

Migrating Your Data Center to an Application Centric Infrastructure

White Paper

April 2015

Contents

Overview and Purpose	3
Cisco Catalyst 6500 in the Current Data Center Topology	3
Cisco Catalyst 6500 in the Data Center Access Layer	5
Migrating EoR Cisco Catalyst 6500 Systems to a Cisco Nexus 9000 System	7
Cisco Catalyst 6500 in the Data Center Aggregation Layer	7
Layer 4 - 7 Service Insertion	8
How the Cisco Nexus 9000 Meets Current and Future Organizational Application Needs	9
Designed for Open Programmability	9
Designed with a Modular Operating System	9
Designed for Hardware Resiliency	9
Designed for Future Scale	10
State-of-the-Art Mechanical Design	10
Data Center Fabrics	11
The Cisco Nexus 9000 Series Switches Will Have Two Modes of Operation	12
NX-OS Standalone Mode	12
ACI Fabric Mode	13
Conclusion	13
Get Started	13
Tools and Services	13
Cisco Nexus Migration Tool	13
References	14

Overview and Purpose

Organizations everywhere recognize that changing application environments are creating new demands for the IT infrastructure, causing a transformation within the data center. There is a need to address key organizational and technological transitions in terms of application delivery and consumption. As such, data center managers are assessing what new capabilities their existing networks needs to support in order ensure it can empower the:

- Highly virtualized, non-stop nature of computing environments that, in turn, demands continuous network availability
- Continued virtualization of server resources resulting in 10 Gb server connectivity requirements that require greater network scalability and performance
- Shift to a networking architecture that offers simplified operations and allows for the creation of an efficient programmable infrastructure

Cisco Nexus[®] 9000 Series Switches, through their dual-mode capabilities, allow you to deploy them as traditional switches within your existing data center network, or as the network foundation for an Application Centric Infrastructure. Cisco[®] Nexus 9000 switches offer both modular and fixed configuration switches that can be deployed into end-of-rack (EoR), top-of-rack (ToR), and middle-of-row (MoR) or spine and leaf architectures. This allows customers to easily integrate them while making your network compatible with future versions through the switches' high-performance and scalability capabilities.

Building your data center network with Nexus 9000 Series Switches allows you to:

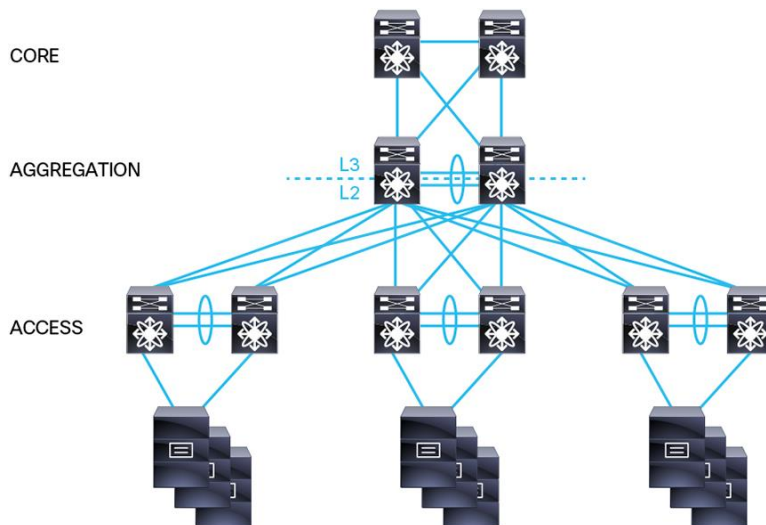
- Support application workloads across a mix of virtualized and non-virtualized server and storage infrastructure
- Provide consistent connectivity, security, and visibility across a range of bare-metal, virtualized, and cloud-computing environments
- Take full advantage of industry-leading port density, scalability, performance, power efficiency, and programmability

It is the purpose of this white paper is to provide guidance to network planners, engineers, and managers who are deploying Cisco Nexus 9000 Series Switches in a data center as replacements for Cisco Catalyst[®] 6500 Series Switches.

Cisco Catalyst 6500 in the Current Data Center Topology

The majority of data center networks are built using a three-layer hierarchical design consisting of access, aggregation, and core. This topology is well defined and was built to provide services to traditional client and server (north-south) traffic flows. This multi-tiered, hierarchical, tree-like structure has scaled out predictably by adding additional pods of equipment and by managing bandwidth over subscription. This three-layer design works well in data centers where the traffic pattern is primarily north-south, such as web (HTTP) applications. Figure 1 depicts the traditional, hierarchical network design.

Figure 1. Traditional Hierarchical Design with Cisco Catalyst 6500 Series Switches



In many data centers the Cisco Catalyst 6500 is deployed at each of these layers. Whether these 6500 switches are deployed in the access (EoR or ToR), aggregation, or core layer, the Cisco Nexus 9000 portfolio offers fixed configuration and modular switches that can be placed in that tier of the architecture.

Cisco Nexus 9000 switches running the Nexus Operating System (Cisco NX-OS) can increase the performance and reliability of these traditional three-tier networks by eliminating oversubscription and bottlenecks. Data center engineers are already familiar with Cisco NX-OS Software, so applying the new switches to solve organizational problems will occur quickly. More powerful physical servers, combined with an increase in the size and number of virtual machines per server, are raising the average volume of traffic per server port. The contention created by sending more traffic per server requires either a redesign that considers lower oversubscription, or removal of oversubscription altogether with non-blocking ports.

As data center network demands have increased, the Cisco Catalyst 6500 has continued to offer improved supervisor engines and line cards. However, the Cisco Catalyst 6500 was originally designed to make forwarding decisions in the supervisor engine. In the newer Cisco Catalyst 10 GE line cards, forwarding decisions are distributed to the line cards to increase performance and to reduce the amount of traffic required to flow through the supervisor engine. To put it mildly, punting data packets to the supervisor for switching can have a negative impact on performance as the backplane does not have enough forwarding capacity to support higher density 10 GE switching to server connections that have become the norm in customer data centers.

Cisco Nexus 9000 switches are designed to support dense, 10 GE to switch to server connections up to 1024 non-blocking 10 GE ports in a single chassis. The Nexus 9000 makes layer 2 and 3 forwarding decisions locally on the line cards. Its fabric supports full line rate on all 10 and 40 GE ports. Data packets do not go to the supervisor. In fact, there are no data ports on the Cisco Nexus 9000 supervisor engines.

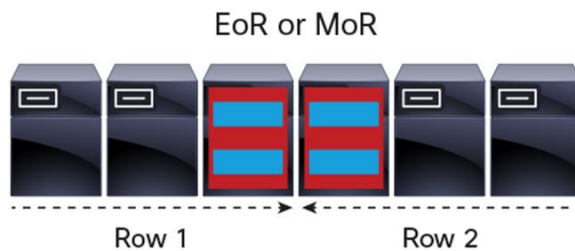
The Cisco Nexus 9000 is designed to support deployment of dense, line rate, 40 GE and 100 GE ports with low latency to meet the needs of bandwidth-intensive applications. Cisco Nexus 9000 switches were designed for the data center and offer up to 83 Tbps of performance with up to 2304 10 Gigabit or 576 40 Gigabit ports in a single chassis, as illustrated later in this white paper, in Table 1.

Cisco Catalyst 6500 in the Data Center Access Layer

The Cisco Catalyst 6500 is commonly deployed in the access layer using an EoR design that provides connectivity to hundreds of servers with 1 Gigabit Ethernet network interface cards (NICs). Each server NIC has cables connected to the EoR (or MoR) Cisco Catalyst 6500. Use existing copper cabling connecting servers to access switches. Then you can migrate from existing 1 GE interfaces on the Cisco Catalyst 6500 switches to a mix of 1/10 GE interfaces on Cisco Nexus 9000 with the same port density using less power and more efficient cooling.

Figure 2 illustrates this setup.

Figure 2. MoR Access Layer Design with Cisco Catalyst 6500 Series



The Cisco Nexus 9500 Series offers several line card options, one of which offers 48 1/10 GE BaseT ports (RJ-45 copper ports), plus an additional 4 x 40 GE QSFP+ ports. The additional 4 x 40GE ports can be used as uplinks or 16 additional 10 GE by utilizing 40 GE to 4 x 10 GE breakout cables. This provides 64 ports of 10 GE in a single line card. It offers the right density to migrate both server and uplink ports for a migration from Cisco Catalyst 6500 switches. Often in the Cisco Catalyst 6500 switch design, ports on the supervisor engine are used as uplinks. Having the additional 40 GE ports on the line card obviates the need to purchase additional higher speed line cards just for uplinks to the aggregation layer.

Figure 3 shows a rear image of the Cisco Nexus 9000 line card, providing 48 1/10G-T ports at line rate for servers and 4 x 40 GE QSFP+ ports for uplinks (or 16 more 10 GE server ports).

Figure 3. Cisco Nexus 9000 Line Card



Because new server deployments have 10 Gbps LAN-on-motherboard (LOM) ports and customers' application environments continue to evolve into more complex, multi-tiered, big-data driven applications, the use of 10 GE downlinks to servers is a common practice. In the Catalyst 6500 the backplane or switching fabric is part of the supervisor engine. The Cisco Catalyst 6500 Series Supervisor Engine 720 provides 40 Gbps per line card while the 6500 Series Supervisor Engine 2T provides 80 Gbps per line card.

When more than four or eight 10 GE ports are used on a single line card the result is oversubscription. When multiple 10 GE servers request data simultaneously, this could result in variations in latency and dropped frames because of oversubscription in the backplane or oversubscription on the uplinks (multiple 10 GE servers are aggregated to a single 10 GE uplink). A single 10 GE port can generate 14.88 million 64-byte packets per second at line rate. Multiply that potential packet flow across multiple ports, then multiply it again by multiple line cards and it quickly becomes necessary for this EoR switch to be capable of hundreds of millions of packets per second (Mpps).

If centralized forwarding and uplinks on the supervisor are used, the level of oversubscription or contention can become a cumbersome traffic engineering exercise. Preventing these variations in delays can result in a complex system to support, and often, many unused ports.

A data center switching design that uses a non-oversubscribed switching platform, such as Cisco Nexus 9000 switches, is much simpler to implement and maintain. Utilizing six fabric modules, the Cisco Nexus 9500 can scale to 1.92 Tbps to each line card (I/O slot).

Using multiple 10 GE links in a port channel does not necessarily solve the potential for congestion and dropped frames. When traffic from multiple bandwidth-intensive applications hash to the same member of a port channel, dropped frames can result. To mitigate this, a trend of using 40 GE uplinks to the aggregation layer is emerging and is highly recommended in order to accommodate the traffic of 10 GE servers.

Using the 4 x 40 GE ports that are included on the 48 1/10 GE port line card (see figure 3) for uplinks, combined with the Cisco innovation of bi-directional optics in a standards-based QSFP+ form factor, make this a cost-effective, high-performance design that can support customer data center needs for several years. (Bi-directional optics are discussed later in this paper).

Additionally, the Cisco Nexus 9000 can help reduce operational expenses as each 40 GE port in the Cisco Nexus 9508 requires less than 12 percent of the power of a Cisco 6500 40 GE port. And each 10 GE port only requires about 11 percent of the power of a Cisco Catalyst 6500 10 GE port. Additional hardware comparisons are provided in Table 1.

Table 1. Cisco Catalyst 6500 and Cisco Nexus 9500 Hardware Specifications

	Cisco Catalyst 6509	Cisco Catalyst 6513	Cisco Nexus 9504	Cisco Nexus 9508	Cisco Nexus 9516 (Q2CY14)
Power per 10 GE port	Approx. 33 W	Approx. 33 W	Approx. 3.5 W	Approx. 3.5 W	Approx. 3.5 W
Power per 40 GE port	Approx. 117 W	Approx. 117 W	Less than 14 W	Less than 14 W	Less than 14 W
Supervisor engines	2 (SUP2T)	2 (SUP2T)	2	2	2
Height (rack units)	14	19	10	13	20
Airflow	Side-Side	Side-Side	Front-Back	Front-Back	Front-Back
Line-card (I/O) slots	7	11	4	8	16
Bandwidth per chassis	2.08 Tbps	2.08 Tbps	15 Tbps	30 Tbps	60 Tbps
Bandwidth per slot	80 Gbps	80 Gbps	1.92 Tbps	1.92 Tbps	1.92 Tbps
Performance per Slot	60 Mpps	60 Mpps			
1-Gbps ports per card	48	48	48	48	48
1-Gbps ports per chassis with two supervisors	336	528	192	384	768
10-Gbps ports per card (non-blocking)	16 (8)	16 (8)	64	64	64
10-Gbps ports per chassis (non-blocking) with two supervisors	116 (32)	180 (48)	576 (576)	1152 (1152)	2308 (2308)
40-Gbps ports per card	4 (2)	4 (2)	36 (36)	36 (36)	36 (36)
40-Gbps ports chassis (non-blocking) with redundant supervisors	28 (14)	44 (22)	144 (144)	288 (288)	576 (576)
Port-to-port latency (between line cards)	Less than 20 ms	Less than 20 ms	Less than 4 ms	Less than 4 ms	Less than 4 ms

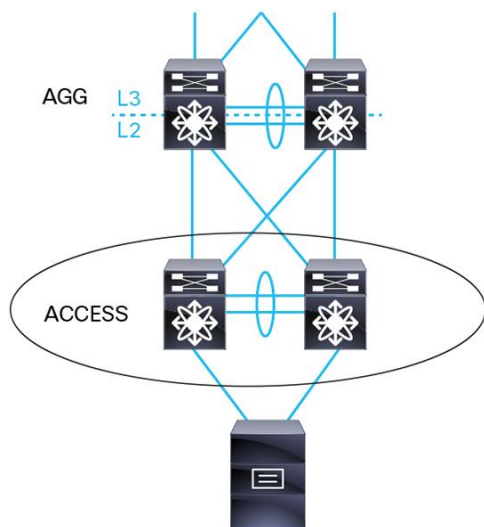
Migrating EoR Cisco Catalyst 6500 Systems to a Cisco Nexus 9000 System

Migration from a Cisco Catalyst 6500 chassis to a Cisco Nexus chassis involves the following high-level steps:

- Plan, design, and build the Cisco Nexus switching system
- Interconnect the current 6500 switching system and Cisco Nexus switching system
- Migrate end-device connection points to the Cisco Nexus switching system

Figure 4 displays a pair of access-layer EoR Cisco Catalyst 6500 switches to be migrated.

Figure 4. Migration of Catalyst 6500 Switches



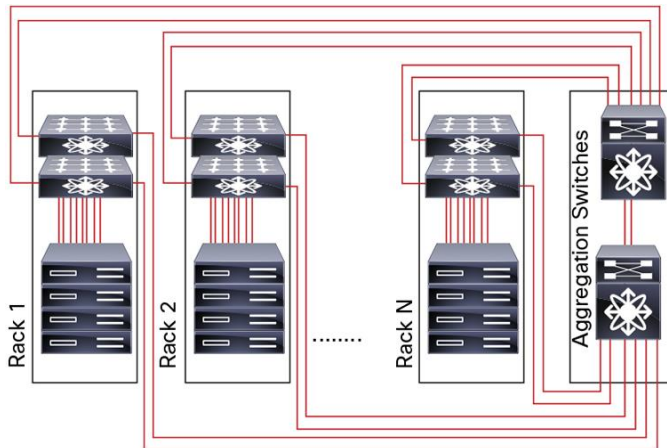
In the Cisco Nexus 9000 [migration guide](#) a pair of access-layer EoR Cisco Catalyst 6500 series switches are migrated to the Cisco Nexus 9500 platform. The guide is intended for implementation engineers and provides a migration process utilizing a Data Center Switching System (DCSS) migration methodology in order to significantly reduce the risk of failure of mission-critical applications during the migration.

Cisco Catalyst 6500 in the Data Center Aggregation Layer

The Cisco Catalyst 6500 is also widely deployed in the aggregation layer to provide both an interconnection point to the core and other pods. As 10 GE hosts create more bandwidth, the use of 40 GE links to connect to the core and access layers is becoming widely adopted and is recommended. 100 GE should also be in the short-term planning for many data centers. 100 GE will maintain the existing bandwidth hierarchy (one GE in access and 10 GE in aggregation, compared to 10 GE in access and 100 GE in aggregation).

Figure 5 shows how the Catalyst 6500 Series resides in the aggregation layer.

Figure 5. Cisco Catalyst 6500 Series in the Aggregation Layer



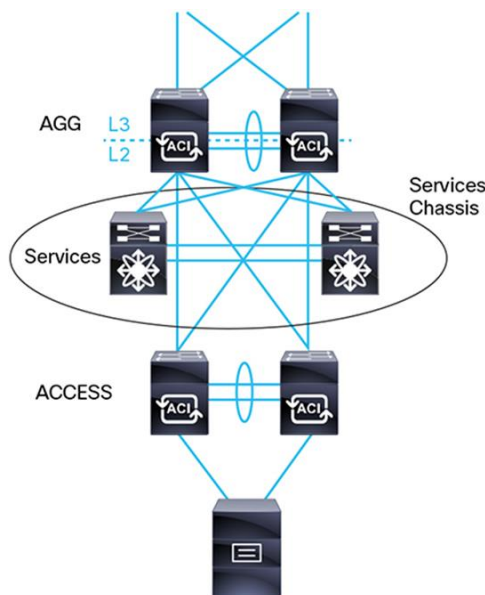
Layer 4 - 7 Service Insertion

Many Cisco Catalyst 6500 switches have service modules to provide integrated services such as firewalling and load balancing. These services are still required in the majority of data center designs. An existing Cisco Catalyst 6500 chassis could continue to be used in the services layer to house the integrated services modules and provide firewalling and load balancing services.

The freed up slots could be repurposed for additional modules or to consolidate services in a centralized services layer. The Cisco Catalyst 6500 chassis are often repurposed in a centralized services layer to provide connectivity to existing service modules such as the Network Analysis Module (NAM-3) and ASA Service Module (ASA-SM).

Figure 6 illustrates how network services are integrated into the architecture through a services layer utilizing Cisco Catalyst 6500 Series service modules. The empty line card slots on the Catalyst 6500 can now be utilized for additional service modules, which otherwise would have required additional chassis in the aggregation layer.

Figure 6. Network Services Integrated into the Architecture



How the Cisco Nexus 9000 Meets Current and Future Organizational Application Needs

Designed for Open Programmability

Many IT organizations are concerned that the network and layer 4 - 7 service deployment is a bottleneck when rolling out new applications. The Cisco Nexus 9000 provides an open, representational state transfer (REST) API to make the network easier to configure and manage. Using standard programming and scripting capabilities, IT organizations use automation to orchestrate application deployments. “Dev/Ops” types of IT organizations that use Agile development methods to facilitate the rapid development and deployment of applications can utilize the open programmatic interfaces of the Cisco Nexus 9000 to deploy these applications onto the network. Refer to the resources links at the end of this paper to learn more about “Dev/Ops” and rapid deployment.

In addition to traditional network management features such as Simple Network Management Protocol (SNMP), Syslog, Remote Monitoring (RMON), and Network Configuration Protocol (NETCONF), Cisco NX-OS Software on the Cisco Nexus 9000 also supports advanced programmability features. This automation of the network configuration allows organizations to rapidly deploy applications and eliminate human configuration errors. Cisco Nexus 9000 programmability features include:

- NX-API to support rapid development of automation tools using API calls
- Python Scripting, which can be utilized in both interactive and non-interactive modes
- Bourne-Again SHell (Bash) access to the underlying Linux system

Designed with a Modular Operating System

Cisco NX-OS is a modular operating system built specifically for the requirements of the data center environment. This OS is built on a 64-bit Linux kernel with isolated processes to provide resilience and modularity. Each service within Cisco NX-OS is designed as a non-kernel space (user space) process that performs a function and has its own virtual address space.

Processes are started and system resources are allocated only when a feature is enabled. Each service instance is run as a separate, independently protected process and can be restarted independently of the other processes. This stateful process restart helps to ensure non-stop operation. This unique capability can enable complete image upgrades without impacting data forwarding.

Our In-Service Software upgrade (ISSU) feature is scheduled to be supported in the second quarter of calendar year 2014 (Q2CY14). Redundant supervisors provide 1+1 redundancy and can perform a non-disruptive supervisor switchover (SSO) in critical failure situations. State and configuration information is constantly synchronized between the supervisors.

The Cisco Nexus 9000 NX-OS allows software maintenance upgrade (SMU) patching (available in Q2CY2014), helping to enable the existing NX-OS to be patched at a binary level. This allows for small fixes without the need for an entirely new image.

The Cisco Nexus 9000 NX-OS will also support both hot and cold patching. In addition, the Cisco Nexus 9500 and 9300 (and Cisco Nexus 3000) will all run from a single image so that customers only have to download and manage a single package from cisco.com for the 9000 and 3000 platforms.

Designed for Hardware Resiliency

Cisco Nexus 9000 switches were designed for reliability. There is no single point of failure in the hardware, including supervisor redundancy, fabric module redundancy, and power supply redundancy.

Designed for Future Scale

The Cisco Nexus 9508 Switch can provide 512 non-blocking 10 GE ports. Each 36 40 Gb line card has 1.44 Tbps of full-duplex connectivity to the fabric, providing full line-rate performance on all ports for all packets sizes.

Consider a Cisco Catalyst 6509 Switch with Supervisor Engine T2. The maximum connection speed to any one line card is 80 Gbps, which means that the entire Cisco Catalyst 6509 chassis (two supervisors and seven 16-port 10 GE cards) can only have 64 10- GE ports without oversubscription. With Supervisor Engine 720, that number would be reduced by half to 32 10 GE ports without oversubscription. This is a best-case scenario because the backplane is not the only limiting factor.

The Cisco Catalyst 6500 has a maximum forwarding capacity of 60 Mpps, therefore the number of ports that could operate at full line rate with small frames would be less the 64 listed above. For example, 64-byte frames multiplied by 60 Mpps is approximately 28.6 Gbps. The layer 2 and 3 switching capacity of the Cisco Nexus 9500 Series is 4.32 Bpps, or approximately 72 times the switching capacity on each line card.

The Cisco Nexus 9000 Series Switches have full support for copper and fiber transceivers, including the new Cisco QSFP BiDi transceiver, which is a short-reach optical transceiver that delivers 40 Gbps over a duplex OM3 or OM4 multimode fiber (MMF). It is MSA-compliant. The Cisco QSFP BiDi transceiver addresses the challenges of fiber infrastructure by providing the capability to transmit full-duplex 40 Gbps traffic over one pair of MMF cables with line card connectors.

Figure 7 is an image of the BiDi Optical Transceiver, providing 40 GE at near 10 GE prices using the same MMF cables.

Figure 7. BiDi Optical Transceiver



With the Cisco QSFP BiDi transceiver, the amount of cabling to support 40 GE is reduced to one pair of fiber strands. This is far more manageable than the existing industry standard defined by the IEEE high-speed study group, which requires four pairs of fiber for 40 GE. In another words, the Cisco QSFP BiDi transceiver allows 40-GE connectivity to reuse the existing, directly connecting 10-Gbps fibers and the existing fiber trunk without the need to add any fibers. Cisco QSFP BiDi transceivers are supported in the Cisco Nexus 9000 Series, as well as all other Cisco products that take advantage of the QFSP+ form factors.

State-of-the-Art Mechanical Design

The Cisco Nexus 9500 platform is designed for power and cooling efficiency. Power and cooling are a significant portion of a data center's operating expenses, leading many customers to make their data centers more efficient and less costly to operate. The Cisco Nexus 9000 Series power supplies are more than 90 percent efficient at 20, 50, and 100 percent of load (platinum rated), providing industry's lowest watts per port.

Air recirculation and mixing of hot air and cold air can cause temperatures to rise significantly, lowering the availability of IT equipment. Many customers purchase special racks or airflow containment kits to eliminate the mixing of hot and cold air. Cisco Nexus 9000 switches have front-to-back airflow, which helps to keep hot and cold air separated. Additionally, Cisco Nexus 9300 switches have options allowing either front-to-back or back-to-front airflow.

The Cisco Nexus 9500 Series offers innovation in that it does not have a mid-plane. Therefore, line cards connect directly to fabric modules. Without a mid-plane restricting airflow and causing hot spots, the Cisco Nexus 9500 Series Switches provide optimal front-to-back airflow and allow switches to operate using less power.

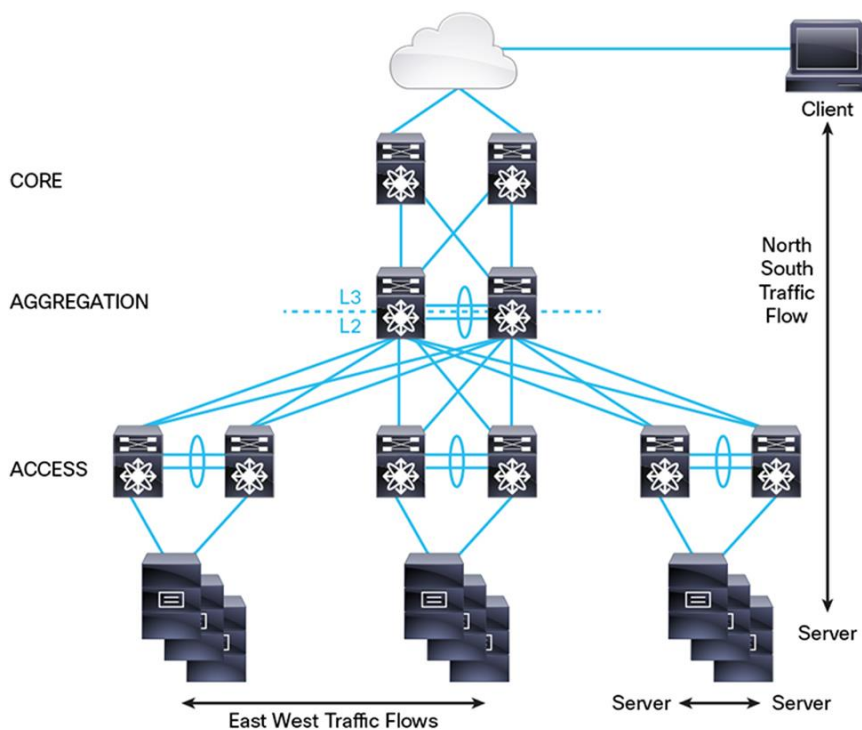
Data Center Fabrics

According to Gartner, “Traffic patterns in the data center are changing from being predominantly client/server to a significant level of server-to-server flows. By 2014, network planners should expect more than 80% of traffic in the data center network to be between servers.”

<https://www.gartner.com/doc/1656316/data-center-network-heading-traffic>.

Figure 8 displays east-west to north-south traffic flows. Services and applications creating east-west traffic flows include SAP HANA, Hadoop, Oracle RAC, server virtualization, desktop virtualization, SharePoint, storage replication, and other tiered applications.

Figure 8. East-West and North-South Traffic



If the majority of the traffic flow is (or will be) east-west then it makes sense to design the infrastructure for the expected usage pattern. The Cisco Nexus 9000 provides a migration path to a data center network fabric. A data center network fabric provides low and predictable any-to-any latency as well as linear scalability. A data center network fabric is typically a spine-leaf design.

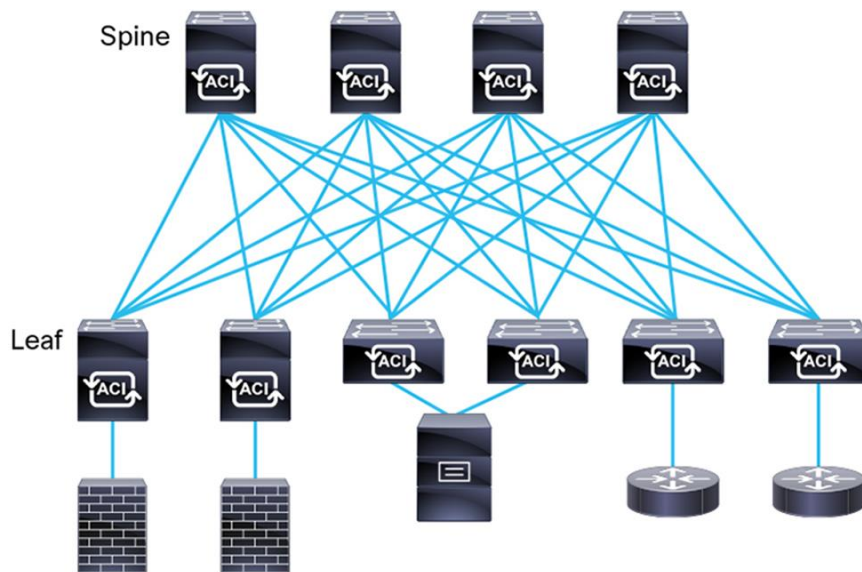
A spine-leaf architecture is foundational to building a highly virtualized multiservice data center. Such fabrics are typically a prerequisite for Software-Defined Networking (SDN).

A two-tier leaf-spine fabric (Figure 9) provides any-to-any connectivity with predictable bandwidth, increased throughput, and consistent low latency. Every leaf switch is connected to every spine switch, providing equidistant behavior. Traffic with the source and destination on the same leaf is handled locally, and all other traffic travels from the ingress leaf to the egress leaf through a spine switch.

Network traffic is spread across multiple physical spines because Equal Cost Multi-Path (ECMP) is enabled across the fabric. The network is simplified by flattening the network architecture and using low-latency non-blocking; performance is not only increased but consistent and predictable.

The Cisco Nexus 9000 allows for the creation of this fabric spine-leaf architecture, utilizing the industry default standard, virtual LAN (VxLAN) technology. This translates into a more scalable, flatter, dynamic, and highly available Ethernet fabric.

Figure 9. Two Tier Leaf-Spine Fabric



The Cisco Nexus 9000 Series Switches Will Have Two Modes of Operation

The Cisco Nexus 9000 Series Switches support two modes of operation: NX-OS standalone mode and Application Centric Infrastructure (ACI) fabric mode. In standalone mode, the switch performs as a typical Cisco Nexus switch with increased port density, low latency and 40 Gb connectivity. In fabric mode, the administrator can take advantage of Cisco ACI (ACI is discussed in greater detail later in this paper).

NX-OS Standalone Mode

NX-OS mode can be used in most designs as it shares the same routing and switching functions as the rest of the portfolio of Cisco Nexus switches. The Cisco Nexus 9000 supports spine-leaf design in NX-OS mode, reducing risk and increasing flexibility in the data center. VxLAN can enable the extension of layer 2 networks over a shared underlying network infrastructure. The layer 2 and 3 boundary is implemented in the access layer, limiting spanning tree protocol to a single rack in a ToR design. VxLAN encapsulation and de-encapsulation is supported in the Cisco Nexus 9000 hardware, providing VxLAN gateway and VxLAN tunnel endpoint (VTEP) advantages without sacrificing performance.

ACI Fabric Mode

In the third quarter of calendar 2014 a software upgrade will be available to evolve the Cisco Nexus 9000 to Application Centric Infrastructure (ACI) mode. ACI is a new paradigm in networking that combines the performance benefits of hardware forwarding with the flexibility of software. This is key as customers deploy mixed environments of hypervisors (VMWare, Microsoft, OpenSource) and physical servers. ACI takes full advantage of the concept of application network profiles (ANPs) to define the connections between endpoints and the layer 4 - 7 services they require to support application deployments today. For more on Cisco Application Centric Infrastructure, visit <http://www.cisco.com/go/acj>.

Conclusion

The Cisco Nexus 9000 Family is a performance, port-density, power-efficiency, and programmability/automation game changer in the data center networking industry. It provides the foundation to deploy next-generation applications and prepare for the future needs of the organization.

The platform is built with the knowledge, experience, and foundations of Cisco award-winning Nexus hardware and software platforms. With an architecture designed for bare-metal servers and multiple different hypervisors, the Cisco Nexus 9000 is the platform of choice for enterprises planning to migrate from the trusted Cisco Catalyst 6500 platform to a next-generation data center switching platform. Start a migration conversation with your Cisco and partner account teams.

Learn more about the [Cisco Technology Migration Program \(TMP\) - "Just Switch It" promotion](#).

Get Started

No two data center networks are ever the same, so understanding your customers' readiness and current environment is critical. To get your customers started, Cisco Advanced Services offers comprehensive services, including:

- Network assessments
- Cloud design and deployment
- Proof of concept planning, and testing
- Migration strategy and planning
- Application dependency mapping

Cisco technical services descriptions can be found at <http://www.cisco.com/go/servicedescriptions/>. Engage a Cisco Services account manager for additional information.

Tools and Services

Cisco Nexus Migration Tool

The Cisco Nexus Migration Tool (NXMT) can be used to convert Cisco IOS® Software commands or configurations from the Cisco Catalyst platform to corresponding NX-OS commands or configurations for the Cisco Nexus platform.

<https://tools.cisco.com/nxmt>

References

Cisco Nexus 9000 Series Switches:

<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html>

Cisco Nexus 9000 Configuration Guide:

<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-installation-and-configuration-guides-list.html>

Cisco Nexus 9000 Migration Guide:

<http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/guide-c07-730158.html>

Cisco QSFP 40-Gbps Bidirectional Short-Reach Transceiver Data Sheet:

<http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-730160.html>

The Cisco QSFP 40-Gbps BiDi transceiver is supported on Cisco switches and routers. For more details, refer to the [Cisco 40 Gigabit Ethernet Transceiver Modules Compatibility Matrix](#) document.

DevOps Organizations:

<http://www.gartner.com/it-glossary/devops>

<http://newrelic.com/devops/what-is-devops>

Data Center Overlays White Paper:

<http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps13386/white-paper-c11-730116.html>

VXLAN Overview: Cisco Nexus 9000 Series Switches:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps13386/white-paper-c11-729383_ns1261_Networking_Solutions_White_Paper.html

Data Center Service Integration: Service Chassis Design Guide:

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/dc_servchas/service-chassis_design.html

Go to <http://www.cisco.com/go/powercalculator> for easy power consumption calculation.



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)