molex

NEXT-GEN DATA CENTER INFRASTRUCTURE RAISES THERMAL MANAGEMENT STAKES

BY JOE DAMBACH | MOLEX INCORPORATED

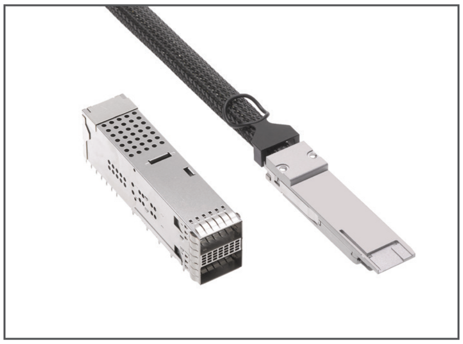

A QSFP-DD cage can receive two optical or electrical cable assemblies.

An array of high-speed, high-density solutions revolving around various small form-factor pluggable modules, such as SFP, SFP+, QSFP, QSFP+ and zQSFP, gained considerable traction in datacenter switch and networking devices. Thermal management strategies are critical as the industry is readying for the launch of next-generation copper and optical QSFP-DD transceivers.

The name 'small' form-factor pluggable I/O is deceptive. These small yet powerful connectors have played a major role in high-density networks handling massive volumes of data generated by mobile phones and devices, streaming content, digital automation and industrial sensors, artificial intelligence, predictive analytics, and a host of other data-driven technologies.

By 2020, analysts predict there may be as many as 200 billion devices in the IoT (Source: IDC, Intel, United Nations).With datacenters undergoing extraordinarily rapid growth, rising demand and data volume are pushing density and power limits. Chips, switches, connectors, cables and optical module technologies are at the core of networks—and impact every branch right down to individuals and corporate enterprises, where delays or downtime can quickly rack up millions of dollars in lost revenue and opportunity.

Bottlenecks that exist now will only get worse without the critical infrastructure needed to support future data rates. Investments in storage, servers and switches will be required to keep up and make that data accessible. Datacenters that do nothing to prepare for future high-speed and high-density data will increasingly face latency problems and ultimately fall behind competitors.

Seamless, reliable connectivity for network, server and storage devices is essential to ensure data flows quickly, efficiently and securely. Datacenters must be able to support faster processing, more bandwidth, and increased density. Making that happen requires orchestration and integrated I/O connectivity to coordinate interoperability for IT equipment from various suppliers. The larger or more complex the workload, the more important it is that components are well balanced and specified so as not to inadvertently create choke points that impair performance.

A portfolio of pluggable I/O products has evolved, all revolving around the various pluggable I/O module formats, with each variant offering a range of bandwidth, distance, and either copper or optical connections. Each specific link can be economically defined using either optical or copper interfaces. Pluggable I/O solutions are designed to support the fastest data rates across the spectrum of distances common in datacenters and telecommunications, with lower latency and insertion loss, and excellent signal integrity, electro-magnetic interference (EMI) protection and thermal management.

SCALABLE AND INTEROPERABLE I/O UPGRADE PATH

High-speed, high-density pluggable I/O solutions provide a highly scalable upgrade path for high-density applications. The industry has come a long way since the SFP form factor introduced 1 Gbps data rates. QSFP optical transceivers added new configuration options and pushed density and speed up to 100 Gbps. Specifically designed for high-density data communications, widely used QSFP+ hot-pluggable transceivers integrate four transmit and four receive channels for greater port density and cost savings over buying higher quantities of traditional SFP+ products. The zQSFP+ system supports high-density requirements, with data rates ranging up to 28 Gbps over four lanes, making it a popular choice and valuable asset in datacenter high-performance computing, switches, routers and storage.

Bandwidth requirements to meet tremendous growth in wireless devices have been a catalyst for higher density designs in server farms. A highly integrated QSFP system combines the effective use of space, power and port density, and will typically comprise connectors, EMI shielding cages, both copper and optical cable assemblies, active optical cables (AOCs), optical loopbacks, and a host connector. Supporting next-generation 100 Gbps Ethernet and 100 Gbps InfiniBand Enhanced Data Rate applications, the zQSFP+ interconnect solution transmits up to 25 Gbps per-serial lane data rates with excellent signal integrity, EMI protection and thermal cooling.

Originally published in EDN.com (November 2017)

molex

NEXT-GENERATION QSFP-DD FORM FACTOR QUADRUPLES SPEEDS

Networks rely on switching technologies to handle and manage data traffic. As data usage rises, the switches at the heart of networks keep data flowing smoothly. High-density solutions relieve some of the pressure on core switches. Connectors must not become bottlenecks. There are already silicon chips on the market that support 256 differential lanes. The missing link is a connector form factor providing adequate density to support this lane count simply in a 1RU box, while managing thermal and signal integrity.

The industry is in the process of developing next-generation QSFP solutions that keep pace with recent developments in switching technologies. In 2016, the QSFP-DD MSA was formed to address the technical challenges of achieving a double-density interface and to ensure mechanical, electrical, thermal and signal-integrity interoperability for module components produced by different manufacturers.

As one of the founder-promoters of the QSFP-DD MSA, Molex collaborated with 52 other companies in addressing the industry need for high-density, high-speed networking solutions. The group recently released specifications for the new form factor that overcome previous technical challenges of specifying a QSFP28 compatible double-density interface.

The Quad Small Form Factor Pluggable Double Density (QSFP-DD) specifications define a module, a stacked integrated cage/connector system and a surface mount cage/connector system. The new form factor expands the standard QSFP four-lane interface by adding a row of contacts providing for an eight-lane electrical interface, each operating up to 25 Gbps with Non-Return-to-Zero (NRZ) modulation or 50 Gbps with Pulse Amplitude Modulation (PAM4). This adaptation allows the QSFP-DD to address solutions up to 200 Gbps or 400 Gbps aggregate per QSFP-DD port. A single switch slot can support up to 36 QSFP-DD modules—providing up to 14.4 Tbps aggregate capacity.

EFFECTIVE THERMAL COOLING AND MANAGEMENT PROTECTS INVESTMENTS

From SFP+ (1 x 10 Gbps) to QSFP+ (4 x 10 Gbps) to zQSFP+ (4 x 25 Gbps)—and now QSFP-DD (8 x 25 Gbps NRZ/50 Gbps PAM4)—as transceiver speeds increase, the energy it takes to drive signals increases and creates more heat. Thermal cooling and management are critical for connector module and overall network power consumption, energy efficiency, performance and longevity.

Advances in heat sink technologies are enabling highly efficient, reliable and resilient thermal management strategies to support both higher density copper and optical connectivity. QSFP optical modules are designed to operate at a maximum temperature of 70 degrees Celsius. Above that point, the performance of optics starts to degrade and the lifetime of the optics diminishes. While the price point has dropped, optical fiber components remain a relatively high cost investment to put at risk due to temperature extremes. An I/O module with fewer than four lanes can typically be cooled using airflow intended to cool other components. The QSFP form factor with four signal lanes requires a combination of airflow and heat sinks. Originally, the zQSFP form factor had no heat sinks and relied on airflow to cool the module. As data density increased, various heat sinks were added to the center and top of the module, and to the cage structure, to improve airflow and cooling.

As zQSFP+ speeds exceed 100 Gbps and power loads over 4W, there are added challenges to managing system and optical module temperatures, so as not to exceed a maximum of 70 degrees Celsius. The overall design must manipulate the airflow through the cage, with heat sinks to transfer the energy from the module to those areas of airflow. Molex has developed and demonstrated innovative thru-flow and internal riding heat sink (IRHS) technology for thermal management on a stacked zQSFP+ module operating at up to 5W or higher. Focusing primarily on enterprise applications with maximum ambient temperatures of about 45 degrees Celsius maximum, Molex testing compares a thru-flow cage optimized for management of the energy generated in the module. The demonstration emulated a 5W optical module by using a wind tunnel containing two 2 x 1 cages—a flow-enhanced cage and a standard cage—with two zQSFP+ modules on each side—and a power supply to drive a fan at a predetermined rate through the test area.

The relative performance of the standard zQSFP+ cage and the flow-enhanced cage is compared via continuous monitoring of the temperature within each of the four modules. Using the enhanced cage design, heat sink and energy transfer strategies, the demonstration achieves a significant improvement of a nine degree Celsius cooling factor using enhanced cages versus standard cages. The QSFP-DD module drives eight lanes in a space only slightly deeper than the QSFP. That means eight lasers creating heat coming out of the module that needs to be cooled and dissipated. By effectively doubling the thermal energy, a 3.5W module becomes a 7W module—a 5W module becomes a 10W module, and so on.

Managing thermal performance at these extremely high heats will be critical moving forward. Advanced thermal management techniques used in the design of the module and cage enable the QSFP-DD to support power levels of at least 7W, with a target range up to 10W. Molex and other MSA partners and contributing member companies are in the process of tooling QSFP-DD products, with cages and modules available for thermal testing in customer environments. Additional work is already underway to develop solutions for cooling the QSFP-DD module operating at up to 12W or higher.

molex

EXTENDING THE LIFE OF EXISTING PLATFORMS

Demand for higher density and better power and thermal management will continue driving development of improved I/O modules. Adoption rates for new network technologies are generally incremental. Industry and multi-vendor investments help assure interoperability of each generation of pluggable I/Os.

Backward compatibility adds a layer of future proofing by extending the life of existing platforms. For example, the QSFP and zQSFP+ share an identical mating interface. While an entirely new interface, the QSFP-DD builds on these prior technologies. Double density refers to the doubling of the number of high-speed electrical interfaces that the QSFP-DD module supports compared with a standard QSFP28 module. In order to accommodate an extra row of contacts, the mechanical interface for the QSFP-DD on the host board is slightly deeper than the standard QSFP28 interface.

Each QSFP-DD port provides 76 electrical contacts to support next-generation switching chips for higher density applications as datacenters expand capacity. The QSFP-DD is ready for the new PAM4 electrical modulation format that supports 50 Gbps, which provides another doubling of speed, resulting in a quadrupling of speed compared to the QSFP28 module. A majority of initial use cases for QSFP-DD modules will support NRZ modulation data rates. PAM4 permits more information in a given bandwidth, but demands more processing power to encode and decode data, which can contribute to latency degradation.Systems designed with QSFP-DD modules will be backward compatible with existing QSFP form factors to provide maximum flexibility for end users, network platform designers and integrators. The port densities are identical, except the QSFP-DD specifications accommodate four more lanes and double the aggregate bandwidth of the QSFP28 in a given panel space. QSFP28 modules can be inserted into four of the eight electrical lanes on the QSFP-DD to allow gradual and incremental upgrade of cables and modules. The high-density QSFP-DD optical transceiver stands to bring exceptional value to manufacturers by allowing them to make more competitive network server, switch and storage products that support ever-increasing data traffic volumes and more complex datacenter and telecommunications networks.

Order No. 987651-6902 USA/0k/GF/2017.11