



25 Gigabit Ethernet for Servers: What Does it Mean?

Introduction

Until recently, the Ethernet speed upgrade path was clearly defined as 10G → 40G → 100G. However, new developments indicate the latest path for *server connections* will be 10G → 25G → 100G with potential for 10G → 25G → 50G → 100G. These developments are being driven by companies such as Cisco, Microsoft, and Intel because 25G provides a more efficient use of hardware and a more logical upgrade path to 100G. 25G Ethernet is also in the process of becoming a standard and has passed the first hurdle in the IEEE standards body with a successful Call for Interest (CFI) in July, 2014. The industry expects that 25G hardware will be available as early as 2015 with the standard finalized by 2016.

40 Gigabit to the Server Will be Limited

While it is impossible to see the future, the current projections by switch and server analysts show that 25G will limit the deployment of QSFP+ 40G ports at the server level. While switch-to-switch connections will stay at 40G, expect that the dominant next generation server speed will be 25G. As a result, Panduit recommends that data centers be planned with 25G at the server. Subsequent upgrade paths will likely include an optional step at 50G and the following step at 100G.

QSFP28 to SFP28 Breakout Cables for Top of Rack with Up to Five Meters of Reach

The IEEE CFI is focused on the Direct Attach Copper (DAC) which includes QSFP28 (based on the IEEE 802.3bj 100G standard) and SFP28 (will be defined by the new IEEE study group). QSFP28 refers to the 100G DAC cable using the QSFP+ form factor, and SFP28 refers to the 25G DAC cable using the SFP+ form factor. While the form factor remains the same, the new standard will require improved cables and connectors capable of handling this additional bandwidth. Existing QSFP+ and SFP+ cable assemblies will not be compatible with QSFP28 and SFP28 ports.

The goal of the IEEE CFI is to use 25G DAC as the low-cost copper server connection for Top of Rack (ToR) switches. Reach is expected to be about the same as existing passive DAC cables or about five meters. There is speculation that some equipment may only run at three meters, but it is still too early in the process to validate or verify that possibility.



Figure 1: QSFP28 to four SFP28 breakout cable

Currently the cabling application is one QSFP28 to four SFP28 cables. The target application for the CFI is 100G switches connecting to four 25 Gigabit in cabinet or adjacent cabinet servers using these breakout cables. There is also expected to be SFP28 to SFP28 cabling, assuming equipment vendors develop a ToR 25G switch. A 25G fiber option is also being developed based on duplex LC that would mate to the SFP28 ports.



40GBASE-T will Include 25GBASE-T and Use the Same Category 8 Channel Already Developed

Because the industry-wide view is that 40G to the server will be limited in scope, the development of the 40GBASE-T standard (IEEE 802.3bq) will expand to include 25GBASE-T to support 25G to the server over twisted pair. The plan is to have a new Call for Interest in the IEEE to officially begin work on 25GBASE-T in November, 2014. It will then roll into 40GBASE-T to become the 25G/40GBASE-T standard. The combined 25G/40GBASE-T standard is expected to reuse much of the work done for 40GBASE-T, such as using the same shielded, 2 connector, 30 meter Category 8 channel. This is expected to help keep to the original timeline of being released in early 2016. Category 7/7A cabling and connectors are not expected to work, due to the 2GHz bandwidth requirements of this standard (Cat 7 is 600MHz, Cat 7A is 1GHz).




Figure 2: Category 8 Channel Maximum Length Configuration

Retaining the 30 meter channel used for 40GBASE-T development limits the complexity of the 25GBASE-T equipment to minimize time to market. It also limits the power per port to under 3W on first generation equipment, and provides a true upgrade path where the same 30 meter channel can be used to support 25GBASE-T, 40GBASE-T, and 50GBASE-T.

25GBASE-T is Good News

The standards for 25GBASE-T and 25G DAC are both scheduled to be released at about the same time in 2016. With a longer 30 meter reach and a cost effective price, 25GBASE-T presents a strong alternative to the 3-5 meter 25G DAC. It will allow customers a more cost-effective alternative for 25G for both Top of Rack (ToR) server connections using point-to-point patch cords, and enable End of Row (EoR) or Middle of Row (MoR) by using the 30 meter structured cabling. This is all achieved with the commonly used RJ45 plug and connector, which offer the familiar advantages of BASE-T:

- Backwards compatibility
- Auto-negotiation
- Cost effective

Conclusion

The dominant next-generation server connection speed is going to be 25G at the server. This may limit 40G to the server, as the migration path from 25G would be to 50G or 100G. However, 40G between switches is expected to remain and will not be affected by this development. A current IEEE study group, driven by the needs of high density cloud computing, is working on 25G to the server using a breakout 100G QSFP28 to four 25G SFP28 copper cabling with the potential for 25G optical. An upcoming IEEE study group will begin working on 25GBASE-T over twisted pair to provide a cost competitive longer reach option for mainstream customers wanting to use EoR and MoR architectures.

Panduit recommends that data centers start planning for the 25G to the server, and will be glad to offer our expertise in choosing the physical infrastructure that best meets your needs.