



**7STARLAKE**

**AV1000**



**Self Contained Rugged Liquid Cooled Micro Cloud System**  
**AV1000 Conduction Liquid Cooled GPU Server**  
**HE4K 4KW Heat Exchanger**

# INDEX

0. Overall introduction
1. What is Conduction Liquid Cooled Plate ?
2. 3rd Gen. Intel Xeon Scalable Platform
3. AI Inference GPU
4. NVMe 48TB



# Introduction



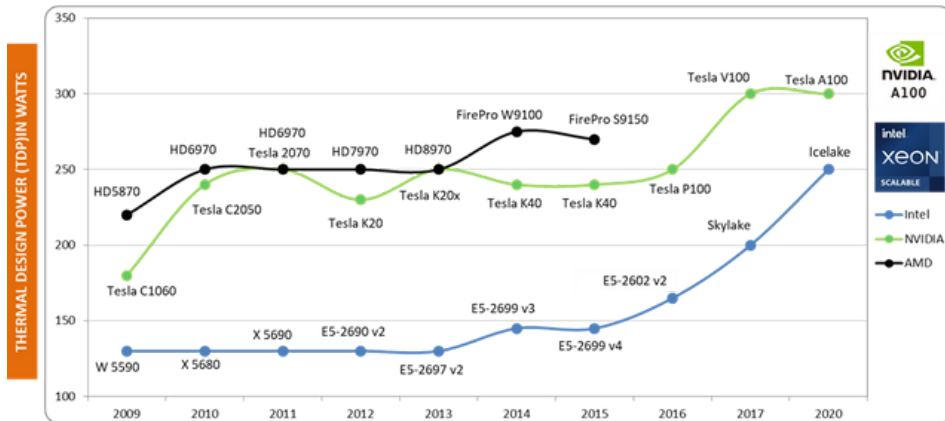
# Reasons to Adopt Rugged Liquid Cooling -The New Era of HPC thermal solution



In the past decade, the demands of cloud, IoT, AI, and edge applications are once again resulting in IT technology changes. To fulfill the business' needs of Data centers, high-performance and high-power GPUs and CPUs are being mandated to simultaneously increase energy efficiency, consolidate operations, it cause CPU and GPU's total power consumption is increasing dramatically. Each year, data centers consume approximately 2% of global power consumption. The Green Computing revolution is well under way; to utilize less energy is the key to make the world a better place.

Take the most popular GPU Server in the market for example, the GPU server is powered by dual Ice Lake Xeon SP processors and integrated with 4 NVIDIA A100 GPU, its TDP will be up to 2KW. It needs an efficient thermal solution to dissipate the generated heat to avoid CPU, GPU throttling.


**GPU and CPU TDP TREND**



Source: Alibaba.com

# Why Liquid Cooling Solution is Important?

In order to reduce the latency and improved the performance, CPUs, GPUs, and other components on board are moving closer, resulting in increased physical density and temperature within the server. Plus the data center rack power densities continue to increase, that will generate a tremendous amount of heat. The heat needs to be removed to protect the equipment, the liquid cooling's benefits are being explored as an important and required solution in data centers.




**Data Center Power Consumption**

PUE (Power Usage Effectiveness) is a ratio that describes how efficiently a computer in data center uses energy. It is also a standard of the power consumption of cooling system, and the ideal PUE ratio is 1.0. Take a look of the overall power consumption of data center, next to the IT systems themselves, cooling system is the biggest energy consumer in data center, that's why cooling system have such a big impact on PUE. One of the major expenses in data centers is also the power usage. Traditional air cooling system needs more energy to drive high RPM motor, and the PUE is up to 1.6. The PUE of conduction liquid cooling can lower to 1.1, Liquid cooling consumes less energy, and it dissipate the heat more efficiently that allows servers to be operated at higher ambient temperatures.

Cooling	Air	Direct to Chip(D2C)	Conduction Liquid	Immersion
PUE	1.6	1.2	1.1	1.05
Noise	High	Low	Low	Low
Performance	<20KW/Rack	20KW~60KW/Rack	20KW~60KW/Rack	60~200KW/Rack
Features	Scalable	Scalable	Ready Deployment	Scalable

# Liquid Cooling Solutions Overview





## 1. Conduction

Using conduction cooling system, the heat is captured at the interface between the copper tubing and CPU and GPU, and then the heat is dissipated to the interface through the copper tubing and the solid conduction plate mounted on the side of the server.

## 2. Direct to Chip (D2C)

Direct to chip system cools processors directly, a liquid coolant is brought via tubes directly to the chip. Heat sinks mounted directly onto the processors and GPU, and heat sinks are connected to small tubing that carried liquid to and from the component.




## 3. Immersion

In an immersion system, the rack hardware is submerged into a tank of non-conductive and nonflammable dielectric fluid. Heat will be dissipated from the components into the fluid via convection to dissipate the heat, as it flows through the rack without any additional active cooling systems or parts required.

# What is C.L.C.P?

7STARLAKE'S UNIQUE CONDUCTION LIQUID COLD PLATE

The liquid cooling can be a good solution to dissipate the heat. But most liquid cooling solutions are using close loop design — Direct to Chip (D2C) integrated pump & cold plate in the system. Some people may worry about liquid cooling solution might be a disaster because of the potential risk of leakage of the liquid. But once we prevent the leakage of the liquid, it becomes an utmost thermal solution.



In order to effectively reduce overall temperature and prevent the leakage of the liquid, 7Starlake innovated a unique heat exchanger which integrating Conduction Liquid Cold Plate (C.L.C.P.) building in most advanced " Gun Drilled ", which with 10 pipes (each pipe 5mm x 5mm x  $\pi$  x 500mm) to dissipate max 4KW heat on the computing system. A gun-drilled liquid cold plate is manufactured by drilling a series of deep holes into an aluminum plate to form the liquid flow path.



The features of gun-drilled cold plates are adaptable to any temperature range; satisfying the flatness requirements; and avoiding the risk of leakage.




▲ Gun Drill C.L.C.P.

# What is C.L.C.P?

7STARLAKE'S UNIQUE CONDUCTION LIQUID COLD PLATE

CLCP includes multi-channel cold liquid inlet/outlet owning high flexibility in adjusting numbers of inlet/outlet 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. 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.



Point	3HR	19HR
CPU Frequency( GHz)	3.20	3.22
CPU T-J (°C)	63	65
GPU Frequency( GHz)	1.95	1.95
GPU T-J (°C)	58.2	59.9
Water Tem. IN	26	26
Water Tem. OUT	35	35
Water Flow(L/H)	467	467
CH1 CPU(°C)	61.5	63.3
CH2 CPU Heatsink(°C)	60	62.2



▲ AV1000 with Liquid Cool System Thermal Testing Result



# What are the advantages of C.L.C.P?

7STARLAKE'S UNIQUE CONDUCTION LIQUID COLD PLATE

## 1. Higher Level of Efficiency

The heatsink is mounted directly on CPU or GPU that will conduct the heat away by direct contact. After that, the heat will be transferred to liquid cold plate which transfers heat more effectively because of its high thermal conductivity, it will efficiently distributes heat over convection surface.

## 2. Longer Lifespan


If the computer becomes too hot, it is possible to destroy and shorten the lifespan of the hardware inside your computer, leading to irreparable damage and potential data loss. Going conduction and liquid cooling keeps the computer operating at a consistently cool temperature.

## 3. Cools off High-Performance GPUs

These days, a high-end GPU can generate two or three times the amount of heat of a CPU. With using the high performance GPUs like RTX6000, they need the outstanding thermal solution - C.L.C.P., which offers an efficient means of cooling things down to successfully combats GPU throttling to maintain peak performance.

## 4. Rackmount Vertically with Heat Exchanger to Save Space

HE4K is a heat exchanger which is designed to dissipate 4kw of heat and can be combined with the C.L.C.P. on AV1000. Integrating C.L.C.P. thermal solution of AV1000 and air cooling solution of HE4K, offers an utmost thermal solution to dissipate heat. With going rackmount vertically, it reduces the space and easily to organize.



# Breakthrough Performance in New AI Inference

1. 3rd Gen Intel® Xeon® Scalable Processors
2. NVIDIA QUADRO RTX6000
3. Support 48TB NVMe



## 1. Outstanding Performance with Powerful 3rd Gen Intel® Xeon® Scalable Processors

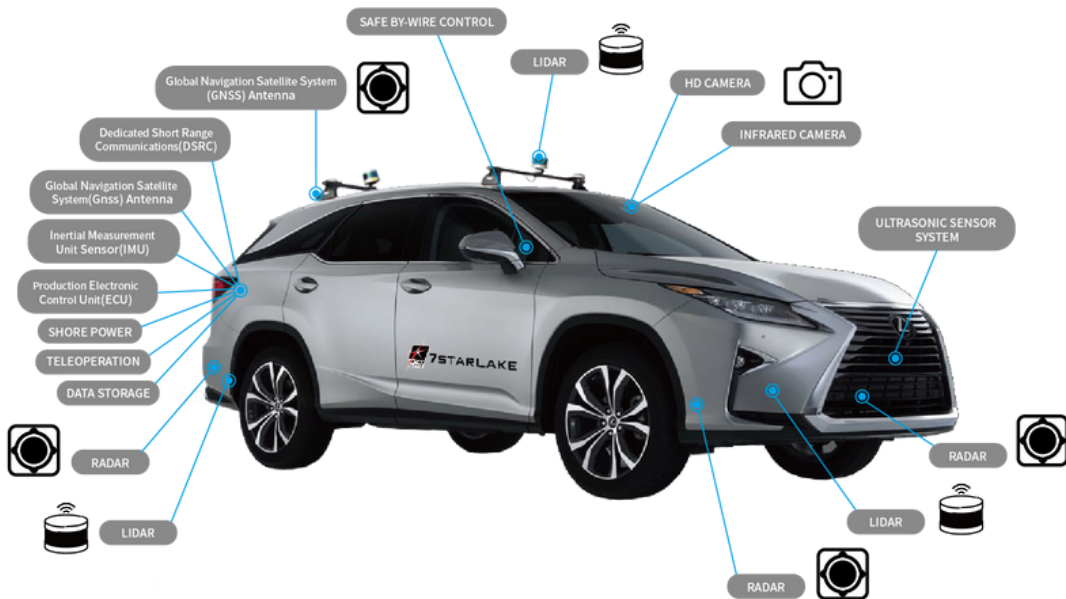

AV1000 are featuring 3rd Gen Intel® Xeon® Scalable processors. The main highlight of Ice Lake-SP processors will be support for PCIe Gen 4 and 8-channel DDR4 memory, offer a balanced architecture with built-in acceleration and advanced security capabilities, designed through decades of innovation for the most in-demand workload requirements. The processors are up to 40 cores and an array of frequency and power supports, and deliver up to 40% better performance compared to the previous generations.



Compared to previous generation

3rd Gen Xeon Scalable features:

- 1.46x average gen-on-gen performance improvement
- Up to 1.60x higher memory bandwidth
- Up to 2.66x higher memory capacity
- Up to 1.33x more PCI Express lanes per processor




# Breakthrough Performance in New AI Inference

1. 3rd Gen Intel® Xeon® Scalable Processors
2. NVIDIA QUADRO RTX6000
3. Support 48TB NVMe

## 2. NVIDIA QUADRO RTX6000

AV1000 supports 2 x NVIDIA Quadro RTX GPUs which are built on the NVIDIA Turing architecture and the NVIDIA RTX platform. With the increasing needs of large AI workloads, the NVIDIA Quadro RTX 6000 passively cooled graphics board features RT Cores and 576 tensor cores for real-time ray tracing, AI, and advanced graphics capabilities which help improving the speed of machine learning applications. Tackle graphics-intensive mixed workloads, complex design, photorealistic renders, and augmented and virtual environments at the edge with NVIDIA Quadro RTX, designed for enterprise data centers to unleash the highest application performance possible on a single server.




**SPECIFICATIONS**

GPU Architecture	NVIDIA Turing
NVIDIA CUDA Cores	4,608
NVIDIA Tensor Cores	576
NVIDIA RT Cores	72
Single-Precision Perf.	14.38 TFLOPS
Half-Precision Perf.	28.75 TFLOPS
Tensor Perf.	119.4 TFLOPS
GPU Memory	24GB GDDR6
Memory Bandwidth	Up to 624 GB/sec
Power Consumption	260W
Form Factor	PCIe 3.0 x16
Thermal Solution	Passive
Compute APIs	CUDA, DirectCompute, OpenCL™™





RTX6000



PCIe Gen4 x 16

PCIe Gen4 x 16



### Fast Storage Access

In addition to NVIDIA Quadro RTX GPUs, AV1000 provides 2 x PCIe 4.0 slots for fast storage access. Combining stunning inference performance, powerful CPU and expansion capability, it is the perfect ruggedized platform for versatile edge AI applications.


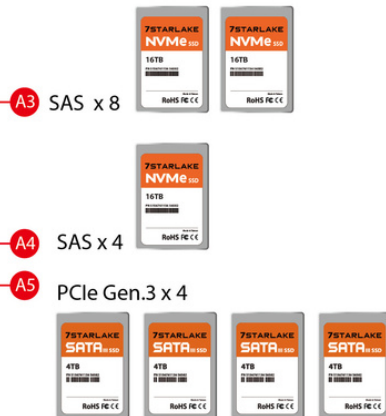

# Breakthrough Performance in New AI Inference

1. 3rd Gen Intel® Xeon® Scalable Processors
2. NVIDIA QUADRO RTX6000
3. Support 48TB NVMe

## 3. Support 48TB NVMe

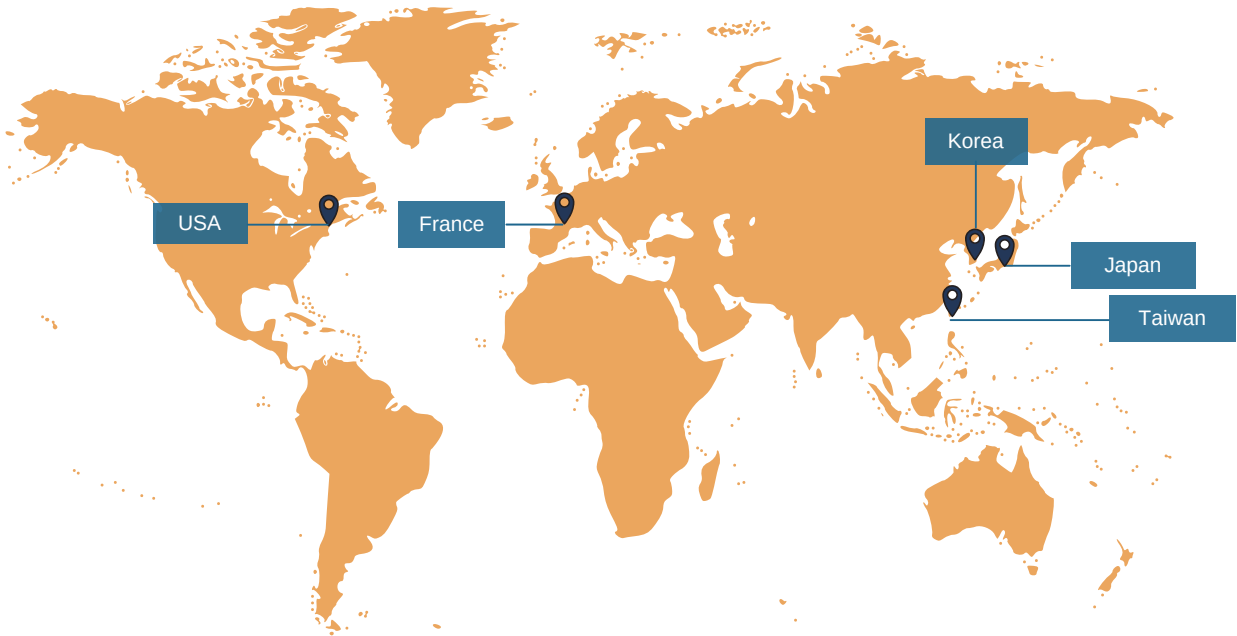
NVMe flash offers great benefits for specific AI use cases like training a machine learning model. The datasets of training of a model can be huge; and paging out the old training data and bringing in new data should be done as fast as possible to keep the GPUs from being idle.

NVMe provides low latency and protocol allowing, offers significantly high performance and superior speed. AV1000 is equipped with 1 x PCIe Gen. 4.0 x 4, 1 x PCIe Gen. 4.0 x 8 to support with 2 x NVMe SSD up to 48TB, which are expected to satisfy to a large extent, the requirements of high speed servers, gaming, graphics and data centers.





# 7STARLAKE



## Contact Us



Taiwan

Tel: +886-2-7744-7738

Korea

Tel: +82-31-304-9916

Japan

Tel: +03-3270-5921

America

Tel: +1-978-276-9787

France

Tel: +33-634024282

[7starlake.com](http://7starlake.com) | [press@7starlake.com](mailto:press@7starlake.com)

©Copyright 7starlake Co, Ltd. All Rights Reserved.