SoftNAS Cloud® Performance Evaluation on AWS

October 25, 2016

Contents

SoftNAS Cloud® Overview 3
Introduction 3
Executive Summary 4
Key Findings for AWS: 5
Test Methodology 6
Performance Summary and Sizing Recommendations 8
EBS Tests 8
S3 Tests 9
Other Observations 9
Key Results 10
Conclusion 14

SoftNAS Cloud® Overview

SoftNAS Cloud® is a software-defined, enterprise-class NAS filer delivered as a virtual appliance for public, private and hybrid clouds. It supports multiple protocols such as NFS, CIFS/SMB, iSCSI, and AFP and makes them available to applications as a POSIX compliant file system. With Active Directory and LDAP integration, SoftNAS Cloud runs in the Amazon AWS, Microsoft Azure, CenturyLink Cloud public cloud platforms, as well as, in your own data center, on VMware vSphere. SoftNAS Cloud is hardware agnostic and requires neither proprietary hardware nor Users having prior storage experience. SoftNAS Cloud can easily scale up or scale down, so a customer only pays for what cloud resources they use.

Introduction

SoftNAS®, Inc. commissioned this 3rd party, independent testing effort, to gather impartial technical insights to help customers with sizing of their cloud storage options. As with any benchmarking tests, the combinations will vary widely, depending on a multitude of variables used in real-world environments. **The information presented in this white paper provides general guidance and should not be taken as an indication nor guarantee of absolute performance metrics.** Customers should always use their own workloads (running on a cloud platform) to ensure satisfactory results. SoftNAS Solutions Architects (SAs) can offer their valuable and extensive expertise to help assess the right infrastructure to match a customer's workloads and cost objectives.

The general guidance covered within this document summarizes a comprehensive set of tests conducted using SoftNAS Cloud. The tests conducted, involved varying cloud platform instance sizes (compute), storage options and networking speeds. Options representative of entry, mid-level, and higher-end configurations were used to provide general guidance. SoftNAS sales experts can provide recommendations on many other options available.

Using the cloud for infrastructure, offers unique capabilities not easily accomplished with traditional on-premises storage options. The flexibility to choose different combinations of compute that host the SoftNAS Cloud virtual storage controller and the options for different storage media, makes cloud options extremely flexible and cost effective. Hence the reason for commissioning this impartial third-party performance testing project to help with the cloud platform choices available to customers.

Unlike conventional on-premises infrastructure, once you purchase a storage appliance, it remains fairly fixed until it is ultimately replaced. Using cloud infrastructure, enables a customer to make flexible and dynamic changes throughout an application workload's

lifecycle. Compute and storage combinations can be easily adapted as use cases change, decrease, or expand. Unlike traditional storage, these changes can occur without creating a new capital expense request; you simply switch the cloud infrastructure combinations to those better suited for your current workloads. The cloud architecture remains malleable and adjustments can occur without any data loss – unlike an on-premises solution.

Cloud infrastructure as a service is constantly changing, evolving, and adapting to new technology as it becomes available. Cloud platform service providers must remain competitive with their offerings, so they are always investing in the latest technology updates (instead of you needing to do so). This fluidity of offerings enables customers to take advantage of the latest computing infrastructure without all the costs and risks. Where else can you get the benefits of running (or switching) to the latest CPUs; getting the best cost or highest performing storage media; or using the latest networking connectivity without a new CapEx budget to overhaul your data center? SoftNAS will always have these evolving resources available to constantly improve our software-based storage solutions.

Executive Summary

Performance of SoftNAS Cloud on AWS is a result of interplay between several variable choices as outlined in Fig 1.

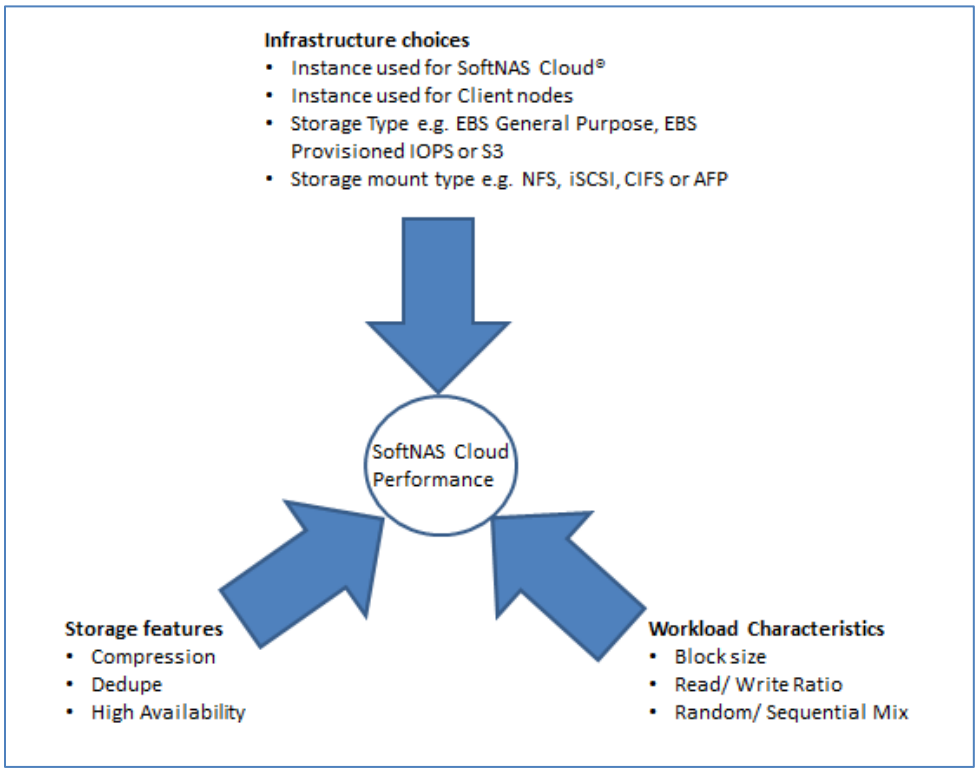

Fig 1: Factors impacting SoftNAS Cloud performance on a public cloud

This White Paper provides users with guidelines on the potential performance results based on the different combinations of controlled tests that were carried out – to characterize the system performance on AWS. Three instance types: from low latency, high I/O instances to high latency and low I/O instances, were selected for these tests. SoftNAS Cloud was installed on each of those instances and then S3 and EBS were used as the storage backend for each instance type. System performance was measured using FIO (Flexible I/O Tester Synthetic Benchmark) for different workload types. This tool was chosen because of its capability to evaluate performance with all of the features available in SoftNAS Cloud (*e.g.,* deduplication, compression, various file protocols, iSCSI and high availability).

Key Findings for AWS:

- Instance sizes have a significant impact on SoftNAS Cloud performance. Surprisingly though, the impact is due more to the cloud platform instance's network and storage throughput limits, than due to its CPU, RAM or SSD characteristics. **Users will need to ensure that the cloud platform instance they choose for SoftNAS Cloud, can meet the network and throughput requirements of the workload, specified by IOPS x Block Size.**

- As important as it is to choose the right AWS instance for SoftNAS Cloud, it is equally important to choose the right instance for client nodes as well. Since client nodes interact with SoftNAS Cloud over the network, the peak storage throughput a client can achieve is limited by its network bandwidth. Ideally, both SoftNAS Cloud and client node should be of the same size.

- SoftNAS Cloud was found to significantly enhance the performance of reads due to its caching feature – especially for sequential reads. Thus, SoftNAS Cloud is a very good fit for workloads such as OLAP, Hadoop, video streaming, log aggregation, etc. - that are read intensive and sequential in nature.

- On a standalone basis, EBS Provisioned IOPS provide double the peak IOPS, as compared to, EBS General Purpose IOPS and is therefore priced at a significant premium. However, when both these EBS options are used with SoftNAS Cloud, EBS General Purpose IOPS performed on par with EBS Provisioned IOPS for read operations. This improvement in performance was due to caching of data by SoftNAS Cloud for reads. **The implication for SoftNAS Cloud users, is that for read intensive workloads, the users can consider using EBS General Purpose IOPS with SoftNAS Cloud instead of using the more expensive EBS Provisioned IOPS.**

- When features such as dedupe and compression are used in SoftNAS Cloud, the performance of write operations was found to drop by 25-30%. This behavior is due to SoftNAS Cloud having to dedicate part of CPU and RAM capacity for inline processing for these features. **As a rule of thumb, SoftNAS Cloud customers should use an instance size one level higher than the base, when these features are enabled.** Furthermore, if replication is also switched on, in addition to the dedupe and compression features, SoftNAS Cloud customers should use larger instance sizes such as c3.8xlarge and above, for high performance workloads.

Test Methodology

- Tests were conducted on SoftNAS Cloud in AWS with various combinations of control parameters:
  - Three instance types: m3.xlarge, m3.2xlarge and c3.8xlarge
  - Three storage options: EBS General Purpose IOPS, EBS Provisioned IOPS and S3
  - Three volume mount options in SoftNAS Cloud: NFS, iSCSI and CIFS
  - Dedupe On/Off for a volume in SoftNAS Cloud
  - Snap Replication On/Off in SoftNAS Cloud

- Topology used for the testbed is outlined in Fig 2 below:

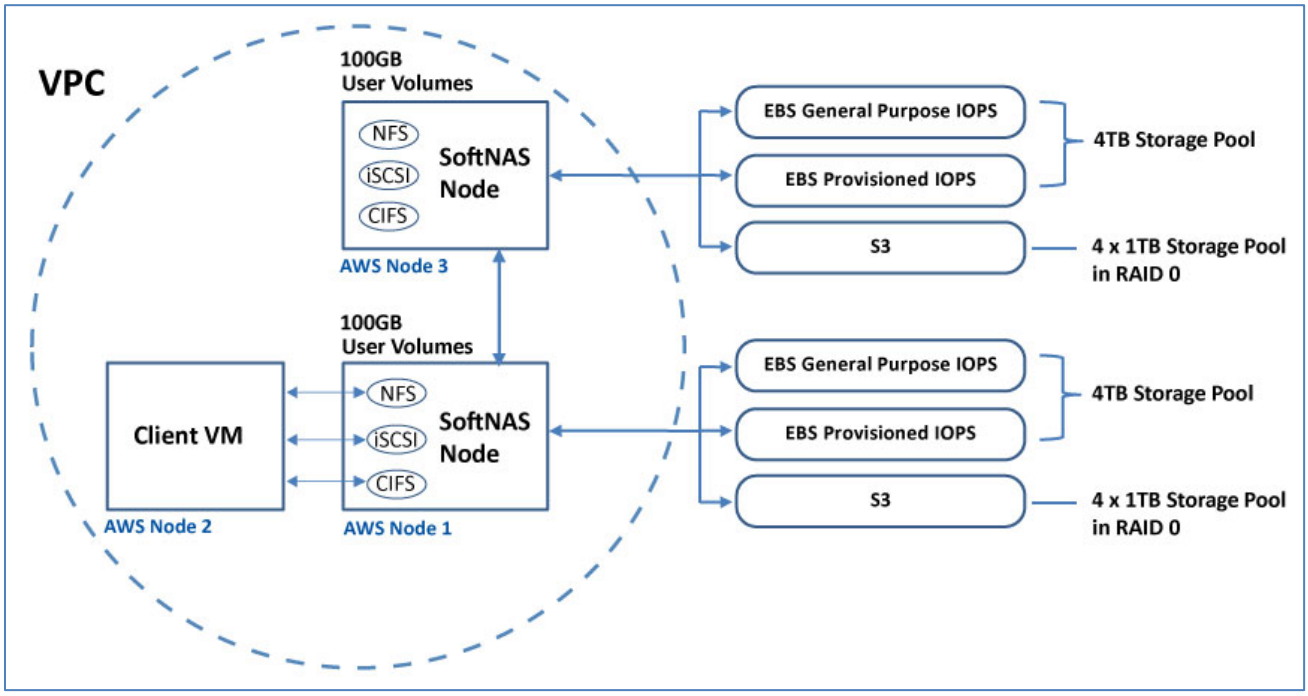

Fig 2: Overview of SoftNAS Cloud test setup on AWS

- Choice of AWS instance types, m3.xlarge, m3.2xlarge and c3.8xlarge, was based on their mix of CPU, RAM and SSD profiles. Specifications of the nodes are summarized in Table 1:

AWS Instance Type	Number of Cores	RAM (GB)	Local SSD (GB)	Max Network Bandwidth (Mbps)	Max Storage Throughput (MB/s)
m3.xlarge	4	15	2 x 40GB	500	62.5
m3.2xlarge	8	30	2 x 80GB	1,000	125
c3.8xlarge	32	60	2 x 320GB	4,000	500

Table 1: AWS Instance Details

- Separate 4 TB storage pools were created in SoftNAS Cloud for EBS General Purpose IOPS, Provisioned IOPS and S3. For a 4 TB volume size, EBS General Purpose IOPS provides peak IOPS of 10,000 while EBS Provisioned IOPS provides peak IOPS of 20,000. However, actual realization of these rated IOPS depends on the network and storage throughput supported by the participating AWS nodes. For example:
  - Storage Throughput = IOPS x Block Size
  - The storage throughput limit should be the greater than the network bandwidth allowed by both the SoftNAS Cloud and Client instances.

- Volumes of 100 GB size were carved out in SoftNAS Cloud for each of iSCSI, NFS and CIFS protocols. In the case of iSCSI, the volumes were not formatted with a file system and raw disks were used for the tests. Dedupe and Replication were enabled for respective tests as needed.

- The AWS instance size of the Client VM was the same size as that of SoftNAS Cloud ensuring that both nodes supported similar network bandwidths and therefore any bottlenecks in performance were not caused by network limitations between the Client and the SoftNAS Cloud instance. **This is a very important consideration.**

- The Client VMs used CentOS for iSCSI and NFS tests, as well as, Windows for the CIFS tests – the volumes were not formatted with a files system and raw volumes were directly used.

- AWS uses a default MTU setting of 9000, which was left unchanged.

- FIO was used to generate load for the tests.
  - 32 KB block size was used for all EBS tests.
  - 1 MB block size was used for S3 tests.
  - 10G FIO file size was used EBS tests.
  - 2G FIO file size was used for S3 tests.
  - IO Queue depth = 128
  - Random and Sequential tests with R/W Ratio of: 0%, 20%, 50%, 80% and 100% were carried out for each combination.
  - Dedupe tests were conducted with 40% dedupable blocks.
  - 3 tests were conducted for each combination and average results were reported.

Performance Summary and Sizing Recommendations

EBS Tests

- Peak Read and Write IOPS obtained for SoftNAS Cloud with EBS Storage and FIO workload of 32KB block size are summarized in Tables 2 and 3:

	Read IOPS	Write IOPS
iSCSI	12,500	5,000
NFS	19,000	5,600
CIFS	5,800	4,400

Table 2: EBS General Purpose IOPS

	Read IOPS	Write IOPS
iSCSI	12,500	10,500
NFS	19,500	12,300
CIFS	5,800	7,300

Table 3: EBS Provisioned IOPS

- On a stand-alone basis, EBS Provisioned IOPS is expected to outperform EBS General Purpose IOPS. However, when these EBS options were used with SoftNAS Cloud, EBS General Purpose IOPS performed on par with the EBS Provisioned IOPS for read operations. This improvement in performance was due to caching of data by SoftNAS Cloud for reads. **The implication for SoftNAS Cloud users is that for read intensive workloads, the users can consider using EBS General Purpose IOPS with SoftNAS Cloud instead of using the expensive EBS Provisioned IOPS.**

- The choice of AWS Instance Size has a big impact on SoftNAS Cloud performance. Surprisingly, the impact is due more to the network and storage throughput constraints of the AWS instances, than to their CPU, RAM or SSD characteristics. Based on the observations from the tests, the recommended AWS instance choices for 32 KB block sizes are summarized in Table 4.

	IOPS	Storage Throughput (MB/s)	AWS Recommended Node
1	< 2,000	<62.5 MB/s	m3.xlarge
2	2,000 – 4,000	62.5 - 125 MB/s	m3.2xlarge
3	4,000 – 20,000	125 - 500 MB/s	c3.8xlarge

Table 4: Recommended AWS Node Sizes for EBS

Note: The recommendations are based on tests conducted with 32 KB block sizes in FIO. Higher IOPS will be observed on smaller nodes when smaller block sizes are used. **Users need to evaluate the throughput limits to refine the choice of nodes.**

S3 Tests

- The S3 tests were conducted with a FIO block size of 1 MB. Since larger block sizes were used, the IOPS obtained for these tests were lower as the performance was determined by AWS' network and storage quotas. However, the storage throughput was consistently high across all the tests due to the large block sizes. Peak Read and Write throughput obtained for S3 are summarized in Table 5.

	Read Throughput(MB/s)	Write Throughput (MB/s)
NFS	600	600
CIFS	350	370

Table 5: S3 Peak IOPS across all tests

- Based on observations from S3 tests, the recommended AWS instance choices for 1 MB block sizes are summarized in Table 6.

	IOPS	Storage Throughput (MB/s)	AWS Recommended Node
1	< 62.5	<62.5 MB/s	m3.xlarge
2	62.5 - 125	62.5 - 125 MB/s	m3.2xlarge
3	125 - 500	125 - 500 MB/s	c3.8xlarge

Table 6: Recommended AWS Node Sizes for S3

Note: The CPU Utilization for some S3 tests on m3.xlarge instances ran higher than recommended levels when the storage throughout reached closer to 60MB/s.
Production workloads should use m3.2xlarge or larger.

Other Observations

- **Right Sizing the Client VMs**: As important as it is to choose the right AWS instance for SoftNAS Cloud, it is equally important to choose the right instance for the client nodes as well. Since client nodes interact with SoftNAS Cloud over the network, the peak storage throughput that client can achieve is limited by its network bandwidth. **Ideally both SoftNAS Cloud and the client nodes should be of the same size.**

- **Impact of Caching**: SoftNAS Cloud was found to significantly enhance the performance of reads due to its caching feature – especially for sequential reads. Thus, SoftNAS Cloud is a very good fit for workloads such as OLAP, Hadoop, video streaming, log aggregation, etc. – that are read intensive and sequential in nature.

- **Impact of Dedupe, Compression and High-Availability (HA)**: When features such as dedupe and compression are used in SoftNAS Cloud, performance of write operations was found to drop by 25-30%. This behavior is due to SoftNAS Cloud having to dedicate part of the CPU and RAM capacity, for inline processing for these features. **As a rule of thumb, users should use an instance size one notch higher than base when these features are enabled.** Furthermore, if replication is also switched on, in addition to dedupe and compression, users should use larger instance sizes such as c3.8xlarge and above for high performance workloads

Key Results

- **EBS General Purpose IOPS:** Figures 4 and 5 outline the results for 100% sequential read and 100% sequential write performance for EBS General Purpose IOPS, for different instance sizes and volume types. The dotted lines indicate the expected IOPS for a given instance based on its network and storage bandwidth limits available from the AWS provided infrastructure. For 32 KB block size, the theoretical peak IOPS for m3.xlarge is 1,952, for m3.2xlarge it is 3,906 and for c3.8xlarge it is 15,625. The theoretical peak IOPS is calculated as (allowed bandwidth in MBps x 1000) / (Block size in KB).

  It is evident from the following results that IOPS on small node types are throttled due to the network bandwidth limits and true unconstrained IOPS can be seen for the c3.8xlarge node.

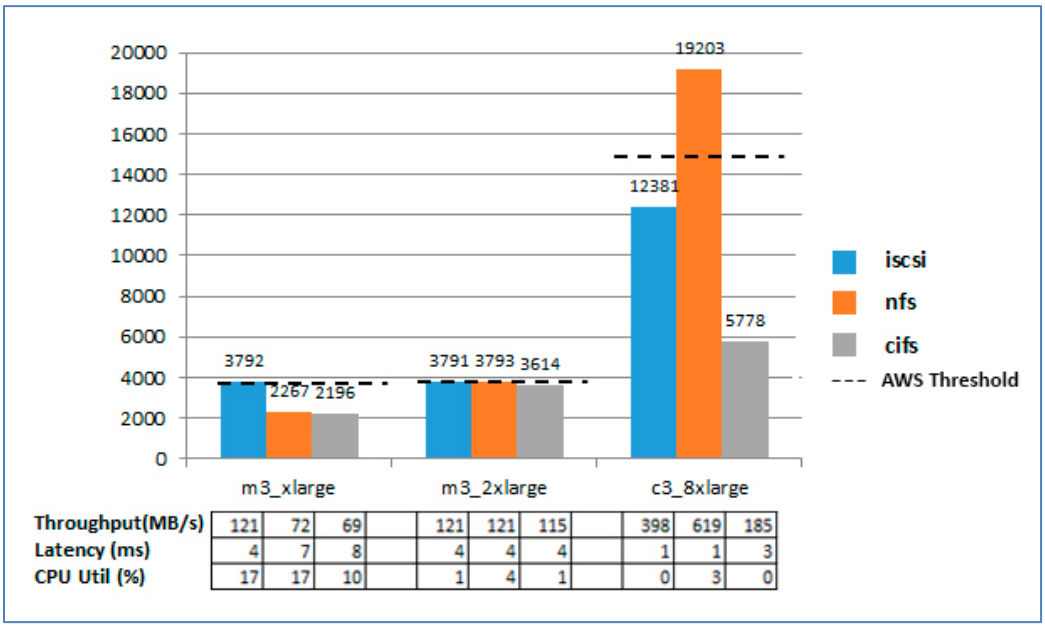

Fig 4: Sequential 100% Read IOPS for EBS General Purpose IOPS

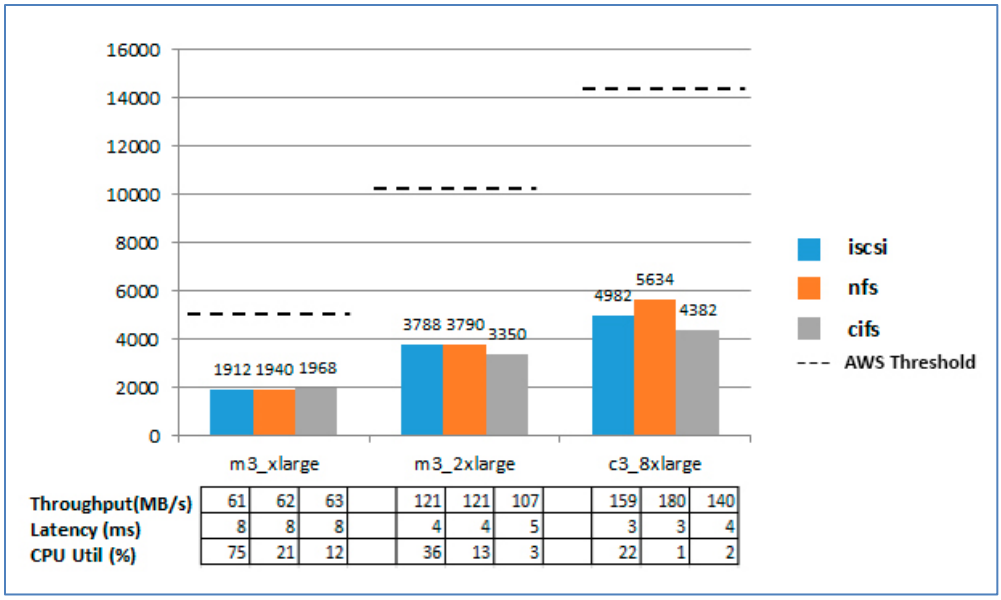

	m3_xlarge			m3_2xlarge			c3_8xlarge		
Throughput(MB/s)	61	62	63	121	121	107	159	180	140
Latency (ms)	8	8	8	4	4	5	3	3	4
CPU Util (%)	75	21	12	36	13	3	22	1	2

Fig 5: Sequential 100% Write IOPS for EBS General Purpose IOPS

- **EBS Provisioned IOPS:** The performance of EBS Provisioned IOPS is throttled on smaller instances due to network and storage throughput limits for those AWS instances. Even on the larger instance sizes, EBS Provisioned IOPS provides similar read IOPS as the General Purpose IOPS, since SoftNAS Cloud boosts the performance of the General Purpose IOPS through caching, to bring it on par with the Provisioned IOPS.

  The key value-add of Provisioned IOPS is for writes where it outperforms the General Purpose IOPs (by approximately 2x) on larger instances. Therefore, their use is justified only where high write IOPS are needed.

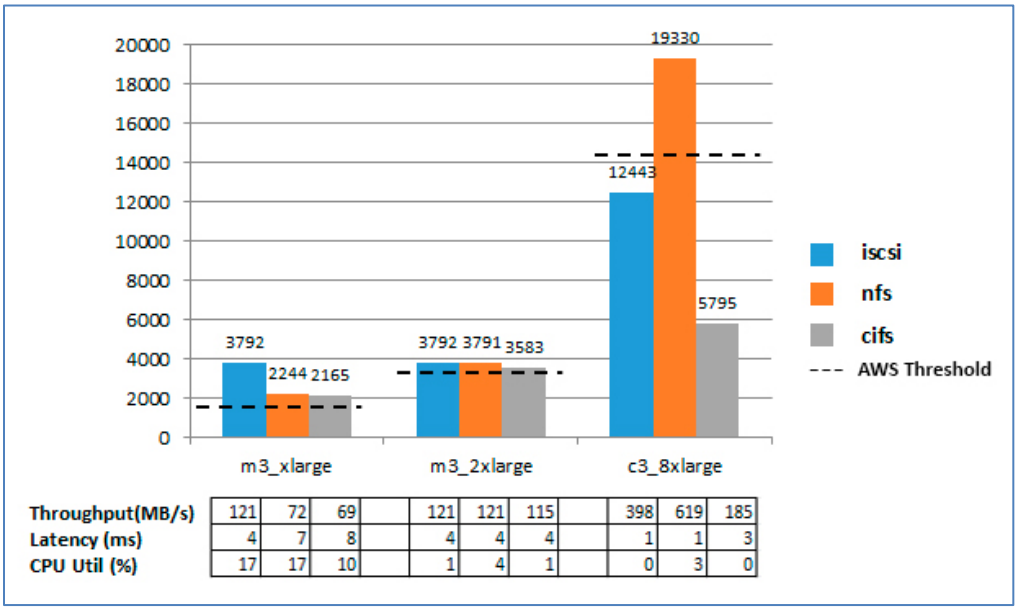

	m3_xlarge			m3_2xlarge			c3_8xlarge		
Throughput(MB/s)	121	72	69	121	121	115	398	619	185
Latency (ms)	4	7	8	4	4	4	1	1	3
CPU Util (%)	17	17	10	1	4	1	0	3	0

Fig 6: Sequential 100% Read IOPS for EBS Provisioned IOPS

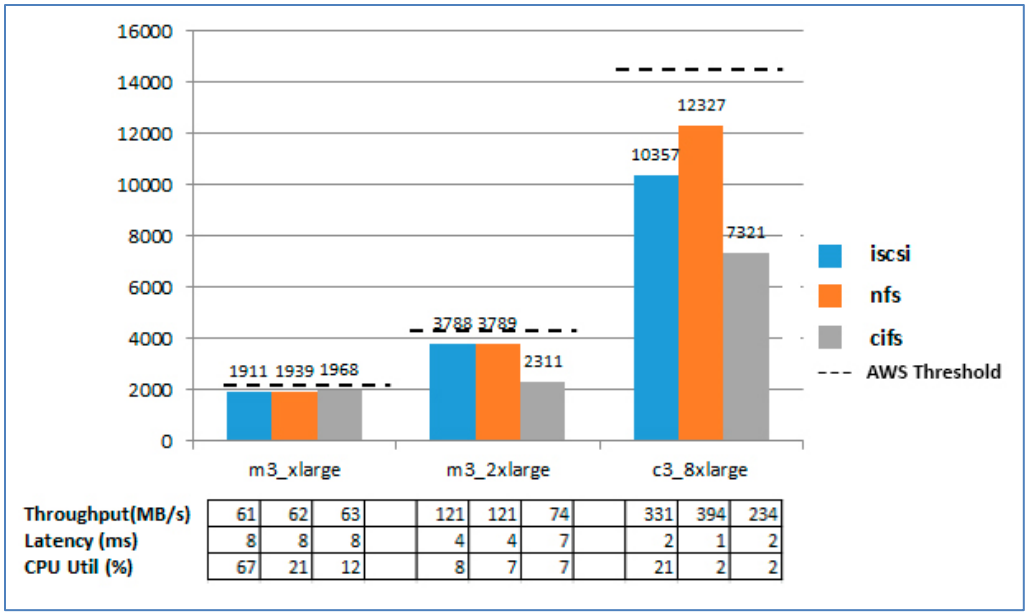

Throughput(MB/s)	61	62	63		121	121	74		331	394	234
Latency (ms)	8	8	8		4	4	7		2	1	2
CPU Util (%)	67	21	12		8	7	7		21	2	2

Fig 7: Sequential 100% Write IOPS for EBS Provisioned IOPS

- **S3 IOPS:** The performance of S3 storage is limited by number of puts and gets that are allowed by AWS for the instance types. The maximum storage throughput is achieved for larger block sizes and therefore S3 Tests used a higher block size of 1 MB. The IOPS obtained for these tests, were lower as the performance was determined by AWS network and storage quotas. However, the storage throughput achieved was high across all the tests again due to the larger block size used.

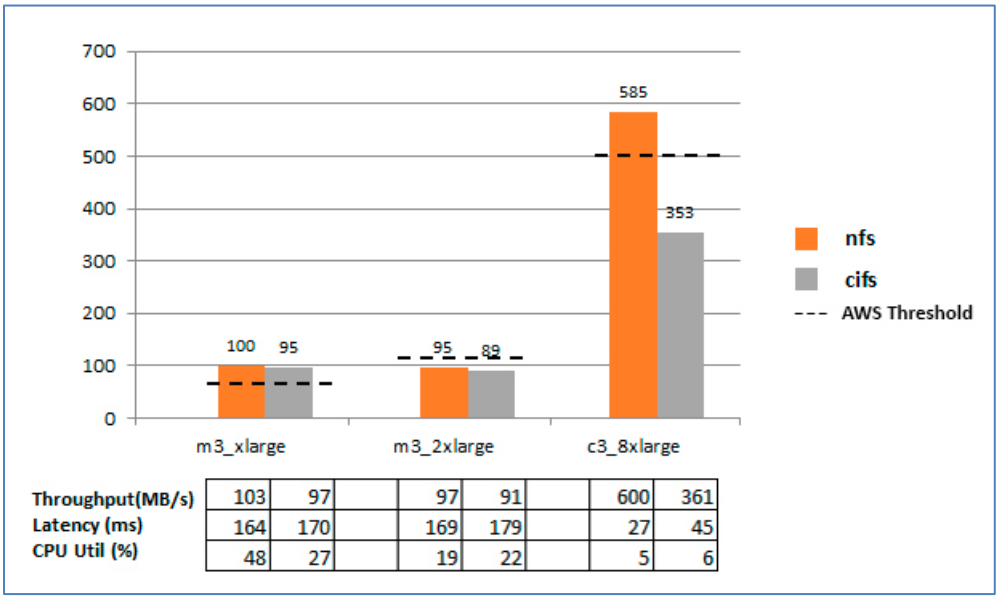

Throughput(MB/s)	103	97		97	91		600	361
Latency (ms)	164	170		169	179		27	45
CPU Util (%)	48	27		19	22		5	6

Fig 8: Sequential 100% Read IOPS on S3

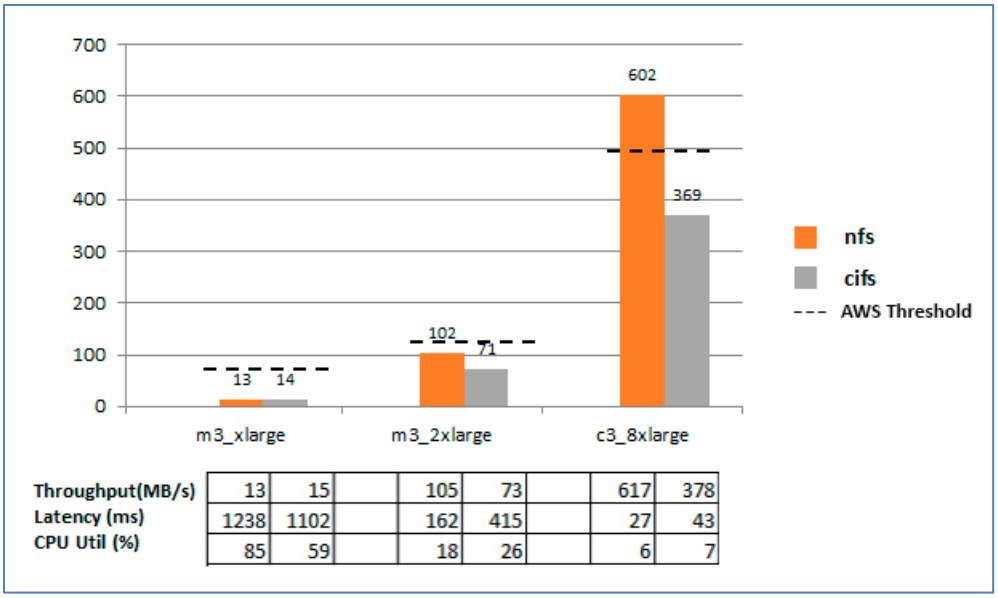

	m3_xlarge			m3_2xlarge			c3_8xlarge	
Throughput(MB/s)	13	15		105	73		617	378
Latency (ms)	1238	1102		162	415		27	43
CPU Util (%)	85	59		18	26		6	7

Fig 9: Sequential 100% Write IOPS on S3

- **CPU Utilization Analysis**: To help the users choose the right instances where multiple features such as: HA, Dedupe and Compression were switched on, sequential write tests were conducted with EBS General Purpose IOPS and NFS volume type and the CPU Utilization was measured during the tests. As observed in Figure 10, the CPU Utilization peaks for m3.xlarge, is moderately high for m3.2xlarge and is well within limits for c3.8xlarge.

  Therefore, m3.xlarge is not recommended for use when multiple features are used. While m3.2xlarge is moderately loaded on average for a block size of 32KB, it will likely get saturated for 8 KB block size. It is therefore recommended for users to use c3.8xlarge in these scenarios.

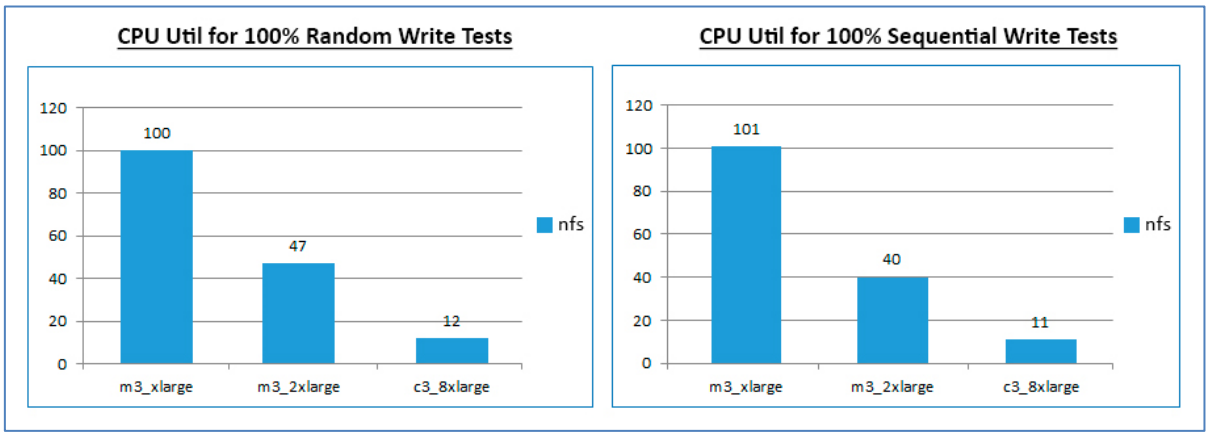

Fig 10: CPU Utilization on SoftNAS node with HA, Dedupe and Compression switched with NFS on EBS General Purpose IOPS

Conclusion

SoftNAS Cloud provides a robust set of features that make using AWS significantly easier, more secure and allows the customer, in some cases, to use less expensive AWS options like EBS General Purpose IOPS instead of using the more expensive EBS Provisioned IOPS.

While testing can't cover every real-world scenario that a customer might have, we believe that the test results show a fair and balanced representation of some of the most common customer workloads. SoftNAS, highly recommends that a customer engage and work closely with one of our Solutions Architects (SAs) in order to properly design and configure their workloads on AWS using SoftNAS Cloud.

As mentioned in the introduction, AWS infrastructure improvements continuously occur. As faster options come to market, the software-based SoftNAS Cloud solutions flexibility allows our customers to use the best of breed ingredients as they become available. These test results reflect offerings that are currently available from SoftNAS and AWS that will change over time.