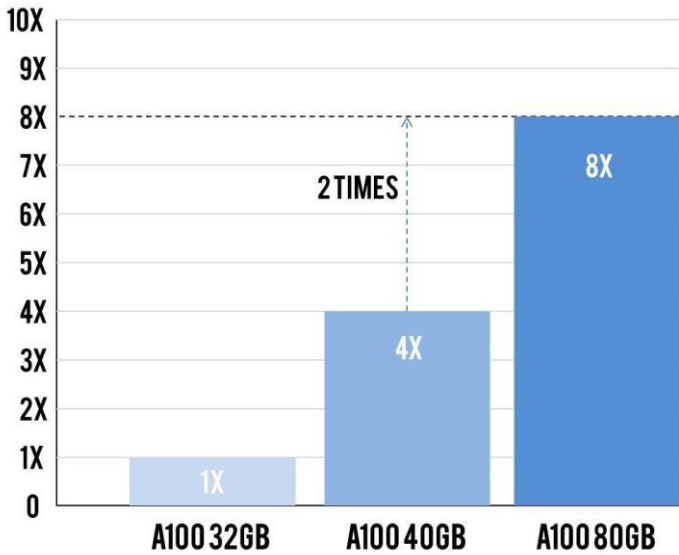


NVIDIA

NVIDIA A100 80GB PCIe

FP64	9.7 TFLOPS
FP64 Tensor Core	19.5 TFLOPS
FP32	19.5 TFLOPS
Tensor Float 32 (TF32)	156 TFLOPS 312 TFLOPS
BFLOAT16 Tensor Core	312 TFLOPS 624 TFLOPS
FP16 Tensor Core	312 TFLOPS 624 TFLOPS
INT8 Tensor Core	624 TOPS 1248 TOPS
GPU Memory	80GB HBM2e
GPU Memory Bandwidth	1,935GB/s
Max Thermal Design Power (TDP)	300W
Multi-Instance GPU	Up to 7 MIGs @ 10GB
Form Factor	PCIe
Interconnect	NVIDIA® NVLink® Bridge for 2 GPUs: 600GB/s PCIe Gen4: 64GB/s
Server Options	Partner and NVIDIA-Certified Systems™ with 1-8 GPUs

A100 is part of the complete NVIDIA data center solution that incorporates building blocks across hardware, networking, software, libraries, and optimized AI models and applications. Representing the most powerful end-to-end AI and HPC platform for data centers, it allows researchers to rapidly deliver real-world results and deploy solutions into production at scale.



A100 introduces groundbreaking features to optimize inference workloads. It accelerates a full range of precision, from FP32 to INT4. Multi-Instance GPU (MIG) technology lets multiple networks operate simultaneously on a single A100 for optimal utilization of compute resources. And structural sparsity support delivers up to 2X more performance on top of A100's other inference performance gains. A100 80GB's increased memory capacity doubles the size of each MIG and delivers up to 1.25X higher throughput over A100 40GB.

-Quad 10GbE X710-T4

The Intel Quad Port 10GbE Ethernet Converged Network Adapter X710T4 is the foundation for server connectivity, providing broad interoperability, critical performance optimizations, and increased agility for Telecommunications, Cloud, and Enterprise IT network solutions. The X710T4 offers four ports of wire-speed 10GbE over CAT6A cabling and is backward compatible to 1GbE and 100MB.



Quad 10GbE x710-t4

Intel Ethernet Converged Network Adapter X710-T4 network adapter - PCIe 3.0 x8 - 10Gb Ethernet x 4

Product Description	Intel Ethernet Converged Network Adapter X710-T4 network adapter - PCIe 3.0 x8 - 10Gb Ethernet x 4
Device Type	Network adapter
Form Factor	Plug-in card - low profile
Interface (Bus) Type	PCI Express 3.0 x8
Ports	10Gb Ethernet x 4
Dimensions (WxDxH)	16.7 cm x 6.87 cm
Data Link Protocol	100Mb LAN, GigE, 10 GigE
Data Transfer Rate	10 Gbps
Network / Transport Protocol	TCP/IP, UDP/IP, iSCSI, SMB, NFS
Compliant Standards	IEEE 802.3, IEEE 802.1Q, IEEE 802.1p, IEEE 802.3ad (LACP), IEEE 802.1as, IEEE 802.1Qbg

-HighPoint RAID SSD-7502

The SSD7502 is the industry's fastest Dual-Port PCIe Gen4 RAID controller, and features bootable RAID support for Windows and Linux based platforms. The low-profile SSD7502 incorporates a state of the art Gen4 PCIe Switch Chipset to ensure each M.2 port has x4 lanes of dedicated bandwidth. HighPoint's 7500 Series combine dedicated, cutting edge PCIe Gen 4 x16 host connectivity with our industry proven RAID technology to deliver unbeatable storage performance.

RAIDSSD-7502

Low-Profile Bootable Gen4 NVMe Solution for Linux & Windows platforms

Works with any PCIe 4.0/3.0 x16 slot

Up to 2 off-the-shelf M.2 SSD's

RAID 0, 1, single disk

Dedicated PCIe 3.0 x16 bus bandwidth; Dedicated x4 bandwidth for each NVMe M.2 SS




Conduction Liquid Cooling Plate (C.L.C.P.)

Most liquid cooling solutions are using close loop design — Direct to Chip (D2C) integrated pump & cold plate in the system. Users may worry about potential risk of liquid leakage.



7Starlake highly values system reliability. In the pursuit of stability and power, our experienced team has successfully optimized the thermal design, bringing out an unprecedented model AV2000. Instead of normal D2C design, 7Starlake innovated an unique heat exchanger integrating **Conduction Liquid Cold Plate (CLCP)** on the computing system.

CLCP includes multi-channel cold water inlet/outlet owning high flexibility in adjusting numbers of inlet/outlet (up to 4in 4out) by request. When coolant flows through top sink, liquid can absorb the heat and take it away from the heat sources quickly to the heat exchanger. When heated



liquid flushes into the heat exchanger, it will be cooled by 9 units of 12 x 12 cm active fan which can run at 2K~3K RPM.

Leveraging both liquid-cooling and air-cooling’s strong points; these features accomplish higher rack density and efficiency, comprehensive reduction in power use, and increase of overclocking potential.



AV2000

Item	Description
CPU	2 x Intel®Xeon® Platinum 8380
GPU	2 x NVIDIA A100
System Size	480 x 600 x 132mm (W x L x H)
CLCP	450 x 600 x 20 mm (W x L x H)
Material of CLCP	AL6063
Gun Drilled Hole∅	10mm

Leveraging both liquid-cooling and air-cooling’s strong points; these features accomplish higher rack density and efficiency, comprehensive reduction in power use, and increase of overclocking potential.




4. How AV2000 Support 32TB S-SATA RAID 10, AND 2TB NVMe RAID 1

PCIe Gen 4 is the upper version of PCIe Gen 3, which surpassed PCIe Gen 2 and Gen 1. The bandwidth provided by PCIe Gen 4 is double as compared to PCIe Gen 3. It is backward compatible with prior generations of PCIe. PCIe Gen 4 is expected to satisfy, to a large extent, the requirements of high speed servers, gaming, graphics and data centers, where number of servers and solid-state devices are increasing as per market demand. PCIe Gen 4 is the new evolution over the PCIe Gen 3. The PCIe Gen 4 provides the data rate of 16 G/Ts as compared to 8G/Ts provided by PCIe Gen 3. It's architecture is fully compatible with all the previous generations of PCIe.

7Starlake is devoted to uncovering the unique way of optimising storage. AV2000 supports 2 slot PCIe x 8 and 4 slot PCIe x16 expansions that provide extreme PCIe Gen 4 speed for GPUs at an exceptional value.

AV2000 can make the best use of two PCIe x 8 host connections that associate with four PCIe gen 4 U.2 NVMe. AV2000 installed with SSD7502 which is the industry's fastest Dual-Port PCIe Gen4 RAID controller, and features bootable RAID support for Windows and Linux based platforms. The low-profile SSD7502 incorporates a state of the art Gen4 PCIe Switch Chipset to ensure each M.2 port has x4 lanes of dedicated bandwidth. In other words, one port is able to back 16 TB. In addition to the 4 x PCIe U.2 NVMe, there are other two on board iPass Connector (labelled as A7) that are capable of offering 16TB each. As a result, AV2000 is able to surmount the limitations of traditional SSDs and supply large data storage.



-2TB NVMe RAID 1

AV2000 offers NVMe RAID Controllers are ideal for professional applications that require a bootable storage solution with host-level redundant RAID capability and native in-box driver support.

AV2000 were designed to maximize uptime of the host platform and ease the of integration and maintenance of NVMe RAID Storage. In addition, a wide range of solution packages are available for the professional VAR; customers can opt for either 2 and 4 dedicated M.2 device ports, integrated Host RAID 1, 0 & JBOD support, optional OOB management when outside of the OS platform, LED RAID status & I/O indication, UEFI/BIOS/CLI/GUI management interfaces.

-32TB RAID 10

AV2000 fully support Intel's Rapid Storage Technology enterprise (RSTe) protects data against hard drive failure when the system is configured for RAID 1, RAID 5, or RAID 10. These RAID settings are fault-tolerant because they copy and store data on multiple hard drives to avoid data loss or system downtime caused by single hard drive failure. A RAID 0 setting, however, increases response time on data intensive applications because it will access data on multiple hard drives simultaneously.

AV2000 RAID 10 combines the benefits of RAID 1 and RAID 0. Read and write performance is increased, but only half of the total space is available for data storage. Four or more drives are required making the cost relatively high, but the performance is great while providing fault tolerance at the same time. In fact, a RAID 10 can sustain multiple drive failures—provided the failures are not within the same subgroup. RAID 10 is ideal for applications with a high input/output demand such as database servers.

5. Specifications

SYSTEM

Processor	Intel® Xeon® ICELAKE Platinum 8380 Processor, 40Cores, 80Threads, Base Frequency 2.3GHz Max Turbo Frequency 3.40GHz
Memory type	Up to 2TB RDIMM, DDR4 3200MHz

GPU

Graphics Card	2 x NVIDIA A100 80GB (2TB/s GPU Memory Bandwidth, HBM2e, 6912 CUDA Cores)
---------------	---

STORAGE

SDD (bootable)	2 x NVMe PCIe GEN 4.0 x 2TB with RAID 1
SDD (Drive Pack)	32TB SAS/SATA with RAID10 hard drive array

LAN

8 x 10GBase-T	2 x Intel X710-T4 Quad Port 10GbE RJ45 PCIe 3.0 x8 Converged Network Adapter
	1 x 1GbEbase-T RJ45 for IPMI LAN

RAID

	1 x 2 Port PCIe Gen4 x8 NVMe RAID Integrated Host RAID 1 & 0
--	--

CAN FD I/O

	1 x PCAN-PCI express FD Dual Channel
--	--------------------------------------

FRONT I/O

X1	DC IN
X2	1 x USB 3.1
X3	1 x VGA
X4	1 x IPMI LAN
X5	1 x 1GbE LAN
X6	1 x COM (RS232)
X7	2 x CAN FD (M12)
X8	8 x10GbE (M12)
Power Button	1 x Power Button with LED backlight

POWER REQUIREMENT

Power Input	DC 24V 2500W(Options for 2500W AC Redundant)
-------------	--

APPLICATIONS

Applications	AV2000 Liquid cooling autonomous AI edge computer workloads and speeds time to insight by optimizing the latest PCIe Gen 4 technologies. Features include 7SL unique Liquid cooling solution Intel 3 rd generation Xeon Scalable processors. NVidia A100 80GB GPU. Modular and flexible I/O, up to 32TB storage and expansion slots.
--------------	---

OPERATING SYSTEM

Operating System	Windows 10 64Bit, Linux by request.
------------------	-------------------------------------

PHYSICAL

Dimension (W x D x H)	480 x 600 x 132 (mm)
-----------------------	----------------------

Weight	20KGS
--------	-------

Chassis	Aluminum Alloy, Corrosion Resistant
---------	-------------------------------------

Finish	Anodic aluminum oxide (Color Iron gray)
--------	---

Cooling	7SL Liquid Cooling
---------	--------------------

ENVIRONMENTAL

Reliability	Designed & Manufactured using ISO 9001 Certified Quality Program.
-------------	---

EMC

Operating Temp.	0 to +60°C
-----------------	------------

Storage Temp.	0 to +70°C
---------------	------------

Relative Humidity	5% to 95%, non-condensing.
-------------------	----------------------------