Power and Cooling Savings with Unified Fabric

What You Will Learn

This document tells the story of one Cisco customer that compared the power and cooling costs of a discrete LAN and SAN design for its 1650-server data center expansion with a Cisco® Unified Fabric supported by Cisco Nexus® 5000 Series Switches. The customer achieved a 41 percent savings in power and cooling costs by implementing a Unified Fabric, amounting to a savings of US$75,114 per year. The Unified Fabric required only one-third the number of network adapters, and the reduced number of access-layer switches that were required freed up rack space equivalent to 172 servers.

Overview

The unified network fabric is now a reality, giving customers a greater range of choices for their data center networks. In the past, separate physical networks were required to handle each type of traffic using technologies including LAN and SAN, and interprocess communication (IPC) mechanisms. Today, Cisco Nexus 5000 Series Switches enable I/O consolidation at the rack level, allowing LAN, SAN, and IPC traffic to be carried over the same link between servers and the access layer, while using the same driver software, management software, and data center best practices for both LAN and SAN.

This application note describes the power and cooling savings that can be achieved with a unified network at the access layer compared to the costs for a discrete LAN and SAN design. It shares the calculations made by one Cisco customer while designing the network for a data center expansion to incorporate 1650 new servers. The original design equipped each server with interfaces and cabling for multiple LAN and SAN connections. When Cisco presented Unified Fabric as an alternative to a discrete LAN and SAN design, the customer calculated the power and cooling savings that it would realize through I/O consolidation at the rack level. The results made a compelling case for adopting Unified Fabric.

- The customer achieved a 41 percent savings in power and cooling costs for the consolidated network's access layer and SAN aggregation layer compared to the costs for a discrete LAN and SAN design. This savings amounts to US$75,114 per year for the 1650 servers and supporting infrastructure.
- The Unified Fabric required only one-third the number of network adapters, not only saving capital and operating expenses, but also eliminating multiple potential points of failure.
- The Unified Fabric required only one-third the amount of rack-level cabling and access ports, reducing the number of interconnects from nine per server to three.

Using Unified Fabric, the network infrastructure power consumption was reduced by 60,216 watts (W). With the customer's servers requiring 500W each, these savings are enough to power an additional 120 servers, or an increase of 7.2 percent. Most organizations would prefer to deploy more computing resources rather than more network equipment. The move to a Unified Fabric delivers more than energy savings: it brings all the benefits of 10 Gigabit Ethernet networking, including the following:

- The move from Gigabit Ethernet to 10 Gigabit Ethernet networking provides a 10X increase in bandwidth. Even when a single 10 Gigabit Ethernet link is used to replace several Gigabit Ethernet connections, the unified network leaves room for future growth. The capability to grow and adapt to rapidly changing business conditions is a strategic benefit that can help maintain a company's competitive edge.

- Unified Fabric supports a "wire once, use later" model in which every server is deployed with standard 10 Gigabit Ethernet and is enabled for LAN, SAN, and IPC protocols as needed through the Cisco Nexus 5000 Series switch configuration. Servers equipped for the unified network can be repurposed later without the need to recable racks or install new I/O adapters.
- A Unified Fabric has fewer points of failure, fewer elements requiring maintenance, and fewer chances for human error. All these factors contribute to increased reliability and availability.

Server Connectivity Requirements

The Cisco customer assessed the cost of deploying 1650 four-rack-unit (4RU) servers in its data center. The initial plan was to populate 165 racks with 10 servers each. End-of-row Gigabit Ethernet and Fibre Channel switches would provide the needed LAN and SAN connectivity.

Using a traditional LAN and SAN architecture, the customer requirements dictate a total of nine data connections to each server for the following purposes (Figure 1):

- Two 4-Gbps SAN connections through two Fibre Channel host bus adapters (HBA); the SAN must support a sustained data rate of 1.3 Gbps per server
- Two 1-Gbps Ethernet connections to the production network supported by two discrete network interface cards (NICs); this network must support a minimum of 1.2 Gbps continuous traffic
- Two 1-Gbps Ethernet connections to the backup network, also supported by two NICs
- One 1-Gbps Ethernet connection for the VMware Service Console, used to manage VMware ESX Server software used on each server; this connection is supported by one of the server's built-in LAN-on-motherboard (LOM) ports
- One 1-Gbps Ethernet connection for the VMkernel port
- One 10/100-Mbps Ethernet connection for the server's lights-out (remote) management functions; this connection is required regardless of whether I/O is unified, so this port, and the switch infrastructure to support it, is factored out of the power and cooling equations for each network design

Figure 1. Using a Traditional LAN and SAN Architecture, Each Server Would Require Nine Data Cables Supported by Six Discrete Interfaces
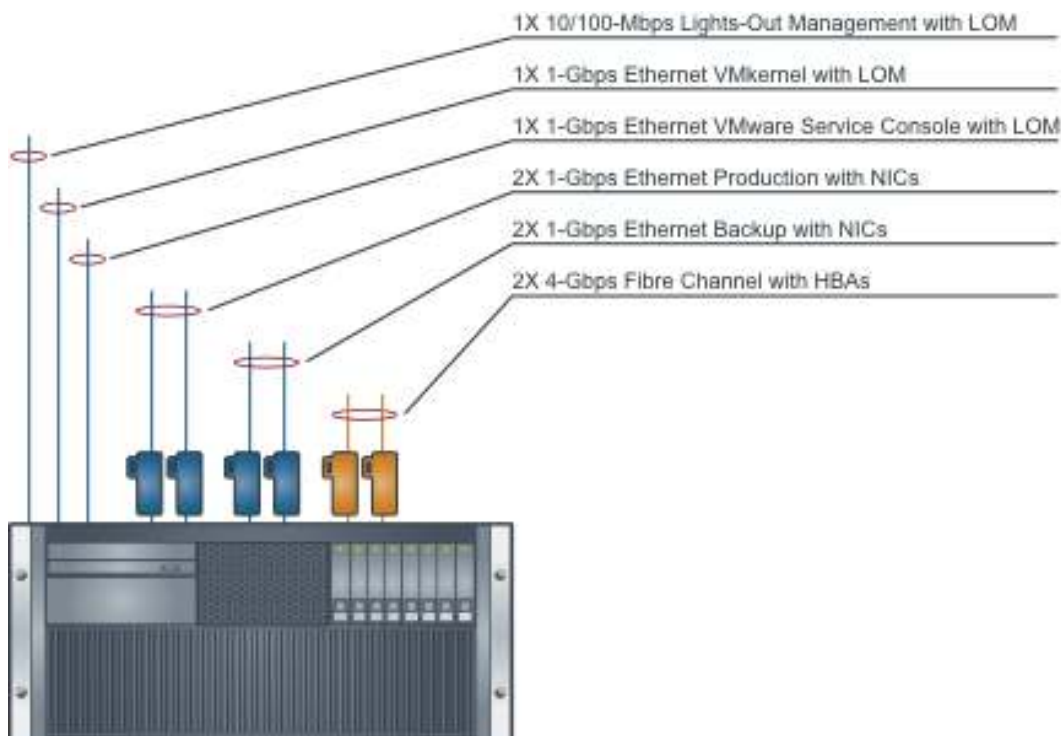

The following sections of this application note compare the details of the discrete LAN and SAN design to the Unified Fabric architecture, assessing the differences in power and cooling costs. The conclusion is evident at this point simply by comparing the two server configurations. With fewer server ports, less upstream network infrastructure is needed. 10 Gigabit Ethernet provides increased throughput and lower latency than Gigabit Ethernet, providing better performance today and room for future growth. The use of fewer cables, switches, and adapters contributes to a more reliable network that can be configured with a "wire once, use later" model.

The number of cables, NICs, and HBAs required by traditional LAN and SAN models hint at the cost and complexity of this approach. Nine cables per server must be installed and routed. NICs and HBAs must be configured, and upstream switch ports and capacity must be provisioned for each. If each NIC requires 3W and each HBA requires 5W, the cost to power even these low-power components amounts to 36.3 kilowatts (kW) when multiplied by the 1650 servers being deployed.

In contrast, Unified Fabric based on 10 Gigabit Ethernet exceeds each server's I/O requirement while carrying LAN and SAN traffic over a single network link. The server end of the Unified Fabric is supported by a converged network adapter (CNA) that presents both a 4-Gbps Fibre Channel HBA and a 10 Gigabit Ethernet NIC to the server operating system, making the existence of the Unified Fabric transparent: the server OS can use the same interfaces and management software, while the CNA merges both traffic types over the 10 Gigabit Ethernet link. When configured with two single-port CNAs for redundancy, each server now requires only three cables: two for the Unified Fabric and one for the server's lights-out management functions (Figure 2). Although today's CNAs use more power than the discrete components cited earlier, the second-generation CNAs used for this analysis consume only 5W according to data provided by the manufacturer.

Figure 2. Using a Unified Fabric (Instead of a Discrete LAN and SAN Design) Requires Only Three Cables and Two CNAs per Server
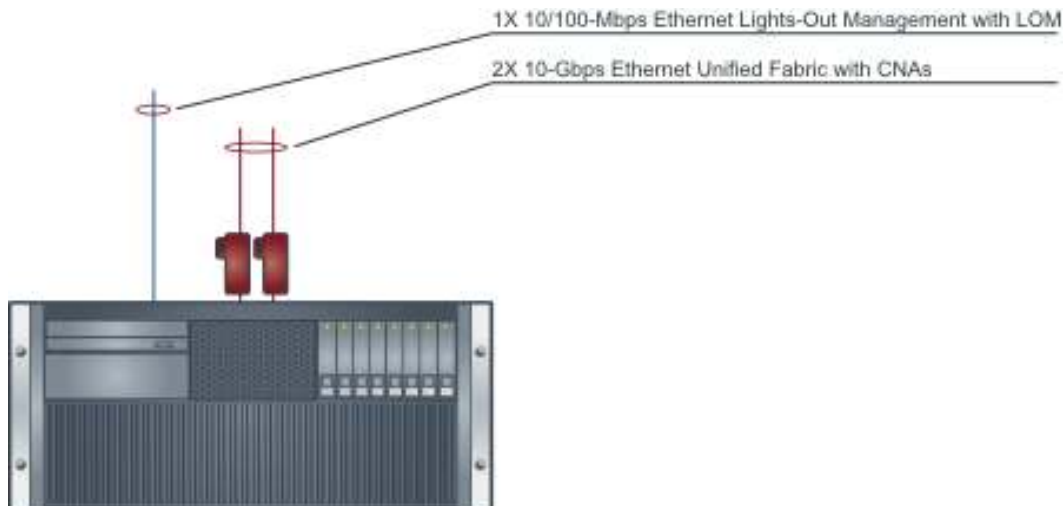
Discrete LAN and SAN Architecture

The traditional approach to supporting each server's I/O and networking requirements is to create separate, parallel LANs and SANs. Each network has its own access, aggregation, and core layers with a sufficient number of ports and upstream bandwidth to handle each server's multiple LAN and SAN cables.
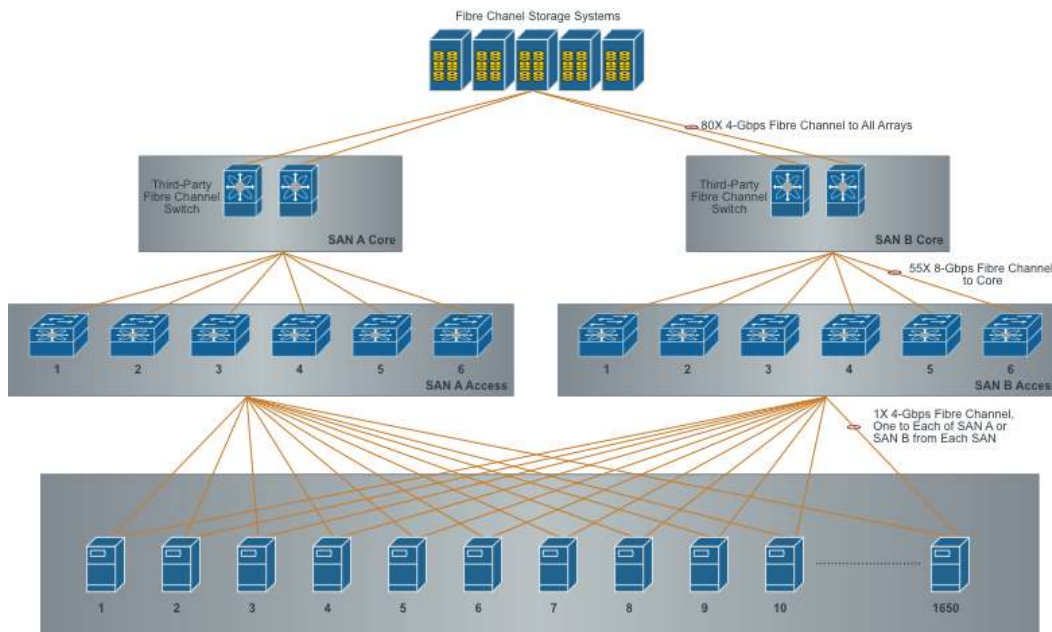
SAN Architecture

The customer's proposed SAN architecture provides connectivity between each server and five dual-ported Fibre Channel storage arrays through two independent SANs (Figure 3):

- Each server is equipped with a pair of 4-Gbps Fibre Channel HBAs. Each port uses a fiber connection to reach one of six third-party Fibre Channel switches in each SAN's access layer. One port connects to one of six SAN A access-layer switches; the other connects to one of six SAN B access-layer switches.
- Each of the 12 access-layer switches connects to the SAN core through 55 8-Gbps Fibre Channel links.
- The SAN core consists of four switches, two supporting each SAN. Each core switch connects to the customer's set of five Fibre Channel storage array through 80 4-Gbps Fibre Channel connections each.

The storage network supports an average of 1.55 Gbps sustained throughput per server port. It requires 16 Fibre Channel switches and a total of 4280 fiber cables.

Figure 3.   The Proposed Discrete SAN Architecture Uses 16 Third-Party Switches in Its Access and Core Layers
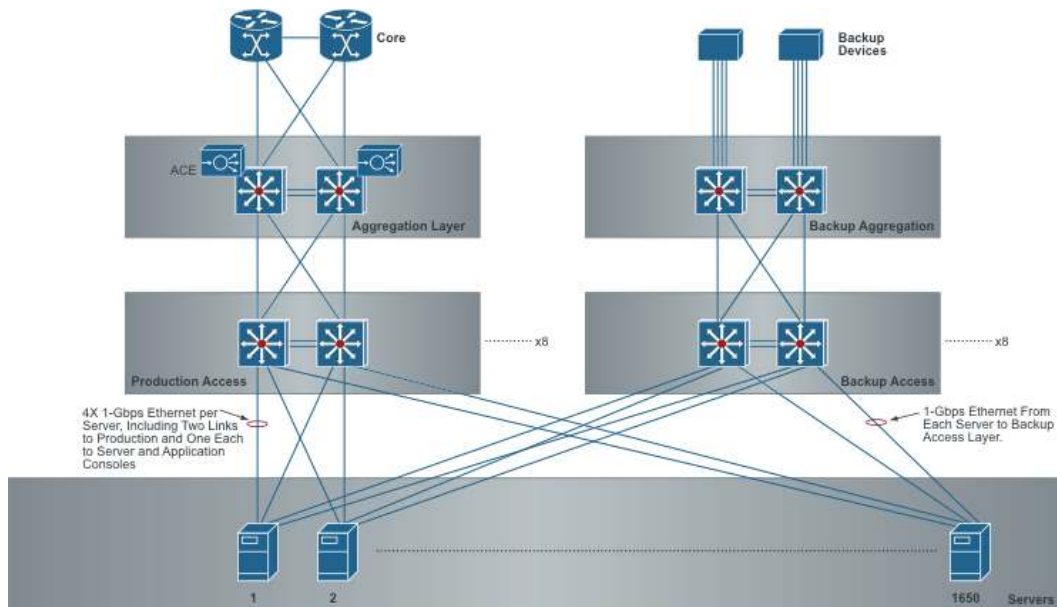

LAN Architecture

The proposed LAN architecture uses Cisco Catalyst® 6500 Series Switches to deliver Gigabit Ethernet connectivity throughout. Each server is configured with two NICs for the production network and two for the backup network, and LOM connections are used for server VMkernel and VMware Service Console connections. The supporting LAN is built to support independent storage and backup networks:

- Each server's two production LAN NICs connect to one of two access-layer Cisco Catalyst 6500 Series Switches in a pair.
- Each server's VMkernel and VMware Service Console ports connect to the access layer.
- Each server's two backup LAN NICs are connected to each of a pair of backup network access-layer switches.
- The production and backup network access layers each consist of eight pairs of interconnected Cisco Catalyst 6500 Series Switches.
- The production LAN aggregation layer is supported by a single pair of Cisco Catalyst 6500 Series Switches that are interconnected as peers. These switches are equipped with a Cisco ACE Application Control Engine Module. These switches connect directly to the LAN core.
- The backup LAN aggregation layer is supported by a single pair of Cisco Catalyst 6500 Series Switches that connect to backup devices.
- The server lights-out management ports are connected to Cisco Catalyst 3750 Series Switches (not shown).

The access and aggregation layer, excluding lights-out-management, uses a total of 34 Cisco Catalyst 6500 Series Switches.

Figure 4. The Proposed Discrete LAN Architecture Is Structured as Two Parallel Networks: One for Production and One for Backups

Unified Fabric Architecture

Cisco proposed an alternative architecture that uses a Unified Fabric to carry all LAN and SAN traffic from servers to Cisco Nexus 5020 Switches in the access layer (Figure 5). The Unified Fabric carries Fibre Channel traffic through Fibre Channel over Ethernet (FCoE), a straightforward, standards-based encapsulation of Fibre Channel into Ethernet. Both LAN and SAN traffic are carried over a common, Ethernet standards-based, Unified Fabric. These standards include IEEE Data Center Bridging that defines a set of extensions to the Ethernet that enhance the network's ability to carry multiple traffic streams over the same physical link.

Figure 5. The Unified Network Proposed by Cisco Uses Cisco Nexus 5020 Switches to Consolidate I/O at the Rack Level, Eliminating Six Cables per Server and the Need for Separate SAN Access-Layer Switches
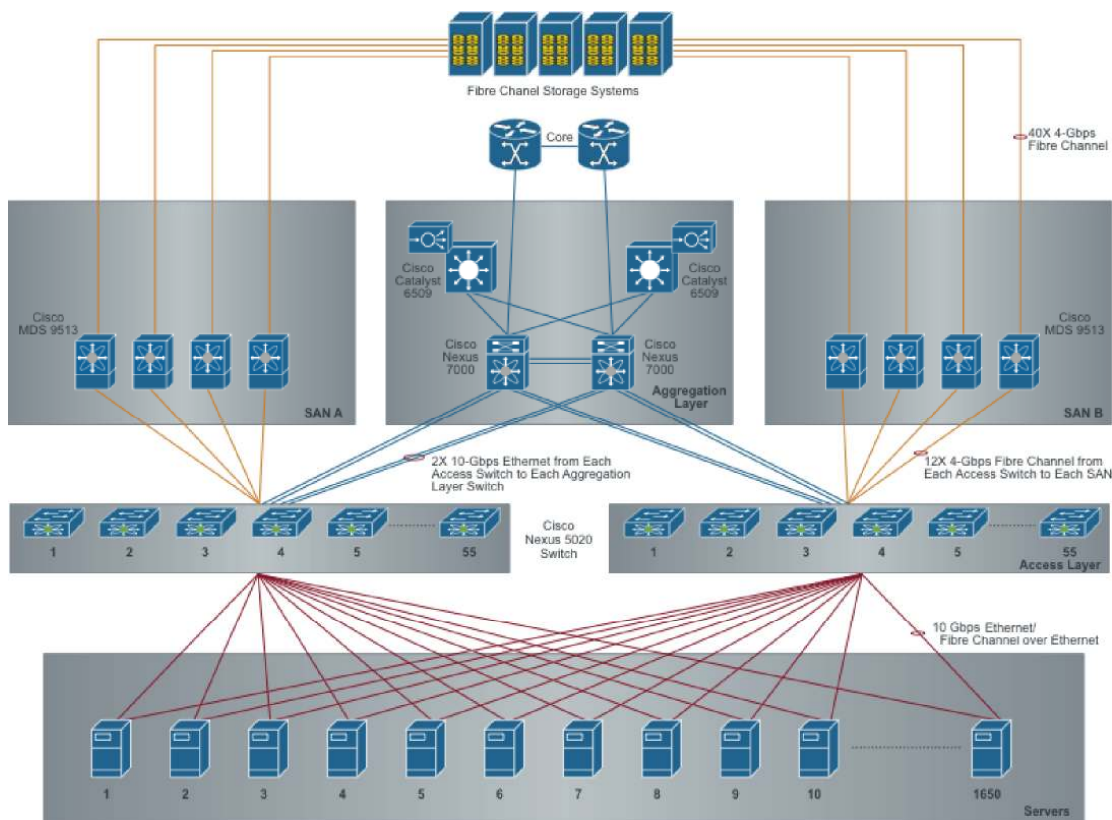
Simplified Server Configuration

The Unified Fabric simplifies each server's I/O configuration. The six NICs and HBAs are replaced by two single-port CNAs that support 10-Gigabit Ethernet and FCoE to the access-layer switches. What previously required a total of nine cables per server now needs only three. All I/O (except lights-out management) is carried over 10 Gigabit Ethernet links, boosting speed and leaving room for future growth in traffic.

Access Layer with Unified Fabric

The access layer is composed of two sets of 55 Cisco Nexus 5020 Switches that replace the entire SAN access layer and 32 of the Cisco Catalyst 6500 Series Switches required by the discrete LAN and SAN design. Cisco Nexus 5020 Switches provide low-latency 10 Gbps Ethernet and FCoE connectivity between servers and between servers and the aggregation layer. They accept FCoE traffic from the servers and connect through native Fibre Channel to the SAN aggregation layer. The switches are deployed in pairs to maintain the Fibre Channel connectivity model: each access-layer switch connects to only one of the two SANs.

Each switch is equipped with 40 fixed 10 Gigabit Ethernet and FCoE ports capable ports that is augmented with two expansion modules that provide Fibre Channel capability. One expansion module provides eight 4-Gbps Fibre Channel links through Small Form-Factor Pluggable (SFP) connectors. The other expansion module provides four 4-Gbps Fibre Channel links through SFP connectors and four 10-Gbps Ethernet links through SFP+ connectors. The switch connectivity is as follows:

- Each server connects to two Cisco Nexus 5020 Switches through 10 Gigabit direct-attached copper cables. This low-cost cabling solution integrates transceivers and Twinax cable for low power draw and low latency.

- Each switch connects to two aggregation-layer switches through two 10 Gigabit Ethernet fiber connections each. This configuration implements an oversubscription ratio of 7.5:1 that can support continuous traffic of 1.33 Gbps per server link between the access and aggregation layers. The port configuration of each switch leaves a total of ten 10 Gigabit Ethernet ports that can be used for future expansion.
- Each switch is equipped with 12 4-Gbps Fibre Channel connections to the SAN aggregation layer. Each of the two switches connected to a server connects to either SAN A or SAN B, with its 12 uplinks distributed across one SAN's four aggregation-layer switches. This configuration supports a 2.5:1 oversubscription ratio and up to 1.6 Gbps continuous bandwidth between each server link and the SAN aggregation layer.

LAN Aggregation Layer

The LAN aggregation layer is composed of two Cisco Nexus 7000 Series Switches, each of which is augmented by the Cisco Catalyst 6500 Series to support service modules. The customer uses Cisco ACE control modules.

SAN Aggregation Layer

The SAN aggregation layer is composed of two sets of four Cisco MDS 9513 Multilayer Directors. Each director accepts 165 4-Gbps Fibre Channel connections from the access layer and then connects to the customer's five storage systems through 40 4-Gbps Fibre Channel connections. This configuration brings the total connectivity to storage devices to 320 4-Gbps connections, or 1.28 terabits per second (Tbps).

Pod Physical Design

Cabling trade-offs suggest a top-of-rack location for access-layer switches rather than the end-of-row configuration that is typical in Gigabit Ethernet environments. While 10GBASE-T cabling can be used to support end-of-row configurations for 10 Gigabit environments, its latency and power draw is significantly higher than for the optical and copper solutions supported by Cisco Nexus 5000 Series Switches. The Cisco Nexus 5020 Switch supports a 10 Gigabit direct-attached copper cabling solution that is available in lengths up to 7 meters (m): ideal for server-to-switch connectivity in single or multiple-rack configurations. For longer runs, such as from the access layer to the aggregation layer, the switch supports multimode, short-reach optical fiber that can span up to 300m.

The benefits of the 10 Gigabit direct-attached copper cabling solution, plus the customer's server density, point to a pod-based rack configuration. Three 48RU racks of 10 servers each are collocated with the two access-layer switches that support them (Figure 6). Each server is connected to each of the two switches in the pod, keeping well within the 7m maximum cable length of the 10 Gigabit direct-attached copper solution. From each pod, 8 fiber cables link to the LAN aggregation layer, and 24 Fibre Channel connections link to the SAN aggregation layer. A total of 55 pods provide space for the 1650 servers to be deployed.

The Cisco Nexus 5000 Series Switches are designed for server rack deployment. They feature front-to-back cooling, with all serviceable components accessible from the front panel. All power and network connections are at the rear of the switch, located next to the server network interfaces that connect to them.

The top-of-rack switch configuration uses rack space that would otherwise be unoccupied by servers. This saves data center floor space by eliminating the need for several racks full of end-of-row switching equipment. In this customer's case, the pod configuration eliminated 480RUs, amounting to 10 rack positions, from the access-layer switch space alone (32 Cisco Catalyst 6500 Series Switches).

Figure 6.   The Pod-Based Model Uses Top-of-Rack Switches to Save 480RUs over the End-of-Row Model and Enables the Use of Low-Cost, Low-Latency, and Low-Power Cabling
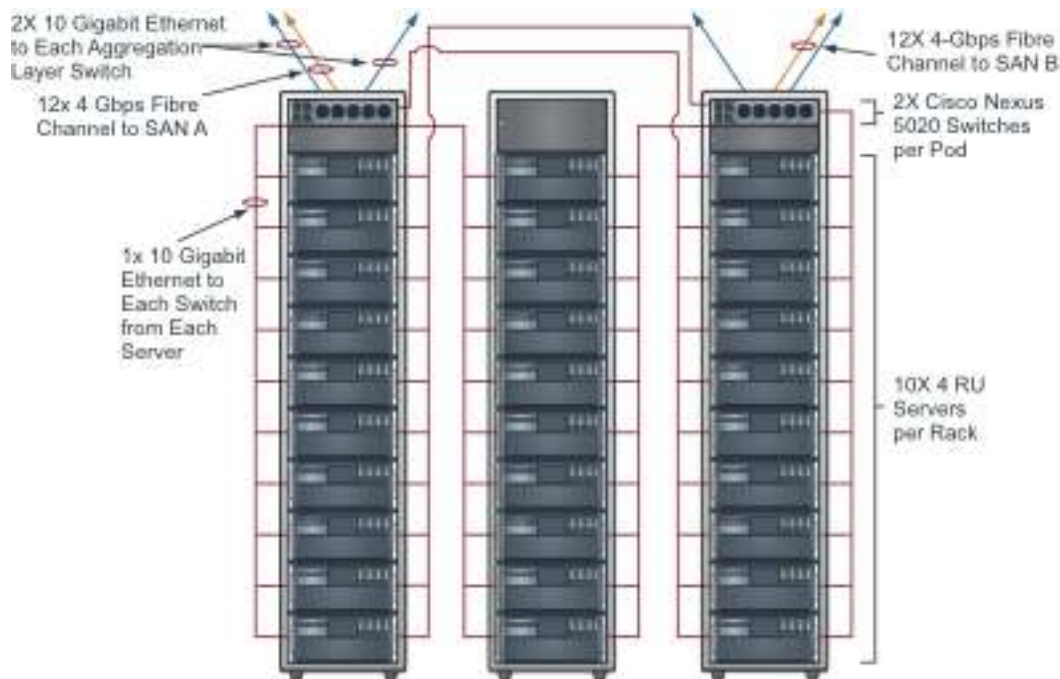
Unified Fabric Power and Cooling Savings

Examining the power and cooling savings for the Unified Fabric, the Cisco customer calculated a 41 percent savings when comparing the access layer and SAN aggregation layer equipment. The total amounted to US$75,115 per year, or US$375,575 over a 5-year period. These amounts were calculated using estimated power draws using vendor power calculators combined with the customer's power cost of US$0.0712 per kilowatt-hour (kWH). The customer focused on calculating the difference between the discrete and the unified networks. The results, summarized in Table 1, are as follows:

- The Unified Fabric eliminates all four NICs from each server. At 3W per NIC, this amounts to a savings of 19,800W.
- The Unified Fabric uses single-port, single–application-specific integrated circuit (ASIC), second-generation converged network adapters in place of Fibre Channel HBAs. A power estimate of 5W each (provided by the manufacturer) makes this an even exchange.
- For each server, six upstream network ports are no longer required, saving 9W per port. The ports eliminated include four ports that used discrete NICs and two ports that formerly used built-in ports. The power savings is based on use of Cisco Catalyst 6500 Series Switches for Gigabit Ethernet connectivity in the access layer. This change saves a total of 89,100W.
- The 10 Gigabit upstream ports, two for each server, are accounted for by the Cisco Nexus 5020 server's calculated power consumption of 480W. These switches add 52,800W to the network's power consumption.
- The Cisco Nexus 5020 Switches serving as both the LAN and SAN access layer allow 12 SAN edge switches to be eliminated. This, plus the change in power consumption between the third-party's SAN core design and the Cisco design using Cisco MDS 9513 Multilayer Directors, adds 4116W to the savings.

Table 1. Calculating Power Savings on a Component-by-Component Basis Demonstrates Annual Power Savings of US$78,308

Components Saved in United Fabric	Power Savings (Watts)
4 NICs per server rated at 3W per NIC	19,800
6 Gigabit Ethernet access-layer ports per server at 9W per switch port (4 NIC and 2 LOM network connections)	89,100
Add total power for 110 Cisco Nexus 5020 Switches calculated at 480W each	−52,800
Eliminate 12 third-party SAN edge switches and replace SAN aggregation layer with Cisco MDS 9513 Multilayer Directors (net power savings shown)	4,116
Total direct power savings	60,216
Power and cooling savings based on power usage effectiveness (PUE) of 2.0	120,432
kWH per year	1,054,984 kWH
Annual customer savings based on US$0.712 per kWH	US$75,114

The direct power savings of 60,216W translates into a much greater savings when cooling costs and other data center inefficiencies are factored in. Power usage effectiveness (PUE) is the ratio of the total facility power (as measured at the electric meter) divided by the IT equipment load. The total facility power is much higher than the IT load because it includes the power to cool the data center, inefficiencies in power distribution including loss in uninterruptible power supplies, lighting, and humidity control.

A study by The Green Grid (The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE) available at http://www.thegreengrid.org/gg_content) reports that many data centers have a PUE of 3.0, indicating that 3W of power is required for every watt delivered to IT equipment. It cites studies that indicate that a PUE of 2.0 can be achieved with proper data center design. The customer used a PUE of 2.0 in the calculations, so the 60,216W savings can be doubled due to the avoidance of data center overhead. This amounts to more than 1 megawatt-hour per year, or a savings of US$75,114.

Savings Beyond Power and Cooling

While the customer's focus was on the power and cooling savings achieved by implementing Unified Fabric at the access layer, a number of other savings became obvious in the process of comparing the two models:

- Although the overall capital cost savings were not evaluated, the customer noted that avoiding the purchase of four Gigabit Ethernet adapters per server alone would save the company US$1,254,000.
- Co-locating the Cisco Nexus 5020 Switches in the server racks saved 480RUs, and 210RUs were saved by eliminating the SAN edge switches. This total of 690RUs is the space equivalent of 172 servers, opening up space for potential future expansion.
- The direct power savings leaves capacity for an additional 120 servers using 500W each, leaving room for a 7.2 percent expansion in server capacity.

Conclusion

Moving from a discrete LAN and SAN design to a Unified Fabric is a strategic move that pays off in power and cooling savings, allowing data centers to allocate more of their energy budget to powering servers that deliver applications to customers.

The 41 percent power and cooling savings calculated by one Cisco customer yields an annual cost savings of US$75,114. A customer implementing a 10,000-server data center with a Unified Fabric could achieve annual savings of US$455,236, and potentially more depending on local power rates.

The power savings is just the beginning of the list of benefits. The move from Gigabit Ethernet to a 10 Gigabit Ethernet–based Unified Fabric reduces the number of NICs required while giving a performance boost to network-intensive applications. Because it uses a single converged network adapter for all I/O, the Unified Fabric enables a "wire once, use later" deployment methodology in which every server can be configured exactly the same, with features such as FCoE enabled as they are needed. This methodology simplifies infrastructure, reduces costs, and reduces deployment times. Simpler infrastructure means greater reliability because there are fewer adapters, cables, switch ports, and switches that require maintenance and that can fail.

The power and cooling savings make a compelling case for unified networking. The additional benefits of Unified Fabric make the Cisco Nexus 5000 Series a product line that makes good business sense.

For More Information

For more information on the Cisco Nexus 5000 Series Switches, visit http://www.cisco.com/go/nexus5000.

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.

Printed in USA

C11-497077-01   06/09