OCI VMware Partnership Sets New Cloud Storage Performance Standards

Fall 2024

Written by: Marc Staimer, theCUBEresearch Sr Analyst and President, Dragon Slayer Consulting

Introduction

Many IT organizations have recognized the high-value of expanding VMware application instances into a VMware based public cloud. This shift can significantly improve customer operations by increasing availability, flexibility, and adaptability, while eliminating the need for on-premises hardware expansion and maintenance. The biggest benefit comes from freeing VMware users from adding, installing, maintaining, and updating more hardware infrastructure on-premises.

Most VMware public clouds offer only vSAN storage, typically using dense bare metal servers with NVMe SSDs. These configurations are highly efficient and are used by many vSphere customers. However, they have limitations, especially when scaling clusters or managing storage latency.

An example of these limitations can be found with cluster scaling. For this situation, compute and capacity resources are tightly coupled. Adding capacity requires adding a new node, which increases both compute and capacity, even if only one is needed. Simply put, compute and capacity do not scale separately.

Cross node storage latency is another notable limitation. Accessing storage within the same node offers low latency, cross-node storage access can increase delays, especially with high VM activity.

This research explores how OCI's VMware service using VMFS with Block Volume Storage, addresses these limitations. Specific recommendations are provided based on use cases.

Premises

IT organizations moving their VMware infrastructure to a public cloud generally want as close to a seamless transition as possible. Minimal changes delivers a sense of familiarity, minimizing learning curves. That makes complete sense. It reduces disruptions while quickly getting everything up and running.

As previously mentioned in the introduction, there are several points of friction that cause IT organizations headaches when moving their VMware workloads to a public VMware cloud.

1. Independent Scaling of Compute and Capacity

   - Scalability choices tend to be quite limiting

     In the vast majority of VMware public clouds, when an application requires more capacity but not necessarily more compute or memory, a new node that includes additional compute and memory must be added to the cluster just to add the capacity. That additional compute and memory are not required and are wasted. Analytic applications such as data warehouses is an example of that type of application.

     When an application requires more compute and memory performance but not necessarily more storage capacity, once again a new node must be added to the cluster even though the capacity is not required and is unnecessary. Relational databases, time series databases, even JSON databases are excellent examples of this situation. They may need more compute and memory resources but not necessarily more storage capacity.

     One major public VMware cloud service provider can dedicate nodes in the cluster as storage nodes only. Each storage node cannot run VMs and reduces the number of hyperconverged infrastructure (HCI) nodes by the number of storage nodes[1]. The max number of storage nodes is 50% of the total in the VMware vSphere cluster. Additionally, a VM can't use Vmotion to move to a storage node in the cluster, nor can it be converted to a regular HCI node in the cluster.

   - Inadequate storage capacity scalability

[1] VMware vSphere clusters are limited to 64 nodes

What this means is that a VMware vSphere cluster's capacity cannot exceed the usable sum total of all of the nodes in the cluster, regardless of the type of node – with a maximum of 64 nodes in the cluster. When the VMs in the cluster or an individual VM needs more capacity than is available in a maxed out cluster, the workarounds get complicated. Capacity can be freed up by deleting data. Not a very palatable workaround. Or the cluster can be upgraded with higher capacity drives – a costly and time consuming manual exercise. Also, not very palatable.

- Storage that does not persist beyond the life of the ESXi or individual VMs

VMs fail. ESXi bare metal cluster nodes fail. VMs are deleted. Users accidentally delete VMs as do malicious actors. The workaround is to replicate or mirror data across multiple nodes in the cluster. That consumes minimally 100% additional storage capacity. For protection against two concurrent failures, 200% additional storage capacity is needed. It's a costly workaround. Some will argue that the amount of consumed storage is much less when volumes are configured as RAID 5 or RAID 6. However, neither of which solves the problem when a node fails. That generally means RAID 50 or RAID 60 needs to be implemented, which is mirroring across nodes again.

- High storage latencies and application response times

Storage latencies tend to be higher than external storage. Causes vary. It's often from the fact that vSAN shares the same hardware resources as vSphere, software defined networking, and the VMs. Those resources are limited by the number of nodes in the cluster, cores, memory, NICs, etc. Latency can also be impacted when VMs access volumes in nodes that transits the network. It can even be caused when there are a lot of highly active VM workloads in the cluster. Regardless of the root cause, the increased latency causes slower application response times. Latency is cumulative. Higher latencies cause slower application response times.

There have been several studies on the value of application response time over the years. The first and most important study came from IBM's "The Economic Value of Rapid Response Time". It showed a deep correlation between application response time and user productivity, quality of work, morale, turnover, personnel cost, time-to-actionable-insights, time-to-action, time-to-market, and ultimately time-to-unique-revenues/profits. The numbers are stunning as seen in the transaction rate versus application response time chart below.

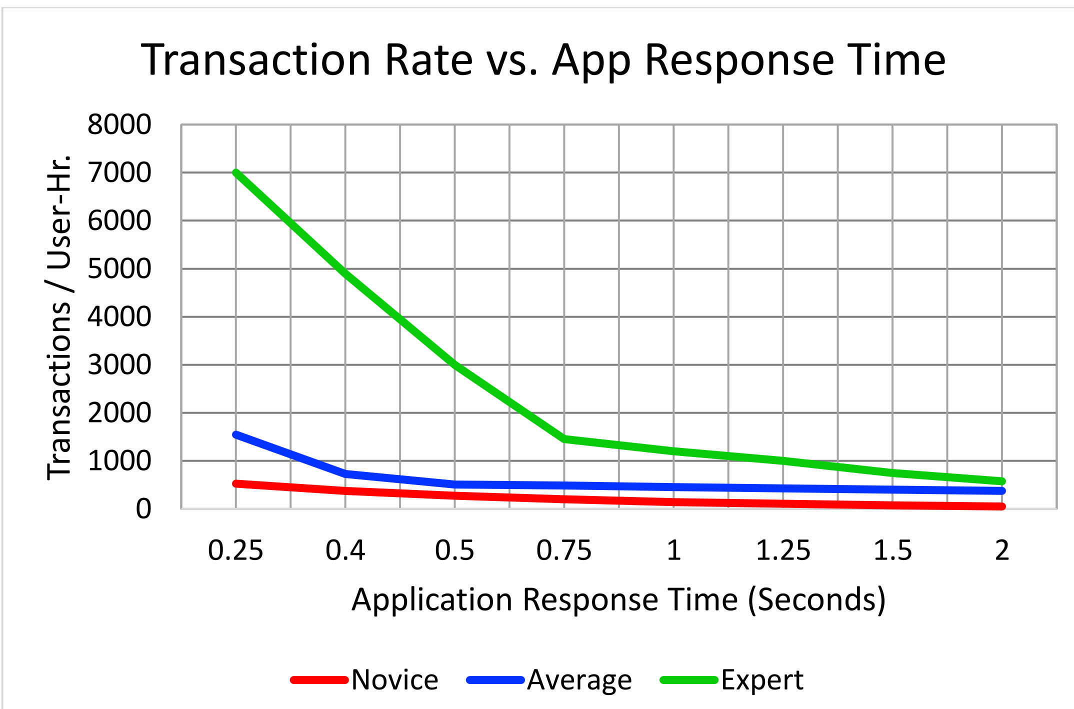

Chart 1: Application Response Time's Effect on User Transaction Rates per IBM's "Economic Value or Rapid Response Times"

The increased productivity from subsecond response times literally jumps off the chart. The research behind it showed adequate user productivity requires a minimum application response time of 2 seconds or less[2]. When response time is more, it falls off a cliff. As application response time gets better, i.e., decreases, productivity begins to increase. It really starts to take off when that response time falls to subsecond. It hits the hockey stick for expert users at ¾ of a

[2] This does not mean the higher storage latencies will cause VM application response times to be greater than 2 seconds, just that it is more likely.

second and 400 ms – .4 seconds – for average and novice users. That 400 ms application response time is also known as the "Doherty Threshold." At the Doherty Threshold or less, users no longer perceive they're waiting on the application. The response time feels as if it is instant. Any response time above the Doherty Threshold enables the user's attention to stray. Application response times at the Doherty Threshold or lower, productivity skyrockets, morale significantly improves, as does quality of work, employee turnover decreases, customer loyalty increases, time-to-actionable insights accelerates, time-to-action quickens, time-to-market shortens, which in turn delivers unique revenues and profits. The savings alone from the increase in productivity can eclipse the total cost of the solution – see more in Appendix A.

These human reactions are based on brain science, a.k.a. neuroscience. It's how human beings work more effectively with their applications. It is not technology specific. The cumulative or total time round trip from their entry or query to when the application kicks back a response. Application response time includes every technology in the path including the server compute and memory hardware, hypervisor, operating system, application network interfaces, network, storage, storage software stack, etc.

A major performance metric that greatly impacts application response time is storage latency. IOPS and throughput also have an impact, just to a lesser extent.

2. Another obstacle occurs when the VMware customer perceives they have already overcome those aforementioned vSAN configuration friction points on-premises. They then realize that moving to a VMware Public cloud will set them back. They were able to solve it on-premises with VMFS and external block storage. Regrettably, that solution is likely not available in their VMware public cloud of choice[3].
3. Last, but certainly not least, is the customer's need to minimize and control their VMware cluster costs in the VMware public cloud. In other words, minimize their cloud spend. This is not just about price, but total cost of ownership (TCO), and performance/TCO.

How Oracle Cloud VMware Solution (OCVS) Overcomes These Points of Friction

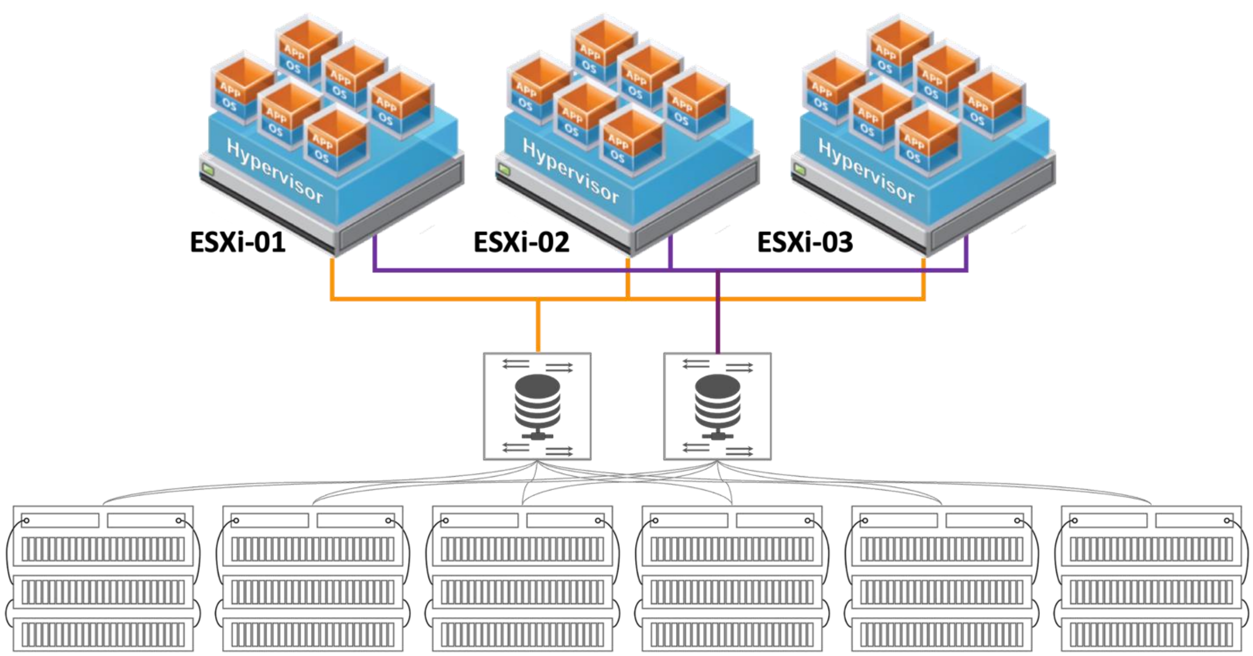

Oracle Cloud VMware Solution (OCVS)

OCVS is a multi-tenant, customer-managed, native VMware-based cloud environment. It is a native VMware ecosystem that provides customers complete control using the same VMware tools they're accustomed to. OCI makes it simple for VMware customers to move or extend their VMware-based workloads to OCI without requiring them to rearchitect their applications or retool their operations.

[3] There may also be differences in CPU to storage ratios in a public VMware cloud than on-premises that might cause additional friction in the cloud.

OCVS provides customers their choice of either Intel or AMD processor-based Compute Shapes. It comes with two distinct types of compute shapes called "Dense I/O" and Standard compute shapes. The Dense I/O shapes are similar to most VMware public clouds in that they use locally-attached NVMe-based SSDs. These shapes are used for VMware vSAN implementations. The OCVS Standard shapes empowers the customer to configure the cores per instance as needed while exclusively attaching to disaggregated as the virtual machine file system (VMFS) datastore.

OCI Block Storage is a highly reliable, high performance, highly redundant, independently scalable, and low-cost affordable block storage for VMFS datastores. It is specifically architected to provide the persistent block storage volumes that live beyond ESXi or a VM's lifespan. It has built-in data redundancy and can scale up to 1 PB per ESXi cluster. That's all good. The question is how OCVS eliminates those aforementioned dense bare metal server vSAN configuration limitations.

Independent Storage Capacity Scaling

OCVS Standard shapes using VMFS and OCI Block Storage enables users to independently scale both compute and storage. When more compute is required, they need only add another bare metal node to the solution cluster up to a maximum of 32 bare metal nodes. Each SDDC supports up to 6 clusters for a total of 192 nodes or hosts.

OCVS in combination with OCI Block Storage frees up ESXi cluster CPU and memory resources that were being consumed by vSAN making them available for VMs and their applications. That does not mean there is no overhead for OCVS management. It ranges up to 20% depending on the customer's configuration.

For the Dense shapes, each SDDC supports up to 6 clusters for a total of 384 bare metal nodes in a cluster.

When additional capacity is required for the OCVS Standard shapes with OCI Block Storage, it's a simple matter to increase the capacity to a volume or add volumes. Each OCVS Standard shape cluster supports up to 32 OCI Block Volumes attachments. Each volume supports up to 32 TBs totaling as much as 1 PB per cluster.

OCVS Standard shapes take advantage of VMware ESXi VMFS that allows multiple bare metal servers to concurrently read and write from the same file system while ensuring each of the VMs' files are locked. VMFS volumes can also be non-destructively scaled up by bonding multiple VMFS volumes together.

OCVS Standard shapes are ideal for those applications demanding high storage capacities but don't necessarily need additional compute or memory resources.

Independent Storage Performance Scaling

OCVS Standard shapes with OCI Block Storage, just like the Dense shapes, are based on high-performance and highly reliable NVMe storage infrastructure. Users can control their performance per volume when they create the volume and anytime thereafter with minimal or no downtime, depending on circumstances[2]. This is tied to a unique OCI Block Storage concept called volume performance units (VPU). VPUs are attainable in increments of 10. There are four performance levels: Lower Cost, Balanced Performance, Higher Performance, and Ultra High Performance as seen in table 1 below.

Elastic Performance Level	Volume Performance Units (VPUs)	IOPS per GB	Max IOPS per Volume	Size for Max IOPS (GB)	KBPS per GB	Max MBPS per Volume
Lower Cost	0	2	3,000	1,500	240	480
Balanced	10	60	25,000	417	480	480
Higher Performance	20	75	50,000	667	600	680
Ultra-High Performance	30	90	75,000	833	720	880
	40	105	100,000	952	840	1,080
	50	120	125,000	1,042	960	1,280
	60	135	150,000	1,111	1,080	1,480
	70	150	175,000	1,167	1,200	1,680
	80	165	200,000	1,212	1,320	1,880
	90	180	225,000	1,250	1,440	2,080
	100	195	250,000	1,282	1,560	2,280
	110	210	275,000	1,310	1,680	2,480
	120	225	300,000	1,333	1,800	2,680

Table 1: OCI Block Storage VPU Scaling

Lower Cost

"Lower cost" OCI Block Storage is not recommended for IOPS intensive VMs. It's best suited for throughput intensive workloads with large sequential I/O, such as video or audio streaming, log processing, time series databases, and data warehouses. The only cost to this option is the storage cost. That's because there are no VPUs and is only available for block volumes.

Balanced

"Balanced" is the default OCI Block Storage performance level. It's used for both block and boot volumes architected to deliver excellent balance between performance and cost for the vast majority of VMware workloads. It comes with 10 VPUs per GB/month.

Higher Performance

"Higher Performance" is generally recommended by OCI for high I/O VMware workloads that don't necessarily require "Ultra High Performance". "Higher Performance" comes with 20 VPUs per GB/month – 2X "Balanced" VPUs.

Ultra-High Performance

When VMware workload I/O demands are at their highest. Comes with 30 – 120 VPUs per GB/month.

Performance	Volume Size	Max Throughput 1 MB block size	Max Throughput 8 KB block size	Max IOPS 4 KB block size
Low Cost	1.44 - 32 TB	480 MB/s	23 MB/s	3000
Balanced	1.44 - 32 TB	480 MB/s	200 MB/s	25000
Higer Performance	1.44 - 32 TB	680 MB/s	680 MB/s	50000
Ultra-High Performance	1.44 - 32 TB	2,680 MB/s	1,350 MB/s	300,000

Table 2: OCI Block Storage Max Performance Comparison

Dynamic Performance Level Scaling

The disaggregated OCI Block Storage permits dynamic performance level scaling configurations for block volumes and boot volumes. That's a huge benefit to the OCVS customers. It means minimal VM application disruptions4 and maximum cost savings. The two forms of dynamic performance scaling are performance based auto-tuning and detached volume auto-tuning.

Dynamic performance scaling dynamically adjusts performance of the volume between customer specified levels of minimum – or default performance – and maximum VPUs/GB. OCI Block Storage monitors volume performance and adjusts based on the volume's throttle operations and guarantees in VPUs/GB, IOPS, and throughput.

Detached volume auto-tuning adjusts and optimizes volume performance based on its attached state. When detached, OCI Block Storage adjusts the performance level to "Lower Cost" (0 VPUs/GB). When reattached it automatically adjusts performance back to the default VPUs/GB setting. This can and does lower storage costs for OCVS users.

Both forms of dynamic performance scaling saves cost without reducing performance. And even more costs can be saved by using them together. One thing to remember about OCI Block Storage SLAs, they apply to "Balanced", "Higher Performance", and "Ultra High Performance" levels, but not "Lower Cost".

What That OCVS Standard Shape OCI Block Storage Independent Scaling Means

- Capacity can be increased without adding addition physical bare metal servers.
- Storage performance can be increased without adding either bare metal servers or additional drives.
- Block volume performance is shared among all connected ESXi hosts.
- VMware cloud spend is minimized without compromising performance.

OCVS Block Storage Volumes That Persist Beyond a VMs Life

Should an ESXi bare metal server node fail, it is a simple matter to reboot all of the VMs on another bare metal server. Block volumes are instantly available. No additional replication is required. No additional storage is consumed. In fact, OCI Block Storage provides automatically multiple copies of data with built-in repair mechanisms built into the service at no additional

Make no mistake, this does not replace backups and/or snapshots. But it does protect the data in the event of a hardware failure.

Faster Application Response Times

As previously discussed, there are several factors that have a major impact on application response times. However, storage latency is by far and away one of the bigger if not the biggest factor.

Many IT storage pros believe VMFS with external storage will have a higher latency than vSAN. Turns out, that's not the case. OCI testing revealed OCVS VMFS latencies can be as much as an order of magnitude (10X) lower than dense vSAN configurations. Distribution balances the workload while lowering latencies.

[4] When a block volume's performance is changed to "Ultra-High Performance" from any other performance level the volume needs to be detached and then reattached.

Consider that delivering a fair apples-to-apples comparison is going to vary for every customer and every configuration. But OCI provided a reasonable comparison test under specific conditions – see Appendix B for test details. The test compared a OCVS dense X7 cluster with 3 hosts equipped with vSAN (122 TB) versus a VMFS datastore in the same cluster configuration using OCI Block Volumes (64 TB spread across 8 volumes of 8 TB each). The tests run were HCI Bench tests using 80 VMs. Below are the results for a 4K blocks with a typical 70% reads with 30% writes test.

	Dense	VMFS	VMFS % faster
Latency (ms)	**12.71**	**1.73**	**86.39%**
Read Latency (ms)	**14.37**	**1.74**	**87.89%**
Write Latency (ms)	**8.89**	**1.73**	**80.54%**
95th Percentile Latency (ms)	**35.51**	**3.13**	**91.19%**

Table 6: vSAN vs OCVS VMFS Latency

The latency is clearly significantly better on VMFS[5] in this configuration and benchmark, primarily because the workload is distributed across multiple volumes. This benefit comes from having a properly designed underlying block volume configuration and is immediately noticeable when overlaying VMs. This is obviously not an absolute. The differences will vary by configuration and workload. However, it is a bit of an eye opener and indicates VMFS with Oracle Block Storage is an outstanding performance alternative.

The OCI Block Storage lower latencies translates into lower application response times. Lower application response times increases user productivity, higher morale, lower turnover, faster time-to-actionable-insights, faster time-to-action, faster time-to-market, and faster time-to-unique-revenues and profits[6].

VMware ESXi Customer Comfortability with OCI Block Storage

Whether the customer is currently using dense bare metal server vSAN configurations on-premises or VMFS with external block storage, OCI is the unique public cloud service provider that provides both. The customer can move to OCI with minimal changes facilitating their move and comfort levels.

OCI Minimizes Their Customer's Cloud Costs

Whether it's compute, storage, or egress fees, previous theCUBEresearch has found OCI to be one of the most cost effective public clouds at minimizing cloud cost. OCVS continues that trend. With independent scaling of compute, storage capacity, and storage performance, with dynamic performance tuning, higher availability with fewer disruptions, distributed shared volumes that reduce latencies, costs are minimized.

The number of bare metal servers required for the ESXi cluster are reduced from not having to run storage on the nodes lowers costs. The elasticity of the storage reduces costs. Higher IOPS and throughput means this is an enterprise ready solution.

Most importantly, those lower latencies translate into lower personnel costs from higher productivity and higher revenues from faster latencies.

Making the Right Choice

Which choice should OCI VMware customers choose for their storage? Dense bare metal server shapes using vSAN configurations or Standard bare metal shapes with VMFS? The dense VMware shapes is generally the comfortable choice. However, when the application is latency sensitive such as eCommerce and transactional databases in general, or needs a higher level of availability, and at the same or even lower cost, OCVS with Oracle Block Storage is the better way to go. It comes down to picking the choice that best meets your requirements. There really is no wrong answer.

Conclusion

Oracle Cloud VMware Solution provides two distinct options for the VMware customer. The first is the Dense shape, which will feel familiar to many VMware customers. It's based on bare metal servers using AMD or Intel sockets and internal NVMe drives with vSAN. The second is unique to OCI. It uses Standard shapes and the unique OCI Block Volume Service. There simply are no comparable disaggregated block volume storage services from other public VMware clouds.

OCVS Standard shapes with OCI Block Storage provides greater storage flexibility, higher performance, and lower latencies. What that translates into is lower customer costs and higher revenues.

[5] The VMFS results are the average mean volume latency of the 8 volumes.

[6] Unique revenues and profits are difficult to estimate because it will vary by customer and industry. However, productivity cost savings can be measured as illustrated in Appendix A.

This means any VMware customer planning on moving their ESXi clusters to a public cloud, should take a long hard look at OCI. Only there can each individual compute of the VMware stack be tuned to meet a range of workload demands. This enables flexible tuning from a small VMware environment to an high performance enterprise grade solution. This is because OCI has changed the VMware public cloud game with its Oracle Cloud VMware Solution.

For More Information

Go to: Oracle Cloud VMware Solution with OCI Block Volumes

Appendix A: Measuring Productivity Cost Savings Based on Faster Application Response Times

Based on IBM's "Economic Value or Rapid Response Times" research, IT organizations have a lot to gain by providing subsecond response times. Especially response times that equal or exceed the Doherty threshold.

IBM's measured calculation productivity gains assumes a generally conservative cost savings of approximately $3,000 USD per month per user at ¼ of a second. That will vary by customer. And the savings will go up the more skilled the users are.

However, as the chart shows, even a small decrease in application response times saves a lot of cost.

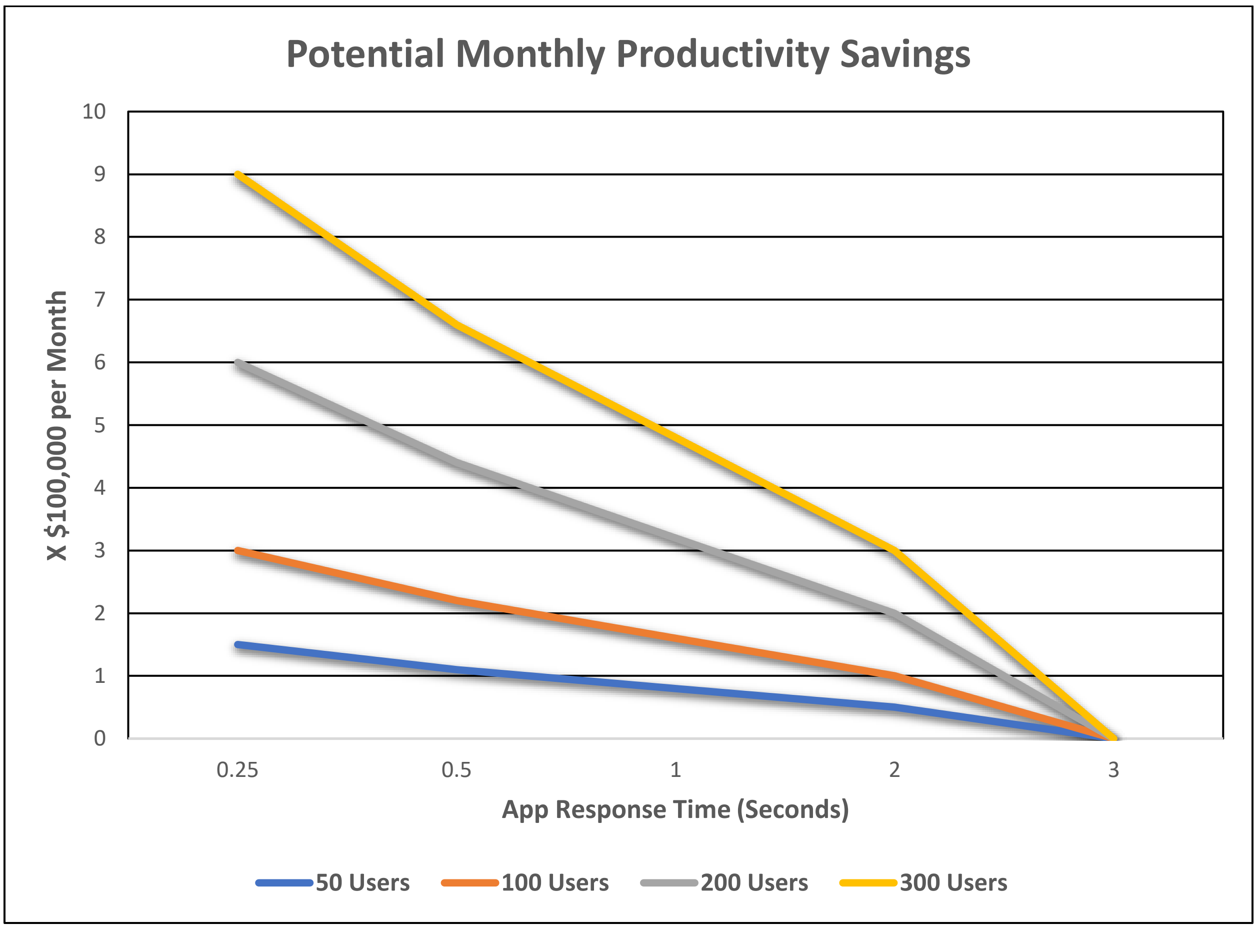

Just to put this in perspective. In the chart above based on ¼ second application response times:

- 50 users can save $1.8 million per year:
- 100 users can save as much as $3.6 million per year:
- 200 users can save as much as $7.2 million per year:
- 300 users can save as much as $10.8 million per year:

Obviously, that is a lot of potential cost savings. Just as obviously, it will vary by customer, skill sets of the users, and the application. Mission-critical and even business-critical applications will have a much bigger impact.

Appendix B: Dense X7 - vSAN vs OCI Block Performance Comparison Tests

Test Scenario

Comparing vSAN performance with OCI Block Volume as a VMFS Datastore. For this test, a Datastore cluster was created with several Block Volumes.

Individual Tests

1) 4K | 70% Read | 100% Random – the Random test simulates the most common workloads
2) 4K | 100% Read | 100% Random – the Random test shows the best realistic IOPS of this given configuration
3) 4K | 100% Write | 100% Random
4) 8K | 50% Read | 100% Random – the Random test simulates the OLTP workloads
5) 8K | 70% Read | 100% Random
6) 16K | 70% Read | 100% Random
7) 128K | 70% Read | 100% Random
8) 256K | 100% Write | 100% Sequential – this test shows the best realistic throughput of this given configuration.
9) 256K | 70% Read | 100% Random

Results:

When introducing more volumes the results showed better performance. When observed from a high level, every block volume performs well but not balanced. Higher cumulative throughput and IOPS are achieved but can be better when the volumes are sized properly.

<table>
<tr><th></th><th>Testing Parameters</th><th>BM. DenseIO2. 52 4K | 70% Read | 100% Random</th><th>BM. DenseIO2. 52 4K | 100% Read | 100% Random</th><th>BM. DenseIO2. 52 4K | 100% Write | 100% Random</th><th>BM. DenseIO2. 52 8K | 50% Read | 100% Random (OLTP)</th><th>BM. DenseIO2. 52 8K | 70% Read | 100% Random</th><th>BM. DenseIO2. 52 16K | 70% Read | 100% Random</th><th>BM. DenseIO2. 52 128K | 70% Read | 100% Random</th><th>BM. DenseIO2. 52 256K | 70% Read | 100% Random</th><th>BM. DenseIO2. 52 256K | 100% Write | 100% Sequential</th></tr>
<tr><td>SDDC Configuration</td><td>Region /Availability domain</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td><td>PHX/AD-2</td></tr>
<tr><td></td><td>OCVS Shape</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td><td>BM.DenseIO2.52</td></tr>
<tr><td></td><td>OCPU count per host</td><td>52</td><td>52</td><td>52</td><td>52</td><td>52</td><td>52</td><td>52</td><td>52</td><td>52</td></tr>
<tr><td></td><td>Number of hosts in SDDC</td><td>3</td><td>3</td><td>3</td><td>3</td><td>3</td><td>3</td><td>3</td><td>3</td><td>3</td></tr>
<tr><td>VSAN Capacity</td><td></td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td><td>122 TB</td></tr>
<tr><td>VMFS Configuration</td><td>OCI Block Volume Size</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td><td>8 TB</td></tr>
<tr><td></td><td>Number of volumes</td><td>8</td><td>8</td><td>8</td><td>8</td><td>8</td><td>8</td><td>8</td><td>8</td><td>8</td></tr>
<tr><td></td><td>Datastore Cluster Capacity</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td><td>64 TB</td></tr>
<tr><td></td><td>Performance</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU/GB: 50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td><td>UHP (VPU /GB:50)</td></tr>
<tr><td></td><td>Multipath</td><td>No</td><td>No</td><td>No</td><td>No</td><td>No</td><td>No</td><td>No</td><td>No</td><td>No</td></tr>
<tr><td></td><td>Attachment Type</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td><td>iSCSI</td></tr>
<tr><td></td><td>Expected IOPS</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td><td>125,000</td></tr>
<tr><td></td><td>Expected Throughput</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td><td>1,280 MB/s</td></tr>
<tr><td></td><td>Storage I/O Control</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td><td>Enabled</td></tr>
<tr><td></td><td>Storage DRS</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td><td>Fully Automated</td></tr>
<tr><td></td><td>DRS</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td><td>Manual</td></tr>
</table>

HCI Bench Configuration	Benchmarking Tool / Version	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407	VDBENCH / 50407
	Workload Type	4K | 70% Read | 100% Random	4K | 100% Read | 100% Random	4K | 100% Write | 100% Random	8K | 50% Read | 100% Random	8K | 70% Read | 100% Random	16K | 70% Read | 100% Random	128K | 70% Read | 100% Random	256K | 70% Read | 100% Random	256K | 100% Write | 100% Sequential
Guest VM Configuration	Number of VMs	80	80	80	80	80	80	80	80	80
	Number of CPU	2	2	2	2	2	2	2	2	2
	Size of RAM in GB	8	8	8	8	8	8	8	8	8
	Number of Drives to Test	8	8	8	8	8	8	8	8	8
	Size of Data Disks in GiB	10	10	10	10	10	10	10	10	10
Testing Configuration	Working-Set Percentage	80	80	80	80	80	80	80	80	80
	Number of Threads / Drive	2	2	2	2	2	2	2	2	2
	Block Size	4K	4K	4K	8K	8K	16K	128K	256K	256K
	Read Percentage	70	100	0	50	70	70	70	70	0
	Random Percentage	100	100	100	100	100	100	100	100	0
	I/O Rate	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL
	Test Time (secs)	600	600	600	600	600	600	600	600	600
	Warmup Time	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL
	Reporting Interval	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL

VSAN Performance Results	VMs	80	80	80	80	80	80	80	80	80
	IOPS	100541.1 IO/S	246127.3 IO/S	65869.3 IO/S	80618.9 IO/S	103089.7 IO/S	98905.5 IO/S	45580.2 IO/S	22447.9 IO/S	7112.7 IO/S
	THROUGHPUT	392.72 MB/s	961.42 MB/s	257.3 MB/s	629.84 MB/s	805.36 MB/s	1545.41 MB/s	5697.53 MB/s	5612.03 MB/s	1778.04 MB/s
	LATENCY	12.71 ms	5.22 ms	19.41 ms	15.83 ms	12.38 ms	12.9 ms	28.02 ms	56.94 ms	179.79 ms
	R_LATENCY	14.37 ms	5.22 ms	0.0 ms	9.15 ms	13.17 ms	15.85 ms	7.93 ms	10.49 ms	0.0 ms
	W_LATENCY	8.89 ms	0.0 ms	19.41 ms	22.41 ms	10.56 ms	6.07 ms	74.7 ms	164.91 ms	179.79 ms
	95th Percentile Latency	35.51 ms	10.06 ms	33.2 ms	31.29 ms	24.58 ms	35.41 ms	85.85 ms	217.77 ms	389.38 ms
	Resource Usage cpu.usage% | cpu.utilization% | mem.usage%	90.19 | 52.96 | 22.01	82.55 | 50.61 | 23.45	82.81 | 49.11 | 22.21	81.69 | 47.28 | 23.59	80.33 | 46.58 | 23.74	81.25 | 47.47 | 23.86	84.85 | 50.9 | 23.97	84.73 | 50.51 | 24.1	83.72 | 47.21 | 24.19

DS-denseX7-BV-1 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	83836.7	91083.9	109147.4	77842.8	85527.7	59191.1	9934.3	5020.2	4215.9
	THROUGHPUT (MB/s)	327.47	355.8	426.35	608.16	668.19	924.87	1241.8	1255.02	1053.99
	LATENCY	1.95 ms	1.78 ms	1.45 ms	2.07 ms	1.91 ms	2.71 ms	16.06 ms	31.72 ms	38.66 ms
	R_LATENCY	1.95 ms	1.78 ms	0.0 ms	2.08 ms	1.92 ms	2.51 ms	20.12 ms	38.93 ms	0.0 ms
	W_LATENCY	1.95 ms	0.0 ms	1.45 ms	2.06 ms	1.89 ms	3.19 ms	6.71 ms	14.96 ms	38.66 ms
	95th Percentile Latency	3.43 ms	3.43 ms	2.7 ms	3.7 ms	3.53 ms	5.15 ms	146.01 ms	196.87 ms	48.79 ms

DS-denseX7-BV-2 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	121972.5	120555.8	122478.8	103089	95158.1	75825.5	9993	4998.3	4236
	THROUGHPUT (MB/s)	476.46	470.91	478.42	805.39	743.43	1184.78	1249.1	1249.59	1058.97
	LATENCY	1.3 ms	1.34 ms	1.32 ms	1.56 ms	1.71 ms	2.13 ms	15.96 ms	31.9 ms	38.37 ms
	R_LATENCY	1.31 ms	1.34 ms	0.0 ms	1.57 ms	1.72 ms	1.99 ms	19.32 ms	39.89 ms	0.0 ms
	W_LATENCY	1.28 ms	0.0 ms	1.32 ms	1.55 ms	1.69 ms	2.44 ms	8.23 ms	13.35 ms	38.37 ms
	95th Percentile Latency	5.26 ms	5.08 ms	2.71 ms	3.43 ms	3.74 ms	9.27 ms	127.66 ms	207.05 ms	46.42 ms

DS-denseX7-BV-3 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	105258	110245.8	118640.6	94466.9	110724.7	73069.1	10012.2	5006	4204.8
	THROUGHPUT (MB/s)	411.17	430.64	463.44	738.02	865.04	1141.71	1251.54	1251.55	1051.19
	LATENCY	1.73 ms	1.53 ms	1.36 ms	1.85 ms	1.51 ms	2.32 ms	15.93 ms	31.85 ms	38.82 ms
	R_LATENCY	1.74 ms	1.53 ms	0.0 ms	1.86 ms	1.51 ms	2.14 ms	19.08 ms	39.28 ms	0.0 ms
	W_LATENCY	1.73 ms	0.0 ms	1.36 ms	1.84 ms	1.49 ms	2.76 ms	8.66 ms	14.6 ms	38.82 ms
	95th Percentile Latency	3.13 ms	3.22 ms	2.6 ms	3.27 ms	2.81 ms	5.65 ms	122.69 ms	199.97 ms	50.9 ms

DS-denseX7-BV-4 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	124117.9	123261.3	124851.7	107873.7	101934.1	75855	10017.7	5020.9	4431.2
	THROUGHPUT (MB/s)	484.83	481.51	487.71	842.77	796.35	1185.24	1252.2	1255.24	1107.79
	LATENCY	1.31 ms	1.32 ms	1.29 ms	1.51 ms	1.59 ms	2.12 ms	15.93 ms	31.76 ms	36.69 ms
	R_LATENCY	1.32 ms	1.32 ms	0.0 ms	1.53 ms	1.6 ms	2.0 ms	19.1 ms	38.97 ms	0.0 ms
	W_LATENCY	1.28 ms	0.0 ms	1.29 ms	1.49 ms	1.55 ms	2.42 ms	8.63 ms	15.02 ms	36.69 ms
	95th Percentile Latency	4.53 ms	5.22 ms	5.61 ms	3.43 ms	3.69 ms	6.86 ms	120.16 ms	191.37 ms	43.95 ms

DS-denseX7-BV-5 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	88484.9	78770.5	101661	76878.6	66665.2	54595.4	10027.9	5012.9	4198.6
	THROUGHPUT (MB/s)	345.65	307.71	397.11	600.59	520.83	853.06	1253.48	1253.2	1049.6
	LATENCY	1.96 ms	2.15 ms	1.67 ms	2.21 ms	2.48 ms	3.04 ms	15.92 ms	31.82 ms	38.96 ms
	R_LATENCY	1.95 ms	2.15 ms	0.0 ms	2.22 ms	2.48 ms	2.82 ms	19.27 ms	38.74 ms	0.0 ms
	W_LATENCY	1.96 ms	0.0 ms	1.67 ms	2.2 ms	2.46 ms	3.55 ms	8.19 ms	15.77 ms	38.96 ms
	95th Percentile Latency	5.06 ms	6.82 ms	3.18 ms	8.24 ms	6.36 ms	11.79 ms	127.01 ms	186.35 ms	50.02 ms

DS-denseX7-BV-6 Performance Test Results	VMs	10	10	10	10	10	10	10		10
	IOPS (IO/S)	120935.8	121361	122142	117022.8	119124.5	76222.1	10018.2	5006	4247.3
	THROUGHPUT (MB/s)	472.41	474.07	477.12	914.24	30.67	1190.98	1252.26	1251.45	1061.82
	LATENCY	1.33 ms	1.32 ms	1.32 ms	1.36 ms	1.34 ms	2.11 ms	15.93 ms	31.86 ms	38.28 ms
	R_LATENCY	1.33 ms	1.32 ms	0.0 ms	1.36 ms	1.34 ms	1.93 ms	19.23 ms	39.85 ms	0.0 ms
	W_LATENCY	1.33 ms	0.0 ms	1.32 ms	1.36 ms	1.33 ms	2.51 ms	8.32 ms	13.31 ms	38.28 ms
	95th Percentile Latency	4.63 ms	6.44 ms	3.28 ms	6.93 ms	4.59 ms	10.12 ms	123.53 ms	208.64 ms	47.75 ms

DS-denseX7-BV-7 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	71235.6	80768.8	94770.6	67144.9	76335.3	49984.8	10004.5	5053.6	4244.3
	THROUGHPUT (MB/s)	278.27	315.51	370.21	524.56	596.37	781.02	1250.58	1263.37	1061.08
	LATENCY	2.26 ms	2.04 ms	1.71 ms	2.39 ms	2.2 ms	3.22 ms	15.95 ms	31.53 ms	38.3 ms
	R_LATENCY	2.26 ms	2.04 ms	0.0 ms	2.39 ms	2.21 ms	2.99 ms	19.24 ms	38.95 ms	0.0 ms
	W_LATENCY	2.26 ms	0.0 ms	1.71 ms	2.38 ms	2.19 ms	3.77 ms	8.36 ms	14.3 ms	38.3 ms
	95th Percentile Latency	5.94 ms	4.08 ms	3.55 ms	4.74 ms	4.34 ms	7.94 ms	128.66 ms	195.14 ms	45.5 ms
DS-denseX7-BV-8 Performance Test Results	VMs	10	10	10	10	10	10	10	10	10
	IOPS (IO/S)	101383.6	106528.5	108720.2	87973	87697.6	67845.1	10017.1	5018.5	4239
	THROUGHPUT (MB/s)	396.04	416.13	424.68	687.28	685.13	1060.09	1252.14	1254.61	1059.77
	LATENCY	1.7 ms	1.69 ms	1.59 ms	1.91 ms	1.95 ms	2.57 ms	15.93 ms	31.79 ms	38.32 ms
	R_LATENCY	1.7 ms	1.69 ms	0.0 ms	1.92 ms	1.96 ms	2.37 ms	19.03 ms	38.29 ms	0.0 ms
	W_LATENCY	1.7 ms	0.0 ms	1.59 ms	1.91 ms	1.94 ms	3.02 ms	8.78 ms	16.7 ms	38.32 ms
	95th Percentile Latency	3.09 ms	3.03 ms	2.73 ms	3.4 ms	3.43 ms	4.64 ms	124.01 ms	175.51 ms	45.0 ms
	Resource Usage cpu.usage% | cpu.utilization% | mem.usage%	93.1 | 54.74 | 22.1	93.36 | 54.18 | 23.05	96.82 | 57.51 | 24.2	94.56 | 54.72 | 23.12	94.17 | 54.43 | 23.21	92.83 | 52.76 | 23.39	85.99 | 44.88 | 23.5	85.67 | 44.4 | 23.9	86.99 | 44.18 | 24.11
Combined VMFS Performance	Cumulative IOPS (IO/S)	817225	832576	902412	732292	743167	532588	80025	40136	34017
	Cumulative Throughput (MB /s)	3192	3252	3525	5721	4906	8322	10003	10034	8504
	Average Throughput (MB /s)	399	407	441	715	613	1040	1250	1254	1063