CRITICAL I/O WHITE PAPER

25Gb Ethernet -
An Easy Upgrade from 10GbE

Abstract

This paper looks at the tradeoffs and challenges of implementing higher performance optical Ethernet networks in rugged embedded systems, and suggests that for many applications, 25GbE is the most straightforward path to dramatically increased Ethernet performance

Revised 6/10/22

CRITICAL IO

25Gb Ethernet – An Easy Upgrade from 10GbE

Communication networks in data intensive real-time systems are faced with the need to move data at increasing rates between sensors and processing elements that may include general purpose CPUs/DPUs, Graphics Processing Units (GPUs) and special purpose processing hardware such as FPGA processing elements. And the increasing quantity, resolution, and bandwidth of the sensors drives the required networking bandwidths higher every year.

10 Gigabit Ethernet (GbE) has been a standard in embedded real time systems for many years. However, the increasing bandwidth requirements for applications such as IR, video, LiDAR, RADAR, ELINT and other sensor systems are driving the need to deploy Ethernet at higher data speeds on field-deployed mobile military platforms. When 10GbE is not fast enough, upgrade possibilities include 25Gb, 40Gb, 50Gb, or 100Gb Ethernet. While all these rates are now in common use in data center and high-performance computing environments, the size, power, and environmental constraints have limited their use in deployed systems. But new hardware implementations are now emerging that make these rates usable in rugged embedded systems.

This paper looks at the tradeoffs and challenges of implementing these faster Ethernet networks in rugged embedded systems, and suggests that for many applications, 25GbE is the most straightforward path to dramatically increased Ethernet performance in these systems.

Evolving Ethernet Standards

Ethernet is unique in high-speed networking technologies in that its "spec" bit rate is exactly equal to its usable bandwidth. In other words, overheads are already accounted for in the specified bit rate. So upgraded "spec" bit rates always yield a proportional and predictable increase in usable data rate, as shown in the table 1.

Table 1. Usable Ethernet Data Rates

"Spec" Bit Rate	Usable Bit Rate Per port, each direction	Usable Data Rate Per port, each direction
10GbE	10Gbps	1.25 GB/s
25GbE	25Gbps	3.125 GB/s
40GbE	40Gbps	5.0 GB/s
100GbE	100Gbps	12.5 GB/s

The 10 Gigabit Ethernet standard encompasses several different physical layer (PHY) standards. In the commercial world, endpoints and switches may use a variety of PHY types by using pluggable transceiver modules such as SFP+. The most common 10GbE physical layer standards are shown in Table 2.

Table 2. Common Optical 10GbE Physical Layer Options

10GbE Optical Physical Layer	Lanes	Typical Transceiver	Typical Connector	Typical Cable	Maximum Distance
10GBase-SR	**1**	**SFP+ or rugged RJ**	**LC**	**MMF**	**300m***
10GBase-LR	1	SFP+ or rugged RJ	LC	SMF	10km
10GBase-ER	1	SFP+	LC	SMF	40km

*Maximum distance varies according to specific MMF cable type

Upgrade Options: 25GbE, 40GbE, 100GbE

The most plausible upgrade options to consider are 25GbE, 40GbE, and 100GbE. In evaluating these options for upgrading from 10GbE, one must consider:

- Performance – How much of a performance boost is needed?
- Physical Compatibility – NIC modules, transceiver form factor, connectors, cables
- Functional Compatibility – Protocols, offloads, features
- Software Compatibility – Drivers, applications
- Reliability – Fewer lanes results in better reliability

Figure 1 graphically compares the performance gains as compared to 10Gb Ethernet.

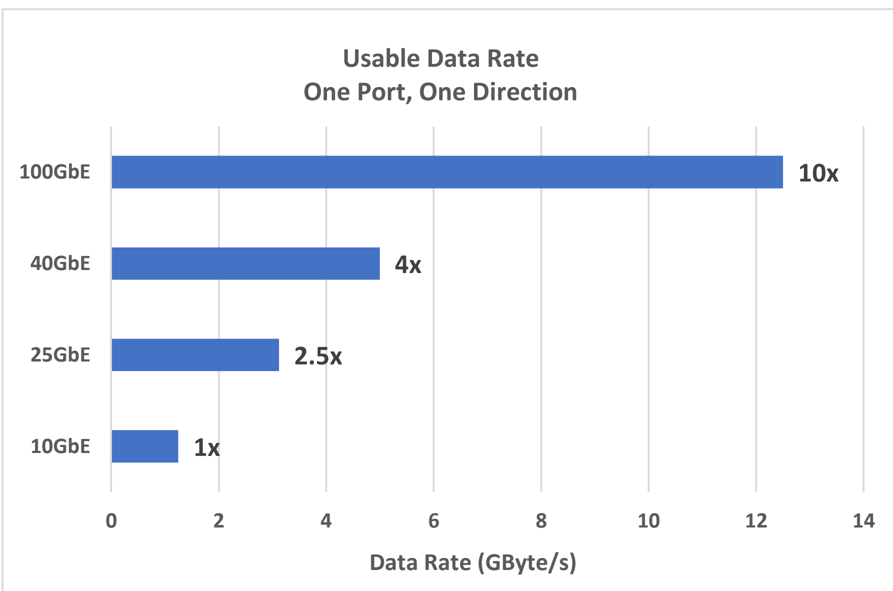

Figure 1: Performance Gain vs. 10Gb Ethernet

In looking at the tradeoffs between these upgrade options, it is useful to look at a bit of Ethernet history, and how these faster Ethernet versions were developed. The initial specifications for 40GbE and 100GbE were developed in tandem. The initial MMF versions of 40GbE were built on 10GbE technology, with 40GbE implemented as four lanes of 10Gb. Similarly, initial MMF versions of 100GbE were implemented as 10 lanes of 10Gb. These choices were made because 10Gb was the practical limit at the time for cost-effective transceiver technology. Essentially, both 40GbE and 100GbE were initially implemented using the same transceiver technology as was used for 10GbE.

However, technology marched on, and highly cost-effective 25Gb transceivers became available. This allowed the more recent versions of 100GbE to use four 25Gb lanes, rather than ten 10Gb lanes. Fewer lanes mean simpler and cheaper cabling and transceivers, and improved reliability due to fewer potential points of failure. Unfortunately, most 40GbE implementations are stuck in the past, and continue to use four 10GbE lanes with the associated higher cabling and transceiver costs due to the high lane counts. Both 40GbE and 100GbE most often use QSFP style transceivers.

25Gb Ethernet is a more recent innovation that is built on the cost-effective 25Gb transceivers developed for newer versions of 100GbE. 25GbE uses a single lane of 25Gb signaling. This achieves 2.5x the throughput of 10GbE with a very minimal increase in cost because 25GbE retains the single lane implementation, cabling, and connectors used for 10GbE. Both 10GbE and 25GbE most often use SFP style transceivers in commercial applications. Rugged 10 and 25GbE embedded applications often use very small form factor “RJ” style transceivers.

Figure 2 compares the typical physical interface implementations of 40GbE vs. 10/25GbE.

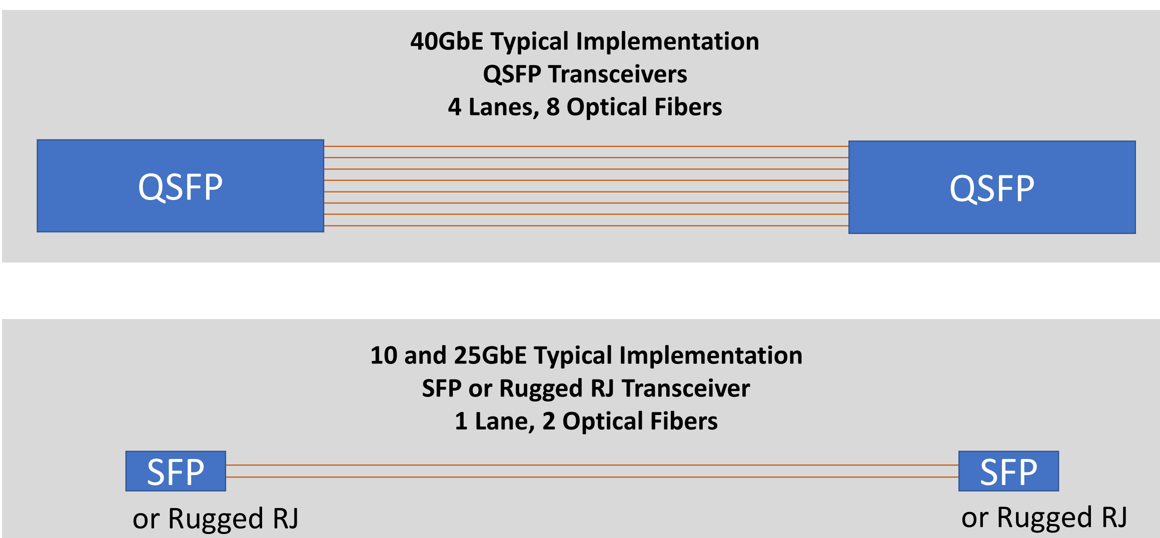

Figure 2. Comparison of 40GbE vs 10 and 25GbE Physical Interfaces

Trading off 25GbE vs. 40GbE

40GbE and 100GbE physical layers are considerably more complex than those used for 10GbE or 25GbE. 40/100GbE physical layers are always hampered by one or more of the following attributes: a) more than one lane, and/or b) more than one TX and RX fiber, and/or c) more than one transceiver wavelength, and/or d) possible bidirectional operation and/or different coding methods. Therefore, the multiple possible combinations of (a), (b), (c) and (d) have yielded dozens of different physical interface standards and pseudo-standards for 40/100 GbE. And this is not even considering additional implementation variations due to different connector types and optical fiber types. These variations are illustrated in Table 3 below. Observe the number of different physical layer standards for 10GbE and 25GbE, vs 40GbE and especially for 100GbE. While having lots of different options may sound good, in reality it leads to significant challenges in ensuring interoperability. As an example, table 4 lists just some of the many different 40GbE physical layer options.

Table 3. Comparison of the Number of Physical Layer Options

Ethernet Rate	Number of common standard and pseudo-standard optical physical layer options	Predominant Transceiver Type	Predominant Connector Type	Predominant Cable Type
10GbE	**3**	**SFP+**	**LC**	**MMF**
25GbE	**3**	**SFP28**	**LC**	**MMF**
40GbE	~7	QSFP+	MPO/MPT	MMF
100GbE	>30!	QSFP28	MPO/MPT or LC	SMF

Note: SFP+ and SFP28 are physically compatible, as are QSFP+ and QSFP28

Table 4. Proliferation of 40GbE Physical Layer Options.

40GbE Optical Physical Layer	Lanes/ Freqs*	Typical Transceiver	Typical Connector	Typical Cable	Maximum Distance
40GBase-SR4	**4/1**	**QSFP+ or rugged optical engine**	**MPO**	**MMF**	**300m**
40GBase-SR2-BiDi	2/2	QSFP+	LC	MMF	150m
40GBase-SWDM4	1/4	QSFP+	LC	MMF	400m
40GBase-LR4	1/4	QSFP+	LC	SMF	10km
40GBase-ER4	1/4	QSFP+	LC	SMF	40km
40GBase-LX4	1/4	QSFP+	LC	SMF	2km
40GBase-PLR4	4/1	QSFP+	MPO	SMF	10km

*More than 1 frequency per lane indicates WDM is used

The multiple options for 100G and 40G can be dealt with in the data center environment by swapping of QSFP modules. But a corresponding approach is challenging to implement in rugged embedded systems where pluggable modules are not always desirable. The large physical size and less than rugged physical design of QSFP modules is a further complicating factor.

10GbE and 25GbE transceivers, on the other hand, are most often implemented with SFP+ or SFP28 modules which are physically much smaller than QSFPs. Furthermore, there are rugged non-pluggable “RJ” versions of these transceivers now available that are well suited for use in rugged embedded systems. Figure 3 compares the physical sizes of the QSFP/28, SFP+/28, and rugged RJ transceiver packages.

Figure 3. Transceiver Package Size Comparison.

25GbE May be the Clear Upgrade Path

If 10GbE (at 1.25 GB/s) is not fast enough for an application, and that application bandwidth needs are less than 2.5x the 10GbE level performance, 25GbE may be the clear choice. 25Gb Ethernet provides 3.125 GB/s of usable bandwidth (per port, in each direction), using the same basic cabling, connectors, and transceiver packaging, and with the same ruggedization options as 10GbE.

The optical physical layer options for both 10GbE and 25GbE are simple, consistent, and easily adaptable to rugged transceiver packaging. Nearly always, the interface consists of one transceiver lane, a single RX fiber, a single TX fiber, and a single laser operating on a single wavelength, and using common LC connectors and multi-mode optical fiber (MMF). The 25GbE single-lane architecture offers better reliability as compared to four-lane 40Gb/100Gb implementations. The most common 25GbE physical layer options are shown in Table 5. Note that these exactly match the 10GbE options that were shown earlier in Table 2.

Table 5. Common 25GbE Physical Layer Options Match 10GbE Options

25 GbE Optical Physical Layer	Lanes	Typical Transceiver	Typical Connector	Typical Cable	Maximum Distance
25GBase-SR	**1**	**SFP28 or rugged RJ**	**LC**	**MMF**	**100m***
25GBase-LR	1	SFP28 or rugged RJ	LC	SMF	10km
25GBase-ER	1	SFP28	LC	SMF	40km

*Maximum distance varies according to specific MMF cable type

Critical I/O's 25GbE Ethernet Upgrade Options

Critical I/O's 25GbE optical products are completely compatible with the 10GbE optical products. Same drivers, protocols, and software support. They leverage the same transceiver packaging, the same LC connectors, and the same MMF cables.

The XGE 25G series of 25 GbE NICs provides dual port 25 Gigabit Ethernet (25GbE) connectivity for embedded systems with the high-performance characteristics that are essential for data intensive real-time systems yet maintaining 100% interoperability and compatibility with all standard or enhanced Ethernet infrastructures. Dual port TCP and UDP bidirectional performance of up to 13 GB/s can be achieved using standard XGE 25G drivers. The standard sockets API, as well as a high performance UDP streaming API and a low latency RDMA API are supported.

Critical I/O 1, 10 and 25 GbE XMC products all feature:

- Driver commonality
- Industry leading performance and protocol offloads
- Support for all standard networking protocols
- Support for RDMA direct memory-to-memory data transfer
- Support for UDP direct accelerated UDP memory-to-memory operation
- Driver support for VxWorks, Linux, and Windows
- Rugged air-cooled and conduction cooled XMC hardware

The XGE 25G XMCs provide a balanced architecture which offers hardware acceleration for large data transfers with the flexibility of a programmable protocol processor to significantly improve network performance. These hardware offloads reduce the CPU cycles and burden required for networking, maximizing usable network bandwidth without sacrificing host CPU processing capability.

Figure 4. Critical I/O's Air Cooled and Rugged Conduction Cooled XMCs Support 1, 10, and 25 GbE