



How to guide

Azure NetApp Files: Getting the Most Out of Your Cloud Storage

Introduction and How-to guide

Chad Morgenstern, NetApp
March 2020

Abstract

This document provides an overview of performance testing strategies and best practices, outlining the field of system performance analysis and its applicability to cloud deployment in Azure NetApp Files.



TABLE OF CONTENTS

1. Introduction	3
2. Workload Tools	3
3. Workload Profiles.....	4
4. Caveats to this Approach	4
5. Setting Expectations	4
6. The Tuning Options.....	8
7. Oracle	13
8. Network Performance: Accelerated Networking	14
9. Volume Bandwidth	15
10. Greatest Throughput; Lowest Latency	17
11. Typical Causes of Poor Performance and Mitigation Factors	19
12. Where to Find Additional Information.....	19

1. Introduction

Over the years, I've been involved with a significant number of storage performance proof of concepts, evaluations, and competitive comparisons. Whether as part of a performance escalation team as a senior member of a support organization, or as a senior engineer in one of our performance engineering groups, knowing how to run, test, and evaluate a storage solution has been critical to the success of my organization—and my team. In this paper, I digest the lessons learned from the field of system performance analysis and experiences with Azure and the Azure NetApp Files environment and use it to guide you, the reader, down the path of application happiness.

Let's talk about what we are going to talk about. To start, we will discuss workload generators: what to consider when selecting a workload generator and when perhaps workload generators are less than ideal. From there we will move on to a discussion of expectations, what they should be—and they can be big—when working with Azure NetApp Files. In that section, expect to learn about the tested limits of a single storage access point as well as the tested limits of a single Azure Virtual Machine. We won't stop there, however; this section continues into the wonderful world of Oracle. Read on to learn about the goodness of Oracle Direct NFS, with a comparison to traditional NFS.

But wait, there's more! Following the expectations section, we take a deep dive into the world of performance testing and tuning. In the end, when you come up for air, you will be as nerdy as I am—and your storage will make you very happy.

Read on my friend.

2. Workload Tools

A synthetic workload generator, often called a benchmarking tool by laymen, is a tool used to generate and measure storage performance. Simplicity and scalability are two big reasons to go with a synthetic generator rather than setting up the actual application. Using tools such as [Vdbench](#) or [fio](#), testing can begin quickly and you can evaluate extremes of load without too much hassle. Though the tool you choose is mostly a matter of preference. Over the years, I have had the most experience with Vdbench; as a result, it's really my go too tool. That said, I understand that ongoing support of Vdbench is questionable and eventually I too will most likely switch over to fio. The strengths of a good generator include:

- The ability to scale out a workload from one to many machines.
- The ability to control data patterns (such as how duplicatable the I/O pattern is).
- Incremental readouts showing status of the test throughout the run.
- The use of variable workload profiles to mimic, for example, logging (in one stream) and database reads and writes (in another).
- The ability to control I/O parallelism using whatever mechanism the tool may choose.

3. Workload Profiles

First and foremost, I can tell you that the closer you can match the evaluation to the actual workload, the better chance you have for a successful outcome. In other words, you must know your application's workload profile. That said, there's often a benefit to evaluating "the four corners" of an environment.

You may have noticed that I always seem to start out testing 8KiB random read/write and 64KiB sequential workloads read and write workloads—that setup is what I mean by "the four corners". Generally speaking, random I/O is performed with smaller operations—4KiB and 8KiB—whereas sequential workloads leverage the largest operation size possible. Typically, applications perform random I/O in small units; 8KiB is typical for databases, and databases are a typical random workload we see. Random I/O is often response-time sensitive, and the time to retrieve each individual operation matters, too. Sequential workloads, on the other hand, are all about time to completion—the time to download the entire file for example. As such, sequential workloads tend to slurp down entire sets of content all at once and do so with the largest operation size possible – thus large block.

In summary, know your application and use your knowledge to simulate its workload. Otherwise: 8KiB random and 64KiB sequential.

4. Caveats to this Approach

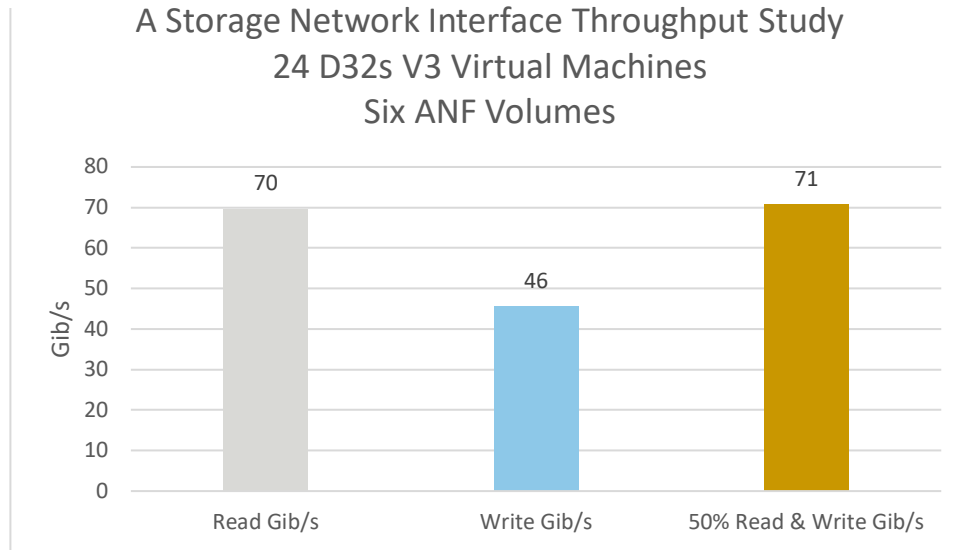
Even the best workload generator cannot mimic the actual application with 100% accuracy, which is another reason you need to know your application. SQL Server will often write its transaction log in bursts, which the tools so far mentioned cannot easily reproduce. Both SAP HANA and EPIC perform writes in long bursts, alternating periods of pure reads followed by intervals of both reads and extreme writes. Oracle may be configured to run over Direct NFS ; as such, it can perform reads and writes in massive parallelism using tens to hundreds of unique NFS network flows. Only Oracle itself can do perform that operation. Each of these scenarios may be tested using application specific benchmarking tools such as the Oracle SLOB or benchmark factories TPC-C for SQL Server, so you're not completely out of luck. While synthetic workload generators may only get you so far, a good synthetic workload generator is useful even in scenarios like the one above (Oracle) to glean the limits of the storage and the network over which it is attached.

5. Setting Expectations

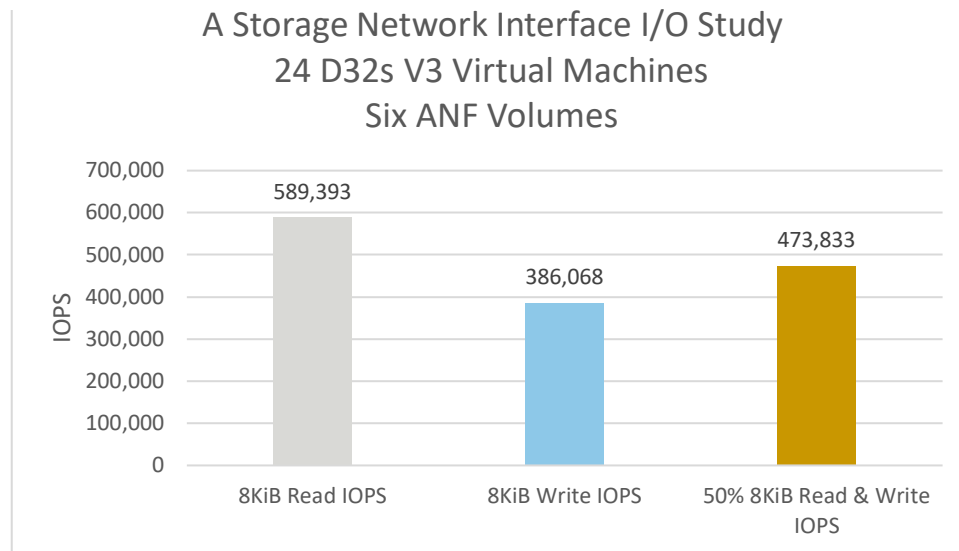
Azure NetApp Files Storage IP Limits

Each Virtual Network is allocated a single delegated subnet, within which one or more storage IP addresses may be assigned by the Azure NetApp Files service.

Through a given storage IP address, testing has shown that Azure Virtual Machines within the VNET may drive up to 71 gibibytes per second (Gib/s) of sequential I/O.



Through a given storage IP address, testing has further shown that Azure Virtual Machines within the VNET may drive up to 600,000 8KiB IOPS as shown below.



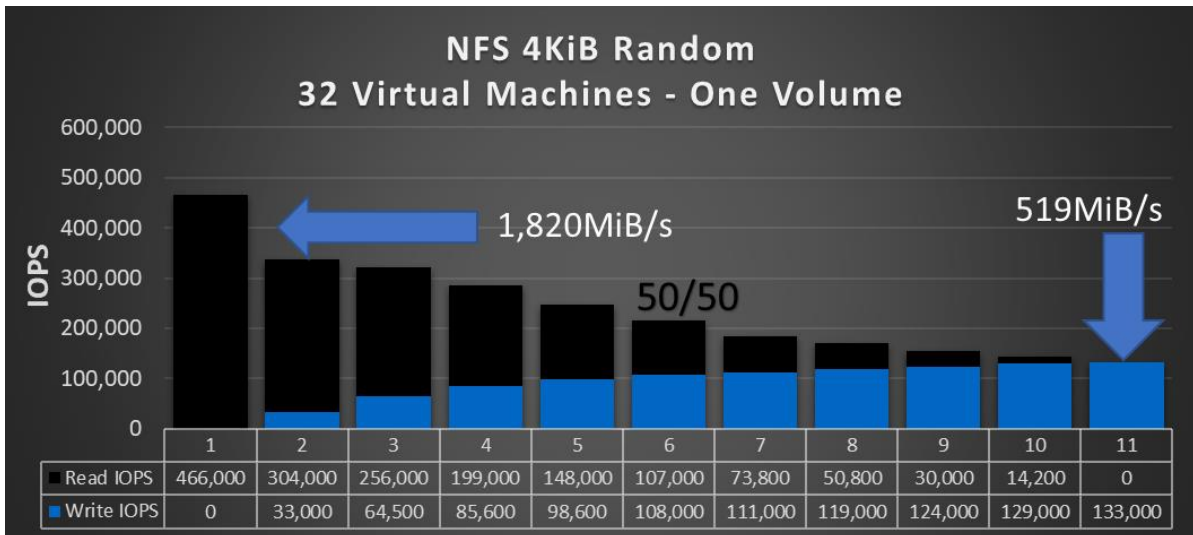
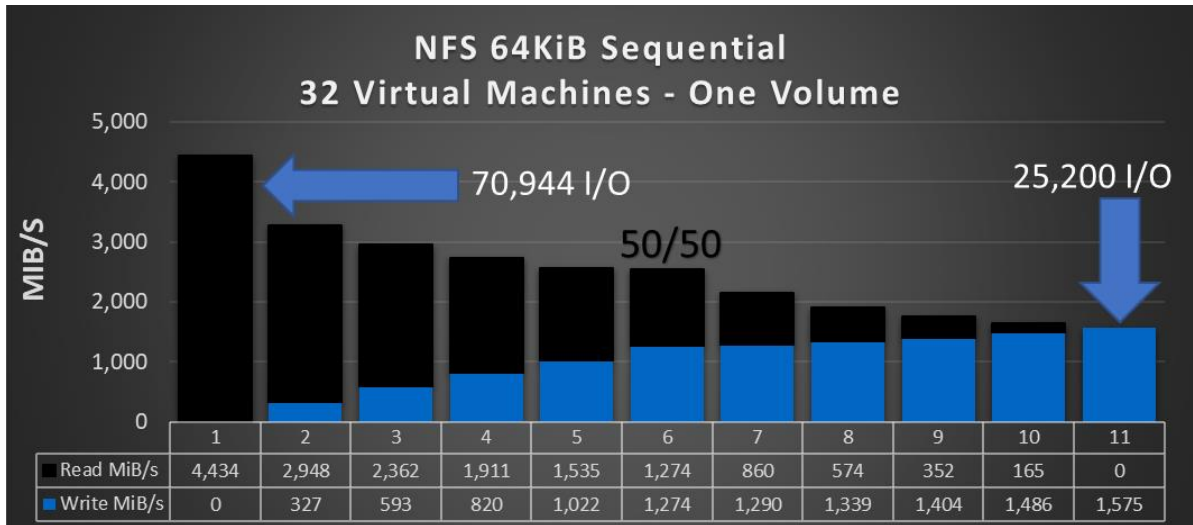
For additional bandwidth, consider peering Virtual Networks, each with its own delegated subnet. Doing so creates a pool of storage end points accessible from any given virtual machine or set of virtual machines. The benefit of this approach is bandwidth or IOPS up to multiples of that shown above. In terms of cost, each GiB transferred between peered networks costs \$0.02, which means that the cost of operations increases with the increased performance.

Azure NetApp Files Single Volume Limits (Scale Out)

The following tables demonstrate the upper limits of a single Azure NetApp Files volume. To reach these limits, 32 separate virtual machines (Sles12SP4) were configured with default mount options. As elsewhere in this paper, the workload was driven using FIO and a 1TiB working set. Unlike elsewhere in this paper, this set of tests focused on range testing, shifting from 100% read to 100% write. In the charts below, the column labeled as “1” represents a 100% read workload, whereas the column labeled “11” represents 100% write workloads. Each additional column represents a 10% change—for example, 90% read and 10% write, and so on.

NFSv3 (Scale Out)

The first graph represents a 64 kibibyte (KiB) sequential workload and the second 4KiB operation size random. As shown, a single volume is capable of handling between ~4,500MiB/s of sequential reads and ~1,600MiB/s of sequential writes. Equally, a volume can achieve anywhere from ~470,000 random reads to 135,000 random writes.

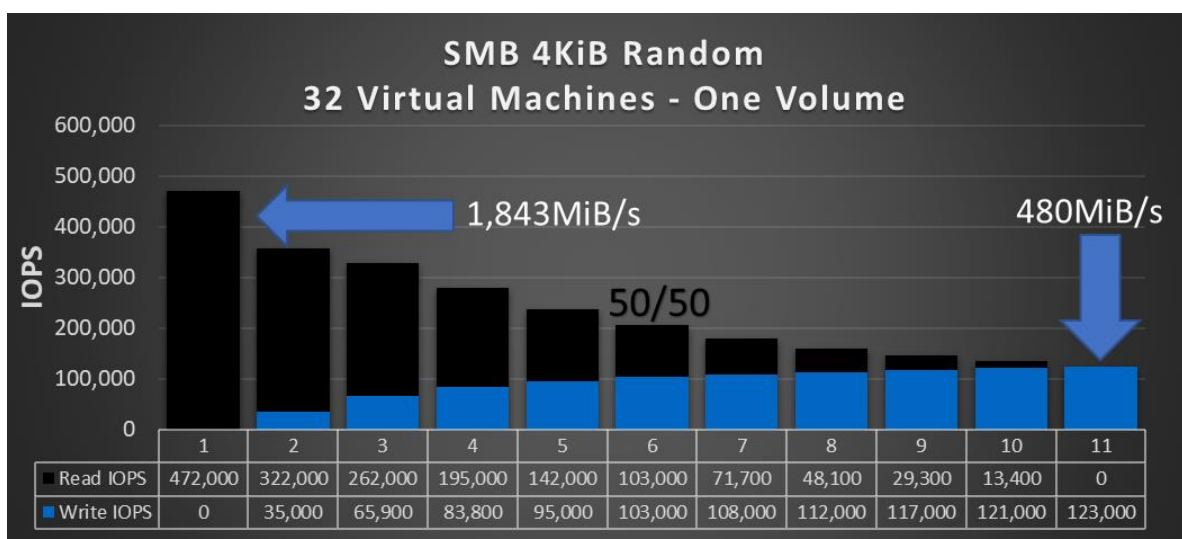
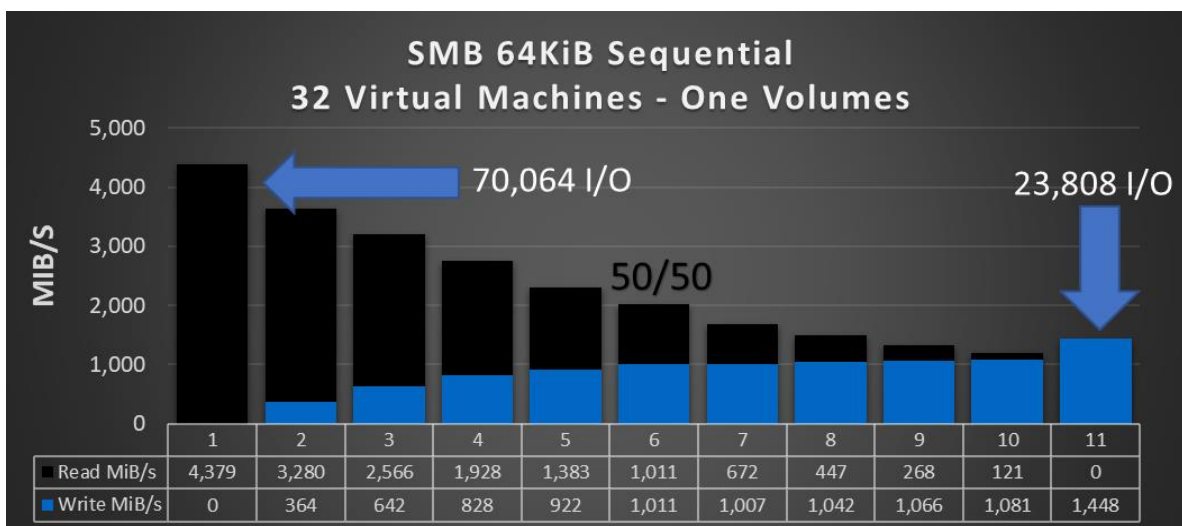


SMB (Scale Out)

For the sake of consistency, and because performing multinode testing with fio with the Windows OS is not supported, the following tests were also performed using 32 SLES12SP4 virtual machines. Though not included in this paper, methodology validation was confirmed by comparing FIO results collected variously against single Windows 2012, Windows 2016, and Windows 2019 single virtual machines against the results captured from single SLES12SP4 Linux VM.

This work is not meant to validate the use of SMB from the Linux operating system. The use of Linux for SMB testing was simply the most expedient choice.

You will notice that the scale out performance of NFS is on par with SMB.



Azure NetApp Files: Single Azure Virtual Machine limits

The following graphs show the single instance bandwidth and I/O limits captured from D32 virtual machines running around the world. They represent pure workloads, namely:

- 100% read
- 100% write
- 100% random
- 100% sequential

Before looking at the graphs, it is important to understand the meaning of the y-axis, which represents application concurrency. Azure NetApp Files offers adjustable bandwidth limits, but it's up to the client/application to consume that bandwidth. Here are the basics:

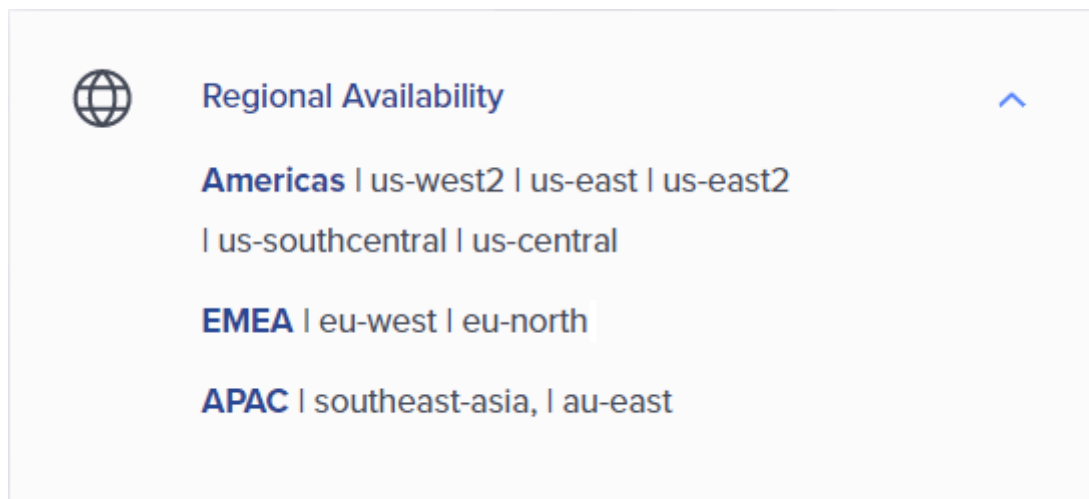
- When bandwidth is consumed, it's called throughput.
- The amount of throughput an application can generate is determined by two factors: concurrency and latency ("round-trip time").

By viewing the concurrency (the y-axis value), you can see that a given amount of concurrency resulted in a graphed amount of throughput. Tools like FIO allow you to determine the upper levels of bandwidth consumable from the environment, while demonstrating the amount of concurrency needed from the application to achieve it. As shown below, lower amounts of concurrency result in lower amounts of throughput. In other words, it's up to the client/application to achieve the numbers shown below.

6. The Tuning Options

In this section, I'll walk you through tuning options to help drive client workloads farther (Oracle dNFS, NFS Nconnect mount option, SMB Multichannel). Then, we'll move onto a study in both load testing and queuing theory.

Of the results shown below, similar tests were run in nine out of 10 Azure ANF Regions; the 10th and newest uk-south region had not come online during the testing cycle.



Each set of tests was repeated 3x to account for run-to-run variability. Rather than use Little's Law (explained later on) to pinpoint the best `iodepth` for both maximum `iorate` with the lowest possible latency, specific `iodepth` ranges were selected in order to compare all regions and zones to find differences across the board.

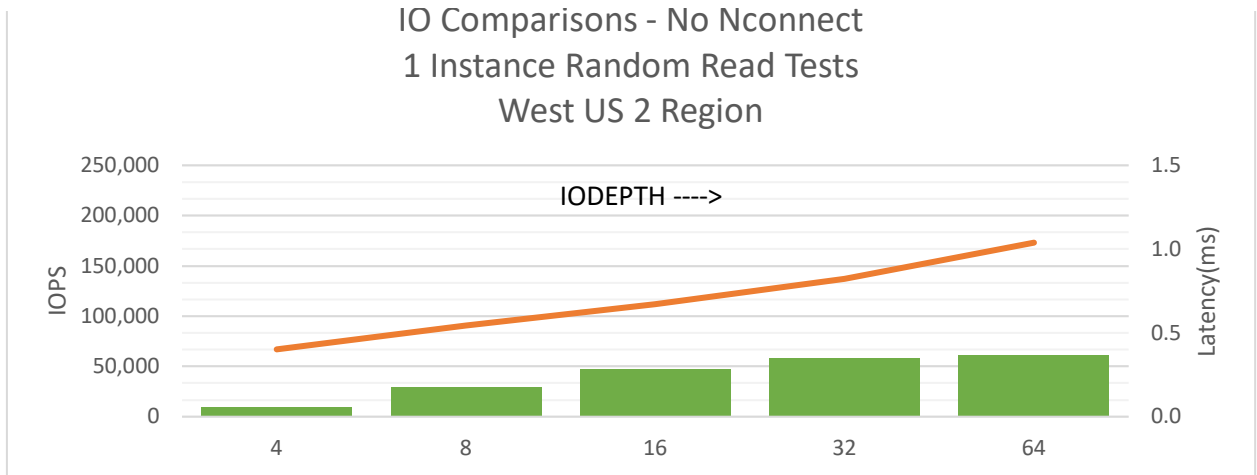
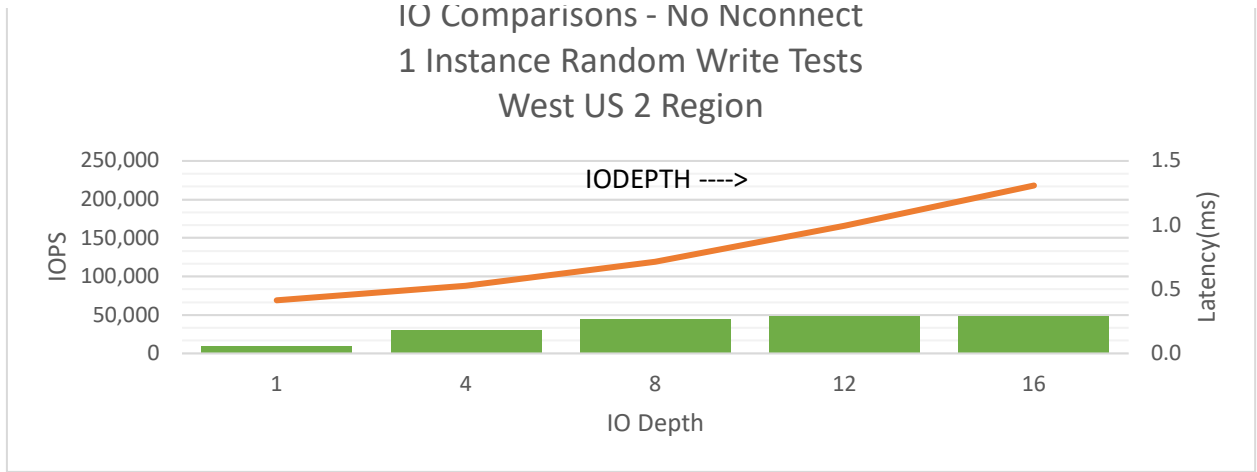
When reading the graphs, notice that the `iorate` tends to flatline while latency continues to increase like a hockey stick (in other words, asymptotically). This latency increase is expected and is primarily due to limits of a single network flow from client to storage.

From region to region and zone to zone, customers may expect similar performance for their workloads, where there are latency differences that may attributed to random Virtual Machine placement even within an availability zone.

Generally speaking, each virtual machine establishes one network connection (network flow) to the storage endpoint when using the NFS protocol storage protocol. As such, the first set of graphs document the single VM/single network flow scenario.

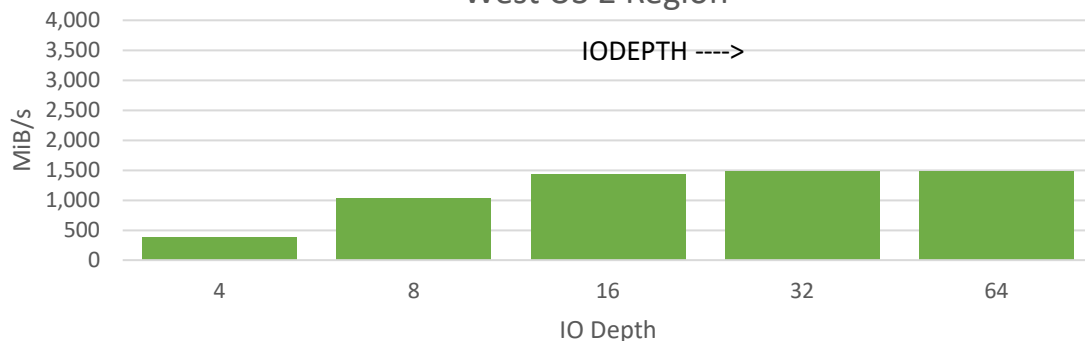
A change has come to the Linux 5.3 kernel, enabling what amounts to single client scale out networking for NFS. NetApp sees the significant performance value of this feature—as shown below—expect an official support statement very early 2020 for NFSv3. **Timelines for Azure NetApp File support of nconnect with NFSv4(.1) are unknown, stand by for more information as we learn more.** Currently adopted by both SLES12SP4 and newer SUSE releases, as well as the Ubuntu19.10 release, the `nconnect` NFS mount option allows for the specification of up to 16 network flows (on which all I/O from each NFS mount will flow). Although implemented in a different way, this feature is similar in concept to both SMB multichannel and Oracle Direct NFS.

Testing in the us-west2 region, a widely used region, has proven that when using traditional NFS mount options, a single D32s_v3 Azure Virtual Machine can drive up to 65,000 4KiB read or 50,000 4KiB write IOPS against a single Azure NetApp Files storage endpoint, regardless of how many volumes are mounted thereon.

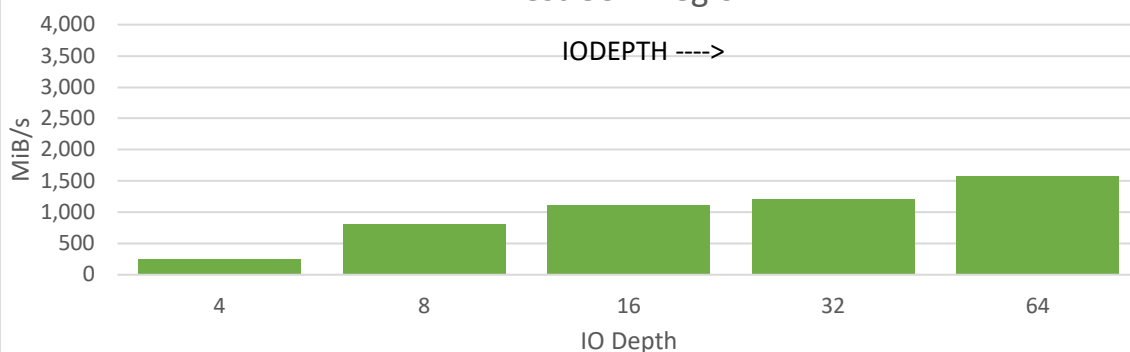


Likewise, testing has shown a 1,500MiB/s upper limit for both sequential read and writes.

Throughput Comparisons - No Nconnect 1 Instance Sequential Write Tests West US 2 Region

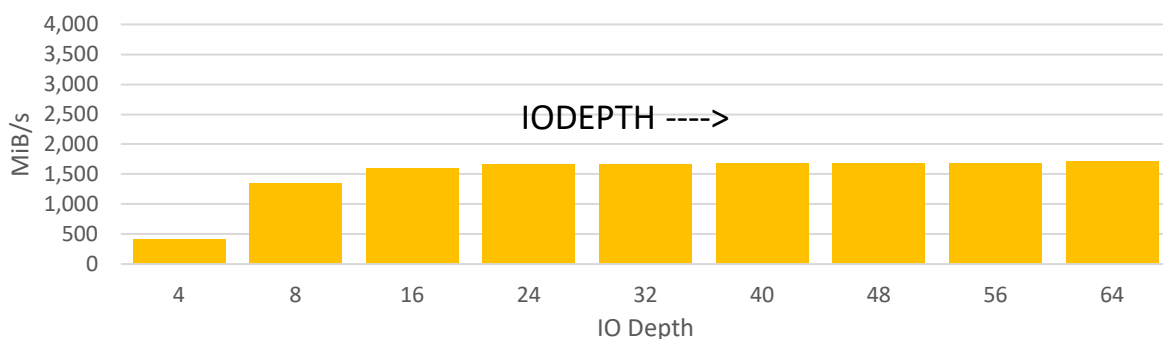


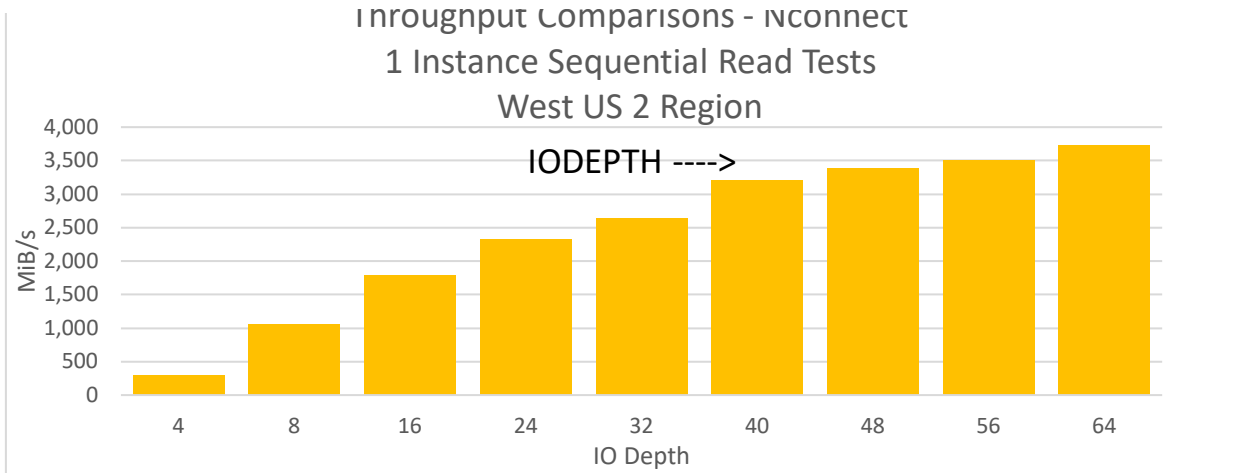
Throughput Comparisons - No Nconnect 1 Instance Sequential Read Tests West US 2 Region



Scaling out network flows—nconnect and NFS or multichannel and SMB—from a single virtual machine significantly increases the storage capabilities of the given instance, as shown below. Presently, these options are supported in SLES12SP4 onward as well as UBUNTU19.10. Further adoption of this feature—present in the Linux 5.3 Kernel—is expected in additional Linux flavors.

Throughput Comparisons - Nconnect 1 Instance Sequential Write Tests West US 2 Region





Notice that, in the graphs above, the sequential write throughput was unimpacted by the introduction of nconnect, while the read throughput increased as compared to the standard mount option configuration by nearly 2x. The write throughput was unimpacted for two reasons:

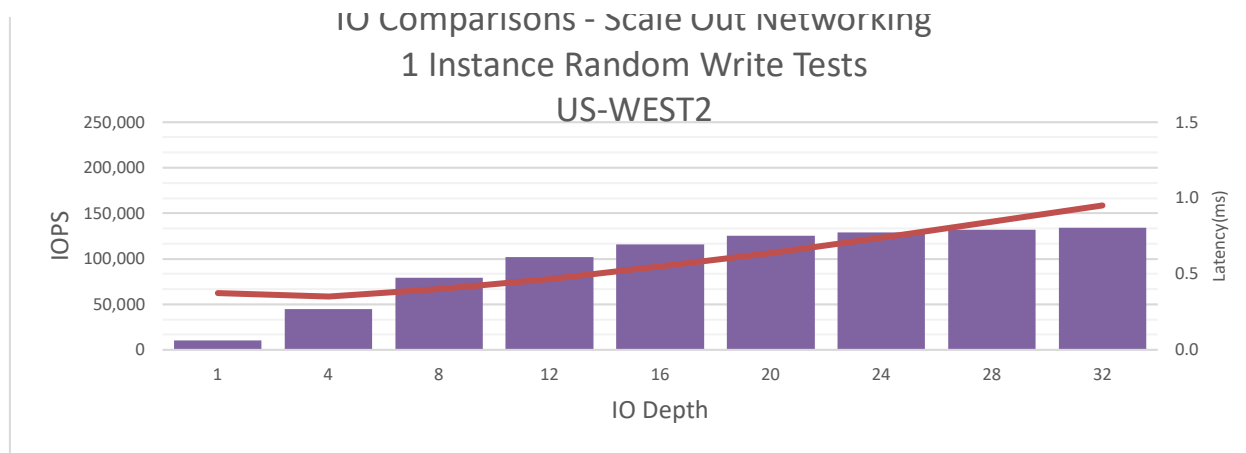
- Testing has shown that a single Azure NetApp Files volume can be driven to no more than ~1500MiB/s.
- The D32 instance type is rate limited on outbound traffic at 16,000Mbps (or 2,000MiB/s).

Because Azure does not impose a VM network bandwidth limit on inbound I/O (reads), the single VM read throughput was unrestricted within Azure.

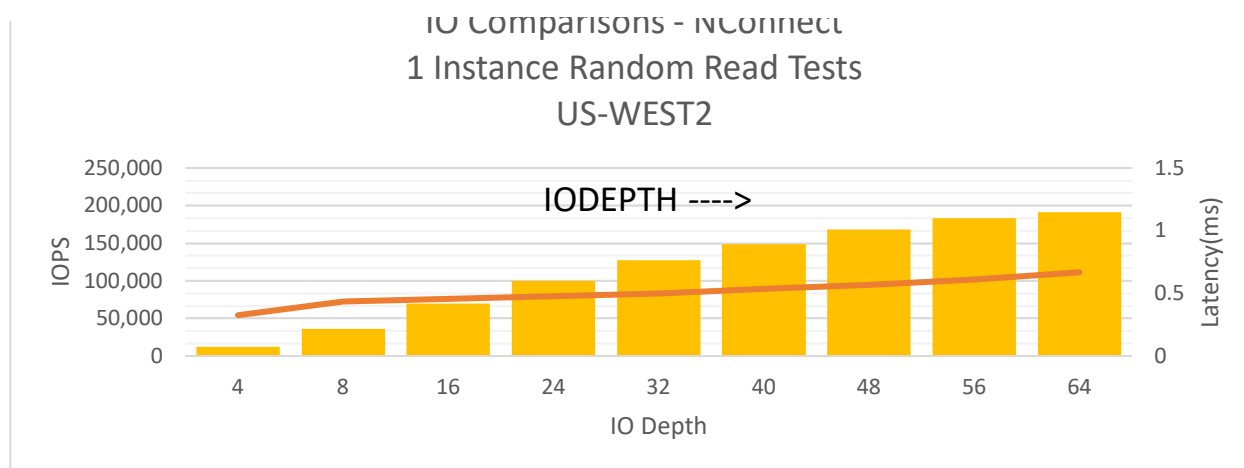
If we look at that from a different angle: Using nconnect for NFS (or SMB Multichannel) increases the storage bandwidth limit of D32s_v3 between 2x and 4.5X above the block storage bandwidth limit of the machine, as shown below.

Size	vCPU	Memory: GiB	Temp storage (SSD) GiB	Max data disks	Max cached and temp storage throughput: IOPS / MBps (cache size in GiB)	Max uncached disk throughput: IOPS / MBps	Max NICs / Expected network bandwidth (Mbps)
Standard_D2s_v3	2	8	16	4	4000 / 32 (50)	3200 / 48	2 / 1000
Standard_D4s_v3	4	16	32	8	8000 / 64 (100)	6400 / 96	2 / 2000
Standard_D8s_v3	8	32	64	16	16000 / 128 (200)	12800 / 192	4 / 4000
Standard_D16s_v3	16	64	128	32	32000 / 256 (400)	25600 / 384	8 / 8000
Standard_D32s_v3	32	128	256	32	64000 / 512 (800)	51200 / 768	8 / 16000
Standard_D48s_v3	48	192	384	32	96000 / 768 (1200)	76800 / 1152	8 / 24000
Standard_D64s_v3	64	256	512	32	128000 / 1024 (1600)	80000 / 1200	8 / 30000

The scale out networking features have an equally impressive effect upon both random writes and random reads. Testing has shown that the addition of `nconnect=16` to the default list of mount options (available from SLES12 SP4 onward and Ubuntu19.10 as well as any distribution using the 5.3 Linux kernel) increases write limits by 230% to 134,000 IOPS.



With the addition of the `nconnect` mount option, random reads increased by over 300% to 190,000 IOPS from the 60,000 IOPS seen previously.



Apart from increased performance, these features decrease compute costs as well. Not even the D64s_v3 is able to deliver the I/O or throughput that the D32s_v3 can. Using `nconnect` or `multichannel`, the D32s_v3 instance drives 200%-300% higher throughput and 150% more I/O than the much larger and more expensive

D64S_v3 is capable of achieving.

Size	vCPU	Memory: GiB	Temp storage (SSD) GiB	Max data disks	Max cached and temp storage throughput: IOPS / MBps (cache size in GiB)	Max uncached disk throughput: IOPS / MBps	Max NICs / Expected network bandwidth (Mbps)
Standard_D2s_v3	2	8	16	4	4000 / 32 (50)	3200 / 48	2 / 1000
Standard_D4s_v3	4	16	32	8	8000 / 64 (100)	6400 / 96	2 / 2000
Standard_D8s_v3	8	32	64	16	16000 / 128 (200)	12800 / 192	4 / 4000
Standard_D16s_v3	16	64	128	32	32000 / 256 (400)	25600 / 384	8 / 8000
Standard_D32s_v3	32	128	256	32	64000 / 512 (800)	51200 / 768	8 / 16000
Standard_D48s_v3	48	192	384	32	96000 / 768 (1200)	76800 / 1152	8 / 24000
Standard_D64s_v3	64	256	512	32	128000 / 1024 (1600)	80000 / 1200	8 / 30000

For more information on service levels and volume expectations, please see “[Azure NetApp Files – Performance So Good You’ll Think You’re On Premises](#)”.

7. Oracle

As discussed in the [orafaq](#) link regarding Direct NFS, **Oracle Direct NFS (dNFS)** is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP). Direct NFS is built directly into the database kernel—just like ASM, which is mainly used with DAS or SAN storage.

A good guideline is to use Direct NFS when implementing NAS storage and ASM when implementing SAN storage.

Direct NFS provides faster performance the operating system's NFS driver can provide because Oracle bypasses the operating system and generates exactly the requests it needs (no user configuration or tuning required). Data is cached just once in the user space, which saves memory (no second copy in the kernel space). Performance is further improved by load balancing across multiple network flows.

Direct NFS is the default option in Oracle 18c and has been the default for RAC for many years.

Grid Infrastructure Installation and Upgrade Guide for Oracle Solaris

- Configuring Storage Device Path Persistence Using Oracle ASMFD
- Using Disk Groups with Oracle Database Files on Oracle ASM
- Configuring File System Storage for Oracle Database
 - Configuring NFS Buffer Size Parameters for Oracle Database
 - Checking TCP Network Protocol Buffer for Direct NFS Client
 - Creating an oranfstab File for Direct NFS Client
 - Enabling and Disabling Direct NFS Client Control of NFS**
 - Enabling Hybrid Columnar Compression on

Enabling and Disabling Direct NFS Client Control of NFS

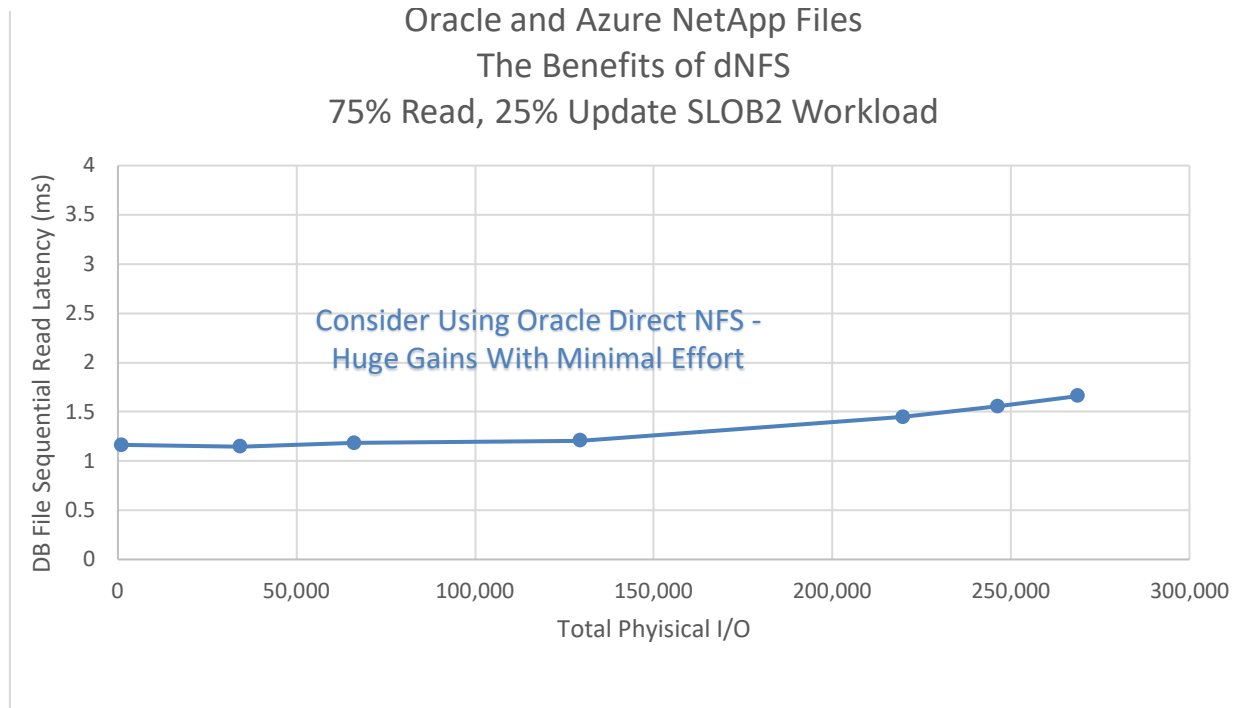
Use these commands to enable or disable Direct NFS Client Oracle Disk Manager Control of NFS.

By default, Direct NFS Client is installed in an enabled state. However, if Direct NFS Client is disabled and you want to enable it, complete the following steps on each node. If you use a shared Grid home for the cluster, then complete the following steps in the shared Grid home:

- Log in as the Oracle Grid Infrastructure installation owner.
- Change directory to `Grid_home/rdbms/lib`.
- Enter the following command:

```
$ make -f ins_rdbms.mk dnfs_on
```

As shown below, by using dNFS (available since Oracle 11g), an Oracle database running on an Azure Virtual Machine can drive significantly more I/O than the native NFS client.



Enabling or disabling dNFS is as simple as running two commands and restarting the database.

Enable: `cd $ORACLE_HOME/rdbms/lib ; make -f ins_rdbms.mk dnfs_on`

or

Disable: `cd $ORACLE_HOME/rdbms/lib ; make -f ins_rdbms.mk dnfs_off`

8. Network Performance: Accelerated Networking

It is a best practice to enable [accelerated networking](#) on Azure Virtual Machines, if supported. Accelerated networking is NOT automatically enabled when resizing upward from a virtual machine if the VM does not support accelerated networking.

9. Volume Bandwidth

An inappropriately sized quota is the number one reason for poor performance in Azure NetApp Files. Please familiarize yourself with the service level documentation posted [here](#) and then read on for a more thorough look at concept of volume bandwidth.

Capacity Pools

Capacity Pools are associated with a service level, and volumes are associated with Capacity Pools. Let's go ahead and create a Capacity Pool called "My-Capacity-Pool" associated with a Premium service level. The Capacity Pool is allocated 500TiB of capacity from which the volumes will be created.

 Add pool  Refresh

NAME	↑↓ CAPACITY	↑↓ SERVICE LEVEL	↑↓
My-Capacity-Pool	500 TiB	Premium	...

Volumes are carved out of Capacity Pools

The volume "volume1" is then created within the Capacity Pool and assigned a quota of 1024GiB as shown below. The volume level bandwidth is derived from the GiB quota. Please see the [link](#) to see the current ranges of quota assignments supported. Collectively, the quotas assigned to all of the volumes within a Capacity Pool cannot surpass the capacity of the pool itself.

NAME	↑↓ QUOTA	↑↓ EXPORT PATH	↑↓ SERVICE LEVEL	↑↓ CAPACITY POOL	↑
volume1	1 TiB	nfs://10.1.0.4/volu...	Premium	My-Capacity-Pool	

How much bandwidth does the volume get?

As stated above, volume bandwidth is derived from the volume quota assignment—a 1024GiB quota gets the 64MiBps of bandwidth (see [here](#) for more info). The math works out like this:

$$\text{Quota} \times \text{Service Level} = \text{Bandwidth}$$

$$1\text{TiB} \times 64 \frac{\text{MiB}}{\text{s}} = \frac{64\text{MiB}}{\text{s}}$$

Please take note: The documentation refers to bandwidth in terms of MiB/s per TiB of capacity. This expression is for the sake of simplicity only. Bandwidth is really allocated in units of KiB/s per of GiB of capacity. If the link above shows 64MiB/s of bandwidth per TiB of capacity, you could just as well say 64KiB/s per 1GiB of capacity.

What if that isn't enough bandwidth?

Let's say that the amount of bandwidth granted at 1TiB capacity is not sufficient; for example, if your application scales out, meaning that 950MiBps may be a better fit. Volume bandwidth assignments may be changed dynamically, but getting the service level right takes a bit of math.

First convert MiBps into KiBps because bandwidth is allocated in units of KiBps.

How many KiBps is 950MiBps?

$$\frac{\text{MiB}}{\text{s}} * \frac{1024\text{KiB}}{\text{MiB}} = \text{Required Bandwidth in KiB}$$

$$\frac{950\text{MiB}}{s} * \frac{1024\text{KiB}}{\text{MiB}} = \text{Required Bandwidth in KiB}$$

$$\frac{972,800\text{KiB}}{s} = \text{Required Bandwidth in KiB}$$

Then, find the required amount of storage quota at the Premium service level.

To find the amount of volume quota needed to satisfy 972,800KiBps, use the following formula. The result will be in GiB, which is perfect. We'll see why in a second.

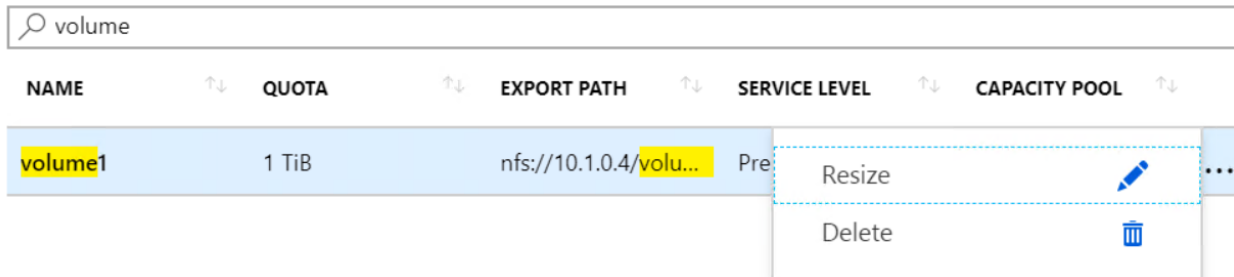
$$\text{Required Bandwidth} \div \frac{\text{Bandwidth}}{\text{GiB}} = \text{Required Quota GiB}$$

$$\frac{972,800\text{KiB}}{s} \div \frac{64\text{KiB}}{s} \text{ per GiB} = \text{Required Quota GiB}$$

$$15,200\text{GiB} = \text{Required Quota GiB}$$

Resize the volume.

Volumes may be resized dynamically as shown below; select the horizontal ellipses button on the right to pull up the resizing menu.



Enter the calculated quota derived in the previous step. If you select a quota that is too large, an error will appear immediately on the screen. Note that the capacity pool has an available capacity within which you must work (as shown below).

10. Greatest Throughput; Lowest Latency

The goal of storage performance testing is to see if you can get the throughput that the application needs. In the event that you're benchmarking two storage technologies, and both meet your application's performance needs, you can consider the performance evaluation done and move required capabilities and total cost of ownership evaluations. If it were always desirable to select the technology with the most bandwidth, the biggest and baddest machine types would always be preferable. But the fact of the matter is that they, simply, aren't.

Little's Law

Synthetic workload generators commonly offer a flag (called, for example, `iodepth` in `fiio`) that allows you to define the amount of outstanding I/O. When used correctly, this flag allows you to home in directly on the amount of I/O you want to hit, which is where Little's Law comes in.

Little's Law is a queuing theory used to understand a queuing system. The formula has three components:

1. Outstanding I/O (a.k.a., queue depth): number of requests submitted to the storage device for processing. When it comes to a workload generator, such as `Vdbench` and `fiio`, each thread represents one outstanding I/O.
2. I/O rate or IOPS: Input and output operations per second.
3. Response time (in other words, latency or roundtrip time): This is roundtrip time for an operation.

The formula behind Little's Law looks like this:

$$\text{outstanding } \frac{I}{O} = \text{IOPS} \times \text{response time}$$

Find the Response Time

To find the response time, turn to the workload generator. The goal is to find the response of the entire I/O operation. Essentially, you're looking for is the latency floor. To that end, you should configure your generator to perform the desired I/O mixture using one thread and one job.

```
[global]
name=fio-test
directory=/mnt/fio/data #This is the directory where files are written
ioengine=libaio #Async threads, jobs turned over to async threads and core moves on
direct=1 #Use directio, if you use libaio and NFS this must be set to 1 enabling directio
numjobs=1 #To match how many users on the system
nrfiles=4 #Num files per job
runtime=10 #If time_based is set, run for this amount of time
group_reporting #This is used to aggregate the job results, otherwise you have lots of data to parse
time_based #This setting says run the jobs until this much time has elapsed
stonewall
bs=64K
rw=rw
#rw=randrw sequential io, choose randrw for random io
rwmixread=100 sequential io, choose randrw for random io
iodepth=1 #<-- Modify this to get the i/o they want (latency * target op count)
size=10M #Aggregate file size per job (if nrfiles = 4, files=2.5GiB)
#ramp_time=20 #Warm up
[test]
```

This process can be short, as the screenshot below shows. In the image below, take note of the 1.9ms average response time for queue depth (outstanding I/O of 1). Remember this response time for later.

```
fio-sles15pl-centrall-a-4r24:/opt/fio-parser # vim ./fio-config
fio-sles15pl-centrall-a-4r24:/opt/fio-parser # fio ./fio-config
test: (g=0): rw=rw, bs=(R) 64.0KiB-64.0KiB, (W) 64.0KiB-64.0KiB, (T) 64.0KiB-64.0KiB, ioengine=libaio, iodepth=1
fio-3.3
Starting 1 process
Jobs: 1 (f=4): [ 0: pid=2596: Tue Oct 29 19:09:09 2019]
test: (groupid=0):
  read: IOPS=100, BW=32.6MiB/s (34.2MB/s) (326MiB/sec)
  slat (usec): min=8, max=100, avg=27.21, stdev= 8.23
  clat (usec): min=1496, max=53025, avg=1882.46, stdev=777.46
  lat (usec): min=1518, max=53042, avg=1909.96, stdev=777.28
```

Before moving on, we will need to convert the response time shown above into seconds, which is the response time unit required by Little's Law. What we will find that $1.909\text{ms} = \mathbf{0.001909\text{sec}}$.

$$\text{response time (s)} = \text{response time (us)} \times \frac{1\text{sec}}{1000\text{us}}$$

$$\text{response time (s)} = \mathbf{1909}(\text{us}) \times \frac{1\text{sec}}{1000\text{us}}$$

$$\text{response time} = \mathbf{.001909\text{sec}}$$

Convert MiB/s to IOPS

Plug in the desired number of IOPS and the observed response time to determine the value you will enter for the thread count. Let's look at an example wherein **200MiB/s** of sequential reads are needed.

First, we need to convert throughput into an I/O rate.

- I/O size: The NetApp storage service has a maximum I/O size of 64KiB. The maximum I/O for NetApp storage could change in the future.

The conversion of throughput in MiB/s to IOPS looks like the example below; note that the process involves first converting MiB/s to KiB/s, which is the unit of I/O size. The result of the following workflow is an I/O rate of 3,200 operations/sec.

$$IOPS = \frac{\text{throughput} \frac{\text{MiB}}{\text{sec}} \times 1024 \frac{\text{KiB}}{\text{MiB}}}{64\text{KiB}}$$

$$IOPS = \frac{\mathbf{200} \frac{\text{MiB}}{\text{sec}} \times 1024 \frac{\text{KiB}}{\text{MiB}}}{64\text{KiB}}$$

$$IOPS = \mathbf{3,200} \frac{\text{operations}}{\text{sec}}$$

Applying Little's Law

Having identified the minimum response time (0.001909s) and the desired I/O rate (**3,200 operations/sec**), we are now ready to put Little's Law to work. When we plug in the variables, we see that **6.11** threads are needed to achieve the desired I/O rate, which amounts to 200MiB/s.

$$\text{outstanding} \frac{I}{O} = IOPS \times \text{response time}$$

$$\text{outstanding} \frac{I}{O} = \mathbf{3,200} \frac{\text{operations}}{\text{sec}} \times \mathbf{0.001909\text{sec}}$$

$$\text{outstanding} \frac{I}{O} = \mathbf{6.11}$$

As the screen shot below shows, a queue depth of 6 results in **3,200 IOPS** and **201MiB/s** of throughput. Spot on!

```

mchad@mchad-jump-host: ~
fio-sles15pl /opt/fio-parser # vim fio-config
fio-sles15pl /opt/fio-parser # fio fio-config
test: (g=0): rw=rw, bs=(R) 64.0KiB-64.0KiB, (W) 64.0KiB-64.0KiB, (T) 64.0KiB-64.0KiB, ioengine=libaio, iodepth=6
fio-3.3
Starting 1 process
obs: 1 (f=4): [R(1)] [70.0%] [r=201MiB/s, w=0KiB/s] [r=3219, w=0 IOPS] [eta 00m:03s]
  
```

Understanding Little's Law

Using Little's Law, we were able to efficiently home in on the lowest possible thread count required to hit our throughput target. What's more, we now know the lowest possible latency at which the desired throughput can be achieved. There's more to Little's Law than configuring a workload generator, however.

Keeping in mind that there are only two variables associated with this queuing formula:

- 1) Outstanding I/O (queue depth or concurrency)
- 2) Response time (latency or response time)

I don't consider I/O rate to be a variable; I/O rate is more of an effect than a cause. If you tweak either of the two variables, your I/O rate will change.

Little's Law explains mathematically that you can overcome the effects of latency (decreased I/O rate) by increasing concurrency, which results in an increased I/O rate. Equally, if you can manage to decrease the latency while keeping concurrency fixed, your I/O rate will also go up. It's easy to prove: it's in the math. Go ahead and try it yourself.

11. Typical Causes of Poor Performance and Mitigation Factors

- ❖ Fully consumed volume bandwidth: Check the amount of bandwidth made available to the volume as provisioned—incorrectly configured service levels are the number one cause of poor performance. For more information on this subject, see [“Azure NetApp Files – Performance So Good You'll Think You're On Premises”](#).
- ❖ Fully consumed network flow bandwidth: Check the limits in the expectation section of this paper for per flow limits. If the per flow limit has been reached, scale out.
- ❖ Insufficiently sized receive window in a higher latency environment: As round-trip time increases, corresponding increases are required in TCP Window Size. Think Little's Law. Check the values of the following and adjust as needed:
 - net.core.rmem_max,
 - net.core.rmem_default
 - net.ipv4.tcp_rmem
 - The formula to identify the bandwidth available to the virtual machine is:
 - $Maximum\ consumable\ bandwidth\ in\ bytes = \frac{TCP\ window\ size\ in\ bits}{Response\ Time\ in\ seconds} \div 8 \frac{bits}{bytes}$
 - The formula to identify the optimal TCP window size in bytes is:
 - $TCP\ windows\ size\ in\ bytes = \frac{Bandwidth \frac{bits}{second} \times Response\ time\ (s)}{8\ bits}$
- ❖ Sub optimal application code: Address the code
- ❖ None of the above: contact support

12. Where to Find Additional Information

To learn more about Azure NetApp Files, please go to cloud.netapp.com.