

GENERATIONAL AMAZON EC2 M7A PERFORMANCE GAINS VERSUS PRIOR-GENERATION M6A INSTANCES

Powered by 4th Gen AMD EPYC™ Processors


October 2023

AT A GLANCE

Amazon EC2 M7a instances powered by 4th Gen AMD EPYC™ processors deliver strong performance and price-performance uplifts compared to Amazon EC2 M6a instances powered by 3rd Gen AMD EPYC processors on a variety of key workloads.¹

PERFORMANCE HIGHLIGHTS

These charts highlight some of the best performance and price-performance uplifts of Amazon EC2 M7a instances powered by 4th Gen AMD EPYC processors normalized to equivalent-sized Amazon EC2 M6a instances powered by 3rd Gen AMD EPYC processors. These charts show selected workload uplifts on instances. Detailed results, including additional benchmarks, appear inside this brief.



AMAZON EC2 M7A INSTANCES

Amazon EC2 M7a instances are designed to deliver the best x86 performance and price performance within the Amazon EC2 general purpose family. M7a instances are powered by 4th-generation AMD EPYC processors with a maximum frequency of up to 3.70 GHz.² Per Amazon, these instances deliver up to 50% greater performance on average compared to Amazon EC2 M6a instances powered by 3rd-Gen AMD EPYC processors and offer new processor capabilities such as AVX-512, VNNI, and BFloat16.³

Amazon EC2 M7a are available in a range of sizes from 1 vCPU core and 4 GiB memory to 192 vCPU cores and 768 GiB of memory with network bandwidth from 12.5 GiB to 50 GiB and 10-40 GiB of Elastic Block Store (EBS) bandwidth. Amazon EC2 M7a instances feature DDR5 memory and deliver 2.25x more memory bandwidth compared to M6a instances to enable high-speed access to data in memory. M7a instances are built on the AWS Nitro System and are ideal for applications that benefit from high performance and high throughput such as financial applications, application servers, simulation modeling, gaming, mid-size data stores, application development environments, and caching fleets. See [Amazon EC2 M7a Instances*](#) to learn more.

IN THIS BRIEF

- Performance Results on Key Workloads **Page 2**
- Instance Configurations..... **Page 7**
- AMD EPYC 9004 Series Processors..... **Page 7**
- For Additional Information **Page 8**

PERFORMANCE RESULTS ON KEY WORKLOADS

This section displays the performance uplifts of Amazon EC2 M7a instances powered by 4th Gen AMD EPYC processors versus Amazon EC2 M6a instances powered by 3rd Gen AMD EPYC processors. All of the results presented in this section are the average of three test runs performed by AMD for each of the listed benchmarks/workloads. All of the charts in this section display the performance or price-performance uplifts between Amazon EC2 M7a instances normalized to the performance or price-performance of the equivalent-sized Amazon EC2 M6a instance, where the performance and price-performance of the equivalent Amazon EC2 M6a instance will always equal 1.00x. The testing, instance configurations, and pricing information used for price-performance comparisons was current for the US-EAST (Ohio) region as of September 23rd, 2023 and is subject to change. Please visit [Amazon EC2 pricing*](#) and the [AWS Cost Calculator*](#) for current instance pricing information.

JAVA PERFORMANCE

The Server Side Java® 2015 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. Figures 1-4 show Server Side Java 2015 single- and multiple-instance performance and price-performance uplifts of Amazon EC2 M7a instances versus Amazon EC2 M6a instances, respectively, with all results normalized to the M6a instances.




Figure 1: Amazon EC2 M7a vs. M6a
Server Side Java 2015 Composite performance uplift




Figure 2: Amazon EC2 M7a vs. M6a
Server Side Java 2015 Composite price-performance uplift




Figure 3: Amazon EC2 M7a vs. M6a
Server Side Java 2015 MultiJVM performance uplift




Figure 4: Amazon EC2 M7a vs. M6a
Server Side Java 2015 MultiJVM price-performance uplift

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 M7a or M71 instance. Testing retained key NGINX server parameters at their default values, including the number of worker processes and cache manager/loader. Figures 5 and 6 showcases the performance uplift of Amazon EC2 M7a instances versus Amazon EC2 M6a instances in requests per second (rps) achieved, with all results normalized to the M6a instances.




Figure 5: Amazon EC2 M7a vs. M6a NGINX performance uplift




Figure 6: Amazon EC2 M7a vs. M6a NGINX price-performance uplift

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 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. Figures 7 and 8 showcase the performance uplift of Amazon EC2 M7a instances versus Amazon EC2 M6a instances in queries per hour at size (qphsize) achieved and price-performance uplift of these instances, respectively, with all results normalized to the M6a instances.




Figure 7: Amazon EC2 M7a vs. M6a MySQL TPROC-C performance uplift




Figure 8: Amazon EC2 M7a vs. M6a MySQL TPROC-C price-performance uplift

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 9-12 show the Redis 2015 GET and SET performance and price-performance uplifts of Amazon EC2 M7a instances versus Amazon EC2 M6a instances, respectively, with all results normalized to the M6a instances.




Figure 9: Amazon EC2 M7a vs. M6a Redis GET performance uplift




Figure 10: Amazon EC2 M7a vs. M6a Redis GET price-performance uplift




Figure 11: Amazon EC2 M7a vs. M6a Redis SET performance uplift




Figure 12: Amazon EC2 M7a vs. M6a Redis SET price-performance uplift

RELATIONAL DATABASE MANAGEMENT SYSTEMS

Relational Database Management Systems (RDBMS) continue to be the foundation for business-critical applications. HammerDB is a decision support benchmark that executes complex queries that require processing large volumes of data. The HammerDB TPROC-H workload is an open source workload derived from the TPC-H Benchmark Standard and as such is not comparable to published TPC-H results, as the results do not comply with the TPC-H Benchmark Standard. Figures 13 and 14 showcase the performance uplift of Amazon EC2 M7a instances versus Amazon EC2 M6a instances across several TPROC-H metrics running on MS SQL Server and the price-performance uplift of these instances, respectively, with all results normalized to the M6a instances.




Figure 13: Amazon EC2 M7a vs. M6a MS SQL Server TPROC-H performance uplift




Figure 14: Amazon EC2 M7a vs. M6a MS SQL Server TPROC-H price-performance uplift

FINANCIAL TRANSACTIONS

TPC Benchmark™ E is an on-line transaction processing (OLTP) benchmark that is more complex than previous OLTP benchmarks such as TPC-C because it includes varied transaction types and a more complex database and overall execution structure. It mixes 12 twelve concurrent transactions of various types and complexity. Transactions can execute on-line or when triggered by price or time criteria. The database consists thirty-three tables with diverse columns, cardinality, and scaling properties. This benchmark simulates a stock brokerage firm but is not limited to a particular business segment; it represents how many industries executes and reports financial transactions. The HammerDB TPROC-E workload is an open-source workload derived from the TPC-E Benchmark Standard and as such is not comparable to published TPC-E results, as the results do not comply with the TPC-E Benchmark Standard. Figures 15 and 16 showcase the Amazon EC2 M7a instances performance uplift versus Amazon EC2 M6a instances in transaction per second (tpsE) achieved and the price-performance uplift of these instances, respectively, with all results normalized to the M6a instances.




Figure 15: Amazon EC2 M7a vs. M6a MS SQL Server TPROC-E performance uplift




Figure 16: Amazon EC2 M7a vs. M6a MS SQL Server TPROC-E price-performance uplift

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 17-24 showcase the performance uplift of Amazon EC2 M7a instances versus Amazon EC2 M6a instances running FFmpeg encoding and transcoding jobs and the price-performance uplift of these instances, respectively, with all results normalized to the M6a instances.




Figure 17: Amazon EC2 M7a vs. M6a FFmpeg encode performance uplift (raw to vp9)




Figure 18: Amazon EC2 M7a vs. M6a FFmpeg encode price-performance uplift (raw to vp9)




Figure 19: Amazon EC2 M7a vs. M6a FFmpeg encode performance uplift (raw to H.264)




Figure 20: Amazon EC2 M7a vs. M6a FFmpeg encode price-performance uplift (raw to H.264)




Figure 21: Amazon EC2 M7a vs. M6a FFmpeg transcode performance uplift (H.264 to vp9)




Figure 22: Amazon EC2 M7a vs. M6a FFmpeg transcode price-performance uplift (H.264 to vp9)




Figure 23: Amazon EC2 M7a vs. M6a FFmpeg transcode performance uplift (vp9 to H.264)




Figure 24: Amazon EC2 M7a vs. M6a FFmpeg transcode price-performance uplift (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, 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.⁴

INSTANCE CONFIGURATIONS

Table 1 shows the available Amazon EC2 M7a instance configuration options as of the date this Performance Brief was published. Please see [Amazon EC2 M7a Instances](#) for current information.

SIZE	vCPU	MEMORY (GiB)	STORAGE (GB)	NETWORK (Gbps)	EBS (Gbps)
m7a.medium	1	4	EBS-Only	Up to 12.5	Up to 10
m7a.large	2	8	EBS-Only	Up to 12.5	Up to 10
m7a.xlarge	4	16	EBS-Only	Up to 12.5	Up to 10
m7a.2xlarge	8	32	EBS-Only	Up to 12.5	Up to 10
m7a.4xlarge	16	64	EBS-Only	Up to 12.5	Up to 10
m7a.8xlarge	32	128	EBS-Only	12.5	10
m7a.12xlarge	48	192	EBS-Only	18.75	15
m7a.16xlarge	64	256	EBS-Only	25	20
m7a.24xlarge	96	384	EBS-Only	37.5	30
m7a.32xlarge	128	512	EBS-Only	50	40
m7a.48xlarge	192	768	EBS-Only	50	40
m7a.metal-48xl	192	768	EBS-Only	50	40

Table 1: Amazon EC2 M7a instance sizes

Table 2 shows the available Amazon EC2 M6a instance configuration options as of the date this Performance Brief was published. Please see [Amazon EC2 M6a Instances](#) for current information.

SIZE	vCPU	MEMORY (GiB)	STORAGE (GB)	NETWORK (Gbps)	EBS (Gbps)
m6a.large	2	8	EBS-Only	Up to 12.5	Up to 10
m6a.xlarge	4	16	EBS-Only	Up to 12.5	Up to 10
m6a.2xlarge	8	32	EBS-Only	Up to 12.5	Up to 10
m6a.4xlarge	16	64	EBS-Only	Up to 12.5	Up to 10
m6a.8xlarge	32	128	EBS-Only	12.5	10
m6a.12xlarge	48	192	EBS-Only	18.75	15
m6a.16xlarge	64	256	EBS-Only	25	20
m6a.24xlarge	96	384	EBS-Only	37.5	30
m6a.32xlarge	128	512	EBS-Only	50	40
m6a.48xlarge	192	768	EBS-Only	50	40
m6a.metal	192	768	EBS-Only	50	40

Table 2: Amazon EC2 M6a instance sizes

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](#)

REFERENCES

1. Testing performed in August and September, 2023. Cloud performance results presented are based on the test date in the configuration and are in alignment with AMD internal bare-metal testing factoring in cloud service provider overhead. 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.
2. Maximum boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. EPYC-18
3. Please see the M7a tab at <https://aws.amazon.com/ec2/instance-types/m7a/>.
4. AMD Infinity Guard features vary by EPYC™ Processor generations. 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-183

AUTHORS

Mike Thompson and Anthony Hernandez contributed to this Performance Brief.

RELATED LINKS

- [Amazon EC2 M7a Instances*](#)
- [Elastic Fabric Adapter*](#)
- [Amazon FSx for Lustre*](#)
- [AWS ParallelCluster*](#)
- [AWS Nitro System*](#)
- [AMD EPYC™ Processors](#) and the [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

Amazon EC2 M7a instances are powered by 4th Gen AMD EPYC processors and offer up to 192 CPU cores, 768 GB of high-bandwidth DDR5 memory, and 300 Gbps of Elastic Fabric Adapter (EFA) network bandwidth for fast, low-latency inter-node communications. These instances offer efficient scaling and are ideal for applications that benefit from high performance and high throughput such as financial applications,

application servers, simulation modeling, gaming, mid-size data stores, application development environments, and caching fleets.

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.

AMAZON EC2 M7A INSTANCES

Amazon EC2 M7a instances are SAP-certified and ideal for applications that benefit from high performance and high throughput such as financial applications, application servers, simulation modeling, gaming, mid-size data stores, application development environments, and caching fleets.

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 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's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

COPYRIGHT NOTICE

©2023 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Infinity Fabric, and combinations thereof are trademarks of Advanced Micro Devices, Inc. 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. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.