


Classic Network Design Using Cisco Nexus 9000 Series Switches

White Paper

November 2013



| | |
|--------------------------------------------------------------------------------------------------------------------------|-----------|
| <u>Basic Data Center Network Design</u> | 3 |
| Classic Network | 3 |
| Cisco FabricPath | 4 |
| <u>Multitier Designs</u> | 8 |
| Data Center Core | 8 |
| Traditional Three-Tier Architecture with 1/10 Gigabit Ethernet Server Access | 9 |
| Option 1: Cisco Nexus 9500 Platform at the Core and the Aggregation Layer | 9 |
| Option 2: Cisco Nexus 9500 Platform at Core and Cisco Nexus 9300 Platform at the Aggregation Layer | 11 |
| Pod Design with vPC | 11 |
| Pod Design with VXLAN | 14 |
| Traditional Cisco Unified Computing System and Blade Server Access | 15 |
| <u>Connectivity with Fabric Extenders</u> | 16 |
| Fabric Extender Connectivity Options | 17 |
| <u>For More Information</u> | 17 |

Basic Data Center Network Design

The data center infrastructure is central to the overall IT architecture. It is where most business-critical applications are hosted and various types of services are provided to the business. Proper planning of the data center infrastructure design is critical, and performance, resiliency, and scalability need to be carefully considered.

Another important aspect of the data center design is the flexibility to quickly deploy and support new services. Designing a flexible architecture that can support new applications in a short time frame can result in a significant competitive advantage.

The basic data center network design is based on a proven layered approach that has been tested and improved over the past several years in some of the largest data center implementations in the world. The layered approach is the foundation of a data center design that seeks to improve scalability, performance, flexibility, resiliency, and maintenance.

Classic Network

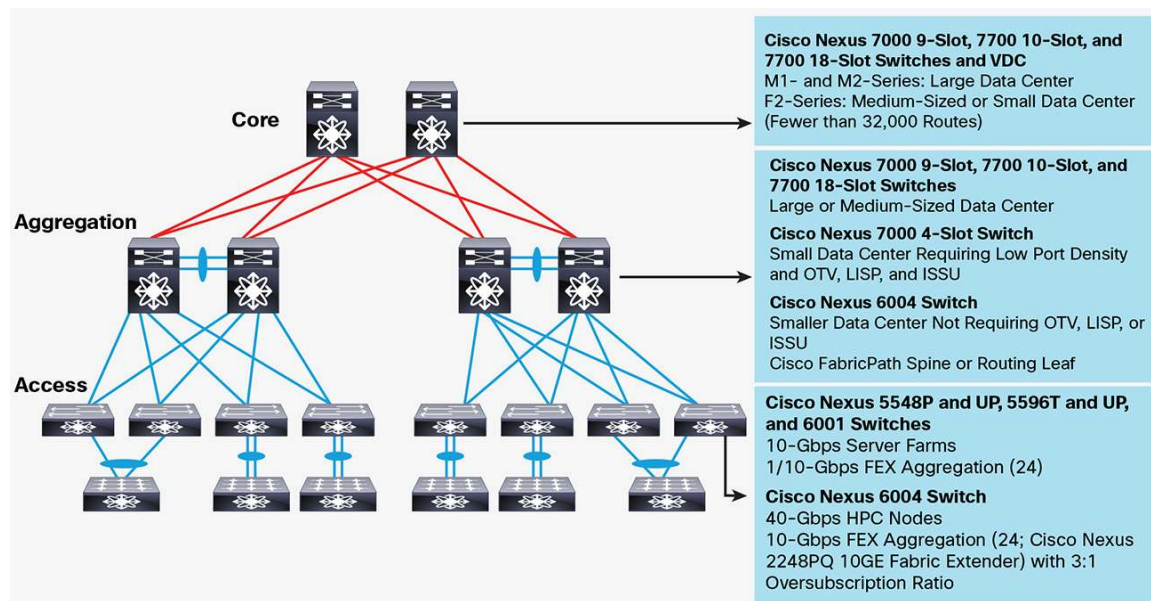
A classic network in the context of this document is the typical three-tier architecture commonly deployed in many data center environments. It has distinct core, aggregation, and access layers, which together provide the foundation for any data center design (Table 1).

Table 1. Classic Three-Tier Data Center Design

| Layer | Description |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Core | This tier provides the high-speed packet switching backplane for all flows going in and out of the data center. The core provides connectivity to multiple aggregation modules and provides a resilient Layer 3 routed fabric with no single point of failure (SPOF). The core runs an interior routing protocol, such as Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP), and load-balances traffic between all the attached segments within the data center. |
| Aggregation | This tier provides important functions, such as service module integration, Layer 2 domain definitions and forwarding, and gateway redundancy. Server-to-server multitier traffic flows through the aggregation layer and can use services, such as firewall and server load balancing, to optimize and secure applications. This layer provides the Layer 2 and 3 demarcation for all northbound and southbound traffic, and it processes most of the eastbound and westbound traffic within the data center. |
| Access | <p>This tier is the point at which the servers physically attach to the network. The server components consist of different types of servers:</p> <ul style="list-style-type: none">• Blade servers with integral switches• Blade servers with pass-through cabling• Clustered servers• Possibly mainframes <p>The access-layer network infrastructure also consists of various modular switches and integral blade server switches. Switches provide both Layer 2 and Layer 3 topologies, fulfilling the various server broadcast domain and administrative requirements. In modern data centers, this layer is further divided into a virtual access layer using hypervisor-based networking, which is beyond the scope of this document.</p> |

Figure 1 shows a classic design using the current Cisco Nexus® product portfolio, including Cisco Nexus 7000 Series Switches and 2000 Series Fabric Extenders (FEXs). You can use this three-tier design to migrate to the new Cisco Nexus 9000 Series Switches.

Figure 1. Current Cisco Nexus Portfolio Scenarios for Transitioning to Cisco Nexus 9000 Series



Many types of services, primarily firewalls and load balancers, can be integrated into these designs. Careful planning is needed for a smooth migration from this type of hardware and topology combination to the new Cisco Nexus 9000 Series hardware and topology combination.

The main features of the new Cisco Nexus 9000 Series are support for FEX, virtual PortChannel (vPC), and Virtual Extensible LAN (VXLAN). The data center architecture can be deployed in a classic design in which existing designs variations are supported, such as the following:

- Data center pods
- Large-scale multitier designs
- VXLAN fabric

Cisco FabricPath


Cisco FabricPath provides another possible combination of technology in which a move to the Cisco Nexus 9000 Series affects the topology. Cisco FabricPath allows the creation of simple, scalable, and efficient Layer 2 domains that apply to many network scenarios. Cisco FabricPath brings the stability and scalability of routing to Layer 2.

With Cisco FabricPath, the switched domain no longer has to be segmented, providing data center-wide workload mobility. Because traffic is no longer forwarded using Spanning Tree Protocol, the bisectional bandwidth of the network is expanded, providing enhanced scalability and a completely nonblocking environment. This type of topology can also be transitioned to the Cisco Nexus 9000 Series, but without the use of Cisco FabricPath in the end state.

A spine-and-leaf topology would need to be planned into the Cisco Nexus 9000 Series design. A spine node is a node that connects to other switches in the fabric, and a leaf node is a node that connects to servers (Figure 2). The current Cisco Nexus portfolio remains the same, focusing on the Cisco Nexus 7000, 6000, and 5000 Series chassis.

From a design perspective, the main point here is that other than the hardware transition, everything will more or less remain the same with the insertion of the Cisco Nexus 9000 Series chassis.

Figure 2. Design Variant Using Cisco FabricPath in a Spine-and-Leaf Topology



Calculations for oversubscription ratios, MAC address scaling, port densities, etc. still apply in the transition from the current Cisco Nexus portfolio to the Cisco Nexus 9000 Series platforms.

In planning a migration, care needs to be taken from a Layer 2 perspective (for example, with Rapid per-VLAN Spanning Tree [RPVST+], vPC, and Cisco FabricPath). Current policies, such as those for access control lists (ACLs) and virtual LAN ACLs (VACLs), must be considered in terms of traffic flow and so on and any other application touchpoints.

The Cisco Nexus 9000 Series provides features optimized specifically for the data center:

- High 10-Gbps densities
- 40-Gbps port densities
- Reliability
- Performance
- Scalability
- Programmability
- Ease of management

With their exceptional performance and comprehensive feature set, Cisco Nexus 9000 Series Switches are versatile platforms that can be deployed in multiple scenarios, including the following:

- Layered access-aggregation-core designs
- Leaf-and-spine architecture
- Compact aggregation-layer solutions

Cisco Nexus 9000 Series Switches deliver a comprehensive Cisco® NX-OS Software data center switching feature set. Table 2 lists the current form factors.

Table 2. Cisco Nexus 9000 Series Switches

| Device Model | Line Cards and Expansion Modules | Description | Deployment |
|----------------------------|----------------------------------|-------------------------------------------------------------------|--------------------------------------------------------------------|
| Cisco Nexus 9508 Switch | N9K-X9636PQ | 36-port 40-Gbps Enhanced Quad Small Form-Factor Pluggable (QSFP+) | End of row (EoR), middle of row (MoR), aggregation layer, and core |
| | N9K-X9564TX | 48-port 1/10GBASE-T plus 4-port 40-Gbps QSFP+ | |
| | N9K-X9564PX | 48-port 1/10-Gbps SFP+ plus 4-port 40-Gbps QSFP+ | |
| Cisco Nexus 9396PX Switch | N9K-C9396PX | Cisco Nexus 9300 platform with 48-port 1/10-Gbps SFP+ | Top of rack (ToR), EoR, MoR, aggregation layer, and core |
| Cisco Nexus 93128TX Switch | N9K-C93128TX | Cisco Nexus 9300 platform with 96-port 1/10GBASE-T | Top of rack (ToR), EoR, MoR, aggregation layer, and core |

With new business services and applications requiring new data center infrastructure designs, it is important to consider the implications for these new designs for current services and applications. Factors such as Layer 2 mobility, bandwidth and latency, and symmetrical paths through load balancers and firewalls are important to plan for to help ensure successful migration of business services from the current setup to a new data center infrastructure.

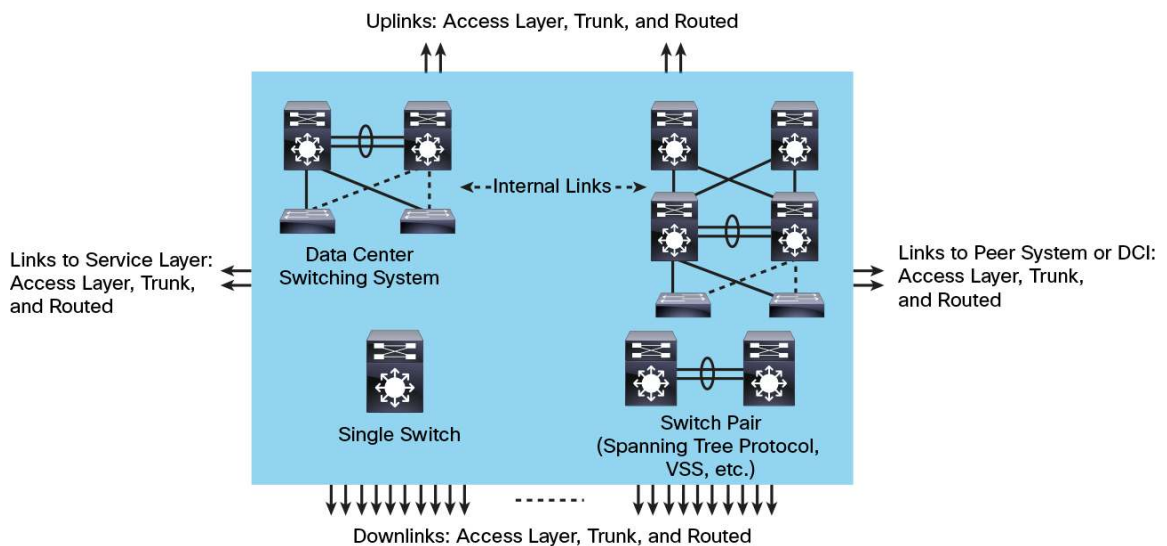
A data center switching system (DCSS) consists of one or many switches (of any kind) that are interconnected so that they collectively provide Layer 1 through 3 connectivity to servers and Layer 4 through 7 devices and the applications that connect to them (Figure 3). Even in the case of Cisco Catalyst® 6500 Series service modules such as the Cisco Catalyst 6500 Series Firewall Services Module (FWSM), Cisco Application Control Engine (ACE), and Cisco Catalyst 6500 Series Network Analysis Module (NAM), which share a chassis, an internal connection still exists between the switch and the service modules.

A DCSS provides network connectivity at Open Systems Interconnection (OSI) Layers 1 through 3 between the end devices using VLANs, switched virtual interfaces (SVIs), Virtual Routing and Forwarding (VRF), routing, ACLs, etc. and consists of the following:

- One or more switches
- Inter-switch links (ISLs; Layer 2 or 3)

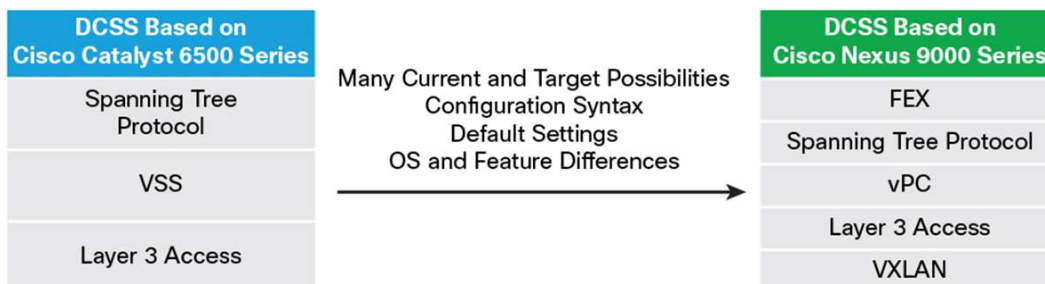
- Interfaces
 - Uplinks to data center core or edge
 - Downlinks to servers and hosts
 - Services links to Layer 4 through 7 devices
 - Data center interconnect (DCI) links to peer DCSSs

Figure 3. Data Center Switching System



In migrating your data center to the Cisco Nexus 9000 Series, you need to consider not only compatibility with existing traditional servers and devices; you also need to consider the next-generation capabilities the Cisco Nexus 9000 Series, including 10/40-Gbps connectivity, Layer 2 mobility, new features, high performance, and programmability (Figure 4).

Figure 4. Data Center Migration



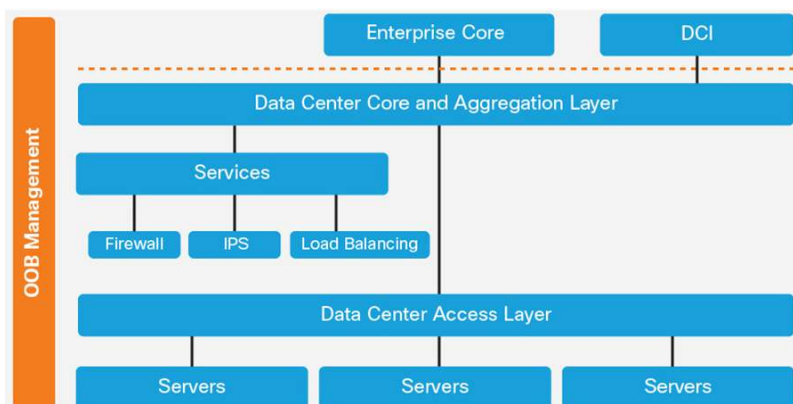
This document provides guidance in the planning, design, and deployment of a data center infrastructure based on the Cisco Nexus 9500 platform.

Multitier Designs

This section discusses network architectures based on a multitier model (Figure 5):

- The data center core interconnects all the building blocks.
- The enterprise core building block is used to connect the rest of the enterprise network, such as campus, WAN, and Internet building blocks located in other data centers.
- The core building block also has direct connectivity to other data centers.

Figure 5. Building-Block View of the Data Center



In a large data center, a single pair of core switches typically interconnects multiple aggregation-layer modules using 10 Gigabit Ethernet Layer 3 interfaces.

Figure 4 shows the core, aggregation, and access layers, but in a more complete picture of the data center, other components are connected to the typical tiers.

Data Center Core

The core provides a fabric for high-speed packet switching between multiple aggregation modules. This layer serves as the gateway to the campus core, where other modules connect (for example, the extranet, WAN, and Internet edge). All links connecting the data center core are terminated at Layer 3 and typically use 10 Gigabit Ethernet interfaces to support high throughput and performance and to meet oversubscription ratios.

The data center core is distinct from the campus core, with a different purpose and different responsibilities. The data center core is not necessarily required, but it is recommended when multiple aggregation modules are used for scalability. Even when a small number of aggregation modules are used, a campus core may be appropriate to connect the data center fabric.

When determining whether to implement a data center core, consider the following:

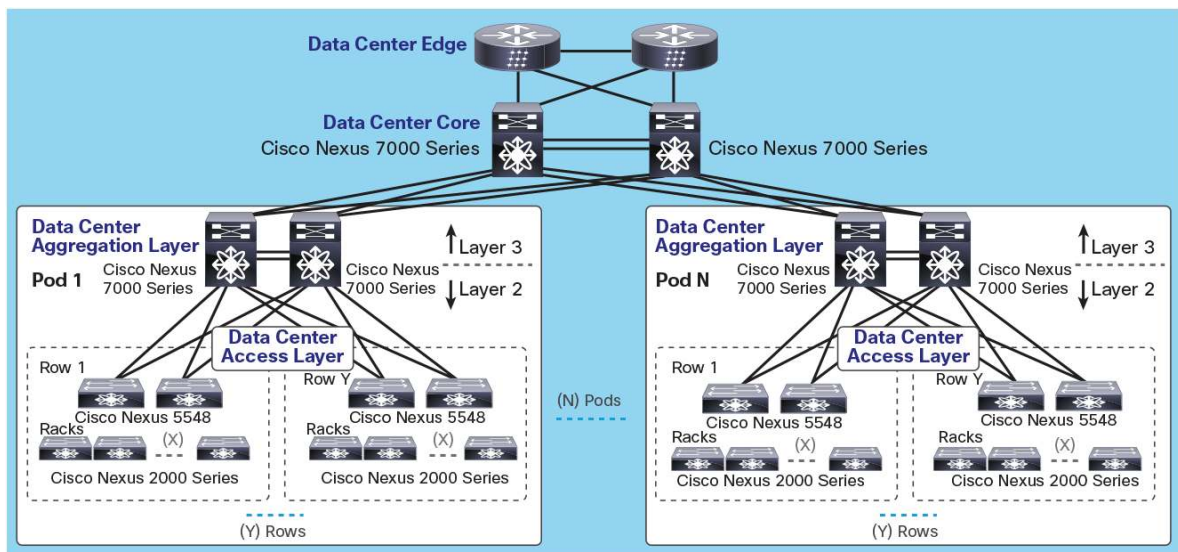
- Administrative domains and policies: Separate cores help isolate campus distribution-layer and data center aggregation-layer administration and policies, such as quality-of-service (QoS) policies, access lists, troubleshooting, and maintenance.
- 10 Gigabit Ethernet port density: A single pair of core switches may not support the number of 10 Gigabit Ethernet ports required to connect the campus distribution-layer and the data center aggregation-layer switches.
- Future impact: The potential business impact of implementing a separate data center core at a later time may make implementing it during the initial implementation stage a preferable approach.

Traditional Three-Tier Architecture with 1/10 Gigabit Ethernet Server Access

In a typical data center design, the aggregation layer requires a high level of flexibility, scalability, and feature integration, because aggregation devices constitute the Layer 3 and 2 boundary, which requires both routing and switching functions. Access-layer connectivity defines the total forwarding capability, port density, and Layer 2 domain flexibility.

Figure 6 depicts Cisco Nexus 7000 Series Switches at both the core and the aggregation layer, a design in which a single pair of data center core switches typically interconnect multiple aggregation modules using 10 Gigabit Ethernet Layer 3 interfaces.

Figure 6. Classic Three-Tier Design



Option 1: Cisco Nexus 9500 Platform at the Core and the Aggregation Layer

In this design, the Cisco Nexus 9500 platform (Figure 7) replaces the Cisco Nexus 7000 Series at both the core and the aggregation layer.

The Cisco Nexus 9508 8-slot switch is a next-generation high-density modular switch with the following features:

- Modern operating system
- High density (40/100-Gbps aggregation)
- Low power consumption

The Cisco Nexus 9500 platform uses a unique combination of a Broadcom Trident-2 application-specific integrated circuit (ASIC) and an Insieme ASIC to provide faster deployment times, enhanced packet buffer capacity, and a comprehensive feature set.

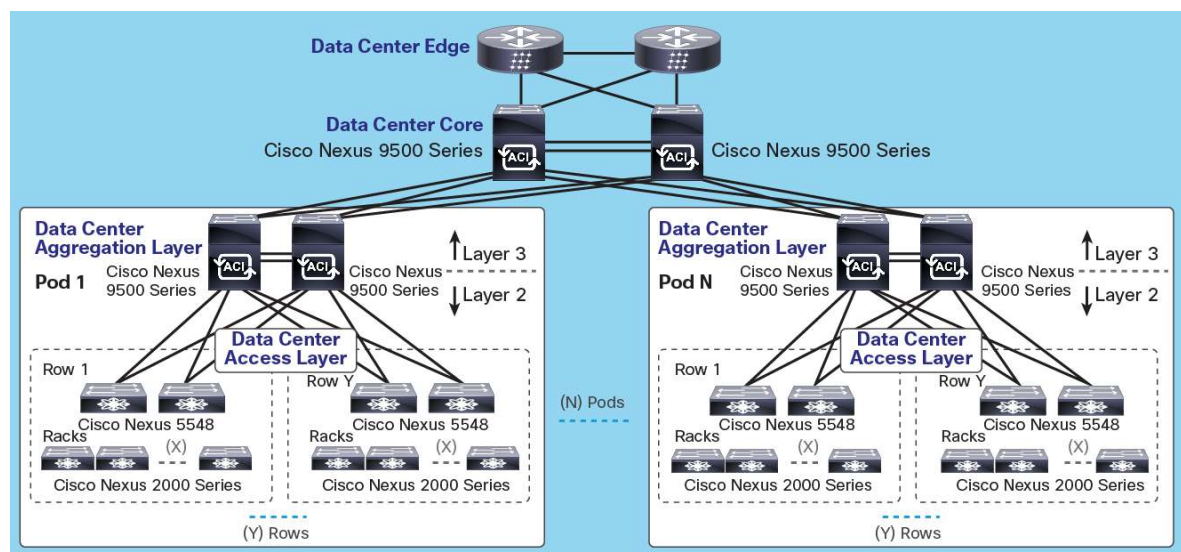
The Cisco Nexus 9508 chassis is a 13-rack-unit (13RU) 8-slot modular chassis with front-to-back airflow and is well suited for large data center deployments. The Cisco Nexus 9500 platform supports up to 3456 x 10 Gigabit Ethernet ports and 864 x 40 Gigabit Ethernet ports and can achieve 30 Tbps of fabric throughput per rack system.

The common equipment for the Cisco Nexus 9508 includes the following:

- Two half-slot supervisor engines
- Four power supplies
- Three switch fabrics (upgradable to six)
- Three hot-swappable fan trays

The fan trays and the fabric modules are accessed through the rear of the chassis. Chassis have eight horizontal slots dedicated to the I/O modules.

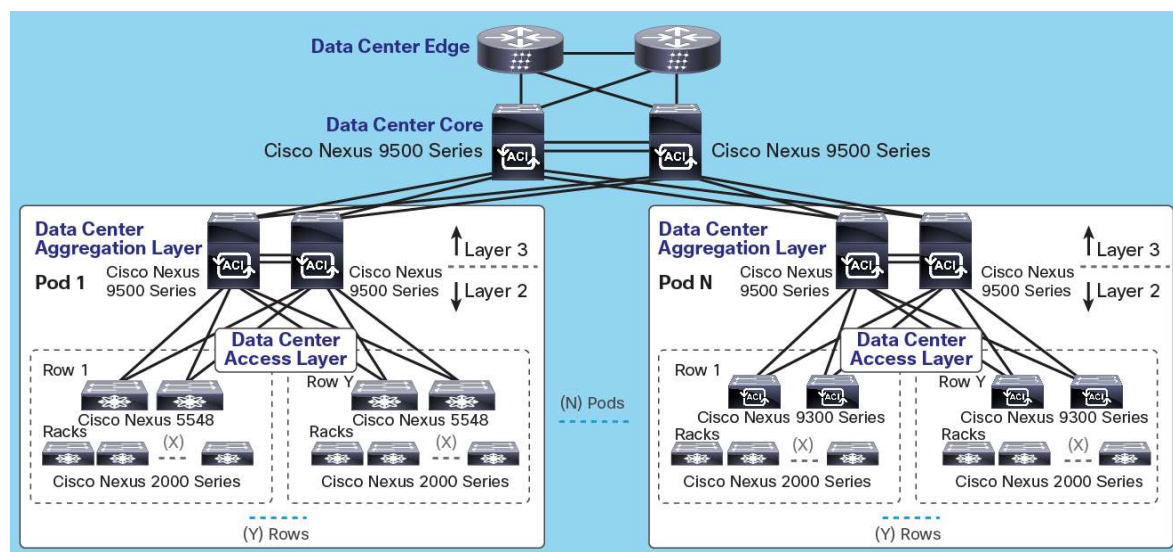
Figure 7. Cisco Nexus 9500 Platform-Based Design



Cisco Nexus 9508 Switches can be fully populated with 10, 40, and (future) 100 Gigabit Ethernet modules with no bandwidth or slot restrictions. Online insertion and removal of all line cards is supported in all I/O slots.

Option 2: Cisco Nexus 9500 Platform at Core and Cisco Nexus 9300 Platform at the Aggregation Layer

Figure 8. Cisco Nexus 9500 and 9300 Platform-Based Design



Depending on growth in the data center, a combination of the Cisco Nexus 9500 platform at the core and the aggregation layer and the Cisco Nexus 9500 platform at the core with the Cisco Nexus 9300 platform at the aggregation layer can be used to achieve better scalability (Figure 8). The Cisco Nexus 9300 platform is currently available in two fixed configurations:

- Cisco Nexus 9396PQ: 2RU with 48 ports at 10 Gbps and 12 ports at 40 Gbps
- Cisco Nexus 93128TX: 3RU with 96 ports at 1/10 Gbps and 8 ports at 40 Gbps

In both options, the existing Cisco Nexus 7000 Series Switches at the core and the aggregation layer can be swapped for Cisco Nexus 9508 Switches while retaining the existing wiring connection.

Currently, Fibre Channel over Ethernet (FCoE) support is not available for this design.

Pod Design with vPC


A vPC allows links physically connected to two different Cisco Nexus 9000 Series Switches to appear as a single PortChannel to a third device. A vPC can provide Layer 2 multipathing, which allows the creation of redundancy by increasing bandwidth, enabling multiple parallel paths between nodes and load balancing of traffic where alternative paths exist.

The vPC design remains the same as described in a vPC design guide with the exception that the Cisco Nexus 9000 Series does not support vPC active-active or two-layer vPC (eVPC). Refer to the vPC design and best practices guide for more information:

http://www.cisco.com/en/US/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf.

Figure 9 shows a next-generation data center with Cisco Nexus switches and vPC. There is a vPC between the Cisco Nexus 7000 Series Switches and the Cisco Nexus 5000 Series Switches, a dual-homed vPC between the Cisco Nexus 5000 Series Switches and the Cisco Nexus 2000 Series FEXs, and a dual-homed vPC between the servers and the Cisco Nexus 2000 Series FEXs.

Figure 9. vPC Design Considerations with Cisco Nexus 7000 Series in the Core



In a vPC topology, all links between the aggregation and access layers are forwarding and are part of a vPC.

Gigabit Ethernet connectivity makes use of the FEX concept. Spanning Tree Protocol does not run between the Cisco Nexus 5000 Series Switches and the Cisco Nexus 2000 Series FEXs. Instead, proprietary technology keeps the topology between the Cisco Nexus 5000 Series Switches and the fabric extenders free of loops. Adding vPC to the Cisco Nexus 5000 Series Switches in the access layer allows additional load distribution from the server to the fabric extenders to the Cisco Nexus 5000 Series Switches.

An existing Cisco Nexus 7000 Series Switch can be replaced with a Cisco Nexus 9500 platform switch with one exception: Cisco Nexus 9000 Series Switches do not support vPC active-active or two-layer vPC (eVPC) designs. The rest of the network topology and design does not change. Figure 10 shows the new topology.

Figure 10. vPC Design with Cisco Nexus 9500 Platform in the Core


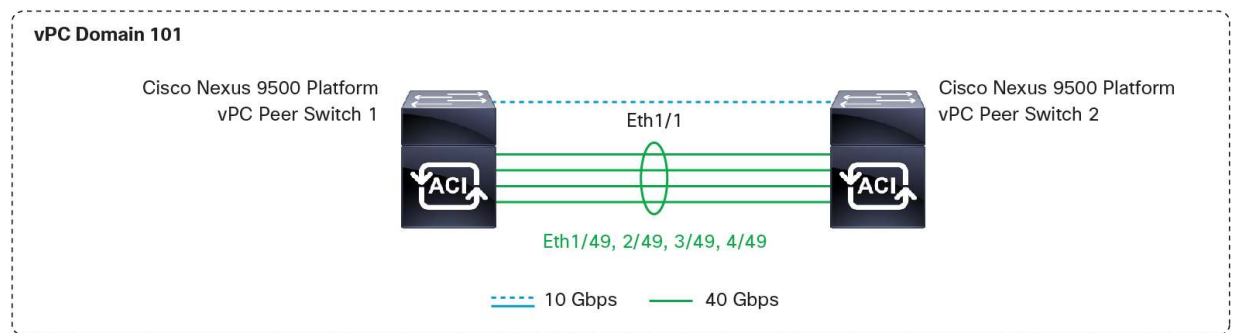


Figure 11 shows the physical peering from the Cisco Nexus 9500 platform.

Figure 11. Peering Between Cisco Nexus 9500 Platforms




Pod Design with VXLAN

The Cisco Nexus 9500 platform uses VXLAN, a Layer 2 overlay scheme over a Layer 3 network. VXLAN can be implemented both on hypervisor-based virtual switches to allow scalable virtual machine deployments and on physical switches to bridge VXLAN segments back to VLAN segments.

VXLAN extends the Layer 2 segment-ID field to 24 bits, potentially allowing up to 16 million unique Layer 2 segments in contrast to the 4000 segments achievable with VLANs over the same network. Each of these segments represents a unique Layer 2 broadcast domain and can be administered in such a way that it uniquely identifies a given tenant's address space or subnet. Note that the core and access-layer switches must be Cisco Nexus 9000 Series Switches to implement VXLAN.

Figure 12. VXLAN Design with Cisco Nexus 9000 Series



In Figure 12, the Cisco Nexus 9500 platform at the core provides Layer 2 and 3 connectivity. The Cisco Nexus 9500 and 9300 platforms connect over 40-Gbps links and use VXLAN between them. The existing FEX switches are single homed to each Cisco Nexus 9300 platform switch using Link Aggregation Control Protocol (LACP) PortChannels. The end servers are vPC dual-homed to two Cisco Nexus 2000 Series FEXs.

Traditional Cisco Unified Computing System and Blade Server Access

In a multilayer data center design, you can replace core Cisco Nexus 7000 Series Switches with the Cisco Nexus 9500 platform, or replace the core with the Cisco Nexus 9500 platform and the access layer with the Cisco Nexus 9300 platform. You can also connect an existing Cisco Unified Computing System™ (Cisco UCS®) and blade server access layer to Insieme hardware (Figures 13 and 14).

Figure 13. Classic Design Using Cisco Nexus 7000 and 5000 Series and Fabric Extenders

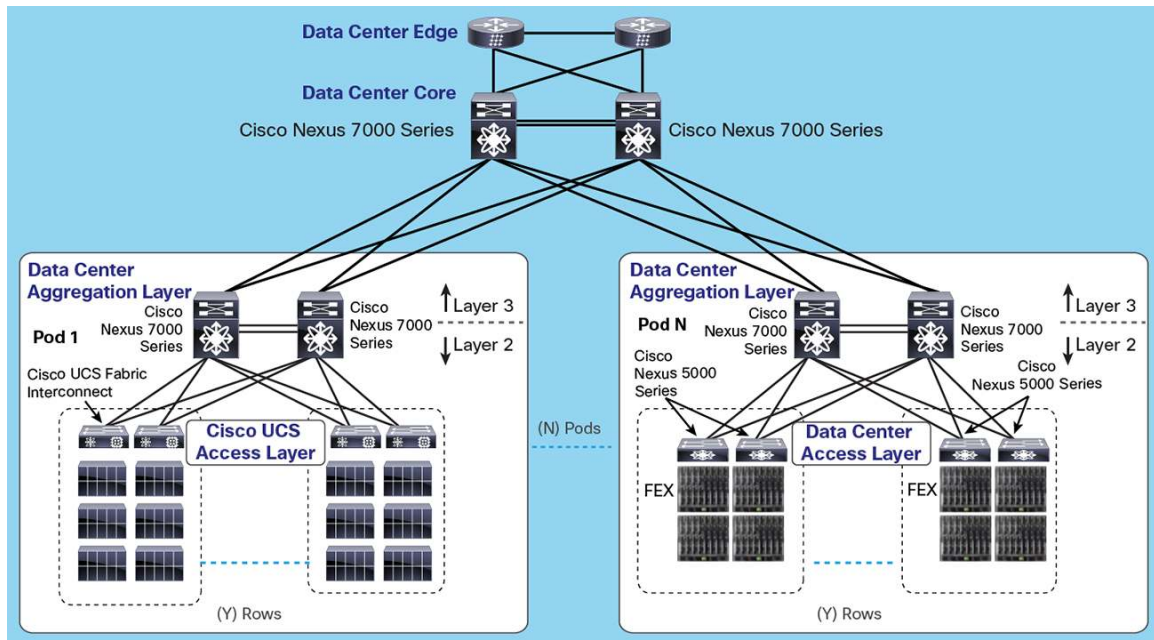
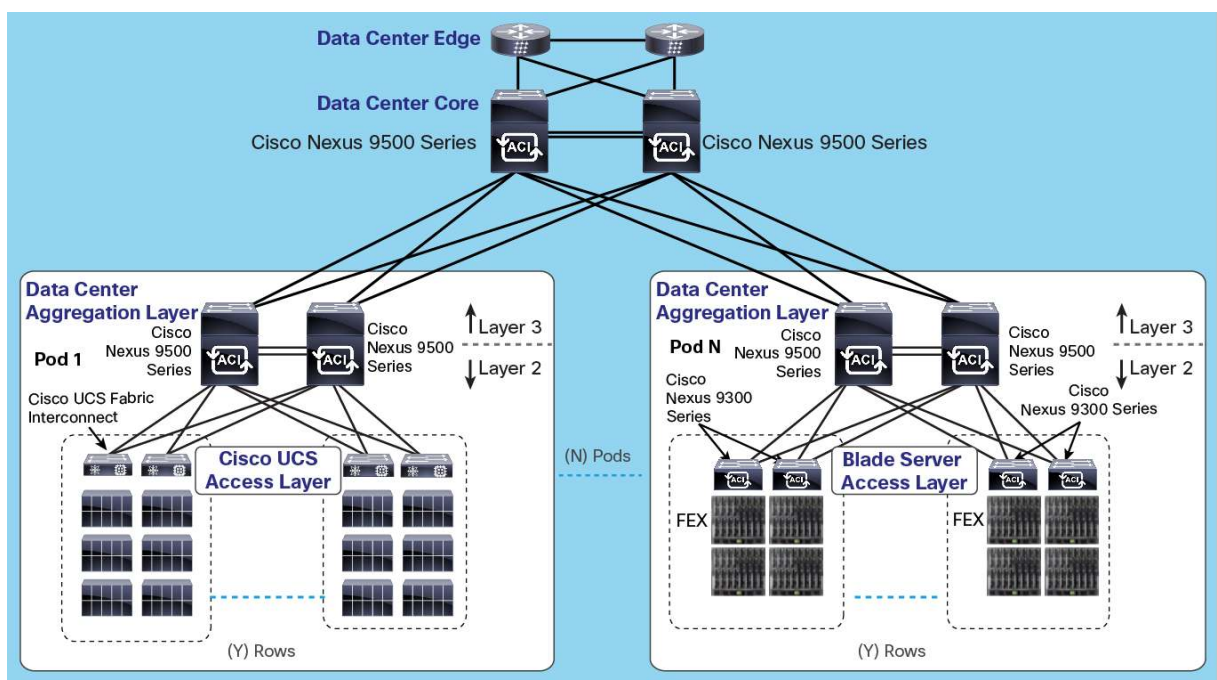


Figure 14. Classic Design Using Cisco Nexus 9500 and 9300 Platforms and Fabric Extenders



Connectivity with Fabric Extenders

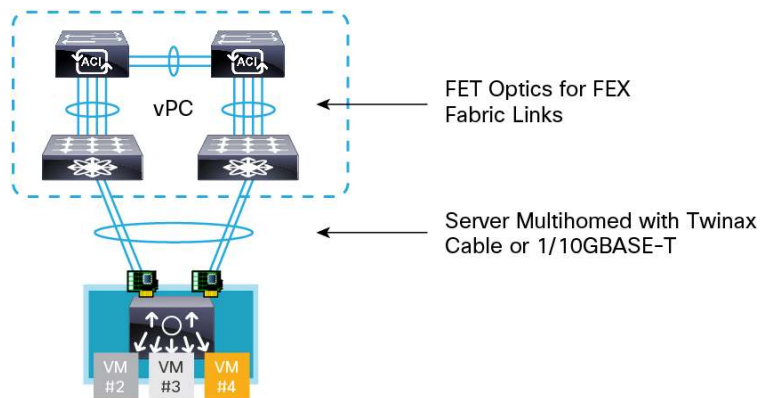
Standalone versions of the Cisco Nexus 9000 Series Switches support Cisco Nexus 2000 Series FEX connectivity using various topologies. It is important to understand the underlying port density requirements and any other scalability requirements that may affect the design: MAC address scaling, virtual links, flow capacity, bandwidth, etc.

The following FEX hardware is supported:

- Cisco Nexus 2248TP GE Fabric Extender
- Cisco Nexus 2248TP-E Fabric Extender
- Cisco Nexus 2232PP 10GE Fabric Extender
- Cisco Nexus 2248PQ 10GE Fabric Extender

Fabric extender transceivers (FETs) also are supported to provide a cost-effective connectivity solution (FET-10G). In this scenario, a pair of access-layer switches connect to the servers, and the access-layer switches connect to the aggregation-layer switches (Figure 15).


Figure 15. Fabric Extender Connection Design



Fabric Extender Connectivity Options

Figure 16 shows the initial supported connections between the fabric extenders and the Cisco Nexus 9000 Series Switches.

Figure 16. Cisco Nexus 9000 Series Fabric Extender Connectivity Support



For More Information

For more information about the Cisco Nexus 9000 platform, please visit <http://www.cisco.com/go/nexus9000>.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)