

Next Steps Toward 10 Gigabit Ethernet Top-of-Rack Networking

Important Considerations When Selecting Top-of-Rack Switches

table of contents

+ Advantages of Top-of-Rack Switching	2
+ How to Get from Here to There	2
+ Juniper Can Help	4
+ Take Your Data Center to the Next Step in Its Evolution	5
+ Transform Your Data Center for the Future	5
+ It's Easy to Get Started	6

You can expand virtualization and deploy high-speed storage confidently with 10GbE top-of-rack switching. Learn how to prepare for a 10GbE data center evolution.

Many IT professionals are feeling the growing pains associated with important data center trends of the past decade, including pervasive virtualization, powerful multi-core servers, cloud-based applications and services, and storage technologies that measure IOPS in millions.

These advances have enabled high-performance, highly scalable compute models. From a networking perspective, these changes have introduced a requirement for more bandwidth. Now the technology is available to deliver 10 Gigabit Ethernet (GbE) connectivity from these servers to the switch.

As networking professionals make the move to 10GbE, they are faced with a number of questions: The first involves the types of applications running in the data center. Another result of changing technology in the data center is the far greater volume of server-to-server (or east-west) traffic compared with server-to-client (or north-south) traffic. It is inefficient to force traffic between servers through a legacy GbE switch, which can easily become overtaxed and create a bottleneck that slows application performance. Add powerful storage to the mix, which can send traffic on different interfaces and with different patterns, and it's easy to understand the need for a bigger L2 "pipe." Outdated switches and tree-like north-south network architectures cannot process packets fast enough to keep up with data-hungry virtualized servers and the storage that's trying to feed them at rates of up to 16 Gbps.

As problematic as they can be, however, these conditions can also be seen as opportunities. You can expand virtualization and deploy high-speed storage confidently by deploying 10GbE switching.

Once you have accepted the need to move to 10GbE switching, your next decision involves choosing switching architecture: top of rack (ToR) or end of row (EoR). At the access layer, ToR has an advantage over EoR in terms of power consumption, ease of scale when growing server PoDs, and a reduction of cable management complexity. It's easy to run cable from new compute or storage equipment to an existing switch and bring the new equipment online — that is, it should be easy. Often, however, cabling at an EoR Layer 2 (L2) switch is so messy and complicated that it undermines the agility advanced technologies were supposed to bring. The cabling chaos also contributes to higher operational expenses in the data center as the staff hours required to manage the cabling increase. Likewise, as the number of servers grows, the resulting cable sprawl increases the chance that manual cabling errors could lead to network downtime. Finally, the EoR model is difficult to sustain at scale because the switches' footprint and energy draw become costly.

This guide will help you take the next steps toward a ToR data center evolution.

Signs That Your Data Center Is Ready

How do you know when your data center is ready for its next evolution? Here are some good indicators:

- Application owners or end users complain about poor performance.
- Your service-oriented architecture delivers disappointing results because of latency.
- You delay moving more applications to your private cloud due to doubts about switch throughput.
- In some racks, 10 or more virtual machines share a single I/O port.
- Your Fibre Channel network is four or more times faster than your Ethernet network.

Advantages of Top-of-Rack Switching

With a ToR design, you can eliminate cabling nightmares, bottlenecks and rigidity while building a network foundation for modern data center applications that also provides a clear path for future growth. “ToR” is a data center network design in which a feature-rich 10GbE switch handles L2 and L3 processing, data center bridging and Fibre Channel over Ethernet (FCoE) for an entire rack of servers. This approach contributes to an agile infrastructure because the ToR switch supports multiple I/O interfaces, including GbE, 10GbE and 40GbE. Additional advantages include:

- **Support for server-to-server application traffic:** Some ToR switching architectures optimize the connectivity between ToR switches to improve server-to-server traffic latency and throughput to deliver better application performance.
- **Future-proofing:** Emerging standards such as 100GbE can leverage the same fiber cable architecture when you are ready to deploy them.
- **Cabling simplicity:** ToR designs eliminate the EoR switch, cabling messes and the single L2 pathway of a north-south configuration.
- **Better support for multi-core servers:** Large multicore servers get the bandwidth, low latency and low-jitter support they need to provide high performance and reliability.
- **Evolution to a flat fabric-based architecture:** This can provide greater scale, simplicity and agility.

How to Get from Here to There

If you are feeling growing pains associated with advanced virtualization and storage technologies, a 10GbE switching architecture might be the answer. While 10GbE results in significant gains in performance and simplicity, there are significant implications for the network. Whether you need connectivity at the access layer or need to scale your data center network with a fabric-based switching architecture to keep pace with evolving business requirements, you can take a few simple steps now to move in that direction, starting with an evaluation of ToR switches. What characteristics do you need in a 10GbE ToR switch? It depends on where you hope to take your network, as shown in Table 1.

Table 1. Characteristics to look for in a ToR switch

If your ideal end state is...

- Flexibility, greater virtual machine density, low latency, high performance and scalable bandwidth
- Storage interconnect, long-term lossless Ethernet, moving toward a converged fabric

Look for ToR switches that can...

- Support both GbE and 10GbE so that you can continue using existing switches during the transition
- Serve as either a standard Ethernet switch or a ToR L2 switch
- Support multiple protocols (GbE, 10GbE, 40GbE) and transit switching for iSCSI

If your ideal end state is...

Juniper Networks QFX5100 line of switches.



QFX5100-24Q



QFX5100-48S



QFX5100-96S

- Highly nimble switch optimized for top-of-rack deployments of various densities and speeds, as well as virtualized data center environments
- Universal building block for Juniper's mixed 1/10/40GbE Virtual Chassis, Virtual Chassis Fabric and QFabric architectures
- Supports open architectures such as Spine and Leaf and Layer 3 fabrics, as well as virtualized network environments including Juniper Contrail and VMware NSX Layer 2 gateway services
- High-density platforms offer 72 ports of 10GbE or 32 ports of 40GbE in a 1 U form factor, or 96 ports of 10GbE in a 2 U platform
- Optional uplink modules provide up to eight additional 40GbE QSFP+ ports
- Up to 2.56 Tbps Layer 2 and Layer 3 performance and as low as 550 ns latency

Juniper Can Help

Whether your end-state vision is a fully converged network and storage fabric or a highly virtualized data center with maximum flexibility and performance, Juniper ToR technologies can help.

Take Your Data Center to the Next Step in Its Evolution

If your data center has a lot of GbE equipment, Juniper QFX5100 switches can maximize throughput and performance while allowing an easy migration to 10GbE. As your data centers add more 10GbE servers, the high-density QFX5100 platforms easily accommodate the extra capacity. With up to 48 10GbE ports in a single rack-unit device, or 96 10GbE ports in a two rack-unit platform, QFX5100 switches can easily clear the bottlenecks that impede virtualization expansion and deployment of advanced storage technologies.

The QFX5100 is also the universal building block for Juniper Networks' unique Virtual Chassis, Virtual Chassis Fabric and QFabric technologies. In a Virtual Chassis configuration, up to 10 switches can be interconnected and operated as a single logical device running a single OS. QFX5100 switches can participate in the same Virtual Chassis configuration as existing Juniper EX4300, QFX3500 and QFX3600 switches, providing additional configuration flexibility for mixed GbE and 10GbE environments.

In a Virtual Chassis Fabric configuration, up to 20 QFX5100 switches can be deployed in a spine-and-leaf architecture that is ideal for high-performance, low-latency data center environments. Up to four QFX5100s can be deployed in the spine while up to 18 QFX5100, EX4300, QFX3500 and QFX3600 switches can be deployed as leaf nodes, providing any-rack-to-any-rack deterministic throughput while simplifying network operations through a single point of management.

Transform Your Data Center for the Future

If your goals for data center transformation are more ambitious, the QFX5100 represents a strong ToR choice. It lets you connect directly to network-attached storage and can serve as a standalone switch for highly demanding 10/40GbE environments. If you want to scale to more than the 10 switches allowed in by Virtual Chassis technology, the QFX5100 switches provide a smooth transition to Virtual Chassis Fabric spine-and-leaf deployments that can accommodate up to 20 switches operating as a single, logical device.

If additional scalability is required, the QFX5100 offers a fabric-ready edge solution that provides Node functionality in a Juniper Networks QFabric System. The QFabric System, composed of multiple components working together as a single switch that can support more than 6,000 10GbE ports, delivers any-to-any connectivity and simplified operations, allowing you to seamlessly converge Ethernet and storage networks to support the flexibility you need for highly virtualized and cloud environments.

It's Easy to Get Started

If your data center is showing signs of age, 10GbE ToR switching might be the right next move. ToR designs provide immediate benefits, such as simplified cabling, a smaller switching footprint and reduced energy consumption. They also provide a 10GbE foundation that delivers high performance and throughput for virtualized data centers and high-speed storage solutions. To make the transition to a ToR design a smooth one, look for a switch that supports multiple protocols and can easily scale — a switch such as the Juniper Networks QFX5100. Juniper technologies, including Virtual Chassis, Virtual Chassis Fabric, QFabric and the Junos OS, offer management simplicity and the demanding performance that today's data centers require.

Learn more about ToR switching and Juniper solutions by visiting www.juniper.net/us/en/products-services/switching/.