10Gb Ethernet Cabling Options

SFP+, DAC, and 10GBase-T

Technology Brief

10GBase-T is taking longer than expected, and cabling with the familiar RJ45 connector and Cat 6A cables is still a year or more away. But SFP+ connectors and DAC fill the gap nicely – especially in the data center.

Summary

For 10Gb Ethernet (10GE) switches and NICs, SFP+ connectors and the Direct Attach Cables (DAC) that attach to them have become widely available. SFP+ and DAC offer certain advantages over 10GBase-T: lower latency, lower power, lower cost and the technology is available and stable today. 10GBase-T offers a different set of advantages – familiarity, longer distance, and use of structured wiring and patch panels. SFP+ and 10GBase-T will offer data center designers an interesting choice in the 2011 timeframe.

The Search for a 10GE Physical Standard

Initial products supporting 10 gigabit Ethernet made use of fixed optics – in which the optical components are soldered directly to the circuit board. This was a valuable step in getting these pioneering products to market, but there are drawbacks to fixed optics. They do not allow a mix of different optical transceivers to handle different distance requirements, and each switch must carry the full cost and power consumption of optics on every port, whether all the ports are needed or not. On early products this was hardly the top issue, though, as the price of 10GE ports started at $10,000 per port or above.

Clearly the advent of pluggable optics made 10GE more practical, but the rapid succession of different optical standards prevented customers from gaining confidence that a standard would not become obsolete. Early pluggable optics standards included XenPak, X2, and XFP. Each standard included fewer electronics in the optical transceiver, moving functions to the circuit board where their cost and size could be absorbed into the latest switch and NIC (Network Interface Card) electronics. More importantly, each successive standard used less power and used less space on the faceplate.

The latest step in this progression is SFP+, and there is a strong indication that this is the final standard with which mass adoption can occur. The name SFP+ reveals that this standard is derived from the SFP (Small form-factor pluggable) standard, which is the most commonly used optical interface for Fibre Channel and 1 gigabit Ethernet. In fact, both SFP+ and SFP standards use the same physical dimensions and SFP transceivers are supported by SFP+ equipment.

SFP+ interfaces take approximately the same space on a switch front panel as the ubiquitous RJ45 connector, which is the most common connector for 1 gigabit Ethernet, Fast Ethernet and initial 10 megabit Ethernet products. The industry has always gravitated towards Ethernet switch products containing 48 ports, and with SFP+ interfaces, switches can be built with 48 ports of 10 gigabit Ethernet in a single rack-unit height. This was not possible with fixed optics or any of the other modular standards.

SFP+ modules, like their predecessors, come in different types to drive signals across fiber optic cables with different maximum distances. The most common, and lowest in cost, is 10GBase-SR (SR, or short reach), which can span 300 meters. Other types can reach as far as 80 kilometers.

Direct Attach Cable (DAC) (twinax)

In the past two years an unprecedented innovation occurred. Instead of connecting two SFP+ interfaces together using optics, a direct way of connecting without optical interfaces was created, using cables with ends built to behave as SFP+ optical modules. Instead of installing an optical transceiver at each end plus a length of fiber optic cable, a cable was invented with each end physically resembling an SFP+ transceiver, but with none of the expensive electronic components. A small component is required to identify cable type to the SFP+ host, but its cost is negligible. This innovation, called either Direct Attach Cable (DAC) or twinax, is a low cost solution for shorter distances, while keeping the high-density equivalent to RJ45 connectors.

10GBase-T

In the early days of Ethernet there were several PHY standards, including various forms of coaxial cable, but Ethernet only gained widespread adoption once it was run on twisted pair cable, and as point-to-point rather than an unstructured loop. It became practical to run an Ethernet cable to each desktop alongside the phone cable, and make all the connections in a wire closet. The connector used for twisted pair was the RJ45 connector. When Ethernet speeds increased to 100 Mbps, and then to Gigabit speeds, the same approach was maintained. While the cable requirements continued to evolve to maintain the 100 meter distance standard, the basic approach using wire closets, twisted pair wiring, and RJ45 connectors continued to be the de-facto standard. Today the PHY parts cost less than two dollars for an interface that can run at 10, 100, or 1,000 megabits per second.

When speeds reached 10 gigabits per second, though, the physics involved in pushing a signal across 100 meters of cable turned out to be much more difficult. Early PHY chips from several vendors show promise, but have not yet reached the cost, size, power or latency figures that will enable widespread adoption. Vendors continue to exploit advances in silicon manufacturing to improve 10GBase-T, and a few products have reached the market, but widespread adoption is still at least 2 years away.

10GBase-T requires the newer Cat 7 or Cat 6A to reach 100 meters, but can work on Cat 6, Cat 5E, or even Cat 5 cable at reduced distances.

Structured Wiring

Best practice for wiring Ethernet networks has evolved to what is called *structured wiring*. Each end point is terminated at a patch panel, and the patch panels are interconnected using short patch cables. For example, in a data center each server can be wired back to a central patch panel and switches are wired to adjacent patch panels. This allows connections between servers to be reconfigured easily. Patch panels have become widespread for 1GBase-T with RJ45 connectors, but are not practical using DACs, and are prohibitively expensive using fiber optics.

Standards vs. MSA

SFP+ interfaces are not governed by IEEE or other standards, but by MSA – Multi Source Agreement, an agreement among vendors on a specification to ensure interoperability. This is also true for XenPak, X2, XFP, and SFP, and has proven to be an effective way to provide multiple sources for products. 10GBase-T is entirely defined by an IEEE standard, **802.3an-2006**. This is the same standards body which defines most aspects of Ethernet. The optical interfaces for SFP+, such as 10GBase-SR, are defined by IEEE standards.

Comparing Cost

The electronic components required to drive 10GBase-T cost approximately $50 to $100 per port. This translates into a $200 to $400 per premium per port, or $400 to $800 per 10GBase-T cable. Using process technologies with 40 nanometer (nm) gate geometries, followed by 32 and 22 nm, the cost will be lowered over time, however 40 nm is just entering mass production for components of this type.

SFP+ components cost much less – currently $20 per port or lower – translating to $40/port or $80 per cable. The cables are also low in cost – approaching the costs of Cat 6A cables and decreasing in cost as the volume is rapidly grows. The list price of SFP+ based 10 gigabit Ethernet switches has dropped below $500 per port, and together with dropping NIC prices and DAC prices the overall solution cost is now quite affordable.

Comparing Latency

The 10GBase-T PHY standard uses block encoding to get data to cross the cable without errors. This block encoding requires a block of data to be read into the transmitter PHY, a mathematical function run on the data, and the encoded data is sent across the link. The reverse happens on the receiving side. The standard specifies 2.6 microseconds for the transmit–receive pair, and the size of the block of data suggests that latency cannot be improved below approximately 2.0 microseconds per link.

SFP+ uses simpler electronics without block encoding, and typical latency is approximately 300 nanoseconds per link.

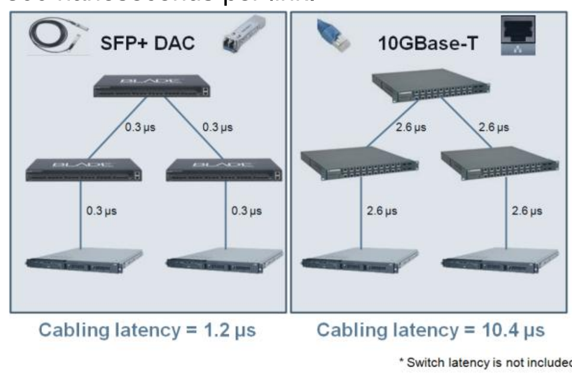

For each technology the latency of the physical media must be added – in fiber or wire the speed is roughly 5 nanoseconds per meter.

Comparing Distance

10GBase-T can reach 100 meters using the latest Cat 6A or Cat 7 cables. The standard has been engineered to allow for patch panels and jumper cables as well.

The MSA which defines direct attach cable specifies a maximum distance of 8.5 meters (about 28 feet). There are no patch panels used with DACs.

SFP+ is designed for optics, however. For longer runs, fiber optic transceivers and fiber cables can be intermixed with DACs, allowing distances of 300 meters at a reasonable cost and up to 80 kilometers at a higher cost.

A new generation of active DACs will allow longer distances with only a small increase in price. These active cables contain electronic amplification within the cable ends. Distances up to 25 meters have been proposed by cable vendors.

Comparing Power and Heat

Power is a very finite commodity in many environments, especially data centers. It is important to note that for every watt of power consumed, typically two additional watts of power are needed for cooling, to remove the heat generated. 10GBase-T PHY components today require 4 to 6 watts per port at each end. SFP+ PHY electronics, while driving DACs, consume about 1.5 watt per port at each end.

Usage in High Performance Computing

Architects of high performance computing (HPC) clusters have been early adopters of 10 gigabit Ethernet, taking advantage of both the increased bandwidth and reduced latency compared to 1 gigabit Ethernet. However, the latency introduced by the 10GBase-T PHY is a poor match for HPC's requirements. To connect two servers in a typical HPC installation there is one cable from server to top of rack, another from top of rack to core, then an additional two hops to reach the other servers, for a total of 4 hops, or 10.4 microseconds of latency. This is a significant performance penalty for many HPC applications, and exceeds the latency caused by the new generation of sub-microsecond switches. Four hops using DAC will create latency of approximately 1.2 microseconds – a better fit for HPC.

Usage in the Data Center

Data center best practices for wiring are undergoing a dramatic change. The priority in the past had been to create a patch panel that could quickly and easily connect any server to another by moving a jumper cable. Several trends are changing design priorities. Using virtualization, it is easy to move virtual machines and applications across the data center rather than changing the network. Architects are also introducing much larger layer 2 switching domains, now that there is less risk of a misbehaving server taking down an entire layer 2 domain. Data center designers are thinking more in terms of groups of servers and switches as single management entities. Bladed servers are internally networked together, establishing the first tier of the network directly attached to the servers. Finally, new and more capable switches specifically designed for the "top of rack" are making it easy to provide for changing topologies within the switches.

As a result, many data center designers are doing away with the concept of patch panels, and instead wiring directly from the servers to the top of rack switches. This is a perfect fit for DACs, as the distance is typically no more than a few meters. Uplinks from the top of rack switches to the network core can be a mix of SR optics and fiber cable and DAC (where distance permits). Since there are far fewer uplinks than server connections, this produces a cost effective solution. The advantages in cost, power and heat dissipation are multiplied by the high number of server connections.

Usage in the Wiring Closet

Wire closets contain switches and patch panels which provide connectivity to large populations of desktop computers. One of the key criteria for wire closet designs is the need to efficiently support adds, moves and changes. Using clearly labeled patch panels to make it easy to change user connectivity has become standard practice, so staying with RJ45 connections to desktops makes sense. And cable installers know how to field terminate cables and install RJ45 connectors. At this point, few desktop users need even 1 gigabit so it will be some time before 10 gigabit to the desktop becomes economically justified. When that happens, 10GBase-T is the right technology for the wire closet.

Another attractive characteristic of 10GBase-T is the reverse compatibility that it offers. Using the RJ45 connector, the 10GBase-T port can connect to 1 gigabit and even 100 megabit interfaces. DACs can't be used to connect to legacy equipment. However, SFP 1 gigabit transceivers can be inserted into SFP+ ports and will work at 1 gigabit, connecting to legacy ports over fiber cable. SFP transceivers supporting 1GBase-T can also be inserted to connect to legacy ports at slower speeds. Either of these solutions will add cost, due to the additional transceivers.

Usage in the Campus LAN

Campus LANs are used to tie together multiple buildings, typically spanning longer distances. In this case neither 10GBase-T nor DAC is a good choice, and most network designers use fiber optics to provide long distance capability.

Conclusions

Comparing 10GBase-T to SFP+ with DAC:

10GBase-T advantages:

- Longer distance – 100 meters vs. 8.5
- Familiar RJ45 connectors and Cat 5/6/7 cables
- Use of patch panels and structured wiring
- Backward compatibility to 1 gigabit Ethernet or 100 megabit Ethernet

SFP+ with DAC advantages:

- Significantly lower overall cost, when you include switch, NIC and cable
- Lower latency – 300 ns per hop vs. 2.6 us per hop
- Lower power and lower heat
- Freely intermix fiber and DAC to meet distance requirements

For data centers, the advantages of SFP+ with DAC are a very good match for today's requirements and emerging trends. That's why SFP+ with DAC is being adopted rapidly as best practice for new data centers.

For wiring closets, 10GBase-T will be the obvious choice once the demand for bandwidth becomes more acute and once the price and power for 10GBase-T technology comes down.

Both technologies should find an important place in the future of network design and best practices.

	SFP+	10GBaseT
Distance	8.5m DAC to 80km fiber (25m DACs in future)	100m
Latency	0.3 µs	2.6 µs
Power	1.5W/port	4-6W/port
Cost	$40/port	$400/port
Use when:	• Lower solution cost • Lower latency • Lower power • Flexibility to mix DAC for racks and fiber for longer distances • Flexibility to mix Gigabit & 10G	• Familiar RJ45 connectors • 100m distance with copper • Structured wiring • Flexibility to mix 100Mb, Gigabit and 10G

Definition of terms:

10GBase-T – IEEE standard 802.3an-2006 for operation at 10 gigabits per second over Cat 5/6/7 twisted pair cable at distances up to 100 meters

10GBase-SR – IEEE standard 802.3ae-2002 for operation at 10 gigabits per second over fiber optic cable at distances up to 300 meters

DAC – Direct Attach Cable (also known as twinax). Some vendors incorrectly call these 10GBase-CU or 10GBase-CR cables, but there is no IEEE standard at this time.

SFP – Small form-factor pluggable (also known as mini GBIC), for Fibre Channel and 1 gigabit Ethernet

SFP+ – Small form-factor pluggable, for 10 Gigabit Ethernet

XFP – 10 Gigabit, small form-factor, pluggable – larger than SFP+

X2 and XenPak – 10 Gigabit pluggable, older and larger than XFP or SFP+

MSA – Multi-Source Agreement

NIC – Network Interface Card

PHY – physical layer standard, or the circuitry which drives it

RJ45 – connector widely used for Ethernet

Cat 5 / Cat 5E / Cat 6 / Cat 6A / Cat 7 – different grades of twisted pair Ethernet cable

MKT091016