WHITE PAPER

EDA APPLICATIONS GET A PERFORMANCE BOOST WITH IMMERSION LIQUID COOLING SYSTEMS

Supermicro, in collaboration with 3M, Intel, Kaori, and Samsung, demonstrate having servers immersed in dielectric liquid boosts EDA performance, more than just cooling efficiency.

TABLE OF CONTENTS

SUPERMICRO

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry's broadest array of product configurations, to fit all computational needs.

Executive Summary

Immersion liquid cooling systems have proven their effectiveness in multiple applications, focusing primarily on improving power efficiency. This means a tangible reduction in utility costs and an enhancement in the ability to run greener in the data center. For software applications, such as EDA used in the semiconductor industry that could be running on per-core licensing models, the number of cores and clock speeds required to achieve the same level of computing power determines the runtime cost-effectiveness. With this type of licensing model, licensing cost reduction could be achieved by utilizing a CPU with fewer cores of higher clock speeds at equivalent computing power. However, the lower core-count CPU at a comparative level of computing power comes at a substantially higher processor case temperature, making conventional air-cooling extremely inefficient.

Supermicro, joined by 3M, Intel®, Kaori, and Samsung, demonstrates a proof-of-concept single-phase immersion liquid cooling system with our BigTwin server family of products. Supermicro's partner Kaori developed an immersion-cooling tank filled with 3M™ Fluorinert™ Electronic Liquid FC-40. The test results show that Intel high-frequency chips with a lower core count delivered high LINPACK performance, while Samsung DRAMs and Solid-State Drives achieved uncompromised electrical characteristics. Supermicro's experimental outcome supports the idea that utilizing high clock speed chips with a lower core count to drive for the same level of computing performance could benefit from the savings on software licensing

Enhancing Heat Removal Capability with Immersion Liquid Cooling Systems

The heavy demand for computing power drives the development of newer processor chips with ever-increasing TDP. Compared to older generations, the upcoming processors efficiently run at well over 250 Watts TDP. As a result, this has the potential of collectively exceeding rack-level power capacity when clustered. One way to mitigate this is by spreading out these “hot silos” areas to distribute heat sources better. However, compromising compute density could result in additive burdens on the network infrastructure. The extra cost and heat generated come with the need to rely more on fiber and high-speed networking. Supermicro considers that densely clustered compute nodes would still be the commonly accepted configuration in the data center and for high performance computing (HPC) applications as long as the significant challenges in effectively and economically removing heat can be resolved.

Air as a medium for cooling will generally be less efficient and less effective than liquid. For example, one way to think of this is that conventional air-cooling essentially cools an object immersed in flowing air, utilizing air as an agent to remove the heat. When comparing an object immersed in a flowing liquid, the liquid conjunction introduces more intimate and closer contact than the gaseous conjunction. The higher heat transfer coefficient with liquid than gas manifests the adequate number of molecular-scale particles densely engaging with the object’s surface.

Liquid agents can carry more heat than air does of equal volume and be more thermally conductive than air. However, the main reason conventional cooling systems utilize air as an agent is not its heat transfer capability but its nonconductivity and compatibility with most of the widely used electronic devices in the electronic hardware industry.

With climate change, those used-to-be trivial problems such as airborne particles, gaseous pollution, global warming, and volcanic activities make air-inlet management on data center premises more costly. Moreover, given the ever-surging power density at the rack level and the growing TDP at the chip level, it is foreseeable that one day air-cooling systems may eventually hit the wall that may only be breached by advanced cooling technologies like immersion liquid cooling systems.

TYPICAL LIQUID COOLING SOLUTIONS

Liquid cooling solutions are beneficial in improving cooling efficiency, lowering overall wattage consumption, and mitigating cost. Various field-proven liquid cooling techniques range from dissipating the heat directly from the source to physically dissipating the heat off-premises.

Liquid Engagement at Different Stages

In the heat transfer and exchange process, there are stages where liquid coolants could intervene as a cooling agent:

- Direct-To-Chip Liquid Cooling Solutions - a liquid coolant removes heat directly from high wattage processors. Then the liquid agent may exchange heat with another liquid agent at the rack level or the facility level.
- In-Row or In-Rack Liquid Cooling Solutions - rack-level or row-level liquid cooling systems where air-to-liquid heat exchange occurs at different system integration levels. The rear door heat exchanger (RDHx) is one typical example of a rack-level liquid cooling solution.
- Immersion Liquid Cooling Solutions - the whole or part of a system is immersed in nonconductive inert dielectric liquids. Then these liquids transfer heat while remaining in a liquid state via a liquid-to-liquid heat exchanger or convert to a gaseous state while carrying the latent heat. Then the gaseous agent is transformed back into a fluid via a condenser.

Visit the Supermicro Liquid Cooling Solutions page for more detail.

EDA Performance Boost Beyond Power Efficiency

Immersion liquid cooling systems offer efficiency in operating IT and data center equipment at cooler temperatures, thus lowering total utility costs. Apart from the immediate benefit, additional tips brought tangible cost savings to the data center by saving licensing fees. The software has typically been priced on a per-core or per-socket model at the commercial level. Per-core licensing model becomes dominant as the market demands more compute power per socket or more cores. Even industries that used to stick to the old school per-socket pricing model have started to charge higher licensing fees depending on the number of cores per CPU. As clock speeds increase and higher wattage CPUs become a norm, IT workloads that leverage more cores maximizing software license utilization out of the less per-socket fees, can no longer play the old trick after software companies pulled the plug. Thus, gearing up to the same level of computing power using fewer cores can be translated into a dollar value by reducing the number of runtime licenses required. Low-core high-frequency CPUs will be a rewarding substitute as long as the excessive heat driven by running the chips faster can be resolved. Given the surging wattages in the next generation of CPUs, a conventional approach using air-cooling starts losing its edge.

Supermicro developed a proof-of-concept unit based on the X11 generation of our BigTwin® 2029BT-HNTR 2U 4-Node high-density server system equipped with Intel® Xeon® Gold 6250 processors. The baseline comparison is the same system using the Intel® Xeon® Gold 6244 processors. Both configurations are populated with Samsung M393A2K40DB3-CWE 16GB 3200Mbps RDIMM and PM983 960GB SSD. We immersed our system using the Intel® Xeon® Gold 6250 in a Kaori liquid cooling tank filled with 3M™ Fluorinert™ Electronic Liquid FC-40. The system using the Intel® Xeon® Gold 6244 is cooled by conventional air-cooling. The CPU calculation power performance is benchmarked with LINPACK, while memory and SSD performance are benchmarked with Stream and FIO.

Figure 1 - BigTwin 2029BT-HNTR

The resulting conclusions found that the Intel® Xeon® Gold 6250 in a low temperature chamber generated notably 8% higher LINPACK results, averaging from the four server nodes compared to the Intel® Xeon® Gold 6244 in our air-cooled test environment. Notably, the experiment would have to be performed in an unrealistic low temperature chamber at freeze point to utilize the full speed of the Intel Xeon Gold 6250 processors.

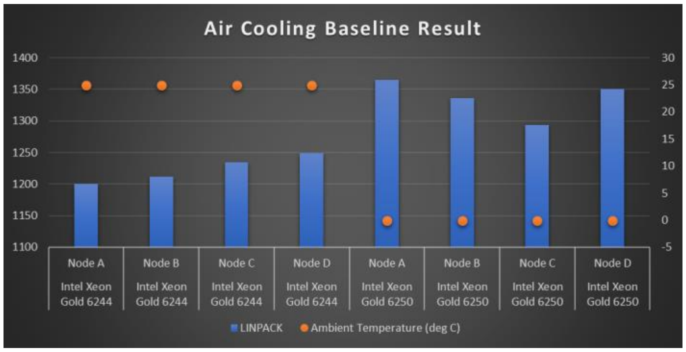

Figure 2 – Air Cooling Baseline Result

In the case of the immersion-cooled Intel® Xeon® Gold 6250, the system at room temperature achieved slightly higher LINPACK results than the previous low temperature chamber results. Here the immersion liquid cooling configuration delivered 12% lower power consumption than the previous low temperature chamber. However, the actual power savings would have achieved approximately 54% when accounting for the excessive power required to bring the temperature to freeze point in the air-cooled configuration.

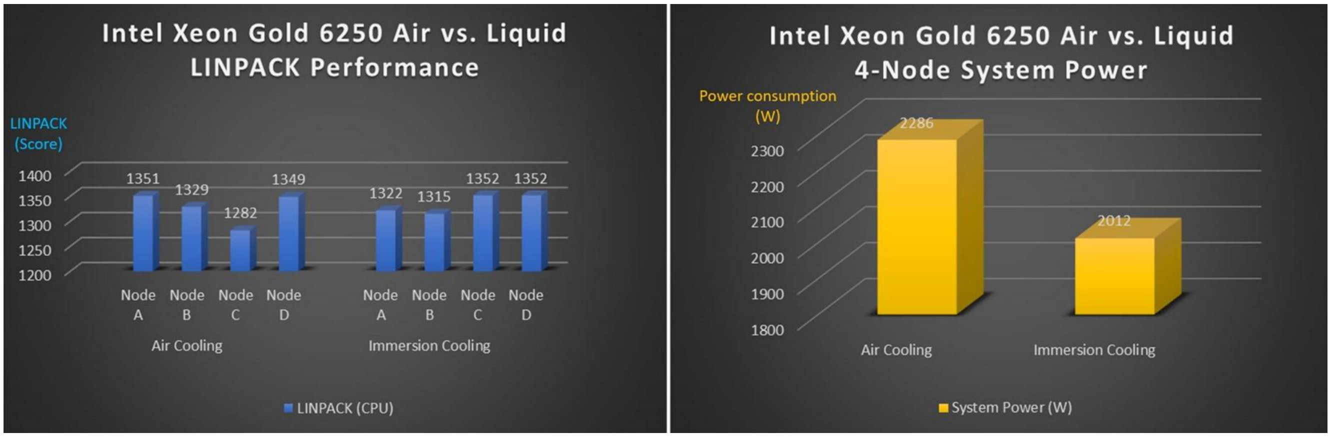

Figure 3 – Intel® Xeon® Gold 6250 Air vs. Liquid Performance Result

The memory performance results from the Stream tool demonstrated a consistent memory bandwidth output as opposed to the baseline configurations. In addition, consistency could also be observed in the NVMe Random and Sequential Read/Write performance based on the FIO tool.

The Intel® Xeon® processors and the Samsung DIMMs and SSDs demonstrated excellent compatibility with the 3M™ Fluorinert™ Electronic Liquid FC-40 without any signs of damage or degradation. In addition, the Kaori immersion liquid cooling tank with the plated heat exchange system provided a solid stand-alone test bench in the lab environment.

Processors with higher frequencies will bring along higher LINPACK results. Utilizing lower core count processors at higher frequencies to achieve similar performance levels as higher core count processors at lower frequencies could bring cost saving benefits in the runtime licensing, a cherry on top of the significant power consumption merits.

Conclusion

Supermicro delivers the highest performance systems to data center operators. New technologies in the electronic design field, AI, and data analytics will require the fastest and most capable CPUs and GPUs for these new workloads. In addition, the higher wattage systems may require liquid cooling for those data center environments pushing the limits of air cooling.

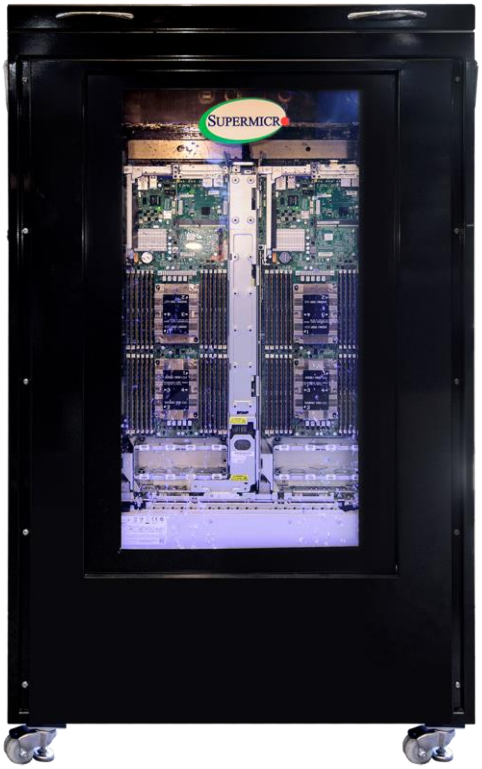

Figure 4 - Supermicro Proof-Of-Concept Immersion Liquid Cooling Demo Unit Jointly Developed with 3M, Intel, Kaori, and Samsung

Supermicro, while teaming up with Intel, Samsung, Kaori, and 3M, was able to demonstrate a unique high-performance CPU with high heat is workable in an actual environment and got effective power consumption reductions when adopting immersion cooling systems, even in some extreme environments without air condition.

Meanwhile, it will reduce server system OPEX by lower total power consumption. Additional benefits include the potential savings in software runtime license fees due to the demonstration that similar performance levels can be achieved with low core-count, high clock speed CPUs. These key points will be conducive to high-performance servers implemented in the future, especially for some fields such as the semiconductor industry, which is already familiar with the characteristics of fluorochemical for a long time.

	Configuration 1	Configuration 2
Processor	Intel® Xeon® Gold 6250	Intel® Xeon® Gold 6244
# of Cores/CPU	8	8
# of Threads/CPU	16	16
Processor Base Frequency	3.90 GHz	3.60 GHz
Max Turbo Frequency	4.50 GHz	4.40 GHz
Processor TDP	185 W	150 W
Tcase	60deg C	74deg C
System	SYS-2029BT-HNTR	SYS-2029BT-HNTR
CPU/System	8	8
Memory/System	16GB 3200 DDR4 RDIMM Samsung M393A2K40DB3-CWE x 96	16GB 3200 DDR4 RDIMM Samsung M393A2K40DB3-CWE x 96
NVMe Drive/System	2.5" Gen3 NVMe 960GB Samsung PM983 x 16	2.5" Gen3 NVMe 960GB Samsung PM983 x 16
SATA SSD/System	2.5" SATA3 240GB Samsung PM883 x 8	2.5" SATA3 240GB Samsung PM883 x 8
SIOM/System	AOC-MGP-i2M x 4	AOC-MGP-i2M x 4

System Configurations

Cooling Method	Processor	CPU1 Temp.	CPU2 Temp.	Ambient Temp.	Node	CPU (Gflops)	Memory (MB/s)	NVMe RR (IOPS)	NVMe RW (IOPS)	NVMe SW (MB/S)	NVMe SR (MB/S)
Air Cooling	**Intel Xeon Gold 6244**	76	83	25	Node A	1200	192767	438000	46000	1260	3274
		78	82	25	Node B	1212	192790	437000	46000	1261	3275
		77	80	25	Node C	1234	191808	436000	46000	1260	3275
		78	78	25	Node D	1249	190817	437000	46000	1260	3274
	Intel XeonGold 6250	74	77	0	Node A	1351	194881	437000	46000	1260	3273
		74	83	0	Node B	1329	194026	436000	46000	1261	3275
		83	84	0	Node C	1282	194746	437000	46000	1259	3274
		77	85	0	Node D	1349	194415	437000	46000	1261	3274
Immersion Liquid Cooling	**Intel XeonGold 6250**	79	80	28	Node A	1322	194805	437000	46000	1260	3275
		85	83	28	Node B	1315	194706	437000	46000	1259	3275
		85	81	28	Node C	1352	194687	437000	46000	1261	3274
		88	89	28	Node D	1352	194880	436000	46000	1260	3274

Table 1 - Memory and NVMe Test Results