

Select Google Cloud N2 VMs Featuring 3rd Gen Intel[®] Xeon[®] Scalable Processors with Intel QuickAssist Technology and Get up to 3.21x the NGINX Performance



Intel QAT



Small VMs:

Up to 2.99 Times the Performance Using Intel QAT

High CPU N2 VMs with 4 vCPUs vs. N2D Standard VMs with 4 vCPUs



Medium VMs:

Up to 3.00 Times the Performance Using Intel QAT

Standard N2 VMs with 16 vCPUs vs. N2D Standard VMs with 16 vCPUs



Large VMs:

Up to 2.91 Times the Performance Using Intel QAT

on High CPU N2 VMs with 48 vCPUs vs. N2D Standard VMs with 48 vCPUs

Standard and High CPU VMs Using Intel QAT Outperformed AMD EPYC Processor-Based VMs without Intel QAT

Providing safe, encrypted network connections using SSL/TLS is a smart business move because it puts your consumers at ease about sending their data over the Internet. NGINX is an open-source web server application that serves as a reverse proxy, load balancer, and mail proxy. If your company runs NGINX web servers on Google Cloud, you can improve performance dramatically by choosing N2 VMs featuring 3rd Gen Intel Xeon Scalable processors.

Installing Intel QuickAssist Technology (Intel[®] QAT) on Google Cloud N2 Virtual Machines can boost NGINX web server performance. Intel QAT and Intel QAT Engine (OpenSSL Engine) use vectorized instructions to speed cryptographic operations and allow more users to connect at a time, which in turn speeds hardware and software. We tested NGINX performance for four series of VMs: N2D and T2D, both powered by AMD EPYC processors, and Standard N2 VMs and High CPU N2, both featuring 3rd Gen Xeon Scalable processors. Our results show that choosing one of the N2 VMs featuring 3rd Gen Xeon Scalable processors and adding Intel QAT would yield the best performance for your NGINX workloads.

VMs Featuring 3rd Gen Xeon Scalable Processors with Intel QAT Improved Performance at Every VM Size

Figure 1 shows results at smaller vCPU counts. Using Intel QAT with N2 VMs with 3rd Gen Intel Xeon Scalable processors boosted NGINX performance significantly over both AMD EPYC-processor-based VMs. N2 VMs with Intel QAT achieved a connection rate up to 2.99 times that of N2D VMs at 4 vCPUs and up to 2.98 times that of N2D VMs at 8 vCPUs.

Normalized NGINX Performance for Small VMs

Relative connections per second | Larger is better


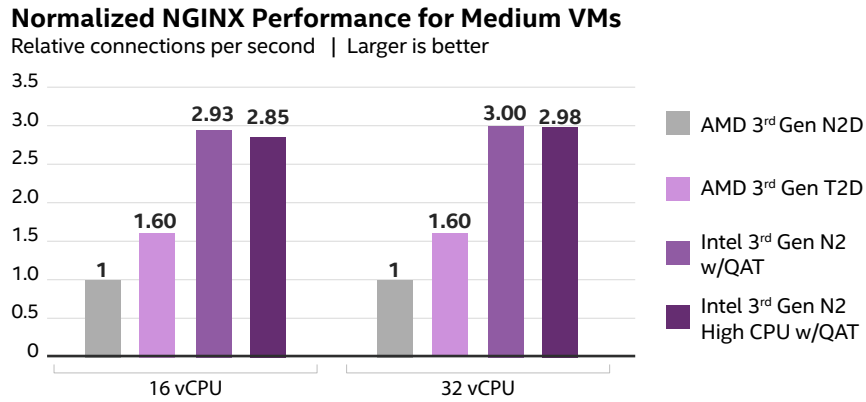


Figure 1. Relative NGINX performance, in connections per second, achieved by small N2D and T2D VMs without Intel QAT and N2 VMs with Intel QAT.


N2 VMs featuring 3rd Gen Intel® Xeon® Scalable processors with Intel QAT achieved a connection rate up to 2.93 times that of N2D VMs at 16 vCPUs and up to 3.00 times that of N2D VMs at 32 vCPUs. (See Figure 2.)

Figure 2. Relative NGINX performance, in connections per second, achieved by medium N2D and T2D VMs without Intel QAT and N2 VMs with Intel QAT.



N2 VMs featuring 3rd Gen Intel Xeon Scalable processors with Intel QAT achieved a connection rate up to 2.91 times that of N2D VMs at 48 vCPUs and up to 2.54 times that of N2D VMs at 64 vCPUs. (See Figure 3.)

Figure 3. Relative NGINX performance, in connections per second, achieved by large N2D and T2D VMs without Intel QAT and N2 VMs with Intel QAT.



Conclusion

Google Cloud N2 VMs featuring 3rd Gen Intel Xeon Scalable processors and using Intel QAT delivered much stronger NGINX performance than both N2D and T2D VMs with AMD EPYC processors.

Learn More

To begin running your NGINX workloads on Google Cloud N2 VMs with 3rd Gen Intel Xeon Scalable processors, visit https://cloud.google.com/compute/docs/general-purpose-machines#n2_machines.

All tests by Intel on 2/12/22 on GCP/us-centra1-a. All tests: Ubuntu 20.04.3 LTS 5.11.0-1028-gcp, v1.24.0.intel-17-g4e2441b6, gcc (Ubuntu 7.5.0-3ubuntu1~18.04) 7.5.0, ldd (Ubuntu GLIBC 2.27-3ubuntu1.4), Client Server:n2-standard-128, Number of Clients:2, Run Iterations:3, Cipher: AES128-GCM-SHA256. All QAT testing: async mode Nginx w/ QATEngine. All ICX VMs used ICX x86_64 CPUs; CLX VMs used CLX x86_64 CPUs. VM details: n2-standard-4: 4 vCPUs, 16GB RAM; n2-highcpu-4: 4 vCPUs, 4GB RAM; n2-standard-8: 8 vCPUs, 32GB RAM; n2-highcpu-8: 8 vCPUs, 8GB RAM; n2-standard-16: 16 vCPUs, 64GB RAM; n2-highcpu-16: 16vCPUs, 16GB RAM; n2-standard-32: 32 vCPUs, 128GB RAM; n2-highcpu-32: 32 vCPUs, 32GB RAM; n2-standard-48: 48 vCPUs, 192GB RAM; n2-highcpu-48: 48 vCPUs, 48GB RAM; n2-standard-64: 64 vCPUs, 256GB RAM; n2-highcpu-64: 48 vCPUs, 64GB RAM.

