



Why the FCoE iSCSI debate continues

By Jeff Asher, SNIA ESF member, NetApp.

A Look Back at FCoE and iSCSI History

There are two entrenched standards for block storage protocols over Ethernet networks. FCoE was ratified in 2009, while iSCSI was ratified in 2004. Of course, various vendors and early adopters supported these protocols before ratification, so the history of these protocols is a couple of years longer than it looks, respectively. While iSCSI simply encapsulates the SCSI protocol in IP, FCoE operates lower in the network stack and to do so required many enhancements to Ethernet. While iSCSI runs on any IP network (mostly Ethernet these days), FCoE requires Data Center Bridging and Converged Network Adapters all running at 10 Gbps or faster.

All of the data center bridging enhancements that make FCoE possible, like lossless Ethernet, benefit all of the protocols using Ethernet as the transport protocol. DCB doesn't just make FCoE possible, but it improves iSCSI at the same time (see the SNIA-ESF blog, [How DCB Makes iSCSI Better](#)). So given that modern servers, networks, and storage may all be connected by hardware capable of running FCoE, that same network is also able to run iSCSI, as well as other network traffic.

Nothing precludes them from running simultaneously on the same network either. The leading storage vendors that offer both FCoE and iSCSI target systems allow administrators to present the same LUN over either protocol with little effort, so a transition from one protocol to the other is not difficult.



Strengths and Weaknesses Which network protocol is the right one?

Each protocol has strengths and weaknesses when judged relative to each other. FCoE has higher throughput at lower host CPU utilization than iSCSI and FCoE doesn't have to process the TCP/IP stack as iSCSI does. iSCSI is relatively simple to setup and troubleshoot when compared to FCoE because zoning is not a factor and IP connectivity (although not optimized for storage traffic) is likely in place already. Also, while FCoE has a comprehensive set of existing tools available to ease troubleshooting, there aren't as many qualified people to use them in most enterprises.

Ease of use, plus the ability to use low cost NICs and switches, gives iSCSI a cost advantage. (However, if you check out our SNIA-ESF webcast, "[How VN2VN Will Help Accelerate Adoption of FCoE](#)," you'll hear about new technologies that reduce

the costs of deploying FCoE.) FC, and by extension FCoE, are perceived to be enterprise-grade, suitable for all workloads; and while iSCSI is being widely adopted at the enterprise level, it is still perceived by some not to be ready for Tier-1 applications. The graph below is excerpted from the report "Intel 10GbE Adapter Performance Evaluation" prepared by Demartek for Intel in September 2010.

This data is consistent with the rest of the report findings and is only intended to be representative of the results from comparative iSCSI and FCoE testing. The report is interesting reading and I recommend looking at it for more information. This graph shows IOPS and CPU utilization for JetStress tests running against NetApp storage over multi-path

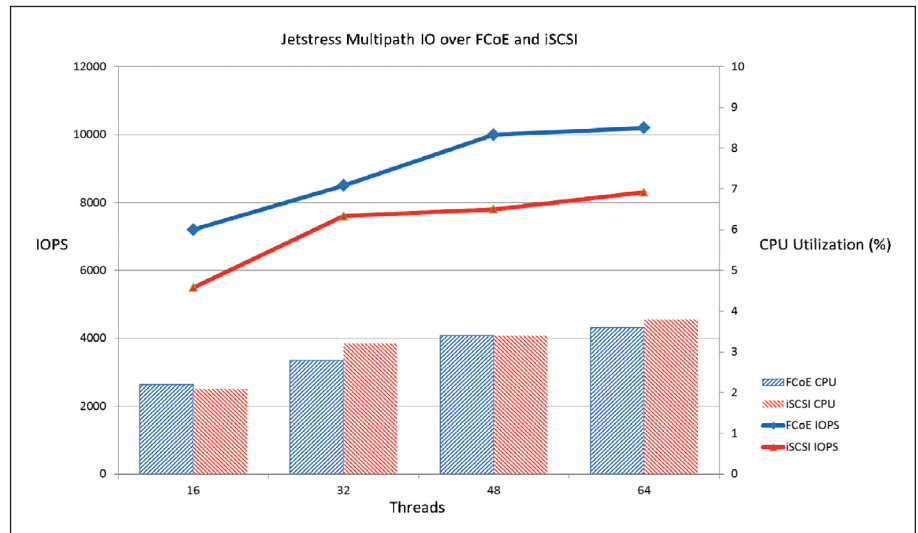
iSCSI and FCoE. Note that latencies were all similar and running the tests against EMC storage showed similar results. Many other factors must be considered, but according to industry pundits- as well as my own personal experience - in the majority of cases either protocol is adequate for the task at hand, and that is to effectively transfer block data across an Ethernet network.

Maximizing throughput

The reality is, most servers, applications, and storage arrays simply won't take advantage of FCoE's superior performance or any storage protocol running over 10GbE. iSCSI and NAS protocols are very fast and are typically sufficient to meet most application requirements. But this is not meant to be a SAN vs NAS post - besides years of history, thousands of happy end users, and billions of continued investment show that both work well enough to meet most business needs. The commonly deployed storage systems and hosts are simply not configured with enough hardware to saturate multiple 10 gigabit network links. While this is rare today, it is going to become more common to see systems capable of saturating 10GbE pipes in the near future, especially as flash memory, either in all-flash arrays or tiered storage systems, find more application.

(Hear more on the impact of flash in our SNIA-ESF webcast, "Flash - Plan for the Disruption"). At least as it relates to spinning media disk systems - network bandwidth increases faster than storage system throughput can keep up. So consider the storage system to be the bottleneck or limiting factor when evaluating storage network performance. After all, in most data center environments, the ratio of servers and applications to storage systems is high.

So, it's reasonable to expect the storage system to be the bottleneck. The absolute throughput of FCoE and iSCSI, when pushing a storage system to its limits, is not sufficient alone to be used as the sole basis for the decision between the two protocols except, for a few edge cases. Bottom line: Whether the storage system is the



bottleneck or the network is the bottleneck the performance relationship between FCoE and iSCSI does not change.

These edge cases tend to be extremely IO intensive database workloads and big data applications, such as Hadoop. Citing the graph above, FCoE is about 15-20% faster on identical hardware than iSCSI. Granted this is a single graph of a single test, but the data is consistent across tests performed by IBM using Emulex network interfaces.

If absolute throughput and efficiency (both network and CPU) are the only criteria when deciding between block protocols, FCoE looks like the choice. Since these cases are rare - because complexity, supportability, and even politics are almost always considered - the decision is not so obvious. Again, beyond the scope of this article, NAS protocols should be considered when determining the proper protocol for an application also.

Is there a clear winner?

While FCoE can claim technical superiority, iSCSI has the edge in cost and supportability. The number and range of systems supporting iSCSI connectivity is greater, particularly at the entry level. What's more, the availability of people that can

troubleshoot end-to-end connectivity for iSCSI is also much greater. (The "ping" command diagnoses most iSCSI connectivity problems.) Also, do a resume search on Monster or LinkedIn and the number of people that can configure VLANs dwarfs the number that can properly zone a Fibre Channel network.

Greater familiarity reduces the support and operating cost of iSCSI.

IDC predicts that FCoE revenue will ramp very quickly through 2016. (If available to you, see the IDC Worldwide Enterprise Storage Systems 2012-2016 Forecast Update.)

As customers decide to transition existing Fibre Channel networks to an Ethernet infrastructure, deploying FCoE would be a comfortable choice due to existing IT expertise and functional expectations of the Fibre Channel protocol.

Both iSCSI and FCoE are capable storage protocols and choosing one over the other will likely be dependent upon budget, IT skill set, and application requirements. For more information about SNIA's Ethernet Storage Forum please visit: <https://www.snia.org/forums/esf>

The commonly deployed storage systems and hosts are simply not configured with enough hardware to saturate multiple 10 gigabit network links. While this is rare today, it is going to become more common to see systems capable of saturating 10GbE pipes in the near future, especially as flash memory, either in all-flash arrays or tiered storage systems, find more application