To 400Gb/s and Beyond: High-speed Interconnects Evolve to Meet the Needs of Data Center Communications

by **Contributed Article** on March 28, 2017

by Jim David – Portfolio Director; Amphenol High-Speed Interconnects

High-speed interconnects adapt to data centers' ever-changing signal, bandwidth, and density requirements.

The emergence of large scale and mega-data centers has driven the need for not only increased data management capability, but also increased the complexity, design, and configuration of these enormous facilities. Meeting data centers' increasing demand for higher signal speeds, higher bandwidth, and higher density interconnections, high-speed connector and cable assembly suppliers are continually developing technologies to accommodate today's data center requirements. As early data centers expanded and new generations of server racks were designed, there was an identified need to have common connector interfaces that could be used to link the various network components (e.g., switches, routers, and storage devices) to each other for applications within a server rack, rack to rack within a row or aisle of servers, or between multiple rows of servers within the data center. Standardization of these different interfaces and form factors, while readily apparent, was not easily accomplished due to the multiplicity of possible interfaces, suppliers, and hardware manufacturers.

Starting in the 1990s, the Small Form Factor (SFF) committee was established as one of the first organizations to formulate and release industry standards that define the connector and cable interfaces as well as give definition of interoperability specifications and electrical signaling protocols for network hardware and components. This work enabled different connector manufacturers and network component suppliers to design and produce devices that are interoperable with one another under what is known as a multi-source agreement (MSA). There are multiple MSA agreements in the market today, and some examples follow below.

One of the first high speed interfaces was dubbed the small form-factor pluggable (SFP) device. This standard was developed to enable both copper and fiber optic networking cabling and connections between what were typically switch and server components. They were usually identified by their rated data transmission speed (e.g., 1Gb/s SFP). As data transmission speeds increased to 10Gb/s per lane, the SFP standard was upgraded to SFP+ to not only accommodate the higher speed but also include the use of an EEPROM, which was added to allow for automatic identification of cable configuration, data transmission speed, and other pertinent information once the cable was plugged into the hardware. However, even with this upgrade, the dominant high speed standard remained primarily a single-lane interface.

QSFP reaches the 100Gb/s milestone

With the success and widespread adoption of the SFP+ interface, the need for still higher bandwidth quickly emerged and led to the development of the Quad SFP (QSFP) interface. This effort resulted in a four-lane interface with each lane running at 10Gb/s or 40Gb/s aggregate bandwidth. Fast forward several more years, and with the increasing need for greater bandwidth capability going unabated came the emergence of a 25Gb/s lane signal speed capabilities. Extended to the QSFP interface, this now gave hardware designers the capability of adding IO ports that had 100G aggregate bandwidth capability.

These leaps in bandwidth, density, and capability have not been without their challenges for connector and cable suppliers. High-performance signal integrity, PCB footprint compatibility, and EMI shielding, as well as the capability to have both copper and optical based solutions, had to be maintained while accommodating increased power and heat dissipation demands.

High-speed signal integrity performance was improved through the strategic used of advanced plastics and metal materials within the connector design, as well as highly engineered PCB, wire stripping, wire management, and wire termination techniques on the cable assembly side.

Enhanced floating heat sink designs were also incorporated into QSFP connectors to improve heat dissipation performance.

OSFP boasts speed and density

Building on the successful implementation of the QSFP interface (100Gb/s per port, universal adoption across a wide range of applications), the next step in the evolution of high-speed interconnection was the octal (8x) small form factor pluggable (OSFP) interface. OSFP was developed as an eight-lane interface x 25Gb/s per channel, thus doubling the aggregate bandwidth capability to 200Gb/s per port. The OSFP standard is supported by an MSA development group consisting of 49 member companies. This connector design features a total of 60 contacts per port on a 0.6mm pitch, with 16 high-speed differential pairs and 10 power/control contacts arranged in two rows. It allows for copper-based cabling solution and delivers significant heat dissipation capabilities up to 15W. This allows for the use of not only short reach optical applications (generally 100m or less), but also long-reach and long-haul applications as well (kilometers in length). This higher heat dissipative capability allows designers the luxury of not having to limit the optical cable or transceiver that is plugged into the port or having to worry whether the port has the necessary heat dissipation properties.

The OSFP system is also designed to accommodate the pulse amplitude modulation (PAM-4) communications protocol that essentially doubles the connection bandwidth through multi-level signaling. While the PAM-4 protocol requires more complex encoding hardware on either end of the data connection, its data bandwidth boost results in an aggregate data transmission of 400Gb/s per port by effectively doubling each of the eight 25Gb/s lanes to 50Gb/s per lane.

The initial OSFP interface connectors and cables are expected in mid-2017 with single (1 x 1) and eight port (2 x 4) configurations. Following the deployment of those standard configurations, other custom configurations to meet application specific needs are likely to follow.

DD-QSFP delivers increased bandwidth with backward compatibility

A different development path to the 400Gb/s per port capability involves increasing the density of the standard 100Gb/s QSFP module by a factor of two. Designated the double-density (or DD-QSFP) interface, this connector design features a total of 76 contacts on a 0.8mm pitch with 16 high-speed differential pairs and 13 power/control contacts arranged in two rows. Akin to the OSFP interface, the DD-QSFP possesses the ability to accommodate eight lanes of 25Gb/s data transmission (200Gb/s) or, using PAM-4 modulation, eight lanes of 50Gb/s (400Gb/s) in a single DD-QSFP port.

Advantages of the DD-QSFP interface include double the bandwidth for the same given space and backwards pluggable compatibility with existing QSFP hardware. The heat dissipative capability of the DD-QSFP interface is not as high as that of the OSFP system, which may limit its use in some optical transmission applications.

DD-QSFP interface connection devices are scheduled for implementation in 2017- 2018, and will feature single-port, ganged, and stacked-board connector configurations. Follow-on developments are likely to include surface mount versions later in 2017 or 2018.

RCx: Making the intra-rack connections within data center equipment

The RCx interface is an emerging MSA between several manufacturers that offers a high-density, low-cost, passive copper connector and cabling system that is specifically designed to provide short-run (up to 3m) intra-rack connectivity for 25Gb/s, 50Gb/s, and 100Gb/s Ethernet applications.

The product design goals include offering a more cost effective, simple, and flexible interconnect system. The RCx design eliminates the need for some components (e.g., EEPROMs, heat sinks, and light pipes), and does not include any provisions for optical transmission. The RCx interface is also intended to streamline the electrical interconnect design for switches, servers, and adapter cards. Modular and flexible by design, the system allows for multiple cable assembly configurations (single, dual, and quad) as well as multiple splitter/breakout cable configurations that can accommodate a wide range of interconnect applications. Each lane supports 25Gb/s signal transmission.

The RCx1 (single-lane) cable design features two differential pairs, four surrounding ground contacts, and two ID pin contacts on a 0.8mm pitch. The RCx2 (dual-lane) and RCx4 (quad-lane) configurations are multiples of the RCx1 configuration. A mechanical keying slot in the center of the bottom of the cable connector shield prevents mismating of the connector orientation, as well as improper lane-shifting errors.

A wide variety of RCx receptacle configurations are available, including 1 x 2 and 1 x 8 SMT versions, as well as 2 x 8 press-fit versions. Depending on the specific application requirements, stacked 2 x 8 and ganged 1 x 8 receptacles can also be combined in several configurations to create multiple-row connection options. RCx connectors can also be used in belly-to-belly applications. The RCx interface standard was developed to operate at cable lengths up to 3m without forward error correction (FEC). Longer cables can be deployed if the appropriate FEC protocol is enabled.

When compared to legacy 100G QSFP and SFP28 interfaces typically used for intra-rack connections, the RCx interface requires less board real estate behind the faceplate and delivers higher linear signal density along the PCB edge. Compared to QSFP28 breakout cables, the RCx4 breakout cables and connector system require much less PCB space versus conventional SFP+. The RCx interface connectors and cable assemblies are currently available in sample and pre-production volumes, with full production capacities anticipated in the third quarter of 2017.

Jim David – Portfolio Director; Amphenol High Speed Interconnects