The Top 4 Networking Challenges Facing Your Embedded Systems

CURTISS-WRIGHT

Embedded networks serve as the backbone of a platform's technology architecture, and act as the front line against remote security attacks. Ensuring that all modules and systems are connected for seamless communication by a trusted, reliable network is critical to mission success.

In this overview, we examine the four biggest challenges facing today's embedded networks. In upcoming papers, we'll explore each of these challenges in greater depth, along with the technology solutions available to address them.

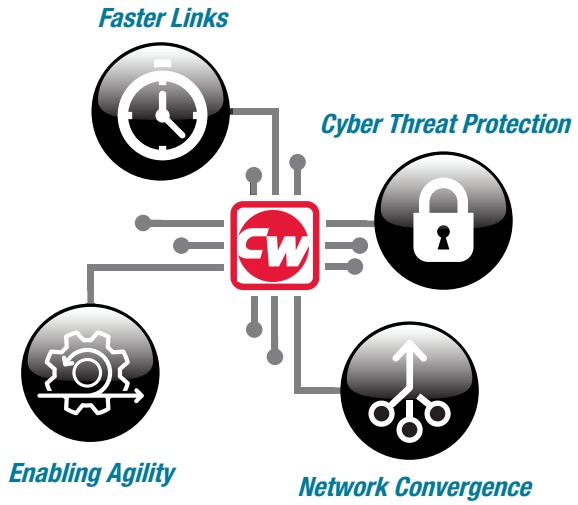

Challenge #1: Delivering Faster Links

Designers of rugged military and aerospace systems continue to look to the commercial world for networking technology, leveraging available Ethernet switch chips to connect between processing modules and deploying PHY devices to link systems over fiber optic or copper cables. Using the latest high-performance networking technology in deployed systems based on modular COTS standards isn't always practical, however. While the latest commercial devices support IEEE standards for 100G Ethernet and beyond - based on 25G and 50G SerDes technology - OpenVPX systems using the standard MultiGig-RT2 connector have so far been limited to 10G SerDes signals. To take advantage of these high-speed standards in modular systems, backplane and connector technology needs to catch up.

Terabit throughput and sub-microsecond latency are the ambitious design goals for today's most demanding systems. In the past, meeting the performance needs for defense electronic applications often necessitated the use of custom hardware and FPGAs. Today, commercial silicon delivers levels of performance and functionality that are hard to beat. The technical challenge faced by many designers is how to implement high-speed interfaces in board-and-backplane modular systems. Modular systems based on VPX backplanes have been proven and deployed using 10G and 40G Ethernet (either by combining four 10G SerDes lanes or through use of the VITA 66 fiber optic standard), but moving to 100G Ethernet and beyond will require the use of four lanes of 25 Gbps (or faster) SerDes technology.

The use of Ethernet at 10Gbps and above has also driven a growing need for optical interconnect solutions. A wide array of commercial products are available for implementing optical links. However, no single approach has emerged as the de-facto standard for rugged systems. As a result, few rugged COTS products integrate optical interfaces, leaving it to integrators to convert from electrical to optical in their systems.

Designers using high-performance and feature-rich commercial Ethernet technology to connect their embedded systems can also deploy a growing range of techniques for managing traffic flows on their networks. Faster networks can carry more traffic from more sensors and applications on a single link, but care must be taken to ensure critical data isn't delayed or dropped.

Challenge #2: Network Convergence

Faster Ethernet links also mean it's now practical to use a single network to carry data from multiple different network services over a single cable. Much as enterprise computer networks now carry voice-over-IP on the "computer" network, Ethernet in aerospace and defense systems can now carry a mix of voice, video and data from various sensors and applications. This "converged" network can replace multiple single-purpose cables, providing substantial SWaP benefits and increased flexibility when adding new capabilities to a platform.

As network designers seek to connect multiple sources over the same data pipe, a big concern is the potential for interference or contention that would delay real-time traffic. For time-sensitive or safety-critical applications, deterministic performance may be essential.

One solution for getting closer to real-time networking is to implement policies that ensure critical data has the highest priority. Another approach for avoiding network congestion is to divide the network into fixed time slices where each application or host is given exclusive access to the network during its assigned slice. These approaches can help provide service guarantees, including bounded end-to-end latency on networks that use standards-based Ethernet.

In recent years, a range of IEEE standards have been ratified to enable "time-sensitive networking" (TSN). TSN mechanisms being implemented in new switches have emerged to address the need for real-time scheduling, some of which include:

- 802.1Qbv: Enhancements for Scheduled Traffic – enables time-slicing
- 802.1Qch: Cyclic Queueing and Forwarding – reduces jitter
- 802.1Qcc: Stream Reservation Protocol – allocates capacity to avoid congestion

Whether they implement the latest TSN standards, or vendor-specific QoS mechanisms, today's embedded switches can enable approaches that carry traffic from a variety of sources, combining mission-critical and latency-sensitive real-time data on converged platform wide networks.

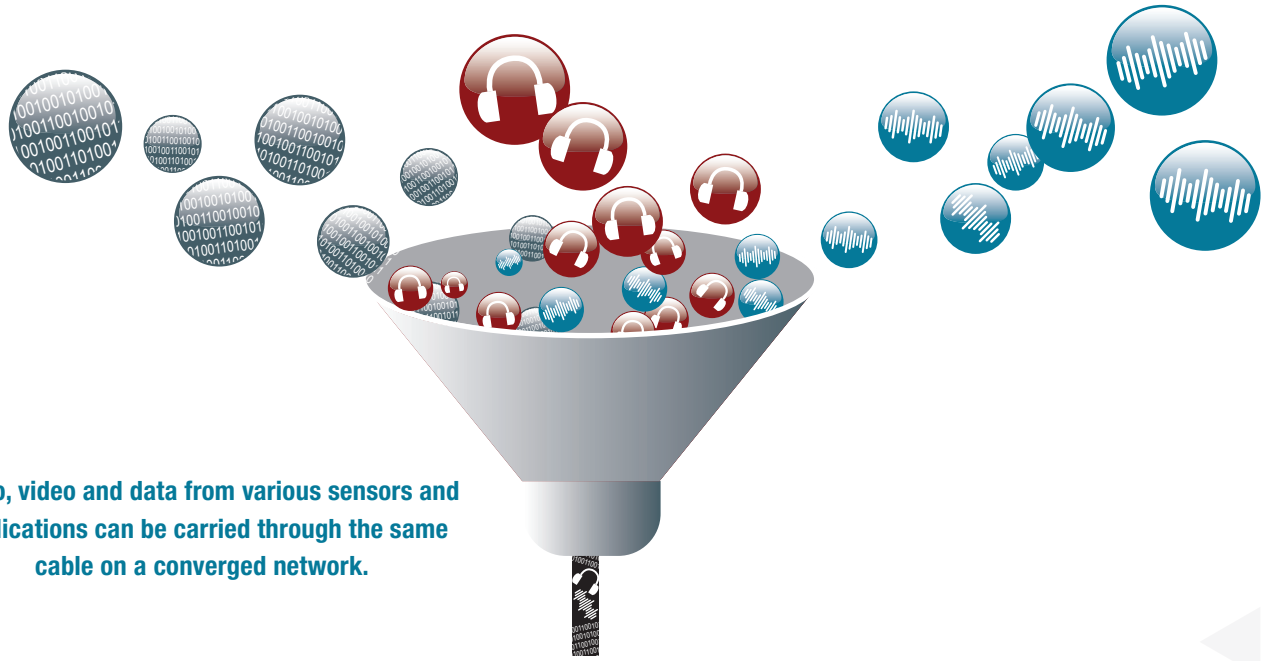

Audio, video and data from various sensors and applications can be carried through the same cable on a converged network.

Challenge #3: Protecting Networks against Cyber Threats

Network security in embedded computing is getting more scrutiny these days. In a constantly evolving threat environment, where new attacks arrive virtually every day, system architects must design networks to be as secure as possible. That requires a constant review process to enable the necessary adaptation, modification, and updates to keep systems safe.

Network security involves providing protections against all devices that are connected or could have access to the network. In this area, embedded architectures are catching up to enterprise networks. In the enterprise environment, where there has always been the risk of an unauthorized person connecting on a port in an office or conference room, the need to lock down the network is well understood. In comparison, airborne or ground vehicle tactical networks typically have been very controlled, with no network ports exposed. Physical access to ports in the past was easy to control. Today, however, we are seeing embedded networks connecting more devices and making more connection ports available. Aboard commercial jetliners, for example, Ethernet might be available at every seat, and Wi-Fi might be provided for entertainment. This expanded offering makes Trusted Computing approaches imperative.

One tool for securing the network is white-listing, or limiting access to only trusted devices. This could be as simple as enabling each port only to allow traffic from a known MAC address. While simple to implement, MAC addresses can be changed and spoofed. Trusting a device just because it has the right address turns out not to be a very robust security solution.

A more advanced technique to keep out unknown users involves IEEE 802.1x for port-based network access control (PNAC). IEEE 802.1x enables the network to authenticate a network endpoint using a cryptographic exchange. Instead of trusting a MAC address, trust is based on a certificate or other credentials. While hard-coding and 802.1x enable control over what devices can access the network, IEEE 802.AE MACsec (for Layer 2) and IPsec (for Layer 3 traffic) uses encryption to protect data on the move and prevent someone from snooping into that data.

Keep in mind, too, that network security doesn't end with the architecture or initial implementation; it needs to be revisited on a regular basis. Patches likely will be necessary to address vulnerabilities at each iteration, at every software upgrade, and every time a new device gets added.

Network security involves providing protections against all devices that are connected or could have access to the network.

Challenge #4: Enabling Agility

As technology continues to rapidly evolve, system integrators must be able to quickly add new functionality to deployed military and aerospace platforms. Adding new hardware point solutions to provide every new capability results in unnecessary cost and SWaP consumption, not to mention a cluttered and challenging in-vehicle experience for platform operators. Plus, the logistics and time required to recall platforms to a service depot for a hardware upgrade adds complexity and delays. The ability to add new capabilities to existing hardware or address bugs, security issues and vulnerabilities through software upgrades greatly reduces these challenges.

In a similar vein, being able to run networking as a software application on a single board computer, for example, instead of having a separate router card allows for speedy, simplified updates and SWaP reduction.

Initiatives to support this hardware/software convergence, such as the C4ISR/EW Modular Open Suite of Standards (CMOSS), will make it easier and more cost-effective to upgrade capabilities or keep pace with commercial technology by eliminating complex integration challenges, lack of competition, and proprietary interfaces.

Learn More

Curtiss-Wright Defense Solutions White Papers

- Copper or Fiber for Military & Aerospace Networks?
- Choosing a Rugged Ethernet Switch/Router Solution
- Latency: Understanding Delays in Embedded Networks
- Linking Outside the Box: Connecting Embedded Systems to Wide-Area Networks
- Multi-Gigabit Ethernet – Beyond 1000BASE-T
- Staying Connected: High Availability Embedded Networking

Authors

Andrew McCoubrey
Senior Product Manager
Curtiss-Wright Defense Solutions

Mike Southworth
Senior Product Manager
Curtiss-Wright Defense Solutions

© 2019 Curtiss-Wright. All rights reserved. Specifications are subject to change without notice.
All trademarks are property of their respective owners I B45.0619