


Rack-Level I/O Consolidation with Cisco Nexus 5000 Series Switches

Introduction


Best practices for I/O connectivity in today's data centers configure each server with redundant connections to each of two networks: two or more Ethernet network interfaces connect to IP networks, and two or more Fibre Channel Interfaces connect to storage networks. This approach has created a proliferation of cabling that imposes significant costs (Figure 1). Each network requires a dedicated set of switches, interfaces, transceivers, cabling, and the ongoing expense of managing, servicing, powering, and cooling the equipment needed to support them. There are hidden costs as well. To have sufficient expansion slot capacity to support all these interfaces, IT departments often must upgrade to larger, more expensive servers solely for their I/O capacity, expending capital and consuming valuable data center rack space.

Figure 1. Today's Data Center Best Practices Equip Each Server with a Redundant Pair of Ethernet and Fibre Channel Adapters, Resulting in a Proliferation of Access-Layer Cabling and Switching



Cisco Nexus™ 5000 Series Switches change this situation by delivering a unified network fabric that supports Ethernet, Fibre Channel, and even interprocess communication (IPC) networks over a single 10 Gigabit Ethernet/Fibre Channel over Ethernet (FCoE) link (Figure 2). Servers can be configured with a single pair of converged network adapters (CNAs), sometimes called multifunction adapters, that interface with both the LAN and storage stacks in the operating system. This new type of adapter does not require any change to LAN or SAN behavioral models. FCoE is compatible with Fibre Channel management models, configuration methods, and tools, making the transition to FCoE a quick and easy proposition.

Figure 2. Cisco Nexus 5000 Series Switches Consolidate Ethernet, Fibre Channel, and IPC Networks onto a Single Network Link, Simplifying Rack-Level Cabling and Switching




The benefits of this rack-level I/O consolidation include the following:

- Multiple NICs and HBAs can be consolidated into one CNA and use up to the 10-Gbps bandwidth limit per port. Reducing the number of adapters simplifies server configuration and conserves expansion slots. Total cost of ownership (TCO) savings can be significant; for more information, please see [Cisco Nexus 5000 Series Switches: Decrease Data Center Costs with Consolidated I/O](#).
- All server-to-access-switch cabling can be replaced with the low-cost, low-latency Small Form-Factor Pluggable Plus (SFP+) direct-attach 10 Gigabit copper solution, which eliminates costly fiber transceivers, large cable bundles, and typically 90 percent of the long runs to end-of-row Fibre Channel switches.
- The Cisco® Nexus 5000 Series Switches connect directly to data center storage networks through native Fibre Channel interfaces, protecting existing investments in storage equipment, management software, and staff training.
- Fewer upstream ports due to aggregation of adapter ports at the access switch means fewer modular SAN switch ports required.

This document explores the benefits of I/O consolidation in two deployment scenarios: high-density racks of two-rack-unit (2RU) servers, and racks of blade servers as they begin to support 10 Gigabit Ethernet.

I/O Consolidation with High-Density 2RU Server Pods

The Cisco Nexus 5020 Switch has the port density and the I/O consolidation capability to dramatically reduce the number of cables and switches needed to support racks filled with 2RU servers. Figure 3 illustrates one possible scenario: a pod composed of two racks, each with 20 2RU servers and a single Cisco Nexus 5020.

Figure 3. I/O Consolidation Scenario with Two 42RU Racks of 20 2RU Servers Each

Each server is configured with two CNAs that carry the server's LAN and SAN traffic over 10 Gigabit Ethernet network links. One CNA per server connects with the switch in its own rack; the other CNA connects to the switch in the second rack. CNAs can provide a pair of 10 Gigabit Ethernet interfaces, so if redundancy is not a concern, a single CNA per server with a cable to each of the two switches is another potential configuration. For cabling within the pod, the SFP+ direct-attach 10 Gigabit copper solution provides a low-cost, low-power, and low-latency alternative to fiber or 10GBASE-T copper. This solution integrates Twinax cabling with SFP+ transceivers to deliver only 0.25 microsecond of latency per link, while consuming only 0.1 watt (W) of power per transceiver, helping to lower both capital and operating costs.

Any available ports on the Cisco Nexus 5020 can be used as LAN uplinks; however, in this scenario the 40 fixed switch ports are all occupied by connections to servers. In this case, ports on expansion modules are used to provide the number of uplinks needed to establish a desired oversubscription ratio. Attaching to native Fibre Channel SANs requires 4-Gbps Fibre Channel ports that are available only on expansion modules.

Cisco offers three expansion module options: a 6-port 10 Gigabit Ethernet/FCoE expansion module; an 8-port 1-, 2-, or 4-Gbps Fibre Channel expansion module; and a 4-port 10 Gigabit Ethernet/FCoE module plus a 4-port 1-, 2-, or 4-Gbps Fibre Channel expansion module.

The most common scenario in this two-rack, 40-server configuration uses two of the Ethernet plus Fibre Channel expansion modules. This scenario would provide a total of eight Ethernet and eight Fibre Channel uplinks. If each port were populated, this would support an oversubscription ratio of 5:1 on each of the two networks. The four most feasible combinations of the two expansion modules are summarized in Table 1.

Table 1. Different Combinations of Expansion Modules Yield Different Oversubscription Ratios for the LAN and SAN

Expansion Module Combinations	Number of Ethernet Links and Oversubscription Ratio	Number of Fibre Channel Links and Oversubscription Ratio
2X 4-port 10 Gigabit Ethernet plus 4 Fibre Channel	8 5:1	8 5:1
1X 4-port 10 Gigabit Ethernet plus 4 Fibre Channel 1X 8-port Fibre Channel	4 10:1	12 3.33:1
1X 6-port 10 Gigabit Ethernet 1X 8-port Fibre Channel	6 6.67:1	8 5:1
1X 4-port 10 Gigabit Ethernet plus 4 Fibre Channel 1X 6-port 10 Gigabit Ethernet	10 4:1	4 10:1

Benefits of Consolidated I/O

An IT department can consolidate I/O in this way using a number of different methods:

- For servers already configured with compatible 10 Gigabit Ethernet NICs, a driver can be used to implement both IP and Fibre Channel stacks and pass both traffic streams over the same link to the access-layer switch.
- In existing Gigabit Ethernet environments, Gigabit Ethernet NICs are often used along with a pair of Fibre Channel HBAs; in these environments, Fibre Channel and several Gigabit Ethernet links can be consolidated onto a single pair of CNAs that use one pair of cables per server.
- In organizations rolling out new servers, the cost of deploying multiple 10 Gigabit Ethernet NICs and multiple Fibre Channel HBAs can be avoided by configuring servers from the start with CNAs to support redundant, consolidated I/O.

Regardless of the route by which I/O consolidation is achieved, the benefits are numerous:

- All in-rack I/O connections can be supported by low-cost, low-latency, and low-power SFP+ direct-attach 10 Gigabit copper cabling between servers and the top-of-rack switch. The need for more expensive fiber connections is limited to LAN connections to the aggregation layer, or Fibre Channel connections to the SAN core.
- The number of in-rack cables is reduced by two or more per server, which also reduces the number of adapters and transceivers along with their power draw and the load they place on the data center cooling infrastructure.
- With aggregation accomplished in server racks, or groups of server racks, only a small number of fiber connections need to extend to the aggregation layer. This feature helps reduce the overall number of switches, freeing valuable data center rack space while reducing capital and operating costs.

I/O Consolidation in Blade Server Racks

A growing number of options exist for 10 Gigabit Ethernet connectivity to blade servers. Today, some blade systems support internal switches that provide 10 Gigabit Ethernet connectivity to each blade and a small number of uplinks from the chassis. This situation is likely to change rapidly, with 10 Gigabit Ethernet appearing on blade server motherboards, and consolidated network adapters appearing in mezzanine-board form factors. The alternative approach is for network pass-through connections to deliver each blade's network interfaces to the chassis and then extended directly to the access-layer switch.


Network pass-through configurations are preferred over built-in switches because they allow all traffic to be managed by a uniform data center network infrastructure, eliminating the need for another, inconsistent layer of management in the data center network. Blade switches force fixed oversubscription ratios at the chassis level, and they also are not likely to provide the features needed to extend the unified network fabric to each blade; pass-through connections eliminate these concerns.

Hypothetical Blade Server with Ethernet Pass-Through Connections

Figure 4 shows a hypothetical blade server configuration using pass-through networking. Two blade systems are installed in a rack, each with 16 blades and two 10 Gigabit Ethernet links per blade extended to the chassis. I/O is consolidated on each blade in either of two ways:

- A mezzanine card can provide hardware-based consolidation just as CNAs provide in rack-mount servers.
- For blades with fixed 10 Gigabit Ethernet-capable Ethernet NICs, software drivers can provide both an Ethernet and a Fibre Channel stack to the OS, consolidating I/O and passing both traffic classes to the 10 Gigabit Ethernet link.

Figure 4. I/O Consolidation Scenario for a Hypothetical Blade Server with an Ethernet Pass-Through Configuration



In this example, 32 of the fixed Cisco Nexus 5020 ports on each switch are occupied with network links to blade systems, leaving 8 fixed ports available for network uplinks. This configuration yields an Ethernet network oversubscription ratio of 4:1. Note that this configuration allows full line-rate connectivity between all blades in the rack with no oversubscription in the rack.

Connectivity to the Fibre Channel SAN is provided by configuration of expansion modules for the Cisco Nexus 5020. A logical choice would be the addition of one 8-port 1-, 2-, or 4-Gbps Fibre Channel expansion card to the switch. This configuration would support Fibre Channel oversubscription ratios that can be adjusted from 4:1 to 16:1 based simply on the number of uplinks used on the 8-port module.

Conclusion

The unified network fabric supported by Cisco Nexus 5000 Series Switches dramatically simplifies rack-level networking by consolidating Ethernet and Fibre Channel traffic onto a single network link. I/O consolidation at the rack level helps reduce the number of adapters, cables, transceivers, and upstream ports required through the use of Cisco Nexus 5000 Series Switches in top-of-rack configurations. The cost savings potential makes I/O consolidation at the rack level a compelling business proposition; for more Information, see [Cisco Nexus 5000 Series Switches: Decrease Data Center Costs with Consolidated I/O](#).

For More Information

- <http://www.cisco.com/go/nexus>
- <http://www.cisco.com/go/dc>
- http://www.cisco.com/en/US/products/ps9670/prod_white_papers_list.html



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.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCI, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0903R)