


# Get up to 3x Better NGINX Performance with AWS Instances Featuring 3<sup>rd</sup> Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable Processors



NGINX



**Up to 3.02x  
the NGINX  
Connections  
per Second with  
C6i Instances  
Enabled by Intel  
Xeon Scalable  
Processors  
with Crypto  
Acceleration**

*vs. C6a Instances with  
AMD EPYC processors*

## **AWS C6i Instances with Crypto Acceleration Handled More Connections per Second than C6a Instances with AMD EPYC Processors**

Encrypting data as it travels makes it possible for customers to safely engage in business with your organization, but encryption does require system resources and can affect overall performance. Organizations running NGINX web servers in the cloud can benefit from instances that support greater user demands. Rather than adding numerous instances to cover demand, selecting the right instances from the start can ensure your organization handles peak performance that allows your users to connect without delay.

AWS offers compute-optimized C6i instances enabled by 3<sup>rd</sup> Gen Intel Xeon Scalable processors with features that boost NGINX performance: Intel Quick Assist Technology (Intel QAT) with Crypto Acceleration. Intel QAT and Intel QAT Engine (OpenSSL Engine) accelerate hardware and software based on vectorized instructions to speed cryptographic operations and allow more users to connect at a time.

We tested two types of AWS instances: C6i instances with 3<sup>rd</sup> Gen Intel Xeon Scalable processors that use Crypto Acceleration and C6a instances with AMD EPYC processors. We found that C6i instances enabled by 3<sup>rd</sup> Gen Intel Xeon Scalable processors delivered significantly higher performance than C6a instances with AMD processors across vCPU counts. These results show that no matter the size of your workload, C6i instances could help more users safely connect.

## **C6i Instances with Crypto Acceleration Handled More Connections**

Different businesses require different VM sizes to meet their specific workload needs. We tested AWS compute-optimized instances with 4 vCPUs, 8 vCPUs, 16 vCPUs, 32 vCPUs, and 64 vCPUs. Compared to the C6a instances with AMD EPYC processors, the C6i instances with 3<sup>rd</sup> Gen Intel Xeon Scalable processors and Crypto Acceleration offered dramatically improved performance, handling up to 3.02 times the NGINX connections per second (Figure 1).



### Normalized NGINX Performance of C6i Instances vs. C6a Instances

Relative connections per second | Higher is better




Figure 1. Relative NGINX performance, in connections per second, that C6i instances handled compared to C6a instances. Higher numbers are better.

As Figure 1 shows, the performance differential between C6i instances and C6a instances was relatively stable across vCPU counts, ranging from 2.68 times the connections per second at 16 vCPUs to 3.02 times the connections per second at 64 vCPUs.

### Conclusion

Whether you're planning for future growth or meeting an existing need for encrypting web traffic with NGINX, these test results show that choosing C6i instances enabled by 3<sup>rd</sup> Gen Intel Xeon Scalable processors with Crypto Acceleration can better meet your needs. Compared to AWS C6a instances with AMD EPYC processors, C6i instances achieved up to 3.02 times the connections per second on an NGINX workload, making it the smarter choice for hosting NGINX in the cloud.

### Learn More

To begin running your NGINX workloads on AWS C6i instances visit <https://aws.amazon.com/ec2/instance-types/c6i/>.

All tests by Intel on AWS/us-west-2b from 03/2022-04/2022. All tests: Ubuntu 20.04.4 LTS 5.13.0-1019-aws, v1.24.2.intel-13-g5ae1948f, gcc (Ubuntu 9.4.0-1ubuntu1~20.04.1) 9.4.0, ldd (Ubuntu GLIBC 2.31-0ubuntu9.7) 2.31, Client Server: c6i.32xlarge, Number of Clients:2, Run Iterations:3, Cipher: AES128-GCM-SHA256. All QAT configs: async mode Nginx w/ QATEngine. VM Instance details: c6i.xlarge: ICX x86\_64 CPUs, 4 vCPUs, 8GB RAM, 4 worker processes; c6a.xlarge: Zen3 x86\_64 CPUs, 4 vCPUs, 8GB RAM, 4 worker processes; c6i.2xlarge: ICX x86\_64 CPUs, 8 vCPUs, 16GB RAM, 8 worker processes; c6a.2xlarge: Zen3 x86\_64 CPUs, 8 vCPUs, 16GB RAM, 8 worker processes; c6i.4xlarge: ICX x86\_64 CPUs, 16 vCPUs, 32GB RAM, 16 worker processes; c6a.4xlarge: Zen3 x86\_64 CPUs, 16 vCPUs, 32GB RAM, 16 worker processes; c6i.8xlarge: ICX x86\_64 CPUs, 32 vCPUs, 64GB RAM, 32 worker processes; c6a.8xlarge: Zen3 x86\_64 CPUs, 32 vCPUs, 64GB RAM, 32 worker processes; c6i.16xlarge: ICX x86\_64 CPUs, 64 vCPUs, 128GB RAM, 64 worker processes; c6a.16xlarge: Zen3 x86\_64 CPUs, 64 vCPUs, 128GB RAM, 64 worker processes



Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://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 above 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 Corporation. 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.

Printed in USA 0822/JO/PT/PDF US001

