



ethernet alliance

2012 Ethernet Alliance TeraFabric Plugfest

Contributors:

Greg McSorley, Amphenol

Chauncey Schwartz, QLogic

Patrick Strick, NetApp



Introduction

Given the speed of technology development and the fiscal constraints that today's data center executives face, there is no doing one thing at a time on one technology. Data centers are taking small steps migrating to higher speeds and new technology. Counting on doing all of their work with a single technology, speed or product is not feasible. Today's data center team must successfully operate in a world in which everything is happening at once—and Ethernet enables converged networking for those that need multiple protocols on the same wire at the same time.

What today's data-center managers seek, consequently, is confidence in proven, comprehensive interoperability across Ethernet vendors, network layers and technology generations ... confidence that whatever Ethernet device that is implemented anywhere across or even beyond their infrastructure will work with all of the disparate elements along the network ... confidence that devices will not only function properly within an infrastructure but go the next step to actually help elevate operational performance and productivity. And the global Ethernet ecosystem is actively striving to make such assurance in end-to-end interoperability a practical, every-day reality for data-center managers.

The 2012 Ethernet Alliance TeraFabric Plugfest is emblematic of this commitment. A tremendous range of key stakeholders from across the Ethernet community came together for the first-of-its-kind event Oct. 22-26, 2012, at the University of New Hampshire InterOperability Laboratory. Whereas such industry plugfests typically are focused on one area of the network, the scope of the TeraFabric Plugfest was sweeping. Participants included Amphenol Corporation (NYSE: APH), Applied Micro (NASDAQ: AMCC), Broadcom Corporation (NASDAQ: BRCM), Cisco Systems (NASDAQ: CSCO), CommScope, Inc., Dell Networking and Dell EqualLogic (NASDAQ: DELL), Emulex Corporation (NYSE: ELX), Foxconn International Holdings Ltd. (HKG: 2038), Intel Corporation (NASDAQ: INTC), Ixia (NASDAQ: XXIA), JDS Uniphase Corp (NASDAQ: JDSU), Leviton Manufacturing Co., NetApp Inc. (NASDAQ: NTAP), Nexans S.A. (EN Paris: NEX), Panduit Corp., PLX Technology, Inc. (NASDAQ: PLXT), QLogic Corporation (NASDAQ: QLGC), Samtec, Inc., Siemon Company, TE Connectivity (NYSE: TEL) and Volex (LSE: VLX). Together, the TeraFabric Plugfest participants—a cross-section of Ethernet switch, storage, adapter, test equipment and cabling industry leaders—demonstrated the interoperability of the complete range of Ethernet and related technologies, including Data Center Bridging (DCB), 10GBASE-T and 40 Gigabit per second (Gb/s) Ethernet. Furthermore, with such a wide array of high-speed Ethernet devices and solutions accessible, the event culminated with the successful construction of an integrated network capable of supporting a 1 Terabit per second (Tb/s) data flow.



The takeaway for data-center managers and the Ethernet industry from the TeraFabric Plugfest: Not only is the next generation of converged Ethernet networks here—the technology is sound, truly interoperable and ready for real-world deployment and broad-scale adoption—the generation beyond that is already in sight. The global Ethernet ecosystem is already anticipating its next frontiers of development.

The Ethernet Industry's Most Comprehensive Interoperability Testing Event To Date

Open standards-based testing in Ethernet typically has been focused on individual testing of a specific technology domain, and this strategy has enabled Ethernet's solution providers to keep a technology moving forward through its stages of maturity. The stage we have entered now, however, demands a different approach.

In today's real-world data centers, all of the components must work together as part of a single, seamless environment. The possibilities for lowered capital and operational expenditures (CAPEX and OPEX) for scaling, unified visibility and management and reduced maintenance complexity have long made the notion of consolidating disparate operations on the same infrastructure appealing, and Ethernet's inherent flexibility and ease of use make it the ideal solution for cost-effectively executing convergence strategies. Ethernet's increasingly sophisticated capabilities and higher speeds—to 10Gb/s and now 40Gb/s in the last two years—have made data-center convergence a legitimate option without harming operational quality.

Interoperability is essential to this vision. Ethernet interoperability fosters flexible, extendible and scalable data centers that can be easily and cost-effectively molded and remolded to meet emerging business needs. Ethernet interoperability ensures data-center managers more choices, at reduced prices and lower risk. And, for these reasons, assurance of end-to-end Ethernet interoperability drives Ethernet market growth. Through this interoperability, data centers can maintain current investments and take small steps with newer technology as they need to without concerns about the new technology working with the existing technology.

The Ethernet Alliance's multi-vendor test events draw together the Ethernet ecosystem to achieve and prove interoperability. Ethernet developers and

manufacturers convene at these “plugfests” to efficiently improve early technology prototypes and verify interoperability with other products in an environment that is safe, secure and confidential, as participants are required to sign non-disclosure agreements. The leading innovative minds across multiple Ethernet disciplines come together to link together various vendors’ equipment and demonstrate that traffic can be sent and received as intended. Sometimes, necessary corrections can be implemented on the spot at the plugfest if issues arise, as all of the participants can put their heads together and brainstorm on diagnoses and fixes. If a problem cannot be overcome at the plugfest, the vendor at the very least returns to its own facilities with more insight into its cause and potential solutions. Months of debugging in isolation can sometimes be avoided with one week of interoperability testing at a plugfest. The plugfest is, effectively, a giant laboratory for the best minds in the Ethernet ecosystem.

The October 2012 Ethernet Alliance TeraFabric Plugfest at the University of New Hampshire InterOperability Laboratory presented an unprecedented opportunity for the Ethernet ecosystem. This plugfest united an unusually broad range of current and next-generation Ethernet components in a controlled test setting to illustrate the benefits that a converged network can bring to the data center. The plugfest was intended to mirror the challenging environment that is already found in most enterprise data centers, which require ultra-flexible and massively scalable solutions.

The TeraFabric Plugfest delivered the most comprehensive and cohesive Ethernet interoperability test experience to date. Testing was performed for 10 Gigabit and 40 Gigabit Ethernet connections. This also included testing where a 40 Gigabit Ethernet link was used as four 10 Gigabit Ethernet links via a “breakout” cable. Also, both converged, DCB/Fibre Channel over Ethernet (FCoE) connections and non-converged Ethernet connections were tested.



Figure 1: 2012 Ethernet Alliance TeraFabric Plugfest participants spanned a cross-section of Ethernet switch, storage, adapter, test equipment and cabling industry leaders.



The following three sections describe the results related to the major 2012 Ethernet Alliance TeraFabric Plugfest testing goals:

- demonstration of interoperability of physical cables,
- testing of IEEE 802.1™ DCB and
- construction of a high-density, terafabric network.

Interoperability of Physical Cables

Cabling is an integral part of the Ethernet ecosystem, and one of the valuable opportunities that plugfests present is for cabling vendors to test their various assemblies with other leading Ethernet technology providers from around the world. In turn, data-center managers glean a greater sense of confidence in the range of choices they have in media for interconnecting devices in their infrastructures.

At the TeraFabric Plugfest, cabling vendors participated in validating the signal integrity performance of their direct-attach passive copper cable assemblies in accordance with the cable-assembly parameters specified for 40GBASE-CR4. The assemblies utilized the Style-1 40GBASE-CR4 MDI connectors, which is the QSFP connector. As noted, testing was also completed using a 40Gb/s Ethernet breakout cable in the QSFP form factor to four 10Gb/s Ethernet copper cables. Testing included:

- insertion loss,
- insertion loss deviation,
- differential input return loss,
- differential output return loss,
- mode conversion and
- integrated crosstalk noise.



Figure 2: Thorough signal integrity testing was performed at the 2012 Ethernet Alliance TeraFabric Plugfest on a variety of media, such as this 40Gb/4x10Gb breakout cable.

In these testing of passive cables, a network analyzer is applied at both ends of a cable; a signal is sent through the cable, and various parameters are measured at the points of compliance. The measurement of the signal is then compared against the applicable IEEE Std 802.3ba-2010™ specifications for 40GBASE-CR4.

Testing of active copper and active optical-cable assemblies enabled via the non-retimed XLPP (40 Gigabit Parallel Physical Interface) specified in IEEE Std 802.3ba-2010 were also conducted and included the following:

- eye mask compliance,
- differential input return loss,
- differential output return loss,
- differential to common mode input return loss and
- common mode output return loss.

Most notably, the TeraFabric Plugfest cabling tests demonstrated the robustness and interoperability among 40Gb/s Ethernet implementations, as well as between the different rates of Ethernet.

Testing Data Center Bridging

Data Center Bridging (DCB) refers to enhancements made to IEEE 802.1 for use in data center environments. Specifically, DCB is intended to eliminate loss due to queue overflow for selected traffic (lossless Ethernet) and to provide quality of service to ensure high-priority workloads get the bandwidth they require. The standards defined



to accomplish DCB are IEEE802.1Qbb™ Priority Flow Control, IEEE 802.1Qaz™ Enhanced Transmission Selection and Data Center Bridging Exchange Protocol and IEEE 802.1Qau™ Congestion Notification. For more details concerning DCB, visit the Ethernet Alliance web site at http://www.ethernetalliance.org/wp-content/uploads/2011/10/document_files_DCB_Whitepaper_v2.pdf.

Typically, data centers have housed a range of application-specific networks based on a variety of link-layer technologies—such as storage on Fibre Channel, high-performance computing (HPC) on InfiniBand and local area network (LAN) connectivity on Ethernet. DCB is intended to enable data center managers to build lossless Ethernet networks including support of Fiber Channel over Ethernet (FCoE), Internet Small Computer System Interface (iSCSI) DCB and RDMA over Converged Ethernet (RoCE) protocols. DCB converged Ethernet networks can be used to cost-effectively shift more and more of their mission-critical storage applications onto the same physical Ethernet infrastructure

The DCB tests at the TeraFabric Plugfest included Enhanced Transmission Selection (ETS) and Priority-based Flow Control (PFC) on multiple switches, as well as multiple classes of PFC. The tests sought to verify whether all devices had properly configured FCoE and iSCSI traffic via the correct priority, whether any priority group’s traffic utilized more than its configured amount of bandwidth and whether devices properly halted and resumed transmission of appropriate traffic classes in response to a “PFC PAUSE” operation. In the tests utilizing multiple switches, five minutes of both “READ” and “WRITE” operations had to successfully transpire between the devices, and the devices had to show up in each other’s management system.

The TeraFabric Plugfest’s testing around DCB delivered a series of key industry breakthroughs:

- demonstration of end-to-end 10GBASE-T in DCB architectures spanning the 10GBASE-T maximum length of 100 meters;
- the industry’s first multi-vendor public test of the Data Center Bridging Exchange (DCBX) protocol in IEEE 802.1Qaz™ “Standard for Local and metropolitan area networks—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks—Amendment 18: Enhanced Transmission Selection for Bandwidth Sharing Between Traffic Classes,” and
- the industry’s first testing of 40Gb/s Ethernet in a multi-vendor DCB environment, including direct-attached passive cables, and enabled by the XLAUI interface - active copper and optical cables. Initial 40Gb/s Ethernet testing included passive cables up to 7 meters, active cables up to 7 meters and AOCs up to 25 meters.



The Big TeraFabric Build

Finally, the TeraFabric Plugfest’s array of Ethernet technology developers, solutions providers and equipment manufacturers brought it all together to successfully construct and transmit data through a 1Tb/s-capable fabric. Multiple Ethernet speeds, technologies and protocols were stitched together to provide the assembled Ethernet ecosystem with a real, live look into tomorrow’s terabit-enabled future.

This was no gimmicky stunt. Rather, the event’s “big build” was an attempt to emulate what a real-world data center is striving to achieve—an integrated network of proven and leading-edge solutions, sourced from multiple vendors, operating at different speeds and all working together seamlessly to deliver unprecedented network performance. And it all depends on interoperability. As illustrated in Figure 3, 10Gb/s Ethernet for iSCSI and FCoE, 40Gb/s Ethernet trunk lines and 10Gb/s-to-40Gb/s breakout cables were used to connect initiator servers and target storage, and data analyzers and generators simulated a real-world movement of data through the network. This marked the first step in the process of demonstrating a 1Tb/s-capable fabric; future Ethernet Alliance plugfests will increase bandwidth over the network and continue to measure throughput over the network.

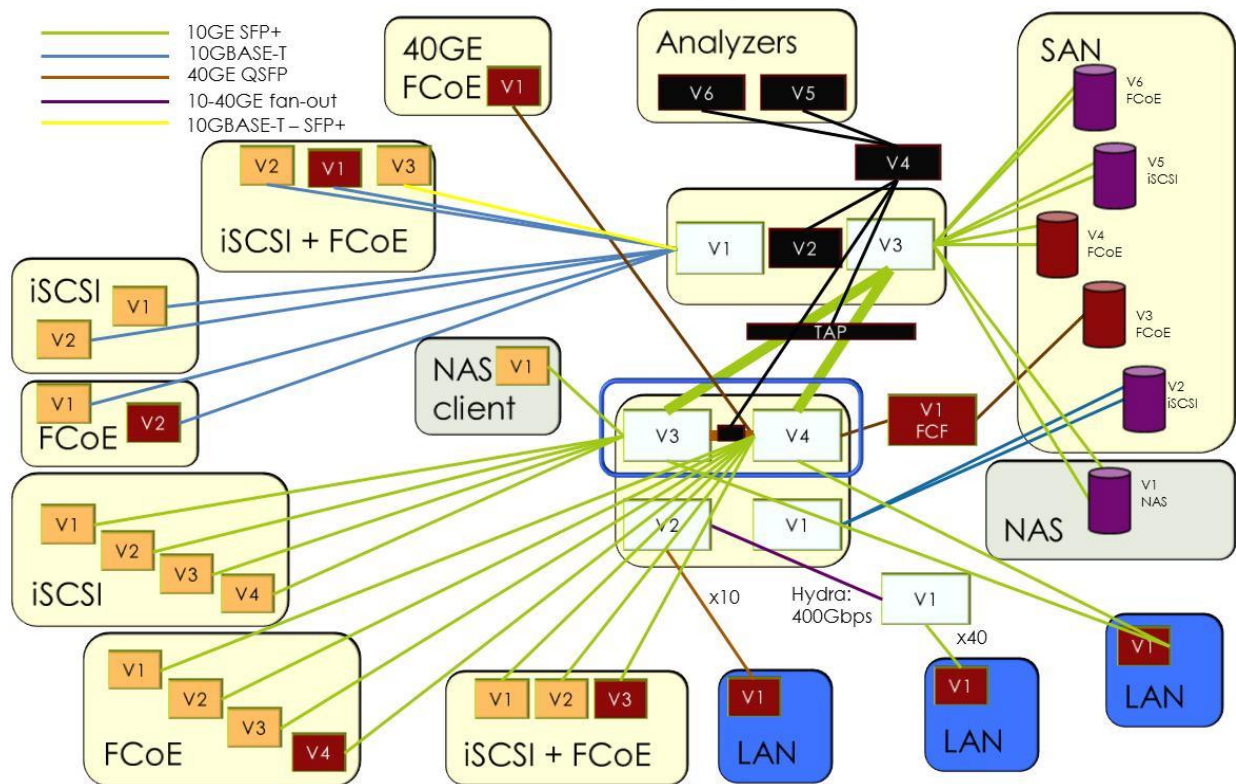



Figure 3: The 2012 Ethernet Alliance TeraFabric Plugfest featured a 1Tb/s-capable fabric comprised of equipment from the event’s array of participating vendors (Vn).

The Ethernet Alliance earlier in 2012 had hosted a Technology Exploration Forum to pulse successful companies from a variety of industries—energy, media, communications and software, among them—on their perspective on the Ethernet interoperability challenge. The lasting impression from that event was that today’s data-center managers require assurance of multidimensional interoperability: between end points in a network, among different sections of the network, across a vendor’s generations of product offerings and even between different chips within the same device.

The 1Tb/s-capable fabric that was successfully constructed at the TeraFabric Plugfest offered data-center managers the type of interoperability assurance that is most meaningful to them. Across network layers and boundaries, technology generations and competitive vendors, end-to-end Ethernet interoperability is possible at unprecedented speeds.



Figures 4 and 5: Ethernet Alliance plugfests function effectively as giant laboratories for the best minds in the Ethernet ecosystem.

Conclusion

Not only did the TeraFabric Plugfest achieve its previously defined interoperability-testing goals, it also provided the Ethernet Alliance and global Ethernet ecosystem with a sense of its next range of needs. The frontiers of interoperability testing to be explored in future events will include more direct-attached testing, testing with more lengths of cables, more tests of DCB at 40Gb/s Ethernet and more broad-based tests with 40Gb/s and then, as it becomes more available in data center environments, 100Gb/s Ethernet.

Even in a technology space as established as Ethernet networking, the pace of



innovation demands rigorous commitment to ensuring interoperability if service interruptions, unnecessary costs and operational complexities are going to be avoided. IEEE 802[®] is constantly being enhanced with new features and capabilities, and—as data-center managers are too well aware—mere compliance to a standard does not necessarily translate into seamless interoperability. Plus, Ethernet is characterized by such a tremendous breadth of products, with multiple vendors to select from for core and edge switches, routers, network interface cards (NICs), physical media and storage. Data-center managers require assurance that the entire array of Ethernet components will work together flawlessly—across both green-field networks and legacy infrastructures—if they are to feel the confidence to pursue their most aggressive strategies for growth.

To encourage Ethernet’s ongoing global expansion and advancement and fuel the trend toward convergence of multiple technologies onto an Ethernet fabric, the Ethernet Alliance will continue to pursue events such as the TeraFabric Plugfest and real-world interoperability demonstrations at public events that engender confidence in end-to-end Ethernet interoperability.