Get competitive logistic regression performance with servers with AMD EPYC™ 75F3 processors

Compared to servers powered by 3rd Gen Intel Xeon Platinum 8380 processors, a cluster of 3rd Gen AMD EPYC 75F3 processor-based servers offered comparable logistic regression performance at a lower solution cost to process more data per hour for each dollar spent

Data analytics can be instrumental for organizations when targeting their marketing, grouping their inventory, or detecting security breaches. To run these compute-intensive workloads, such as logistic regression (LR), organizations can benefit from server solutions that can quickly and efficiently process data. Choosing 3rd Gen AMD EPYC™ processor-based servers could deliver fast analysis at a better value for your organization.

We compared the LR performance of a cluster of four servers, each powered by two 32-core 3rd Gen AMD EPYC 75F3 processors, to a second cluster of four servers, each powered by two 40-core 3rd Gen Intel® Xeon® Platinum 8380 processors. Both clusters ran VMware® vSphere® 7.0 Update 3d and used VMware vSAN™. In hands-on testing with the big data benchmark suite HiBench, the two clusters completed an LR workload targeting an Apache Spark™ analytics engine in nearly the same time. The two solutions also offered comparable throughput in terms of a MB per hour rate, but the AMD EPYC processor-based solution offered lower hardware, software, and support costs.

Because it delivered the comparable performance at a lower price, the AMD EPYC processor-based cluster processed more MB per hour per dollar. Organizations that choose these servers with AMD EPYC 75F3 processors for LR clustering workloads could see performance like that of the Intel Xeon Platinum 8380 processor-based cluster but at a lower CapEx, for a better overall value.

Process 72.6% more MB/hour per dollar

Based on LR performance and solution cost*

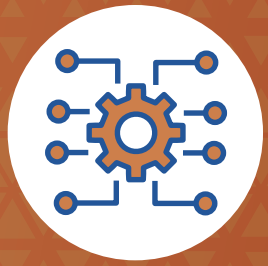

Spend 42.9% less

On hardware and support*

**on the AMD EPYC 75F3 processor-powered servers we tested vs. the Intel Xeon Platinum 8380 processor-powered servers we tested*

How we approached testing

We compared the following four-node VMware vSAN clusters:

- Supermicro® AS-1124US-TNRP servers powered by dual AMD EPYC 75F3 processors
  - For one server, the total cost of hardware plus three years of support and labor and a one-year warranty was $21,305.90—a total of $85,223.60 for a four-node cluster[1]
- Supermicro SYS-620U-TNR servers powered by dual Intel Xeon Platinum 8380 processors
  - For one server, the total cost of hardware plus three years of labor and support and a one-year warranty was $30,338.40—a total of $121,353.60 for a four-node cluster[2]

Other than the processors, we configured the server clusters identically. Each of the servers in both clusters had one 240GB 6Gbps SATA SSD to use for the hypervisor and four PCIe 4.0 NVMe SSDs for the vSAN storage. We also equipped each server with 1,024 GB of PC4-3200 RAM across 16 memory modules.

During testing, we first ran an LR clustering workload using HiBench, targeting a Spark database on each server under test. We recorded the total time each configuration took to complete the workload, as reported by Spark. We then divided the dataset size (52 GB) by that time to calculate the throughput of megabytes per second (MB/s) that each solution processed. Using the MB per second throughput rates, we calculated the average throughput in MB per hour (MB/hr) for simpler reporting. For more details about our configurations, testing methodology, and CPU utilization, see the science behind the report.

About AMD EPYC 75F3 processors

Part of the third-generation of EPYC 7003 Series Processors, the EPYC 75F3 has 32 cores and 64 threads of computing power. According to AMD, the processor features PCI Express® 4.0 I/O connectivity and supports up to eight DDR4 memory channels per socket.[3] The third generation of AMD EPYC processors can also offer AMD Infinity Guard security features, such as Secure Encrypted Virtualization (SEV), Secure Nested Paging (SEV-SNP), Secure Memory Encryption (SME), and more.[4]

To learn more about the EPYC 75F3 processor, visit https://www.amd.com/en/products/cpu/amd-epyc-75f3.

Comparing times to insight

To conduct our data analytics testing, we used a HiBench LR clustering workload. LR is a machine learning algorithm that uses a decision tree to predict a categorical response and sort data into similar groups or clusters. When a solution can complete an LR workload quickly, it can deliver insights fast, empowering decision-makers to act quickly.

The 3rd Gen AMD EPYC 75F3 processor-powered cluster completed the workload in 1,676 seconds. This was comparable to the cluster with 3rd Gen Intel Xeon Platinum 8380 processors, which took 1,650 seconds to finish the workload, or 1.5 percent less time (see Figure 1).

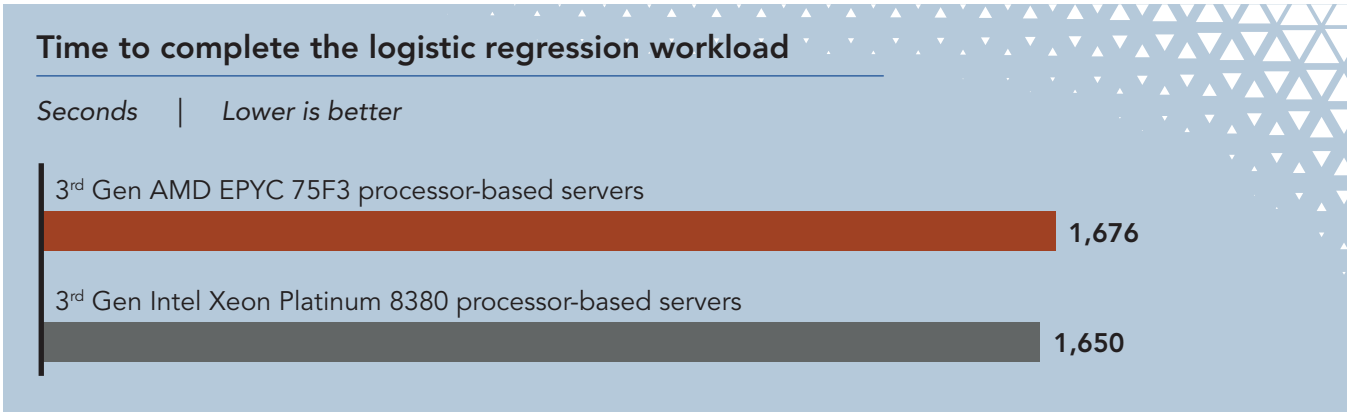

Figure 1: Time in seconds to complete the HiBench LR workload on a 52GB dataset. Less time is better. Source: Principled Technologies.

Real-world benefits for finance

For financial organizations, using Spark to perform data analytics can help pinpoint credit fraud, identify if a customer is in danger of default (default propensity modeling), personalize their marketing, and more.[5,6,7] Accomplishing these goals quickly could provide these organizations with enough time to make well-informed decisions, communicate with clients, and keep assets secure.

With the comparable throughput from the AMD EPYC 75F3 processor-based servers we tested, organizations could categorize that data in about the same amount of time as they could with the Intel Xeon Platinum 8380 processor-based servers we tested. But you don't have to work at a financial institution to understand that the 42.9 percent lower cost of the AMD EPYC processor-based cluster and 72.6 percent better performance per dollar can make it a worthwhile investment.

Comparable throughput while processing data

Dividing the size of the dataset (52 GB) by the time each cluster took to finish the LR workload, we arrived at each cluster's throughput (the rate of MB per second each solution processed). The Intel Xeon Platinum 8380 processor-powered cluster handled 32.3 MB per second while the AMD EPYC 75F3 processor-powered cluster handled 31.8 MB per second—meaning that the two clusters delivered comparable throughput, with the Intel Xeon processor-based cluster handling 1.5 percent more data per hour. Both solutions needed less than a half-hour to complete the workload.

We then calculated a MB per hour rate for the purpose of visualizing the throughput data (Figure 2) and for calculating the performance-to-cost ratio. (Note: The MB per second and MB per hour calculations represent the rate of processing the dataset, rather than disk or networking throughput.)

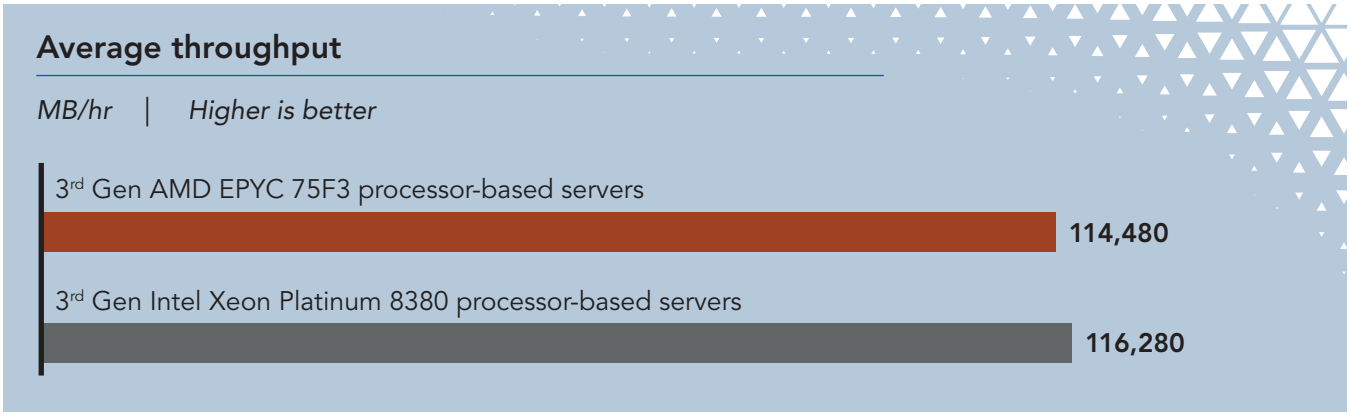

Figure 2: Average data processing rate (MB per hour) during the HiBench LR workload. Higher is better. Source: Principled Technologies.

About VMware vSAN

For organizations looking to reduce the complexity and footprint of their data centers, hyperconverged infrastructure (HCI) can help. As part of their HCI portfolio, VMware offers software-defined storage with vSAN that eliminates the need for bulky, expensive, external arrays and instead brings compute and storage resources together. According to VMware, vSAN is "an enterprise-class storage virtualization software that provides the easiest path to HCI and hybrid cloud."[8]

To learn more about VMware vSAN, visit https://www.vmware.com/products/vsan.html.

Get better performance per dollar

To arrive at our performance-to-cost ratio, we first determined the list prices of hardware with three years of support for both solutions, which we show on page 2. The costs in Figure 3 include the hardware, three-year support, and VMware software licensing. The AMD EPYC processor-based solution cost 42.9 percent less than the Intel Xeon processor-based solution.

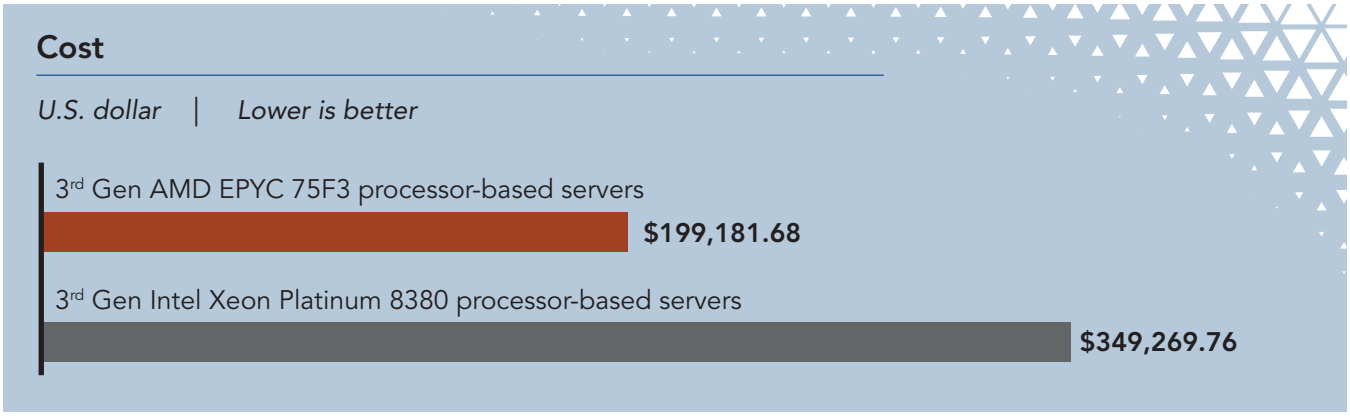

Figure 3: Total cost in USD of the hardware, three-year support, and software for the two solutions we tested. Lower is better.
Source: Principled Technologies.

We then divided the processing rate (MB/hr) by the system cost to arrive at a performance-to-cost ratio. The AMD EPYC processor-based solution handled 0.57 MB per hour for each dollar of the list price, while the solution with Intel Xeon processors offered a ratio of 0.33 MB per hour for each dollar (Figure 4). This means that for every dollar an organization spends on hardware, three-year support, and software for the four-node cluster with eight 3rd Gen AMD EPYC 75F3 processors, they could get 72.6 percent higher LR application throughput than each dollar spent on the four-node cluster with eight 3rd Gen Intel Xeon Platinum 8380 processors.

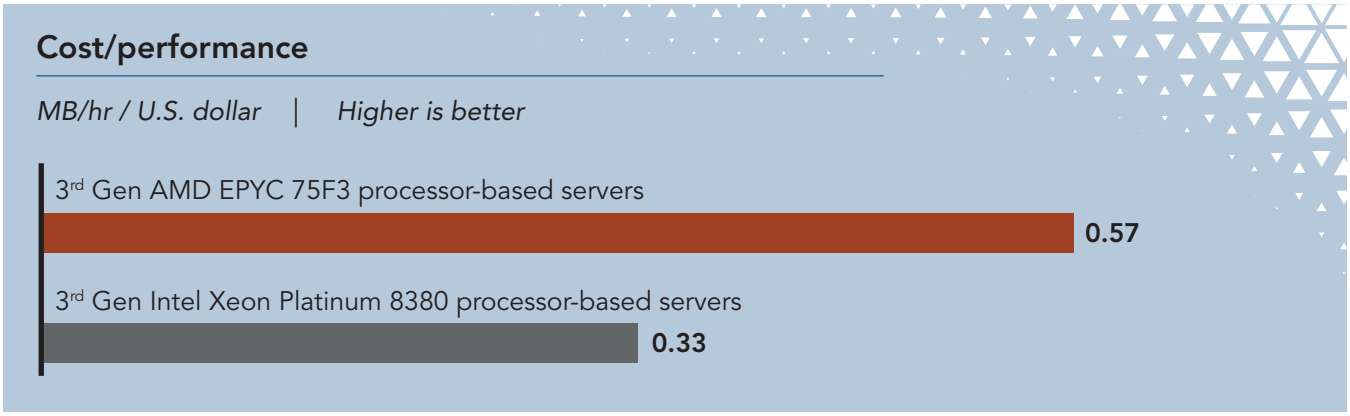

Figure 4: The MB processed per hour per dollar for both solutions we tested with the LR workload. Higher is better.
Source: Principled Technologies.

Conclusion

Servers that can run LR machine learning algorithms quickly could help deliver important insights for organizations that rely on data analytics. When we compared the LR clustering performance of a cluster with 3rd Gen AMD EPYC 75F3 processors to that of a cluster with 3rd Gen Intel Xeon Platinum 8380 processors, we found their performance to be comparable (less than 2 percent difference in runtime and throughput). However, our price analysis shows that the hardware, three-year support, and software costs of the AMD EPYC 75F3 processor-powered Supermicro AS-1124US-TNRP server cluster was 42.9 percent less than that of the Intel Xeon Platinum 8380 processor-powered Supermicro SYS-620U-TNR server cluster. Using the comparable performance and lower solution cost, the AMD EPYC 75F3 processor-based cluster offered 72.6 better LR performance per dollar spent.

---

1. We received a quote from Supermicro on February 2, 2022 for the hardware and support cost of the server minus drive costs. To arrive at the total cost, we added this amount to a drive cost quote we had received from Supermicro on August 9, 2021.
2. We received a quote from Supermicro on February 2, 2022 for the hardware and support cost of the server minus drive costs. To arrive at the total cost, we added this amount to a drive cost quote we had received from Supermicro on August 9, 2021.
3. AMD, "AMD EPYC™ 75F3," accessed May 11, 2022, https://www.amd.com/en/products/cpu/amd-epyc-75f3.
4. Server OEMs and cloud providers must enable AMD Infinity Guard features for use. In addition, security features can vary by EPYC processor generations. Learn more about Infinity Guard at https://www.amd.com/en/technologies/infinity-guard.
5. Chris D'Agostino, "Credit Fraud Prevention with Spark and Graph Analysis." Spark Summit, June 11, 2016. YouTube video, 18:59, accessed May 11, 2022, https://youtu.be/0VO-ts0dsbI.
6. Level Up Education, "How are Big Companies using Apache Spark," accessed May 11, 2022, https://medium.com/tao_66792/how-are-big-companies-using-apache-spark-413743dbbbae.
7. Lawton, George, "Logistic regression," accessed May 11, 2022, https://www.techtarget.com/searchbusinessanalytics/definition/logistic-regression.
8. VMware, "What is vSAN?" accessed May 11, 2022, https://www.vmware.com/products/vsan.html.

This project was commissioned by AMD.

Read the science behind this report at https://facts.pt/QH4ttw4 ▶

Principled Technologies®

Facts matter.®