



25Gb Ethernet: The new standard for data center connectivity

Accelerate network performance and reduce costs of server-to-switch connections



Data center bandwidth requirements are growing at double-digit rates. On the demand side, data centers are expanding with a 20 percent growth in servers, and virtualization density is increasing by 30 percent,¹ driving increased bandwidth requirements. On the supply side, multi-core processing power, NVMe and PCIe 3.0 I/O buses, and high-speed flash storage require networking that can keep pace with increased throughput.

To meet these needs, networking and the Ethernet industry are moving in a new direction. Discussions with enterprise customers that previously centered on 10GbE, 40GbE and 100GbE networking have expanded to include 25GbE and 50GbE technology as well. In particular, many of these discussions center on moving from 10GbE to 25GbE for top-of-rack (ToR) switching. This paper explains why, and shows how changes in demand, technology and networking economics are working together to alter the networking landscape.

Networking at the crossroads

As organizations push more traffic into their enterprise and private data centers, and with bigger “pipes” coming into the enterprise, more bandwidth is also required as data moves down into the fabric or ToR. ToR architectures utilizing 10GbE to the server are among the largest number of connections in data centers today. But these architectures are rapidly outgrowing 10GbE capabilities, often requiring multiple ports.

Organizations would need to deploy twice as many 10GbE switches to keep up with required network bandwidth, along with additional cables, space, power and cooling. That would mean a significant increase in capital and operating expenses. The traditional next step up to meet demands, 40GbE, is not cost-effective or power-efficient enough for many organizations. Moreover, as organizations look ahead, there is no simple, economical path from 10GbE to 40GbE to 100GbE — it requires organizations to make three product transitions.

To better address high-performance network and scalability requirements, industry leaders including QLogic and Dell are increasingly implementing 25GbE for their customers. These implementations capitalize on the 25GbE specification adopted by the 25 Gigabit Ethernet Consortium. The specification makes use of single-lane 25 Gbps Ethernet links and is based on the existing IEEE 100GbE standard (802.3bj).

With the adoption of the 25GbE specification, a new generation of switch and adapter chips is becoming available. In the last few years, the underlying high-speed communication technology has advanced to the point that 25 GHz signaling interfaces have become economically viable, and all of the various physics-related challenges in signal integrity have found reliable solutions.

The economics of 25GbE

The availability of 25GbE switches and adapters has changed the economics of upgrading server-to-switch connectivity. New 25GbE chips use the same class of cables as 10GbE but deliver 2.5 times the performance, while using the same

number of lanes as 10GbE. That means 25GbE switches can be manufactured very cost-effectively. Using a single lane on a switch or adapter, 25GbE ports also compare favorably to 40GbE technology that requires four lanes.

Deploying 25GbE networks also enables organizations to significantly reduce the required number of switches and cables — along with space, power and cooling — compared to similar bandwidth solutions using 10GbE and 40GbE technology. Fewer network components reduce ongoing management and maintenance costs while more efficiently utilizing servers already in place. Most important, with 25/50/100GbE multi-rate connectivity, enterprises now have a path to 100GbE that is nondisruptive, can be implemented today and provides economical steps along the way.

Considering the economics, it is no surprise that the move to 25GbE (and 50GbE) is accelerating — a recent five-year forecast by industry analyst Dell’Oro Group predicts 25GbE will deliver a significant amount of Ethernet bandwidth by 2018 (see Figure 1).

Speed migration on servers—total market (Included in Dell’Oro Group’s server report)

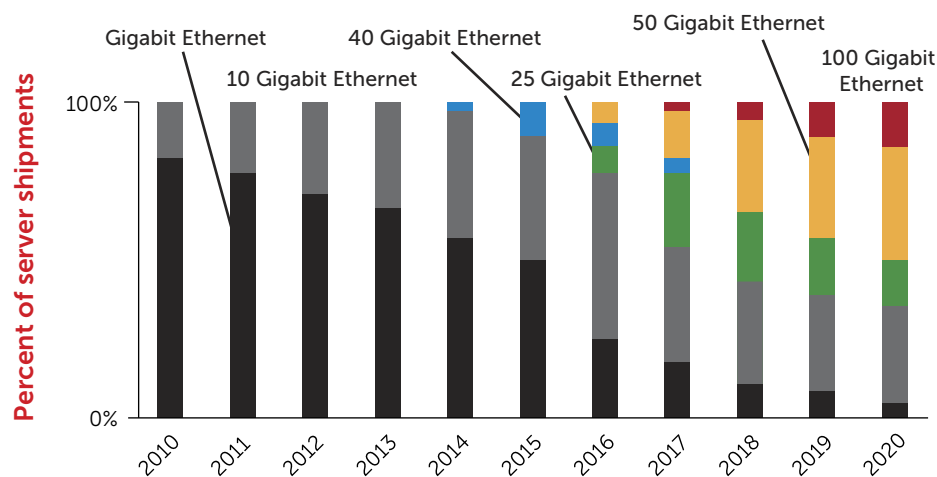


Figure 1. Projected growth of 25GbE and 50GbE networking²




Figure 2. Breakout of four 25GbE connections per 100GbE port

The impact of ASIC development

A 100GbE switch previously required multiple application-specific integrated circuits (ASICs) because the number of ports per ASIC were limited. The new generation of 100GbE switch implementations is based on high-density, industry-standard ASICs that can support up to 32 ports of 100GbE or 128 breakout ports of 25GbE in a single chip, a huge leap forward. With the QSFP28 form factor, the four lanes can be used to enable four 25GbE server connections from each 100GbE port (see Figure 2).

The move from 10 lanes to 4

Until recently, a majority of available 100GbE implementations used 10 lanes of 10GbE.

Using four-lane variants such as QSFP28 is more economical in several ways:

- Recent innovations make four 25 Gbps lane transceivers less expensive than ten 10 Gbps lanes because the transceiver is simpler and less costly to manufacture.
- The power required to run that transceiver is much less than required for a typical 10-lane transceiver.



- Short run (150M or less) multi-mode fiber (MMF) cabling runs are also less expensive because instead of requiring 10 transmit (TX) and 10 receive (RX) fiber strands, only 4 pairs of fiber for transmit and receive are needed.
- Moving from 10GbE to 40GbE typically requires a forklift upgrade to thicker, more expensive cables, but a 25GbE direct attach copper connection does not.
- For fiber connections, moving from 10GbE to 40GbE may require an upgrade to four times the number of fibers (MPO), but a 25GbE connection does not because it is the same as 10GbE (single TX and single RX, not four each for TX and RX).

The new standard for ToR upgrades

New single-lane 25GbE is not only economical compared to 40GbE, it also provides more available bandwidth than 40GbE (see *table*). While 25GbE may seem like a step back from 40GbE, it actually delivers more.

25GbE is an easier upgrade path from 10GbE as it fits into the existing model and requires half the number of PCIe lanes compared to 40GbE, leading to better PCIe bandwidth utilization and lower power consumption. The 25GbE physical interface specification also supports a variety of form factors, allowing flexible configuration options.

Partnering for your success

Enterprise organizations continue to push the envelope as they deploy the latest technologies to keep up with exponential data growth. As members of the 25 Gigabit Ethernet Consortium, QLogic and Dell are dedicated to bringing the benefits of the new technology to their customers. QLogic was the first to demonstrate an end-to-end 25GbE solution, and Dell introduced the industry's first high-density, 100GbE multi-rate switch in a 1RU fixed-form factor for aggregation and access layers (see *sidebar on Dell Networking switches, next page*). Both Dell and QLogic 25GbE technologies are backward-compatible with 10GbE, allowing customers to build and cross-connect a heterogeneous-speed Ethernet network.

Considering the uses and benefits, new switches based on single-lane 25GbE technology represent an important step forward in the performance and economics of networking: two-and-a-half times the performance of existing 10GbE, with greater port density and lower costs than 40GbE. The 25GbE technology empowers organizations to cost-effectively scale network bandwidth in support of next-generation server solutions in data center environments.

About QLogic

QLogic is a global leader and technology innovator in high-performance server and storage networking connectivity products.

About Dell

Dell integrates QLogic connectivity into Dell Networking switches, PowerEdge servers and other solutions. Dell listens to customers and delivers worldwide innovative technology, business solutions and services that give them the power to do more.

At a glance: Reshaping data center networks

Factors enabling 25GbE advances include:

New silicon

- ASIC development now allows high-density 100GbE switches in affordable form factors
- 32x100GbE switching bandwidth is now available on a single chip
- The new switches offer the ability to run 4x25GbE off of a 100GbE port

New optics

- MMF optics have been reduced to 4 RX and TX fibers (100GBASE-SR4)
- New 100GbE optical modules use a x4 electrical interface
- These optical modules require less power than currently comparable solutions

New economics

- 25GbE reduces CapEx with fewer ToR switches at lower cost than 40GbE
- It also reduces OpEx because 25GbE switches require less power, cooling and space
- 25 Gbps Ethernet adapters adhere to standards, helping protect technology investments

Example: Ethernet switch with 3.2 Tbps capacity and 128 ports

Port speed	Lane speed (Gbps)	Lanes per port	Usable ports	Total BW (Gbps)
10GbE	10	1	128	1280
25GbE	25	1	128	3200
40GbE	10	4	64	2560
100GbE	25	4	32	3200

Bandwidth comparison for 25GbE and other Ethernet speeds



Dell Networking Z9100-ON 100GbE fabric switch

Next-generation fixed-form factor
10/25/40/50/100GbE



Industry's first 100GbE multi-rate
1RU switch

- Multi-rate switching with 32 ports 100GbE (QSFP28), 32 ports 40GbE, 128 ports 1GbE, 128 ports 25GbE or 64 ports 50GbE
- Additional 2 fixed 1/10GbE SFP+ ports

Both the Dell Networking Z9100-ON and S6100-ON switches are built to support future-ready, high-end data center applications and a range of switching speeds that provide flexibility for the life of the investment. Both switches also support third-party operating systems and the Open Networking Install Environment (ONIE).

Dell Networking S6100-ON multi-rate switch

Fully modular 2RU fixed-form factor
10/25/40/50/100GbE switch




Industry's first multi-rate modular
in-rack switch

- 3 unique modules providing additional multi-rate flexibility:
 - 16x40GbE QSFP+ enabling high-density 40GbE and up to 64 ports QSFP+ in 2RU
 - 8xQSFP28 ports, allowing cost-efficient 10/25/40/50/100GbE
 - 4xQSFP28 and 4xCXP, bringing both cost-effective multi-rate switching and CXP interconnects for legacy 100GbE
- Up to 32x100GbE, 64x40GbE, 64x50GbE, 128x25GbE, 128x10GbE or 128x1GbE

QLogic FastLinQ QL4521x-DE Series

25 Gbps Ethernet-to-PCIe 3.0
Intelligent Ethernet Adapters



With superior price/performance versus 10GbE, the powerful QLogic® FastLinQ™ QL4521x-DE Series is optimized for use across enterprise, managed service provider (MSP) and large public cloud deployments.

- Leverages QLogic fifth-generation technology to deliver true 25 Gbps Ethernet performance
- Increases virtual machine density and accelerates multi-tenant networks with full offload for tunneling protocols
- Enables organizations to achieve new levels of performance in physical, virtual and cloud environments

Integrated, advanced networking helps eliminate I/O bottlenecks, conserve CPU cycles and enable greater virtual machine density in the next generation of software-defined data centers.

Learn More

For more information about Dell 25GbE switches, contact your Dell representative or visit: Dell.com/networking

To learn more about QLogic 25GbE technology, contact your QLogic representative or visit: www.qlogic.com/Products/adapters/Pages/25Gb-Ethernet.aspx

¹Infonetics Research: *Data Center Strategies North American Enterprise Survey*, May 2015

² Dell'Oro Group: *Server Report*, March 2016.

© 2016 Dell, Inc. ALL RIGHTS RESERVED. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose without the written permission of Dell, Inc. ("Dell").

Dell, the Dell logo and products — as identified in this document — are registered trademarks of Dell, Inc. in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

