

FCIA Market Positioning and FAQ for FCoE

Industry drafted standard statement:

“Fibre Channel (FC) is the technology standard for datacenter SAN environments, and Fibre Channel over Ethernet (FCoE) is an example of continued innovation that will enhance investments in FC-based architectures and business solutions. FCoE expands Fibre Channel into the Ethernet environment, combining two leading technologies - FC and Ethernet - to provide more options to end users for SAN connectivity and networking. We expect FCoE to deliver leading performance without the latency and complexity issues of TCP, to provide greater flexibility in network infrastructure and protocol decisions through network convergence, and to enhance FC investments in the datacenter.”

FCoE:

- Provides a method to encapsulate Fibre Channel over the Ethernet transport
- Is similar to other Fibre Channel encapsulations defined in the T11 Back Bone project
- Transmits FC frames over native Ethernet rather than TCP to ensure good performance
- Takes advantage of proposed enhancements to Ethernet in support of today's demanding datacenter environments
- Expands FC into the Ethernet environment, providing more options for SAN connectivity and networking
- Allows servers using multi-protocol attached NICs to leverage & extend existing FC networks
- Allows network convergence for attachment of FC SANs, IP LANs, and RDMA clustering on a single network structure.
- Helps to ensure FC remains the technology of choice for SAN connectivity

Frequently Asked Questions (FAQ):

Q: What is FCoE?

A: FCoE stands for “Fibre Channel over Ethernet”. FCoE is a proposed standard within the INCITS T11 FC-BB-5 committee for transporting Fibre Channel traffic over Ethernet. FCoE operates at the same OSI layer as IP, making FC over Ethernet analogous to IP over Ethernet.

Q: Why another standard?

A: There are many forces at work giving FCoE serious momentum. FCoE offers a seamless protocol connection to existing FC SAN's, bringing the ability to mix FC and FCoE, or to migrate to FCoE from FC. The ability for Host systems to consolidate their external connections onto a single type of connection also brings tremendous value to businesses both large and small.

Q: What effect will FCoE have on iSCSI?

A: FCoE is being designed to make its implementation as light-weight as possible. This means it does not need the overhead of TCP, as does iSCSI. With iSCSI a gateway function is needed to terminate/originate the TCP/IP connections. This being said, FCoE works best with new features being worked on by the 802.3 committees to provide such capabilities as lossless frame transmission and multi-pathing. This means iSCSI will still be useful for quite some time to allow transmission of block data over WAN infrastructures. iSCSI will also be useful wherever there is a mix of existing and new Ethernet capabilities.

Q: When should I expect to see products using FCoE?

A: Best estimates to date are late 2008 to early 2010.

Q: Where is FCoE in the standards process?

A: The project has been approved and meetings have been scheduled to agree upon details of the standard. There is wide agreement on most fundamental aspects of the technology. The wide technical interest assures that the standard is likely to meet its scheduled completion date of June 2008.

Q: How long will it take to get through the standards process?

A: The standard is scheduled to be forwarded in June of 2008.

Q: Should I wait to implement my SAN based on FCoE?

A: Absolutely not! The FC SAN you buy today will be completely compatible with the products based on the new FCoE standard. Even with FCoE, the architecture of your SAN in the future will include a high percentage of existing FC equipment.

Q: How does FCoE affect the existing FC roadmap?

A: FC will continue to evolve along the existing roadmap speeds. As an example, 8GFC, with products deploying in 2008, is still a vital part of the FC story.

Q: How does FCoE affect my existing FC SANs?

A: One of the motivations behind FCoE is the transparent transport of the FC protocol throughout the infrastructure. Because of this, all existing FC deployments will be able to include or exclude FCoE, making the implementation of an FCoE solution seamless.

Q: How does FCoE relate to all the other IP type of protocols like FC-BaseT, FCIP, IPFC, iSCSI?

A: As shown in figure 1, FCoE replaces the TCP/IP stack needed by FCIP, iFCP and iSCSI. FC-BaseT is at a layer even lower than shown in Figure 1, as FC-BaseT is a physical connection mechanism and not a protocol.

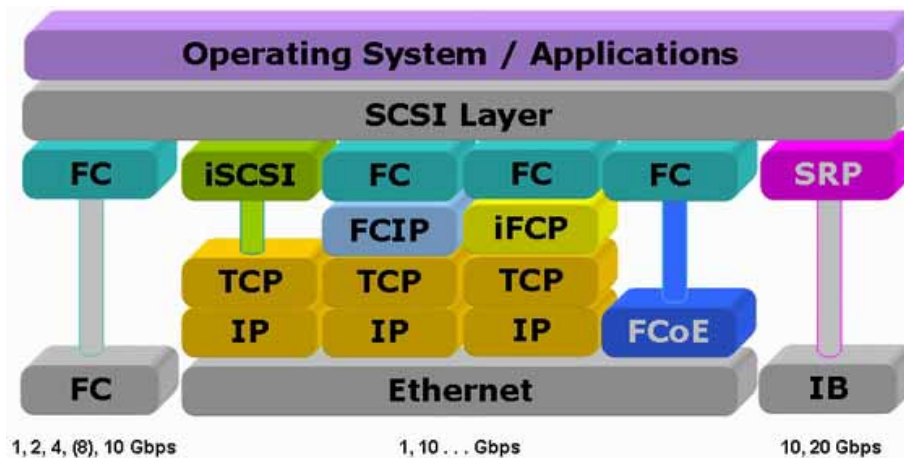


Figure 1- FCoE Relation to ISO Layers

Q: How does the standard committee plan to ensure high performance over FCOE?

A: The standard's entire purpose is to ensure there are efficient implementations that will perform as well or better than standard Fibre Channel and do this while leveraging Ethernet enhancements already underway in 802.3.

Q: Does FCOE require any changes and enhancements to Ethernet? What are they?

A: Yes. IEEE 802.3 has defined a new feature called PAUSE that will be required in Ethernet gear, as well as potentially using other Ethernet enhancements that 802.3 is currently working on. To what extent these other enhancements will be required for FCoE is currently under discussion by T11.

Q: Are the I/O rate and bandwidth equivalent to the numbers shown in the FCIA roadmap?

A: 1GFC is 1.0625 Gb/s, while 1Gbe is 1.25 Gb/s. Both use 8b10b encoding. 10GFC and 10Gbe both operate at the same line rate of 10.52 Gb/s and use 64b/66b encoding. Aside from line-speed

differences or similarities, FC and Ethernet each have their own efficiencies. Generally, FC is more efficient since it was designed for the transport of block data.

Q: Does Infiniband provide more than FCoE is planning to provide?

A: Yes. IB is predominately shipping SDR (Single Data Rate) at 10Gbps (4 X 2.5 Gbps) in 2007 with DDR (Double Data Rate) 20G (4 x 5 Gbps) starting to ship in the spring of 2007. It is expected that most IB shipped in 2008 will be predominately DDR capable. The IB roadmap specifies QDR (Quad Data Rate) at 40Gbps (4 X 10 Gbps) in three to 5 years. However, these solutions are virtually all server-to-server connections and not storage related. There is no SCSI protocol mapping available (SCSI command sets are THE language of enterprise storage - FC for FCP, iSCSI, SAS, etc.), nor is it in the works, for IB-SCSI, thereby relegating it to the server.