SOLUTION BRIEF

SUPERMICRO X14 HYPER SYSTEMS WITH SINGLE SOCKET DEMONSTRATE 4X PERFORMANCE AND EFFICIENCY COMPARED TO PREVIOUS GENERATION OF DUAL SOCKET SYSTEMS

Single Socket Systems with the Intel® Xeon® 6 Series Processor Show Dramatic Increase in Performance Compared to Earlier Generations of Dual Socket Systems

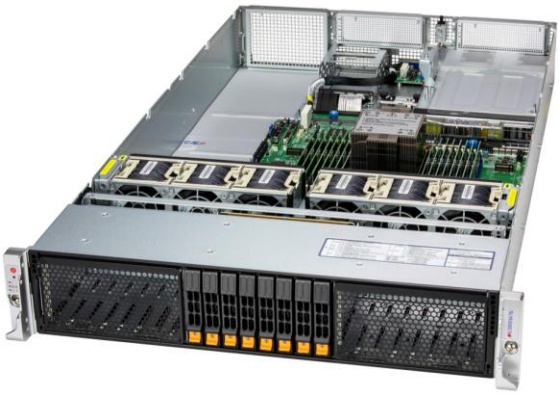

Supermicro X14 Hyper System – Single-Socket

TABLE OF CONTENTS

Executive Summary

Content Delivery Networks (CDN) enhance user experience and improve network efficiency when strategically placed at the edge. Well-designed standard rack servers can be distributed geographically, reducing the physical distance from the server to the end-user device, leading to a more responsive experience and reduced backhaul traffic. In addition, by caching frequently accessed content on high-

speed storage and in memory, a more significant portion of the content can be available at the edge. The new Supermicro X14 Hyper SuperServer SYS-212H-TN has just a single socket and can make an excellent CDN server in terms of storage and network fanout, performance, and energy efficiency. These traits also allow the Supermicro Hyper servers to be great software-defined storage nodes. The performance and efficiency gains of the new Supermicro Hyper servers are impressive compared to earlier generations of Supermicro servers with older Intel Xeon processors.

System Details

Supermicro X14 Hyper UP (single socket) systems, powered by the groundbreaking Intel Xeon 6 with P-cores as a 1-socket solution, deliver unmatched performance, flexibility, and scalability to meet the demands of modern computing workloads. The system can scale the memory capacity to meet the most intensive tasks, up to 2TB of 6400MT/s ECC DDR5 RDIMM. It also offers a wide range of storage configurations from 8x up to 24x SATA/SAS/NVMe SSDs, providing exceptional storage capacity and efficiency. With PCIe Gen 5 x16 support, it can deliver up to 6 x 400Gbps I/O for different bandwidth requirements, and customers can enjoy the unprecedented capability, ensuring seamless data throughput and optimal performance for high-bandwidth applications. Whether for miscellaneous enterprise workloads, the Supermicro X14 Hyper UP sets a new benchmark in server technology.

These tests use the Supermicro Hyper server, the SYS-212H-TN, containing the Intel® Xeon® 6741P for these tests and comparisons. By comparing the performance and performance per watt running Varnish Enterprise software with similarly configured X11 generation of Supermicro servers equipped with the 2nd Gen Intel Xeon processors, it is easy to see the benefits of the new systems and the new Intel architecture. In the X11 generation, a DP configuration was necessary to connect enough NVMe drives and network cards. With X14 UP, the architecture can be simplified while still supplying enough compute, memory, and PCIe capability for the use case.

The performance of the Supermicro X14 configuration was compared to that of a Supermicro X11 server containing dual CPUs, the 2nd Gen Intel Xeon 6252N processor. The Supermicro server used for this benchmark is the Supermicro SuperServer SYS-1029U-TN10RT (dual processor) and the Supermicro Hyper server, the SYS-212H-TN. A summary of the systems used is below in Table 1.

Generation		
System Type (# Processor)	Dual CPUs	Single CPU
System	SYS-1029U-TN10RT	SYS-212H-TN
CPU	2x 2nd Gen Intel Xeon 6265N (48 cores total)	1x Intel Xeon 6741P (48 Cores)
Memory	384GB DDR4-2933	512GB DDR5-6400
Storage	10x Gen3x4 NVMe	16x Gen5x4 NVMe
Network	200Gbps (2x Gen3x16 100GbE)	800Gbps Line Rate (2x Gen5x16 2x200GbE)

Table 1 - Summary of Supermicro Systems Used for Tests

Software Details

Varnish Enterprise 6.0.13r10 was used for the content delivery software in these tests. Varnish Enterprise is a flexible, efficient, and high-performance HTTP(S) cache that can be deployed on standard x86 servers. Varnish Enterprise is built on top of the open-source Varnish Cache and is enhanced with additional features. This NUMA-aware software uses an efficient in-core TLS implementation using standard OpenSSL APIs and Linux async I/O for disk access. Additionally, Varnish's behavior can be heavily customized through Varnish Configuration Language (VCL). In the interest of software consistency, all systems tested were running RHEL 8 with kernel 6.6.65 with identical kernel tunings.

Test Environment and Methodology

For the testing discussed in this paper, the systems under test were each connected to a non-blocking network switch appropriate for the speed of their networking devices (100GbE for X11, 400GbE for X14), along with an additional system to act as a content origin and several systems to act as clients, where each client has a 100GbE NIC and would run the *wrk* benchmark. *Wrk* is open source and is a modern HTTP benchmarking tool capable of generating significant load when run on a single multi-core CPU. It combines a multithreaded design with scalable event notification systems such as *epoll* and *kqueue*. The number of clients in use was determined by the network line rate of the system under test; the X11 system is limited to at most 200Gbps, so used two clients; the single-socket X14 is limited to at most 800Gbps, so eight clients were used. A diagram of the test environment can be seen below.

The multiple client systems were coordinated using parallel-ssh, and during the tests, the network bandwidth of the systems under test was measured with *dstat*. For these tests, all communication between the client and the system under test is HTTPS traffic. When power measurements are reported, the values are measured at the power distribution unit and are the total consumption of the server, including all CPUs, memory, networking, storage, etc.

In addition, this testing is intended to cover two important use cases for CDNs.

- **Video-on-demand** (VoD) is a workload where the content is pre-recorded and provided to the caching layer for later distribution to end users. Many examples of this type of video service offer access to a library of pre-recorded TV programs and movies. This type of workload will use the system's NVMe storage and DRAM for caching. Typically, the DRAM is used for transient buffers and the most popular subset of content.

- **Live video** or Live-Linear is a workload where the content is distributed to the caching layer for immediate distribution to end users. This workload can be in the form of TV programming streamed out linearly to players or a service streaming a popular live sporting event.

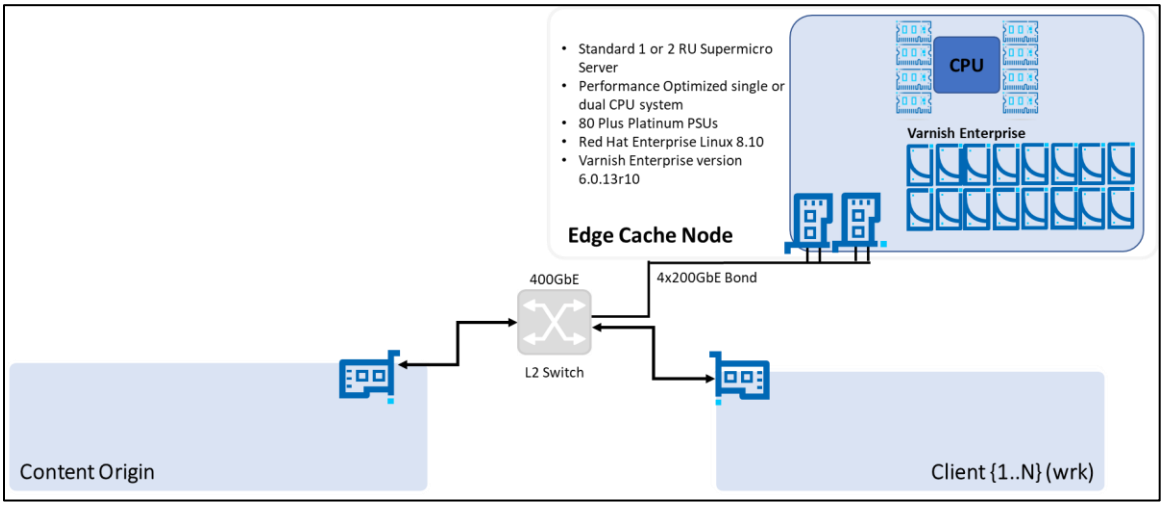

Figure 1 - Benchmark Setup

Benchmarks Results

Varnish Enterprise was run on the systems and placed under load with wrk, and the raw performance and performance per watt were measured. The X14 systems offer substantially higher performance and efficiency than the prior-generation servers. These new systems allow a CDN operator to deliver more traffic to end users with less energy per bit of throughput.

Live-Linear:

For the Live-Linear test, the single-socket Supermicro X14 server with an Intel Xeon 6741P processor showed up to a 4.03X performance improvement over the dual-socket X11 system with two 2nd Gen Intel Xeon processors when looking at HTTPS network throughput.

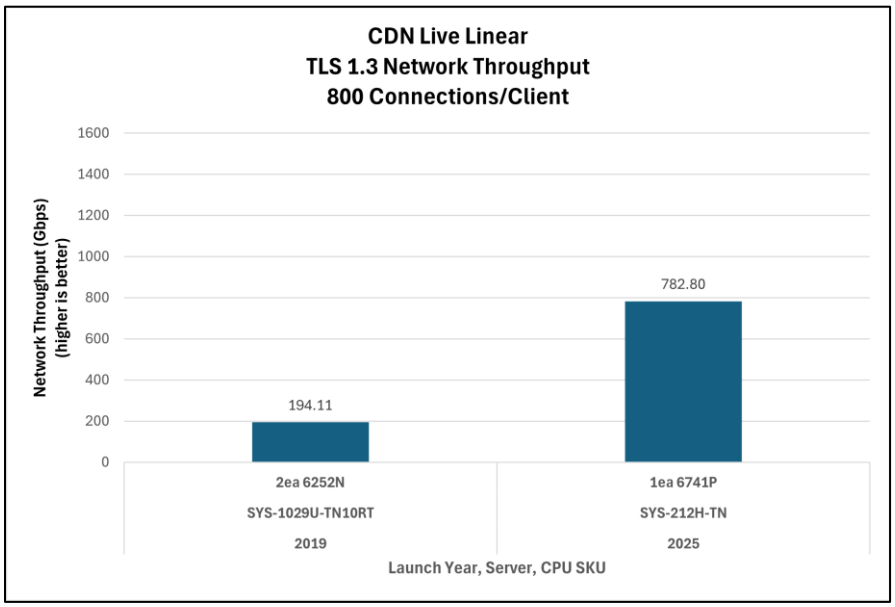

When looking at the network throughput per watt, the Supermicro X14 Hyper server with a single Intel Xeon 6741P processor showed up to a 2.77X advantage over a system with dual 2nd Gen Intel Xeon processors.

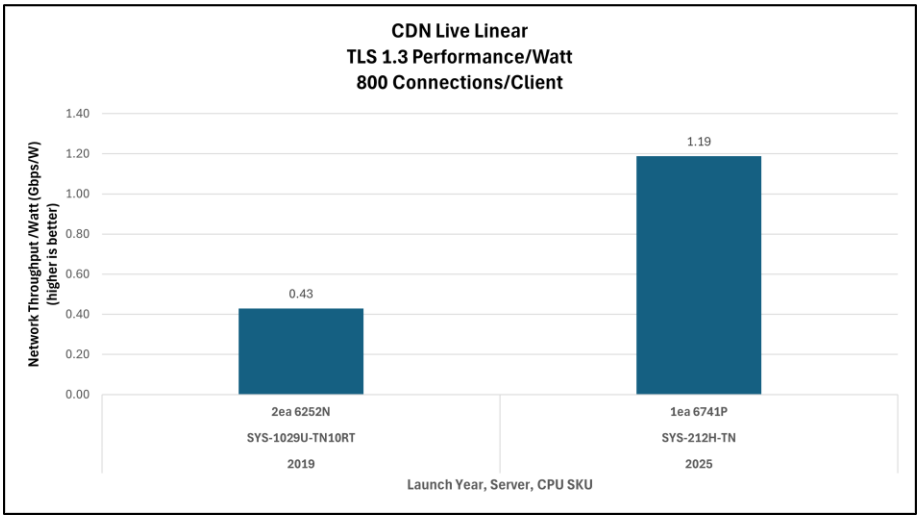

Video-On-Demand:

Another measure of performance in a CDN environment is Video-On-Demand performance. The benefits become apparent by comparing the network throughput of the single-socket Supermicro X14 Hyper UP system with that of an older-generation system. The Supermicro system with the Intel Xeon 6741P processor shows a 3.25X performance improvement over the older generation system with dual 2nd-generation Intel Xeon processors.

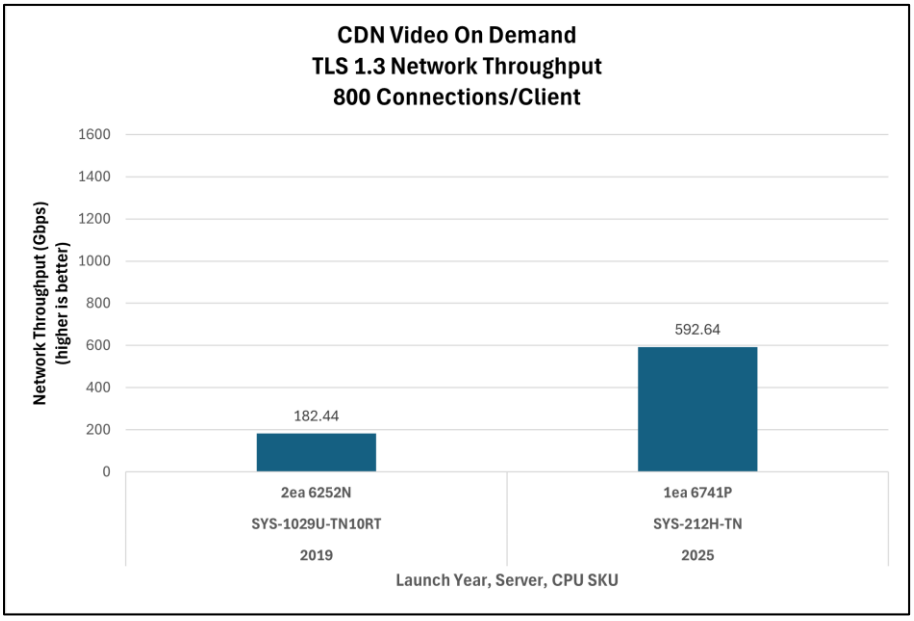

When measuring power consumption for throughput tests, the Intel Xeon 6 6741P processor system showed over a 2X gain, meaning that less power is used for each gigabit of data transmitted.

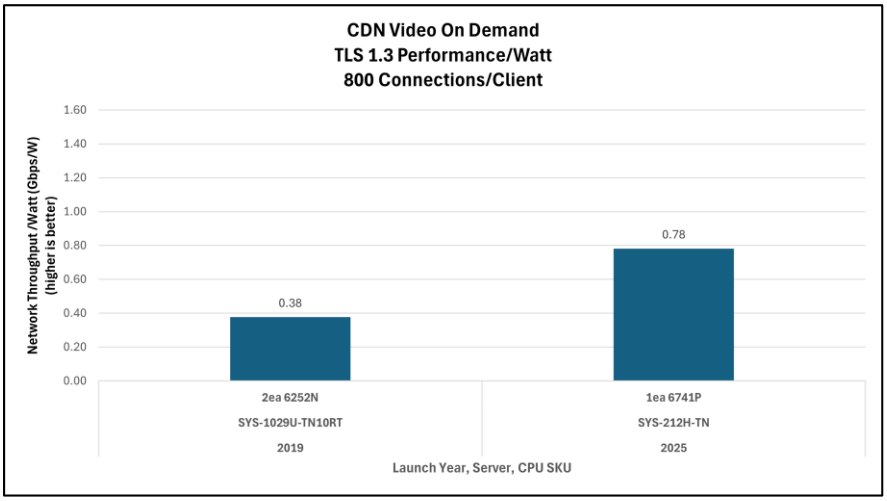

Summary

These results show a significant improvement in the performance of systems based on the Intel Xeon 6 processor with P-cores over systems from about 6 years prior. With these improvements, an operator can choose to shrink the physical size of a PoP, serve additional users, or deliver higher bitrate video content while improving energy efficiency.

These performance and efficiency improvements result from a combination of enhancements in the new Supermicro SuperServers and the processors installed in them. The X14 single-socket platform with Intel Xeon 6 with P-cores processor outperforms the dual-socket platform and offers the flexibility to support up to 8 channels of DDR5-6400 (up to 1.3x more channels at up to 2.2x higher speed per channel), up to 136 lanes of PCIe Gen5 per processor (up to 2.83x more PCIe lanes at 4x higher speed per lane when compared to 2nd Generation Intel Xeon Scalable Processors), and are available in a variety of configurations tailored to critical use cases such as content delivery networks, edge computing, and software-defined storage.

Content Delivery Networks will excel with the Supermicro X14 Hyper, using the new Intel Xeon 6 6500P/6700P Series, especially the Intel Xeon 6 with P-cores as a 1-socket solution. The configurations tested offer higher performance per watt, reducing energy consumption, lowering costs, and improving the user experience compared to similar 2nd Generation Xeon Scalable configurations. Organizations that use the Supermicro SYS-212H-TN with efficient software, such as Varnish Enterprise, can see improved performance and reduced costs due to lower power demand per bit for each server compared to prior generations.

Further Information

https://www.supermicro.com/en/products/system/hyper/2u/sys-212h-tn

https://www.varnish-software.com/

https://www.intel.com/content/www/us/en/products/details/processors/xeon.html

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.

For more information: www.supermicro.com

INTEL

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, visit www.intel.com

VARNISH SOFTWARE

Varnish Software is at the forefront of high-performance data, web and content delivery solutions, offering Varnish Enterprise as its flagship product. Varnish Enterprise is designed to meet the complex needs of large-scale, mission-critical deployments, providing unparalleled speed, scalability, and reliability for digital content delivery.

For more information: www.varnish-software.com

Notices and Disclaimers

Test by Intel as of 01/22/2025.

Single Intel® Xeon® 6741P: 1-node, 1x Intel(R) Xeon(R) 6741P, 48 cores, 300W TDP, HT On, Turbo On, Total Memory 512GB (8x64GB DDR5 6400 MT/s [6400 MT/s]), BIOS 1.1, microcode 0x1a000080, 2x MCX755106AS-HEAT, 16x 1.8T KIOXIA KCMY1RUG1T92, 1x 894.3G Micron_7450_MTFDKBG960TFR, Red Hat Enterprise Linux 8.10 (Ootpa), Kernel 6.6.65.

Dual Intel® Xeon® Scalable 6252N: 1-node, 2x Intel(R) Xeon(R) Gold 6252N CPU @ 2.30GHz, 24 cores, 150W TDP, HT On, Turbo On, Total Memory 384GB (12x32GB DDR4 3200 MT/s [2933 MT/s]), BIOS 4.2, microcode 0x5003605, 2x MCX516A-CDAT, 10x 3.7T INTEL SSDPE2KX040T8, 1x 477G INTEL SSDPEKNW512G8, Red Hat Enterprise Linux 8.10 (Ootpa), Kernel 6.6.65.

Software: Workload: wrk master 02/07/2021, varnish-plus-6.0.13r10 revision c79290209cb5a2825479d95841a67fddad8dd9c2, OpenSSL 1.1.1k FIPS 25 Mar 2021, keep alive, 800 connections per client, throughput measured with 100% Transport Layer Security (TLS) traffic with 98.8%(Live Linear) and 100% (Video-On-Demand) target cache hit ratio.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.