



Two hot trends for 2013

Members of SNIA's Ethernet Storage Forum offer their views on the past, present and future, with a focus on consolidation, object storage and 10GBASE-T.

A 2012 review and what to expect in 2013.



By Jason Blossil

As we come to a close of the year 2012, I want to share some of our successes and briefly highlight some new changes for 2013. Calendar year 2012 has been eventful and the SNIA-ESF has been busy. Here are some of our accomplishments:

10GbE: With virtualization and network convergence, as well as the general availability of LOM and 10GBASE-T cabling, we saw this is a "breakout year" for 10GbE. In July, we published a comprehensive white paper titled "10GbE Comes of Age." We then followed up with a Webcast "10GbE – Key Trends, Predictions and Drivers." We ran this live once in the U.S. and once in the U.K. and combined, the Webcast has been viewed by over 400 people!

NFS: has also been a hot topic. In June we published a white paper "An Overview of NFSv4" highlighting the many improved features NFSv4 has over NFSv3. A Webcast to help users upgrade, "NFSv4 – Plan for a Smooth Migration," has also been well received with over 150 viewers to date. A 4-part Webcast series on NFS is now planned. We kicked the series off last month with "Reasons to Start Working with NFSv4 Now" and will continue on this topic during the early part of 2013. Our next NFS Webcast will be "Advances in NFS – NFSv4.1 and pNFS." You can register for that here.

Flash: The availability of solid state devices based on NAND flash is changing the performance efficiencies of storage. Our September Webcast "Flash – Plan for the Disruption" discusses how Flash is driving the need for 10GbE and has already been viewed by more than 150 people.

We have also added to expand membership and welcome new membership from Tonian and LSI to the ESF. We expect with this new charter to see an increase in membership participation as we drive incremental value and establish ourselves as a leadership voice for Ethernet Storage. As we move into 2013, we expect two hot trends to continue – the broader use of file protocols in datacenter applications, and the continued push toward datacenter consolidation with the use of Ethernet as a storage network. In order to better address these two trends, we have modified our charter for 2013. Our NFS SIG will be



renamed the File Protocol SIG and will focus on promoting not only NFS, but also SMB / CIFS solutions and protocols. The iSCSI SIG will be renamed to the Storage over Ethernet SIG and will focus on promoting data center convergence topics with Ethernet networks, including the use of block and file protocols, such as NFS, SMB, FCoE, and iSCSI, over the same wire. This modified charter will allow us to have a richer conversation around storage trends relevant to your IT environment. So, here is to a successful 2012, and excitement for the coming year.

Object Storage is a big deal (and Ethernet matters)



By Ingo Fuchs

A significant challenge in managing large amounts of data (or Big Data) is a lack of what I like to call "total data awareness". It's a situation where you know (or suspect) that you have data – you just can't find it. When you think about many current IT environments, they are often not built for total data awareness. This starts with core

elements of the IT infrastructure, such as file systems. Traditional file systems and access methods were not designed to store hundreds of millions or billions of files in a single namespace. This leads to admins storing data in multiple file systems, multiple shares, complex directory structures – not because the data should be logically organized in that way, but simply because of limitations in file system architectures. This issue becomes even more pressing when data sits in multiple locations, maybe even across on-premise and off-premise, cloud-based storage.

Is object-based storage the answer?

Think about how you find data on your computer. Do you navigate complex directory structures, trying to remember the file name of the file that hopefully has the data you are looking for – or have you moved on and just use search tools like Spotlight? Imagine you have

hundreds of millions of files, scattered across dozens or hundreds of sites. How about just searching across these sites and immediately finding the data you are looking for? With object storage technology you have the ability to store data in objects, along with metadata that describes the object. Now you can just search for your data based on metadata tags (like a filename – or even better an account number and document type) – as well as manage data based on policies that leverage that metadata. However, this often means that you have to consider interfacing with your storage system through APIs, as opposed to NFS and CIFS – so your applications need to support whatever API your storage vendor offers.

CDMI to the rescue?

Today, storage vendors often use proprietary APIs. This means that application vendors would have to support a plethora of APIs from a number of different vendors, leading to a lack of commitment from application vendors to support more innovative, object-based storage architectures. A key path to solve this issue is to leverage technology and standards that have been specifically developed to provide this idea of a single namespace for billions of data sets and across locations and even managed services that might reside off-premise.

Relatively new on the standards side you have CDMI (<http://www.snia.org/cdmi>), the Cloud Data Management Interface. CDMI is a standard developed by SNIA (<http://www.snia.org>), the Storage Networking Industry Association, with heavy involvement from a number of leading storage vendors. CDMI not only introduces a standard interface to ingest and retrieve data into and out of a large-scale repository, it also enables applications to easily manage this repository and where the data sits.

CDMI is the new NFS

Forgive the provocation, but when it comes to creating and managing large, distributed content repositories it quickly becomes clear that NFS and CIFS are not ideally suited for this use case. This is where CDMI shines, especially with an object-based storage architecture behind it that was built to support multi-petabyte environments with billions of data sets across hundreds of sites and accommodates retention policies that can reach to "forever".

CDMI and NFS have ethernet in common

One of the key commonalities between CDMI and NFS is that they both are ideally suited to be deployed in an Ethernet infrastructure. CDMI, specifically, is a RESTful HTTP interface, so it runs on standard Ethernet networks. Even for object storage deployments that don't support CDMI, practically all of these multi-site, long-term repositories support HTTP (and thus Ethernet) through proprietary APIs based on REST or SOAP.

Why does this matter

Ethernet infrastructure is a great foundation to run any number of workloads, including access to data that sits in large, multi-site content repositories that are based on object storage technologies. So if you are looking at object storage, chances are that you will be able to leverage existing Ethernet infrastructure.

How is 10GBASE-T being adopted and deployed?



By David Fair

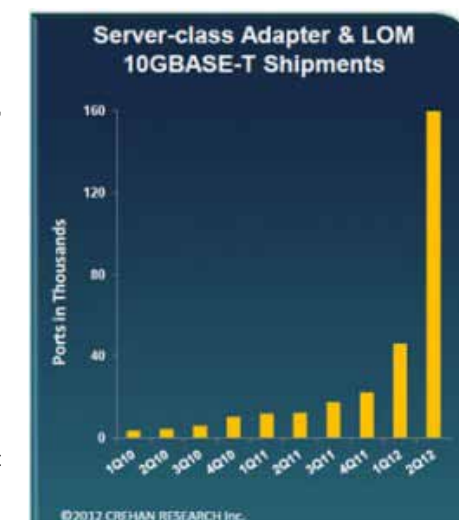
For nearly a decade, the primary deployment of 10 Gigabit Ethernet (10GbE) has been using network interface cards (NICs) supporting enhanced Small Form-Factor Pluggable (SFP+) transceivers. The predominant transceivers for 10GbE are Direct Attach (DA) copper, short range optical (10GBASE-SR), and long-range optical (10GBASE-LR). The Direct Attach copper option is the least expensive of the three. However, its adoption has been hampered by two key limitations:

- DA's range is limited to 7m, and
- because of the SFP+ connector, it is not backward-compatible with existing 1GbE infrastructure using RJ-45 connectors and twisted-pair cabling.

10GBASE-T addresses both of these limitations

10GBASE-T delivers 10GbE over Category 6, 6A, or 7 cabling terminated with RJ-45 jacks. It is backward-compatible with 1GbE and even 100 Megabit Ethernet. Cat 6A and 7 cables will support up to 100m. The advantages for deployment in an existing data center are obvious. Most existing data centers have already installed twisted pair cabling at Cat 6 rating or better. 10GBASE-T can be added incrementally to these data centers, either in new servers or via NIC upgrades "without forklifts." New 10GBASE-T ports will operate with all the existing Ethernet infrastructure in place. As switches get upgraded to 10GBASE-T at whatever pace, the only impact will be dramatically improved network bandwidth.

Market adoption of 10GBASE-T accelerated sharply with the first single-chip 10GBASE-T controllers to hit production. This integration become possible because of Moore's Law advances in semiconductor technology, which also enabled the rise of dense commercial switches supporting 10GBASE-T. Integrating PHY and MAC on a single piece of silicon significantly reduced power consumption. This lower power consumption made fan-less 10GBASE-T NICs possible for the first time. Also, switches supporting 10GBASE-T are now available from Cisco, Dell, Arista, Extreme Networks, and others with more to come. You can see the early market impact single-chip 10GBASE-T had by mid-year 2012 in this analysis of shipments in numbers of server ports from Crehan Research.



Server-class adapter & LOM 10GBASE-T shipments

Note, Crehan believes that by 2015, over 40% of all 10GbE adapters and controllers sold that year will be 10GBASE-T.

Early concerns about the reliability and robustness of 10GBASE-T technology have all been addressed in the most recent silicon designs. 10GBASE-T meets all the bit-error rate (BER) requirements of all the Ethernet and storage over Ethernet specifications. As I addressed in an earlier SNIA-ESF blog, the storage networking market is a particularly conservative one. But there appear to be no technical reasons why 10GBASE-T cannot support NFS, iSCSI, and even FCoE.

Today, Cisco is in production with a switch, the Nexus 5596T, and a fabric extender, the 2232TM-E that support "FCoE-ready" 10GBASE-T. It's coming – with all the cost of deployment benefits of 10GBASE-T.