

GOOGLE CLOUD C3D CONFIDENTIAL INSTANCE PERFORMANCE PROTECTION WITH PERFORMANCE


Powered by 4th Gen AMD EPYC™ Processors
 August 2024

AT A GLANCE

Google Cloud C3D instances powered by 4th Gen AMD EPYC™ processors experience only small performance impacts across several key workloads when enabling [AMD Secure Encrypted Virtualization \(SEV\)](#), which uses the AMD Secure Processor to issue one key per virtual machine that isolates guests and the hypervisor from one another, thereby defending data while it is in use.¹

PERFORMANCE HIGHLIGHTS

These charts highlight some of the best performance uplifts of Google Cloud c3d-standard-16 instances powered by 4th Gen AMD EPYC processors running popular workloads. In these charts, performance with AMD SEV disabled always equals 1.00x so that you can easily see the performance impact of enabling AMD SEV to defend your sensitive data.



GOOGLE CLOUD C3D INSTANCES

Google Cloud C3D instances are part of the [general-purpose machine family](#)* that support up to 360 vCPUs and up to 2,880 GB of memory, where each CPU uses two threads per core and each vCPU represents a single thread. These instances are powered by 4th Gen AMD EPYC processors with a maximum frequency of 3.7 GHz and are optimized to deliver optimal, reliable, and consistent performance. By default, C3D instances support up to 100 Gbps network bandwidth and up to 200 Gbps per VM Tier_1 networking performance. C3D VMs are available in sizes from 4 to 360 vCPUs and up to 2,880 GB of memory in the following configurations:

- standard, for general use.
- highcpu, for applications that do not require large amounts of memory.
- highmem, for applications that require larger amounts of memory.
- standard with lssd, which offers local NVMe storage.
- highmem with lssd.

IN THIS BRIEF

- Performance Results on Key Workloads [Page 2](#)
- AMD EPYC 9004 Series Processors..... [Page 5](#)
- Instance Configurations..... [Page 5](#)
- For Additional Information [Page 5](#)

PERFORMANCE RESULTS ON KEY WORKLOADS

This section displays the performance uplifts of Google Cloud c3d-standard-16 instances powered by 4th Gen AMD EPYC processor running the same workloads with identical configurations with AMD Secure Encrypted Virtualization disabled and then enabled. All of the results presented in this section are the average of at least three test runs performed by AMD for each of the listed benchmarks/workloads. All of the charts in this section display the performance impacts between Google Cloud c3d-standard-16 instances with AMD SEV enabled normalized to the performance of the same Google Cloud c3d-standard-16 instance with AMD SEV disabled. The testing and instance configurations used were current for the us-central1 (Iowa) region as of May 30th, 2024, and all tests ran on Ubuntu® 22.04.3 LTS instances using 16 vCPUs with 8 physical cores and 64 GB memory per instance.²

Note: All results in all charts are normalized to the Google c3d-standard-16 instance with AMD SEV disabled.

JAVA PERFORMANCE

The Server Side Java® benchmark enables performance measurements of server-side Java based applications. Server Side Java® simulates a company with an IT infrastructure that handles a mix of point-of-sale requests, online purchases, and data-mining operations. With the rapid adoption of Java across the industry in the last two decades, this benchmark is relevant to all audiences including Java Virtual Machine (JVM) vendors, hardware developers, Java application developers, researchers, and members of the academic community. Figure 1 compares the Google Cloud c3d-standard-16 Server Side Java multiple-instance max ops/sec performance with and without AMD SEV enabled, with all results normalized to non-confidential (AMD SEV disabled) configuration.




Figure 1: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) Server Side Java multi-instance max ops/sec performance impact

WEB SERVER PERFORMANCE WITH NGINX™

NGINX™ is a popular web server that can also be used as a reverse proxy, load balancer, mail proxy, and HTTP cache. AMD tested NGINX throughput in connections per second as a high-performance web server in conjunction with the WRK web (http) client. These tests used a single NGINX server instance on a single c3d-standard-16 instance. Testing retained key NGINX server parameters at their default values, including the number of worker processes and cache manager/loader. Figure 2 compares the Google Cloud c3d-standard-16 NGINX performance with and without AMD SEV enabled, with all results normalized to non-confidential (AMD SEV disabled) configuration. These tests measured NGINX 1.18.0 (Ubuntu) requests per second (rps) using WRK 4.2.




Figure 2: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) NGINX performance impact

MYSQL TPROC-C

TPC Benchmark C™ is an on-line transaction processing (OLTP) benchmark that defines a set of functional requirements that can be run on any transaction processing system, regardless of hardware or operating system. The HammerDB 4.2 benchmark tool was used to build and generate the TPROC-C workload. The HammerDB TPROC-C workload is an open-source workload derived from the TPC-C™ Benchmark Standard and as such is not comparable to published TPC-C results, as the results do not comply with the TPC-C Benchmark Standard. Figure 3 compares the Google Cloud c3d-standard-16 MySQL 8.0.28 TPROC-C performance with and without AMD SEV enabled, with all results normalized to non-confidential (AMD SEV disabled) configuration.




Figure 3: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) MySQL TPROC-C performance impact

IN-MEMORY DATABASE PERFORMANCE WITH REDIS™

Redis™ is an in-memory data structure store used as a distributed, in-memory key-value database, cache, and message broker, with optional durability. Redis supports different kinds of abstract data structures, such as strings, lists, maps, sets, sorted sets, HyperLogLogs, bitmaps, streams, and spatial indices. Redis works with an in-memory dataset to achieve top performance. Depending on the use case, Redis can persist the data either by periodically dumping the dataset to disk or by appending each command to a disk-based log. Figures 4 and 5 show the Google Cloud c3d-standard-16 Redis 7.4.0 with jemalloc 5.3.0-0 GET and SET performance with and without AMD SEV enabled, with all results normalized to non-confidential (AMD SEV disabled) configuration.




Figure 4: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) Redis GET performance impact




Figure 5: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) Redis SET performance impact

FFMPEG

FFmpeg is a multimedia framework that can encode, decode, transcode, stream, filter and play just about any type of video in virtually any format, from legacy to ultramodern. FFmpeg works on a wide variety of operating systems, environments, and configurations. Figures 6-9 showcase the Google Cloud c3d-standard-16 FFmpeg 4.4.2-Ubuntu0.22.04.1 encoding and transcoding performance at 1080p resolution with and without AMD SEV enabled, with all results normalized to non-confidential (AMD SEV disabled) configuration




Figure 6: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) FFmpeg encode performance impact (raw to vp9)




Figure 7: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) FFmpeg encode performance impact (raw to H.264)




Figure 8: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) FFmpeg transcode performance impact (H.264 to vp9)




Figure 9: Google Cloud c3d-standard-16 and c3d-standard-16 (SEV) FFmpeg transcode performance impact (vp9 to H.264)

AMD EPYC 9004 SERIES PROCESSORS

AMD EPYC 9004 Series Processors continue to redefine processor standards for modern datacenters. 4th Gen AMD EPYC processors are built on the innovative x86 architecture and “Zen 4” core. 4th Gen AMD EPYC processors deliver efficient, optimized performance by combining high frequencies, the largest-available L3 cache, 128 lanes of PCIe® 5 I/O (1P) and up to 160 lanes (2P), and synchronized fabric and memory clock speeds, plus support for up to 6 TB of DDR5-4800 memory. Built-in security features, such as AMD Infinity Guard technology, Secure Memory Encryption (SME), and Secure Encrypted Virtualization (SEV-SNP) help protect data while it is in use.¹ See [AMD EPYC™ Processors](#) for additional information.

INSTANCE CONFIGURATIONS

Table 1 shows the available standard Google Cloud C3D instance configuration options as of June 14th, 2024. Please see [C3D machine series](#)* for current information, including instance types with varying CPU/RAM ratios.

MACHINE TYPES	vCPU	MEMORY (GIB)	DEFAULT EGRESS BANDWIDTH (GBPS)	TIER_1 EGRESS BANDWIDTH (GBPS)
c3d-standard-4	4	4	up to 20	N/A
c3d-standard-8	8	8	up to 20	N/A
c3d-standard-16	16	16	up to 20	N/A
c3d-standard-30	30	30	up to 20	up to 50
c3d-standard-60	60	60	up to 40	up to 75
c3d-standard-90	90	90	up to 60	up to 100
c3d-standard-180	180	720	up to 100	up to 150
c3d-standard-360	360	1,440	up to 100	up to 200

Table 1: Google Cloud C3D instance sizes (standard only)

FOR ADDITIONAL INFORMATION

Please see the following additional resources for more information about 4th Gen AMD EPYC features, architecture, and available models:

- [AMD EPYC™ 9004 Series Processors](#)
- [AMD EPYC™ Products](#)
- [AMD Documentation Hub](#)
- [General-Purpose Machine Family for Compute Engine*](#)

REFERENCES

1. AMD Infinity Guard features vary by EPYC™ Processor generations and series. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <https://www.amd.com/en/technologies/infinity-guard>. GD-183A
2. Testing performed in May, 2024. Results may vary due to changes to the underlying configuration, and other conditions such as the placement of the VM and its resources, optimizations by the cloud service provider, accessed cloud regions, co-tenants, and the types of other workloads exercised at the same time on the system.

AUTHORS

Allen Leibovitch and Anthony Hernandez contributed to this Performance Brief.

RELATED LINKS

- [General-Purpose Machine Family for Compute Engine*](#)
- [AMD EPYC™ Processors](#)
- [AMD Documentation Hub](#)

**Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.*

BOOST PERFORMANCE WITH 4TH GEN AMD EPYC

Google Cloud C3D instances are powered by 4th Gen AMD EPYC processors and support up to 360 vCPUs, 2880 GB of memory, and 100 Gbps of network bandwidth. These instances are powered by Google's Titanium system of offloads that includes Google's custom Infrastructure Processing Unit. C3D instances support Google Kubernetes Engine (GKE) and are optimized for consistent performance.

GOOGLE CLOUD C INSTANCES

The C family of instances can be the optimal choice for workloads such as web, app and ad servers, databases and caches, game servers, analytics, media streaming, and network appliances.

AMD EPYC 9004 FOR GENERAL COMPUTING

4th Gen AMD EPYC processors deliver blazing per-core performance thanks to fast CPU frequencies and low latency memory. AMD EPYC processors provide high bandwidth between nodes with support for PCIe® Gen 5 network devices and accelerators that greatly benefit general computing applications.

DISCLAIMERS

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18u

COPYRIGHT NOTICE

©2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices. Java is a registered trademark of Oracle and/or its affiliates. NGINX and the NGINX logo are trademarks of F5, Inc. in the US and other countries. Redis is a trademark of Redis Labs. Ubuntu is a registered trademark of Canonical, Ltd. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure.