



Hardware Assets

Ethernet's Staying Power

Ethernet Strengthens Its Military Connections

Armed with proven longevity and popularity, Ethernet is capturing military mindshare without much fight. A rich set of military-applicable Ethernet specs seal the deal.

Jeff Child,
Senior Editor

The military's affection toward Ethernet, although far from new, is reaching a whole new level. Ethernet offers just the kind of longevity and ubiquity that the military loves, and its general strengths are compelling: it's been around in some form for 25 years, used by the vast majority of computer networks worldwide, and offers constant ramp-ups in performance. Those strengths, combined with the wealth of system specification work accomplished over the past few years within the embedded bus-board community, make Ethernet a no-brainer for a broad number of defense system designs. Among them are shipboard data communications, avionics suites, mobile C3 centers and back-end data collection links for C4ISR systems.

As the most popular means of transporting Internet Protocol (IP) packets, Ethernet is pushing into an even wider scope. Public switched telephone networks (PSTN) for example are beginning to transition from a circuit switching to a packet switching methodology. Beyond the LAN, Ethernet/IP has become the networking standard of choice in telecommunication WANs and MANs, but also in data communications, wireless networks, home networks and factory floor networks.

Ethernet's performance curve has an impressive history, from 10 Mbit/s

10BaseT and 10Base-2 (thin cable alternative to 10 BaseT) to 100 Mbit/s 100BaseT, and then on to Gigabit Ethernet within the last few years. Continuing the upward ramp, last summer the 10 Gbit Ethernet Standard IEEE 802.3ae was ratified. Along the way, Ethernet has held on to its dominance in computer networks, except for a brief period in the early 90s when Asynchronous Transfer Mode (ATM) seemed on its way to taking over all aspects of networking, but ultimately did not. Ethernet's success exemplifies the power of evolutionary solutions versus more revolutionary approaches.

Still new to even the commercial communications realm, the 10 Gbit Ethernet version is significantly different in some respects from earlier Ethernet standards, primarily in that it will only function over optical fiber, and only operate in full-duplex mode. That means it has no need for the carrier-sensing multiple-access with collision detection (CSMA/CD) protocol that defines slower, half-duplex Ethernet technologies. In every other respect, such as packet format, 10 Gigabit Ethernet remains true to the original Ethernet model.

Navy Shifts to Ethernet

Several years ago, ATM looked so attractive that the Navy chose it for its next-gen high-speed bus. Today the Navy plans to replace ATM SONET technology

planned for the Shipwide Area Networks (SWAN) over to Gigabit Ethernet. In late 2000, the Navy's Space, Command and Control, and Information Warfare directorate in the Office of the Chief of Naval Operations (OPNAV), endorsed Gigabit Ethernet, supported by a mesh-network topology, as the Navy's standard shipboard network architecture.

That reversed a 1997 plan that endorsed ATM SONET as a replacement for standard Ethernet and FDDI networks. ATM networks are currently installed on aircraft carriers, amphibious ships, Ticonderoga-class Aegis cruisers and Los Angeles-class attack submarines. OPNAV officials plan to have 10 Mbit/s standard Ethernet replaced by new 10/100 Mbit/s fast Ethernet as the standard design for the IT-21 network, now called the integrated shipboard network system (ISNS).

Ethernet Over the Backplane

Of particular interest to the military realm is Ethernet's expanding role in embedded system backplanes. Both PICMG and VITA have established specs for using Ethernet as a backplane transport. The PICMG 2.16 architecture overlays an Ethernet-based packet switching architecture on top of CompactPCI to create an embedded System Area Network (SAN). The PICMG 2.16 standard uses a dual redundant star architec-

ture using Gigabit Ethernet paths and managed by high-performance Ethernet switch fabrics. In effect, this provides node-to-node connectivity within the backplane that resembles a typical

Ethernet: Back-End Link for C4ISR Data Collection

One possible role of Ethernet is as a back-end storage link within data collection applications such as C4ISR (command, control, communications, computers, intelligence, surveillance and reconnaissance) systems. Ethernet's reach probably won't extend to the radar and sonar edges of C4ISR data collection. In radar, for instance, once a channel is open to a radar sensor suite, the system begins transferring massive amounts of data in one direction, keeping the channel open. That's why Fibre Channel is the preferred technology for radar and sonar data collection. Using standard Ethernet wouldn't make sense in this role. Ethernet would have to spend time on packetizing the data—putting in all the headers, footers and so forth—and then de-packetizing it on the other end.

In contrast, core C4ISR installations do data collection in a different way and need a different back-end infrastructure. So while the size of the radar sensor data seems to lend itself to Fibre Channel, the more general C4ISR communication I/O space is gravitating toward Gigabit Ethernet or just 10/100 Ethernet. Another alternative may be the emerging iSCSI standard. iSCSI uses Ethernet as a transport for data from servers to storage devices or storage-area networks. The block-level storage protocol lets users create a separate storage network using Ethernet.

The iSCSI protocol puts standard SCSI commands into TCP and sends those SCSI commands over standard Ethernet. Because iSCSI uses SCSI commands, relying on Ethernet only to transport the SCSI commands, operating systems see iSCSI-connected devices as SCSI devices and are largely unaware that the SCSI device resides across the room or across town.

Ethernet computer network. The PICMG 2.16 CompactPCI P3 connector has two Gigabit Ethernet ports for improved performance and redundancy.

An example implementation of PICMG 2.16 is Performance Technologies' IPnexus CPC6600 (Figure 1). The system is an embedded Ethernet switch providing Layer 2/Layer 3 Gigabit Ethernet to every slot in a PICMG 2.16 chassis. Aimed at applications such as wireless infrastructure equipment, military communications systems and VoIP media gateways, the switch provides 24 Ethernet ports, which support 10 Mbit, 100 Mbit, and Gigabit Ethernet. The system features an advanced fast filter processor for wire speed Layer 2-7 packet classification and filtering.

Meanwhile, VITA has its own Ethernet backplane scheme called VITA 31.1-2003: Gigabit Ethernet on VME64x. This standard defines a pinout and interconnection methodology for implementing a 10/100/1000 Base-T Ethernet switched network on a VME64x backplane. The spec achieved official ANSI standard status this month. The CompactPCI P3 connector and the VME64x P0 connector are identical and have the same placement on the backplane. The VITA 31.1 spec adopts the PICMG 2.16 P3 connector pinout for use on VME64x boards. It also borrows the definition of the fabric card described in PICMG 2.16. As a result, PICMG 2.16-compliant systems and VITA 31.1 systems can use the same switched fabric boards.

Ethernet on a Mezzanine

Also within PICMG's arsenal are other specs that integrate Ethernet, but with a more telecom-specific aim. One of the first fabric specifications finalized for the new AdvancedTCA family of specifications was PICMG 3.1 Ethernet/Fibre Channel. Ratified in March, PICMG 3.1 maps the signaling protocols defined for Ethernet by IEEE 802.3-2003 and the Fibre Channel FC-PI (Physical Interfaces) specifications onto the Fabric Interface defined by PICMG 3.0. The Ethernet definitions support speeds up to 10 Gbit Ethernet connections.

Both Ethernet and Fibre Channel links may exist on a single PICMG 3.0 channel depending on the data rates supported. It's unclear whether these schemes will have applicability in defense applications. For high-end backplane computing, the military is unlikely to stray from VME. That said, the DoD's heavy emphasis on network-centric communications makes PICMG's next-generation CompactPCI, called AdvancedTCA, worth a look.

In contrast, Ethernet via a mezzanine is probably more palatable to military designers. PICMG recently revised its PICMG 2.15, PCI Telecom Mezzanine Card (PTMC) specification, adding two new options to augment the PTMC connector configurations. One enhances TDM capacities by extending the TDM (H.110) bandwidth and adding Ethernet links, and the other combines ATM (UTOPIA Level 2) capabilities with Ethernet links. Both of these new options standardize use of



Figure 1

Performance Technologies' IPnexus CPC6600 is an example of a PICMG 2.16 implementation. The embedded Ethernet switch offers Layer 2/Layer 3 Gigabit Ethernet to every slot in a PICMG 2.16 chassis. The switch provides 24 Ethernet ports, which support 10 Mbit, 100 Mbit and Gigabit Ethernet.

Ethernet MDI links on the PMC connector at 10 Mbits/s, 100 Mbits/s and 1 Gbit/s. Because PTMC is based on the same four-connector standard as PCI Mezzanine Card (PMC), it works on the same form-factors as PMC, including VME.

Commercial Realm Leads Avionics Charge

For its part, the avionics realm is also riding the Ethernet wave. Here the charge is being lead by the commercial avionics realm. Airbus Industries developed its Avionics Full Duplex Switched Ethernet (AFDX) scheme for use in the new A390 Airbus airliner, due to enter service in 2006. In principle AFDX is a switched Ethernet solution, but with each data link having a fixed bandwidth. Packets are sent and received in sequence, both as point-to-point and one-to-many communications. Two redundant networks are employed, closely monitored to ensure that data loss or errors cannot affect performance or give rise to single point failure conditions while also allowing for signal loss or corruption.

AFDX provides for electrical and logical redundancy—in other words dual routes and switches for every VL. Faulty packets are identified at the end-system by ordinal integrity—similar to TCP/IP. But while TCP/IP can leave conventional connections hanging—waiting for missing frames—that's not acceptable in an aircraft. Since only the most up-to-date data is needed, an ordinal integrity algorithm is used that only allows good packets within a count of +1 or +2 to pass.

Meanwhile, ARINC 644, an effort under ARINC's Aircraft Data Network (ADN) Working Group, provides a top-level set of standards for the Ethernet-based data networks for new airplane designs. ARINC 664 is essentially a rather broad set of best practice documents defining how Ethernet can be used in the data highway environment of commercial aircraft. Last fall the group incorporated the Internet Protocol version 6 (IPv6) into ARINC 664. ■■

10 Gigabit Ethernet Alliance
Newport Beach, CA.
(949) 250-7155.
[www.10gea.org].

AFDX on AIM-Online
Valley, NE.
(866) AIM-1553.
[www.afdx.com].

PICMG
Wakefield, MA.
(781) 246-9318.
[www.picmg.org].

Performance Technologies
Rochester, NY.
(716) 256-0200.
[www.pti.com].

VME Industry Trade Association
Scottsdale, AZ.
(490) 951-8866.
[www.vita.com].