ıılııılıı
CISCO
The bridge to possible

The Key Role of 50G in Today's NG Networks

Contents

Introduction

50G SFP transceivers deliver double the data rate of 25G SFP transceivers in the same form factor.

The soaring popularity of data-intensive applications in Next-Generation (NG) networks, like the Internet of Things, streaming video, and cloud computing, has caused bandwidth demand to skyrocket. Network operators need to upgrade capacity from an installed base of 25G SFP ports (SFP28; see table 1) while minimizing cost. Based upon the relatively recent speed boost from 10G to 25G in the SFP transceiver family, 50G would be the next SFP data rate increase.

SFP56 modules take advantage of mature 50G technology already widely deployed in 400G transceivers, which use eight lanes of 50G in a QSFP-DD package (QSFP56-DD). SFP56 technology is available today as an approach for network operators to double the data rates of their existing SFP28 links.

Table 1.   Naming conventions and data rates for pluggable transceivers

	Data rate	ASIC support available?	Number of electrical lanes
SFP	10 G	Yes	1
QSFP+	40 G	Yes	4
QSFP28	100 G	Yes	4
SFP28	25 G	Yes	1
QSFP56-DD	400 G	Yes	8
QSFP56	200 G	Yes	4
SFP56	50 G	Yes	1

Another important consideration is that switch manufacturers today can build in SFP56 ports that are compatible with 1G, 10G, 25G, and 50G transceivers. This development provides great flexibility. A single switch, router, or server can leverage both legacy lower-speed connections and new high-speed connections, just by plugging in the appropriate transceiver.

SFP56 transceivers deliver double the data rate of SFP28 transceivers in the same form factor and therefore dramatically reduces the cost per bit. The SFP56 form factor minimizes power consumption and maximizes faceplate density. The transceivers directly interface with the latest generation of switch ASICs that are equipped with a 50G serializer/deserializer (SerDes). Best of all, most SFP56 ports are also backward compatible with existing SFP+ and SFP28 transceivers.

Cisco is building out a comprehensive portfolio of SFP56 products, including SMF, MMF, DAC, and industrial-temp transceivers, known as our SFP-50G product line.

Figure 1.
Cisco SFP-50G Transceiver

Getting to 50G

Bandwidth demand has led the industry to concentrate on reaching the highest possible data rates as quickly as possible. This focus, in turn, has driven the evolution of switches and optical modules. From the user point of view, it seems logical that data rates should grow monotonically. In reality, progression to higher speeds starts with multilane transceivers based on the incumbent peak data rate and then back fills with single-lane transceivers once the technology is ready. When the 25G SerDes was released, for example, the fastest transceivers on the market were QSFP+ (4x10G). Rather than leveraging the new 25G SerDes and ASICs to produce SFP28 switch ports and the modules to go with them, suppliers immediately focused on developing QSFP28 (4x25G) transceivers. The SFP28 market only matured later.
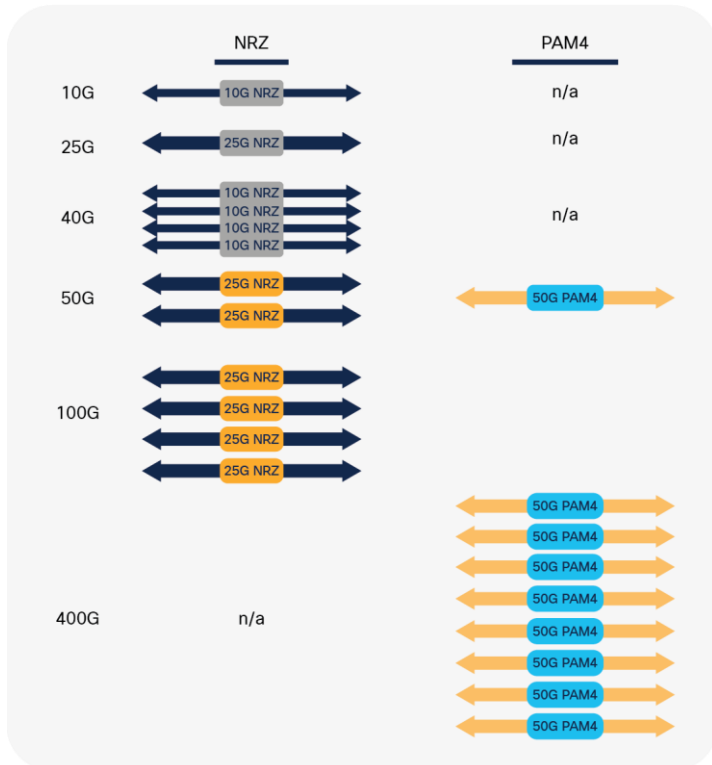

Figure 2.
To increase data rates between releases of new switch ASICs, industry turns to multilane transceivers. Single-lane versions are more compact and less expensive and consume less power.

Today, the same dynamic is in play with the 50G SerDes: the initial focus was on QSFP56-DD and QSFP56, but now the SFP56 market is expanding.
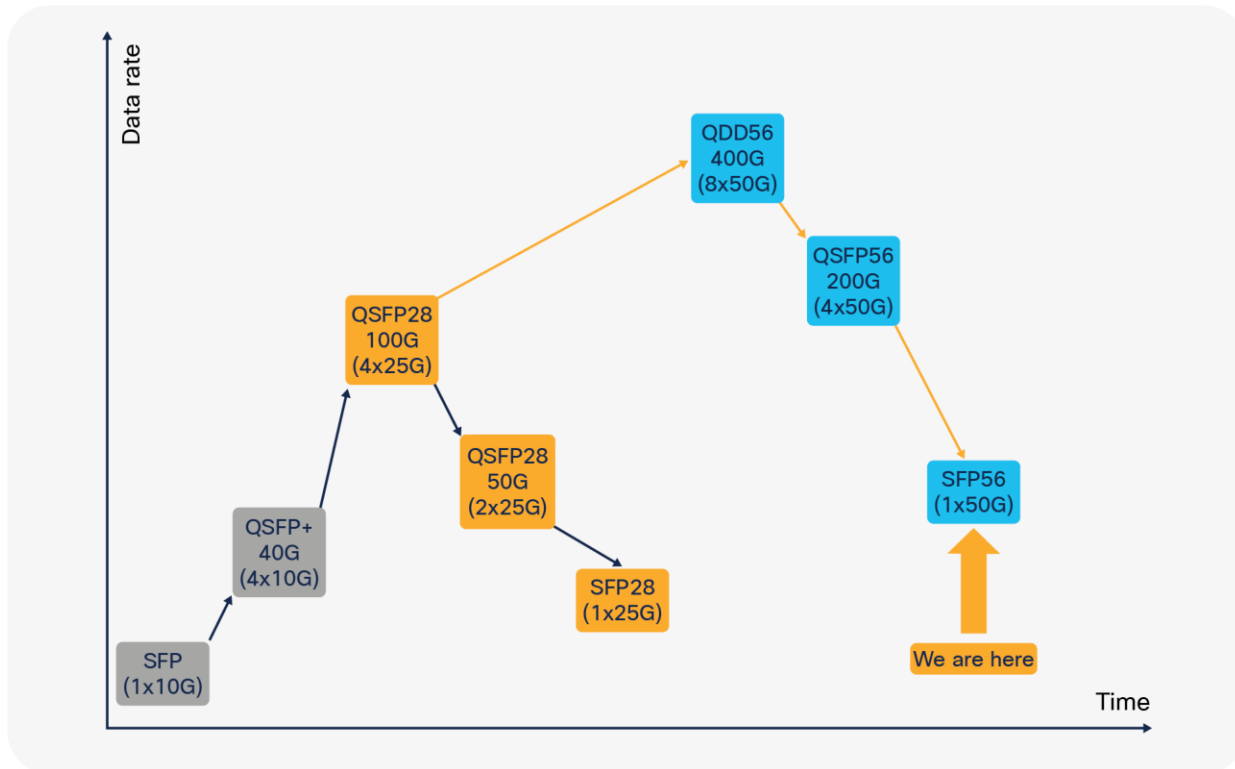

Figure 3.
Data rate increases depend on the availability of the Serialization/Deserialization (SerDes) ASICs. Historically, the higher-density interfaces always come before the single-lane interfaces.

What is SFP56?

SFP56 is a 50G small-form-factor pluggable Ethernet transceiver. The SFP56 has a 256b/257b coding overhead on top of a 50-Gb/s payload, which increases the physical line rate of the signal to 26.5625 Gbaud.

The SFP56 electrical interface is based upon single-lane 50G four-level Pulse-Amplitude Modulation (PAM4). Instead of the binary signal found in NRZ modulation, PAM4 encodes pairs of binary 1s and 0s into a single symbol with four possible amplitude levels (see figure 4). In other words, PAM4 signaling makes it possible to transmit two bits per symbol. Decoding simply involves reversing the process.
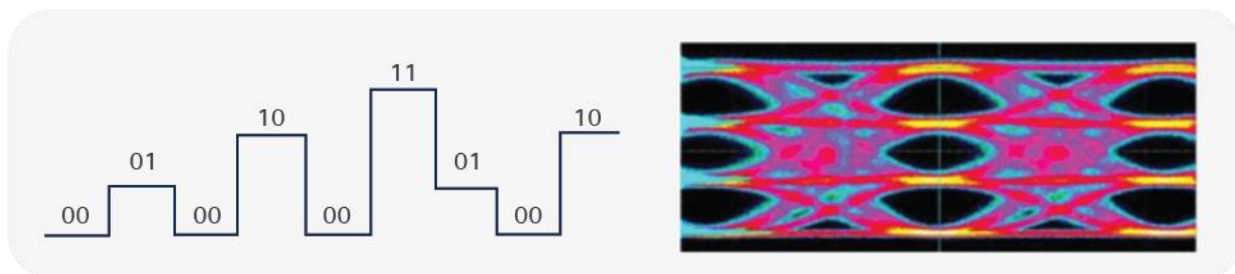

Figure 4.
Four-level Pulse-Amplitude Modulation (PAM4) signaling encodes two bits per symbol, for a data rate that is double the line rate.

Because PAM4 transmits two bits in each symbol slot, it cuts the signal bandwidth in half. The modulation format makes it possible to use a 25-GHz rate to transmit data at 50 Gb/s. In the case of SFP56, our physical line rate of 26.5625 Gbaud can be used to achieve an actual data rate of 51.125 Gb/s. The SFP56 electrical interface is designed for a signaling rate of up to 56 Gb/s to accommodate OTN overhead.

This PAM4 electrical modulation is implemented in the switch with the 50G SerDes, which is used by many new switch ASICs, including the Cisco Silicon OneTM family (see figure 5).
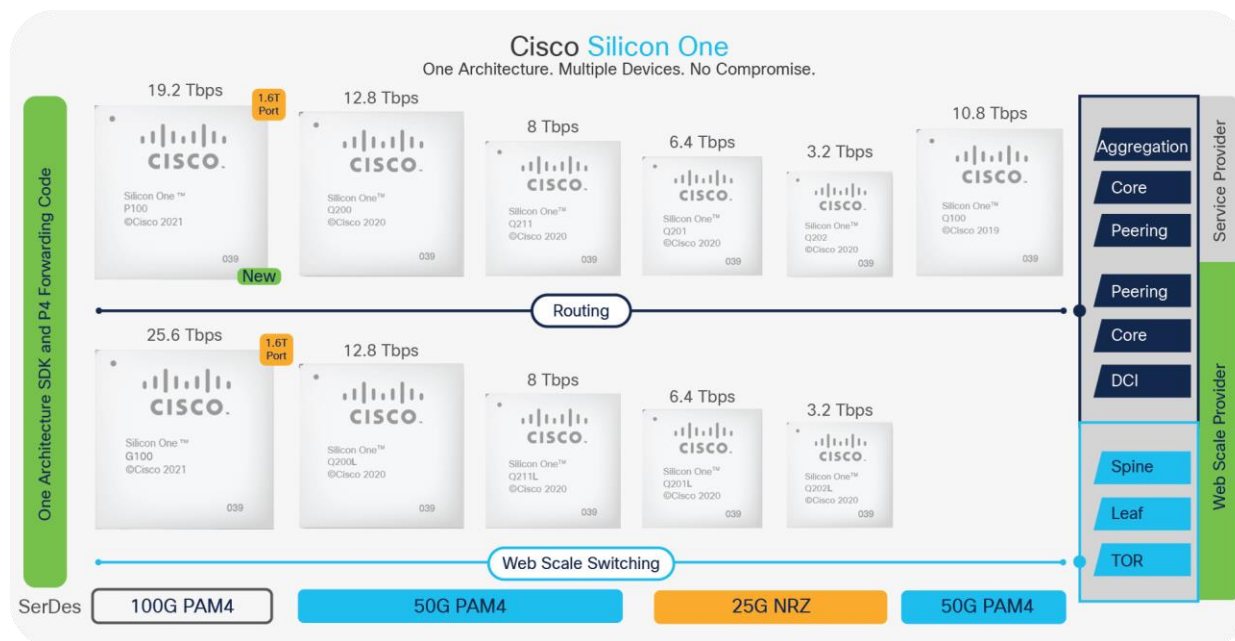

Figure 5.
A glance at the Cisco Silicon One family shows the number of products available with the 50G SerDes (blue).

The Cisco SFP-50G product family

Cisco has released the SFP-50G transceiver portfolio. The product line encompasses a wide variety of high-density and low-power 50G Gigabit Ethernet connectivity options for data center, high-performance computing network, service provider access, and campus network applications. It includes:

- SFP-50G Direct Attached Cables (DAC)—Copper patch cords designed for server-to-server and server-to-switch connections, up to 5 m long

- SFP-50G Short Link (SL) Modules—Optical patch cords for distances of up to 30 m, designed as replacements for active optical cables

- SFP-50G Short Reach (SR) Modules—Transceivers for OM3/OM4/OM5 multimode fiber, reaches of up to 100 m, depending on network configuration

- SFP-50G Long Reach (LR) Modules—Designed for single mode fiber, capable of links as long as 15 km (with FEC)

Not every organization is ready to make the shift to 50G immediately, so the product line also includes the ability to interface with 25G. Notably, the SFP-50G is "dual rate," which means it supports both 50G and 25G data rates. This provides an efficient and cost-effective migration path by allowing customers to initially connect with equipment at the other end of the line at 25G; when they're ready to upgrade to 50G, they only need to change a software parameter.
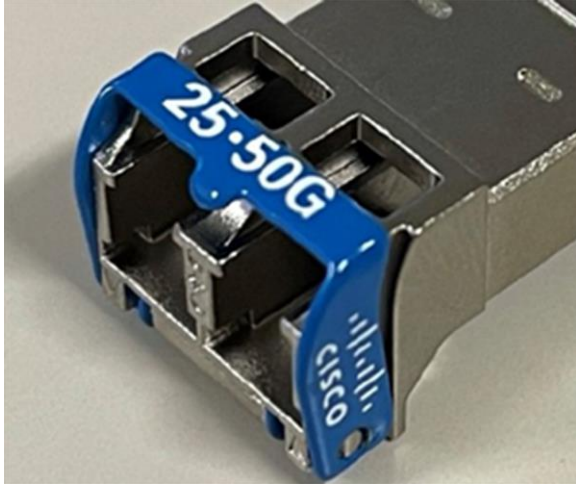
Figure 6.
Dual-rate SFP transceiver lets network operators flexibly manage their migration to 50G.

All products are compliant to the relevant IEEE standards (see table 3). They are interoperable with other IEEE-compliant 25G and 50G interfaces as appropriate. Each product has been tested and certified on Cisco® SFP56 ports to guarantee performance and reliability.

Table 2.   Standards for 50G optical interfaces

Standard	Description	Reach
IEEE P802.3cd	Amendment of IEEE standard 802.3 802.3™ 2012 (December 5, 2018)	SR, FR, LR

Conclusion

50G–The practical upgrade path for network operators

The SFP56 form factor is a key element of the data rate roadmap for service providers. It doubles the data rate compared to 25G modules. Historically, that kind of speed increase in an identical form factor has significantly reduced the cost per bit. SFP56 modules use PAM4 signaling, which enables them to directly interface to switch ASICs with 50G QSFP28, which is an advantage in many applications.

The comprehensive SFP56 product portfolio includes transceivers for single-mode fiber and multimode fiber for various reaches, as well as active and passive copper cables for short-reach interconnects.

Learn more

To learn more about how SFP-50G products can help you upgrade capacity to stay ahead of customer demand, contact your Cisco representative.

www.cisco.com/go/optics

Printed in USA   C11-3109901-00   08/22