White Paper

Communication service providers (CoSPs), Telco, Network operators
Network architects, Engineers, CxO

Tests Show Energy Management Impact of Intel-based Varnish CDS

Content delivery (CD) services are growing with increased market demand for data rich content; energy management testing of Varnish Enterprise CD software on Intel servers show excellent energy efficiency

Meeting the demand for video, data rich content, AR/VR content and live streaming is driving up the demand for content delivery (CD) software. This, in turn, is increasing demand for electricity to power CD servers.

Intel and Varnish have been working together to reduce energy consumption in CD servers using Varnish Enterprise CD software running on Intel® architecture servers.

For a CD server, energy efficiency is typically measured in the amount of content delivered per unit of energy consumed, (for example, bits per joule or bits per kWh) or in terms of throughput per unit of power (bits per second per watt).

In a test paper published in Sept. 2023[1], Varnish Software, an Intel® Network Builder ecosystem member, demonstrated the energy savings of its optimized Varnish Enterprise CD Software. These changes, combined with power saving features in Intel® architecture processors help to maximize energy efficiency.

Laying Groundwork with Energy Efficiency Tests

In the original tests (completed in Sept. 2023), the two companies tested the scalability of both edge node performance and energy efficiency across four different CD performance levels. The headline was the achievement of 1.2 Tbps of data throughput with energy consumption of only 1.18 Gbps/Watt[2].

These results were impressive, but there is no industry benchmark for comparison. In this paper, the tests are put into context of cost savings and greenhouse gas reduction to show how increased energy efficiency delivers significant impact when compared to legacy infrastructure over a representative hardware lifecycle. That context will involve a series of assumptions on useful life, cost of energy, energy generation mix, etc., but our intent is to show the process so that the reader can substitute in values that are appropriate to their locality or existing equipment.

Varnish Enterprise Edge Content Delivery Software

The testing described in this paper used the Varnish Enterprise Content Delivery software (see Figure 1) which adapts the web cache and HTTP(s) features of Varnish Enterprise for deployment at the network edge. The benefits of deploying content delivery at the network edge include significantly reduced latency and reduced backhaul network traffic to origin or mid-tier servers.

Varnish Enterprise is a CD software solution optimized for deployment on multi-access edge computing (MEC) servers or from a telco cloud data center, network aggregation site or wireless base stations. Varnish Enterprise can be deployed in virtualized, containerized, or bare metal environments. This deployment flexibility makes it an ideal solution for providing CD services over 5G networks.

Table of Contents

Laying Groundwork with Energy Efficiency Tests........... 1
Varnish Enterprise Edge Content Delivery Software 1
Test Setup 2
August 2023 Test Results 3
Energy Consumption Modeling...4
Conclusion.........................6

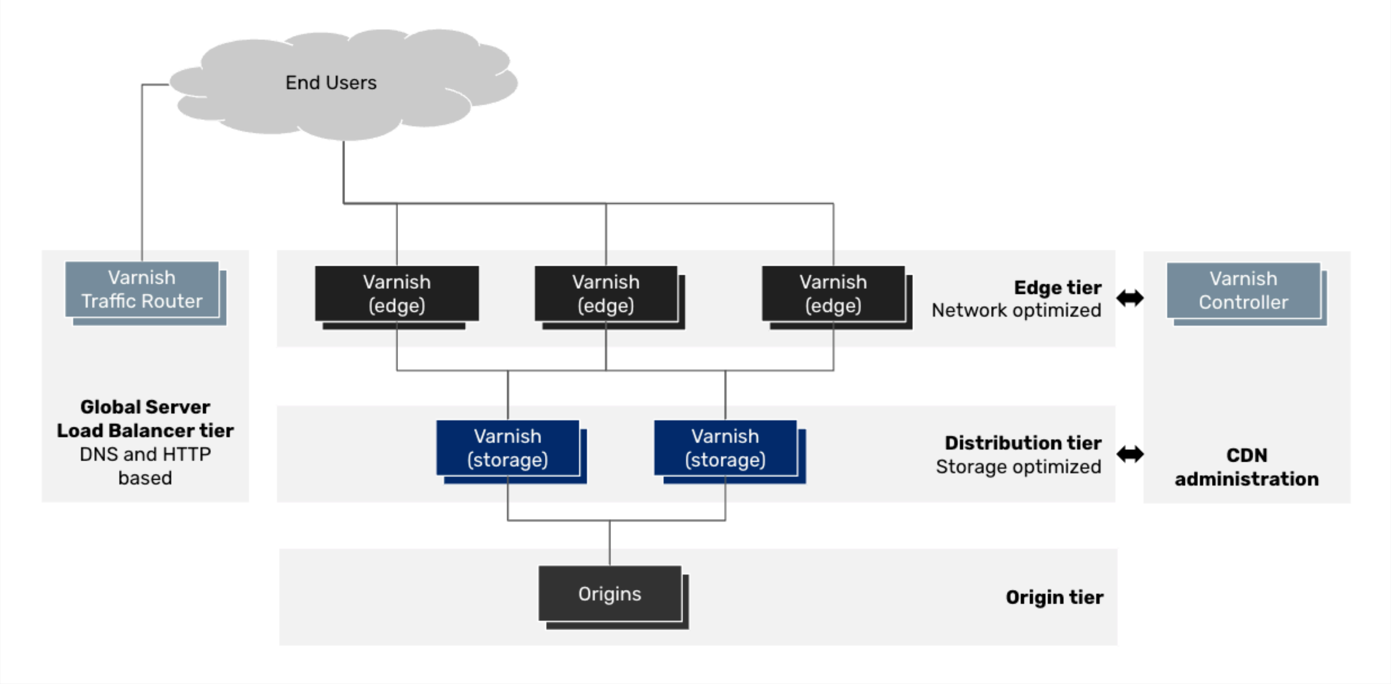

Figure 1. Varnish Enterprise deployment block diagram.

Varnish Enterprise is built on top of open source CD software frameworks with enterprise resiliency features and is enhanced with robust features for high performance and scalability. It allows companies to deliver low-latency, high bitrate content even during periods of peak demand. Varnish Enterprise provides capabilities for web and API acceleration, streaming, and private CD server deployments.

The software is also optimized for non-uniform memory access (NUMA) architectures. Historically, NUMA has been associated with multi-processor systems, but in modern processor architectures, both single and dual processor systems may use NUMA, allowing the processor(s) and operating system to be aware of the relative distance between a CPU and memory or I/O. For these tests, all Intel® Xeon® Scalable CPU-based systems are leveraging sub-NUMA clustering (SNC) to split the CPU into two or four NUMA regions. This is beneficial as it allows the operating system to attempt to localize work and helps to prevent the system from becoming bottlenecked by Linux memory management, which currently scales with NUMA zones. By enabling SNC, it's possible to reduce the impact of that Linux behavior, enabling higher total throughput.

Varnish has also incorporated an in-core TLS implementation, built upon OpenSSL. This TLS implementation makes efficient use of standard OpenSSL APIs, and, when combined with the NUMA awareness features, VCL and threading model, results in more efficient caching software. When combined with the tight integration of Intel® Xeon® D processors or Intel® Xeon® Scalable processors with larger caches, greater memory bandwidth and greater PCIe bandwidth than the previous generations; and a well-designed server, an end user can observe excellent performance and energy efficiency.

Test Setup

In September 2023, Intel and Varnish Software teamed up to test five server configurations to measure network throughput and power consumption and to quantify how improvements in these areas could reduce both energy costs and greenhouse gas emissions.

The configurations selected for these tests demonstrate excellent performance and efficiency over a wide range of performance levels and power budgets, and all are commercially available. These configurations are only a subset of what is possible with Intel® Xeon® D and 4th Gen Intel® Xeon® Scalable based servers, with many more options available to tailor the solution to a particular deployment.

Network Line Rate	Processor(s)	Memory	Storage	Server
Up to 50Gbps (2x25GbE)	1 each Intel® Xeon® D-2733NT	128GB (4x32GB DDR4 3200 @ 2666MT/s)	2 each Intel® P5510 Gen4x4 3.84TB	1RU Supermicro SYS-110D-8C-FRAN8TP
Up to 400Gbps (1x 2x200GbE)	1 each Intel® Xeon® Gold 5418N	512GB (8x64GB DDR5 4800 @ 4000 MT/s)	10 each Samsung PM1743 15.36TB	1RU Supermicro SYS-111C-NR
Up to 400Gbps (1x 2x200GbE)	1 each Intel® Xeon® Gold 6428N	512GB (8x64GB DDR5 4800 @ 4000 MT/s)	10 each Samsung PM1743 15.36TB	1RU Supermicro SYS-111C-NR
Up to 800Gbps (2x 2x200GbE)	2 each Intel® Xeon® Gold 6438N	512GB (16x32GB DDR5 4800)	12 each Samsung PM1743 15.36TB	2RU Supermicro 221H-TNR
Up to 1.6Tbps (4x 2x200GbE)	2 each Intel® Xeon® Platinum 8480+	512GB (16x32GB DDR5 4800)	12 each Samsung PM1743 15.36TB	2RU Supermicro 221H-TNR

Table 1. Test configurations listed by network data rate.

In each of the system under test (SUT) configurations, all memory channels were populated, the optimal number of NVMe devices were installed, and networking cards commensurate with the line rate were installed, except in the Intel® Xeon® D-based system, which used its integrated networking. For additional detail on any of the system configurations, please refer to the system configuration footnote.

In all tests, network throughput was measured at the SUT with dstat and power consumption was measured at the PSU inlets with a Raritan iPDU, and includes all power consumed by the server, from the processors to the cooling fans. The systems were subjected to analogs of video on demand and live-linear workloads using WRK, an open source HTTP(S) benchmarking tool. In the live-linear tests, the NVMe drives were installed (and thus consuming power), but in an idle state as data was cached only in memory.

August 2023 Test Results

Figure 2 shows the relationship between network throughput and total system power consumption under load. In this chart, we can see that some Intel® Xeon® Scalable-based configurations are more efficient than others, but that with these systems, energy efficiency is at or above 1Gbps / 1 Watt to a power footprint as low as 336 watts under load. Some energy consumers within a server do not change as much with the performance capability of the server and cannot be amortized across as much work. As a result of those costs, achieving 1Gbps/ 1Watt was not possible with the smallest configuration, based upon Intel® Xeon® D, but that system was able to demonstrate excellent efficiency within a small power envelope, 153 watts under load.

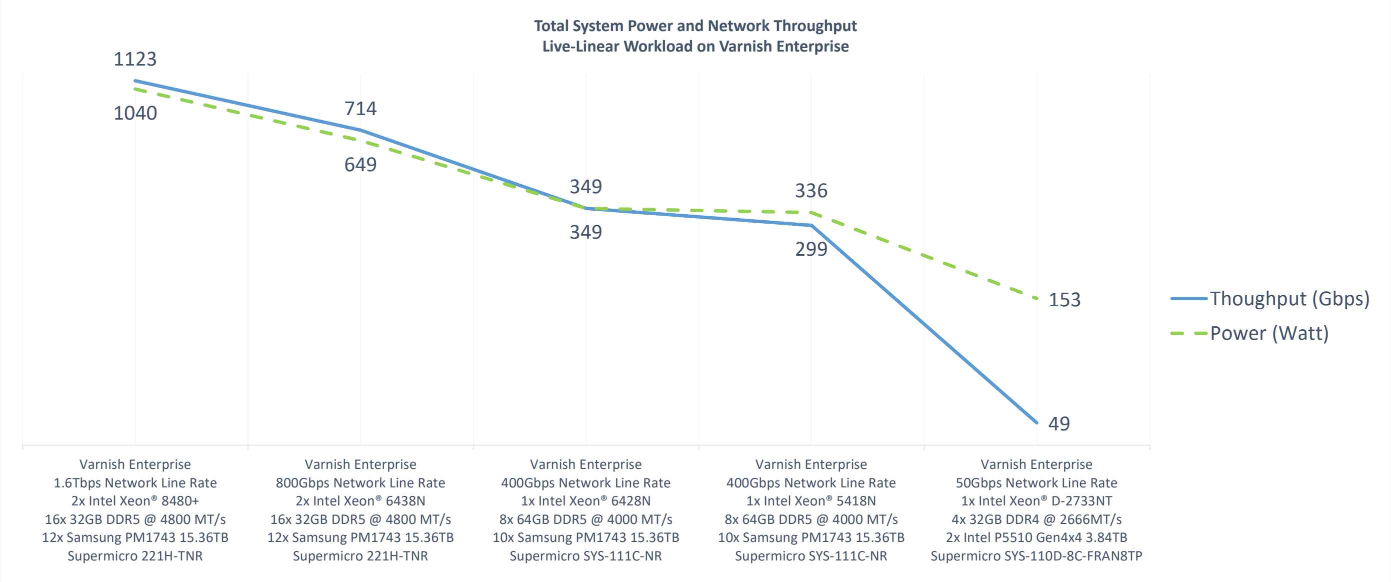

Figure 2. Throughput and power.

Energy Consumption Modeling

Figure 3 shows the estimated total energy consumed per year is based on a load assumption that the system will be at typical load and power 68% of the time, low load and 66% of typical power 16% of the time, high load and 125% of typical power 16% of the time.

This energy consumption estimate includes the TOR switch necessary to accommodate the network requirements. The equation used for analysis is:

Estimated Annual Energy Consumption = 365 * 24 * ((.16*.66*typical_power) + (.68*1*typical_power)+(.16*1.25*typical_power))

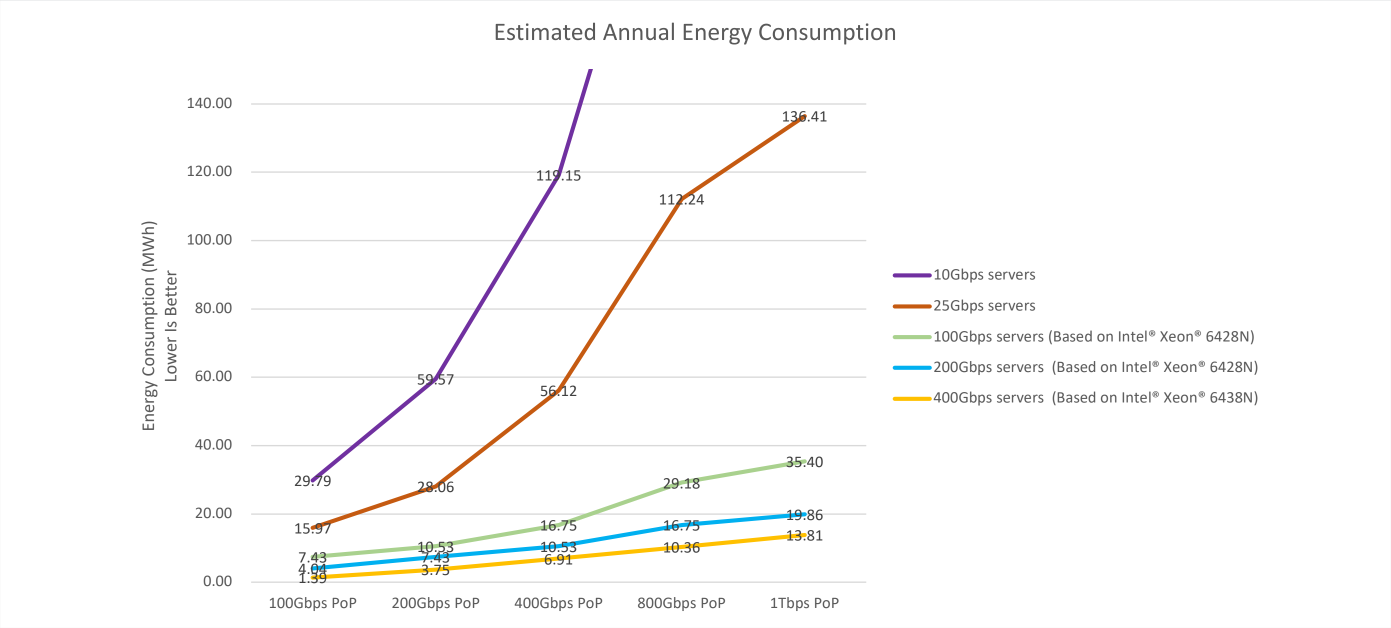

Figure 3. Graph of estimated annual energy consumption.

Estimated Annual Power Usage (MWh)

	100Gbps PoP	200Gbps PoP	400Gbps PoP	800Gbps PoP	1Tbps PoP
10Gbps servers	29.36	55.26	107.06	214.12	269.38
25Gbps servers	17.27	31.08	58.71	117.42	148.50
100Gbps servers (Based on Intel® Xeon® 6428N)	6.48	9.50	15.54	27.63	33.67
200Gbps servers (Based on Intel® Xeon® 6428N)	6.56	6.56	9.67	15.89	18.99
400Gbps servers (Based on Intel® Xeon® 6438N)	12.78	12.78	12.78	16.06	19.34

Table 2. Estimated annual energy consumption.

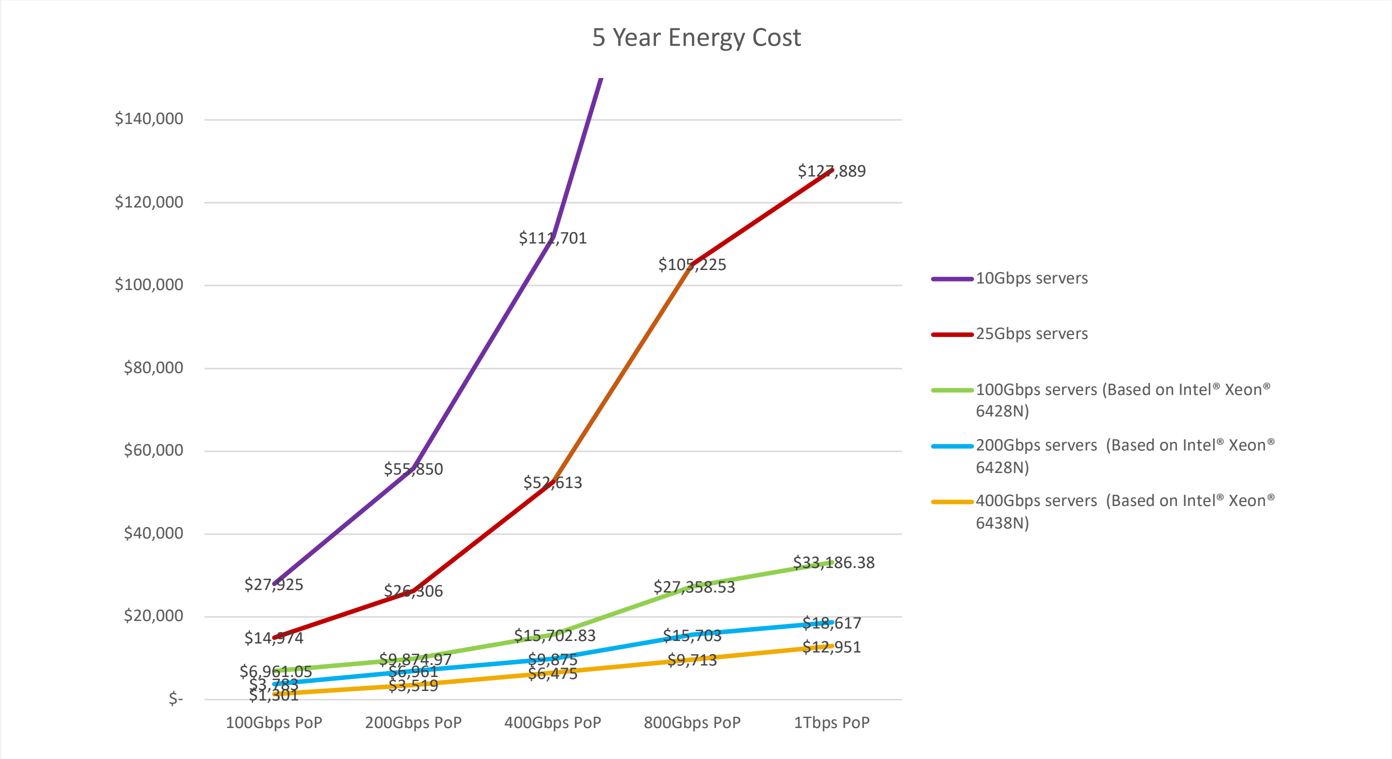

Figure 4. Estimated 5 year energy cost.

In these estimates, we are accounting for the minimal amount of servers necessary to meet a throughput requirement, and depending upon the deployment there may be additional considerations such as a minimum cache capacity.

For this study, we assume a five-year service life, an energy cost of $.0125/kWh and a power usage effectiveness (PUE) of 1.5. Given these assumptions, the total energy consumed and total cost over that service life was extrapolated in Figure 4 below. The equation used is:

5-year energy consumption = 5 *estimated annual energy consumption * PUE

5-year energy cost = 5-year energy consumption * energy cost ($.0125/kWh or $12.5/mWh)

Another way to look at this is the carbon emissions involved in generating the electricity to power the equipment. For this, we assume the PoPs will be in the US, and are using the 2022 US average mix of energy sources and their emissions[3].

Estimated 5-year Carbon Footprint (tCO2e)

	100Gbps PoP	200Gbps PoP	400Gbps PoP	800Gbps PoP	1Tbps PoP
10Gbps servers	26.65	50.16	97.18	194.35	244.51
25Gbps servers	15.67	28.21	53.29	106.58	134.79
100Gbps servers (Based on Intel® Xeon® 6428N)	5.88	8.62	14.11	25.08	30.56
200Gbps servers (Based on Intel® Xeon® 6428N)	5.96	5.96	8.78	14.42	17.24
400Gbps servers (Based on Intel® Xeon® 6438N)	11.60	11.60	11.60	14.58	17.55

Table 3. Estimated 5 year carbon footprint.

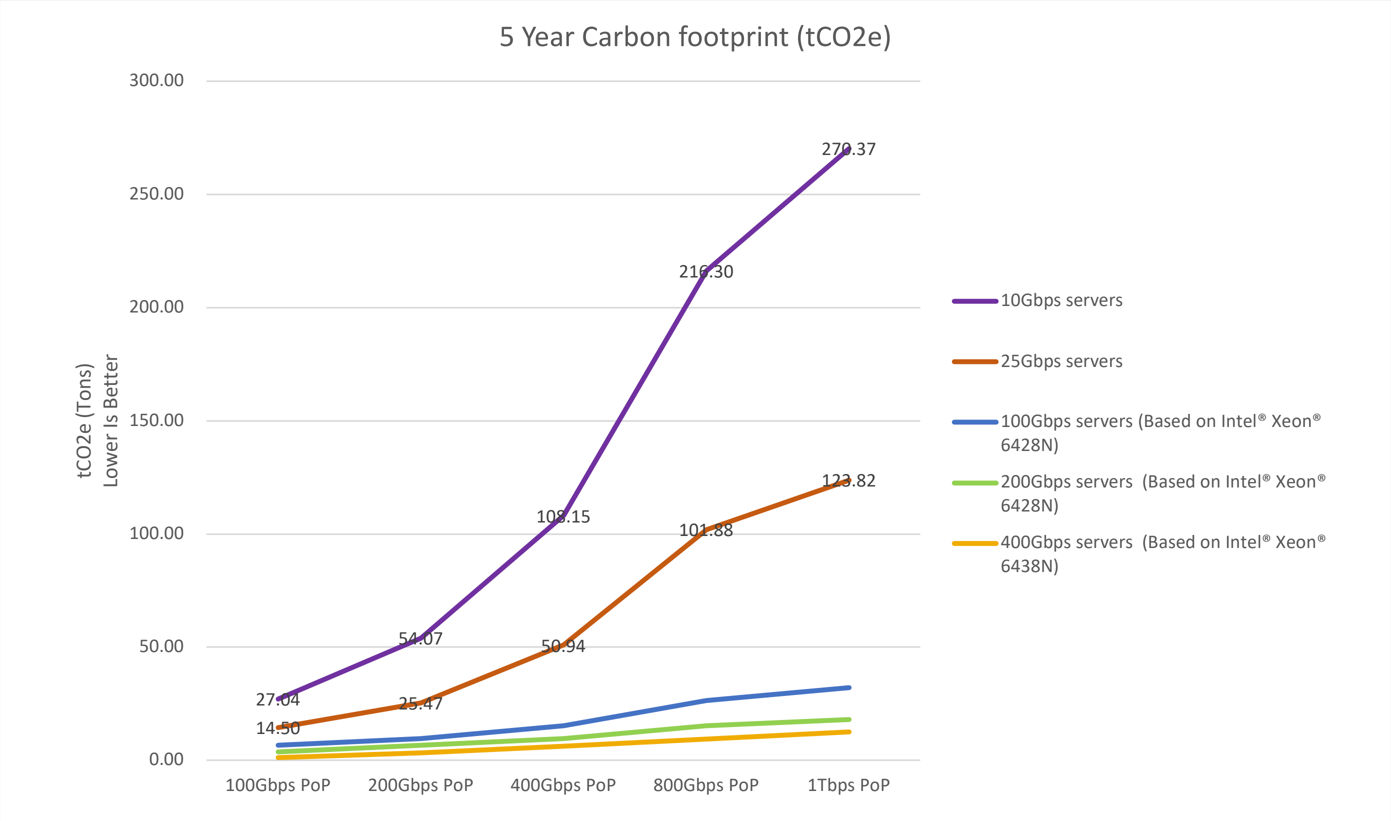

Figure 5. Graph of estimated 5 year carbon footprint.

From that energy mix, we assume that gas is used for 33% of generation, at 182.54 kgCO2e/MWh; oil for 36% of generation, at 246.77 kgCO2e/MWh; coal for 10% of generation at 324.63kgCO2e/MWh; and either nuclear or renewables for 21% of generation, at 0kgCO2e/MWh; for an overall USA energy generation carbon emissions estimate of 181.54kgCO2e/MWh.

Conclusion

As we can see in this study, there is an opportunity to reduce energy consumption by replacing older equipment with a well selected new system running Varnish Enterprise across a wide range of performance requirements and power footprints. In the case of a 400Gbps PoP that is nearing the end of its five-year life cycle, replacing a solution capable of 25Gbps per server with one capable of 200Gbps per server can save an estimated 45.59 MWh per year; and replacing servers capable of 40Gbps each with the same 200Gbps system can save an estimated 31.77 MWh per year. If those systems were to be left in service for another five years instead of being replaced, the increased energy cost alone would be $42,700 for the 25Gbps system and $29,000 for the 40Gbps capable system. From a greenhouse gas, perspective, energy savings in the upgrade would reduce the carbon footprint by an estimated 41.38 and 28.84 tCO2e respectively.

Learn More

Varnish home page

Varnish Enterprise

Intel® Network Builders

Varnish Enterprise Shows Up To 1.2 Tbps Data Rate, Up To 1.18 Gbps/Watt Efficiency

---

[1] https://www.intel.com/content/www/us/en/content-details/788717/varnish-enterprises-shows-up-to-1-2-tbps-data-rate-up-to-1-18-gbps-watt-efficiency.html

[2] https://www.intel.com/content/www/us/en/content-details/788717/varnish-enterprises-shows-up-to-1-2-tbps-data-rate-up-to-1-18-gbps-watt-efficiency.html?wapkw=varnish

[3] https://ctprodstorageaccountp.blob.core.windows.net/prod-drupal-files/documents/resource/public/Conversion_factor_introductory_guide.pdf

Notices & Disclaimers

Single Intel® Xeon® D-2733NT: 1-node, 1x Intel® Xeon® D-2733NT, 8 cores, HT on, Turbo ON, Total Memory 128GB (4 slots / 32GB / 3200 MT/s @ 2666 MT/s), BIOS 1.2a, microcode 0x1000211, 2x Intel Ethernet (integrated), RHEL 8.8, 2x Intel P5510 SSDPF2KX038TZO 3.84TB, kernel 4.18.0-477.15.1.el8_8.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-18), OpenSSL 1.1.1k FIPS 25 Mar 2021, mlx5_core 4.18.0-477.15.1.el8_8.x86_64, Varnish Enterprise 6.0.11r4 revision 676b15e5f7393eb5d5700df47ea504055db032d4, wrk master 02/07/2021 (keep alive, 192, 400 OR 3200 total connections). Throughput measured with 100% Transport Layer Security (TLS) traffic with 100% target cache hit ratio. Test by Intel as of 8/22/2023.

Single Intel® Xeon® 5418N: 1-node, 1x Intel® Xeon® 5418N, 28 cores, HT on, Turbo ON, SNC 2, Total Memory 512GB (8 slots / 64GB / 4800 MT/s @ 4000 MT/s), BIOS 1.1, microcode 0x2b000181, 1x NVIDIA MCX755106AS-HEAT, RHEL 8.8, 10x Samsung PM1743 MZ-WL015T0 15.36TB, kernel 4.18.0-477.15.1.el8_8.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-18), OpenSSL 1.1.1k FIPS 25 Mar 2021, mlx5_core 4.18.0-477.15.1.el8_8.x86_64, Varnish Enterprise 6.0.11r4 revision 676b15e5f7393eb5d5700df47ea504055db032d4, wrk master 02/07/2021 (keep alive, 384, 800 OR 6400 total connections) Throughput measured with 100% Transport Layer Security (TLS) traffic with 100% target cache hit ratio. Test by Intel as of 8/22/2023.

Single Intel® Xeon® 6428N: 1-node, 1x Intel® Xeon® 6428N, 32 cores, HT on, Turbo ON, SNC 2, Total Memory 512GB (8 slots / 64GB / 4800 MT/s @ 4000 MT/s), BIOS 1.1, microcode 0x2b000181, 1x NVIDIA MCX755106AS-HEAT, RHEL 8.8, 10x Samsung PM1743 MZ-WL015T0 15.36TB, kernel 4.18.0-477.15.1.el8_8.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-18), OpenSSL 1.1.1k FIPS 25 Mar 2021, mlx5_core 4.18.0-477.15.1.el8_8.x86_64, Varnish Enterprise 6.0.11r4 revision 676b15e5f7393eb5d5700df47ea504055db032d4, wrk master 02/07/2021 (keep alive, 384, 800 OR 6400 total connections) Throughput measured with 100% Transport Layer Security (TLS) traffic with 100% target cache hit ratio. Test by Intel as of 8/22/2023.

Dual Intel® Xeon® 6438N: 1-node, 2x Intel® Xeon® 6438N, 32 cores, HT on, Turbo ON, SNC 2, Total Memory 512GB (16 slots / 32GB / 4800 MT/s), BIOS 1.3, microcode 0x2b000461, 2x NVIDIA MCX755106AS-HEAT, RHEL 8.8, 12x Samsung PM1743 MZ-WL015T0 15.36TB, kernel 4.18.0-477.15.1.el8_8.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-18), OpenSSL 1.1.1k FIPS 25 Mar 2021, mlx5_core 4.18.0-477.15.1.el8_8.x86_64, Varnish Enterprise 6.0.11r4 revision 676b15e5f7393eb5d5700df47ea504055db032d4, wrk master 02/07/2021 (keep alive, 768, 1600 OR 12800 total connections). Throughput measured with 100% Transport Layer Security (TLS) traffic with 100% target cache hit ratio. Test by Intel as of 8/22/2023.

Dual Intel® Xeon® 8480+: 1-node, 2x Intel® Xeon® 8480+, 56 cores, HT on, Turbo ON, SNC 4, Total Memory 512GB (16 slots / 32GB / 4800 MT/s), BIOS 1.3, microcode 0x2b000461, 4x NVIDIA MCX755106AS-HEAT, RHEL 8.8, 12x Samsung PM1743 MZ-WL015T0 15.36TB, kernel 4.18.0-477.15.1.el8_8.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-18), OpenSSL 1.1.1k FIPS 25 Mar 2021, mlx5_core 4.18.0-477.15.1.el8_8.x86_64, Varnish Enterprise 6.0.11r4 revision 676b15e5f7393eb5d5700df47ea504055db032d4, wrk master 02/07/2021 (keep alive, 768, 1600 OR 12800 total connections). Throughput measured with 100% Transport Layer Security (TLS) traffic with 100% target cache hit ratio. Test by Intel as of 8/22/2023.

0324/CP/H09/PDF ♻ Please Recycle 359570-001US