

Cisco Nexus 5000 and Cisco Nexus 2000: Resource Optimization for Small and Mid-size Data Centers

Executive Summary

As all data centers continue to evolve their architecture to better support the adoption of new applications, demands, and operational limits, their infrastructure administrators are pursuing two main principles: better use of resources and simplification. Both are aimed to reduce the total cost of ownership (TCO) of the acquired solutions, including application software, servers, storage, network, and facilities.

These requirements tend to be harder on small and midsize data centers that usually have limited operations personnel, and therefore must rely more on the easiness and flexibility of the deployed solutions.

With the introduction of the Cisco Nexus® series of data center switches, Cisco brought a few revolutionary concepts to optimize these kinds of environments: **fabric extenders** with the Cisco Nexus 2000 Series, which allow architects to have significant cable reduction and centralized control-plane management in their data centers; and **converged fabric**, through the use of Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) to drastically reduce server cabling and provisioning operations through a “wire once and walk away” approach.

Recently the next generation of Cisco Nexus 5000 Series Switches, the Cisco Nexus 5500 platform, extended its applicability with Layer 3 capabilities, making it possible for small and midsize data centers with up to 100 physical servers to build a new, simplified network architecture. This document describes the main advantages of this architecture in these environments.

Cisco Nexus 5500 Switches

The Cisco Nexus 5500 Switches continue the evolution of the Cisco Nexus Family of data center switches. The Cisco Nexus 5548UP Switch offers - in a single rack unit - 32 fixed Unified Ports that can function as 10 Gigabit Ethernet Small Form-Factor Pluggable Plus (SFP+) ports with DCB and FCoE capabilities, Gigabit Ethernet ports or native 1-, 2-, 4-, and 8-Gbps Fibre Channel (depending on the inserted optics) ports with one expansion slot for added flexibility (Figure 1). The currently available expansion modules for this slot include:

- Sixteen-port 10 Gigabit Ethernet interfaces with DCB and FCoE capabilities (that can work as Gigabit Ethernet as well)
- Eight-port 10 Gigabit Ethernet interfaces with DCB and FCoE capabilities (that can work as Gigabit Ethernet as well) plus 8-port native 1-, 2-, 4-, and 8-Gbps Fibre Channel interfaces to allow servers to access external storage devices
- Sixteen unified ports that can function as Gigabit Ethernet, 10 Gigabit Ethernet with DCB and FCoE capabilities, or native 1-, 2-, 4-, and 8-Gbps Fibre Channel (depending on the inserted optics) interfaces

Figure 1. Cisco Nexus 5548UP Switch



Layer 3 Module for Cisco Nexus 5548P Switch

The recently introduced Layer 3 functions in the Cisco Nexus 5548UP expand the applicability for this switch. In fact, this field-upgradable hardware module (which does not consume the expansion slot) permits the Cisco Nexus 5548UP to act like a true “distributed modular system” with its scalability associated with the number of associated Cisco Nexus 2000 Series Fabric Extenders.

At first customer shipment (FCS) the Cisco Nexus 5548UP with Layer 3 module and license will be able to implement:

- Default gateway with redundancy protocols such as Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP)
- Layer 3 routed ports on its interfaces for connections to a WAN router or campus core switches
- Static routing and routing protocols such as Routing Information Protocol Version 2 (RIPv2), Open Shortest Path First Version 2 (OSPFv2), Enhanced IGRP (EIGRP), and Border Gateway Protocol (BGP)
- Multicast protocols such as Interior Group Management Protocol (IGMP) (all versions), Protocol Independent Multicast Version 2 (PIMv2) sparse mode, and Multicast Source Discovery Protocol (MSDP)
- Access control lists (ACLs) to enforce the security police between hosts
- Virtual Route Forwarding lite (VRF-lite), which permits multiple independent routing instances (for different applications or customers) inside the same switch

More details about the Cisco Nexus 5548UP Switch is available at:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-618603.html.

The Cisco Nexus 5548P also has a low-latency cut-through design that provides predictable, consistent traffic performance regardless of packet size, traffic pattern, or enabled features on 10 Gigabit Ethernet interfaces with less than 2 microseconds for switched traffic and 4.8 microseconds for routed traffic.

Further scalability, with the same features, can be accomplished with the Cisco Nexus 5596UP Switch, which can provide a density of 96 Gigabit Ethernet or 10 Gigabit Ethernet connectivity ports in 2 rack units (2RU).

Cisco Nexus 2000 Series Fabric Extenders

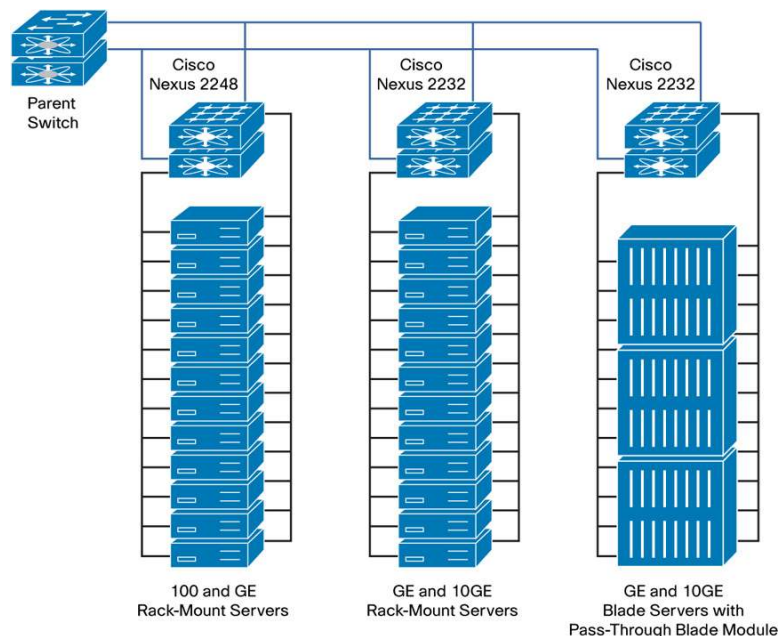
The Cisco Nexus 2000 Series Fabric Extenders (FEX) increase the connectivity and configuration options of data centers because they simplify their architecture and operations to meet the customer's business and application needs (Figure 2). Through the connection to a parent switch such as the Cisco Nexus 5548UP or Cisco Nexus 5596UP, the Cisco Nexus 2000 Series permits the deployment of a highly scalable server-access network with connectivity options such as Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet.

Figure 2. Cisco Nexus 2000 Series Fabric Extenders



Acting like a “remote line card”, the Cisco Nexus 2000 extends its parent switch capabilities directly into the servers rack and forms a distributed virtual switch that can drastically reduce the data center cabling and allow the administration of thousands of ports under the same control plane (that is, the parent switch command-line interface). This setup is illustrated in Figure 3.

Figure 3. Cisco Nexus 2000 Series Fabric Extender Architecture



The Cisco Nexus 2000 Series provides two types of ports: server attachment ports (host interfaces) and uplink ports (fabric interfaces), which are used for connectivity to the upstream parent Cisco Nexus switch. Table 1 lists the specifications of the Cisco Nexus 2000 Series Fabric Extenders.

Table 1. Specifications of Cisco Nexus 2000 Series Fabric Extenders

Description	Specification
Cisco Nexus 2224TP	24 Fast Ethernet or Gigabit Ethernet host interfaces and two 10 Gigabit Ethernet fabric interfaces (SFP+)
Cisco Nexus 2248TP	48 Fast Ethernet or Gigabit Ethernet host interfaces and four 10 Gigabit Ethernet fabric interfaces (SFP+)
Cisco Nexus 2232PP	32 Gigabit Ethernet or 10 Gigabit Ethernet (FCoE-enabled) host interfaces (SFP+) and eight 10 Gigabit Ethernet fabric interfaces (SFP+)

Note: Fabric extenders can be mixed and matched to a parent switch to provide several connectivity options.


Design Principles with Cisco Nexus 5500 Switches and Cisco Nexus 2000 Series Fabric Extenders

The latest Cisco Nexus 5000 Series Switches, the Cisco Nexus 5500 platform with the Layer 3 module can manage up to eight Cisco Nexus 2000 Fabric Extenders independently of their models or the number of connected uplinks:

- Up to four uplinks per Cisco Nexus 2224TP and 2248TP
- Up to eight uplinks per Cisco Nexus 2232PP

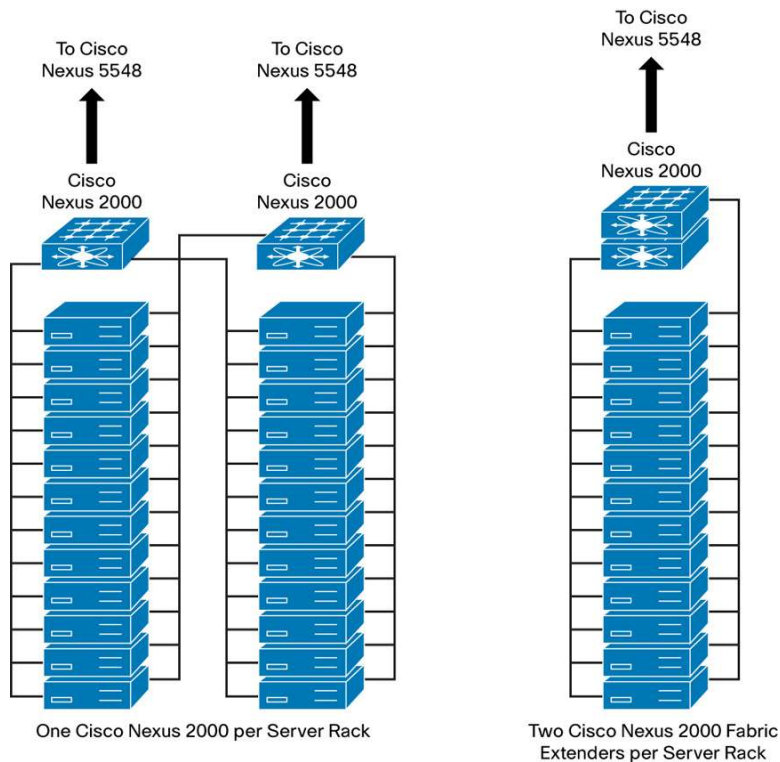
Both devices support Virtual PortChannel (vPC) technology, which permits servers directly connected to two Cisco Nexus 2000 Switches (2224TP, 2248TP, or 2232PP) associated to different parent switches to have up to 16 active interfaces, as shown in Figure 4.

Figure 4. vPC to the Server



The Cisco Nexus 2000 Series Switches can deliver significant cabling savings and operating flexibility by using the fabric extender in a distributed top-of-rack line-card design. Common data center designs usually accommodate one or two Cisco Nexus 2000 Switches in each rack (Figure 5).

Figure 5. Cisco Nexus 2000 Rack-Mount Installation Options




Several metrics must be accounted for in order for the customer to decide how many servers can be accommodated inside of a single rack, such as server physical characteristics (size in rack units, power, cooling, blade, or rack), rack power feeding, cooling capacity, and rack size.

Solution

The Cisco Nexus 5500 with Layer 3 capabilities and Cisco Nexus 2000 can provide a complete, scalable, robust, and highly available architecture for small and midsize data centers that have up to 100 physical servers.

As depicted in the architecture of Figure 6, a network infrastructure can be built on the Cisco Nexus 5500 and Cisco Nexus 2000 to provide organizations their own optimized data center network ready for new demands such as video and server virtualization.

Figure 6. Cisco Nexus 5500 and Cisco Nexus 2000 Architecture for Small and Midsize Data Centers



Using the Cisco Nexus 5500 and Cisco Nexus 2000 as a consolidated switching platform offers several advantages to data center managers, such as:

- Simplicity and flexibility
- High availability
- Scalability
- Embedded security
- Cost efficiency

Simplicity and Flexibility

The Cisco Nexus 5500 and Cisco Nexus 2000 solution offers a high degree of infrastructure simplification through the consolidation of Ethernet Layer 2 switching access, Layer 3 IP routing, and storage area network (SAN) into a single platform. This scalable solution will be compatible with future network solutions for a smooth transition between the different server access technologies, such as Fast Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet (for LAN) and 1-, 2-, 4-, or 8-Gbps Fibre Channel (for SAN).

As described, this solution allows the connected servers to access external storage devices through a great variety of methods (that is, Small Computer System Interface over IP [iSCSI], (Network File System [NFS], and Common Internet File System [CIFS]). More specifically, if the 10 Gigabit Ethernet-connected servers are equipped with Converged Network Adapters (CNAs), they can access Fibre Channel and FCoE storage devices through the same cabling. This simplified approach will work if the storage targets are directly attached to the Cisco Nexus 5500 interfaces or if they are connected to a SAN switch (that is, Cisco MDS 9000).

The use of a single management platform is also an advantage for a converged infrastructure like the one proposed here. Through the Cisco Nexus 5500 command-line interface (CLI) or the Cisco Data Center Network Manager (DCNM) graphical tool, the infrastructure administrators can manage both the LAN and the SAN infrastructure with the same GUI (Figure 7). Using the Role-Based Access Control (RBAC) feature in the Cisco Nexus 5500, you can create a network administrator role for network-related activities and a storage administrator role for Fibre Channel configurations - and you even can restrict their scope of configuration actions to some virtual LANs (VLANs) or virtual SANs (VSANs).

Figure 7. Cisco Data Center Network Manager




Applications that use multicast traffic (that is, video streaming and financial feeds) are also supported by the architecture because both Cisco Nexus 5548UP and Cisco Nexus 5596UP can detect IGMPv1, IGMPv2, and IGMPv3 join messages to announce multicast group members to the WAN through IGMPv2 sparse mode, and also restrict multicast traffic to the member's ports only.

The Cisco Nexus 5500 and Cisco Nexus 2000 solution provides traffic differentiation with quality of service (QoS); for example, eight hardware queues per port, policing, Layer 2 class-of-service (CoS) trust, strict-priority queuing, and Weighted Round Robin (WRR) scheduling.

The proposed design also allows the insertion of network services appliances such as firewalls (for example, Cisco ASA 5500 Series Adaptive Security Appliances) and application-delivery controllers (for example, Cisco ACE 4710 Application Control Engines). Enforcement of traffic by using VLANs (Figure 8) allows the deployment of both routed and bridged appliances.

Figure 8. Network Service Appliance Insertion



WAN accelerators such as the Cisco Wide Area Application Services (WAAS) can be deployed in the proposed data center design using two traffic interception methods: **Inline cards** (positioning the WAAS appliances between the edge routers and the Cisco Nexus 5500 switches); or **Web Cache Control Protocol version 2** (WCCPv2) configured on the WAAS appliances and edge routers. More details are found at:

http://www.cisco.com/en/US/prod/collateral/contnetw/ps5680/ps6870/prod_white_paper0900aecd806d976a_ps6474_Products_White_Paper.html

Appliances that support interface aggregation can take advantage of the high availability and bandwidth scaling provided by vPC technology.

Moreover, the solution supports Switched Port Analyzer (SPAN), which enriches management capabilities by providing data center managers the tool to have full visibility of the applications traffic.

High Availability

You can achieve outstanding high availability for the server connections with the following features and configurations:

- Redundant and hot-swappable power and fan modules available in every Cisco Nexus 5500 and Cisco Nexus 2000
- Operating system (Cisco Nexus Operating System [NX-OS]) modularity and serviceability
- In-Service Software Upgrade (ISSU) on Cisco Nexus 5500 (and consequently on all of its associated Cisco Nexus 2000 Fabric Extenders)
- Configuration rollback if undesirable change results
- Virtual PortChannel (vPC) to all the servers with dual-homed connections to two Cisco Nexus 2000 (or two Cisco Nexus 5500) providing sub-second failover recovery and higher bandwidth availability
- Server default gateway high availability through the use of first-hop redundancy protocols such as HSRP and VRRP
- Routing to outside IP networks with fast convergence from protocols such as OSPF, EIGRP, and BGP

Scalability

As shown in Figure 6, each Cisco Nexus 5548UP (equipped with a Layer 3 module) and its associated Cisco Nexus 2000 Fabric Extender constitute a distributed modular system with up to eight fabric extenders per Cisco Nexus 5500.

With this infrastructure, the data center administrators can use the Cisco Nexus 2224TP and Cisco Nexus 2248TP models to scale the number of Fast Ethernet and Gigabit Ethernet server interfaces under the same control plane. The number of 10 Gigabit Ethernet ports available for host connections can be also scaled, according to the customer needs, with the addition of Cisco Nexus 2232PP to the topology.

Security

This solution can provide a high level of security to the applications hosted on the connected servers through the use of the following features:

- Isolation and traffic control of hosts that belong to the same VLAN through the use of Private VLANs
- Protection from rogue Dynamic Host Configuration Protocol (DHCP) servers through DHCP snooping

-
- Authentication, authorization, and accounting (AAA) for identity verification, access granting, and action tracking of users managing Cisco Nexus 5500 Switches
 - Management account enforcement through a local database or external servers using RADIUS or TACACS+
 - Encrypted management protocols such as Secure Shell (SSH) Protocol, Simple Network Management Protocol Version 3 (SNMPv3), and Secure File Transfer Protocol (SFTP)
 - ACLs to control traffic through interfaces, inside or between VLANs

Cost Efficiency

The solution provides acquisition and operational cost reduction and efficiency:

- Ethernet and SAN consolidation in a single platform offers significant savings on hardware, cabling, and implementation services costs.
- For in-rack or adjacent-rack server cabling, the Cisco Nexus 5500 (and Cisco Nexus 2232PP) platforms support SFP+ direct-attach 10 Gigabit Ethernet copper, an innovative solution that integrates transceivers with Twinax cables into an energy-efficient, low-cost, and low-latency solution.
- Cisco Nexus 5500 and Cisco Nexus 2000 connection is possible with low-cost 10 Gigabit Ethernet Fabric Extender Transceivers (FET) or Twinax cables (up to 10 meters), and customers will save on the cost of expensive 10 Gigabit Ethernet optics.

Conclusion

The Cisco Nexus 5500 with its Layer 3 module and the Cisco Nexus 2000 Fabric Extenders bring new capabilities for small and midsize data centers. Acting as a distributed modular system, they offer the simplicity of management of hundreds of Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet ports and several advantages over other designs such as I/O consolidation, cable reduction, and provisioning simplicity.

When compared to conventional architectures, the proposed solution can drastically reduce the TCO of a data center infrastructure, minimizing the number of deployed devices, cables, and (provisioning and management) operations.

For more information about the products and solutions discussed in this document, please visit:

- <http://www.cisco.com/go/nexus5000>
- <http://www.cisco.com/go/nexus2000>
- <http://www.cisco.com/go/asa>
- <http://www.cisco.com/go/ace>
- <http://www.cisco.com/go/waas>
- <http://www.cisco.com/go/dcnm>
- vPC Design Guide: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/C07-572831-00_Dsgn_Nexus_vPC_DG.pdf



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 of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)

Printed in USA

C22-661448-00 06/11