

MODERNIZE DATA CENTER VIRTUALIZATION WITH AMD EPYC™ PROCESSORS



together we advance_data center computing

August 2024

MODERNIZE DATA CENTER VIRTUALIZATION WITH AMD EPYC™ PROCESSORS

CONTENTS

INTRODUCTION	3
MODERNIZATION ECONOMICS	4
What About Intel?	5
What About Migration Risk?	5
What About Increased Blast Radius?	6
Software Reliability	6
Hardware Reliability	7
A COMPELLING CHOICE	8

INTRODUCTION



The pressure on IT managers today is immense. How are you going to turn on new revenue streams by rolling out new applications? What is your strategy for supporting a newly mobile workforce without the cost of cloud hosting? What plans do you have to keep containers and cloud-native applications on premises? And where are you going to put a range of new GPU-intensive servers for your company's AI initiatives?

Surrounding this pressure is heightened uncertainty: The average age of a data center server is more than five years. Older infrastructure often requires more servers to deliver a given level of performance than newer ones. Older servers may use process technology that today seems outdated, can be less energy efficient, and can be vulnerable to unplanned downtime and security threats. During a time when generational performance gains were only incremental, it was hard to justify server refresh.

That has all changed since the arrival of servers powered by AMD EPYC™ processors. With consistent double-digit performance increases generation over generation, we've changed the calculation for data center modernization. [ROM-236](#), [MLN-003](#), [EPYC-038](#) We have set more than 300 performance records, [EPYC-022E](#) AMD EPYC processors power the most energy-efficient x86-architecture servers in the world, delivering exceptional performance and helping to reduce energy costs. [EPYC-028D](#)

What if we could offer you a modernization approach that would pay for itself in about two months, help you in achieving your sustainability goals, and open up floor space in your data center for new servers for artificial intelligence and machine learning? The economics are compelling.

MODERNIZATION ECONOMICS


Modernizing your data center is the key to alleviating the pressures on your IT organization. Consider a data center with 1000 servers each with two 28-core Intel® Xeon® 8280 processors. These were top-of-the-line processors when they were launched in 2019 and are often still running in data centers that have stretched out their refresh cycle to five years.

Running VMware vSphere® software, these servers each produce a VMmark® benchmark score of 9.02 per server—so in the aggregate, the data center of 1000 servers can deliver a score of 9,020 (Figure 1). A thousand of these 2-socket servers have a total of 56,000 cores, each of which has to be licensed to run vSphere software.

To achieve roughly the same performance as the older servers requires only 336 servers each running two 48-core AMD EPYC 9474F processors. This set of servers has 32,256 cores, saving 23,744 VMware licenses, and resulting in a reduction of ~42% in licensing costs, which can lead to a CapEx payback period of only about two months. Moving to the servers running EPYC processors brings you energy savings of up to 47% and up to a 36% savings in five-year total cost of ownership. Imagine cutting the number of servers in your data center by almost two thirds and letting the licensing savings pay for it in just about two months! [SP5TCO-073A](#)

While AMD has established a reputation for AMD EPYC processors having the highest core counts in the x86-architecture server market (up to 128 cores), inquiring minds might ask why we chose

Figure 1: It only takes 336 servers running AMD EPYC processors to replace 1000 servers with Intel Xeon processors

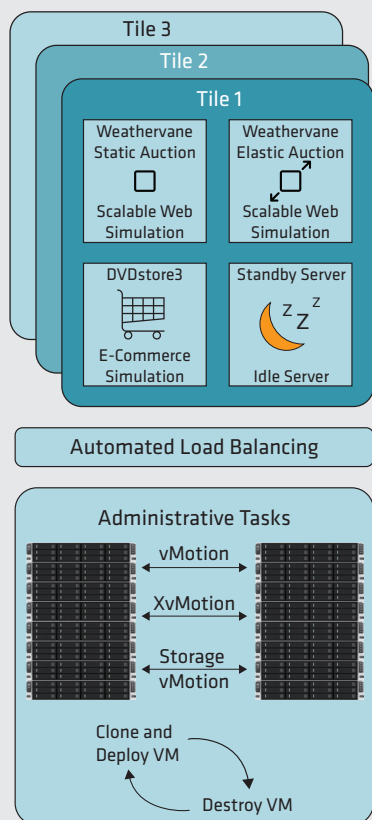


THE VMARK 3.1 BENCHMARK AT A GLANCE

The VMmark benchmark is developed and maintained by VMware, and results are reviewed and posted publicly at vmmark.com.

The benchmark is designed to measure how a cluster performance by running a predefined set of workloads defined as a 'tile.' Each tile runs a two-tier e-commerce application (DVDstore3), two instances of an auction site (Weathervane), and a standby server that gets moved by the tile's infrastructure operations. Each tile also simulates administrative functions that are typical in real-world situations, including vMotion, storage vMotion, XvMotion, and VMware Dynamic Resource Scheduling (DRS) load balances the workloads.

During a test run, the number of tiles is increased until a performance threshold is reached, and the workload and infrastructure performance measurements are combined to produce an overall VMmark 3.1 score that is usually cited along with the number of tiles at which the score was reached.



a 48-core CPU for this example? The 9474F CPU is the top of our high-frequency line of processors denoted by the 'F' in the product number. These CPUs are designed to have lower core counts that run at higher frequencies in order to deliver high performance per core. We accomplish this with our hybrid, multi-chip architecture that allows us to space active cores in the CPU package to optimize thermal characteristics so they can run at higher frequencies. With lifecycle software licensing costs easily dominating CPU costs, it makes sense to use our high-frequency processors with per-core-licensed software so that you can optimize performance for each license. As shown above, with fewer, but high-performance cores, you can get roughly the same performance out of your cluster with much lower licensing costs.

WHAT ABOUT INTEL?

If the first example were your data center, you would be wise to compare our upgrade proposal to what you could achieve by modernizing with current-generation Intel Xeon processors (see Figure 2 on next page). And how about considering 64-core Intel Xeon Platinum 8592+ CPUs, since each 2-socket server produces about the same VMmark score (27.52) as 2-socket servers running 48-core AMD EPYC 9474F processors (26.95)?

To achieve the same ~9020 aggregate VMmark score, you'd need only 328 servers powered by the Intel Xeon processors. But these would require licensing software for 41,984 cores, an approximate 30% increase in licensing costs over the AMD EPYC choice. In addition, choosing AMD EPYC processors gives you up to 24% lower five-year TCO and up to 33% lower hardware costs. [SP5TCO-073A](#)

The facts are compelling. The first example shows how you can simplify your data center, make room for new servers, and pay off the modernization in just about two months. The second example shows that the smart choice is to move to AMD EPYC processors, and we make it easy for you to undertake the migration.

WHAT ABOUT MIGRATION RISK?

You may ask about the risk of migrating thousands of virtual machines from Intel Xeon to AMD EPYC processors. The work to migrate is the same as moving from older Intel Xeon processors to newer ones. It always requires a cold migration to move between generations of processors, and that is true for migration between Intel Xeon generations as it is to move from Intel Xeon to AMD EPYC processors.

If this is news to you, it may be that you are using VMware Enhanced vMotion Compatibility (EVC) to live migrate to new servers without a cold boot. But what EVC does is to make virtual machines on newer CPUs emulate your older CPUs in order to avoid a cold migration. But this can cause you to miss the primary reason

328 Servers

2P Intel Xeon Platinum 8592+

VMmark® 3.1 matched pair performance score of 27.52 per server

Total of 9,027 VMmark 3.1 matched pair performance score



41,984 cores



~2,787k kWh per year



41,984 vSphere licences

2P Intel Xeon
8592+

VS

2P AMD
EPYC™ 9474F

336 Servers

2P AMD EPYC™ 9474F

VMmark 3.1 matched pair performance score of 26.95 per server

Total of 9,055 VMmark 3.1 matched pair performance score



32,256 cores



~2,830k kWh per year



32,256 VMware licences

Up to 24%
lower five-
year TCO

Up to 33%
lower hardware CAPEX

Up to 23%
fewer cores

Up to 23%
lower VMware
licensing cost

AUTOMATED VIRTUAL MACHINE MIGRATION

Migrating virtual machines can be challenging without tools to help. That's why we worked with VMware to create tools that simplify, automate, and accelerate the process:

- **STEP 1:** Identify processors. Select the AMD EPYC™ processors that are best suited to existing virtual machine deployments and workloads. Be sure to identify comparable processors and consider factors that affect TCO.
- **STEP 2:** Download the VMware Architecture Migration Tool. Install and configure the software for the migration session.
- **STEP 3:** Identify and tag virtual machines. Create a list of the virtual machines to be migrated to the new platform. Tag the virtual machines as 'readyToMigrate' so they can take part in the migration and future rollback processes.
- **STEP 3:** Execute the migration. The virtual machines are shut down and migrated from the old cluster to the new cluster. Once the migration is complete, a status summary appears.
- **STEP 4:** Validate the migration result. Check that the virtual machines were migrated to the correct target platform. Note that virtual machines can be rolled back from the migrated cluster at any point in time after a migration operation is complete.

Figure 2: When you choose to modernize, it pays to choose servers with AMD EPYC processors

for modernizing, which includes using the latest processor features for more performance and the latest security features and firmware fixes.

The good news is that we have worked with VMware to create tools to automate the process and help minimize application downtime. The resulting VMware Architecture Migration Tool (VAMT) provides an easy-to-use, automated process for migrating virtual machines between clusters with the same or different x86 processor architectures within a VMware vCenter deployment. This PowerShell script uses VMware PowerCLI to move virtual machines among x86-architecture systems, including between Intel processor-based systems and AMD EPYC processor-based systems.

WHAT ABOUT INCREASED BLAST RADIUS?

A common concern is that, by increasing core counts and the number of virtual machines per server, more VMs could be affected by a single-system outage, putting application availability at risk. The term "Blast radius" refers to the size of a failure domain in terms of number of virtual machines that would need to be recovered.

Two advances in technology mitigate blast radius concerns: the capability for software to mitigate hardware failures; and increasing reliability of CPU hardware.

SOFTWARE RELIABILITY

Most modern software architectures are built to be resilient to hardware failures. Whether they are traditional, multi-tier Web architectures or modern containerized platforms, applications are built to run multiple instances of each component at once, distributed across different servers. For example, in the VMmark benchmark itself, the DVDstore3 component runs three instances of the store's front-end Web servers. If one fails, the application is still operational. The Weathervane Web application is an

IMPORTANT AMD EPYC RAS FEATURES

AMD EPYC processors include a wide range of reliability, availability, and serviceability features, all of which have the goal of maintaining correct operations and providing constant uptime. Two of the most important areas for increasing reliability include correcting memory errors and detecting I/O errors:

- **AMD ADVANCED MEMORY DEVICE CORRECTION (AMDC)** goes beyond standard ECC DRAM by using a type of ECC that allows large groups of bits to be corrected with negligible performance impact. This is a chipkill mechanism that helps prevent silent data corruption that can cause application problems. AMDC helps increase server availability and enables DIMMs with errors to remain in service.
- **DRAM ADDRESS AND COMMAND PARITY WITH REPLAY** helps to overcome transient memory bus errors by replaying the request, helping to maintain high levels of service.
- **DATA POISONING** passes information on uncorrectable errors to the CPU by routing a 'poison' bit along with the data so the CPU can report the problem to the operating system, helping prevent the data from being used by applications.

example of an elastic application where eight application instances are the baseline and it scales up to six more in response to greater workloads—so if any one of these instances fail, a user session may be lost, but the application will still provide services to the majority of users.

In addition, VMware vSphere offers options that can help increase application availability.

- **VSPHERE HIGH AVAILABILITY (HA)** This feature monitors application components, and if it detects a failure—due to hardware or software—it restarts the components on secondary hosts.
- **VSPHERE FAULT TOLERANCE (FT)** runs a live shadow instance of virtual machines on a secondary server that mirror the primary VMs to prevent data loss and downtime in the event of an outage. If an outage occurs, vSphere FT triggers a failover process that makes the mirror the primary server so the application is continuously available.

HARDWARE RELIABILITY

AMD EPYC processors support a wide range of reliability, availability, and serviceability (RAS) features. Our goals for each area are as follows:

- **RELIABILITY:** Give correct answers, and avoid, correct, or recover from errors with a goal of providing constant uptime.
- **AVAILABILITY:** Survive partial failures, and minimize planned and unplanned downtime.
- **SERVICEABILITY:** Make problems effortless to diagnose, easy to service, and simple to deploy.

With every AMD EPYC processor generation, we work with top-tier silicon fabrication vendors to target improvements in CPU reliability.

AMD EPYC processors go beyond checking their internal functions: they actively participate in overall reliability of the platform that uses them. With newer DDR generations, the fault rate of memory chips has increased, and EPYC processor features are designed to mitigate single and multi-bit errors. Our advanced approach to error correction even allows memory DIMMs that might otherwise cause a server outage to remain in service, helping maintain continuous application availability (see sidebar).

A wide range of trends all contribute to reducing concerns about blast radius. Modern applications have resiliency built in. VMware offers formidable tools for increasing availability and even fault tolerance. And at a CPU and a platform level, AMD EPYC processors offer high availability through vigilance over internal functions and resilience from external failures such as memory and I/O errors.

A COMPELLING CHOICE



The pressure on IT managers continues every day with demands to support new revenue streams, a hybrid workforce, cloud-native applications, and new AI initiatives—all with limited budget flexibility.

Recognizing that software licensing costs make it increasingly costly to continue running on old infrastructure, data center modernization is a solution that can lower costs, increase performance, and make room to support AI-specific servers.

If you modernize to servers running AMD EPYC processors, you can reduce the number of servers and slash licensing costs which can quickly pay back your investment. And if you compare the choice of using the latest AMD EPYC processors with the latest Intel Xeon processors, you'll find impressive total cost of ownership benefits when you move to AMD EPYC processors. We help you advance your [data center sustainability goals](#), even as you push the limits of high-performance computing.

Data center modernization is a solution to many of your challenges, and the choice of AMD EPYC processors is a compelling one.

END NOTES

For details on the footnotes used in this document, visit www.amd.com/en/legal/claims/epyc.html.

- EPYC-022E: For a complete list of world records see amd.com/worldrecords.
- EPYC-028D: SPECpower_ssj® 2008, SPECrate®2017_int_energy_base, and SPECrate®2017_fp_energy_base based on results published on SPEC's website as of 2/21/24. VMmark® server power-performance / server and storage power-performance (PPKW) based results published at <https://www.vmware.com/products/vmmark/results3x1.html?sort=score>. The first 105 ranked SPECpower_ssj®2008 publications with the highest overall efficiency overall ssj_ops/W results were all powered by AMD EPYC processors. For SPECrate®2017 Integer (Energy Base), AMD EPYC CPUs power the first 8 top SPECrate®2017_int_energy_base performance/system W scores. For SPECrate®2017 Floating Point (Energy Base), AMD EPYC CPUs power the first 12 SPECrate®2017_fp_energy_base performance/system W scores. For VMmark® server power-performance (PPKW), have the top 5 results for 2- and 4-socket matched pair results outperforming all other socket results and for VMmark® server and storage power-performance (PPKW), have the top overall score. See <https://www.amd.com/en/claims/epyc4#faq-EPYC-028D> for the full list. For additional information on AMD sustainability goals see: <https://www.amd.com/en/corporate/corporate-responsibility/data-center-sustainability.html>. More information about SPEC® is available at <http://www.spec.org>. SPEC, SPECrate, and SPECpower are registered trademarks of the Standard Performance Evaluation Corporation. VMmark is a registered trademark of VMware in the US or other countries.
- EPYC-038 Based on AMD internal testing as of 09/19/2022, geometric performance improvement at the same fixed-frequency on a 4th Gen AMD EPYC™ 9554 CPU compared to a 3rd Gen AMD EPYC™ 7763 CPU using a select set of workloads (33) including est. SPECrate®2017_int_base, est. SPECrate®2017_fp_base, and representative server workloads.
- MLN-003 Based on AMD internal testing as of 02/1/2021, average performance improvement at ISO-frequency on an AMD EPYC™ 72F3 (8C/8T, 3.7GHz) compared to an AMD EPYC™ 7F32 (8C/8T, 3.7GHz), per-core, single thread, using a select set of workloads including SPECrate®2017_int_base, SPECrate®2017_fp_base, and representative server workloads.
- ROM-236 Based on AMD internal testing, average per thread performance improvement at ISO-frequency on a 32-core, 64-thread, 2nd generation AMD EPYC™ platform as compared to 32-core 64-thread 1st generation AMD EPYC™ platform measured on a selected set of workloads including sub-components of SPEC CPU® 2017_int and representative server workloads.
- SP5-028D: SPECrate®2017_fp_base comparison based on published scores from www.spec.org as of 03/31/2023. Comparison of published 2P AMD EPYC 9554 (1230 SPECrate®2017_fp_base, 720 Total TDP W, 128 Total Cores, 9.61 Perf/C, 1.708 Perf/W, <http://spec.org/cpu2017/results/res2023q1/cpu2017-20230116-33525.html>) is 1.21x the performance [1.13x Perf/C] of published 2P Intel Xeon Platinum 8490H (1020 SPECrate®2017_fp_base, 700 Total TDP W, 120 Total Cores, 8.5 Perf/C, 1.457 Perf/W, <http://spec.org/cpu2017/results/res2023q1/cpu2017-20230310-34571.html>) [at 1.17x the performance/W]. Published 2P AMD EPYC 9534 (1160 SPECrate®2017_fp_base, 560 Total TDP W, 128 Total Cores, 9.06 Perf/C, 2.071 Perf/W, <http://spec.org/cpu2017/results/res2023q1/cpu2017-20230116-33521.html>) is shown for reference. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.
- SP5TCO-073A: As of 06/18/2024, this scenario contains many assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision making over actual testing. The Server Refresh & Greenhouse Gas Emissions TCO (total cost of ownership) Estimator Tool compares the selected AMD EPYC™ and Intel® Xeon® CPU based server solutions required to deliver a TOTAL PERFORMANCE of ~9020 units of VMmark3 matched pair performance based on the published scores (or estimated if indicated by an asterisk) for Intel Xeon and AMD EPYC CPU based servers. This estimation reflects a 5-year time frame. This analysis compares a 2P AMD 48 core EPYC_9474F powered server with a VMmark 3.1 score of 26.95 @ 26 tiles, <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2024-05-14-Supermicro-A5-2125H5-TNR.pdf>; compared to a 2P Intel Xeon 64 core Platinum_8592+ based server with a VMmark 3.1 score of 27.52 @ 28 tiles, <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2024-04-16-Fujitsu-PRIMERGY-RX2540M7.pdf>; versus legacy 2P Intel Xeon 28 core Platinum_8280 based server with a VMmark 3.1 score of 9.02 @ 9 tiles. Environmental impact estimates made leveraging this data, using the Country / Region specific electricity factors from the '2020 Grid Electricity Emissions Factors v1.4 - September 2020', and the United States Environmental Protection Agency 'Greenhouse Gas Equivalencies Calculator'. Results generated by: AMD EPYC™ Server Refresh & Greenhouse Gas Emission TCO Estimation Tool - version 1.51 PRO. VMmark is a registered trademark of VMware in the US or other countries. For additional details, see <https://www.amd.com/en/claims/epyc#SP5TCO-073A>

© 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe® is a registered trademark of PCI-SIG Corporation. SPEC® and SPECpower® are trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners. VMware, VMmark, and vSphere are a registered trademarks of VMware in the US or other countries. 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. LE-87402-00 08/24

