

FCoE Plug-Fests Ease the Concerns of Deployment Interoperability

Overview

Fibre Channel over Ethernet (FCoE) is considered the promising next-generation storage protocol that is based on Ethernet interface and maps Fibre Channel (FC) natively over Ethernet. It provides seamless connection to FC storage area networks (SAN) and enables converged networks in data centers. This is a significant step of input/output (I/O) consolidation in data centers based on Ethernet technology. Ever since the standard body first discussed FCoE, its implementation has been ubiquitous.

Multi-vendor interoperability plug-fest is the most effective way of proving new technology and bringing the ecosystems together. Since the very first discussions surrounding FCoE technology, the Fibre Channel Industry Association (FCIA) and the Ethernet Alliance (EA), both of which are non-profit international marketing organizations that focus on different technologies, have hosted many plug-fest events to expedite the adoption of new technologies. All the key players in the FC and Ethernet industries, such as ATTO, Brocade, Broadcom, BNT, Cisco, Dell, EMC, Emulex, Fulcrum Microsystems, HP, IBM, Intel, JDSU, Marvell, Mellanox, NetApp, and QLogic, work diligently in collaboration to ensure interoperability for the new FCoE and Data Center Bridging (DCB) standards.

Interoperability Challenges in FCoE Networks

Each plug-fest provides a forum for testing the interactions between multi-vendor storage arrays and converged network adapters (CNA) in a converged FCoE environment including DCB exchange (DCBx) switches. The key goal here being to prove and demonstrate the ecosystem interoperability by imitating various network architectures, both simple and complicated, in the real world.

Bringing together the interoperability among FCoE ecosystems is much more challenging than in conventional Ethernet networks, which are direct plug-n-play, because of the following realities in storage networking:

- The storage network is a fully controlled environment for security reasons. Storage protocols, such as FC, concentrate on strict link initializations and management, total link level flow control, and exchange steps; all of which complicate the link communications and make interoperability hard to converge. This well-known fact for FC continues for FCoE implementation.
- Proprietary implementation further complicates interoperability. The plug-in issues in the current FC networks will continue in FCoE network deployment.
- FCoE is an evolving technology. Early adopters implemented various protocols that could differ significantly from one another. The manufacturer's frequent firmware upgrades constantly break down the interoperability of FCoE ecosystems and complicate the deployment of FCoE networks.



Verification of Interoperability of FCoE Networks

JDSU plays a key role in FCoE plug-fests with its flagship product Xgig, well-known as the most trusted brand of protocol analyzer in storage industry, provides complete visibility into protocol communications with insightful hints for debugging issues. Xgig was used in every test case during FCoE plug-fest to smooth out interoperability issues among ecosystem partners. Xgig proved itself again the de facto tool for verifying interoperability in FCoE networks.

Plug-fest is also a good place to preview the issues in the FCoE pre-deployment. Helping ecosystem partners troubleshoot during interoperability tests gained JDSU valuable insight into what to expect when deploying FCoE networks in data centers. Issues such as the ones listed below have been identified as those most commonly seen during FCoE ecosystem installation:

- errors in virtual local area network (VLAN) settings
- discrepancies with versions of FCoE protocols
- FCoE Initialization Protocol FIP automatic VLAN assignment errors
- illegal exchange sequences
- failure to maintain link losslessness
- incorrect percentage of bandwidth utilization based on priority classes

Furthermore, these are link level negotiation issues that only hardware-based systems such as Xgig can detect, capture, and debug.

To facilitate a relatively smooth installation of a unified FCoE network, pre-deployment architects have an available installation guide [1], derived from the knowledge collected at the plug-fests. This joint effort among Cisco, Qlogic, and JDSU illustrates the potential network architecture scenarios, discusses interoperability issues, and defines verification processes.

Sample Test Case [2]

Figure 1 shows a very simple but typical FCoE network architecture. One CNA host is connected to FC storage through an FCoE switch. The JDSU Xgig Analyzer is placed in-line to monitor the traffic. This scenario can be used to verify FCoE initialization processes using FIP protocol, FCoE mapping, and Priority Flow Control (PFC).

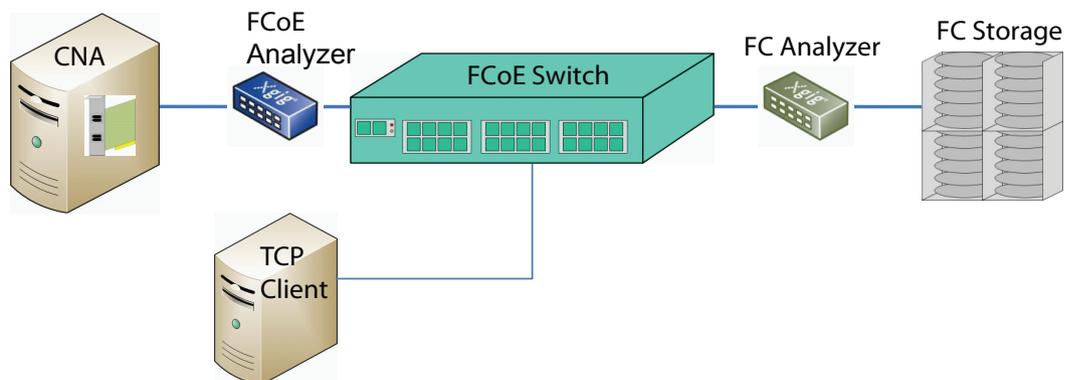


Figure 1. The simplest FCoE network architecture with a JDSU Xgig Analyzer installed in-line

For example, to detect and verify PFC behavior, the JDSU Xgig protocol analyzer was configured to capture and trigger on a PFC Pause frame. Figure 2 shows the first captured PFC frame is pausing priority 3 for 3355 μ s, a second PFC frame is sent to extend the pausing state, and a third PFC frame with a zero pause time is sent to release it. The traffic from the initiator stopped at the first pause and resumed after the pause release.

mm:ss.ms_us_ns (R)	Delta Time	Port	Protoc	Summary
00:00.000_478_005		Initiator	FC	FC4SData; SCSI FCP; Offset = 0x00062000; Len = 0x0800;
00:00.000_479_765	1.760	Initiator	FC	FC4SData; SCSI FCP; Offset = 0x0007C000; Len = 0x0800;
00:00.000_481_526	1.760	Initiator	FC	FC4SData; SCSI FCP; Offset = 0x00077800; Len = 0x0800;
00:00.000_483_286	1.760	Initiator	FC	FC4SData; SCSI FCP; Offset = 0x00077000; Len = 0x0800;
00:00.000_483_927	0.642	Switch	MPCP	PFC; Pause(3) = 3355 us;
00:00.000_520_125	36.197	Switch	MPCP	PFC; Pause(3) = 3355 us;
00:00.000_535_403	15.278	Switch	MPCP	PFC;
00:00.000_536_248	0.845	Initiator	FC	FC4SData; SCSI FCP; Offset = 0x00011800; Len = 0x0800;
00:00.000_538_008	1.760	Initiator	FC	FC4SData; SCSI FCP; Offset = 0x00062800; Len = 0x0800;

The initiator pauses the transmission when receiving the pause frame, and resumes when the pause is released

Figure 2. An initiator stops transmitting when receiving a Pause frame.

Xgig Expert software was used to detect all pause issues automatically by reading through the captured trace. As Figure 3 shows, depending on the severity of the issue, Expert ranked the debug hits with different categories and numbers. This sample proved that each device was actually pausing when requested. However, different devices took more or less time to stop transmitting frames after receiving the pause. Some devices exceeded the time allowed in the specification.

	Description	Timestamp	Source	Destination	Ports	Type	Total ...	Value 1
3	Bad TCP Retransmission (already ACK'ed)	000:02:41.035_244_;	10.1.1.100 ; 0000	10.1.1.60 ; 0000	1,3,1 / 1,3,2	TCP	539553	Seq Num 0
3	Frame Received while PFC Class Paused	000:02:41.039_600_ (VLAN 101 ; 0e:fc:00:d5:00:02	VLAN 101 ; 00:d5:00:D9:AE:8	1,1,1 / 1,1,2	GigE	124	Time since Pause 3.661 (us)

Figure 3. Xgig Expert reports devices that do not react quickly enough to the Pause requests.

Other Deployment Verification Concerns

Except for interoperability, which is a rudimentary step in deploying FCoE networks, other critical aspects must be tested to ensure a robust installation, especially as most FCoE networks are combined with virtualization. The JDSU Xgig platform based on the same hardware with different functionality licenses can emulate hundreds of virtual ports and thousands of virtual clients/servers. The benchmarking traffic generation provides the essential tools for verifying performance, scalability, and robustness in installed FCoE networks.

Conclusion

Ethernet is everywhere. Facilitated with FCoE technology, future data centers will be simpler, cheaper, and greener. However, storage managers are conservative and slow to accept new technology. Smooth transitions and backward compatibility are key to their acceptance. FCoE technology guarantees seamless migration by design. Industry-wide plug-fests help bring up the ecosystem successfully and ease the concerns of deploying next-generation networks.

References

1. Installation Guide—Storage Networking (Unified Fabric Pilot), by Cisco, Qlogic, and JDSU, 2010; available upon request
2. Ethernet Alliance White Paper: Data Center Bridging Plug-fest, May, 2010; http://www.ethernetalliance.org/files/document_files/DCB%20Whitepaper_v21.pdf.

Test & Measurement Regional Sales

NORTH AMERICA TEL: 1 866 228 3762 FAX: +1 301 353 9216	LATIN AMERICA TEL: +1 954 688 5660 FAX: +1 954 345 4668	ASIA PACIFIC TEL: +852 2892 0990 FAX: +852 2892 0770	EMEA TEL: +49 7121 86 2222 FAX: +49 7121 86 1222	WEBSITE: www.jdsu.com/test
---	--	---	---	--