



Prepare images in Kubernetes for machine learning faster with a Dell EMC cluster powered by AMD EPYC 7543 processors

A cluster of Dell EMC PowerEdge R7525 servers backed by a Dell EMC PowerStore 5000T storage array performed better with 3rd Gen AMD EPYC 7543 processors than an identical cluster with 2nd Gen AMD EPYC 7532 processors

Companies using Kubernetes can benefit from data center solutions that provide high performance on containerized applications. To test the performance an organization could expect from Dell EMC™ PowerEdge™ R7525 servers, we set up a three-server cluster with a Dell EMC PowerStore 5000T array in two configurations: one in which 3rd Gen AMD EPYC™ 7543 processors powered the servers and one in which 2nd Gen AMD EPYC 7532 processors powered the servers.

We set up a VMware environment on both clusters using Update 2 of VMware® vSphere® 7.0 with VMware Tanzu™ Kubernetes Grid (TKG) Service, which integrates Kubernetes support directly into vSphere. We designed and ran a synthetic workload on Kubernetes containers that emulated image processing tasks that are typical at the beginning of an image-based machine learning workflow. The Dell EMC PowerEdge R7525 server cluster with the 3rd Gen AMD EPYC processors handled 19,668 images in 13.9 percent less time than the server cluster with the 2nd Gen AMD EPYC processors, delivering a 13.1 percent higher frames per second (FPS) rate. Additionally, when we compared the price of the hardware and support for each cluster, we found that they cost the same. Organizations running containerized workloads on Kubernetes with TKG could get more image processing work done with the performance benefits of Dell EMC PowerStore 5000T storage-backed PowerEdge R7525 servers powered by 3rd Gen AMD EPYC 7543 processors.



Process over 19K images in 13.9% less time*



Handle 13.1% more frames per second*



Better performance for the same price*†

*vs. the same server cluster with AMD EPYC 7532 processors †Based on the total cost of hardware and support

How we approached testing

In our data center, we compared the following clusters:

- A three-node cluster with Dell EMC PowerEdge R7525 servers, 3rd Gen AMD EPYC 7543 processors, and a Dell EMC PowerStore 5000T storage array. As of April 26, 2021, the price for hardware plus Basic Next Business Day support (36 months) for one server with 3rd Gen processors was \$35,718.01, or \$107,154.03 for a three-server cluster.¹ The price of the Dell EMC PowerStore 5000T with ProSupport and Next Business Day Onsite Service (36 months) was \$425,865.93.² Adding these prices together, we found that the total cost of this cluster configuration came to \$958,885.89.
- A three-node cluster with Dell EMC PowerEdge R7525 servers, 2nd Gen AMD EPYC 7532 processors, and a Dell EMC PowerStore 5000T storage array. As of April 26, 2021, the price for hardware plus Basic Next Business Day support (36 months) for one server with 2nd Gen processors was \$35,718.01, or \$107,154.03 for a three-server cluster.³ The price of the Dell EMC PowerStore 5000T with ProSupport and Next Business Day Onsite Service (36 months) was \$425,865.93.⁴ Adding these prices together, we found that the total cost of this cluster configuration came to \$958,885.89.

Aside from the processors we tested, we configured the server clusters identically, equipping each with two 480GB 6Gbps SATA M.2 SSDs, 512 GB of PC4-3200 RAM, a dual-port 1Gb Ethernet adapter, and a dual-port Fibre Channel adapter for storage connectivity. We configured and created a 2TB volume on the Dell EMC PowerStore 5000T storage system, and then mapped the volume to all PowerEdge R7525 servers as a shared datastore for TKG deployment. The vSphere with TKG cluster in our testing had 44 guaranteed-large worker nodes and a pod on each node with 3 CPU threads and 10 GB of memory reserved. For more details about our configurations and testing methodologies, see the [science behind the report](#).

About AMD EPYC 7543 processors

These 32-core processors use AMD Infinity Architecture and are part of the AMD EPYC 7003 Series. The latest offering from AMD, 3rd Gen EPYC processors offer increased I/O with up to 32MB L3 cache per core, 7nm x86 hybrid die core, and new security features such as Secure Encrypted Virtualization - Secure Nested Paging (SEV-SNP) and Encrypted State (SEV-ES).⁵ According to AMD, the EPYC 7543 model is well suited for workloads such as analytics, CAE/CFD/FEA, and media streaming.⁶

Learn more at <https://www.amd.com/en/processors/epyc-7003-series>.



Less time to process images

Before organizations can use their image-based machine learning algorithms to analyze data, they must first prepare the data. The workload we used mimicked simple image processing tasks that a company might run on Kubernetes with TKG during this preparation phase. With a solution that can sort and prepare images faster, organizations working in areas such as manufacturing, healthcare, and security could reach insights sooner. (See [page 5](#) for more real-world benefits.) As Figure 1 shows, Dell EMC PowerEdge R7525 servers powered by 3rd Gen AMD EPYC 7543 processors and backed by the Dell EMC PowerStore 5000T array took 36 minutes and 41 seconds to process 19,668 images—13.9 percent less time than the same servers and storage array with 2nd Gen AMD EPYC 7432 processors, which took 42 minutes and 37 seconds.




Figure 1: Time to process a set of 19,668 images from the synthetic benchmark we used (minutes:seconds). Lower is better. Source: Principled Technologies.



About Dell EMC PowerEdge R7525 servers

According to Dell Technologies, these servers offer the following specifications:⁷

- Up to 128 high-performance 3rd Gen AMD EPYC cores in two sockets
- Support for 24 Gen4 NVMe™ direct connections
- 32 DDR4 RDIMM or LRDIMM memory module slots
- Up to eight PCIe® Gen4 slots
- Automated server life cycle management
- OpenManage Integration for VMware vSphere, which delivers server life cycle support in VMware vCenter® environments with hardware monitoring and alerts, VMware ESXi™ deployment, and firmware management⁸

To learn more, visit <https://www.dell.com/en-us/work/shop/povw/poweredge-r7525>.

About the Dell EMC PowerStore 5000T

During testing, we backed each cluster configuration with a Dell EMC PowerStore 5000T storage array. According to Dell Technologies, “PowerStore T models provide organizations with all the benefits of an enterprise unified storage platform for block, file and vVol data, while enabling flexible growth with the intelligent scale-up and scale-out capability of appliance clusters.” The 5000T model offers unified block, file, and VMware vSphere vVols support, 898.56 TB of raw capacity, and 25 NVMe drive slots.⁹

Learn more at <https://www.delltechnologies.com/en-us/storage/powerstore-storage-appliance/powerstore-t-series.htm>.

Because our workload was CPU-intensive, the configurations only minimally utilized the storage array, which indicates that the Dell EMC PowerStore 5000T had the resources available to potentially run other workloads or support more clusters. For more information on how the workload impacted the storage array during testing, see the [science behind the report](#).

A higher FPS processing rate

Figure 2 illustrates the image processing rate each solution achieved in our testing. The cluster with 2nd Gen AMD EPYC 7432 processors and the PowerStore 5000T storage array achieved a rate of 8,263 FPS. Meanwhile, the 3rd Gen AMD EPYC 7543 processor-powered Dell EMC PowerEdge 7525 cluster with the PowerStore 5000T storage array processed the images at a rate of 9,349 FPS—an improvement of 13.1 percent.




Figure 2: The image processing rate for each solution in frames per second (FPS). Higher is better. Source: Principled Technologies.

About VMware vSphere with Tanzu

To run our workload, we used Tanzu Kubernetes Grid. VMware states that TKG is a “CNCF-certified, enterprise-ready Kubernetes runtime that streamlines operations across a multi-cloud infrastructure”¹⁰ that also “provides a consistent, upstream-compatible implementation of Kubernetes, that is tested, signed, and supported by VMware.”¹¹ TKG is part of the VMware Tanzu portfolio, which enables organizations to “build, run and manage modern apps on any cloud—and continuously deliver value to your customers,” as well as to “simplify multi-cloud operations and free developers to move faster with easy access to the right resources,” according to VMware.¹²

The VMware Tanzu portfolio includes products that could enhance a user’s TKG deployment, including the following:

- **Tanzu Mission Control:** From a single pane of glass, Tanzu users can provision and manage their TKG clusters, secure apps and clusters with policy-based security features, and provide self-service Kubernetes to developers¹³
- **Tanzu Service Mesh:** Allows users to secure applications through network encryption and authorization policies as well as auto-scale and load balance applications¹⁴
- **Tanzu Observability:** Offers health and performance monitoring of applications and cloud infrastructure and alerting for on-premises and cloud-based Kubernetes clusters and applications¹⁵

For more information, visit <https://tanzu.vmware.com/tanzu>.




Real-world benefits for retail

For stores that use computer vision to prevent shoplifting, keep their shelves properly stocked, or optimize their marketing, accomplishing tasks in less time during the machine learning workflow could translate to achieving those goals faster. During the image preprocessing phase, which prepares images for machine learning algorithms, your company could save time by running vSphere with TKG workloads on Dell EMC PowerEdge R7525 servers with 3rd Gen AMD EPYC 7543 processors backed by a Dell EMC PowerStore 5000T storage array. We saw a boost in performance when a cluster of these servers processed images in 13.9 percent less time—and at a 13.1 percent higher FPS rate—compared to the same servers with 2nd Gen AMD EPYC processors and the same storage array. With this increased performance, your store could get important insights in less time, enabling you to deliver a safe and personalized shopping experience to your customers.

About VMware vSphere 7.0 Update 2

We used VMware vSphere 7.0 Update 2 for our test environment. According to VMware, this vSphere 7.0 update includes a CPU scheduler that is optimized for AMD EPYC processors, as it is “designed to take advantage of the multiple last-level caches (LLCs) per CPU socket offered by the AMD EPYC processors.” VMware also states that vSphere 7.0 Update 2 with the CPU scheduler can achieve near-optimal performance on most applications and benchmarks on AMD EPYC processors.¹⁶

For more information, visit <https://www.vmware.com/products/vsphere.html>.



Conclusion

In a vSphere 7.0 Update 2 environment with Tanzu Kubernetes Grid, we compared the image processing performance of Kubernetes workloads on two cluster configurations of Dell EMC PowerEdge R7525 servers: one cluster with 3rd Gen AMD EPYC 7543 processors and the other cluster with 2nd Gen AMD EPYC 7532 processors. We backed both clusters with a Dell EMC PowerStore 5000T storage array, which showed minimal utilization during testing. We found that the cluster with the 3rd Gen AMD EPYC processors handled 13.1 percent more frames per second, processing 19,668 images in 13.9 percent less time than the cluster with the 2nd Gen AMD EPYC processors. These results indicate that without paying an increased price for newer hardware, an organization using vSphere with TKG to run their Kubernetes workloads could get better image processing performance by choosing Dell EMC PowerEdge R7525 servers powered by 3rd Gen AMD EPYC 7543 processors. By backing the servers with a Dell EMC PowerStore 5000T, they could also get additional storage support to run other workloads.

-
- 1 "PowerEdge R7525 Rack Server," accessed April 26, 2021, <https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r7525/>.
 - 2 On January 26, 2021, we received a quote from a third party reseller for the Dell EMC PowerStore 5000T.
 - 3 "PowerEdge R7525 Rack Server," accessed April 26, 2021, <https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r7525/>.
 - 4 On January 26, 2021, we received a quote from a third party reseller for the Dell EMC PowerStore 5000T.
 - 5 "AMD EPYC 7003 Series Processors," accessed April 26, 2021, <https://www.amd.com/en/processors/epyc-7003-series>.
 - 6 "AMD EPYC 7543," accessed April 26, 2021, <https://www.amd.com/en/products/cpu/amd-epyc-7543>.
 - 7 "OpenManage Integration for VMware vCenter," accessed April 30, 2021, <https://www.dell.com/support/kbdoc/en-us/000176981/openmanage-integration-for-vmware-vcenter>.
 - 8 "PowerEdge R7525," accessed April 26, 2021, https://i.dell.com/sites/csdocuments/Product_Docs/en/PowerEdge-R7525-Spec-Sheet.pdf.
 - 9 "Dell EMC PowerStore Storage Family," accessed April 26, 2021, <https://www.delltechnologies.com/en-us/collaterals/un-auth/data-sheets/products/storage/h18143-dell-emc-powerstore-family-spec-sheet.pdf>.
 - 10 "VMware Tanzu Kubernetes Grid," accessed May 18, 2021, <https://d1fto35gcffzn.cloudfront.net/tanzu/tkg/TKG-Solution-Overview.pdf>.
 - 11 "VMware Tanzu Kubernetes Grid Documentation," accessed April 26, 2021, <https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/index.html>.
 - 12 "VMware Tanzu," accessed April 26, 2021, <https://d1fto35gcffzn.cloudfront.net/tanzu/VMware-Tanzu-Solution-Brief-0121.pdf>.
 - 13 "Multi-cluster Kubernetes management across clouds," accessed April 30, 2021, <https://tanzu.vmware.com/mission-control>.
 - 14 "Enterprise-class service mesh," accessed April 30, 2021, <https://tanzu.vmware.com/service-mesh>.
 - 15 "Enterprise observability for multi-cloud environments," accessed April 30, 2021, <https://tanzu.vmware.com/observability>.
 - 16 "Performance Optimizations in VMware vSphere 7.0 U2 CPU Scheduler for AMD EPYC Processors," accessed May 17, 2021, <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/performance/vsphere70u2-cpu-sched-amd-epyc.pdf>.

Read the science behind this report at <http://facts.pt/oYvyggs> ►



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.

This project was commissioned by Dell Technologies.