

# Using Dell EMC Unity Storage in an Equinix Cloud-Connected Data Center

## A Detailed Review

### Abstract

This whitepaper discusses use cases for running Dell EMC Unity storage in a public cloud colocation facility, validates storage latency and performance, and provides recommendations.

October 2019



## Revisions

| Version | Date          | Description                        |
|---------|---------------|------------------------------------|
| 0.01    | February 2019 | Initial draft                      |
| 0.18    | May 2019      | Updated to include Microsoft Azure |
| 0.24    | October 2019  | Minor edits                        |

## Acknowledgements

This paper was produced by the following members of the Dell EMC Presales team:

Author: Ryan Tassotti

This whitepaper was possible thanks to the contributions from:

- Rajiv Juneja                      Dell EMC
- Alex Seymour                    Dell EMC
- Gabriel Lageyre                Equinix
- Lee Sharping                    Equinix
- Mischa Jampolsky               Equinix

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Equinix, Platform Equinix, and IBX are registered trademarks of Equinix, Inc. Equinix Cloud Exchange Fabric, ECX Fabric, and Equinix International Business Exchange are trademarks of Equinix, Inc.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Table of contents

|  |    |
|--|----|
| Revisions.....   | 2  |
| Acknowledgements.....  | 2  |
| Executive summary.....   | 4  |
| 1 The Case for Storage in a Cloud-Connected Colocation Facility..... | 5  |
| 1.1 Use Cases .....  | 5  |
| 1.2 Why Validate by Testing?.....                                    | 5  |
| 2 Amazon Web Services Architecture and Testing.....                  | 6  |
| 2.1 Architecture.....  | 6  |
| 2.2 AWS Testing Configuration .....                                  | 6  |
| 2.3 The Tests.....   | 9  |
| 3 Amazon Web Services Testing Results .....                          | 14 |
| 3.1 Windows Disk Benchmarking Results.....                           | 14 |
| 3.2 Windows SQL Benchmarking Results.....                            | 16 |
| 3.3 Linux Disk Benchmarking Results.....                             | 17 |
| 3.4 Linux Oracle Benchmarking Results .....                          | 19 |
| 3.5 Unity performance observations.....                              | 19 |
| 4 Microsoft Azure Architecture and Testing .....                     | 22 |
| 4.1 Architecture.....  | 22 |
| 4.2 Azure Testing Configuration .....                                | 22 |
| 4.3 The Tests.....   | 25 |
| 5 Microsoft Azure Testing Results.....                               | 27 |
| 5.1 Windows Disk Benchmarking Results.....                           | 27 |
| 5.2 Windows SQL Benchmarking Results.....                            | 29 |
| 6 Conclusion.....  | 30 |
| 6.1 Observations.....  | 30 |
| 6.2 Recommendations.....   | 30 |
| 6.3 Calls to action .....  | 30 |
| A Technical support and resources .....                              | 31 |
| A.1 Related resources .....  | 31 |

## Executive summary

For some customers, public cloud infrastructure has changed the way they deploy, manage and consume IT. For other customers, public cloud is something that should be explored, because it has the potential to be a suitable location for some of their workloads. However, having workloads in the public cloud can raise issues that are not applicable to an on-premises architecture, such as:

- Data sovereignty: Where is my data stored, and who has access to it?
- Cost predictability: What are my future public cloud charges?
- Performance predictability: How do I maintain consistent performance?

Generally speaking, the public cloud conversation will discuss taking a customer's on premises workload (including compute and storage) and moving it in its entirety into a public cloud provider's infrastructure, which can raise the above issues for some customers. The purpose of this whitepaper is to validate an alternative approach, which is to take advantage of the compute benefits that public cloud provides, while addressing the aforementioned issues by housing the customer's data holdings on a Dell EMC Unity storage array in a public cloud colocation facility, such as an Equinix® Cloud-Connected Data Center.

This whitepaper will outline use cases for public cloud compute instances connecting to a Dell EMC Unity storage array in an Equinix facility, compare storage latency and performance for both public cloud storage and Unity storage, and provide recommendations.

# 1 The Case for Storage in a Cloud-Connected Colocation Facility

When it comes to cloud infrastructure, there are 3 commonly used architectures:

- Private cloud (on premises): The customer's workloads run on onsite infrastructure.
- Public cloud: The customer's workloads are housed offsite.
- Hybrid cloud: Some of the customer's workloads are in public cloud, and some workloads are in their onsite private cloud.

There is also mixed architecture, where for a given set of workloads, a customer will have their compute instances in public cloud, and their data (storage) housed in a public cloud-connected colocation facility, such as an Equinix Cloud-Connected Data Center.

## 1.1 Use Cases

There are four broad use case categories which describe why a public cloud compute customer would want to have their data at rest in a public cloud-connected colocation facility:

- **Data sovereignty:** A customer wants to take advantage of the economics of public cloud compute services and application, however is mandated to keep data stored within a geographical location.
- **Disaster recovery:** A customer has an on-premises storage array and wants to create a disaster recovery capability by performing storage replication to an offsite storage array. If this offsite storage array is hosted in a public cloud-connected colocation facility, then public cloud compute instances can be spun up and attached to the offsite storage replica(s) in a disaster recovery scenario.
- **Predictable TCO:** Storage can be the highest cost component of public cloud. Hosting a storage array in a public cloud-connected colocation facility provides a fixed cost for storage, while leveraging public cloud compute services and applications.
- **Storage performance:** Some deployed data center workloads such as databases require consistent low latencies from the storage system, to ensure a predictable user experience. Hosting a storage array in a public cloud-connected colocation facility provides more consistent and lower storage latencies than public cloud storage.

## 1.2 Why Validate by Testing?

Many customers using an on premises private cloud architecture run workloads that are sensitive to both storage latency and performance, a good example being databases. If a customer has an OLTP database workload that is important to their business or operation, the probability that this customer expects consistent low storage latency and predictable performance is high.

## 2 Amazon Web Services Architecture and Testing

This section will outline the Amazon Web Services (AWS) architecture as tested, and detail the tests performed.

### 2.1 Architecture

The high-level architecture is presented in Figure 1. The architecture consists of:

- Dell EMC Unity storage array, residing in a public cloud-connected colocation facility (Equinix);
- Dell EMC 10Gb networking switch, also in a public cloud-connected colocation facility (Equinix);
- 10Gb AWS Direct Connect IP network link between colocation-connected facility and AWS public cloud;
- EC2 compute instances in AWS that connect to the Dell EMC Unity storage array.

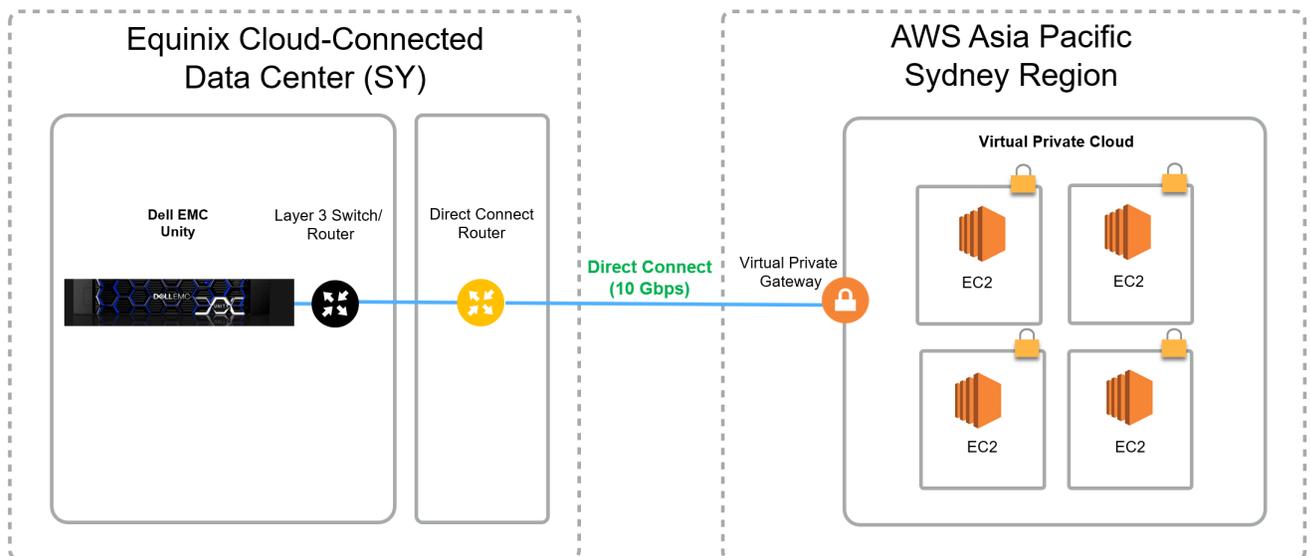


Figure 1 High level AWS architecture as tested

### 2.2 AWS Testing Configuration

This section will outline the as-tested configuration of the AWS EC2 instances, the AWS EBS volumes, the AWS Direct Connect network link, the Dell EMC Unity storage array, and the Dell EMC network switch.

---

**Note:** All AWS resources (EC2, EBS, etc.) were deployed in the *Asia Pacific (Sydney)* Region, specifically the *ap-southeast-2a* Availability Zone. The Dell EMC resources (Unity storage array, network switch) were deployed at an Equinix International Business Exchange™ (IBX®) facility in Sydney.

---

#### 2.2.1 AWS EC2 Windows Instance

The features of the chosen Amazon EC2 Windows instance are as follows:

- m4.4xlarge Instance type
- Microsoft Windows Server 2016 with SQL Server 2016 Standard
- 64 GiB memory
- 16 vCPUs

- Enhanced Networking
- EBS-Optimized *storage*
- Jumbo frames is enabled

## 2.2.2 AWS EC2 Linux Instance

The features of the chosen Amazon EC2 Linux instance are as follows:

- m4.4xlarge Instance type
- Oracle Linux 7 Update 6 (x86\_64)
- 64 GiB memory
- 16 vCPUs
- Enhanced Networking
- EBS-Optimized *storage*
- Jumbo frames is enabled

## 2.2.3 AWS EBS Volumes

The EBS Volumes for attaching to the Windows EC2 instance were created and configured as follows:

- One (1) 30 GiB General Purpose SSD (gp2) volume was configured as the C: drive
- This was included as part of the m4.4xlarge instance deployment
- Not encrypted
- Five (5) 30 GiB General Purpose SSD (gp2) volumes were presented to the host, and a single software RAID5 volume was created, mounted as the D: drive
- Not encrypted

The EBS Volumes for attaching to the Linux EC2 instance were created and configured as follows:

- One (1) 30 GiB General Purpose SSD (gp2) volume was configured as the /root volume
- This was included as part of the m4.4xlarge instance deployment
- Not encrypted
- Five (5) 30 GiB General Purpose SSD (gp2) volumes were presented to the host, and a single software RAID5 device was created, mounted as /home/oracle/raid5.
- Not encrypted

---

Several additional General Purpose SSD (gp2) volumes were presented to the host for the Oracle installation, as per Oracle best practice. No performance tests were run against these volumes, and these volumes were not encrypted.

---

---

**Note:** AWS EBS General Purpose SSD (gp2) volumes have the concept of IO burst credits, which is covered in detail [here](#). In summary, EBS gp2 volumes will perform at a minimum of 3 IOPS per GiB and will burst up to 3,000 IOPS per volume if in credit. **For all tests performed against and results given for EBS volumes in this whitepaper, the EBS volumes were wholly using burst credits**, i.e. the EBS volumes were never tested at their minimum stated performance of 3 IOPS per GiB. It is for this reason that each test iteration was only run for several minutes and were spaced over a 24-hour period.

---

## 2.2.4 AWS Direct Connect

When in an Equinix facility, there are two ways to connect the co-located Unity storage array to the AWS EC2 compute instances:

- Equinix Cloud Exchange Fabric™ (ECX Fabric™); or
- AWS Direct Connect

### Equinix Cloud Exchange Fabric

The major benefit of connecting into public cloud compute via the Equinix Cloud Exchange Fabric is that all of the network links into public cloud providers (Amazon AWS, Microsoft Azure, Google Cloud Platform, and Alibaba Cloud) are handled by a single connection out of the customer's co-located equipment. So, for example, the Unity storage array connects to a network switch in the same rack, and it is this network switch that connects once to the Equinix Cloud Exchange Fabric. From there, connections to the required public cloud providers are established and managed by Equinix. More information can be found [here](#).

*“ECX Fabric brings together cloud service providers and users, enabling them to establish affordable, private, high-performance connections within Platform Equinix®” (Equinix).*

### AWS Direct Connect

To establish a private connection between AWS and the Unity storage array, the customer engages with AWS via the AWS Console to procure a Direct Connect connection. More information can be found [here](#).

*“Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.” (Amazon)*

For the purposes of this testing, an AWS Direct Connect network link was established. Direct Connect routing was provided by Equinix.

The features of the AWS Direct Connect network link as configured are as follows:

- Port Speed is 10Gbps
- Jumbo Frame Capable is true
- *Location is* Equinix SY1 – SY4, Sydney, AUS

The features of the associated Virtual Interface as configured are as follows:

- Type is private
- Jumbo Frame Capable is true
- MTU is 9001

## 2.2.5 Dell EMC Unity Storage

Deliver consistent high performance for workloads with Dell EMC Unity storage. Dell EMC Unity brings you high performance and low latency across a broad range of SAN and NAS use cases.

Dell EMC Unity redefines simple, modern, affordable, and flexible primary storage. Its modern architecture is designed for flash and unified data services, in a footprint starting at 2U. High-density SSD technology and inline data reduction adds to performance and value.

This unified storage comes standard with support for file and block, point-in-time snapshots, local and remote data replication, and built-in encryption. You'll also have deep integration with VMware, Microsoft, and OpenStack ecosystems. Dell EMC Unity is also available as a virtual appliance (Dell EMC UnityVSA), and in the public cloud (Dell EMC Unity Cloud Edition).

Ongoing operation is simple with proactive assistance and automated remote support. Gain intelligent insight into your storage health with CloudIQ, which provides cloud-based proactive monitoring and predictive analytics. More information can be found [here](#).

The configuration of the tested Dell EMC Unity storage array is as follows:

- Dell EMC Unity 400 hybrid storage array
- Dell EMC Unity OE version 4.4.1.1539309879
- 2 x 400GB SSDs
- Configured as RAID 1+1
- NFS share
  - NFS v3
  - 75GB size
  - Mounted as /home/oracle/nfs on the Linux host
- SMB share
  - SMB v3
  - 75GB size
  - Mounted as Y: drive on the Windows host
- Jumbo frames enabled
- 2 rack units

## 2.2.6 Dell EMC Network Switch

The S4148F-ON network switch is a one rack unit (RU), full-featured fixed form-factor top-of-rack (ToR) 10/25/40/50/100GbE switch for 10G servers with small form-factor pluggable plus (SFP+), quad small form-factor pluggable plus (QSFP+), and quad small form-factor pluggable (QSFP28) ports.

Dell EMC Networking OS10 Enterprise Edition is a network operating system (OS) supporting multiple architectures and environments. The networking world is moving from a monolithic stack to a pick-your-own-world. The OS10 solution allows disaggregation of the network functionality. More information can be found [here](#).

The configuration of the tested network switch is as follows:

- Model: S4148F-ON
- OS version: 10.4.2.0
- Connectivity to Unity storage array: 10Gb TwinAx
- Connectivity to Direct Connect router: 10Gb LR SFP+
- 1 rack unit
- Jumbo frames enabled

## 2.3 The Tests

There are four test categories used to validate the viability of the Dell EMC Unity solution in an Equinix facility:

- Windows disk benchmarking

- Windows SQL benchmarking
- Linux disk benchmarking
- Linux Oracle benchmarking

## 2.3.1 Windows Disk Benchmarking

To test disk latency and performance, the Microsoft tool *diskspd* is installed on the AWS EC2 Windows instance. The version used is 2.0.17, and the tool is available [here](#).

*“A feature-rich and versatile storage testing tool, Diskspd combines robust and granular IO workload definition with flexible runtime and output options, creating an ideal tool for synthetic storage subsystem testing and validation.”* (Microsoft)

The four tests that are performed as part of Windows disk benchmarking are:

- 100% read, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -w0
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -w0
```

- 100% write, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -Z1G -w100
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -Z1G -w100
```

- 75/25% read/write, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -Z1G -w25
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -Z1G -w25
```

- 100% read, sequential, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -si -o4 -d120 -Sh -ft -L -w0
```

```
diskspd.exe -b8k -c1G -t4 -si -o16 -d120 -Sh -ft -L -w0
```

---

**Note:** There are two iterations of each test run. The first iteration sets the number of outstanding IOs at 4 (-o4). The second iteration sets the number of outstanding IOs at 16 (-o16).

---

Additional notes about these tests and results:

- Each test iteration is run five times against each storage type, at various times throughout a 24-hour period.
  - Each test iteration is run for two minutes.
  - The results presented are averaged (Results are added then divided by five).
  - Two sets of results are produced for each test iteration: Latency and IOs.
  - Latency figures given are for the 99<sup>th</sup> percentile.
  - IOs figure given is the average number of IOs transferred over the two-minute period.
-

## 2.3.2 Windows SQL Benchmarking

To test synthetic SQL database performance, the open source tool *HammerDB* is installed on the AWS EC2 Windows instance. The version used is 3.1, and the tool is available [here](#).

*“HammerDB is the leading benchmarking and load testing software for the world’s most popular databases supporting Oracle Database, SQL Server, IBM Db2, MySQL, MariaDB, PostgreSQL and Redis.”* (HammerDB)

Notes about this test and results:

- The SQL version is 2017.
- Warehouses created is 16.
- Virtual users is 16.
- Warm up time is 2 mins.
- Run time is 10 mins.
- Each test is run five times against each storage type, at various times throughout a 24-hour period.
- The results presented are averaged (Results are added then divided by five).
- The results given are Transactions per Minute (TPM).
- No database-tuning was performed.

## 2.3.3 Linux Disk Benchmarking

To test disk latency and performance, the open source tool *fio* is installed on the AWS EC2 Linux instance. The version used is 2.0.9, and the tool is available [here](#).

*“Fio was written by Jens Axboe <axboe@kernel.dk> to enable flexible testing of the Linux I/O subsystem and schedulers. He got tired of writing specific test applications to simulate a given workload, and found that the existing I/O benchmark/test tools out there weren’t flexible enough to do what he wanted.”*  
(Github)

The four tests that are performed as part of Linux disk benchmarking are:

- 100% read, random, 8KB IO size

```
fio --name=randread --ioengine=libaio --iodepth=4 --rw=randread --bs=8k --direct=1 --size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based
```

```
fio --name=randread --ioengine=libaio --iodepth=16 --rw=randread --bs=8k --direct=1 --size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based
```

- 100% write, random, 8KB IO size

```
fio --name=randwrite --ioengine=libaio --iodepth=4 --rw=randwrite --bs=8k --direct=1 --size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based
```

```
fio --name=randwrite --ioengine=libaio --iodepth=16 --rw=randwrite --bs=8k --direct=1 --size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based
```

- 75/25% read/write, random, 8KB IO size

```
fio --name=randrw --ioengine=libaio --iodepth=4 --rw=randrw --bs=8k --direct=1 --size=2048m
--numjobs=4 --runtime=120 --group_reporting --time_based --rwmixread=75
```

```
fio --name=randrw --ioengine=libaio --iodepth=16 --rw=randrw --bs=8k --direct=1 --
size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based --rwmixread=75
```

- 100% read, sequential, 64KB IO size

```
fio --name=seqread --ioengine=libaio --iodepth=4 --rw=read --bs=64k --direct=1 --size=2048m
--numjobs=4 --runtime=120 --group_reporting --time_based
```

```
fio --name=seqread --ioengine=libaio --iodepth=16 --rw=read --bs=64k --direct=1 --
size=2048m --numjobs=4 --runtime=120 --group_reporting --time_based
```

---

**Note:** There are two iterations of each test run. The first iteration sets the number of outstanding IOs at 4 (--iodepth=4). The second iteration sets the number of outstanding IOs at 16 (--iodepth=16).

---

Additional notes about these tests and results:

- Each test iteration is run five times against each storage type, at various times throughout a 24-hour period.
- Each test iteration is run for two minutes.
- The results presented are averaged (Results are added then divided by five).
- Two sets of results are produced for each test iteration: Latency and throughput.
- Latency figures given are for the 99<sup>th</sup> percentile.
- Throughput figures given are the average number of MBs transferred over the two-minute period.

## 2.3.4 Linux Oracle Benchmarking

To test synthetic Oracle database performance, the Oracle tool *ORION* (ORacle IO Numbers) is installed on the AWS EC2 Linux instance. ORION is included with the Oracle installation, the version used is 12.2.0.1.0, and more information about the tool is [here](#).

*“Oracle Orion is a tool for predicting the performance of an Oracle database without having to install Oracle or create a database. Unlike other I/O calibration tools, Oracle Orion is expressly designed for simulating Oracle database I/O workloads using the same I/O software stack as Oracle. Orion can also simulate the effect of striping performed by Oracle Automatic Storage Management.”* (Oracle)

The test that is performed, when in the applicable storage subdirectory, is:

- OLTP, 8KB IO size

```
orion -run oltp -testname <storage_name> -hugenotneeded -cache_size 0 -num_disks 15
```

Additional notes about these tests and results:

- ORION tests with random small (8K) IOs at increasing loads to determine the maximum IOPS of the storage it is run against. So eventually at a certain data point, the maximum IOPS for that storage subsystem is reached, and the latency observed at that data point is recorded.
- Each test is run until completed.
- IOPS figure given is the maximum as recorded by ORION.
- Latency figure given is at the maximum IOPS figure as recorded by ORION.

### 3 Amazon Web Services Testing Results

This section provides the results for disk and application benchmarking for AWS EBS storage and Dell EMC Unity storage.

#### 3.1 Windows Disk Benchmarking Results

The latency results with outstanding IOs set to 4 are presented in Table 1. Lower numbers are better.

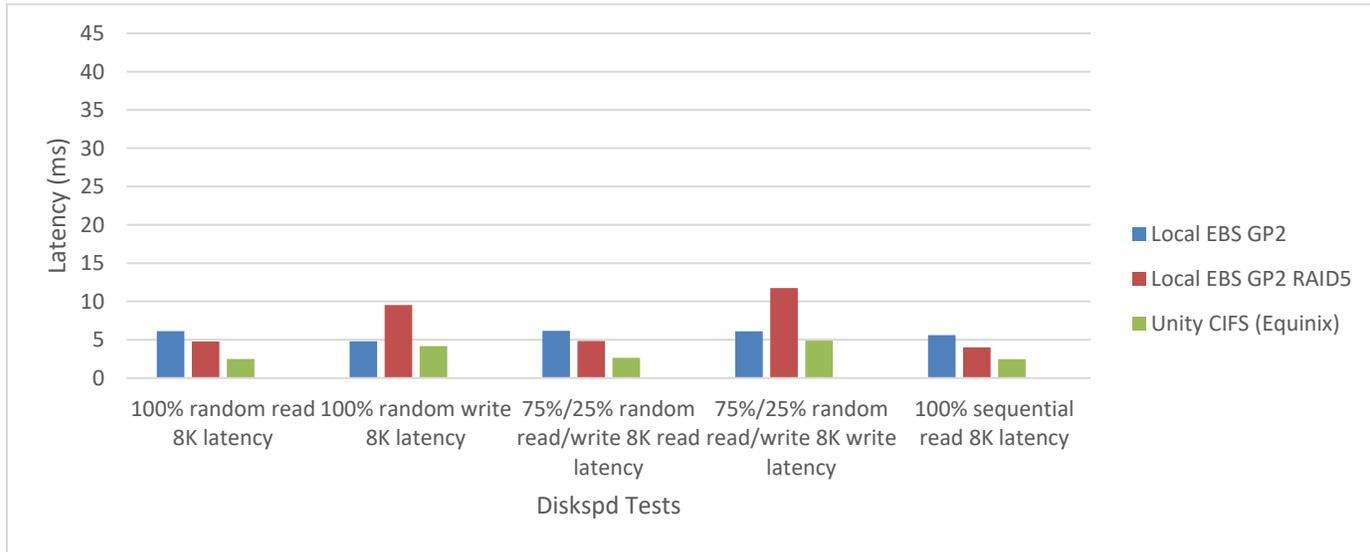


Table 1 *Diskspd* latency observed at host with outstanding IOs = 4

The latency results with outstanding IOs set to 16 are presented in Table 2. Lower numbers are better.

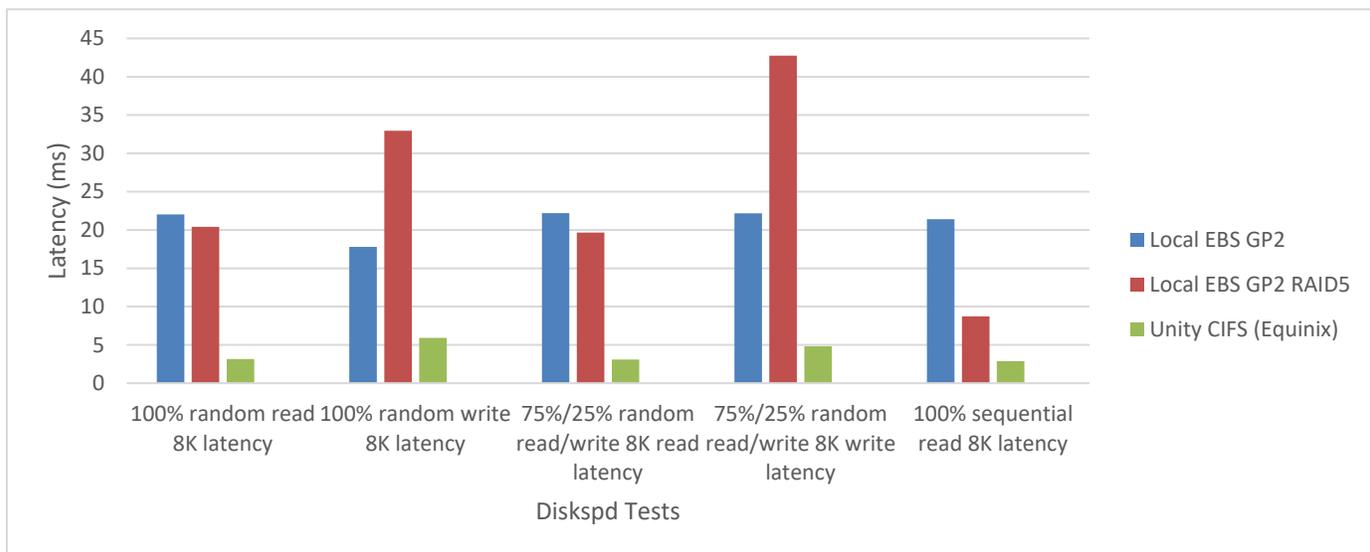


Table 2 *Diskspd* latency observed at host with outstanding IOs = 16

The IOPS results with outstanding IOs set to 4 are presented in Table 3. Higher numbers are better.

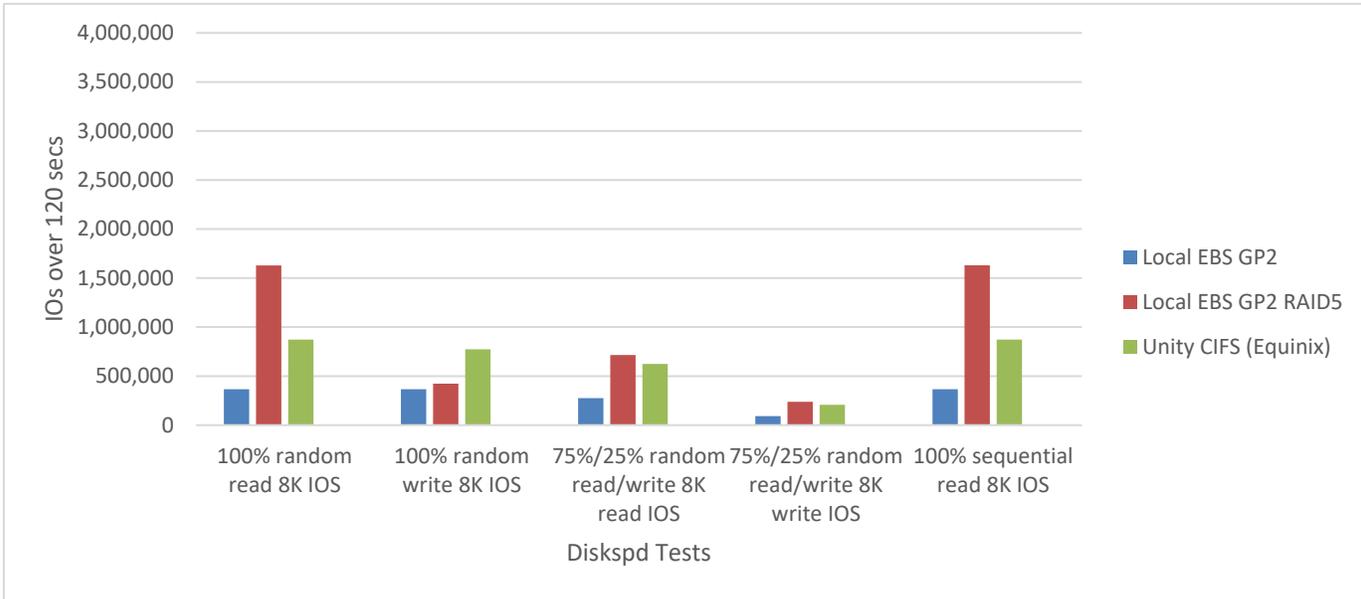


Table 3 *Diskspd* IOs transferred with outstanding IOs = 4

The IOPS results with outstanding IOs set to 16 are presented in Table 4. Higher numbers are better.

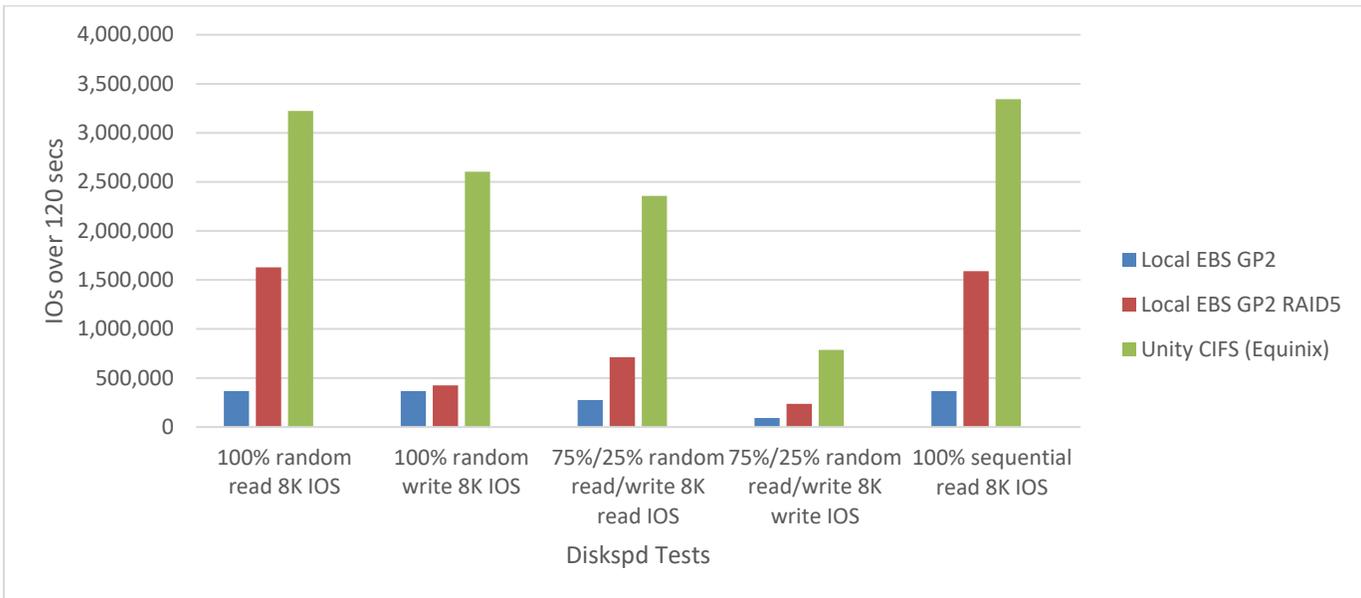


Table 4 *Diskspd* IOs transferred with outstanding IOs = 16

## 3.2 Windows SQL Benchmarking Results

The synthetic SQL benchmarking results are presented in Table 5. Higher numbers are better.

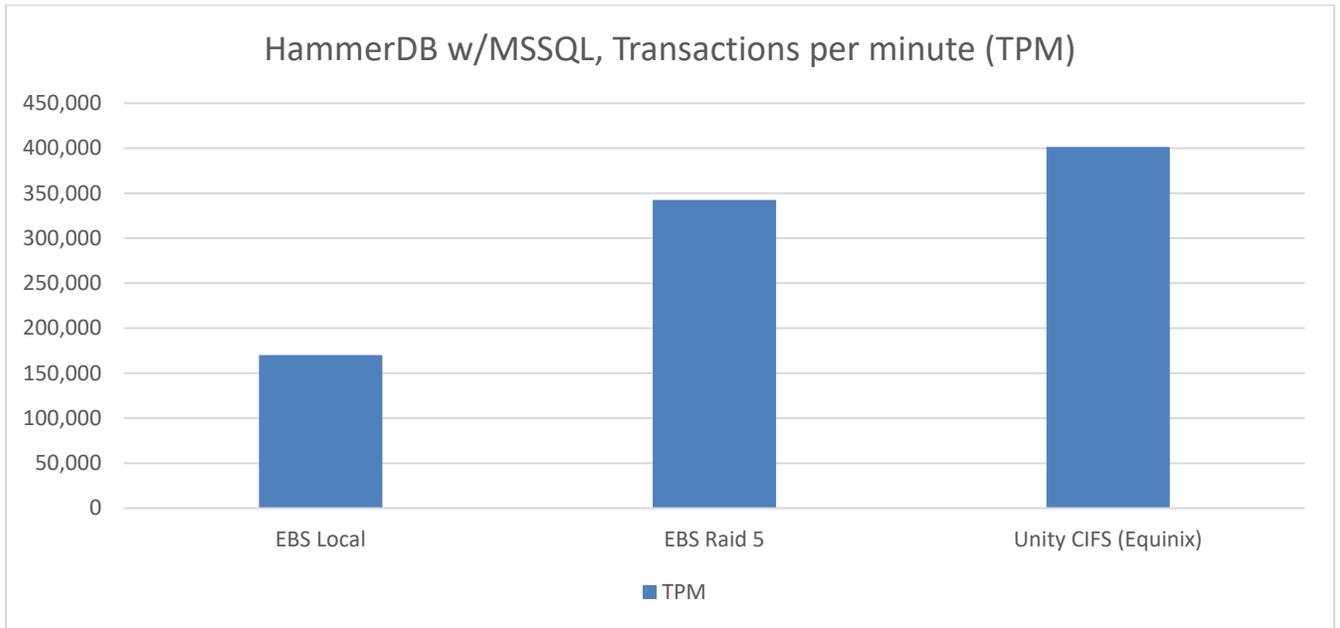


Table 5 *HammerDB* Transactions per Minute

### 3.3 Linux Disk Benchmarking Results

The latency results with outstanding IOs set to 4 are presented in Table 6. Lower numbers are better.

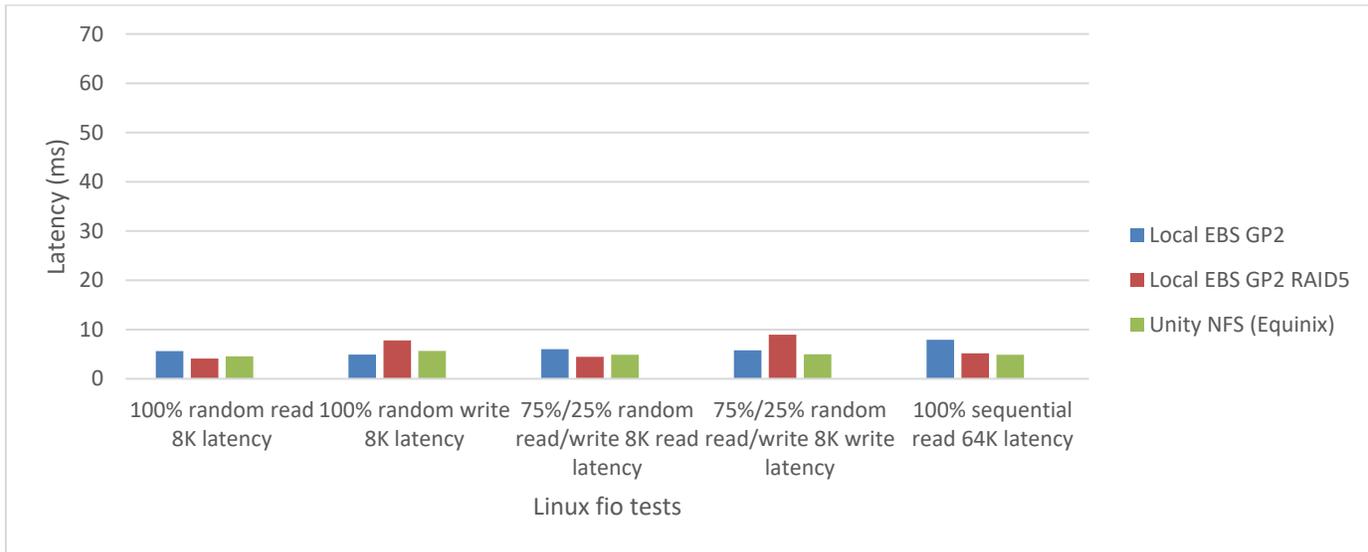


Table 6 *Fio* latency observed at host with outstanding IOs = 4

The latency results with outstanding IOs set to 16 are presented in Table 7. Lower numbers are better.

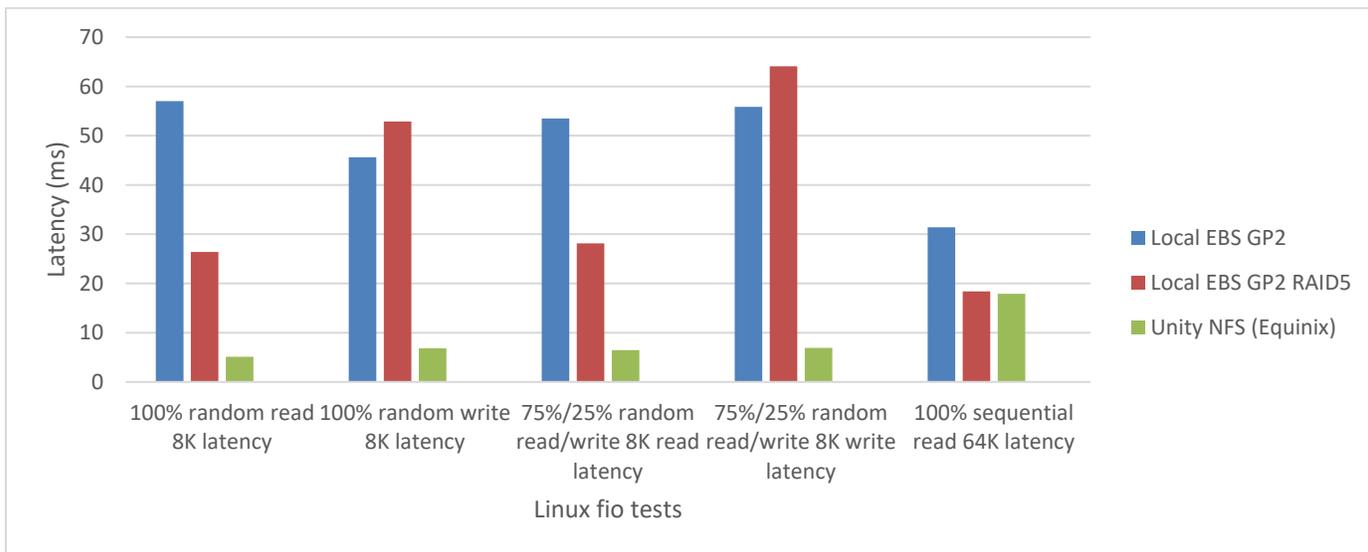


Table 7 *Fio* latency observed at host with outstanding IOs = 16

The throughput results with outstanding IOs set to 4 are presented in Table 8. Higher numbers are better.

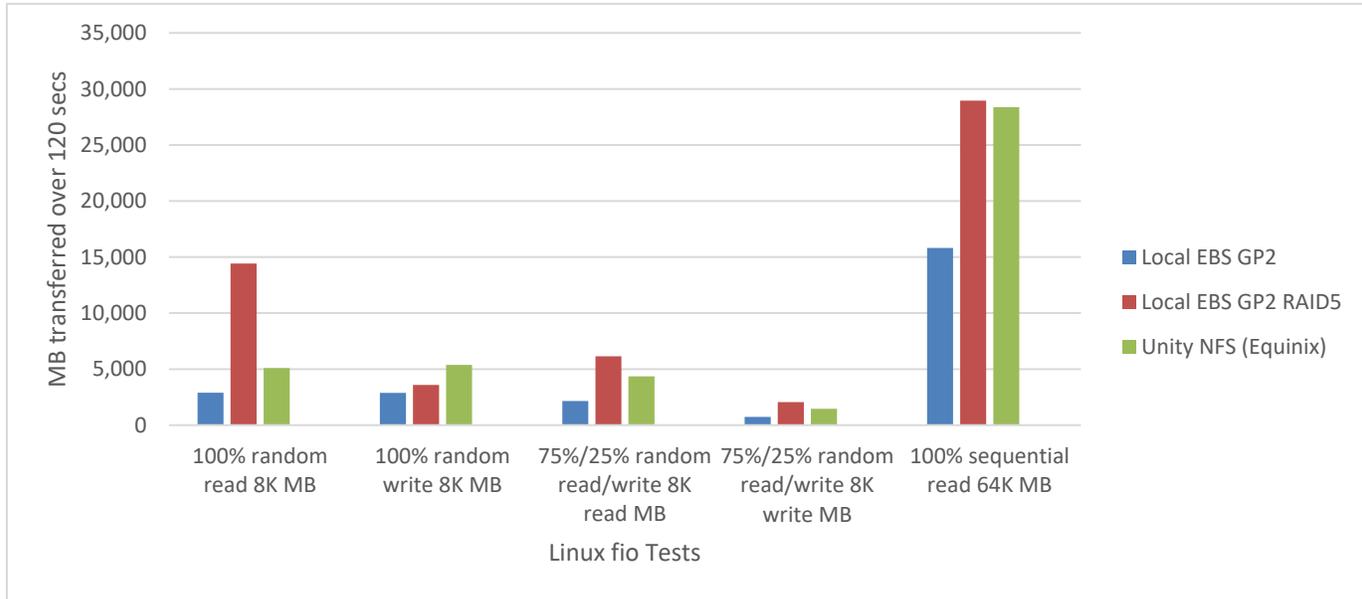


Table 8 Fio MBs transferred with outstanding IOs = 4

The throughput results with outstanding IOs set to 16 are presented in Table 9. Higher numbers are better.

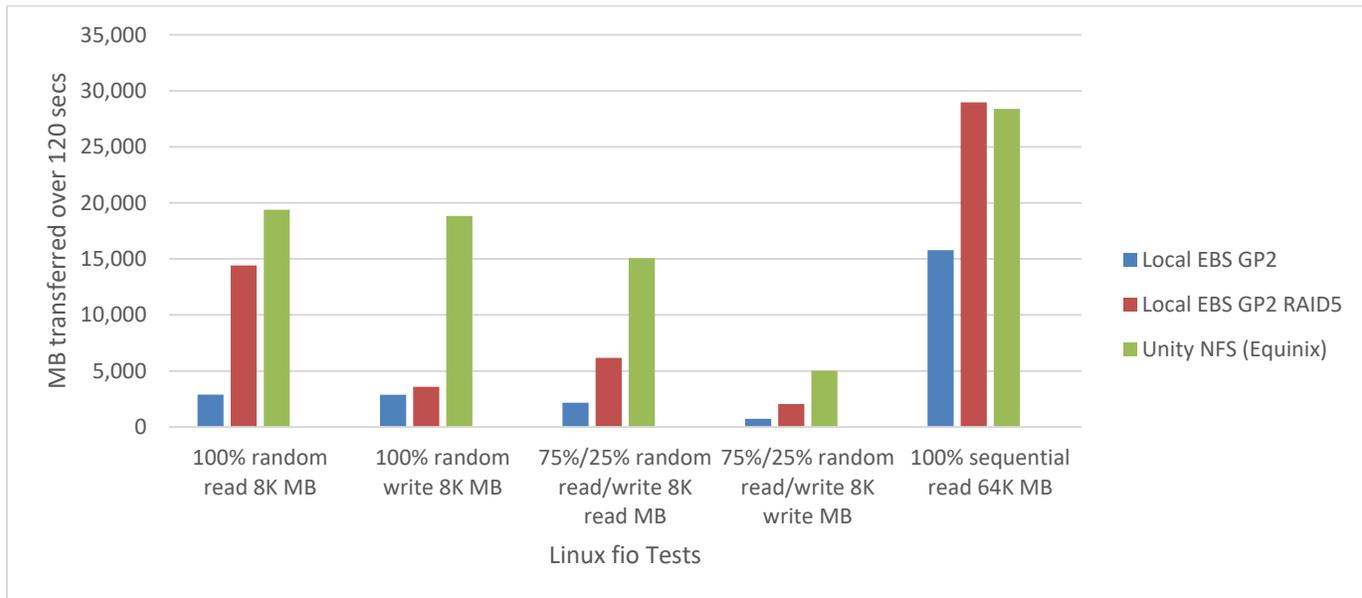


Table 9 Fio MBs transferred with outstanding IOs = 16

### 3.4 Linux Oracle Benchmarking Results

A summary of the maximum IOPS for each storage type tested, and the latency observed at the host at these maximum IOPS, is presented in Tables 10 and 11.

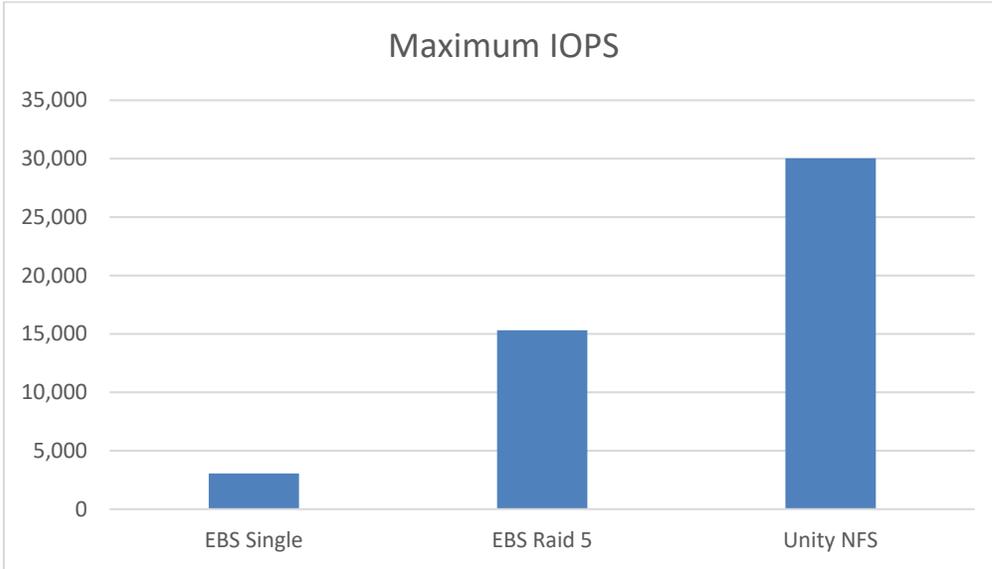


Table 10 *ORION* Maximum IOPS observed per Storage Type

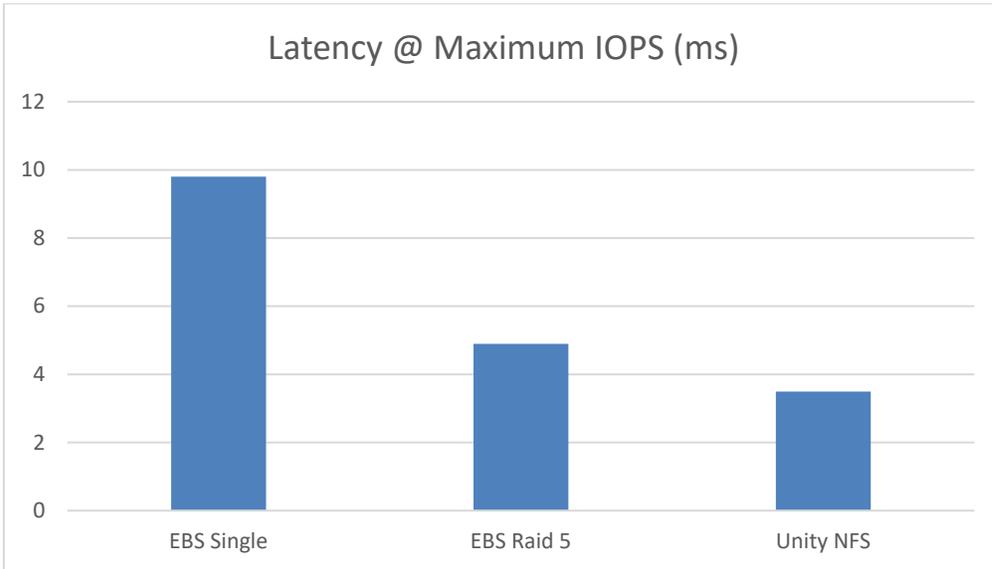


Table 11 *ORION* Latency observed per Storage Type at Maximum IOPS

### 3.5 Unity performance observations

While performance tests were running on the EC2 instances, some of the Unity storage arrays performance charts were captured and are presented in Tables 12-14. Recall that the disk configuration within the tested Unity storage array is 2 x 400GB SSDs configured as RAID1.

Table 12 displays the results of Windows *diskspd* tests from the storage array perspective. It can be observed that during the 8KB random read testing, performance is approx. 220MB/sec.

Table 13 displays the results of Windows *diskspd* tests from the storage array perspective. It can be observed that during the 8KB random write testing, performance is approx. 175MB/sec.

Table 14 displays the results of the *ORION* testing from the storage array perspective, and it can be observed that: (a) 8KB IOs are being served by the array, and (b) *ORION* gradually ramps up IOPS to the storage, to the point where the maximum disk performance is seen, approx. 220MB/sec.

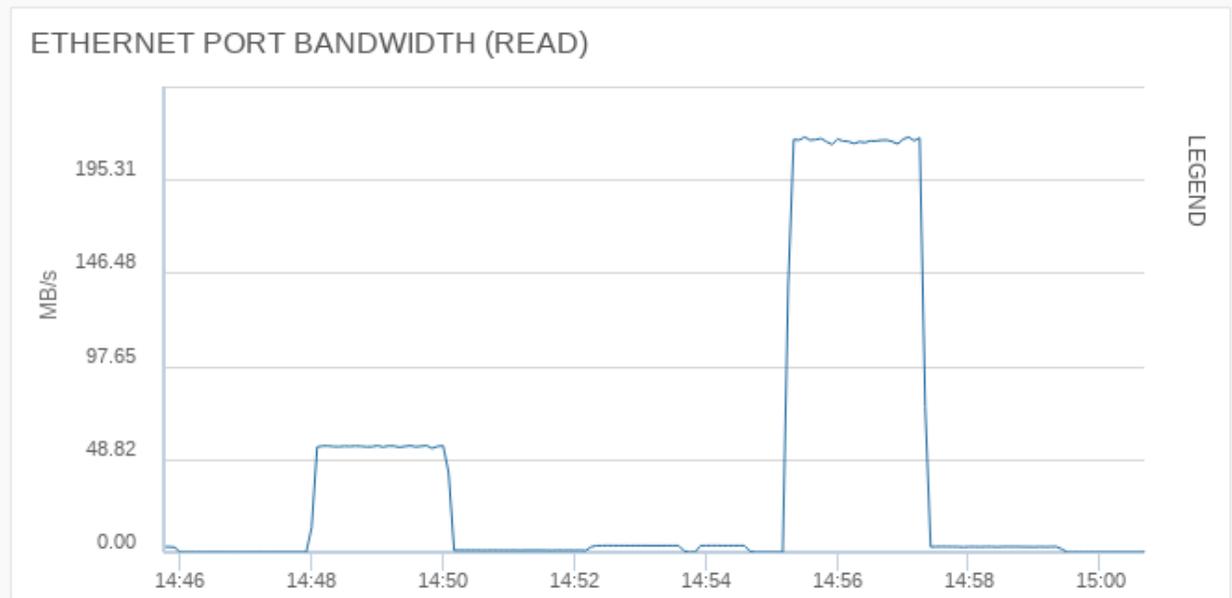


Table 12 Unity Bandwidth during *Diskspd* testing – Read (MB/s)

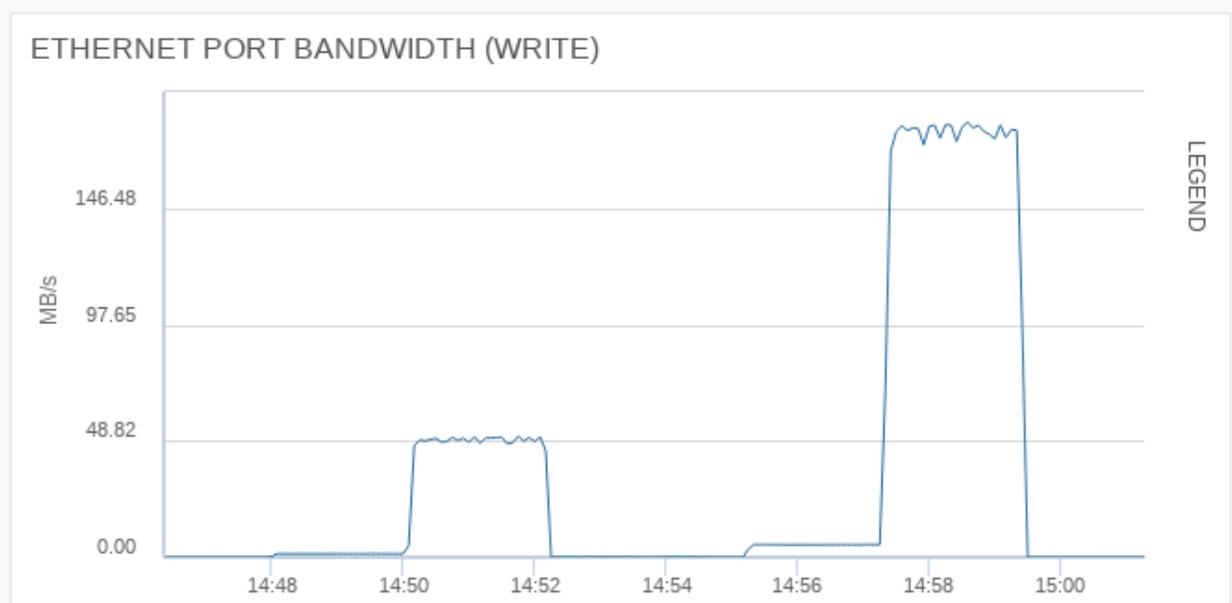


Table 13 Unity Bandwidth during *Diskspd* testing – Write (MB/s)

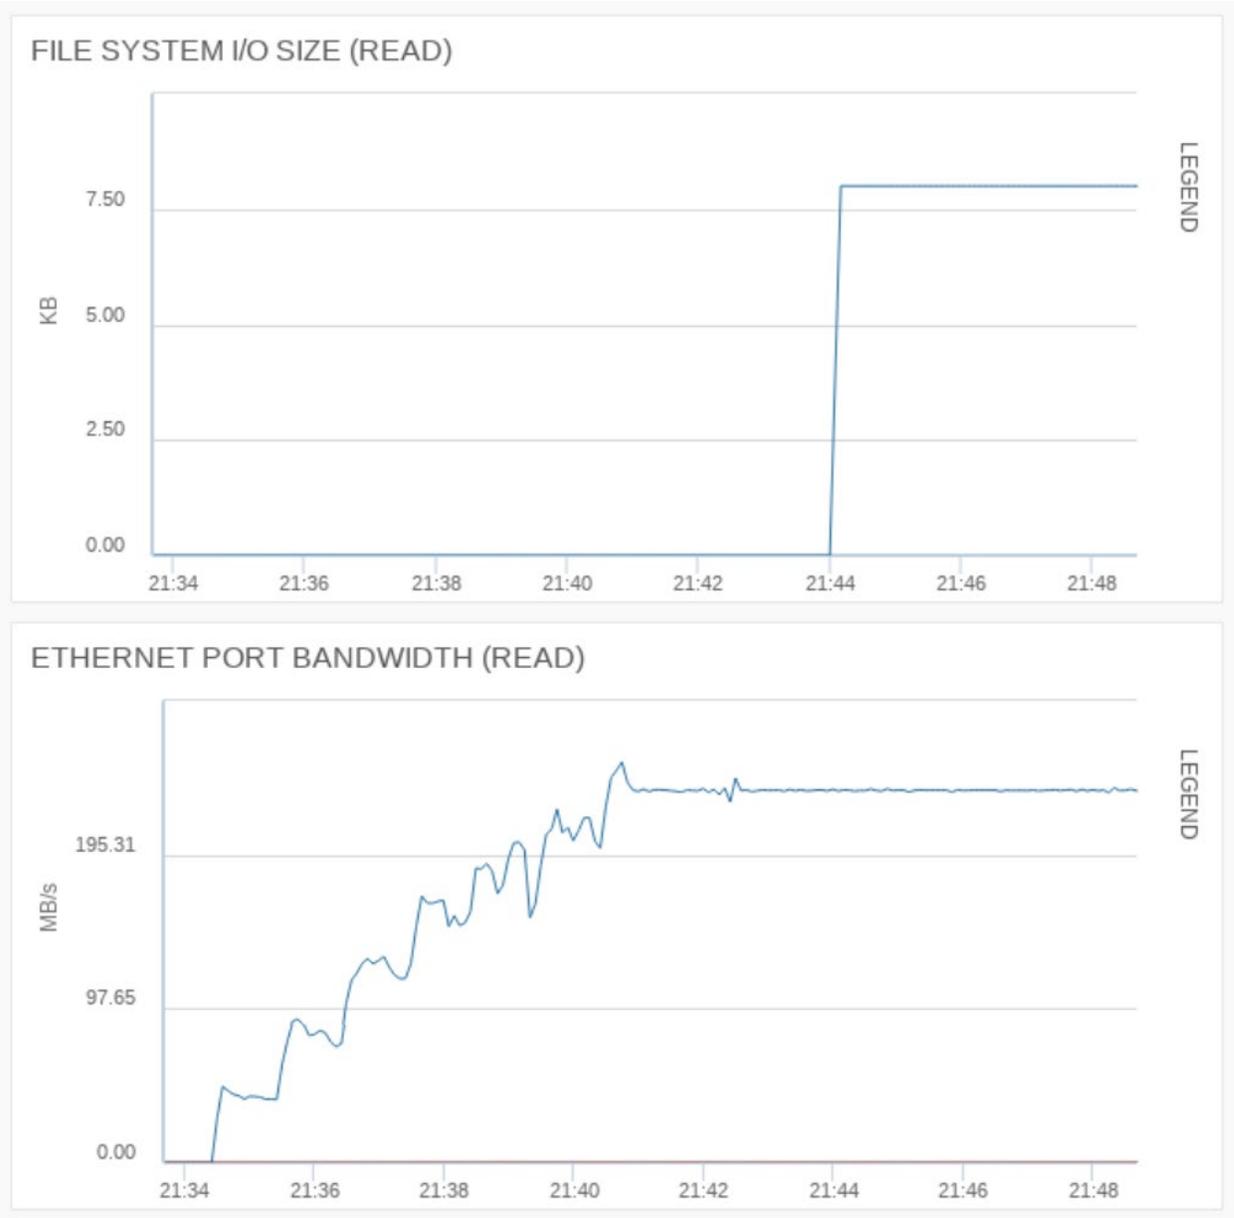


Table 14 Unity IO size and Bandwidth during *ORION* testing

## 4 Microsoft Azure Architecture and Testing

This section will outline the architecture with Azure as tested, and detail the tests performed.

### 4.1 Architecture

The high-level architecture is presented in Figure 2. The architecture consists of:

- Dell EMC Unity storage array, residing in a public cloud-connected colocation facility (Equinix);
- Dell EMC 10Gb networking switch, also in a public cloud-connected colocation facility (Equinix);
- 2Gb Azure ExpressRoute IP network link between a colocation facility and Azure public cloud;
- Virtual machines in Azure that connect to the Unity storage array.

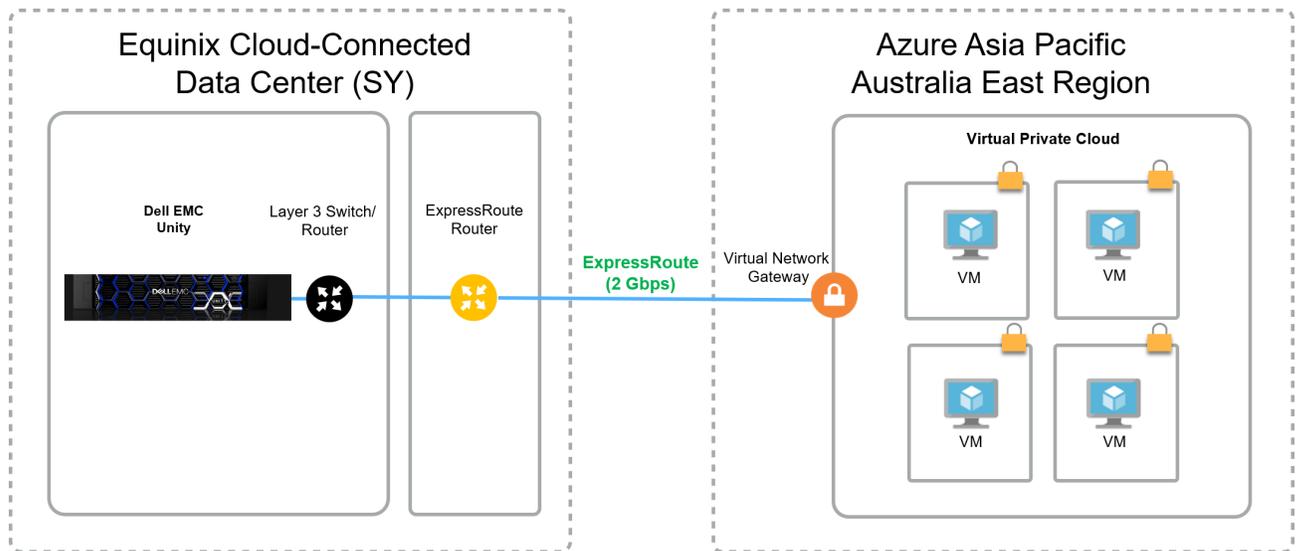


Figure 2 High level Azure architecture as tested

### 4.2 Azure Testing Configuration

This section will outline the as-tested configuration of the Azure virtual machines, the Azure Premium SSDs, the Azure ExpressRoute network link, the Dell EMC Unity storage array, and the Dell EMC network switch.

---

**Note:** All Azure resources (virtual machines, Premium SSDs etc.) were deployed in the (*Asia Pacific*) *Australia East* Region. The Dell EMC resources (Unity storage array, network switch) were deployed at an Equinix Cloud-Connected facility in Sydney.

---

#### 4.2.1 Azure Windows Virtual Machine

The features of the chosen Azure Windows virtual machine are as follows:

- Standard D16s\_v3
- Microsoft Windows Server 2016 with SQL Server 2017 Standard
- 64 GiB memory
- 16 vCPUs
- Accelerated Networking
- Premium Disks supported

## 4.2.2 Azure disks

The Azure disks for attaching to the Windows virtual machine were created and configured as follows:

- A single 1 TiB (P30) Premium SSD was configured as the C: drive
- Five 1 TiB (P30) Premium SSDs were presented to the host, and a single software RAID5 volume was created, mounted as the D: drive.

## 4.2.3 Azure ExpressRoute

When in an Equinix facility, there are two ways to connect the co-located Unity storage array to the Azure virtual machines:

- Equinix Cloud Exchange Fabric (ECX Fabric); or
- Azure ExpressRoute

### Equinix Cloud Exchange Fabric

The major benefit of connecting into public cloud compute via the Equinix Cloud Exchange Fabric is that all the network links into public cloud providers (Amazon AWS, Microsoft Azure, Google Cloud Platform, and Alibaba Cloud) are handled by a single connection out of the customer's co-located equipment. So, for example, the Unity storage array connects to a network switch in the same rack, and it is this network switch that connects once to the Equinix Cloud Exchange Fabric. From there, connections to the required public cloud providers are established and managed by Equinix. More information can be found [here](#).

*“ECX Fabric brings together cloud service providers and users, enabling them to establish affordable, private, high-performance connections within Platform Equinix.”* (Equinix)

### Azure ExpressRoute

To establish a private connection between Azure and the Unity storage array, the customer engages with Azure via the Portal to procure an ExpressRoute connection. More information can be found [here](#).

*“Azure ExpressRoute enables you to create private connections between Azure data centers and infrastructure that's on your premises or in a colocation environment. ExpressRoute connections do not go over the public Internet, offering more reliability, faster speeds, lower latencies, and higher security than typical connections. In some cases, using ExpressRoute connections to transfer data between on-premises and Azure can also yield significant cost benefits. With ExpressRoute, you can establish connections to Azure at an ExpressRoute location (Exchange Provider facility) or directly connect to Azure from your existing WAN network (such as a MPLS VPN) provided by a network service provider.”* (Microsoft)

For the purposes of this testing, an Azure ExpressRoute network link was established. ExpressRoute routing was provided by Equinix.

The Azure ExpressRoute network link is configured are as follows:

- ExpressRoute circuit is 2000Mbps

## 4.2.4 Dell EMC Unity Storage

Deliver consistent high performance for workloads with Dell EMC Unity storage. Dell EMC Unity brings you high performance and low latency across a broad range of SAN and NAS use cases.

Dell EMC Unity redefines simple, modern, affordable, and flexible primary storage. Its modern architecture is designed for flash and rich data services, in a footprint starting at 2U. High-density SSD technology and inline data reduction add to performance and value.

This unified storage comes standard with support for file and block, point-in-time snapshots, local and remote data replication, and built-in encryption. You'll also have deep integration with VMware, Microsoft, and OpenStack ecosystems. Dell EMC Unity is also available as a virtual appliance (Dell EMC UnityVSA), and in the public cloud (Dell EMC Unity Cloud Edition).

Ongoing operation is simple with proactive assistance and automated remote support. Gain intelligent insight into your storage health with CloudIQ, which provides cloud-based proactive monitoring and predictive analytics. More information can be found [here](#).

The configuration of the tested Dell EMC Unity storage array is as follows:

- Dell EMC Unity 400 hybrid storage array
- Dell EMC Unity OE version 4.4.1.1539309879
- 2 x 400GB SSDs
- Configured as RAID 1+1
- SMB share
  - SMB v3
  - 75GB size
  - Mounted as Y: drive on the Windows host
- 2 rack units

## 4.2.5 Dell EMC Network Switch

The S4148F-ON network switch is a one rack unit (RU), full-featured fixed form-factor top-of-rack (ToR) 10/25/40/50/100GbE switch for 10G servers with small form-factor pluggable plus (SFP+), quad small form-factor pluggable plus (QSFP+), and quad small form-factor pluggable (QSFP28) ports.

Dell EMC Networking OS10 Enterprise Edition is a network operating system (OS) supporting multiple architectures and environments. The networking world is moving from a monolithic stack to a pick-your-own-world. The OS10 solution allows disaggregation of the network functionality. More information can be found [here](#).

The configuration of the tested network switch is as follows:

- Model: S4148F-ON
- OS version: 10.4.2.0
- Connectivity to Unity storage array: 10Gb TwinAx
- Connectivity to Direct Connect router: 10Gb LR SFP+
- 1 rack unit

---

**Note:** Jumbo frames were not enabled on any of the components in the Azure testing.

---

## 4.3 The Tests

There are two test categories used to validate the viability of the Dell EMC Unity solution in an Equinix Cloud-Connected Colocation facility:

- Windows disk benchmarking
- Windows SQL benchmarking

### 4.3.1 Windows Disk Benchmarking

To test disk latency and performance, the Microsoft tool *diskspd* is installed on the Azure Windows virtual machine. The version used is 2.0.17, and the tool is available [here](#).

*“A feature-rich and versatile storage testing tool, Diskspd combines robust and granular IO workload definition with flexible runtime and output options, creating an ideal tool for synthetic storage subsystem testing and validation.”* (Microsoft)

The four tests that are performed as part of Windows disk benchmarking are:

- 100% read, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -w0
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -w0
```

- 100% write, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -Z1G -w100
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -Z1G -w100
```

- 75/25% read/write, random, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -r -o4 -d120 -Sh -ft -L -Z1G -w25
```

```
diskspd.exe -b8k -c1G -t4 -r -o16 -d120 -Sh -ft -L -Z1G -w25
```

- 100% read, sequential, 8KB IO size

```
diskspd.exe -b8k -c1G -t4 -si -o4 -d120 -Sh -ft -L -w0
```

```
diskspd.exe -b8k -c1G -t4 -si -o16 -d120 -Sh -ft -L -w0
```

---

**Note:** There are two iterations of each test run. The first iteration sets the number of outstanding IOs at 4 (-o4). The second iteration sets the number of outstanding IOs at 16 (-o16).

---

Additional notes about these tests and results:

- Each test iteration is run five times against each storage type, at various times throughout a 24-hour period.
  - Each test iteration is run for two minutes.
  - The results presented are averaged (Results are added then divided by five).
-

- Two sets of results are produced for each test iteration: Latency and IOs.
- Latency figures given are for the 99<sup>th</sup> percentile.
- IOs figure given is the average number of IOs transferred over the two-minute period.

### 4.3.2 Windows SQL Benchmarking

To test synthetic SQL database performance, the open source tool *HammerDB* is installed on the Azure Windows virtual machine. The version used is 3.1, and the tool is available [here](#).

*“HammerDB is the leading benchmarking and load testing software for the worlds most popular databases supporting Oracle Database, SQL Server, IBM Db2, MySQL, MariaDB, PostgreSQL and Redis.”* (HammerDB)

Notes about this test and results:

- The SQL version is 2017.
- Warehouses created is 16.
- Virtual users is 16.
- Warm up time is 2 mins.
- Run time is 10 mins.
- Each test is run five times against each storage type, at various times throughout a 24-hour period.
- The results presented are averaged (Results are added then divided by five).
- The results given are Transactions per Minute (TPM).
- No database-tuning was performed.

## 5 Microsoft Azure Testing Results

This section provides the results for disk and application benchmarking for Azure Premium SSD storage and Dell EMC Unity storage.

### 5.1 Windows Disk Benchmarking Results

The latency results with outstanding IOs set to 4 are presented in Table 15. Lower numbers are better.

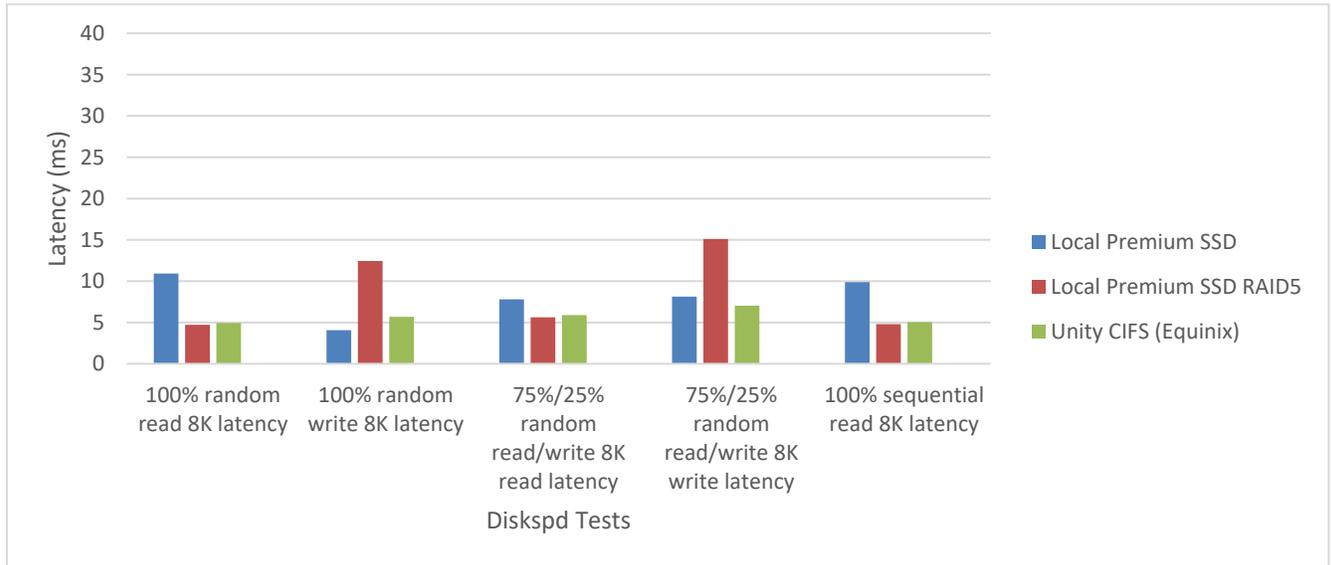


Table 15 *Diskspd* latency observed at host with outstanding IOs = 4

The latency results with outstanding IOs set to 16 are presented in Table 16. Lower numbers are better.

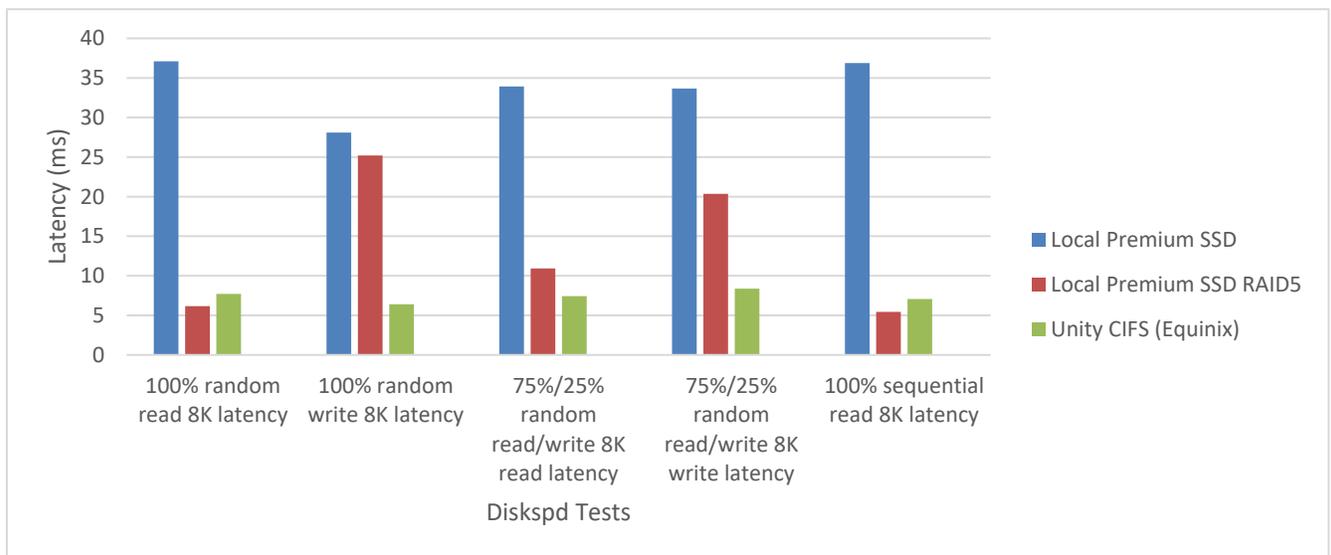


Table 16 *Diskspd* latency observed at host with outstanding IOs = 16

The IOPS results with outstanding IOs set to 4 are presented in Table 17 Higher numbers are better.

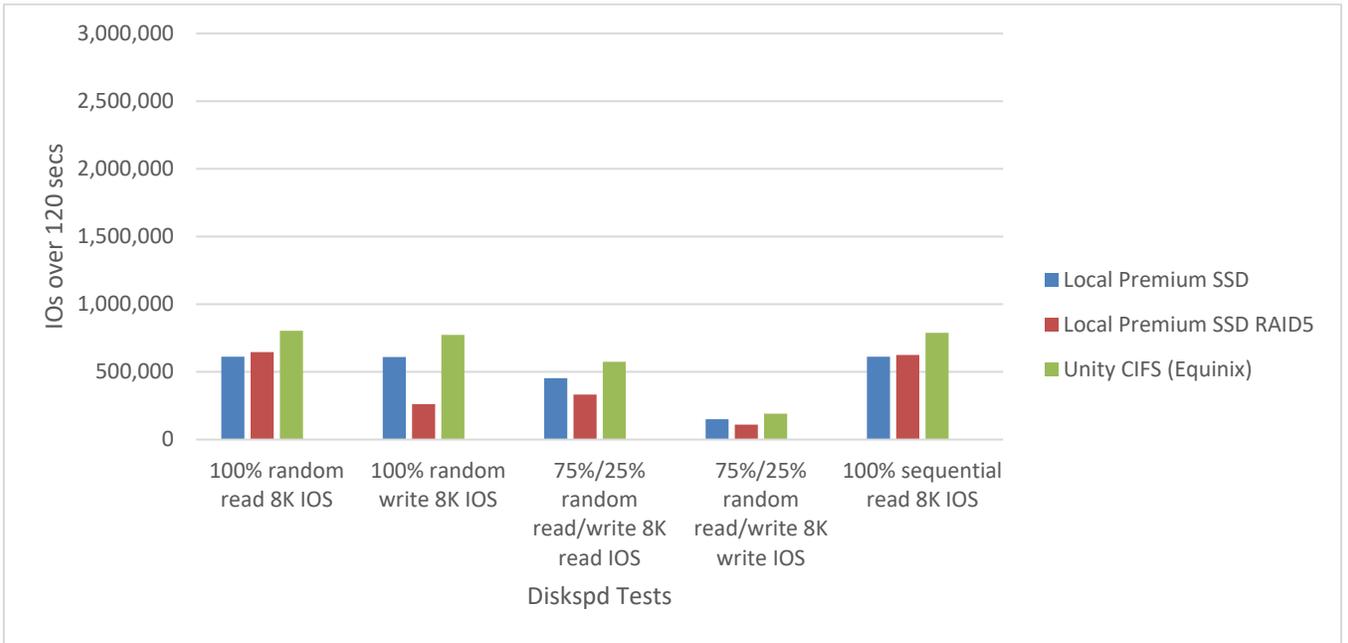


Table 17 Diskspd IOs transferred with outstanding IOs = 4

The IOPS results with outstanding IOs set to 16 are presented in Table 18. Higher numbers are better.

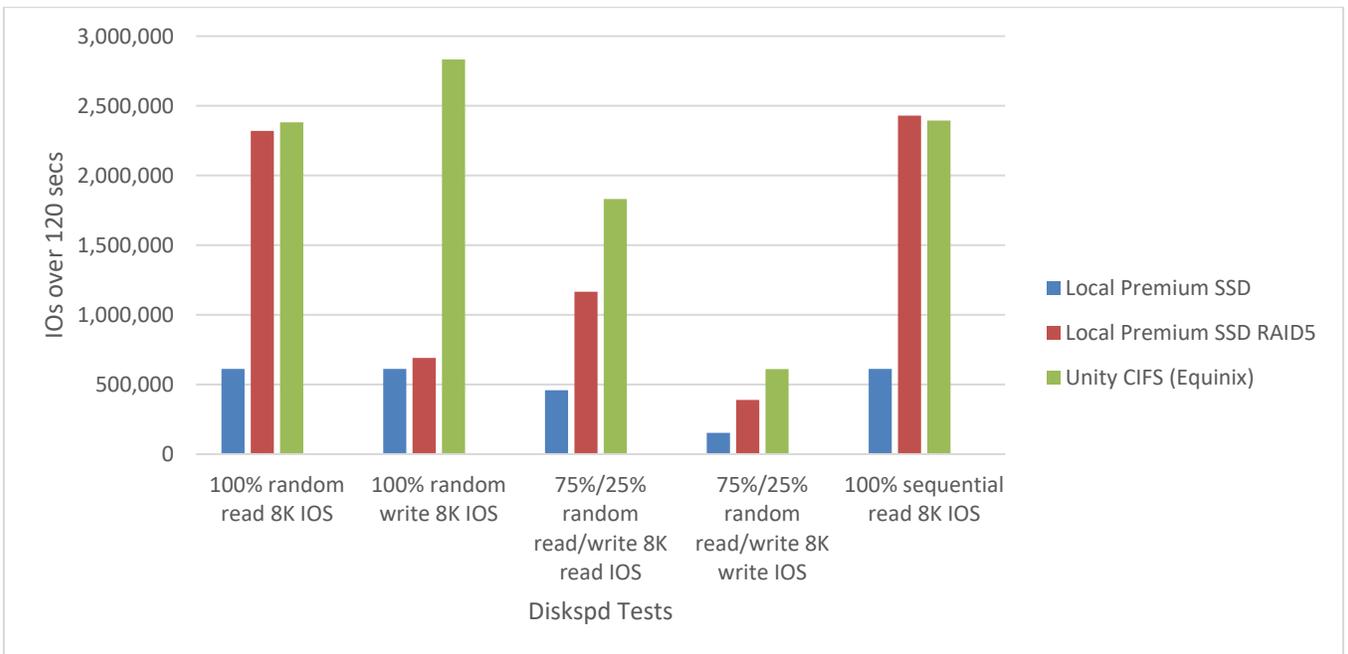


Table 18 Diskspd IOs transferred with outstanding IOs = 16

## 5.2 Windows SQL Benchmarking Results

The synthetic SQL benchmarking results are presented in Table 19. Higher numbers are better.

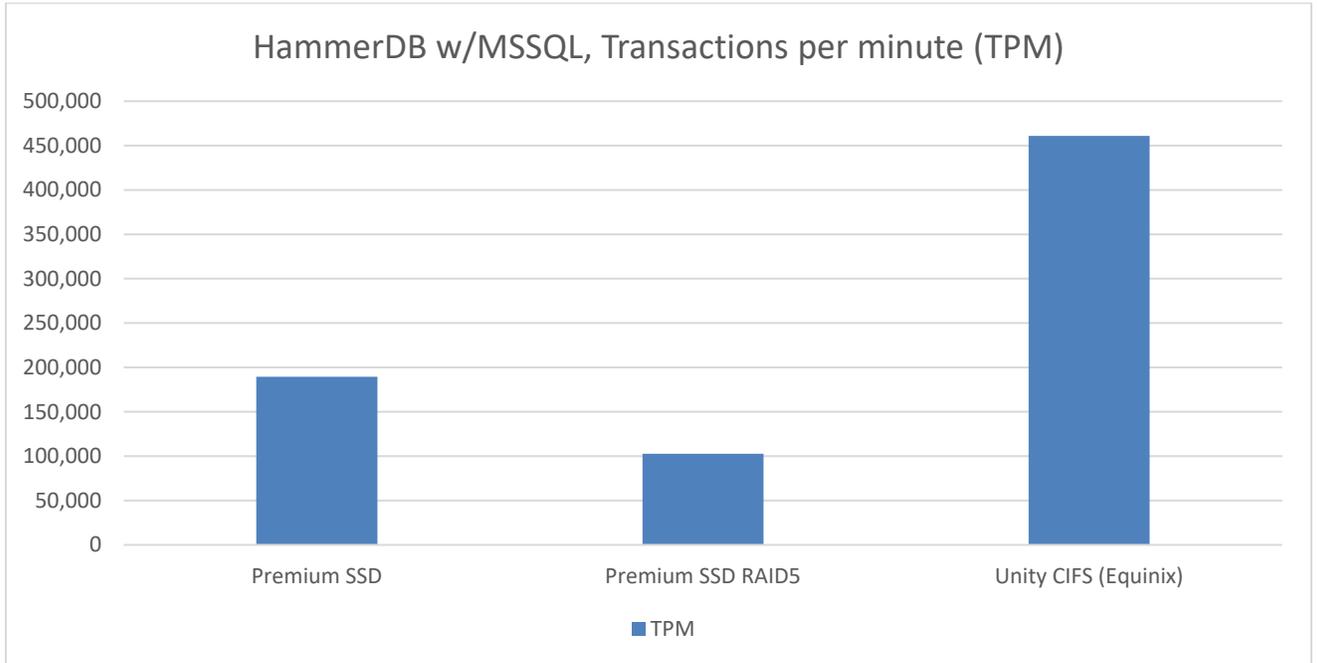


Table 19 *HammerDB* Transactions per Minute

## 6 Conclusion

Connecting Dell EMC Unity storage to public cloud compute resources is not only viable, it can provide a performance improvement to applications running within public cloud infrastructure. And while doing so, issues such as data sovereignty can be addressed.

### 6.1 Observations

For small block workloads within a public cloud compute instance (e.g. databases), the Dell EMC Unity storage array performs at a lower latency, and exhibits higher performance, than both AWS EBS volumes and Azure Premium SSDs. This is seen across both disk and application tests, reads and writes, random and sequential. The benefits of Unity storage are amplified when the number of outstanding IOs increases, a sign that the Dell EMC Unity architecture is well-suited to handling bursty and heavier storage workloads.

### 6.2 Recommendations

- Choose a co-location facility/provider that is as close to the public cloud resources as possible. The further the distance, the greater the latency observed at the host.
- Present storage from the Dell EMC Unity storage array using file protocols (NFS, SMB). Preliminary testing showed iSCSI had higher latencies than the file protocols.
- For AWS, enable jumbo frames at all points in the infrastructure stack to achieve the best performance.
  - You must enable jumbo frames on the Dell EMC Unity storage array, network switch, AWS Direct Connect, and the operating system within the EC2 instance.

### 6.3 Calls to action

- For more information about deploying Unity in a colocation cloud facility, contact your Dell EMC sales or pre-sales representative.
- Review the Unity storage platform at [www.dellemc.com/unity](http://www.dellemc.com/unity)
- Connect with [Dell EMC Shop](#) for the right Unity solution.

## A Technical support and resources

- Dell EMC Unity  
<https://www.dellemc.com/en-au/documentation/unity-family.htm>
- Dell EMC S4148F-ON Network Switch  
<https://www.dell.com/support/home/us/en/04/product-support/product/networking-s4148f-on/>

### A.1 Related resources

- AWS Direct Connect  
<https://aws.amazon.com/directconnect/>
- AWS EBS Volumes  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html>
- AWS EC2 instances  
<https://aws.amazon.com/ec2/>
- Equinix Cloud Exchange Fabric  
<https://www.equinix.com/interconnection-services/cloud-exchange-fabric/>
- Equinix Data Centers and Colocation  
<https://www.equinix.com/services/data-centers-colocation/>
- Fio  
<https://github.com/axboe/fio>
- HammerDB  
<https://www.hammerdb.com/>
- Microsoft Diskspd  
<https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223>
- Microsoft Azure ExpressRoute  
<https://azure.microsoft.com/en-us/services/expressroute/>
- Microsoft Azure Disk Storage  
<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types#premium-ssd>
- Microsoft Azure Virtual Machines  
<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes>
- Oracle ORION  
[https://docs.oracle.com/cd/E18283\\_01/server.112/e16638/iodesign.htm#BABFCFBC](https://docs.oracle.com/cd/E18283_01/server.112/e16638/iodesign.htm#BABFCFBC)