Dell Technologies CSP SAP HANA Workload

As clients and industry increase adoption of cloud service providers for critical applications, this paper examines the consistency of Cloud Infrastructure as a Service (IaaS) offerings and compares this to running SAP HANA on Dell EMC PowerEdge server hardware.

Table of Contents

Executive Summary

Dell Technologies and Krystallize Technologies conducted performance testing to evaluate the cost effectiveness and performance capability of a Dell Technologies PowerEdge Server Solution for an SAP HANA database relative to two major cloud service providers (CSPs). A standard methodology was used to ensure reproducible measures.

The results show that the Dell EMC PowerEdge servers show positive comparative results over CSP 1 and CSP 2, as illustrated in the graphic to the right and

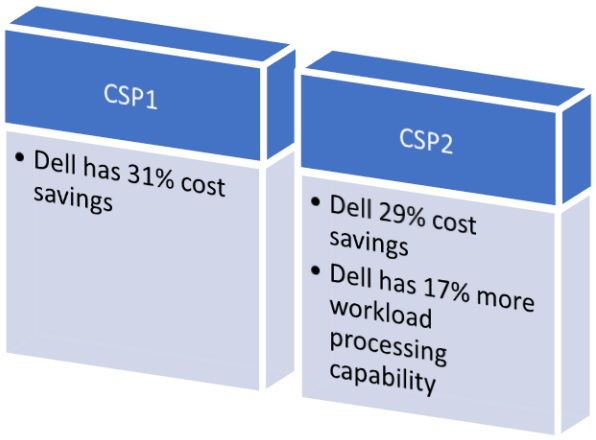

Figure 1.

Figure 1: Cumulative Cost for Platform

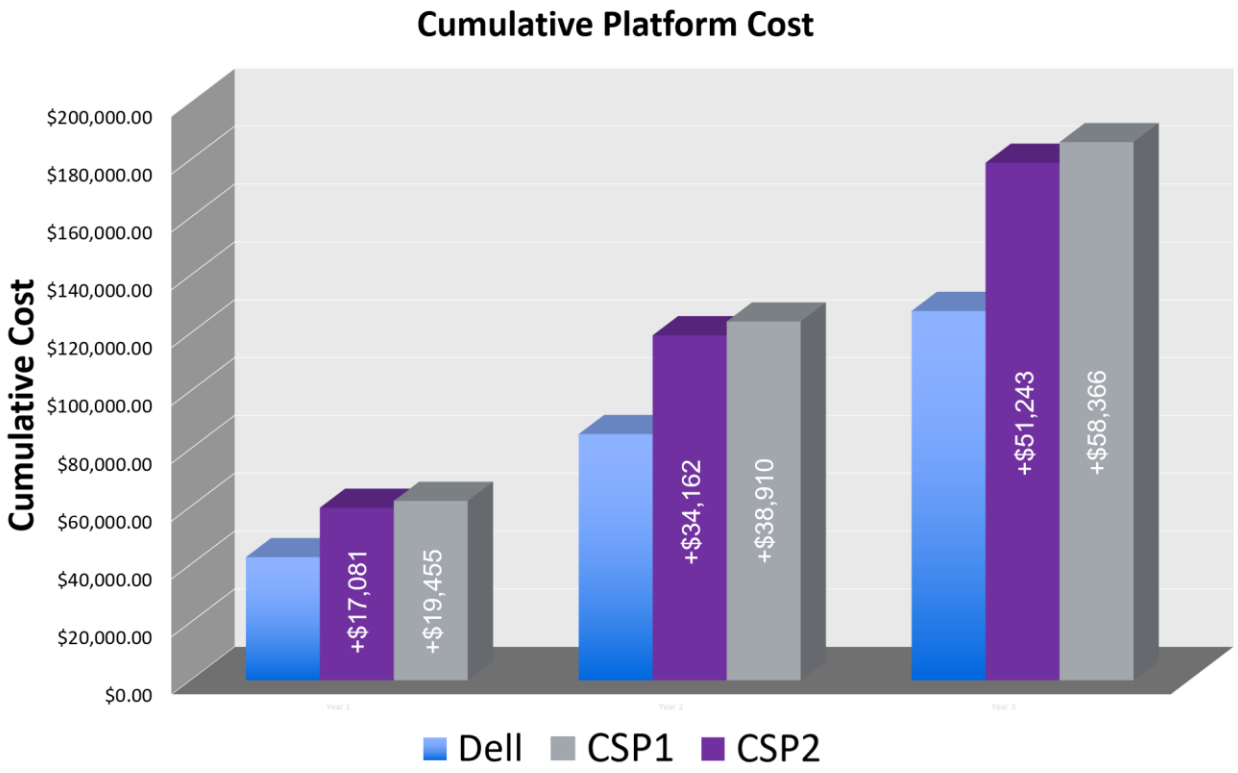

1. Introduction

This study's objective is to evaluate the overall performance capability of two major cloud service providers (CSPs) relative to a Dell Technologies PowerEdge server solution for an SAP HANA database. As the ubiquity of cloud has increased, so too has the acceptability of running critical system workloads with higher system demand on them. Workloads that traditionally used bespoke hardware and dedicated environments have moved to software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). Diverse organizations of all sizes increasingly migrate workloads to cloud-based data center infrastructures to reduce cost and improve agility, but few CSPs provide the transparency to validate consistent cloud quality of service or performance characteristics.

Krystallize Technologies' Service Capability Index (SCI) addresses this need to track service trends and performance by assessing a CSPs ability to deliver consistent workloads. The index is determined with synthetic workloads tuned to the unique needs of an application. Krystallize Technologies used this technique to evaluate the consistency of service delivery and demand fulfillment of major CSPs using a generic workload for commonly used SKUs. A major finding of this testing was the significant variance of service levels across the CSPs.

Deciding an implementation approach for this study involves a range of technical and strategic considerations. Many, seeking the least disruptive path to preserve their existing investment and the customization of their business processes, are choosing on-premise solutions. Taking this into account along with the critical nature of SAP workloads, and the observed variability of numerous CSPs, Krystallize Technologies and Dell Technologies compared the reliability and cost effectiveness of a critical load run on CSPs with the same load run on a dedicated Dell EMC PowerEdge solution.

Krystallize Technologies used steps one (1) through (3) of their CloudQoS™ methodology illustrated in Figure 2 to perform a workload assessment for a set of SAP benchmarks. The Dell Technologies SAP performance lab determined an appropriate set of SAP benchmarks to perform the evaluation.

Figure 2: Krystallize Technologies CloudQoS™ Methodology

In step one (1) a native workload from an existing application is monitored by the Krystallize Technologies SmartAgent as it collects data (the metrics) on usage for each of CPU, memory, network, and storage. In step two (2) these metrics are used to create a synthetic workload that resembles the native workload across the four usage categories—CPU, etc. In step three (3), this synthetic workload is deployed on the potential target platforms to test and assess each platform's price and performance capability.

It is Krystallize Technologies' efficient creation of the synthetic workload that enables a quick and reliable stress tests of alternative cloud and on-premise platforms and a relatively short timeframe for execution of those tests.

The value of having a clear understanding of the performance qualities of the environment are numerous. These include:

- Objectively choosing the right cloud or physical environment for a specific service and workload
- Being able to right-size the environment based on the actual operational metrics
- Avoiding costly and protracted migrations and repatriation in the event the wrong environment is chosen

2. Approach

Dell Technologies invited Krystallize Technologies to evaluate the performance and consistency of SAP workloads on a Dell EMC PowerEdge server and to compare the results with the performance two major CSPs. As described above, Krystallize Technologies leverages synthetic application demand profiles to approximate the demand of client applications as illustrated in Figure 3. Krystallize synthetic benchmarks target sustained peak utilization. Dell setup an initial configuration to create a synthetic workload. With this workload, the Krystallize software synthesized the demand then deployed and evaluated the capability of the different physical or virtual platforms to deliver against that demand. Dell Technologies determined that this process would allow a disciplined approach to measuring and comparing the performance of SAP HANA on physical and virtual environments.

Figure 3: Krystallize Technologies Methodology for Benchmarks

2.1. Workload Identified

Krystallize Technologies used the expertise of Dell Technologies' SAP team to create a baseline benchmark system leveraging common tools. An SAP HANA standalone database was installed on a Dell EMC PowerEdge R940 as listed in Table 1.

Table 1: Source Platform

Vendor	Dell EMC
Model	PowerEdge R940
Number of CPU	4
Number of Cores / CPU	28
Number of Hyperthreads / Core	2
Processor Model	Intel Xeon Scalable Processor 8180M – 2.50 GHz 28 HT Cores
RAM Type and Speed	DDR 2666 MHz 32GB x 24 DIMMS = 768 GB
Storage	Toshiba PX05SMB320Y SSD 3.2 TB x 4
SAP HANA Version	2.00.042.00.1564994110 (fa/HANA2sp04)

For CSPs each virtual CPU (vCPU) represents a hyperthread, so in the case of an EMC Dell PowerEdge server the number of vCPUs present is represented below.

$$vCPU = 4\ Processors \times \frac{28\ Cores}{Processor} \times \frac{2\ Hyperthreads}{Core} = 224\ vCPUs$$

A pair of standard SAP benchmarking tools were used to target a 70% CPU and a consistent storage stressor representing a busy though not overwhelming workload on the platform.

2.2. Metrics Collected

For organizations that have an incomplete understanding of their workload profile, collection times can vary greatly to ensure appropriate data is collected to represent the target load in times of high load. For this exercise, the workload is well known, and the time window was relatively short at one (1) hour. Dell Technologies performed their SAP benchmark and system statistics were recorded from both integrated Dell Remote Access Controller (iDRAC) and Krystallize Technologies SmartAgent 4.2 system statistics collector. The recorded metrics supplied the raw resource usage and performance of the system running SAP HANA.

Three reference runs were performed to ensure that consistent metrics of system demand were collected, as illustrated in Figure 4. It should be noted that the overall storage utilization remained relatively low due to the capabilities of the SSDs used, and as such do not represent a large bottleneck.

Figure 4: Observed Performance

After determining the characteristic workload, that workload was then deployed to the platforms listed in Table 3.

2.3. Demand Profiled

Krystallize Technologies SmartAgent evaluated the system statistics collected to establish a set of target utilization values. In this case, the "profiling" process was performed on the same hardware that the benchmarks were run on to ensure that the stressors best mimic the workload of the application or benchmark being synthesized.

In cases where the same hardware is not available, a surrogate can be identified; but this may increase the time required to baseline and adjust the environment. The Krystallize Technologies SmartAgent runs several scenarios to establish the proper stressor configuration required to mimic the system statistics of the SAP HANA benchmark runs.

The resulting "Synthetic Workload Profile" was created in terms of a set of configuration parameters which were then leveraged to emulate application demand on desired systems. Once this profile was developed, baselines and benchmarks can then be deployed and run on numerous environments without having to install, configure and support SAP HANA. For this exercise, this would include one SKU from CSP 2, one SKU from CSP 1, and a Dell EMC PowerEdge server SKU.

2.4. Baseline Created

The "Synthetic Workload Profile," was then deployed to the original server to create a baseline, which is run over a predetermined time cycle. The primary objective of this exercise is to evaluate service consistency and cost effectiveness over time; and as such the need to observe a continual workload necessitated the baseline be run over a span of time. Because this server was dedicated to the task at hand, many of the variables present in cloud environments were not present. There were no other workloads present to affect service delivery; and there was little to interrupt except for a scheduled test-lab outage that occurred over the weekend. Krystallize Technologies' provided a consistent workload measure, or Service Capability Index (SCI) of 106[1].

2.5. Benchmarks Run

For the purposes of this study, a physical environment and cloud environments that resembled each other as close as possible were chosen. Target selection in the cloud can at times be challenging when attempting to identify equivalents. The original baseline system proved to have a lower Memory to CPU ratio than several the options available, so a target memory of 768 GB was chosen, leveraging the certified SAP platforms by the cloud providers. We attempted to find candidates that matched the baseline system, however, were unable to find an exact match as indicated Table 2.

Table 2: Benchmark Candidate Instances

CSP	Instance	vCPU	Memory (GB)
CSP 1	R5x24.large	96	768
CSP 1	R5 Metal	96	768
CSP 2	M64S	64	1024
CSP 2	M64iS	64	512

[1] The synthetic workload created ran at 79% utilization

To balance the difference in vCPU, the Processors in the Dell EMC PowerEdge R940 were replaced to provide a closer match to the vCPU counts of the CSP 2 and CSP 1 instances. This replacement made it possible to apply the stressor to nearly equivalent machines. Although traditional IT environments tend to look at the makeup of the servers that are providing the service, cloud providers offer a service which is a pre-defined compute shape and is immutable.

To determine the capability of the target platforms being tested, the synthetic workload was pushed to the limits of the environment, which in this case represented 98-99% CPU utilization. Due to the size of the study, we did not require more than one (1) TB for the original test and kept the 1024 GB SSD as a match. A set time over the span of a week was chosen to determine whether there was noticeable change in service over the duration. The environments chosen to be benchmarked are listed in Table 3.

Table 3: Environments Evaluated and Benchmarked

	Dell Technologies	CSP 1	CSP 2
Model	PowerEdge R940	R5.24xlarge	M64s
vCPU	72	96	64
Processor Model	Intel(R) Xeon(R) Gold 6154 CPU @ 3.00GHz	Intel(R) Xeon(R) CPU E7-8890 v3 @ 2.50GHz	Intel(R) Xeon(R) Platinum 8259CL CPU @ 2.50GHz
RAM Type and Speed	DDR 2400MHz 32GB x 24 DIMMS = 768 GB	768 GB Memory	1024 GB Memory
Storage	Toshiba PX05SMB320Y 3.2 TB x 4	G2: 1024 GB	P30: 1024 IOPS limit: 5000 Throughput limit: 200

One of the major challenges with identifying alternatives in the public cloud is the lack of consistent offerings amongst the providers. Dell Technologies offers denser solutions in terms of cores per GB of memory, and our original system configured at 224 cores and 768 GB of memory did not provide a great comparison and the CPUs of the original study were replaced with a lower core count model as illustrated in Table 3.

As the constraint seemed largely around memory, the study evaluated two comparable options available at the time of this study. For CSP 2, this represented the somewhat dated M64s 1TB Storage platform and for CSP 1 the more up-to-date SAP-endorsed R5.24xlarge image. It should be noted that one of the challenges with the cloud providers is that the hardware offered, and the availability of the hardware is at the CSP's discretion and the CSP may use an oversubscription model. This challenge can be region specific.

3. Results

The performance of the various solutions is illustrated by the SCI[2] in Figure 5. In this case Dell Technologies and CSP 1 deliver similar performance in our testing, with an average score of 79. CSP 2 on the other hand trails behind with an average score of 68. In addition, the quality of service with CSP 2 tends to fluctuate notably more than the other two. For critical workloads, this has the potential of decreasing performance consistency and user experience. Figure 5 provides an overview of the scores between the benchmarked systems on a scale of 0-100 kOps/sec.

Figure 5: SCI Benchmark Systems

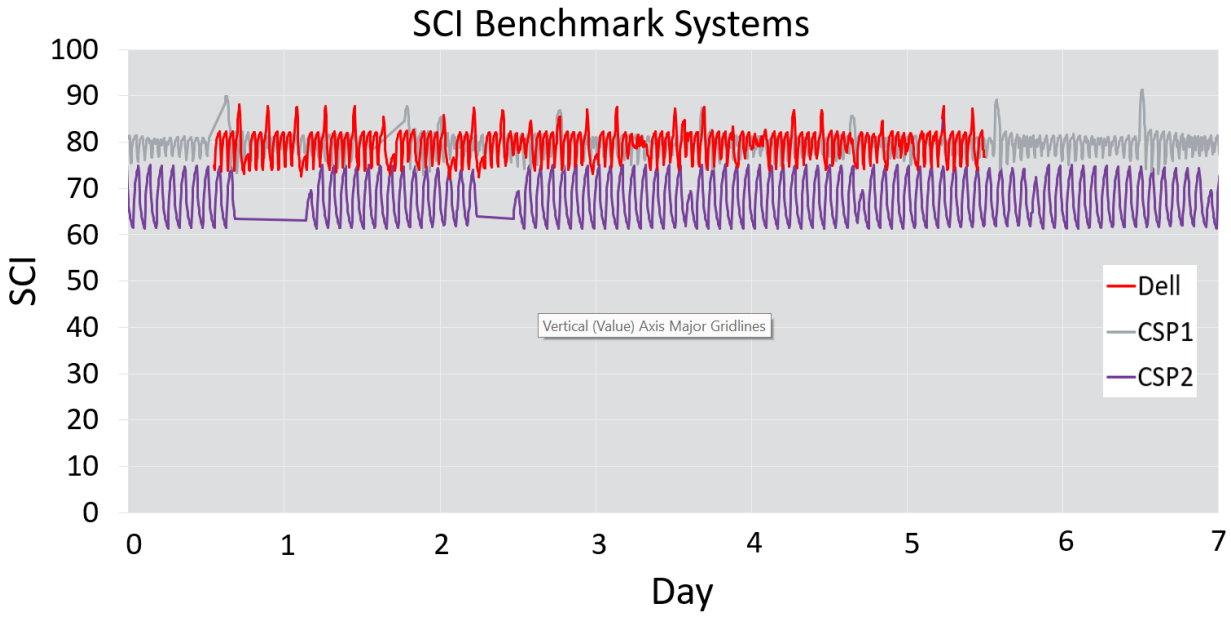

The testing was run for seven (7) days, and aside from two process related interruptions in collection on the first and third day the performance of each environment appeared consistent. Of note in this environment is that Dell Technologies and CSP 1 prove to be similar in terms of performance despite the difference in core count.

Because CSP 1 had more vCPU, we expected higher performance numbers. One reason for this may be due to the decreased processor speed. Using the core speed difference between the two vCPUs, 2.5 GHz for CSP 1 and 3.0 GHz for Dell Technologies and applying the speed ratio, an expected equivalent of CPUs is provided below.

$$CSP\ 1\ vCPU_{Dell} = \frac{2.5\ GHz}{3\ GHz} \times 96\ Cores = 80\ vCPU_{Dell}$$

$$Expected\ CPU\ Output\ = \frac{CSP\ 1\ vCPU_{Dell}}{Dell\ vCPU_{Dell}} = \frac{80 vCPU_{Dell}}{72 vCPU_{Dell}} = 111\%$$

As illustrated above, we expected 11 percent improved performance for the CSP 1 platform, but that is not what was observed.

[2] Service Capability Index, SCI, is the consistent measure which Krystallize Technologies communicates consistent capability measurement.

3.1. Cost Comparison

The following section provides cost and expense details for the different environments. Figure 6 illustrates the cumulative cost breakdown over a three-year period.

Figure 6: Cumulative Cost for Platform

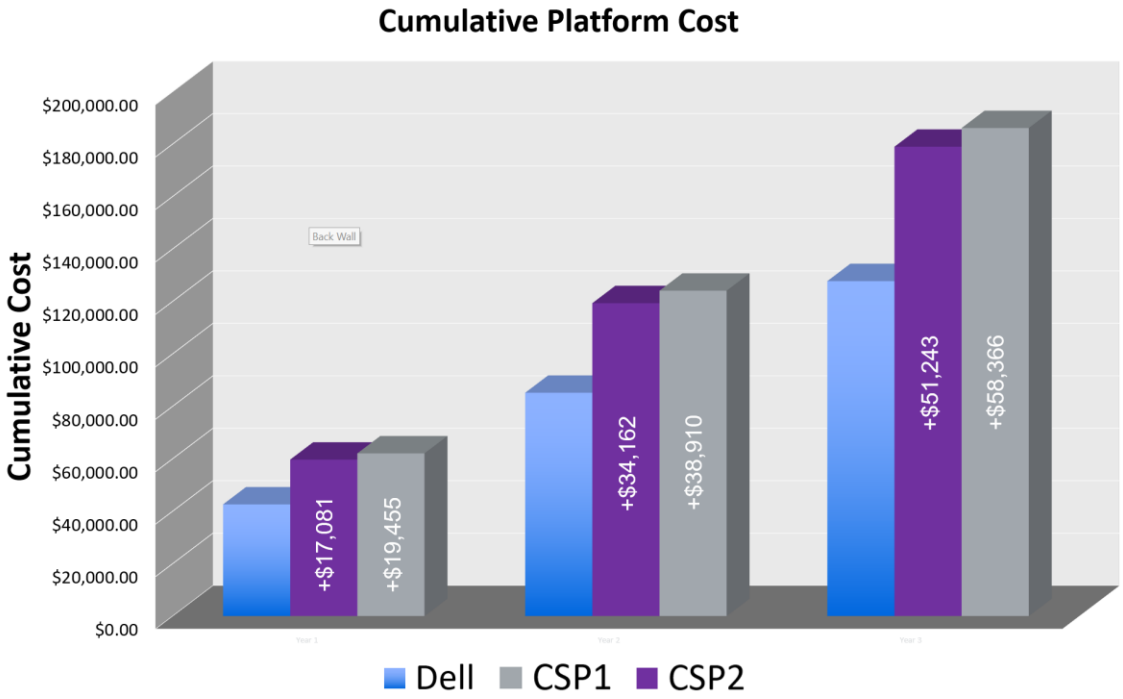

Table 4 provides a breakdown of the costs of the three evaluated systems. For this environment, the Dell EMC PowerEdge server environment is cost adjusted as a 3-year depreciation, with support included. The CSP costs are based on the standard hourly cost obtained on February 19, 2020 multiplied by 24 and 365 to give the yearly cost. On-demand instances will be used in the Western region for both CSP's.

Table 4: Cost Breakdown

	Dell Technologies R940		CSP 1		CSP 2	
	Unit	Annual	Unit	Annual	Unit	Annual
System	$127,702.95	$42,567.65	$6.912/hr	$58,420.44	$6.669/hr	$60,549.12
Storage	Included		$122.88/mo	$1473.60	$102.4/mo	$1228.80
Cost	$42,567.65		$62,022.72		$59,648.80	
3-year cost	$127,702.95		$186,068.16		$178,946.40	
Cost Premium	0%		+31%		+29%	

This indicates a significant cost savings for Dell Technologies over CSP 1 as well as CSP 2.

Note: There can be some variance to the overall costs, average costs, and the compiled costs. With the physical hardware, there will be operational costs for the hosting, which could include electricity, space, and other items. Based on Krystallize Technologies experience, a markup to the numbers in Table 4 for the CSPs can be expected due to other variable factors that were not evaluated in the cost comparison:

- Region of the environment, for the CSP's this can amount to 5 to 10% cost difference
- Use of Reserved Instances or other CSP cost saving methods
- Depreciation versus useful lifetime, most physical environments are depreciated over the 3- or 4-year period, but often for an SAP HANA environment can run for 4 to 8 years, which dramatically changes the per year cost.
- Discount rates on the hardware and software, what is displayed are list prices, the actual negotiated price will usually be lower
- Backups and the operational components of maintaining an on-premise or managed hosting environment, this will usually be company specific
- Data Transfers, public IPs and other required components add incremental costs which are covered in standard on-premise environments can add up to 30% additional cost
- Security monitoring and network

The overall cost can vary based on configuration, but estimated costs included in the above analysis are the direct hard costs of the tested configuration.

4. Conclusion

This paper compares and illustrates the cost-performance of SAP HANA running a benchmark load on physical and cloud provider environments. The Dell EMC PowerEdge server configuration is less expensive and provides better performance for the cost.

This paper provides a detailed repeatable comparison showing the performance and cost comparison of running a large-scale SAP HANA environment on physical hardware either in a hosted facility or in a server room versus running an environment with similar resources in a Cloud Service Provider.

Taking into account the baseline with benchmarks the results were modified to provide an overall score of the efficacy of the environment as illustrated in Figure 7 with three year cost included.

Figure 7: Summary of Scores Between Environments for SAP HANA

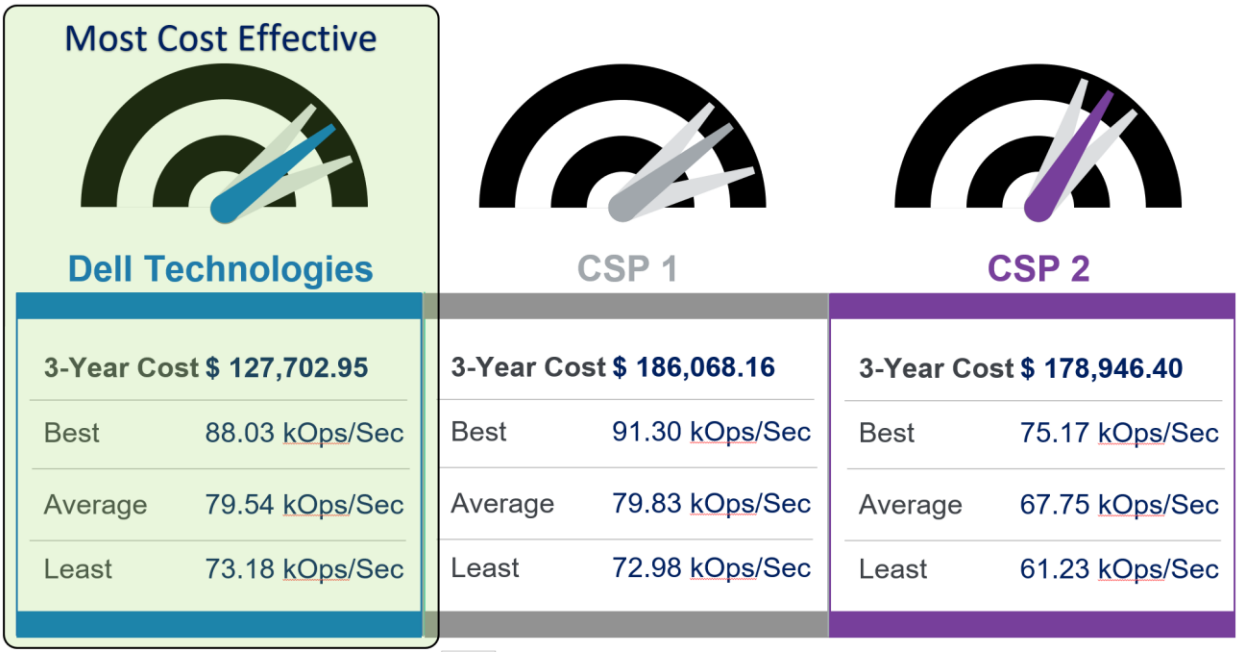

Contrary to the oft cited narrative, CSP hosted platforms are not guaranteed to be more cost effective. The Dell EMC PowerEdge system provides a more cost-effective solution (29% less than CSP 2 and 31% less than CSP 1 over three years) in this comparative study.

This analysis provides users with critical insights into the performance of the three platforms to assist users with determining best-fit/best-cost performance for the SAP HANA application. These results apply to this specific test configuration and are not a comprehensive analysis of the costs of the three platforms in all circumstances.

This document is provided for information purposes only and is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Krystallize and Cloud QoS are trademarks of Krystallize Technologies, Inc. Other names may be trademarks of their respective owners.