Ethernet powers the new Cloud Storage Network

What's Hot | June 2017

http://www.dcseurope.info

SNIA ESF | ETHERNET STORAGE

When storage networks first became popular starting around 1997, the term generally meant block storage on a Fibre Channel (FC) network. The reason was twofold: first, Fibre Channel was seen as the only dedicated storage networking technology and second, anyone building a storage network generally wanted it for high-performance block storage.

By Tim Lustig, SNIA Ethernet Storage Forum Member, Mellanox.

The Traditional Storage Network was Fibre Channel

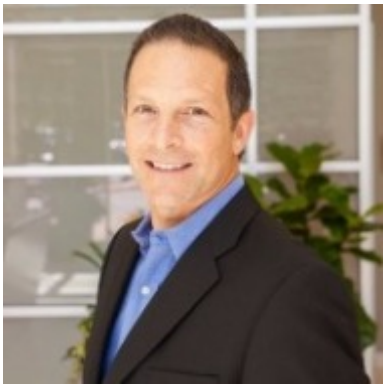

The following factors helped make Fibre Channel the choice of storage networks:

- FC was faster: 1Gb FC when Ethernet was at 100Mb, then 4/8Gb FC while Ethernet ran at 1Gb/s.
- Fibre Channel was a lossless network while Ethernet was lossy.
- Most networked storage was block (fortunate, since FC essentially only does block storage).
- There was no cloud, no Internet of Things, no Facebook, and no GoPro cameras to produce terabytes of daily digital content, so there was less need for fast file and object storage.
- Most storage networks were relatively small.
- It was assumed that the only way to make a storage network reliable and secure was to have it physically separated from the regular network.

FC Storage Area Networks (SANs) largely serviced critical and performance-sensitive applications like Oracle, Siebel, PeopleSoft, and SAP. Since storage was relatively small and the data associated with the applications using them was valuable, customers were not as sensitive to the cost associated with implementing a separate network just for storage.

Enter the Cloud

When the cloud arrived, and grew, and grew, and then grew some more, as did the use of social media it started to blur the boarders between professional and amateur digital content growth. The first all-digital major motion picture (Star Wars Episode II: Attack of the Clones) was released in 2002. Facebook and GoPro launched in 2004, followed by YouTube in 2005. Twitter and Amazon's S3 storage service both launched in 2006. Keeping up with the Kardashians (and the Apple iPhone — coincidence?) debuted in 2007. Instagram started in 2010 and then Snapchat in 2011, all of which added to the digital content explosion. At this point, a shift in block-based to file-based data began and due largely to trends associated with social media we saw a huge swing in content growth of unstructured data.

The world of content and storage was increasingly moving towards a digital and cloud-oriented world. Clouds and cloud storage have several important characteristics:

- Scale: Whether it's 100TB in a private cloud to 20,000PB (20 Exabytes) in a public cloud, the infrastructure will be big and ever growing. Cloud storage networks must be able to scale easily to accommodate rapid growth.
- Cost/Efficiency: Cloud storage is usually distributed, requiring many more nodes and network ports than typical enterprise storage. As a result, the networking infrastructure must be inexpensive and efficient; which is the reason we see public clouds consistently lowering their prices in an effort to remain competitive.
- Focus on file and object: Block storage will always be important, but the amount of file and object content is growing far faster. A 30 percent annual growth in database entries now comes with a 100 percent increase in associated videos, photos, and unstructured files.
- Converged network: When looking for efficiencies, Cloud architects like to limit network technologies use to just one for everything, and since Ethernet is ubiquitous the trend is to use it.

Not surprisingly, all the large cloud providers choose Ethernet to standardize on because it's relatively inexpensive, well understood and quick to deploy. Even in many large enterprise (non-cloud) deployments, customers are saying their data is growing so much faster than their budgets and Ethernet offers them the greatest flexibility.

Ethernet Innovation

Fortunately for the growing number of cloud customers of the world, Ethernet technology was moving forward quickly as well. In 2005 (circa YouTube), Ethernet caught up with FC in throughput, as vendors started shipping 10Gb Ethernet and 8Gb FC around the same time. In 2011 (circa Snapchat), Ethernet delivered 40Gb/s Ethernet. Today, Ethernet has been shipping 100Gb adapters and switches for more than two years and is on a well-defined path to 200Gb/s.

Beyond mere bandwidth, lossless Ethernet debuted in 2010 (circa Instagram), with Data Center Bridging (DCB). The original purpose was to support Fibre Channel over Ethernet (FCoE) but has since expanded to include other protocols such as iSCSI, SMB Direct, and NVMe over Fabrics. An additional development was Remote Direct Memory Access (RDMA), which was added (as RDMA over Converged Ethernet – RoCE) as an Ethernet specification in 2010 and first shipped in 2011. This gave Ethernet an extra level of performance with an enormous increase in efficiency and is the reason behind RoCE's success in several production cloud deployments today.

Ethernet has also added QoS and virtualization solutions to prioritize storage traffic and create logically separate storage networks when needed. It includes strong encryption and authentication options for security, along with robust tools for traffic management and monitoring.

The Ethernet Ecosystem

Ethernet has a huge ecosystem — with at least 6 significant high-speed NIC vendors and dozens of switch vendors. The resulting competition drives innovation at a rapid pace while pushing down per-port prices. This is exactly the technology curve cloud providers need to support their growth and constant price cuts. The new speeds like 25, 50, and 100GbE offer more bandwidth per port, often reducing the number of ports needed. And Ethernet networks can be used for block, file and object storage, as well as compute and management traffic. This enables network bandwidth to be cost-effectively scaled to support server and storage solutions that reside within cloud and web-scale data center environments.

Everything Ethernet in the Cloud, Including the Storage

Nearly all large cloud service providers have implemented Ethernet networking for storage. Some set up separate logical or virtual networks to isolate storage traffic while others simply provision enough Ethernet bandwidth to handle all types of traffic. Typically, they use iSCSI for block, NFS or SMB for file, and an AWS S3 or Swift-type API for object storage. In some cases, distributed storage solutions such as Ceph or HDFS (Hadoop Distributed File System) are also used. Some clouds are still using a mix of 10GbE and 40GbE while others have already migrated to 25/50/100GbE technology.

By now it should be clear that there are many options and variables for building storage networks for the cloud and there is one common theme, it is almost always built on Ethernet.

For additional perspective on how the cloud changes storage, see the following SNIA resources on Hyperscaler Storage at: www.snia.org/hyperscaler. The Ethernet Storage Forum (ESF) is committed to providing vendor-neutral education on the advantages and adoption of Ethernet storage networking technologies. Learn more at www.snia.org/esf.