The 1/10 Gigabit Ethernet Ripple Effect

Ethernet continues to adapt to meet the ever increasing demands of high-performance networking. New standards nearing completion will make possible a unified 1/10 Gigabit Ethernet fabric solution.

by Nauman Arshad
Curtiss-Wright Controls Embedded Computing

Ethernet is hands down the most pervasive networking technology in the world. Nearly all wireline traffic on the Internet starts or ends with an Ethernet connection. As standards have evolved from 10 Mbit/s, 100 Mbit/s and 1 Gbit/s to 10 Gbit/s data rates, Ethernet has continually transformed markets by creating new ones, while defeating competitive network interconnect technologies.

In the process, a multi-billion dollar ecosystem has arisen of companies that have invested in Ethernet's continued ability to adapt to meet the ever increasing demands of high-performance networking. This continued commitment to Ethernet has led to high-volume economies of scale that drive ongoing investments in R&D, technology innovations and standardization activities.

Today, momentum is clearly building for 10 Gigabit Ethernet (GbE) in the commercial space for high-end servers, network backbones and storage applications. The number of 10 GbE ports used in commercial markets is predicted to grow from approximately half a million in 2006 to over 9 million in 2010, a substantial projected increase over the next four years. At these rates, 10 GbE is emerging as one of the fastest growing networking technologies.

With a strong growth curve, 10 GbE is following in the successful footsteps

Standard or Draft	Ratification Date	PHY Types	Medium	Reach (max)
Std. 802.3ae	2002	10GBASE-SR 10GBASE-LR 10GBASE-ER 10GBASE-LX4 10GBASE-SW 10GBASE-LW 10GBASE-EW	MMF SMF SMF MMF/SMF MMF SMF SMF	300m 10km 40km 300m/10km 300m 10km 40km
Std. 802.3ak	2004	10GBASE-CX4	Coax	15M
P802.3an	June 2006	10GBASE-T	Cat 6 UTP or better	100M
P802.3aq	Sept 2006*	10GBASE-LRM	MMF	220M
P802.3ap	March 2007*	10GBASE-KX4 10GBASE-KR	4-lane FR backplane 1-lane FR4 backplane	1M 1M

* Expected ratification date

Figure 1 10 Gigabit Ethernet Standards and Draft Standards.

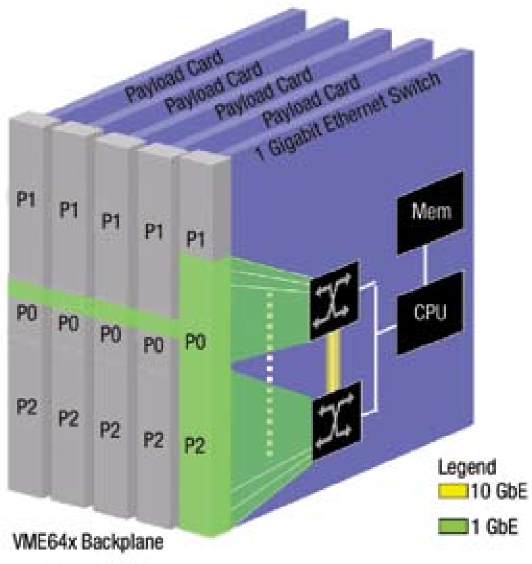

Figure 2 A VME64x High-Performance 1 Gigabit Ethernet Star Topology with 10 Gigabit Ethernet on board stacking.

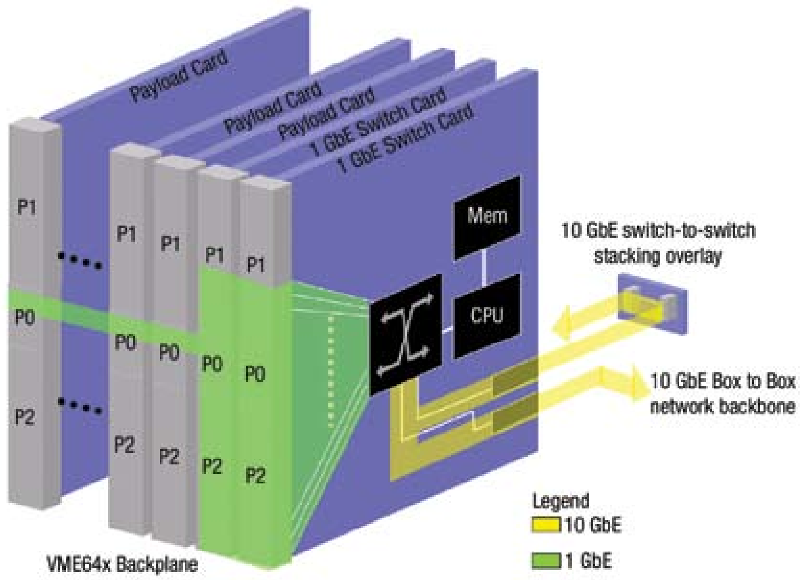

Figure 3 A VME64x High-Performance 1 Gigabit Ethernet Star Topology with 10 Gigabit Ethernet off-board stacking and box-to-box uplink.

of its 1 GbE predecessor. 10 GbE is a disruptive technology that promises to transcend multiple applications and markets over time. This momentum in the commercial space is creating a ripple effect that is finding its way into the emerging market of network-centric infrastructure for rugged applications such as defense.

Implementation of network-centric warfare is driving battlefield communications to embrace bandwidth-hungry secure IP networks. With compulsory DoD requirements for network-ready resources that can seamlessly interface into the Global Information Grid (GIG), switched 1 GbE has emerged as the preferred interconnect for high-bandwidth IP platform networks. The natural progression is to move to multi-speed 1/10 GbE to enable a scalable interconnect solution that can keep up with the increasing needs of high-bandwidth applications over several years. This progression from 1 GbE to 1/10 GbE has already started to appear at the component level.

The market dynamics of hardware components, such as switch chips, PHYs and NICs, are shifting from handling auto-negotiation among 10/100/1000 Mbit/s Ethernet to auto-negotiation between 1 and 10 Gbit/s Ethernet. Leading switch chip vendors such as Broadcom, Fulcrum and Marvell provide multi-speed 1/10 GbE switch chips, so from a switching perspective, low-latency, line-rate 10 Gbit/s technology is available today.

However, one of the major performance hurdles with both 1 and 10 GbE is the CPU performance bottleneck when a 1 or 10 GbE interface terminates at a processor card. As packets arrive through a high-capacity Ethernet interface, more processor compute cycles are diverted to processing the incoming packets that carry a TCP/IP protocol overhead than are focusing on application processing. This bottleneck became apparent with 1 GbE technology but it was not critical for communication over a 1 Gbit/s interface that operated at less than full capacity, since part of the processor's cycles could be used for handling the network interface and part of it for the application.

With 10 GbE this limitation is painfully obvious and a standardized working solution is a required prerequisite for usable IP-based 10 GbE termination. For example, handling a single 1 Gbit/s termination at full speed maxes out current 1 GHz 32-bit processors, let alone a full duplex 10 Gbit/s termination. However, with the evolution of TCP offload engine (TOE) technology, which offloads the CPU from TCP/IP protocol processing, and Remote Direct Memory Access (RDMA), which enables high-performance memory-to-memory data transfers with deterministic low latencies, both 1 and 10 GbE termination bottlenecks are being alleviated.

Recently, vendors such as such as NetEffect have emerged from "stealth" mode to provide Internet Engineering Task Force (IETF) standardized Internet Wide Area RDMA Protocol (iWARP) NIC chip implementations that feature built-in TOE and RDMA functionality. iWARP is not just limited to WAN networks, but is finding its way into embedded designs as a way to enable high-performance blade-to-blade communications over 1 and 10 GbE.

10 GbE Standardization Efforts

Although various 10 GbE standards have been around for a number years (Figure 1), cost and performance hurdles have so far delayed it from becoming widely adopted. The primary factors that limited widespread adoption in the commercial world were due to initial standards that required costly 10 GbE optical components or due to the limited short reach of 15 meters using coax.

In June, the IEEE Standards Association standards board approved a new copper standard, IEEE P802.3an (10GBASE-T), which defines 10 GbE operation over unshielded or shielded twisted-pair, e.g., augmented CAT 6 or better. With the introduction of a lower-cost copper solution, 10GBASE-T shows promise to follow the success and economies of scale of 1000BAE-T. Additionally, improvement in distances obtainable with twisted-pair copper cabling now extends the reach of 10 GbE to 100 meters over low-cost copper, making it ideal for various enterprise backbones.

802.3aq (10GBaseLRM) provides an efficient Small Form-Factor Pluggable (SFP) optical solution for running 10 Gbits/s over multimode-fiber backbones at distances of up to 220m. Although 10GBaseLX4 was the first optical interface standard developed to run 10 Gbits/s over multimode fiber, 10GBaseLRM is expected to establish a leadership position due to its lower cost, smaller size and low power. Its smaller sized connectors are ideal for real-estate limited 6U and 3U VME/VPX boards. Additionally, the use of optical interconnection for box-to-box platform networks is preferred in rugged environments, due to its immunity to EMI and RFI.

Although not used heavily by defense board manufacturers, the current dominant 10 GbE technology for board and backplane routing in the commercial embedded systems space is XAUI (802.3ae), which requires eight pairs of wires (16 pins) per channel. In embedded defense applications, XAUI has mostly been limited to an onboard interconnect while the backplane interconnect has been dominated by switched 1 GbE.

However, an up-and-coming new draft standard under development, 802.3ap, aims to define new 1 and 10 GbE PHY transport mediums that can be used in high-speed backplanes. The 1 GbE PHY defines 1000BaseKX that provides 1.25 Gbits/s over a single lane with 8B10B encoding. The 10 GbE parallel version, 10GBase-KX4, splits the 10 Gbit/s signal into four lanes of 3.125 Gbits/s with 8B10B encoding per lane, while the 10 GbE serial version, 10GBase-KR, defines one lane at a full 10 Gbits/s with 64/66B encoding. Although 802.3ap is not expected to be finalized until March 2007, it shows great promise for physical standards that can be used in 1/10 GbE backplane routing for next-generation, high-performance VITA46 backplanes.

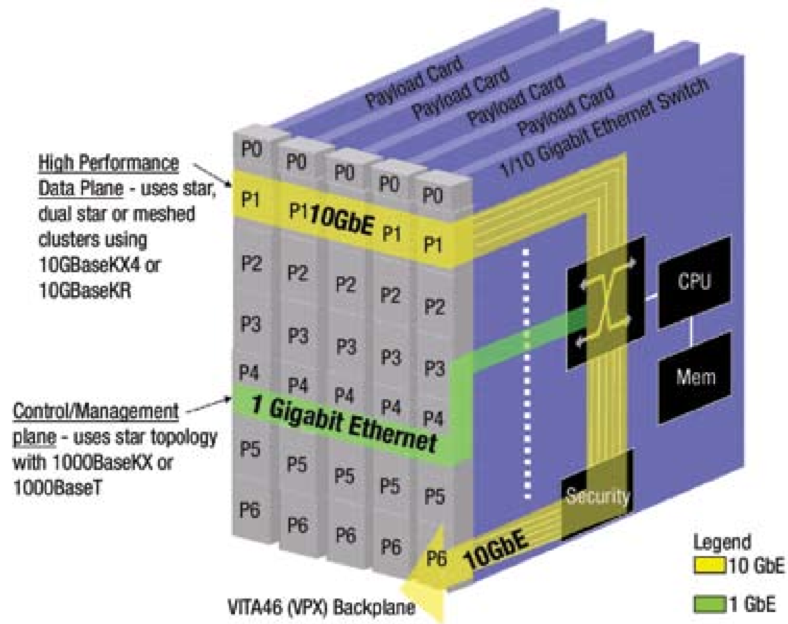

Figure 4 Conceptual 1/10 Gigabit Ethernet System Leveraging next-generation VITA46 and 802.3ap standards.

10 GbE and Legacy VME64x Systems

In embedded defense boards and subsystems, popular backplane standards such as VME64x are limited in their ability to support 10 GbE's higher multi-GHz signaling. Due to this limitation—as well as to the processing overhead bottleneck—current use of 10 GbE for VME64x-based systems has been limited to onboard stacking of switch chips, front panel off-board stacking/cascading of switch blades, or as box-to-box high-performance platform backbones.

An example of onboard 10 GbE stacking can be seen on Curtiss-Wright's VME-680 SwitchBlade, which uses dual 12-port Broadcom BCM5690 multilayer switch chips that are stacked using a single 10 Gbit/s interface (Figure 2). Onboard stacking allows up to 24 x1 Gbit/s ports for larger applications while simplifying management and configuration by using a single IP instead of two separate IP addresses for each onboard multi-layer switch.

Off-board 10 GbE stacking is also possible. In the company's VME-682 FireBlade, the 10 GbE interfaces are routed to the front. With two VME-682s in adjacent slots, the 10 GbE interfaces on each card can be stacked using a simple low-profile 10 GbE front panel overlay (Figure 3). This will enable a 48-port stacked solution in a VME64x chassis. This approach enables larger port counts in existing VME64x systems that can scale from benign to extremely harsh environments.

The off-board 10 GbE ports can also be used to architect high-performance platform-wide backbones. With 10 GbE-capable switches in each box, the switches can be connected using their 10 GbE ports. Meanwhile, inside the box, 1 GbE interfaces can be fanned out to connect each blade within a VME64x box, providing a balanced network topology.

The system-wide 1/10 GbE backbone provides a unified infrastructure that can facilitate easier inter-blade or inter-box management, configuration, monitoring and data traffic transportation. The 10 GbE box uplink port can also be combined with security features on an advanced router card to control the IP traffic flowing into and out of the box through the use of firewalls, filtering and packet inspection.

10 GbE and Next-Generation VITA 46 (VPX)

The VITA 46 (VPX) next-generation backplane standard is specifically targeted at rugged environments. It offers seven rugged Tyco Electronics Multi-Gig RT2 connectors, P0 to P6, to form the backplane for a 6U card and three connectors, P0 to P2, for smaller form-factor 3U cards. P0 is a half-height utility connector used in both 6U and 3U versions, while the other connectors are all full height. VITA 46 provides unmatched pin-count density, with a total of 384 pins per slot. Designed for extremely harsh defense environments, VITA 46—along with the VITA 48 (VPX-REDI) draft standard that defines cooling specifications for rugged systems—offers a unique solution of increased backplane performance, support for high-speed signaling fabrics, increased pin counts, rugged connectors with built-in ESD protection and support for in-field two-level maintenance.

The base specification, 46.0, defines a distributed switching architecture that enables five card mesh clusters, where each payload card can connect to every other card. To apply 10 GbE within the VITA 46 standard, the 46.7 substandard must be completed. The 802.3ap 10GBaseKX4 would be an ideal fit for 46.7, while 10GBaseKR would enable increased scalability in future implementations to provide 40 Gbit/s bandwidth using the same backplane routing.

Furthermore, the completion of another key sub-standard, currently known as V46.20, would introduce centralized switch slots that could support star, dual-star and hybrid (e.g. star + mesh) architectures within VITA 46. This would be valuable for systems that do not need meshed architectures or have dual topology requirements where meshes and star/dual stars are required.

Until the new physical interface 802.3ap 10 GbE specifications, and backplane VITA 46 along with all its dot specifications, achieve the status of official standards, their benefits remain conceptual.

The benefits of a hypothetical VITA 46/802.3ap solution would yield a unified, multi-speed fabric that can support IP traffic, an ideal fit for NCW and GIG interoperability (Figure 4). Furthermore, the unified messaging over IP would simplify application design and portability without the need to learn new fabric protocols. Finally, the massive software ecosystem, tools and trained professionals already available for IP protocols could be leveraged to drastically accelerate the time-to-market and long life cycle support required for defense applications.

With exciting new standards like 802.3ap and VITA 46 getting closer to completion, a unified 1/10 Gigabit Ethernet fabric solution for defense platforms is possible in the near future. Once again, Ethernet's tide is poised to ripple across multiple markets including defense. The ubiquity of 1/10 GbE is a certainty, but the application possibilities are wide open. ◢

Curtiss-Wright Controls
Embedded Computing
Ottawa, Ontario, Canada.
(613) 599-9199.
[www.cwcembedded.com].