Unplanned Obsolescence: EOL for FC?

Present in some forms since as early as the 1980s, Fibre Channel (FC) was standardized in the mid 1990s. While it took a few years to adopt, not unlike any other storage protocol, it has enjoyed an exceptionally long period of use given elevated costs and a relatively slow rate of improvement. Nevertheless, FC addressed the real need for scalable storage area networks (SANs), and has since been widely deployed. Some of this resiliency to competing technologies may be explained by storage centric capabilities and reliability, offloaded hardware implementations and a dedicated ecosystem. A specialized technology, FC has kept largely ahead of Ethernet in terms of storage features and performance.

In fact, FC has been so comfortably entrenched in the SAN market, that when alternative storage networking protocols over Ethernet finally matched it in performance, it was caught unprepared and lacking a competitive response. And as Ethernet storage leaps over in both performance and features, the unplanned obsolescence of FC seems in sight. This crossover has coincided with several other market trends – datacenters, clouds and, in general, converged infrastructure – that are not aligned with FC, hence accelerating the migration away from FC for all new installations. FC's lack of flexibility and capability in enabling general purpose networking and latency sensitive compute applications has further relegated the protocol to a single purpose, and closed off any avenues to move it forward.

This paper starts by giving a brief overview of Fibre Channel's development, and discusses the reasons for its success so far. It then presents the current status of storage over Ethernet, and the reasons for the ongoing transition from FC to Ethernet.

FC Development

A serial protocol meant to replace the 50-pin HIPPI protocol in the HPC space, FC eventually found its niche in serializing parallel SCSI, completing the transition to pure storage when connected directly to disks.

When the first Gigabit Fibre Channel products became available in the second half of the 1990's, Gigabit Ethernet was in its infancy. Until then, Ethernet used a shared media access control model, perhaps good enough for networking but not for providing predictable storage performance. As Ethernet moved to fully switched operation, it built up the feature set – such as QoS and link level flow control – necessary for handling storage applications.

In the meantime, FC showed regular speed increases, doubling every 3 to 4 years. Hence, it developed a lucrative market where the installed base was refreshed on a regular basis. This very profitable vertical enabled high ASP's to the point that the profit dollars of Fibre Channel business were more than half of profit dollars of Ethernet despite the huge disparity in unit count.

FC Strengths

It is important to consider the key benefits which enabled FC to retain its customer base, besides an initial edge in performance:

1. Specialization: FC is dedicated to the purpose of carrying the SCSI protocol. It provides the required infrastructure for storage, and the isolation needed from other traffic, albeit at the expense of running a separate fabric. In addition, specialized management tools integrated with storage were developed to ease the deployment and management of SANs.
2. Hardware implementation: the FC protocol is typically fully offloaded to HBAs, minimizing the impact of storage I/O on the host CPU.
3. Long distance: the use of fiber allowed long distance operation, and remote backup and disaster recovery at a time where Ethernet was still a local area network technology, with limited reach.

However, the high acquisition and operation costs of FC motivated the development of technologies on the Ethernet side, mainly the Internet SCSI (iSCSI) protocol which allowed the transport of SCSI directly over TCP/IP. Recently, a proposal for encapsulating FC frames in Ethernet has been standardized. The following section discusses both alternatives.

Ethernet Alternatives

Work on the Internet SCSI protocol started in the late 1990s at IBM, and the standard RFC appeared in 2004. The motivation for iSCSI in offering block storage networking over the ubiquitous Ethernet technology was to simplify and commoditize SAN connectivity. Since then, iSCSI has gone through the usual development and maturity curve, and it now stands as robust, high performance alternative to FC with a fully featured management toolset.

iSCSI natively supports critical robustness and efficiency features such as end-to-end CRC protection including T10-DIF, and direct placement of storage data into host memory. An Ethernet native, it enjoys the well-known benefits of IP – unlimited routability over subnet boundaries – and TCP – stability, scalability, reliability, and robustness. As a result, iSCSI traffic implements the critical congestion and flow control mechanisms needed to scale to large networks as well as to run over legacy equipment. Additionally, it can coexist with all other Ethernet traffic and can be observed, monitored, debugged and managed identically to any other TCP/IP/Ethernet traffic, with common tools such as *Wireshark* and *tcpdump*. Chelsio's iSCSI hardware stack is a 4th generation high performance pipelined implementation, using a hardened protocol processing engine with over 400,000 ports shipped, and in large scale deployment for half a decade. Furthermore, Chelsio's solution integrates advanced traffic management and QoS functionality which exceed anything available in the FC environment. Hardware based traffic shapers and schedulers allow provisioning resources down to flow level, and ensuring different levels of service accordingly.

Fibre Channel over Ethernet is a recent entrant into the Ethernet SAN arena. FCoE replaces the lower layers of the FC stack with raw Ethernet encapsulation, and in place of the removed flow control and reliability mechanisms, requires "lossless" Ethernet operation. A first disadvantage with respect to iSCSI is that the parts of the FC protocol which remain require the use of specialized and expensive FC Forwarders (FCF switches). While FCoE is going through the pains of early development and deployment, serious issues undermine its long term prospects for success, such as its need for lossless network operation. Not only does this requirement translate into higher costs in DCB (Data Center Bridging) enabled equipment and switches with oversized packet buffers, questions about its suitability

for large scale deployment remain unanswered. In particular, FCoE's lack of network protection and scalability mechanisms raises serious concerns once a more realistic multi-hop environment is considered. In fact, even single hop benchmarks tend to expose its high sensitivity to the slightest network disturbances. The rest of the discussion therefore focuses on iSCSI, which enjoys stronger inherent robustness and is at a higher maturity level.

Performance

The following graph tracks the link speeds available for FC and Ethernet over time. Initially, FC enjoyed a significant advantage over Ethernet. Even when 10Gbps became available in 2003, Ethernet lagged in price competitiveness for a few years. The price advantage has switched since then, while FC fell behind Ethernet in performance. Today, FC is not only lower performance and more expensive overall, but it stands to face the onslaught of 40Gbps and 100Gbps Ethernet, which are expected to arrive in quick succession. On the other hand, the 32G FC standard remains out of sight and the most optimistic estimates put it well after 100Gbps Ethernet is set to be deployed. It is hard to see how FC will be able to compete on these two key metrics.

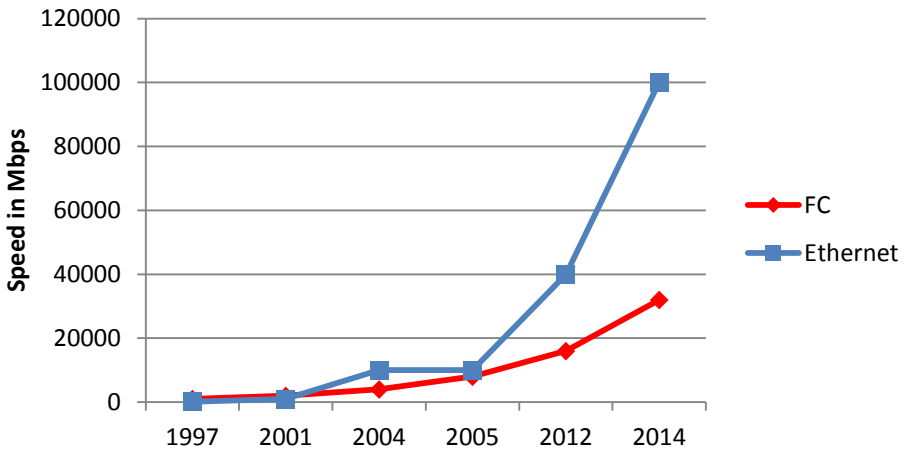

Figure 1 – Fibre Channel vs. Ethernet Link Speeds

The chart above shows that FC, today 16G, has a planned speed bump in 2016 or beyond. In fact, most FC market is currently shipping at 8G or lower however since 16G FC only became available in early 2012. Given Ethernet has already been commoditized at 10Gb, and will be commoditized at 40Gb by 2013, the transition to 16Gb Fibre Channel has been slow and the transition to 32Gb appears unlikely at best since it has already been obviated by 40Gb Ethernet. With 100Gbps Ethernet showing up in 2014, Ethernet is set to enjoy overwhelming superiority for many years.

In terms of I/O per second performance, and counter to initial projections, iSCSI over 10Gbps Ethernet has been shown to provide very high performance (1.1M IOPs compared to the 200K from 8Gb FC). Ethernet is further expected to double this performance in 2013, and then again in 2015. Achieving such high storage benchmarks on Ethernet requires a high speed hardware implementation, such as Chelsio's offload engine. Chelsio's Terminator ASIC is architected to process all the layers of the stack in parallel and in a pipelined fashion, resulting in very high packet processing capacity. The following figure projects the IOPS performance figures for iSCSI and FC, which along with the capacity discussed paint a dim view of FC's prospects for competing against iSCSI.

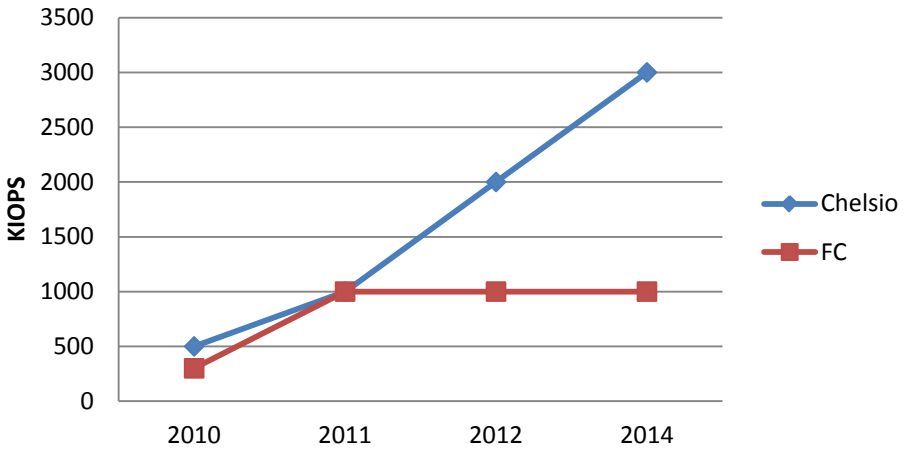

Figure 2 – Fibre Channel vs. Ethernet IOPS

Return on Investment

A dedicated FC SAN infrastructure is typically considered from its positive angle, i.e. the separation of storage traffic from others. However, it also is a liability, since at least an additional separate infrastructure is also needed for networking. With iSCSI and FCoE, it possible to use the same physical interconnect for converged traffic. This can be achieved by separating traffic on different ports, and in case of a QoS enabled adapter, it can even be done on the same port. Therefore, moving from FC to iSCSI results in significant cost savings, in addition to a performance upgrade.

Conclusion

Fibre Channel today stands at an inflection point, with all the evidence available pointing to a downward slope for the aging technology. In the decades since FC's inception, Ethernet based alternatives have been introduced and have matured, while rapidly increasing in performance. This is not a precedent, and will not be the last time Ethernet builds up the feature set and capabilities needed to take over the market of a specialized interconnect.

In the case of FC, the process required a scalable high performance offload architecture to provide high IOPS and bandwidth at high efficiency and reliability levels, a feat which seemed difficult with iSCSI, and perhaps led to a false sense of safety. In fact, the arguments against iSCSI scaling with the performance of Ethernet are as enduring as the ones detractors had against heavier than air flight: the right engine simply blows them away. Chelsio is now uniquely positioned to deliver the high performance iSCSI based solution, and to enable the massive reduction in TCO associated with convergence over Ethernet. Like the descendants of the mighty airships of a century ago, the future iterations of FC may well end up dotting the landscape as a technological curiosity.

Related Links

RFC3720 Internet Small Computer System Interface
CERN FC Overview
T11 FCoE Page
Dell Comparison of iSCSI, FC and FCoE