SUPERMICRO'S LIQUID COOLED SYSTEMS PROVE PERFORMANCE INCREASE AND LOWER ENERGY USE

Supermicro Demonstrates Performance Gains for High Performance Computing Benchmarks When Using 5th Gen Intel Xeon Processors and Liquid Cooling

Figure 1 – Supermicro Hyper System

TABLE OF CONTENTS

Executive Summary

Users of high-performance computing applications are a primary beneficiary of the advancement of CPU performance. Getting higher throughput with each new generation of CPUs is thrilling, but most often, it comes at the cost of higher thermal design power (TDP), as higher performing CPUs require more power. The increase in the TDP of each CPU typically requires (at rack scale) more data center cooling, adding to the TCO.

Liquid cooling solutions reduce the total data center power usage by up to 40%. The power reduction needed at the

server level decreases, as does the power used for the data center fans and heat removal equipment. With liquid cooling, an additional benefit is that the CPUs can run faster when running compute intensive applications. The result is that the measurable performance can be quantified. A Supermicro Hyper system, SYS-221H-TNR, was equipped with a 5th Gen Intel Xeon 8593Q processor, and several CPU intensive benchmarks were executed with liquid and air-cooling solutions.

System Configuration

The system configuration is pivotal for optimizing the performance, stability, and functionality of a computer system. Proper configuration involves the strategic arrangement and tuning of hardware and software components, ensuring the most optimal performance possible.

System	Cooling	Processor	Memory
SYS-221H-TNR	Air	5th Gen Intel Xeon 8593Q	16x 64GB 5600MT/s
SYS-221H-TNR	Liquid (Cold Plate)	5th Gen Intel Xeon 8593Q	16x 64GB 5600MT/s

The 5th Gen Intel Xeon 8593Q has 64 cores and a base frequency of 2.2GHz. The Max Turbo Frequency is 3.9GHz with 320MB of cache and a top TDP of 385W. Intel Turbo Boost Technology works automatically to give your CPU more speed when needed.

The BIOS CPU Performance setting is set to Extreme Performance Mode, Hyperthreading OFF, and the System Fan Settings are set to Full Speed. The Operating System for these tests is Ubuntu 22.04.03 LTS. At the OS level, performance governors were set to performance.
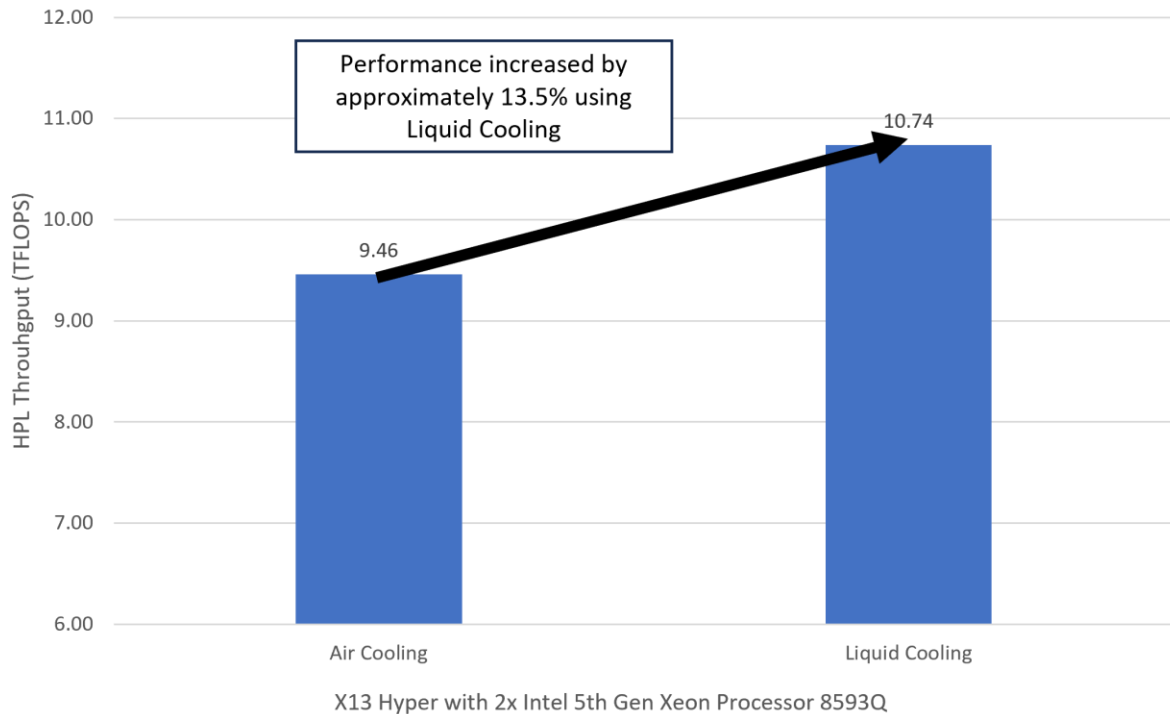
HPL Benchmark Description and Results

Many HPC applications rely on a set of libraries known as Linpack to validate if the system can handle High-Performance Computation. The Linpack Benchmark measures a computer's floating-point rate of execution. It is determined by running a computer program that solves a dense system of linear equations.

The High Performance version of Linpack (HPL), which is based on Linpack, is designed to work on many cores and in distributed environments. Since many applications use similar libraries for a significant portion of their execution time, benchmarking HPL indicates the system's performance under heavy computational workloads. The benchmark's metric is typically measured in FLOP/s (floating-point operations per second), quantitatively measuring a system's raw computational power. Most HPL computational performance primarily depends on the system's central processing units (CPU).

Figure 1 shows the Teraflops (trillions of floating-point operations per second) when running on the Supermicro Hyper system as described above. First, HPL is run on the air-cooled system, and then HPL is run on the liquid cooled system. As can be seen, the liquid cooled system runs HPL 13.5 % faster than the air-cooled system. This means that many HPC type applications will run faster on a liquid cooled system than on an air-cooled system, or larger models can be run simultaneously. In addition, when an entire data center is liquid cooled, up to 40% power savings is possible.

X13 Cooling Solution HPL Benchmark

Performance increased by approximately 13.5% using Liquid Cooling

9.46

10.74

HPL Throuhgput (TFLOPS)

Air Cooling

Liquid Cooling

X13 Hyper with 2x Intel 5th Gen Xeon Processor 8593Q

HPL For Accelerator Introspection Benchmark Description and Results

The HPL-AI benchmark, which is relatively new, stands for "the High Performance LINPACK for Accelerator Introspection." The benchmark is used with mixed-precision arithmetic for a linear system of equations. As AI (low precision) and HPC (high precision) converge, this benchmark allows systems to be compared using just the CPU. The result is Floating Point Operations Per second (FLOPS).

February, 2024

X13 Cooling Solution HPL-AI Benchmark
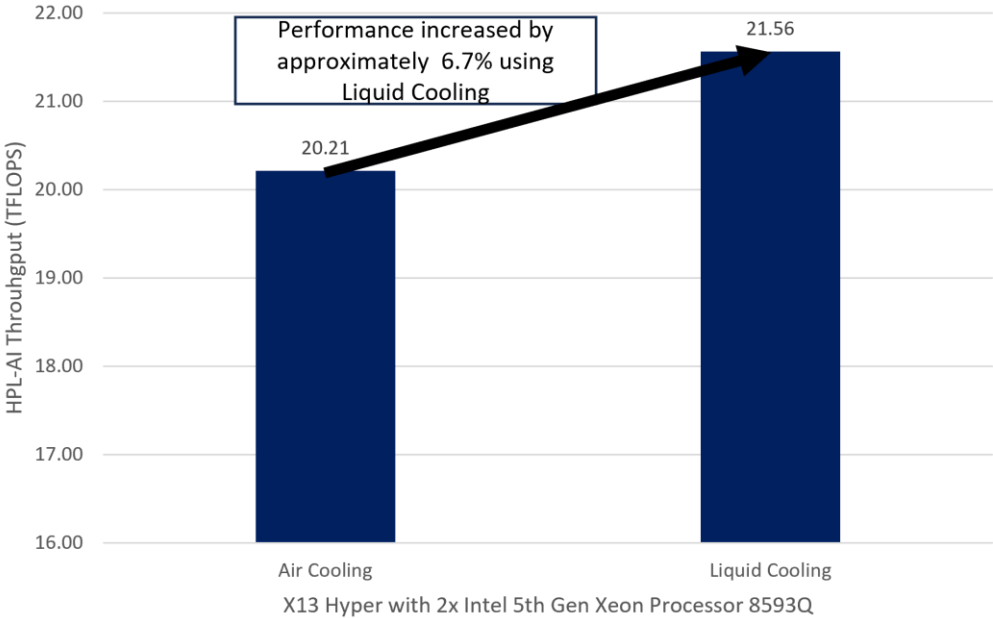
The performance on this benchmark improves by 6.67 % when using the system with liquid cooling.

Conclusion

The HPL and HPL-AI benchmarks significantly improve performance when using Supermicro's liquid-cooling thermal solution on the X13 Supermicro Hyper system with 5th Gen Intel Xeon processor (8593Q). Liquid cooling allows the CPU to run faster with no throttling while using extreme performance mode compared to an air-cooling solution. The results demonstrate the performance gain when using liquid cooling while running CPU-intensive applications due to the ability of the CPU to run at a faster clock rate, translating into additional TCO reduction when liquid cooling is deployed at rack scale.

For More Information:

Supermicro Liquid Cooling Website – www.supermicro.com/liquid-cooling

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs and enables us to build and deliver application-optimized solutions based upon your requirements. See www.supermicro.com

February, 2024