Realize 2.1X the performance with 20% less power with AMD EPYC processor-backed clusters

Three AMD EPYC processor-based two-processor solutions outshined comparable Intel Xeon Scalable processor-based solutions by handling more Redis workload transactions and requests while consuming less power

Handle more Redis work

*Up to 2.1X the requests per second**

*From servers with AMD EPYC 9684X processors vs. servers with Intel Xeon Platinum 8480+ processors

Processing more transactions for Redis could mean more real-time analytics, faster retrieval of files, or caching service to reduce repetitive generative AI computations. This advantage came from the server cluster backed by AMD EPYC™ 9684X processors, which features AMD 3D V-Cache technology designed to help keep Redis data closer to the core.

Spend less on data center power

*Up to 20% less power under load**

*From servers with AMD EPYC 9534 processors vs. servers with Intel Xeon Platinum 8480+ processors

The cost of energy is rising rapidly, so slashing any data center spending could be a boon for your bottom line. In addition, consuming less power could help advance your organization's sustainability goals.

Use less energy to get better performance

*Up to 2X the performance per watt**

*From servers with AMD EPYC 9684X processors vs. servers with Intel Xeon Platinum 8480+ processors

What happens when you combine better Redis performance with less power consumption? You get a more efficient hardware cluster: more processed transactions per watt. More efficient clusters could then contribute to making your data center more efficient on the whole, allowing your organization to meet performance demands with fewer physical servers while also helping to reduce data center sprawl.

Overview

When choosing processors for server clusters that run data-intensive workloads, such as Redis, performance is often the number one consideration. A faster processor could enable data-intensive workloads to handle large volumes of data with low latency for responsive caching, session storage, or real-time analytics/AI workload performance. CPU power consumption plays a significant factor as well due to rising data center power costs.[1]

We created two clusters of four worker nodes to compare Redis performance and power consumption in three CPU scenarios. The four worker nodes in the first cluster used 4th Generation AMD EPYC processors, and the four nodes in the second cluster used 4th Gen Intel® Xeon® Scalable processors. All worker nodes had the same amount of RAM on the same number of DIMMs. We installed Red Hat® OpenShift® to containerize the VMs running workloads for each cluster configuration. In each test scenario, the AMD EPYC processor-based cluster supported better performance—up to 2.1 times the requests per second (RPS). In two of the three scenarios, the AMD EPYC processor-based cluster consumed less power (up to 20 percent less watts) while achieving better performance. Furthermore, each configuration of the AMD EPYC processor-based cluster supported more RPS per watt—up to 2 times the RPS per watt in our second scenario. These outcomes suggest that processors from the AMD EPYC 9004 series in your servers could help your Redis or data-intensive applications and workloads do things like support more messages or detect fraud sooner, all while potentially saving on data center power costs.

We tested the following processors from the EPYC family:

AMD EPYC 9534

This data center processor has 64 cores and 128 threads, with a max boost clock speed up to 3.7Ghz.[2]

AMD EPYC 9684X

This server processor has 96 cores and 192 threads, with a max boost clock speed also up to 3.7Ghz. The processor also features AMD 3D V-Cache™ Technology, which stacks cache vertically for up to 1.15 GB of L3 cache.[3]

AMD EPYC 9174F

This processor has 16 cores and 32 threads, with a max boost clock speed up to 4.4Ghz.[4]

About the AMD EPYC processors

The portfolio of fourth generation of AMD EPYC processors offers a range of choices to help meet the many evolving workload demands in today's data centers. The 9004 series in particular, from which we chose the processors we tested, comes available with 16 to 128 cores and up to 1.15 GB of L3 cache.[5] All three AMD EPYC processors we tested leverage AMD Infinity Fabric interconnect and offer security features with AMD Infinity Guard.

About Red Hat OpenShift

Red Hat OpenShift is an enterprise-grade, commercial distribution of the open-source Kubernetes container orchestration platform. We used OpenShift Virtualization, which is a feature that allows IT teams to run VMs alongside containers on the same platform, which could simplify management and improve time to production. Containerization with OpenShift can also enable organizations to enhance application reliability and reduce maintenance complexity, potentially resulting in lower support costs. To learn more about Red Hat OpenShift, visit https://www.redhat.com/en/technologies/cloud-computing/openshift.

How we tested

Understanding how a generation of processors could impact any organization's operations requires more than just a single view of performance. To help paint that picture, we crafted three comparison scenarios of CPU. The section "Our three test scenarios" explains our reasoning for choosing different processors and the detailed differences between them. In each scenario, we used a different processor from the 4th Generation AMD EPYC 9004 Series. In two scenarios, we compared the AMD EPYC processor-based clusters to a cluster backed by the same model of 4th Gen Intel Xeon Scalable processor. In the last scenario, we chose a second 4th Gen Intel Xeon Scalable processor for the comparison cluster.

In each test scenario, we used two input/output (I/O) profiles. Doing so provides further information on how processors can affect performance, because each individual I/O profile has different characteristics and offers different I/O operation patterns. We first ran a 100 percent read profile, which is common for data analysis or data warehouse workloads. These can sometimes be part of online analytical processing (OLAP) applications or business intelligence (BI) tools. We then ran an 80-20 read-write I/O profile. This read-heavy profile can appear in general use database workloads or in online transaction processing (OLTP) workloads—workloads that require frequent data retrieval with some updating or modification.

We used the in-memory data structure store Redis as a database because it is a versatile platform that has many use cases, such as fraud detection, gaming leaderboards, caching, messaging, and more.[6] Redis performance directly influences application interactions, including data retrieval and balancing the load of backend services through caching. Using Redis' multi-model database capabilities, the platform can also empower developers to construct low-latency microservice architectures that scale based on user demand. And although the general Redis platform works well with a single CPU core, we chose Redis Cluster to scale up our workloads and make use of the multi-core CPUs we tested.

To create the Redis database and workload, we used the benchmarking tool memtier_benchmark. We also captured power metrics during testing using ExTech 380801 power monitors.

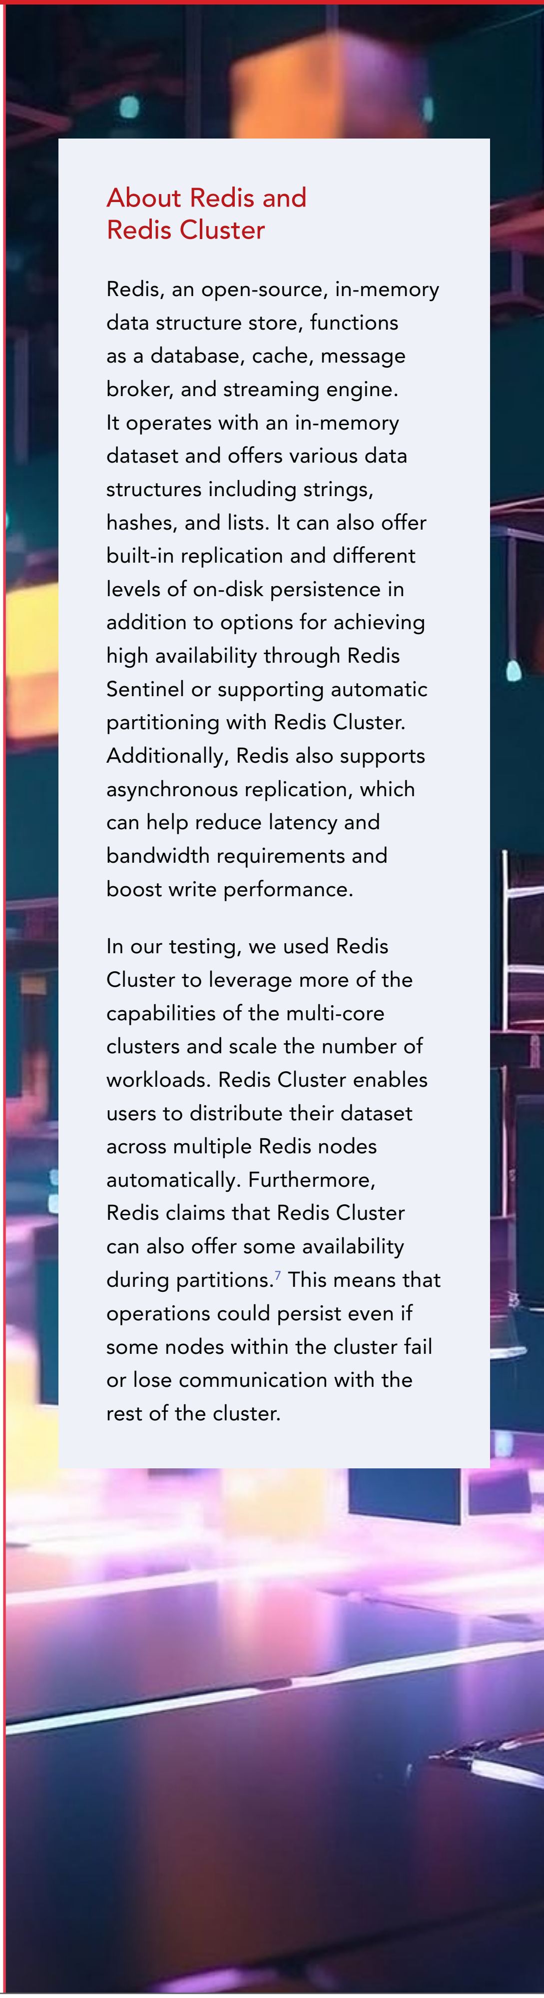

About Redis and Redis Cluster

Redis, an open-source, in-memory data structure store, functions as a database, cache, message broker, and streaming engine. It operates with an in-memory dataset and offers various data structures including strings, hashes, and lists. It can also offer built-in replication and different levels of on-disk persistence in addition to options for achieving high availability through Redis Sentinel or supporting automatic partitioning with Redis Cluster. Additionally, Redis also supports asynchronous replication, which can help reduce latency and bandwidth requirements and boost write performance.

In our testing, we used Redis Cluster to leverage more of the capabilities of the multi-core clusters and scale the number of workloads. Redis Cluster enables users to distribute their dataset across multiple Redis nodes automatically. Furthermore, Redis claims that Redis Cluster can also offer some availability during partitions.[7] This means that operations could persist even if some nodes within the cluster fail or lose communication with the rest of the cluster.

About memtier_benchmark

Redis developed this tool to help users understand the performance of their systems running Redis workloads. According to Redis, memtier_benchmark "provides a robust set of customization and reporting capabilities all wrapped into a convenient and easy-to-use command-line interface."[13] The tool outputs operations per second, latency, and more. To learn more about memtier_benchmark, visit https://redis.com/blog/memtier_benchmark-a-high-throughput-benchmarking-tool-for-redis-memcached/.

Our three test scenarios

Clusters with the closest possible core counts

For the first CPU comparison, we configured a server cluster with 64-core AMD EPYC 9534 processors and a cluster of servers with 56-core Intel Xeon Platinum 8480+ processors. We considered both processors to have the closest core counts we could find, offering similar levels of performance characteristics such as operating speed, power consumption, and more. The Intel processor had a higher max frequency at 3.8 GHz, while the AMD processor had a max frequency of 3.7 GHz. The AMD processor in this configuration typically has a lower retail cost than the Intel processor.[8, 9]

Clusters with processors that typically have similar retail prices

For the second CPU comparison, we configured the AMD EPYC-based server cluster with AMD EPYC 9684X processors. This AMD processor had 96 cores and 192 threads. We did not change the Intel Xeon Platinum 8480+ processor for this comparison. As this AMD processor typically has a retail cost about the same as the Intel processor,[10] the goal for this test scenario was to compare the performance between processors at similar price points.

Clusters with processors that have the same number of cores

For the third CPU comparison, we configured the AMD EPYC-based cluster of servers with AMD EPYC 9174F processors and configured the Intel processor-based server cluster with Intel Xeon Gold 6444Y processors. Both processors had 16 cores and 32 threads. Our aim here was to demonstrate performance when processors from the two OEMs matched on core count. The two processors typically have a similar retail cost, with the Intel processor costing slightly more.[11, 12]

What we found

For each scenario, we gathered the requests per second (RPS) output from memtier_benchmark and the power under load wattage using ExTech 380801 power meters. We then divided the RPS by the wattage to create a performance per watt metric. Table 1 shows the results from each scenario in both I/O profiles.

More RPS could translate to better Redis workload performance, potentially improving response times and the customer experience for the use cases we previously discussed. Using less watts under load could mean spending less on powering your data center, thus helping your departmental bottom line. The performance per watt metric helps quantify the value of each solution's Redis performance.

Table 1: RPS, watts under load, and performance per watt for each cluster configuration in both I/O profiles. For RPS and performance per watt, higher is better. For watts under load, lower is better.

		Intel Xeon Scalable processor-based cluster	**AMD EPYC** processor-based cluster	**Win % for the AMD** processor-based cluster
Test scenario 1				
100% read	RPS	84,355,048	122,389,531	**45%**
	Watts under load	1,002	801	**20%**
	Performance/watt	84,162	152,745	**81%**
80-20 read/write	RPS	76,761,057	112,596,334	**46%**
	Watts under load	1,003	796	**20%**
	Performance/watt	76,499	141,282	**84%**
Test scenario 2				
100% read	RPS	84,355,048	177,763,205	**110%**
	Watts under load	1,002	1,010	**-0.8%**
	Performance/watt	84,162	175,894	**108%**
80-20 read/write	RPS	76,761,057	157,999,207	**105%**
	Watts under load	1,003	1,016	**-1.3%**
	Performance/watt	76,499	155,396	**103%**
Test scenario 3				
100% read	RPS	37,980,229	38,807,375	**2.1%**
	Watts under load	744	709	**4.7%**
	Performance/watt	50,990	54,702	**7.2%**
80-20 read/write	RPS	35,450,081	36,722,284	**3.5%**
	Watts under load	749	710	**5.1%**
	Performance/watt	47,320	51,655	**9.1%**

We assigned each scenario a unique color to help differentiate the comparisons.

Test scenario 1: Clusters with the closest possible core counts

In this scenario, the 64-core AMD EPYC 9534 processor-based cluster supported more RPS than the 56-core Intel Xeon Platinum 8480+ processor-based cluster in both I/O profiles (see Figure 1). Additionally, as Figure 2 shows, the AMD processor-based cluster consumed less power while running the two workloads, resulting in better performance per watt for the AMD processor-backed cluster (see Figure 3). When you consider that the AMD EPYC 9534 processor typically has a lower retail cost than the Intel Xeon Platinum 8480+ processor,[14,15] you could potentially spend less to see better Redis performance while drawing less power.

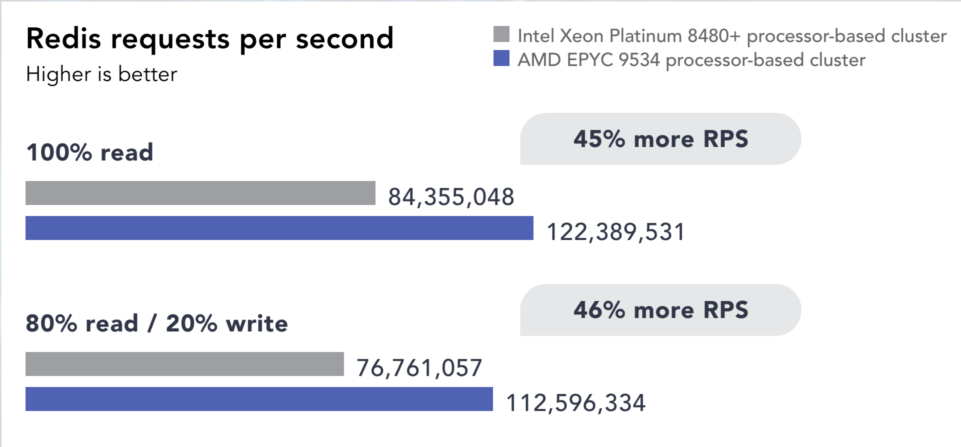

Figure 1: The RPS each solution achieved in the first test scenario under both I/O profiles. Higher is better. Source: Principled Technologies.

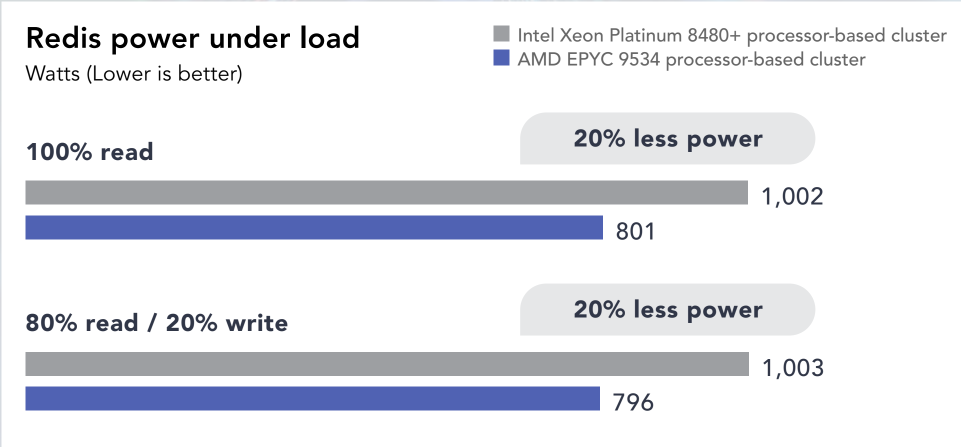

Figure 2: The power each solution consumed while running the workloads in the first scenario. Lower is better. Source: Principled Technologies.

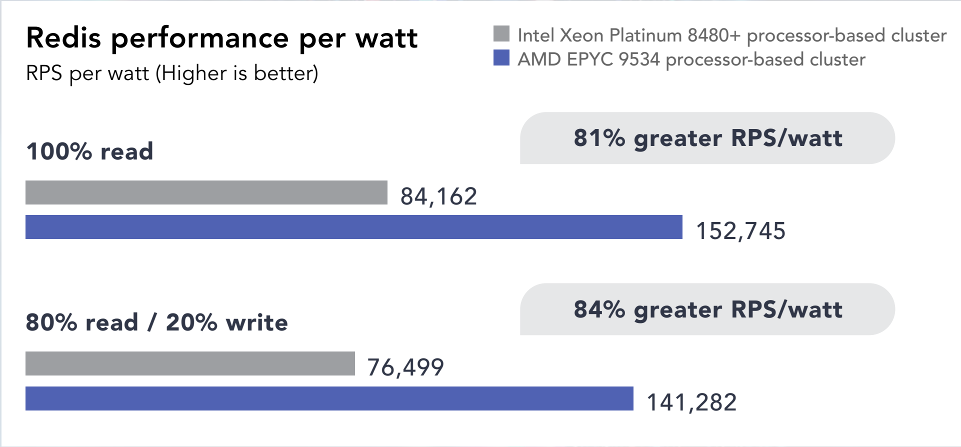

Figure 3: The performance per watt each solution delivered in the first test scenario under both I/O profiles. Higher is better. Source: Principled Technologies.

Test scenario 2: Clusters with processors that typically have similar retail prices

In this scenario, the AMD EPYC 9684X processor-based cluster supported significantly more RPS than the Intel Xeon Platinum 8480+ processor-based cluster in both I/O profiles (see Figure 4). Plus as Figure 5 shows, the AMD processor-based cluster consumed a similar amount of power while running the two workloads. Due to handling more RPS, the AMD processor-backed cluster offered better performance per watt (see Figure 6). Investing in higher-bin AMD EPYC 9684X processors for your servers could yield better Redis performance and thus a better value in terms of performance and power consumption than if you spent a similar amount on Intel Xeon Platinum 8480+ processors for the servers.[16,17]

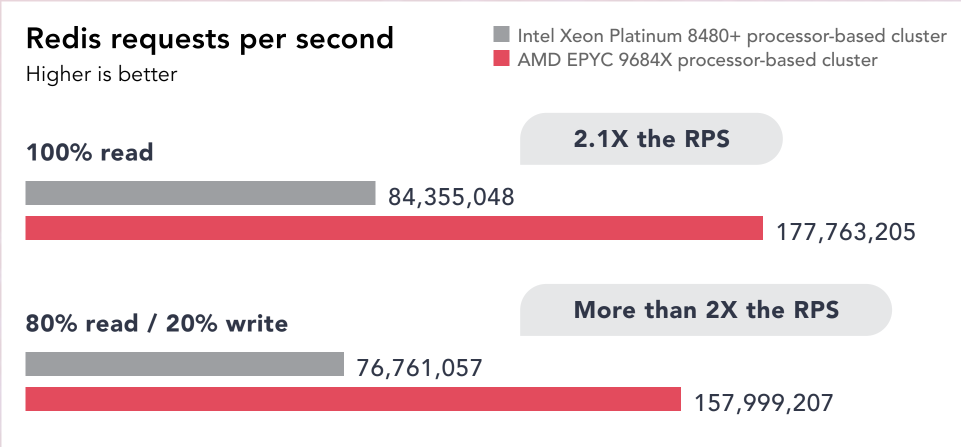

Figure 4: The RPS each solution achieved in the second test scenario under both I/O profiles. Higher is better. Source: Principled Technologies.

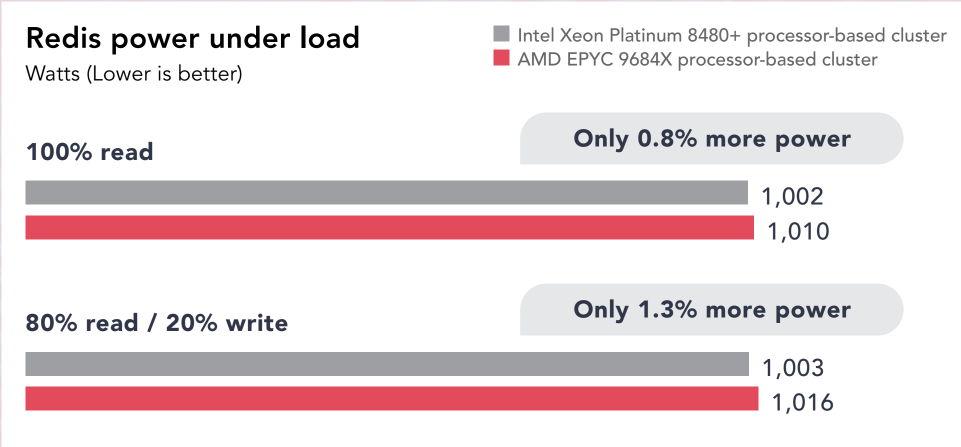

Figure 5: The power each solution consumed while running the workloads in the second scenario. Lower is better. Source: Principled Technologies.

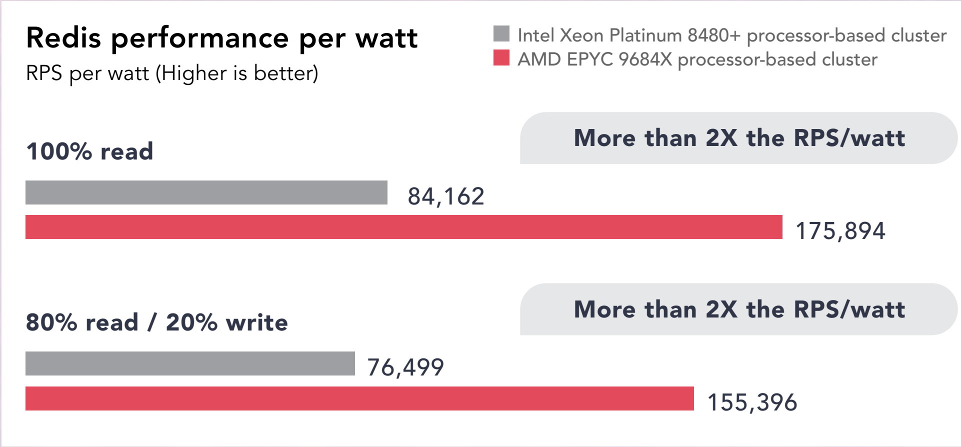

Figure 6: The performance per watt each solution delivered in the second test scenario under both I/O profiles. Higher is better. Source: Principled Technologies.

Test scenario 3: Clusters with processors that have the same number of cores

As the two processors both had 16 cores, the AMD EPYC 9174F processor-based cluster supported better per-core performance than the Intel Xeon Gold 6444Y processor-based cluster. Considering that the cluster with the AMD EPYC 9174F delivered slightly better performance, consumed slightly less power under load, and typically has a per-processor retail cost slightly lower than the Intel Xeon Gold 6444Y processor,[18,19] choosing the AMD processor for your servers could be a smarter investment and give you processors with the same core and thread counts.

Table 2 shows the results for this scenario.

Table 2: RPS, watts under load, and performance per watt for both cluster configuration in both I/O profiles for the third test scenario. For RPS and performance per watt, higher is better. For watts under load, lower is better. Source: Principled Technologies.

		Intel Xeon Gold 6444Y processor-based cluster	**AMD** EPYC 9174F processor-based cluster	**Win %** for the AMD processor-based cluster
Test scenario 3				
100% read	RPS	37,980,229	38,807,375	**2.1%**
	Watts under load	744	709	**4.7%**
	Performance/watt	50,990	54,702	**7.2%**
80-20 read/write	RPS	35,450,081	36,722,284	**3.5%**
	Watts under load	749	710	**5.1%**
	Performance/watt	47,320	51,655	**9.1%**

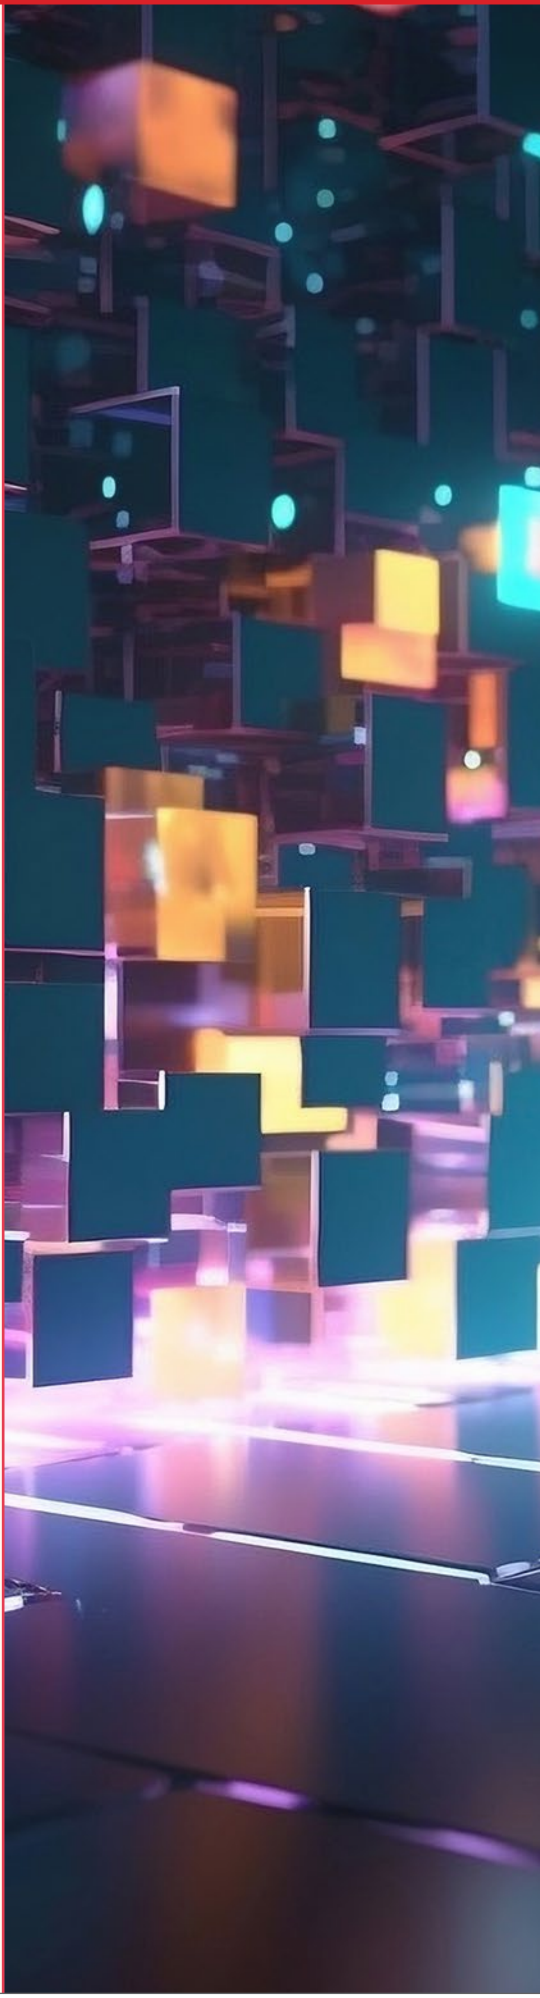

Conclusion

Performance and energy efficiency are significant factors in processor selection for servers running data-intensive workloads, such as Redis. We compared the Redis performance and energy consumption of a server cluster in three AMD EPYC two-processor configurations against that of a server cluster in two Intel Xeon Scalable two-processor configurations. In each of our three test scenarios, the server cluster backed by AMD EPYC processors outperformed the server cluster backed by Intel Xeon Scalable processors. In addition, one of the AMD EPYC processor-based clusters consumed 20 percent less power than its Intel Xeon Scalable processor-based counterpart. Combining these measurements gave us power efficiency metrics that demonstrate how valuable AMD EPYC processor-based servers could be—you could see better performance per watt with these AMD EPYC processor-based server clusters and potentially get more from your Redis or other data intensive applications and workloads while reducing data center power costs.

1. Paul Bubny, “Data Centers See Sharp Rises in Power Costs, Lease Rates,” accessed April 2, 2024, https://www.connectcre.com/stories/data-centers-see-sharp-rises-in-power-costs-lease-rates/.
2. AMD, “AMD EPYC™ 9534,” accessed March 29, 2024, https://www.amd.com/en/products/cpu/amd-epyc-9534.
3. AMD, “AMD EPYC™ 9684X,” accessed March 29, 2024, https://www.amd.com/en/products/cpu/amd-epyc-9684x.
4. AMD, “AMD EPYC™ 9174F,” accessed March 29, 2024, https://www.amd.com/en/products/cpu/amd-epyc-9174f.
5. AMD, “AMD EPYC™ 9004 SERIES PROCESSORS,” accessed March 29, 2024, https://www.amd.com/content/dam/amd/en/documents/products/epyc/epyc-9004-series-processors-data-sheet.pdf.
6. Ajeet Raina, “Redis Use Case Examples for Developers,” accessed April 1, 2024, https://redis.com/blog/5-industry-use-cases-for-redis-developers/.
7. Redis, “Scale with Redis Cluster,” accessed April 4, 2024, https://redis.io/docs/latest/operate/oss_and_stack/management/scaling/.
8. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2064-Core/2.45GHz/AMD/100-100000799WOF_370868.htm).
9. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/Intel%20Xeon%2056-Core/2.0GHz%20-%2016GT%20UPI/INTEL/SRM7H_371122.htm).
10. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2096-Core/2.55GHz/AMD/100-000001254_381212.htm).
11. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2016-core/4.1GHz/AMD/100-000000796_368285.htm).
12. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/Intel%20Xeon%2016-Core/3.6GHz%20-%2016GT%20UPI/INTEL/PK8071305121400_371183.htm).
13. Redis, “memtier_benchmark: A High-Throughput Benchmarking Tool for Redis & Memcached,” accessed April 4, 2024, https://redis.com/blog/memtier_benchmark-a-high-throughput-benchmarking-tool-for-redis-memcached/.
14. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2064-Core/2.45GHz/AMD/100-100000799WOF_370868.htm).
15. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/Intel%20Xeon%2056-Core/2.0GHz%20-%2016GT%20UPI/INTEL/SRM7H_371122.htm).
16. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2096-Core/2.55GHz/AMD/100-000001254_381212.htm).
17. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/Intel%20Xeon%2056-Core/2.0GHz%20-%2016GT%20UPI/INTEL/SRM7H_371122.htm).
18. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/AMD%20EPYC%2016-core/4.1GHz/AMD/100-000000796_368285.htm).
19. Price captured from ServerSupply on April 4, 2024 (https://www.serversupply.com/PROCESSORS/Intel%20Xeon%2016-Core/3.6GHz%20-%2016GT%20UPI/INTEL/PK8071305121400_371183.htm).

Read the science behind this report at https://facts.pt/XbgM65a ▶

Principled Technologies®

Facts matter.®

This project was commissioned by AMD.