Demartek

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

*Dell single-socket AMD EPYC systems running VMware vSAN provide solid TCO benefits.*


## Executive Summary

Dell EMC provides VMware vSAN Ready Nodes that combine a wide range of Dell hardware and VMware vSAN software into a ready-to-order package. These include the new single-socket Dell EMC PowerEdge R7415 servers powered by the AMD EPYC 7000-series processors. These Dell EMC vSAN Ready Nodes are validated and configured to meet hyper-converged workload needs while taking advantage of ***lower single-socket licensing costs*** from VMware.

Dell commissioned Demartek to evaluate the performance of the new Dell EMC PowerEdge R7415 servers that use the AMD EPYC 7000-series of processors. Mixed enterprise workloads were run in this VMware vSAN cluster using one model of the AMD EPYC processor in each of three identical nodes. Then the tests were repeated using different models of AMD EPYC processors in the vSAN cluster nodes.

AMD EPYC processors have 64 MB of Level 3 processor cache, 8 memory channels per socket, support up to 2 TB of main memory and support 128 PCIe 3.0 lanes.

Each VMware vSAN cluster node was configured with two NVMe SSDs for the caching layer and six SATA SSDs for the capacity layer.

Several enterprise workloads were run simultaneously in virtual machines (VMs) across each cluster, including:

> **Microsoft SQL Server** running an online transaction processing (OLTP) workload

> **Microsoft Exchange Jetstress** running an email server storage stress test

> **Iometer "All-in-one"** workload to stress test the storage devices.

In addition, separate, stand-alone tests were conducted with another database workload known as **DVD Store 2** running in a VM in the vSAN cluster.

## Key Findings

> **TCO** – In addition to reduced hardware costs for single-socket servers, VMware vSphere and vSAN licensing costs are $7610 lower for single-socket servers, per node.

> The vSAN mixed workloads achieved comparable throughput performance results of approximately 400 MB/s using only ***single-socket*** AMD EPYC systems as compared to the previous generation ***dual-socket*** Dell EMC PowerEdge R730 systems.

> For the vSAN DVD Store 2 results the single-socket 24-core and 32-core AMD processors achieved **78%** of current generation traditional dual-socket systems for max operations per minute (OPM).

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## Single-Socket vs. Dual-Socket Systems

For many applications, a single-socket server is a good choice for several reasons. Economically speaking, a single socket system has a **lower cost**. This lower cost can be divided into the **hardware cost** and the **software cost**.

The lower hardware cost is due to single-socket motherboards being simpler to design with fewer components and the result of having to pay for only one processor.

Furthermore, a major chunk of the savings comes from software licensing savings. VMware vSphere licensing is based on the number of CPUs (or "sockets"), and often, there are special bundles for single-socket systems. If a single-socket system with many cores – such as the Dell EMC R7415 server with one 32-core AMD EPYC processor – will satisfy your workload requirements, the software licensing savings will be substantial compared to a traditional architecture dual-socket system.

A second advantage of single-socket systems is **power consumption savings**. Electric power is a limited resource and consuming less power to accomplish a certain amount of work saves on the energy bill. Moreover, lower power consumption translates into lower cooling requirements, providing an incremental savings. These savings factor into operational costs over the long term, separate from any initial capital expense savings.

A technical advantage of a single-socket server is that all the **memory is on the local node**, or CPU. This means that there is no "hop" to access memory assigned to a different processor, since all of the memory in a single-socket system is assigned to the one CPU. For some mission-critical applications, memory latency (round-trip time) can be an important performance consideration. Given that the AMD EPYC processors support up to 2 TB of main memory per processor, a fairly large number of applications may be able to run quite well in this type of single-socket environment.

## Total Cost of Ownership Comparison

Software licensing costs for single-socket systems can provide substantial savings.

The chart below shows the VMware HCI Kit Advanced licensing costs for a one-node and three-node cluster, comparing single-socket servers to dual-socket servers. The VMware HCI Kit includes a license for vSAN and vSphere. For a one-node system, a single-socket server can deliver $7,610 in savings on VMware licensing costs. For a three-node cluster, these licensing cost savings are $22,830.

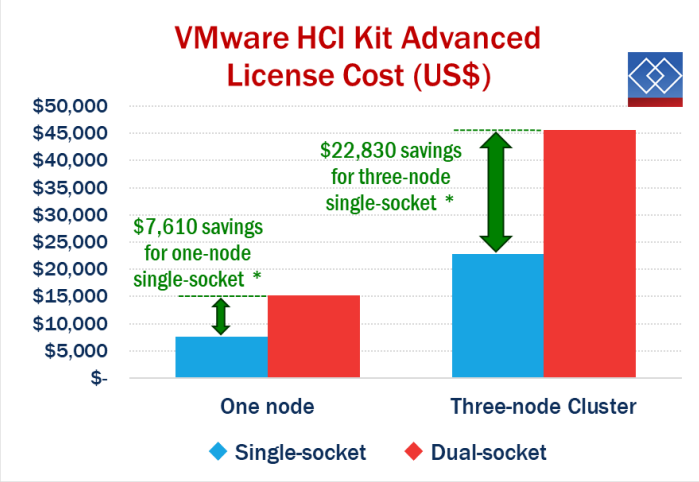


* Based on the VMware HCI Kit cost of $7,610 per CPU, part number HCI-ADV-CPU-C, as of April 2018. This is the list price in US dollars before any volume discounts or promotional discounts are applied.

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## VMware vSAN Configuration

A three-node VMware vSAN cluster was created including the following components:

> Dell EMC PowerEdge R7415 servers with AMD EPYC processors

> 128 GB RAM

> 2x NVMe SSD

> 6x SATA SSD

> VMware vSphere 6.5 with vSAN 6.5 (6.5.0-7515524)

## Processor Models

The tests were run with the same model of processor in each node of the cluster. The tests were repeated with each of the following processors in the vSAN cluster.

> AMD EPYC 7351P, 2.4 GHz, 16c/32t

> AMD EPYC 7451P, 2.3 GHz, 24c/48t

> AMD EPYC 7551P, 2.0 GHz, 32c/64t

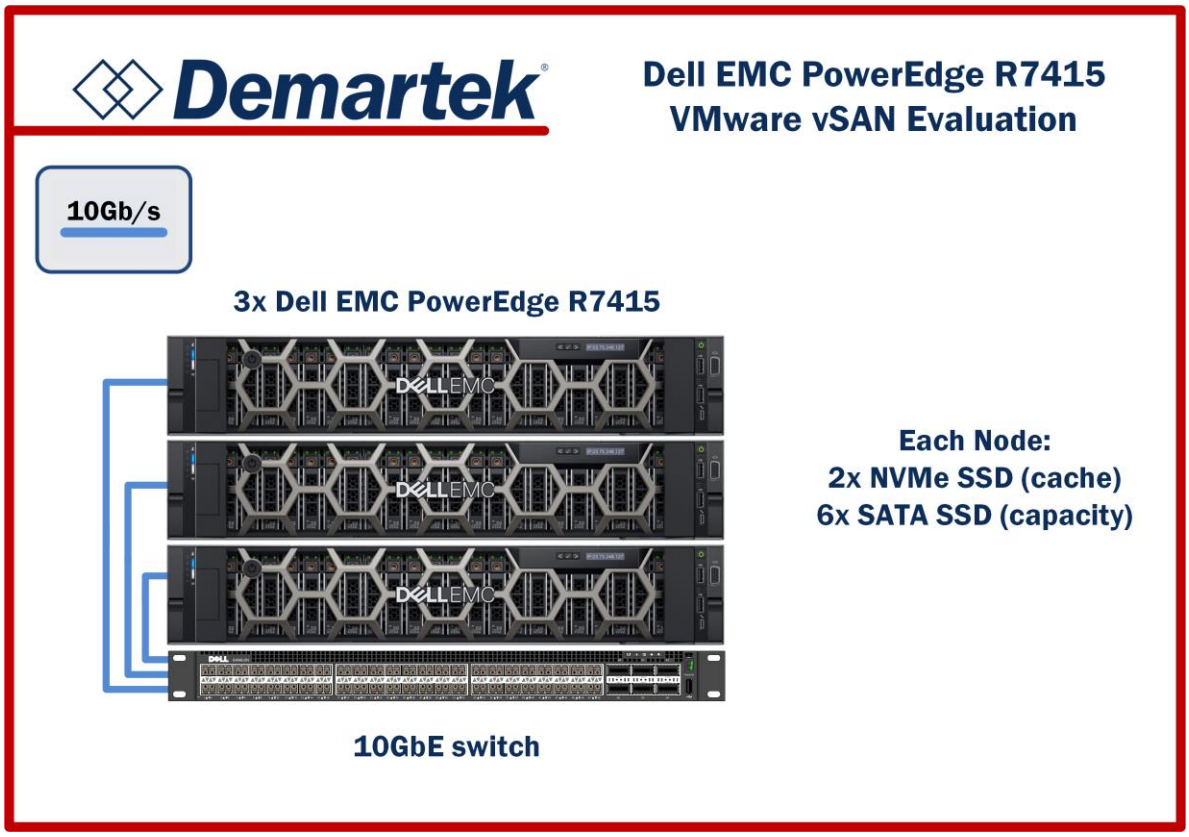

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## Performance Tests

The enterprise workloads included in these tests are typical of those found in enterprise datacenters.

### Microsoft SQL Server OLTP workload

Running in Microsoft SQL Server in a Windows Server guest VM, this workload is a variation of an On-Line Transaction Processing (OLTP) workload that models a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research. This workload is approximately ***90% read and 10% write*** (not counting the log files which are 100% write). The brokerage firm in turn interacts with financial markets to execute orders on behalf of the customers and updates relevant account information.

The benchmark is "scalable," meaning that the number of customers defined for the brokerage firm can be varied to represent the workloads of different-size businesses. The benchmark defines the required mix of transactions the benchmark must maintain.

**Workload Behavior**

- This workload emphasizes ***IOPS*** (smaller block sizes) and has low, relatively steady latency.
- This workload can consume 50% CPU utilization on all or most cores, depending on the test settings.

### Microsoft Jetstress workload

Running Microsoft Exchange Jetstress in a Windows Server guest VM, this workload simulates the Exchange Server email activity by using the same I/O engine that Exchange Server uses.

**Workload Behavior**

- This workload emphasizes steady, somewhat random, mid-range block size (~ 32KB) with approximately 50% read 50% write ratio (± a few percentage points).

### Iometer "All-in-One" Workload

Iometer is an I/O subsystem measurement and characterization tool that is frequently used for synthetic testing of storage devices. Individual I/O parameters such as block size can be used to test "corner cases," or a mix of I/O characteristics can be specified to approximate a more real-world workload.

The "All-in-One" test specification executes a variety of I/O characteristics running concurrently. Six different block sizes (512B, 4KiB, 16KiB, 32KiB, 64KiB and 256KiB) are combined with five different read/write mixes (100% read, 75% read / 25% write, 50% read / 50% write, 25% read / 75% write and 100% write) along with a mix of sequential and random I/O. This results in 29 different I/O configurations, each being executed either 3% or 4% of the time during the test. Alignment is either on sector boundaries or 4K boundaries. This simulates a completely mixed workload and runs at a fairly steady level.

### DVD Store 2 Workload

The DVD Store Version 2 (DS2) is a complete, online e-commerce test application with a back-end database component, a Web application layer, and driver programs. The goal in designing the database component as well as the mid-tier application was to utilize many advanced database features (transactions, stored procedures, triggers, referential integrity) while keeping the database easy to install and understand. The DS2 workload may be used to test databases or as a stress tool for any purpose.

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## Test Results – Mixed Workloads

A series of mixed workloads were run simultaneously on the vSAN cluster with each of the three processor models. The best results were achieved with the single-socket 32-core processor.

Because this report focuses on the processor performance the throughput and IOPS are included here.

The results shown below compare favorably with the previous generation dual-socket Dell EMC PowerEdge R730 results generated last year.

## 32-Core Processor Results – Throughput

The throughput results show that the Exchange Jetstress and Iometer workloads ran at a steady state throughout the tests. The OLTP database workload had a bit more fluctuation but was reasonably steady.

The throughput results are shown for the totals for all three applications and broken out showing the read throughput and write throughput separately.

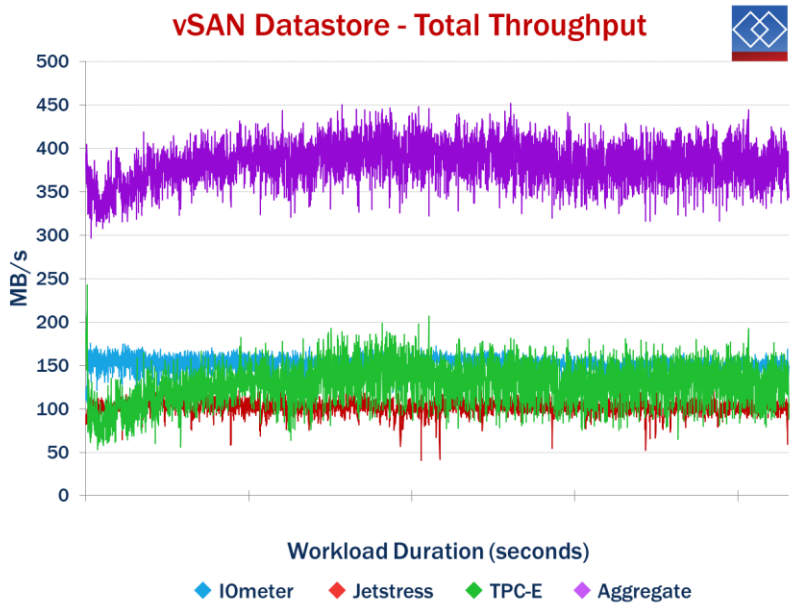


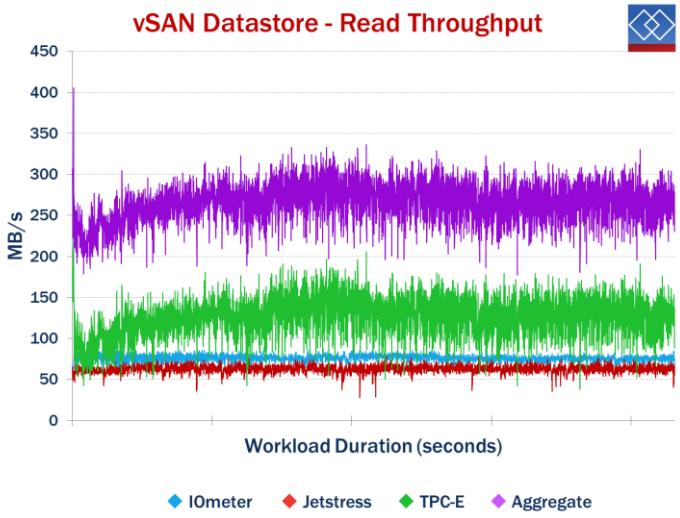


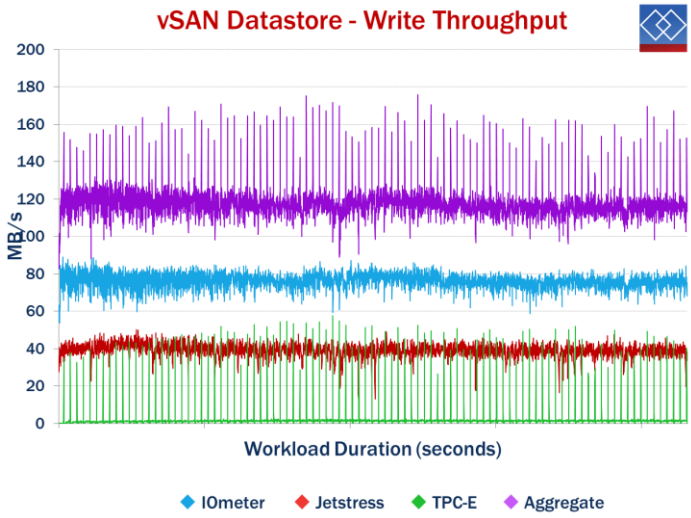

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

## 32-Core Processor Results – IOPS

The Exchange Jetstress and Iometer workloads ran at a steady state in terms of IOPS, meaning that these workloads were highly random in terms of data locality. The OLTP database workload, however, was able to take advantage of the vSAN caching storage tier, showing an increase in performance during the first one-third of the run as the hot data filled the cache.

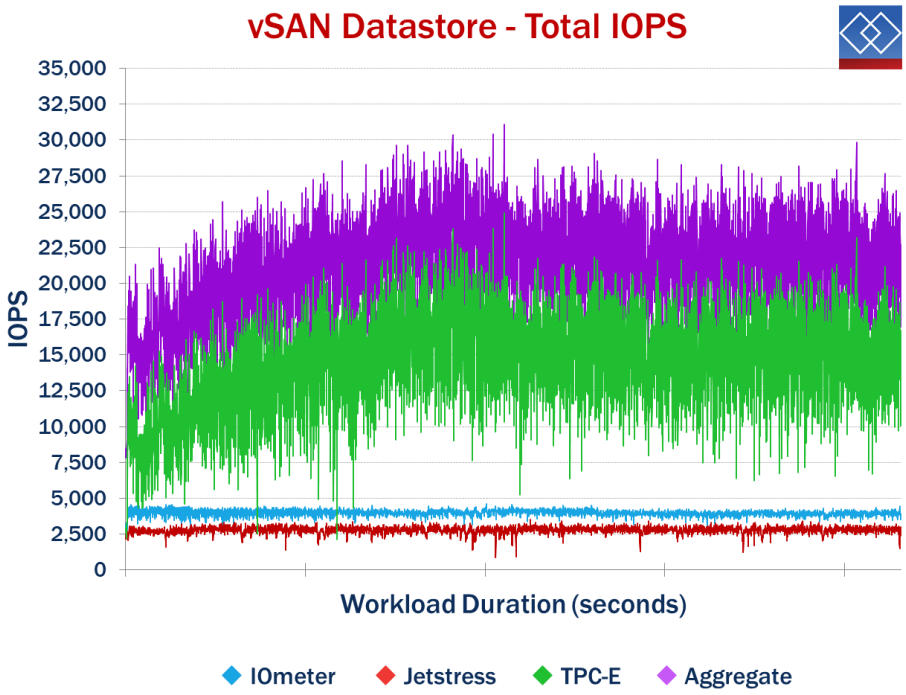


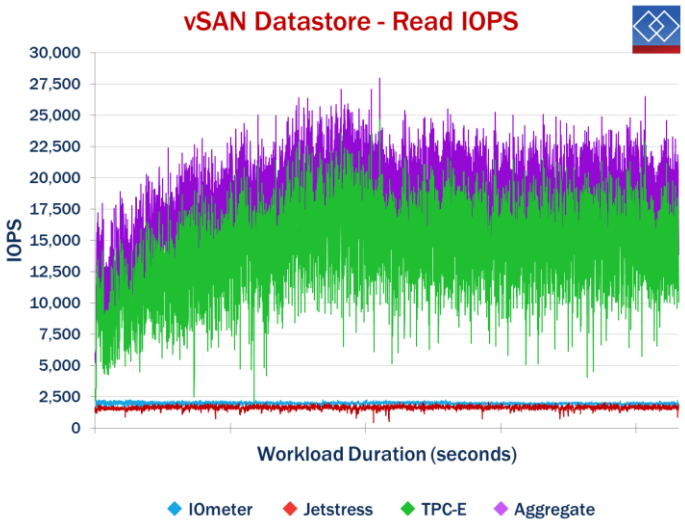


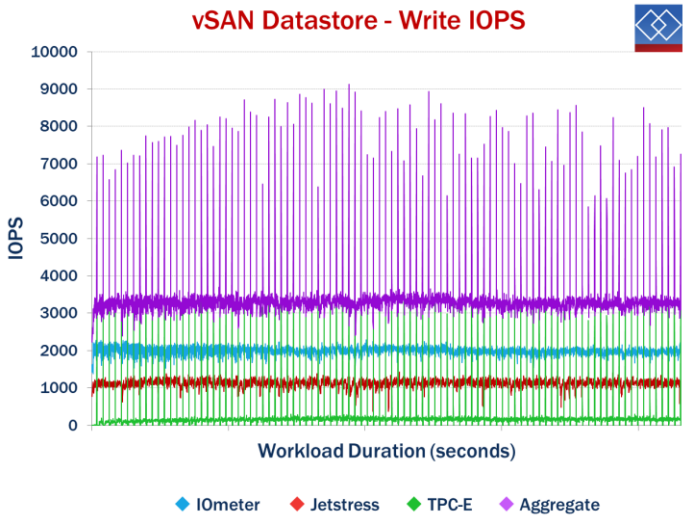

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## Previous Performance Tests

We previously ran a similar, though not identical, mixed workload using **dual-socket Dell EMC PowerEdge R730 servers in a VMware vSAN cluster**. This report was published in September 2017 and is available here: http://www.demartek.com/Demartek_Dell_EMC_vSAN_Toshiba_SSDs_2017-08.html.

These results were achieved for this report with 5x SAS SSDs used as the storage for vSAN. These high-performance SAS SSDs served for both the cache and capacity tiers of the three-node vSAN cluster.

The throughput of these enterprise workloads running on a dual-socket VMware vSAN cluster is similar to the throughput we obtained in our single-socket test. Although the specific mix of application workloads were slightly different and the vSAN storage configuration was updated, our test results show that the single-socket Dell EMC PowerEdge R7415 can achieve similar levels of performance.

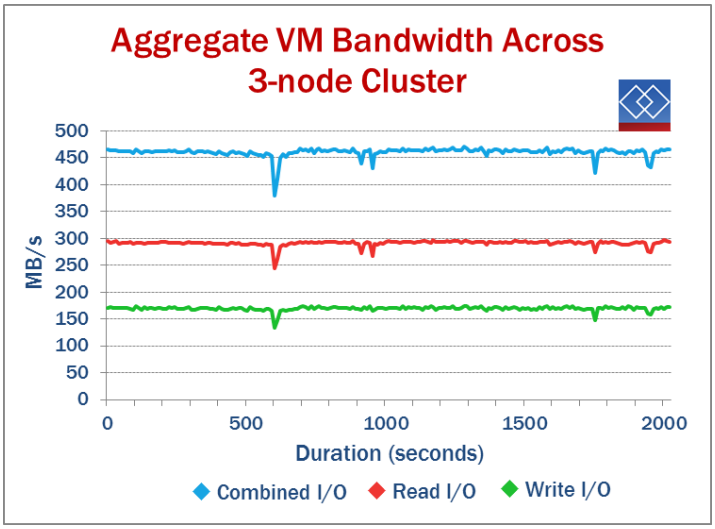


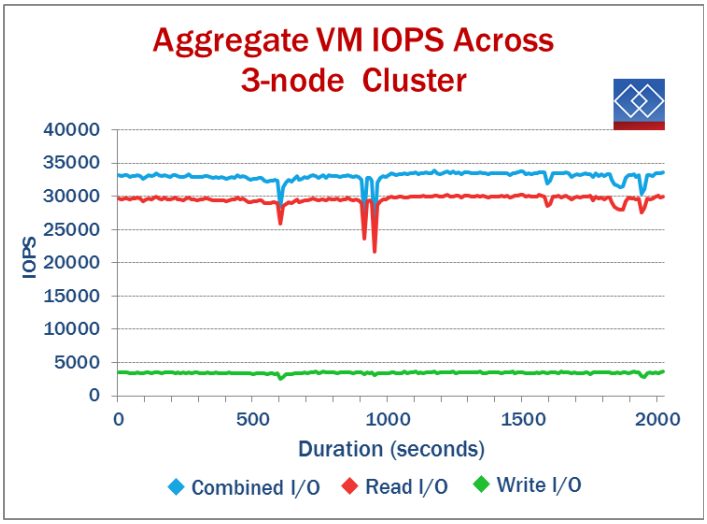

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

## Test Results – DVD Store 2

In reviewing the DVD Store 2 results, there did not appear to be a major difference between the 24-core processor and the 32-core processor. It should be noted that the 24-core processor has a base clock speed of 2.3 GHz. The 32-core processor has a base clock speed of 2.0 GHz.

**DVD Store 2 (DS2) Max Operations Per Minute (OPM)**

| | DS2 Max OPM | Percentage Max OPM |
|---|---|---|
| **Single-socket AMD EPYC 7351P 16-core** | 51844 | **65.8%** |
| **Single-socket AMD EPYC 7451 24-core** | 61662 | **78.3%** |
| **Single-socket AMD EPYC 7551P 32-core** | 61509 | **78.1%** |
| **Current generation traditional dual-socket system** | 78787 | 100% |

# Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance

Demartek
April 2018

## Summary and Conclusion

In our tests, we achieved comparable throughput performance (approx. 400 MB/s) in our mixed-workload tests on the single-socket Dell EMC PowerEdge R7415 running VMware vSAN compared to that obtained in a similar, but not identical, set of tests running on a previous generation dual-socket system (approx. 460 MB/s). In our DS2 tests, we achieved approximately 78% of the performance of a current generation dual-socket system using two different processors in a single-socket configuration.

There are total cost of ownership (TCO) benefits to deploying applications on a single-socket system including VMware vSphere and vSAN licensing savings of $7,610 per node for the single-socket servers. For users running mixed enterprise workloads in VMware vSAN environments, a single-socket system such as the Dell EMC PowerEdge R7415 can be an excellent choice of server.

The most current version of this report is available at http://www.demartek.com/Demartek_Dell_EMC_PowerEdge_R7415_vSAN_Mixed_Workloads_Evaluation_2018-04.html on the Demartek website.



 