

Cloud Storage

Storage has become a critical part of all datacenters today – be it a small datacenters, private clouds or larger public clouds. The tremendous growth in storage has resulted in a new trend towards high- performance, scalable cost-effective storage solutions using 10 GbE networks.

Before we look at the various storage technologies available to you, it is important to understand the attributes that affect your cloud storage solution.

- **Latency:** In many cases, compute resources are wasted because they are blocked on a read/write operation. Storage latency has a direct impact on application performance. It is important to minimize read/write latency from your servers to your storage system in order to fully utilize the compute bandwidth, as well as to improve overall performance.
- **Throughput:** The size of data you can Read or Write per second is one of the most important attributes of any storage solution.
- **Cost:** The storage solution needs to scale over time as your needs grow. While it is easy to do a price-performance analysis today, planning for future technology changes is always harder. Staying with a mainstream technology such as NAS over 10GbE will allow you to scale out and keep your costs down.
- **Scale:** It is easy to design for the above attributes with just local storage attached to the servers. However, as the scale increases that does not work. Cloud Storage solutions need to be able to address the above attributes even at very large scales – with Petabytes of storage.

Designing Cloud Storage Solutions

For quite some time, you could tradeoff between the above attributes and design your storage with either local storage, NFS over 1GbE, or Fiber Channel for higher performance. 4G FC had the best performance but was also expensive.

Now, the amount of data you need to store has grown so fast that the older technologies don't work or you don't get the benefit of economies of scale that one would expect with larger storage clusters. Cloud Storage solutions such as Amazon's S3 are clearly taking advantage of these trends to take storage to the next level.

Keeping the Cloud Storage attributes in mind, let's look at how various storage technologies stack up, with a focus on the interconnect.

Latency

The Read/Write latency of any solution is the sum of the storage system latency and the latency of the interconnect fabric. While 4Gbps FC had better performance, now with iSCSI and NFS over 10GbE you can get far better performance. This can be seen in Figure 1 below.

The latency of the interconnect fabric starts to go up significantly when using legacy Ethernet switches, as each hop adds to the latency of every packet, and in some cases hundreds of packets traverse the network to complete a single read or write operation. Hence it is important to use ultra- low latency switches that can forward packets in sub-microseconds in a larger network.

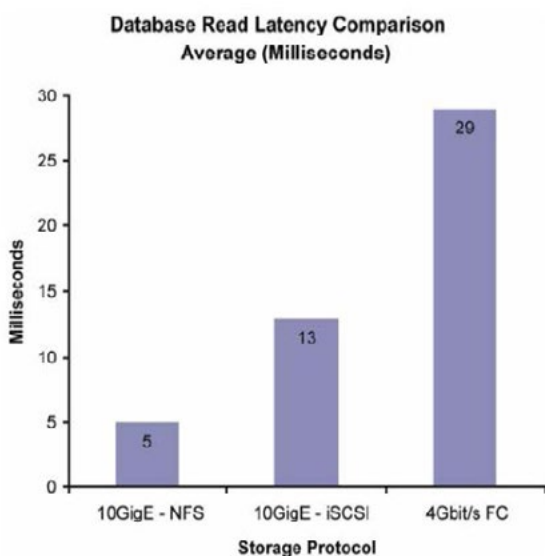


Figure 1: Read Latency Comparison - Source: NetApp

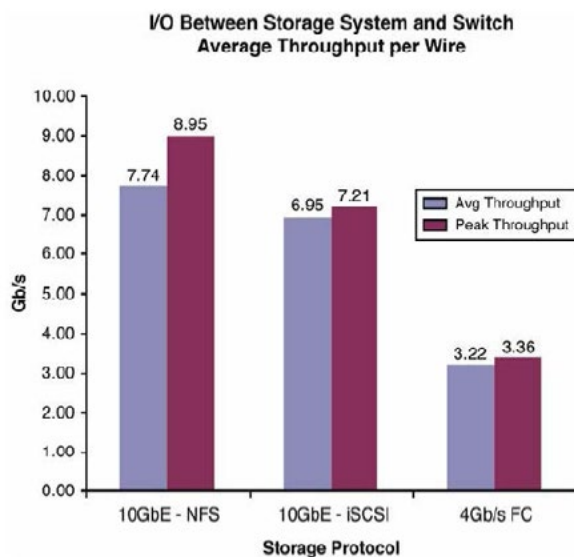


Figure 2: Storage Throughput Comparison - Source: NetApp

Throughput

Servers and Storage devices have much higher performance now and you can easily saturate a 10GbE link with your storage traffic. With NFS and iSCSI over 10GbE, you can now take advantage of the full 10GbE bandwidth.

As can be seen in Figure 2, the Read/Write throughput of iSCSI over 10GbE and NFS over 10GbE far outperforms traditional FC solutions.

Cost & Scale

These attributes go hand-in-hand when designing a Cloud Storage solution. The solution needs to be cost effective today and also when you scale the storage cluster. The costs come from the storage system (including the file system & management), host adapters for all your servers and the interconnect.

NFS or iSCSI over 10GbE can be run using standard host adapters and existing 10GbE Ethernet networks. The larger datacenters have created their own file systems, or adopted an open source file system such as Luster or ZFS in order to scale. You also get improved efficiency in your data center with higher density 10GbE switches as an entire rack can now be wired up with a 1RU 48 port 10GbE switch, offering non-blocking throughput.

FCoE is relatively new and very expensive once you add up all the hardware and software license costs involved. While FCoE is needed to provide interconnects to legacy Fiber Channel SANs, wide- scale deployment has not been seen. The trend is clearly towards NAS (NFS or iSCSI) over 10GbE.

Summary

Various storage technologies have emerged over time. NFS or iSCSI over 10GbE are the most cost effective and have the best performance for building cloud storage solutions today. 10GbE is now the primary interconnect technology for building robust, cost-effective and scalable storage solutions of all sizes.



Arista's 7000 Family of high density and low latency data center switches are well suited to improve your storage latency even with multiple hops, and provide a robust platform to build a single network fabric that connects your compute, cache, database and storage nodes. With up to 48 10GbE ports in a 1RU switch or 384 ports in an 11RU modular switch, you can get non-blocking connectivity to all your nodes. Arista's products deliver industry leading price-performance, and help you keep your costs to a minimum.

Arista EOS provides self-healing resiliency and unmatched robustness for a highly available interconnect fabric. EOS open APIs allow you to integrate the management of your Storage and Network devices on to a single management application of your choice.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document.