

The Adaptive Architecture for the Data Center Network

Abstract

Evolving business needs and increasing application complexity are placing mounting demands on the server environment in the data center. A single approach to server infrastructure is no longer appropriate and is being replaced with multiple server strategies, each designed to meet different functional and cost objectives. This change is, in turn, placing new demands on the data center network that links these servers together, because each server strategy has differing implications on the network design and operations. Cisco® offers a complete portfolio of award-winning switches for the data center, so you can support multiple concurrent server strategies while still maintaining a consistent approach to the overall data center network architecture.

Introduction

As IT becomes further ingrained in business and organizational processes, the typical present-day data center must support myriad application environments. Each application has differing needs in terms of performance, availability, and cost-efficiency. As a result, organizations are exploring different approaches to meeting their particular mix of technical, functional, and cost requirements.

The critical principle to keep in mind is that the data center network should always be architecturally driven, and that the architecture should encompass a portfolio of switches to address the numerous operational and functional requirements within the data center. An effective architecture will offer the following operational characteristics:

- Access layer flexibility
- Capacity and service scalability without causing disputation across domains
- High availability and fabric stability
- Feature consistency
- Easy operations and management
- A simple, deterministic topology
- Effective segmentation (i.e. including the control and data planes) for security policy and regulatory compliance

The corollary to this principle is that the application dictates the server choice and the server choice dictates the network access mechanism. Hence, the data center network architecture should be able to accommodate new technologies or approaches at the edge without compromising its operational characteristics.

Three Basic Approaches

There are three primary approaches for server networking today: end-of-row switching, top-of-rack switching, and integrated switching. These approaches are largely determined by the server form factors being deployed and certain operational objectives. Each design has its advantages and trade-offs and in reality, most data centers will house at least two, if not all three approaches in the access layer of their network architecture.

End-of-row switching is the most traditional approach, where a single large chassis-based switch such as the Cisco Catalyst® 6500 Series Switch is used to support one or more racks. Because the Catalyst 6500 Series can support considerable density, this approach is usually the most cost-effective in terms of delivering the highest level of switch and port utilization, especially when coupled with the rich set of network virtualization services available in the Catalyst 6500 Series. Because it also supports a wide variety of service modules, the Catalyst 6500 Series also simplifies pushing security and application networking services into the access layer, which can be a significant advantage from a compliance and security perspective.

The end-of-row approach is also the most server-independent, so it provides maximum flexibility to support a broad range of servers. In certain scenarios, end-of-row switching can provide performance advantages, because two servers that exchange large volumes of information can be placed on the same line card to take advantage of the low latency of port-to-port switching (as opposed to card-to-card or switch-to-switch, which will be slower).

The primary disadvantage of end-of-row switching is the need to run cable back to the switch. Assuming every server is connected to redundant switches, this cabling can incur considerable costs and add operational complexity. Also, the physical volume of the cable wastes valuable rack space and may impede efficient equipment cooling. And, going forward, any significant move to 10 Gigabit Ethernet for servers will present some challenges.

Top-of-rack switching is a viable choice for dense 1 rack unit (1RU) server environments. In this approach the 1RU Cisco Catalyst 4948 Switch is placed at the top of the rack and all the servers in the rack are cabled to the switch, which then has one uplink to the aggregation layer. In some instances a pair of Catalyst 4948 Switches are used for high-availability purposes. This approach significantly simplifies cable management and avoids the rack space and cooling issues of end-of-row switching. This approach also provides some architectural advantages such as fast port-to-port switching for servers within the rack, predictable oversubscription of the uplink and smaller switching domains (one per rack) to aid in fault isolation and containment.

The common challenge with the top-of-rack approach is suboptimal port utilization – not enough servers to fill the switch. One option is to put one top-of-rack switch server in an adjacent rack: this preserves the advantages of the top-of-rack switch while increasing port utilization. However this approach should be done sparingly and with good discipline to avoid devolving into a poorly executed end-of-row arrangement. Finally, with a top-of-rack approach, attention must be paid to ensure that commonly used or high-volume data pathways do not end up spanning multiple racks and thus multiple switches.

Because top-of-rack switches are traditionally fixed-configuration boxes, they create the need for the services (security, application networking) to be implemented deeper into the network. There are a number of ways to address this, for example as a separate services layer, or perhaps folded into the aggregation layer. Because each data center has unique needs and priorities, it is important to understand the pros and cons of each approach before choosing a strategy.

As a final caution, care must be taken in selecting a top-of-rack switch. It is tempting to deploy a wiring-closet-type switch, because it has the appropriate form factor and port count. However, a top-of-rack switch appropriate for the data-center has much higher demands placed upon it than a wiring closet switch, so while the Cisco Catalyst 4948 may “look” like other switches, its internal design (I/O performance, buffering capacity, how it drops packets) is optimized for the data center environment.

